

Large-scale Measurement Experiments of P2P-TV Systems. Insights on Fairness and Locality

Thomas Silverston, Lorand Jakab, Albert Cabellos-Aparicio, Olivier Fourmaux, Kavé Salamatian, Kenjiro Cho

► To cite this version:

Thomas Silverston, Lorand Jakab, Albert Cabellos-Aparicio, Olivier Fourmaux, Kavé Salamatian, et al.. Large-scale Measurement Experiments of P2P-TV Systems. Insights on Fairness and Locality. Signal Processing: Image Communication, Elsevier, 2011, 26 (7), pp.327-338. 10.1016/j.image.2011.01.007 . hal-00648019

HAL Id: hal-00648019

<https://hal.archives-ouvertes.fr/hal-00648019>

Submitted on 24 Oct 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Large-scale measurement experiments of P2P-TV systems insights on fairness and locality

Thomas Silverston^{a,b,*}, Loránd Jakab^c, Albert Cabellos-Aparicio^c, Olivier Fourmaux^d, Kavé Salamatian^e, Kenjiro Cho^f

^a The University of Tokyo, 7-3-1 Hongo, Bunkyo-Ku, Tokyo 113-8656, Japan

^b Japanese-French Laboratory for Informatics (JFLI/CNRS), Japan

^c Universitat Politècnica de Catalunya, Barcelona, Spain

^d UPMC Univ Paris 6, LIP6/CNRS, UMR 7606, France

^e Université de Savoie, LISTIC Polytech' Annecy-Chambéry, France

^f Internet Initiative Japan, Tokyo, Japan

ARTICLE INFO

Article history:

Received 29 October 2010

Accepted 30 January 2011

Available online 22 February 2011

Keywords:

Peer-to-Peer

Video live streaming

Measurement experiment

Traffic analysis

ABSTRACT

P2P-TV is an emerging alternative to classical television broadcast systems. Leveraging possibilities offered by the Internet, several companies offer P2P-TV services to their customers. The overwhelming majority of these systems, however, is of closed nature, offering little insight on their traffic properties. For a better understanding of the P2P-TV landscape, we performed measurement experiments in France, Japan, Spain, and Romania, using different commercial applications. By using multiple measurement points in different locations of the world, our results can paint a global picture of the measured networks, inferring their main properties. More precisely, we focus on the level of collaboration between peers, their location and the effect of the traffic on the networks. Our results show that there is no fairness between peers and that is an important issue for the scalability of P2P-TV systems. Moreover, hundreds of Autonomous Systems are involved in the P2P-TV traffic and it points out the lack of locality-aware mechanisms for these systems. The geographic location of peers testifies the wide spread of these applications in Asia and highlights their worldwide usage.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

Peer-to-peer video live streaming applications (P2P-TV) emerged recently as a new framework to deliver live video such as television over the Internet. The quick spread of these applications surprisingly shows that user oriented technologies, based on collaboration between

similar users without a central control entity, are capable to deliver delay sensitive multimedia content. As a consequence, the Internet counts today several of these applications such as PPStream [1], PPLive [2], SOPCast [3], TVUPlayer [4] or TVAnts [5]. Millions of users located all over the globe are involved and hundreds of channels can be watched live.

The P2P model, essentially known for its scalability, is a practical solution for broadcasting live events or TV shows to a large number of receivers without any deployment cost as it is the case with content distribution networks (CDNs) [6]. Nowadays, P2P traffic contributes greatly to the Internet traffic [7]. The same observation applies on video streaming traffic generated by platforms

* Corresponding author at: The University of Tokyo, 7-3-1 Hongo, Bunkyo-Ku, Tokyo 113-8656, Japan.

E-mail addresses: thomas@hongo.wide.ad.jp (T. Silverston), ljakab@ac.upc.edu (L. Jakab), acabello@ac.upc.edu (A. Cabellos-Aparicio), olivier.fourmaux@lip6.fr (O. Fourmaux), kave.salamatian@lip6.fr (K. Salamatian), kjc@ijlab.net (K. Cho).

such as YouTube [8]. Thus, P2P-TV applications that combine these two technologies are expected to account for a large part of the Internet traffic.

However, the main problem remains in characterizing the unknown effect of P2P video streaming traffic on the Internet and on Internet Service Providers (ISPs) networks. It was already shown that the traditional P2P file-sharing traffic is a serious threat for ISPs [9]. Nevertheless P2P video streaming traffic, which consumes a lot of bandwidth resources and is very sensitive to the end-to-end delay, is a more intriguing case. Moreover, the fact that television services target a huge number of users spread worldwide further complicates traffic engineering tasks for ISPs. Therefore, it is of a great significance to better characterize the impact of P2P-TV traffic on the Internet and ISPs networks [10].

Numerous P2P-TV measurement experiments focused mainly on the reverse engineering of commercial applications [11–13]. Because most of these experiments studied the traffic from a single measurement point, as it was the case with our previous work [14], the main goal of these works was to infer the underlying mechanisms or architectures used by these proprietary applications. However, these applications are used at the planet-scale and the geographic location of peers, the users' interest in content according to its location, and their Internet access environment have an impact on the behavior of users and the properties of collected traffic.

In this paper, we present our large-scale measurement experiments of P2P-TV systems. We collect the traffic from multiple measurement points located on both sides of the world: in France, Japan, Romania and Spain. Through these measurement experiments, we study the overall P2P-TV networks and extract new characteristics relevant for ISPs and application designers. We study the global organization of the peers and the amount of traffic they exchange between them. We also study the distance among them and discuss the number of ASes involved in the traffic and the impact on ISPs. We analyze the geographic distribution of users and the effect of the content on the P2P-TV population.

The contribution of this paper can be summarized as follows: (i) we perform large-scale measurement experiments of P2P-TV systems that lead to uncovering the following drawbacks; (ii) the lack of fairness between peers in exchanging data. This is an important issue for the viability of P2P systems, which mainly relies by nature on the reciprocal contribution between peers; (iii) the lack of locality-aware mechanisms that makes P2P-TV traffic a dilemma for ISPs since hundreds of ASes are involved in the traffic; (iv) we also provide a comprehensive study on the geographic location of users and the impact of the content on the population of P2P-TV users.

The remainder of this paper is organized as follows. In Section 2, we present our measurement experiments, the experimental testbeds, and the measurement methodology. We also describe the collected traces and their main properties. The results of our measurement experiments and traffic analysis are presented in the following sections. Section 3 analyzes the exchanges of traffic between peers, while Section 4 focuses on the locality of peers. The

geographic distribution of peers is emphasized in Section 5. Section 6 discusses the related work. We conclude this work in Section 7 and present our research perspectives.

2. Measurement experiments

In this section, we present the two sets of experiments that have been conducted to collect the data analyzed throughout the paper. First, we describe our testbed located in France and Japan. Then, we introduce the second testbed in Romania and Spain.

2.1. Japan–France testbed

For this measurement experiment, we passively collected the traffic from multiple points located in France and Japan. We focused on the most popular P2P-TV applications, namely PPStream, TVUPlayer, SOPCast and TVAnts. We selected these applications according to our feedbacks and those from the online community [15]. During our experiments, we measured live soccer games because such events exhibit a real interest to be watched as they happen. There is also a large community of P2P-TV users for this purpose.

Our measurement testbed is described in Fig. 1 and is composed of two distinct parts situated in France and Japan. In each part, we collect packets by using PCs equipped with 1.8 GHz CPU, common graphic card capabilities, and Windows XP. For each of the four measured applications, we performed an experiment involving a different number of PCs according to their availability (five to seven PCs at the same time). Three or four PCs were situated in the UPMC campus network in France and were directly connected to the Internet through a 100 Mbps Ethernet link. We used two to three PCs in the campus network of the University of Tokyo in Japan, also directly connected to the Internet (Ethernet 100 Mbps). During each experiment, all the PCs were running the same P2P-TV application as well as *windump* to collect the packets. All the video bitrates were 400 Kbps.

Table 1 summarizes the collected traces (publicly available on [16]). All the traces have the same duration of 2 h 45 min. This duration is slightly larger than a soccer




Fig. 1. Large-scale measurement experiment testbed from Japan and France. Each PC is directly connected to the Internet within UPMC or the University of Tokyo network.

Table 1

P2P-TV traffic traces from the Japan–France experiment. All the traces have the same duration: 2 h 45 min (165 min).

	Size (GB)	Number of IPs	Similarity (%)	Up. (%)	Down. (%)
PPStream					
France A	3.1	2625	68	82	18
France B	3.1	3317	57	82	18
France C	3.1	3224	59	82	18
Japan X	3.0	3421	54	82	18
Japan Z	3.1	2544	67	82	18
TVUPlayer					
France A	1.7	2122	46	66	34
France B	1.6	1262	66	68	32
France C	1.6	1093	70	67	33
Japan X	1.8	1111	89	67	33
Japan Y	2.2	1034	90	71	29
Japan Z	3.1	1064	92	78	22
SOPCast					
France A	1.0	3755	79	43	57
France B	1.2	4268	73	50	50
France C	1.2	3920	76	53	47
France D	1.0	3925	80	41	59
Japan X	3.2	4269	81	78	22
Japan Z	3.0	4048	83	77	23
TVAnts					
France A	2.7	1854	97	79	21
France B	2.0	1864	97	72	28
France C	2.6	1768	97	78	22
France D	2.7	1887	97	79	21
Japan X	2.4	1855	97	77	23
Japan Y	2.3	1862	97	75	25
Japan Z	2.5	1877	97	78	22

game (105 min) because we wanted to capture the effects that happen at the beginning or the end of the games. For clarity reasons, we refer to the PCs situated in France as *France A–D* and those in Japan as *Japan X–Z*. We measured PPStream with five PCs (nodes) during the UEFA Champion's league between Liverpool and Toulouse (08/15/2007). Six nodes measured TVUPlayer during the qualifying game for the Olympic soccer tournament between Japan and Vietnam (08/22/2007). For SOPCast, we used six nodes during a similar event but with China and Japan (08/03/2007). TVAnts was measured with all the seven nodes for another qualifying game between China and Vietnam (08/23/2007).

2.2. Romania–Spain testbed

In order to extend the scope of our work [17], we performed an additional experiment using PPLive [2]. Indeed, PPLive is also one of the most popular P2P-TV applications and is largely deployed and used especially in China [18]. It was therefore significant to include PPLive in our study. We also took advantage of the opportunity to perform an additional experiment to collect the traffic during a different kind of event. Instead of sports events, we collected the traffic during the “Memorial service of Michael Jackson” on July 7, 2009. Since this artist was famous worldwide, we expected that this huge media event had the potential to attract a large number of users

Table 2

P2P-TV traffic traces from the Romania–Spain experiment. Romania trace is 2 h 43 min (163 min) and Spanish trace is 3 h 23 min (203 min).

PPLive	Size (GB)	Number of IPs	Similarity (%)	Up. (%)	Down. (%)
Romania	1.44	23,713	N/A	61	39
Spain	0.71	3252	N/A	10	90

located all over the world. Just like for sport events, there was a strong interest for users to watch this event live.

Regarding the measurement testbed [19], the traffic has been collected by using two vantage points: one in Spain and other in Romania. Similar to previous experiments, a vantage point is a PC running the PPLive application and watching the live channel broadcasting MJ memorial. The video bitrate was 400 Kbps. Each PC collected the traffic by using Windump, as previously explained. The node in Spain was situated in the campus network at the Technical University of Catalonia, in Barcelona. The campus network (Spain) has high-speed access to the Internet, using a 2 Gbps Ethernet access. The PC was equipped with 100 Mbps network interface, which limits the bandwidth use. The use of a firewall limited also its connectivity. The node in Romania was connected to the Internet through a home broadband network in Cluj-Napoca with a 100 Mbps Ethernet access. It therefore presents similar bandwidth condition as nodes in the Japan–France experiment.

Table 2 summarizes the collected traces. The trace collected in Romania is named as PPLive (RO) while the one from Spain is named as PPLive (ES). The two traces have different lengths: 2 h 43 min for PPLive (RO) and 3 h 23 min for PPLive (ES).

Even though the measured events, number of measurement points or network conditions are different for the two experiments, these two experiments are complementary and help to improve our overall knowledge of P2P-TV systems. The results of the Romania–Spain experiment will extend those from the Japan–France experiment and put some of its results into perspective.

2.3. Data set observations

For the Japan–France experiment, we present in Table 1 some statistical properties of the traces such as their size or number of IPs (i.e. remote peers) encountered in each trace. We also compute a similarity measure (described hereafter) and the ratio of upload and download traffic.

The upload ratio is more important for large traces than for the smaller ones and it testifies to a more important upload activity. Indeed, a peer aims to download the video only once but can upload it several times to remote peers. Regarding the volume of download, it may range from 512 MB for France B to 682 MB for Japan Z with TVUPlayer, while the traces have the same duration (2 h 45 min). This additional amount of download traffic comes from the signaling traffic generated by the remote peers. This observation is an important issue in several

scenarios. An altruistic peer serving video will receive in return a large amount of signaling traffic from those remote peers orchestrating their download. It will waste its bandwidth resource to download the video and could directly affect the video quality if its access link becomes saturated.

In order to understand if the number of measurement points (five to seven) was sufficient to obtain a global view of the P2P network, we computed a “similarity measure”. The similarity measure is defined as the ratio of IPs from a trace that is also present in the other traces of the same application. For instance, France B with PPStream accounts 3317 IPs where 57% of those are also present in the other PPStream traces. For PPStream (five nodes), in average 61% of IPs are also present in its other traces. The average similarity for TVUPlayer is 75% and 79% for SOPCast (six nodes) and 97% for TVAnts (seven nodes). We observed that the more measurement points we had, the more our nodes communicated with similar sets of peers. This validates that the observed high similarity ratio indicates that our nodes encountered the entire population of peers in the network, leading to a global and precise view of the P2P network.

Our experimentation showed that only seven measurements points allowed us to attain a similarity measure of 97%. This lead us to believe that this order of magnitude of measurement points is adequate to measure entirely a P2P network. This is an important observation as it shows that there is no need to measure the network with a very large number of peers (hundreds, thousands) in order to have a precise view of the network. In our experiments the cases with five to six nodes are not as precise as with seven nodes. However, they give a much more precise and global view of the P2P networks when compared with previous studies with only a single measurement point [14].

Regarding the Romania–Spain experiment in Table 2, the PPLive traces present different properties. The trace collected from a home network in Romania has a similar size and traffic ratio as those in the Japan–France

experiment (i.e. SOPCast, TVUPlayer) but presents a very high number of IPs. The Spanish one is much smaller than all the other but still counts an important number of IPs. The campus network in Spain were behind a firewall, which limited its connectivity to remote peers of the Internet. It explains its relative small number of IPs compared to the trace from Romania and also its very low upload activity (10% of the traffic in upload). For this experiment, we did not compute the similarity measure for the PPLive traces. In this case, it was not relevant since there is an order of magnitude in the number of IPs for each of the PPLive traces, i.e. 23,713 for PPLive (RO) and 3252 for PPLive (ES).

The large number of IPs for PPLive can be explained from the fact that, two years after the first experiment, P2P-TV applications – or at least PPLive – is getting more and more popular and used to watch TV online. The nature of the measured event could also be responsible to attract a huge amount of viewers whereas a soccer game may only have regional interest (for the countries involved in the game). Finally, this evolution in number of IPs with PPLive confirms the need for an additional measurement experiment to obtain a more accurate knowledge of P2P-TV systems.

3. Collaboration between peers

In P2P systems, peers are responsible to duplicate the content to others. The overall P2P networks rely on the effective collaboration of users. In this section, we study the amount of traffic exchanged with each remote peers. In order to evaluate the level of collaboration between peers, we compute for each peer the “sharing ratio”. The sharing ratio of a peer is the amount of traffic it uploads divided by the amount of traffic it downloads. The sharing ratio is a key parameter that indicates the level of collaboration between peers.

For all the applications, Fig. 2 shows the cumulative distribution function (CDF) of the sharing ratio (upload/




Fig. 2. Traffic sharing ratio (upload/download). The ratio evaluates the level of collaboration between peers. When ratio is above 1, our peer is altruistic; when equal to 1, the exchange is fair; below 1 our peer take advantage of the resources of others. (a) Japan–France experiment and (b) Romania–Spain experiment.

download) between a peer which we control, and all the peers with which it exchanges data. In Fig. 2(a), we show a single trace for each application of the Japan–France experiment because it is representative for the other traces. Fig. 2(b) shows the two PPLive traces. The X-axes of the figures show the sharing ratio. When the ratio is above 1, our peer is altruistic toward another peer (i.e. our peer sends more data than it receives); when equal to 1, the exchange is fair; when below 1, our peer is taking advantage of the altruism of others.

For PPStream and TVAnts, 10% of peers are altruistic with our controlled peer and transmitted more data than they received. The 90% remaining are beneficiaries of our altruism. This behavior is the same for TVUPlayer at a distinct rate: 30% of peers are altruistic and the remaining 70% benefits from our altruism. SOPCast shows the opposite behavior: the majority of the remote peers transmitted more data than they received (60%). We can distinguish two different cases for PPLive (Fig. 2(b)), whether the traffic has been measured in Romania or in Spain. For PPLive (RO), about 55% of the remote peers are altruistic with our peers and the remaining benefits from our resources, such as with SOPCast. For PPLive (ES), almost the entire population (95%) of peers is altruistic with our peer. This was expected since our node in Spain was situated behind a firewall, limiting its abilities to upload traffic to remote peers on the Internet.

Overall, the traffic exchanges for all the applications are never fair to our controlled peers, with the large majority being mainly consumers (ratio > 1) rather than producers (ratio < 1). SOPCast and PPLive (PPLive RO) presents also an unfair sharing ratio since a large majority of peers are altruistic with our peers.

However, there exist distinct peer behavior traits as those presented in Fig. 2. Looking at Fig. 3, we isolate uplink and downlink traffic between one of our peers and two remote peers on the Internet. We extracted this traffic from a PPStream trace, but it is representative for other traces and applications. In Fig. 3(a), the two curves are not the same, but show that our controlled peer and a remote peer try to contribute in the same proportion

during this exchange. In contrast, in Fig. 3(b) we can see that our peer has contributed a lot of information while it has received virtually nothing in exchange from the other remote peer. Clearly, whatever the applications or controlled peers, the second case (Fig. 3(b)) is the case that is the most likely to happen: unfair exchange between peers without any reciprocity.

These results show that fairness is not achieved in P2P-TV systems. Since each peer is in charge to forward the traffic to other peers, each peer should forward the same amount of data it received. If it is not the case, the peers are not helping the other peers by sharing their resources in the network and this is a major issue for the scalability of these P2P systems.

4. Locality of peers

In this section, we study the distance from the remote peers to ours to fetch the video. We study this distance at the IP level and at the Autonomous System level (AS).

4.1. Distance in IP hops

In order to infer the distance between peers, we investigate the IP datagram of the packets. The IP header has a TTL field, which is decreased at each hop by the Internet routers. This piece of information reveals the distance in number of IP hops from the source to the destination. Given that the default value used for the TTL field of the IP header is usually set to 128 (MS-Windows) or 64 (Linux, FreeBSD) the chances are that, if a packet arrives with a TTL of 120 (or 60), the packet was originally sent with a TTL of 128 (respectively 64) [20]. The distance in number of hops between the source and the receiver of the IP datagram is therefore eight hops (respectively four). Since most of the measured P2P-TV applications have been only built for MS-Windows, 85% of the received datagrams in all the studied traces have a default TTL value set to 128. The other 15% of packets have a default TTL value set to 64. Half of these packets were generated by SOPCast, which is the only application among those we




Fig. 3. Examples of traffic exchange between a controlled peer and a remote peer on the Internet (Ex: PPStream). (a) Fair exchange of traffic between two peers and (b) unfair exchange of traffic between two peers.

measured that has a Linux implementation (default TTL is 64). The remaining packets had their TTL field that has been rewritten by the traversed Network Address Translator systems [21]. For security issues, the TTL field is commonly set homogeneously among all packets forwarded toward the external network.

Fig. 4 shows the CDF of the distance in number of IP hops between the source and the destination. For the Japan–France experiment (Fig. 4(a)), all the applications show similar results and we present only those for TVAnts. From Fig. 4(a), there are two distinct traits for the traces collected in France and those collected in Japan. For our French nodes, 75% of the remote peers are situated from 20 to 30 hops. For the Japanese nodes, 80% of the remote peers are situated from 10 to 16 hops. Clearly, our French peers download the video at further distance than the Japanese peers. For PPLive (Fig. 4(b)), 90% of the remote peers are situated from 15 to 25 hops, whatever the location of the nodes (Romania or Spain). This is about the same distance as for our French nodes.

Several facts can explain this trend. First, P2P-TV applications have been released by Chinese companies and are already very popular in Asia. Most of the peers and source servers are probably situated in Asia. The use of P2P-TV in Europe is still limited and even PPStream or PPLive, which were broadcasting programs of interest for European users (UEFA Champions' league, MJ memorial), have the same behavior as the other applications. Consequently, if the large majority of peers are situated in Asia, the French or European nodes must fetch the content from peers at further distance compared with the Japanese nodes. Second, these applications do not use any locality-aware mechanisms to select the provider peers. Otherwise, the French nodes would have downloaded the video at a closer distance. A possible scenario could be that only a single French node downloads the video from a far distance and then forwards the content to the other French nodes situated in the same network. In this case, we would observe in Fig. 4(a) only a single French peer that downloads far from it and the other get the data at a

closer distance. This scenario also raises feasibility issue related to the real-time constraints of the exchanged traffic that should be investigated.

The long distance for the peers to get the video is a critical issue since it may increase the end-to-end delay to receive the video packets and affect the quality of the video. This also points out the fact that traffic is exchanged through different ISPs and crosses inter-ISP links to reach distant destinations.

4.2. Autonomous Systems

To extend our study of the distance between peers and the locality of the traffic, we investigate the amount of traffic exchanged between different Autonomous Systems (ASes). Even though the network of ISPs can be divided into several ASes, the number of ASes involved in the exchanges provides insight on the traffic between ISPs. To this end, we mapped each IP address to its parent AS by using the `whois` [22] service provided by Cymru [23] and we aggregate the traffic by ASes in upload (Fig. 5(a) and (b)) and download (Fig. 5(c) and (d)). For the Japan–France experiment, since it is about 24 traces, we present one trace per application because the other traces from the same application show similar results.

Hundreds of ASes are involved in the traffic exchanges, which indicate a large dispersion of peers. This is even more pronounced in upload than download. PPStream and TVUPlayer involve approximately 200 ASes in upload, 150 for SOPCast or PPLive (RO). For PPStream, SOPCast and PPLive (RO), 90% of the upload traffic goes toward 50 different ASes. The traffic of TVUPlayer is more spread: 90% of its upload traffic goes to 75 ASes. For TVAnts, only 30 ASes count for 90% of its upload traffic: half as many ASes as the other applications.

For all the other applications, the download traffic comes from a much smaller set of ASes than the upload traffic. For PPStream and TVUPlayer, 90% of the download traffic comes from 30 ASes, 25 ASes for PPLive (RO), 10 ASes for TVAnts and only 3 ASes for SOPCast. As expected,




Fig. 4. Distance in number of IP hops between our controlled peers and remote peers on the Internet. (a) Japan–France experiment (Ex: TVAnts) and (b) Romania–Spain experiment.




Fig. 5. Traffic by Autonomous Systems (ASes). (a) and (c) Japan–France experiment; (b) and (d) Romania–Spain experiment.

the peer behind a firewall (PPLive (ES)) counts only a few ASes in the traffic because of reachability limitation.

From Fig. 4(a), it appears clearly that during the same event, the French peers download the video at longer distance in IP hops compared with the Japanese one. In order to obtain the number of AS hops between our controlled nodes in France and the other peers on the Internet we performed, during the Japan–France experiments, a `traceroute` from our controlled France peers toward each destination. It allows discovering the route packets take to reach the host destinations. From the traceroute results, we obtained the corresponding AS of each network interface of a route by using the whois service. By aggregating the similar consecutive ASes of a route, we get the AS path between our nodes and each destination. The traceroutes have been performed only from our nodes situated in France, since this involves peers that download the video far from them in term of network hops.

We present for a French trace of SOPCast the CDF of the traffic according to the number of AS hops (Fig. 6). French traces of all applications show similar behavior. We observe that 50% of the upload traffic is within a radius of three AS hops whereas only 30% of the download traffic comes from this distance. Fifteen percent of the

traffic is downloaded at a distance ranging from 5 to 7 AS hops while upload traffic never reach destinations beyond five AS hops.

We noticed previously that the nodes in France download the video at farther distance than the Japanese one. Moreover, a large number of ASes are involved in the exchanges. Consequently, for the French nodes, the download traffic has transited into much more ASes compared with the upload traffic.

These results highlight that P2P-TV systems do not consider the locality of peers when exchanging traffic. It is a critical issue for ISPs and it imposes significant traffic engineering challenges. For ISPs, it is important to keep the P2P traffic local. Otherwise it may overload the links between ISPs, which are already considered as the network bottlenecks [24]. These inter-ISP links are also very costly and ISPs aim to limit the use of these links to their customers. Clearly, service providers need real incentives to use such links to convey traffic to peers from other ISPs.

5. Geographic location of peers

In this section, we study the geographic location of peers and the volume of traffic exchanged by countries.

We mapped a peer's IP address to its origin country by querying the free MaxMind GeoIP database [25]. For each trace, we separate the upload and download and present the volume of traffic and the population by countries. Each trace is therefore represented by four stacked histograms as shown in Figs. 7–11. For all these plots, the legends are similar and present the 20 countries that are the most represented in all our traces. The rest of the traffic comes from countries that never generate more than 1% of the traffic and is included under the label "other" on the legend. As an example, if we consider Fig. 7 for PPStream and the trace France B: the ratio of upload peers is distributed as follows: 35% of peers are from China, 15% from Hong Kong, 5% from France, 15% from United Kingdom (UK), 10% from a set of countries labeled as "other", and the rest comes from other countries of the legend. For a given trace, there may be differences between the population in upload or in download since peers may be active only in one side.




Fig. 6. Distance in number of AS hops in the Japan–France experiment (Ex: SOPCast).




Fig. 7. PPStream: geographic location of traffic and peers (Japan–France experiment).

5.1. Influence of the content on the population of peers

Regarding the countries involved in the traffic, as it was expected, the broadcasted event has an effect on the geographic location of peers present in the traces. For instance, we observe an important number of peers from United Kingdom (UK) with PPStream because a soccer game with the team of Liverpool was broadcasted. There are also many peers from China and Japan with SOPCast. To mitigate our observations, it is worth highlighting that the number of French peers is not as important as the number of UK peers with the PPStream games. Moreover, regarding the dataset of the Japan–France experiment, a small ratio of French peers exists in almost every trace of each application even if the broadcasted events did not involve this country. This can be due to the fact that a part of our testbed is located in France; hence attracting several French peers for collaboration. However, we did not observe the same behavior with Japanese peers although the second part of our testbed is situated in this country. For Japanese peers, there is a sharp contrast depending on whether the broadcasted content involved Japan or not.

Surprisingly, in all the traces, there is only a few number of peers from the United States (US), while Chinese peers represent a large part of the population. The presence of US traffic is barely more pronounced for PPLive (15% of download traffic for PPLive (RO) in Fig. 11), while there is still a large amount of Chinese peers. Regarding the Japan–France experiment, the lack of US users is probably due to the unpopular nature of the broadcasted content (soccer games). The schedule of the games was also in the morning or during the working time for US residents, limiting the number of potential viewers. However, from the Romania–Spain experiment, we still observe only a few US users, while the memorial for Michael Jackson was supposed to be of interest for this population of users. Then this second experiment points out the unpopularity of P2P-TV applications in the US.

The limited number of US users may come from the following reasons. First, P2P-TV applications have been




Fig. 8. SOPCast: geographic location of traffic and peers (Japan–France experiment).




Fig. 9. TVUPlayer: geographic location of traffic and peers (Japan–France experiment).




Fig. 10. TVAnts: geographic location of traffic and peers (Japan–France experiment).




Fig. 11. PPLive: geographic location of traffic and peers (Romania-Spain experiment).

launched by Chinese companies and are already very popular in China but suffer from a lack of popularity in other countries and especially in the US. Second, the limited deployment of high-speed Internet access in the US may also slow down the spread of P2P-TV in this area. Practically, the US company Zattoo [26] already launched their application in Europe but it is still not available in the US. The limited Internet broadband access is one of the reasons behind this strategy.

5.2. Population and volume of traffic

There is not necessarily a relationship between the population of peers and the volume of traffic exchanged by country. In the PPStream traces (Fig. 7) China represents 30% of the peers that generate from 10% to 20% of traffic in upload. Chinese peers in France B (SOPCast) represent 40% of the population in upload (Fig. 8) but only 8% of the volume of traffic. This statement is also verified in TVUPlayer, TVAnts or PPLive (respectively in Figs. 9–11). For instance, 5% of the TVUPlayer peers are from Korea and count for 30% of the download traffic. For PPLive (RO), 70% of the upload peers are from China but count only for 40% of the traffic.

In download, SOPCast and PPLive show a particular behavior since they count only a few number of countries. More generally, in download, our nodes fetch the content almost entirely from Asia. Concerning upload, the traffic is spread into much more countries. The upload traffic of our nodes located in Europe (France, Romania) is directed toward Europe whereas the upload traffic of our Japanese nodes is toward Asia. Only the node from Spain does not show the same behavior because it cannot upload a lot of traffic because of the firewall.

5.3. Impact on the collaboration of peers

If we pay more attention to the volume of traffic in both directions, for all the applications, our nodes download the largest part of traffic from Asia. Then, the nodes located in France and Romania upload more to Europe

and those in Japan to Asia. (The use of a firewall for PPLive (ES) biased this observation.) For instance, France C with PPStream (Fig. 7) downloads 70% of its traffic from Asia and 20% from Europe but uploads 30% to Asia and 50% to Europe. For Japan X with SOPCast (Fig. 8), almost all download traffic and 70% of the upload is related to Asia. For PPLive (RO) (Fig. 11), 65% of the download traffic comes from Asia where only 40% of the upload traffic return to Asia.

At the country-level, our nodes do not trade equitably the data. Intuitively, this result can be extended to the peers-level. Peers download the data from a country and forward in turn to another one; they do not reciprocate fairly with their provider peers. This corroborates our previous observation that fairness was not achieved in the P2P-TV systems (Section 3).

Geographic location of peers has a considerable effect on this observation. In fact, the video streams follow a directed path among peers and it is not possible to reciprocate with data to the provider peers. Our French or European nodes are not able to reciprocate with data to the Asian peers because they are situated above in the play-out point of the video. Thus, they are late in the video playback and do not have any data of interest to transmit in return to Asian peers (ahead in the playback time). They can only transmit to other European peers that are in the same playback time of the video. For our Japanese nodes, they are situated in similar play-out points as the other Asian peers and can reciprocate with data. This case is illustrated in Fig. 12. Peer 1 is ahead on the play-out time compared with Peer 2 and has some blocks of interest for Peer 2 (blocks 3 and 4). Since Peer 2 is late, it may not have any block to send in return to Peer 1.

This phenomenon could directly come from the long end-to-end network delay and number of network hops from Asia to Europe since we observed previously that our French nodes download the traffic at a further distance than the Japanese nodes (Section 4). It is also consistent with the fact that the download traffic comes from a larger distance in number of AS hops than the upload




Fig. 12. Uncooperative behavior of peers due to the nature of the content and its *temporal* constraints. Peer 2 is late compared with Peer 1 and does not have any blocks of data to send to Peer 1 in return.

traffic. Then, peers with long delay will be late from others. In that case, late peers could not necessarily collect the data on time, and if a block of data arrives after its playback deadline, it is useless and the peers experience a poor video quality. Moreover, late peers could not be able to reciprocate with data to their provider peers because they have no data of interest to transmit in return (Fig. 12). Differently from BitTorrent and P2P file-sharing systems, when designing P2P video live streaming systems, the delay between peers and their lag should be taken into account and be at the core of the architecture, as it is proposed with Pulse [27].

One could also argue that the non-reciprocal exchange of traffic is due to the asymmetric Internet connection of most of the Internet users with residential connections (DSL). However, neither our controlled nodes which have high-speed symmetric access to the Internet do not reciprocate with the same amount of data to their provider peers. For instance, with SOPCast or PPLive, our French or Romanian peers download a lot from China but do not upload in the same proportion. This behavior is similar for PPStream and TVAnts with Hong Kong, or for TVUPlayer with Korea.

Residential asymmetric Internet access is not the reason that prevents the reciprocal exchange between peers. The multimedia flows impose temporal constraints that result from the continuous nature of the transmission. Peers cannot transmit data in return, not necessarily because they are uncooperative, but because the nature of the content and the temporal constraints make it pointless. Since P2P systems rely on the efficient collaboration between peers, this observation is a critical issue regarding the scalability of P2P-TV systems. It is therefore of a great significance to design an incentive mechanism adapted to the continuous nature of the multimedia flows. Such a mechanism must enforce peers to collaborate in the network, even if the continuous nature of the content being distributed mitigates against the transmission of data in reciprocating manner.

6. Related work

Since the first release of P2P-TV applications, there have been an increasing number of P2P-TV measurement experiments.

The P2P-TV designers themselves have conducted some studies [28–30] in order to describe their systems, to evaluate and improve them. They have a complete knowledge of the system under test. Based on the data

collected by Zattoo [28], Shami et al. presents statistics about the capacities of peers to redistribute the data and the ratio of peers behind NAT systems. They show that most of the Zattoo users are behind NAT systems and cannot contribute to the P2P-TV system. By obtaining the source code of UUSEE, Wu et al. [29] measure the behavior of the UUSEE P2P streaming service. They show that the streaming quality experienced by peers has been decreasing over time, as the source servers' capacities are saturated. Since some channels are popular and some are not, they propose a model to predict the demand in each channel and to provision the server capacities according to the demand over time. Following this study, Liu et al. [30] explore the influential factors to peer longevity and bandwidth contribution level in a large-scale P2P live streaming system. They propose a mechanism to favor these superior peers in network because they would be more likely to contribute in the P2P streaming system. Recently, after collecting 200 GB of data during the 2008 Olympics [31] with UUSEE, Liu et al. evaluate the use of network coding and discuss its potential limitation due to a high signaling overhead.

Some other studies perform active measurements of the systems by actively probing the peers of the network. In their works, Hei et al. [32] and Vu et al. [33] modified a PPLive client to crawl the PPLive network. Vu et al. [33] indicate that the PPLive overlay characteristics differ from those of P2P file-sharing. They show that PPLive network is similar to random graphs and thus more resilient to the massive failure of nodes. By getting information on the video blocks available by each peer, Hei et al. [34] studied the overall video quality for all the users in the PPLive network. They also used their crawler to transmit fake video chunks [35] demonstrating that PPLive had no defense against pollution attacks.

Passive measurement of the traffic is also performed when it is the only feasible way to get knowledge of the P2P-TV systems, which appeared as *black boxes*. Hei et al. [11] was – to the best of our knowledge – the first work to perform passive packet sniffing to analyze PPLive. They focus on the video session lifetime of peers and the download rate of their controlled peers. In our previous work [14], we measured the traffic of four commercial applications in order to infer their underlying mechanisms. We study the properties of the video and signaling traffic, the download policies of peers and the number of peers in their neighborhood. In [12], we provide a deeper analysis of the P2P-TV traffic and study the behavior of P2P-TV users. Recently, Alessandria et al. [13] analyze the behavior of four P2P-TV applications. They capture the traffic and vary the network conditions with their access router to study how the different P2P-TV applications adapt dynamically to these changes. They show that all the tested applications are effective in trying to overcome network impairment but they become aggressive to download the video stream in case of bandwidth bottleneck.

7. Conclusion

In this paper, we present our P2P-TV large-scale measurement experiments that have been conducted in

order to characterize the overall P2P-TV network and study the effect of this traffic on the Internet.

Throughout this study, we quantify relevant parameters for P2P-TV: at the network level, we provide information on the overall volume of traffic, its ratio in upload and download and the number of IPs present in each trace. At the peer level, we indicate the amount of data each peer sends or receives, and compute the ratio of traffic the peers share in the network. For service providers, we provide useful results regarding the locality of the traffic and the number of ASes involved in the exchange. Furthermore, at the user point of view, we study the geographic location of peers and the usage of these applications.

Our results indicate that the P2P-TV traffic is exchanged arbitrarily toward many distinct locations involving many ASes, and this traffic is not kept local. This is an important issue for ISPs because it wastes their network resources by overloading the links between ISPs, increasing their costs to transport the P2P-TV traffic. The broadcasted content has an effect on the population of peers and their geographic location. The current spread of these applications in Asia explains the large presence of Chinese peers in our traces. However, there is not necessarily a relationship between the population of peers and the volume of traffic they generate. We also find out that peers do not reciprocate fairly when downloading the traffic. This is a critical issue concerning the scalability of P2P-TV. The video streams follow a directed path among peers and it is not possible for peers to collaborate and forward data in return to their provider peers.

Some perspectives of this work should entail the design of new P2P architectures or mechanisms adapted to multimedia live streaming. These architectures have to deal with the continuous nature of the transmission and its temporal constraints. For instance, we are currently designing a new incentive mechanism for P2P-TV systems. This mechanism enforces peers to collaborate even if the streaming nature of the content prevent the transmission of data in a reciprocating manner. We also plan to propose a locality-aware mechanism to select closer peers. This would improve the end-to-end delay between peers and reduce the costs for ISPs.

References

- [1] PPStream: <http://www.ppstream.com>.
- [2] PPLive: <http://www.pplive.com>.
- [3] SOPCast: <http://www.sopcast.com>.
- [4] TVUPlayer: <http://www.tvuplayer.com>.
- [5] TVAnts: <http://www.tvants.com>.
- [6] Akamai: <http://www.akamai.com>.
- [7] Cachelogic: <http://www.cachelogic.com>.
- [8] YouTube: <http://www.youtube.com>.
- [9] D.R. Choffnes, F.E. Bustamante, Taming the torrent: a practical approach to reducing cross-isp traffic in peer-to-peer systems, in: ACM SIGCOMM, 2008.
- [10] Y. Zhang, N. Zong, G. Camarillo, J. Seng, R. Yang, Survey of P2P streaming applications, <http://tools.ietf.org/html/draft-gu-ppsp-survey-02>, 2010.
- [11] X. Hei, C. Liang, J. Liang, Y. Liu, K.W. Ross, Insights into PPLive: a measurement study of a large-scale P2P IPTV system, in: IPTV Workshop, International World Wide Web Conference, 2006.
- [12] T. Silverston, O. Fourmaux, A. Dainotti, A. Botta, A. Pescape, G. Ventre, K. Salamatian, Traffic analysis of P2P IPTV communities, Elsevier Computer Networks, 53 (4) (2009) 470–484. doi:10.1016/j.comnet.2008.09.024.
- [13] E. Alessandria, M. Gallo, E. Leonardi, M. Mellia, M. Meo, P2P-TV systems under adverse network conditions: a measurement study, in: IEEE INFOCOM, 2009.
- [14] T. Silverston, O. Fourmaux, Measuring P2P IPTV systems, in: ACM NOSSDAV'07, International Workshop on Network and Operating Systems Support for Digital Audio & Video, 2007.
- [15] MyP2P: <http://www.myp2p.eu>.
- [16] Content: <http://content.lip6.fr/traces/>.
- [17] T. Silverston, O. Fourmaux, K. Salamatian, K. Cho, Measuring P2P-TV systems on both sides of the world, in: 2nd International Workshop on IPTV Technologies and Multidisciplinary Applications (IWITMA), IEEE ICME, 2010.
- [18] Y. Huang, T.Z.J. Fu, D.-M. Chiu, J.C.S. Lui, C. Huang, Challenges, design and analysis of a large-scale P2P-VoD system, in: ACM SIGCOMM, 2008.
- [19] L. Jakab, A. Cabellos-Aparicio, T. Silverston, M. Solé, F. Coras, J. Domingo-Pascual, Corecast: how core/edge separation can help improving inter domain live streaming, Elsevier Computer Networks, 54(18) (2010) 3388–3401. doi:10.1016/j.comnet.2010.07.005.
- [20] <http://www.cpni.gov.uk/Docs/InternetProtocol.pdf>.
- [21] F. Gont, P. Srisuresh, Security implications of network address translators (NATs), <http://tools.ietf.org/html/draft-gont-behave-nat-security-01>.
- [22] IETF RFC 3912: Whois protocol specification <http://tools.ietf.org/html/rfc3912>.
- [23] Cymru: <http://www.team-cymru.org>.
- [24] V. Aggarwal, A. Feldmann, C. Scheidele, Can ISPs and P2P systems co-operate for improved performance? in: ACM SIGCOMM Computer Communications Review (CCR), 2007.
- [25] Maxmind: <http://www.maxmind.com>.
- [26] Zattoo: <http://www.zattoo.com>.
- [27] F. Pianese, D. Perino, J. Keller, E. Biersack, PULSE: an adaptive, incentive-based, unstructured P2P live streaming system, IEEE Transactions on Multimedia, 9(18) (2007) 1645–1660. doi:10.1109/TMM.2007.907466.
- [28] K. Shami, D. Magoni, H. Chang, W. Wang, S. Jamin, Impacts of peer characteristics on P2PTV networks scalability, in: IEEE INFOCOM, 2009.
- [29] C. Wu, B. Li, S. Zhao, Multi-channel live P2P streaming: refocusing on servers, in: IEEE INFOCOM, 2008.
- [30] Z. Liu, C. Wu, B. Li, S. Zhao, Distilling superior peers in large-scale P2P streaming systems, in: IEEE INFOCOM, 2009.
- [31] B.L. Zimu Liu, C. Wu, S. Zhao, UUSec: large-scale operational on-demand streaming with random network coding, in: IEEE INFOCOM, 2010.
- [32] X. Hei, C. Liang, J. Liang, Y. Liu, K.W. Ross, A measurement study of a large-scale P2P IPTV system, IEEE Transactions on Multimedia, 9(18) (2007) 1672–1687. doi:10.1109/TMM.2007.907451.
- [33] L. Vu, I. Gupta, K. Nahrstedt, J. Liang, Understanding overlay characteristics of a large-scale peer-to-peer IPTV system, in: ACM Transactions on Multimedia Computing, Communications and Applications, 6 (4) (2010). doi:10.1145/1865106.1865115.
- [34] X. Hei, Y. Liu, K. Ross, Inferring network-wide quality in P2P live streaming systems, IEEE Journal on Selected Areas in Communications, 25 (9) (2007) 1640–1654. doi:10.1109/JSSAC.2007.071204.
- [35] P. Dhungel, X. Hei, K.W. Ross, N. Saxena, The pollution attack in p2p live video streaming: measurement results and defenses, in: P2P-TV Workshop, ACM SIGCOMM, 2007.