Investigating Overlay Topologies and Dynamics of P2P-TV Systems: The Case of SopCast

Original

Investigating Overlay Topologies and Dynamics of P2P-TV Systems: The Case of SopCast / BERMUDEZ CORRALES, IGNACIO NICOLAS; Mellia, Marco; Meo, Michela. - In: IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS. - ISSN 0733-8716. - STAMPA. - 29:9(2011), pp. 1863-1871. [10.1109/JSAC.2011.111010]

Availability:

This version is available at: 11583/2495554 since:

Publisher:

IEEE-INST ELECTRICAL ELECTRONICS ENGINEERS INC, 445 HOES LANE, PISCATAWAY, NJ 08855 USA

Published

DOI:10.1109/JSAC.2011.111010

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Investigating Overlay Topologies and Dynamics of P2P-TV Systems: the Case of SopCast

Ignacio Bermudez
Politecnico di Torino, Italy
ignacio.bermudez@polito.it

Marco Mellia
Politecnico di Torino, Italy
marco.mellia@polito.it

Michela Meo
Politecnico di Torino, Italy
michela.meo@polito.it

Abstract—Several successful commercial P2P-TV applications are already available. Unfortunately, some algorithms and protocols they adopt are unknown, since many follow a closed and proprietary design. This calls for tools and methodologies that allow the investigation of the application behavior.

In this paper, we present a novel approach to analyze the graph properties and the traffic generated by P2P-TV applications run by customers in operative networks. The proposed methodology allows us to distinguish and investigate three different graphs: the social networks that link users based on their interest, the overlay networks created by peers that are watching the same channel, and the distribution networks that involve the subset of peers that are contributing to the video distribution.

We apply this methodology to the traffic collected for more than one year from three national ISPs in Europe, where SopCast is the largely preferred application. Considering users' behavior, we uncover the attitude to use the P2P-TV application mainly to follow live sport events. P2P-TV systems have then to deal with both flash crowd and sudden peer departures that happen at the beginning and end of an event. Furthermore, channel zapping among channels offering the same event is also relevant. SopCast deals with this by implementing a very robust and greedy overlay topology discovery process in which more than 170 peers are contacted every 60 s. Considering video distribution, we provide evidence that SopCast implements algorithms that as consequence restrict traffic within Autonomous System boundaries. Still, high bandwidth peers must be present to supply the necessary upload capacity to sustain the video service.

I. INTRODUCTION

In the recent years we have witnessed the success of P2P-TV applications, bringing TV channels, some of which live, to the users' home through the Internet. Several commercial P2P-TV systems are available and some are popular among users because they feature cheaper video broadcasting than other solutions, e.g., IPTV or pay-TV. Unfortunately, most of the successful P2P-TV applications rely on proprietary protocols and unknown algorithms, so that the understanding of such systems is intrinsically complex. Thus P2P-TV traffic characterization has become a topic of great interest for the research community and for network operators. Both are interested in understanding the positive and negative aspects of P2P-TV applications, to understand how these complex systems work and to improve their design and effectiveness.

In P2P-TV systems, three different graphs can be identified. The first graph represents the users that run the applications forming a "social network graph". Where are the users? When and for how long do they run the application? Is churning

relevant for P2P-TV systems? These and others are all relevant questions whose answer allows researchers to design more robust applications, e.g., by exploiting natural localization properties of users.

Peers form then an "overlay topology", a second graph where peers are interconnected by logical links. Peers interested in the same channel then form a "swarm"; independent overlay topologies are built for different swarms. Which are the properties of the swarms? Are the peer neighbors carefully selected or are they randomly chosen? By understanding the overlay topology graph properties it is possible to understand the P2P-TV system properties, its robustness to churning or its scalability with respect to the number of peers.

Finally, the subset of the overlay links that are used by peers to exchange the video traffic forms the third graph, the "distribution graph". Is the video data being downloaded from neighbors in the same Autonomous System to reduce the network provider cost? Are neighbors with larger upload capacity preferentially selected to download content from? What is the fraction of high capacity peers in a swarm? Recall indeed that the total available upload capacity plays a key role in the success of P2P-TV content distribution since the video stream must be downloaded at an almost constant rate by each peer.

In this work, we contribute to the characterization of P2P-TV system graph properties by analyzing the traffic in the operative links of the networks of three ISPs in Europe. Using a purely passive methodology, we collect traffic for more than one year. Our results point that in the monitored countries SopCast is by far more popular than other applications such as PPLive and TV-Ants. Interestingly, the usage of P2P-TV applications is discontinuous and associated to events, such as sport events, that are popular but expensive to retrieve through normal TV broadcasting systems. We then focus on SopCast traffic observed during the two months when the UEFA Champions League 2008/09 final matches were held.

Compared to works that rely on active measurements, the adoption of a pure passive methodology gives to our work a surplus value, since we observe the typical usage of the system without interfering it, and moreover our measurements have information related to the users' habits in our analysis.

We propose a general methodology to identify swarms corresponding to TV channels; we observe churning associated to SopCast events, finding out that users stay connected to

the P2P-TV system for the whole duration of the event, but they can frequently change swarm seeking for better channels broadcasting the same event. We study the peer discovery process in the overlay topology, finding out that SopCast implements a simple random discovery which is very robust. Conversely, the distribution graph is severely biased by peer upload capacity and by the Autonomous System a peer belongs to. Results suggest that the implications of traffic burstiness, peer population and their evolution over time might become challenging to face.

Some key aspects we highlight in this paper include:

- Despite the average bandwidth usage of P2P-TV applications is not significant, it can be substantial during periods in which popular events are shown. Today, a few tens of users can contribute to 15% of total aggregate traffic generated by more than 20,000 users on a network access link.
- Geo-locality of social network graphs is deeply affected by cultural and language trait of customers. This biases the traffic distribution graph that is inherently geographically localized.
- Evidence shows that often high-speed residential networks and University networks altruistically serve content to residential peers with highly asymmetric bandwidth. Without the contribution of those peers, the P2P-TV system would not sustain the service at all.

II. RELATED WORK

P2P overlay and system characterization has attracted the interest of several researchers. In [1] a nice summary on overlay properties of several P2P systems for file-sharing is given. This work well defines structured and unstructured overlays, and concludes that the latters have been widely deployed in actual systems.

Several works use active crawling techniques to snapshot overlays in P2P systems for file sharing. [2]–[7]. Recently, [8], [9] propose random-walk strategies based on an unbiased method to visit vertexes by random-walk with equal probability even if the graph has a power-law distribution of the node degree. Based on active crawling, this methodology is constrained to the cases in which the communication protocol is well-known. Therefore, we cannot apply it in the context of SopCast, since the protocol is unknown and encryption mechanisms are adopted.

Similar to our work, in which the P2P protocol is obscured, [10]–[12] have studied Skype overlay using passive methodologies. They start with the assumption that Skype uses super-node peers, which is a well-known characteristic of the Skype system. They show that Skype takes advantage of super-nodes as peers with high uplink capacity. Furthermore, [12] shows also a bias induced by the social network on the Skype overlay.

Several researches have focused on single commercial P2P-TV systems, and investigated their internals using active crawling methodologies too, e.g., [13]–[16]. None of them focuses on the characterization of P2P-TV graphs properties. Only [16] shows the importance of channel swapping, which is particular to P2P-TV systems. However, their motivations are beyond our scope, since they propose a way to make the

TABLE I
PROBES CHARACTERISTICS

Name	Customers	Technology	CC
TP	10,000	ADSL 0.5/6Mb/s	PL
MT	4,000	ADSL 0.5/5Mb/s	HU
IT-ADSL	15,000	ADSL 1/20Mb/s	IT
IT-FTTH	4,000	FTTH 10/10Mb/s	IT

swapping process faster in terms of bootstrap. In our work we show that churning due to channel swapping is not negligible in the case of P2P-TV systems, causing possible inter-overlay pollution.

In this paper, we present results collected passively monitoring actual users, running the application at home at their willing. We do not have control on any peer nor we alter the P2P-TV system under observation. We highlight that, to the best of our knowledge, this is the first work in characterization of overlays of a commercial P2P-TV system. Furthermore, we propose protocol agnostic methods that can be applied to other systems with similar characteristics.

III. EXPERIMENTAL SETUP

Our work is based on the data collected during monitoring experiments performed in the context of the Network Aware Peer-to-Peer Application under WIsE Network (NAPA-WINE) Project, funded by the EU in the Seventh Framework Programme [17]. Several traffic monitoring probes were installed to passively collect packet level traces from ISP network operational links. Traces were collected by running the Tstat [18] traffic analyzer on each probe machine. Through a Deep Packet Inspection (DPI) technique, Tstat [18] is instructed to identify P2P-TV traffic of prominent and popular commercial P2P-TV systems, namely TV-Ants, PPLive and SopCast. Packets belonging to those applications are dumped on output files to be later post-processed. DPI rules have been manually tuned and verified using both laboratories testbed, and experiments in the wild, showing very reliable results [19].

Probes are located on aggregation points (Point-of-Presence, PoP) of three European ISPs. Each vantage point monitors thousands of residential users, accessing the network via DSL or FTTH lines. The main characteristics of the 4 probes are summarized in Tab. I, which reports the name used throughout the paper, the approximate number of aggregated users, the access technology, maximum upload/download capacity offered to users and the country (CC) the probe is placed in. As it can be observed, the set of probes is very heterogeneous: they span over three different countries, using either ADSL or FTTH access technologies. Depending on the type of contract with the ISP and on the quality of the physical medium, ADSL technology offers the users different bitrate ranging from 2 to 20 Mb/s downstream and up to 1024 kb/s upstream. IT-FTTH users enjoy 10 Mb/s Ethernet based full-duplex connectivity. IT-ADSL and IT-FTTH probes are in the same ISP in Italy.

A. Terminology

As already mentioned, in P2P-TV systems three different graphs can be identified: i) in the “social graph”, users are

Fig. 1. P2P-TV Traffic at Different Time-scales in the TP PoP.

vertexes interconnected by edges representing their common interests; ii) in the “overlay graph” peers downloading the same content are vertexes and logical links used to exchange any kind of data are edges; iii) in the “distribution graph” a subset of peers are vertexes and logical links used to exchange actual video information are edges. Considering P2P terminology, the overlay graph is often referred to as the “swarm”, that is the set of peers watching the same video channel. Peers in the distribution graph are referred to as “contributing peers”, since they contribute to the content download. In the following, we will refer to the different graphs using the term *user*, *peer* and *contributing peer*, respectively. A user and a peer are identified by an IP address at the host running the application during an event¹. In the context where several users are using the same public IP address, the identification fails and users with the same address are aggregated in the same address. This issue cannot be solved by just looking at the ports used by the application flows in order to individuate the traffic generated by a specific user, since the protocol opens several UDP ports and randomly. However since the analyzed networks are residential this issue should be negligible.

Finally, considering the monitoring architecture, we further distinguish between *internal* and *external* users/peers, i.e., the peers ran by users inside the monitored PoP and ran by users in Internet. Similarly, we define *incoming* traffic (RX) the one flowing from external peers to internal peers, and *outgoing* traffic (TX) the one flowing in the opposite direction.

¹Being the event duration no longer than 3 hours, the impact of IP addresses reuse is marginal.

B. Overview of the dataset

We start by characterizing the usage and popularity of P2P-TV among users. Fig. 1 reports the P2P-TV average incoming bitrate versus time, for different timescales observed at the TP vantage point. Results are qualitatively similar in other PoPs, whose results not shown here for the sake of brevity. On average, the traffic generated by these applications is marginal. However, the *burstiness* of traffic reflects P2P-TV usage that is concentrated during short periods of time when the amount of generated traffic can reach very high and possibly disruptive peaks. Observe that the volume of traffic generated by P2P-TV typically coincides with the transmission of popular sport events, e.g. UEFA Champions League during Wednesday and Thursday or Premier League (England First Division) on Saturday and Sunday. Less than one hundred users is typically running the P2P-TV application at the same time, corresponding to less than 0.5% of customers connected to the monitored PoP. Still, the download bitrate often exceeds 15% of total PoP incoming traffic during those events. For the sake of comparison, consider that P2P-TV bitrate during peaks is larger than the aggregated YouTube bitrate consumed by customers in the same network. This “bursty” user behavior, which can be pretty difficult to handle, is also very different from normal TV and IPTV usage pattern, that is typically smoother and more evenly distributed during the whole day. Notice also the abrupt drop of traffic in the bottom plot that happens after 20:30, i.e., after the event ends. This hints that flash crowd phenomena are not negligible in P2P-TV systems, as we discuss in the next sections.

During the experiments, we compared the popularity of the three P2P-TV applications we detect (PPLive, TV-Ants and SopCast), and we noticed that SopCast is by large the most popular one. This holds true in all monitored networks. Therefore, in the following, we restrict our analysis to some of the largest traces of SopCast traffic. Since SopCast adopts a proprietary protocol and relies on encryption mechanisms, we avoid cumbersome reverse engineering of the SopCast algorithms and protocols. Instead, we identify both the overlay and distribution graph properties by devising simple methodologies that can be leveraged to study other P2P applications too.

IV. USER GRAPH PROPERTIES

In this section we investigate the SopCast user habits aiming at investigating their impact on the P2P application.

A. Swarm identification methodology

We first would like to identify how many different channels are actually watched by users inside the PoP and how many users watch the same channel at the same time. We verified by running some testbed experiments that SopCast peers form different swarms for different TV channels, e.g., peers of users watching two different channels belong to two different overlay graphs. Given any two internal peers, we count the number of common neighbors; then, we group those peers that have similar neighbors claiming they belong to the same swarm.

Fig. 2. Example of swarming matrix. Darker blocks refer to peers in the same swarm.

More in detail, the methodology is the following. Let a and b denote two internal peers and let $P(a)$ be the set of peers *contacted* by a , i.e., peers which a sent a packet to. The number of common peers between a and b is then $C(a, b) = |P(a) \cap P(b)|$, where $|\cdot|$ is the cardinality operator. We define then the common peer matrix M , as a matrix in which element (i, j) is $M_{ij} = C(i, j)$.

Let V_a be the vector of common peers of a with all other peers, i.e., the a -th row of M . Denote by V_a^T the transposed of V_a , i.e., the a -th column of M . The product,

$$S(a, b) = 2 \frac{V_a V_b^T}{V_a V_a^T + V_b V_b^T} \quad (1)$$

is a measure of the similarity between the neighborhoods of a and b . By iteratively sorting the list of peers, moving closer to each other those with larger similarity, we obtain the *swarming matrix*, i.e., an ordered common peer matrix M' that depicts in a clear way how peers are grouped. For instance, Fig. 2 shows the swarming matrix obtained from a 3 hour trace. It shows that several swarms were active during that event; the largest swarm includes peers with identifier from 0 to 75. The second largest swarm includes peers [103, 128], and so on. This, combined with the fact that most of the users are active in correspondence to sport events, shows that users watch different channels during the same event: they watch to either two concurrent soccer matches or the same match that is available from different channels. Notice that users can change channel during an event, as testified by peer 124 that has a strong similarity with both swarms. We will investigate this in the Sec. IV-B.

Tab. II summarizes the set of largest swarms we will focus on in the remaining of the paper. The table reports the Swarm ID (used in the paper to refer to a given swarm), the overall number of internal peers, the estimated swarm size considering the swarm size as stable and the vantage point in which peers were observed. The last column reports the percentage of external peers that belongs to the same Autonomous System (AS) of the considered PL ISP, which is one of the largest ISP

TABLE II
LIST OF THE LARGEST SWARMS

Swarm ID	Internal Peers	Stable Swarm Size Estimation	CC	PL AS %
0	35	3k–3.5k	PL	32
1	30	4k–5k	PL	31
2	50	4k–5k	PL	28
3	30	7k–9k	PL	27
4	40	2k–3k	PL	27
5	70	10k–12k	PL	25
6	65	12k–13k	PL	23
7	20	11k–14k	PL	21
8	80	6.5k–7k	PL	19
9	10	9k–12k	PL	10
10	5	9k–11k	PL	9
11	10	13k–15k	PL	9
12	10	12k–14k	PL	8
13	10	4k–4.5k	PL	8
14	15	10k–12k	PL	6
15	15	3k–4k	PL	5
16	25	18k–20k	PL	4
17	10	9k–12k	PL	3
18	10	15k–18k	PL	2
19	10	15k–24k	PL	2
11	1	–	HU	N/A
14	1	–	IT	N/A

in Poland. Swarms have been sorted in decreasing percentage of PL AS peers. Note that the bottom two rows refer to two peers connected to swarm 11 and 14 that were present in the HU and IT-FTTH dataset, respectively.

B. Swarm size

Few consideration hold when observing Tab. II: first, the number of internal peers in the monitored PL PoP is smaller than 100. Considering that the total number of customers we monitored is around 10,000, the popularity of SopCast is below 1%. The swarm size, i.e., the total number of peers participating to the swarm, is typically smaller than 10,000 peers. At last, the fraction of peers that are inside the PL AS varies a lot. This suggests that there are some swarms/channels which are very popular among PL AS customers and in Poland in general, while other swarms/channels are less popular. Still they attract a quite large number of users. This means that there is a high localization of peers inside the same AS or country which is naturally induced by cultural preferences of users.

C. User churning and flash crowd

We have already observed in Fig. 1 that users run SopCast only when they are interested in watching a particular event, so that a flash crowd effect is present when an event starts. Users then keep running SopCast for the whole event duration, and then suddenly stop using it at the event end. This is confirmed by Fig. 3 which reports the Cumulative Distribution Function (CDF) of the internal user lifetime. Both separate and aggregate CDFs are reported. Results show that users lifetime is very similar for different events, and it is rather long, e.g., 90% of users have a lifetime longer than 30 minutes.

Fig. 3. User lifetime CDF; the thick line refers to the aggregate statistic.

As highlighted by the horizontal lines, 50% of users have a lifetime within 90 min and 130 min, which corresponds to the typical duration of a soccer event. Only less than 10% of users run SopCast for more than 150 min. Do the users stay connected to the same channel or do they frequently change channel?

To answer this question, Fig. 4 shows the churning percentage over time. Both peers that change channel and peers that close the application are considered. Two different data sets are plotted in top and bottom plots. Results are obtained by computing the swarming matrix for different time slots. Intuitively, the swarm similarity of a peer that changes channel changes in time. In detail, we compute $M(t_n)$ every time period t_n of 5 min; then we compare $M(t_n)$ and $M(t_{n-1})$ to count peers entering or leaving a swarm and peers swapping swarm. Finally, the churning percentage is computed with respect to the number of active peers at time t_n .

Fig. 4 shows that the number of users changing channel is not negligible, e.g., around 8-10% of users changes channel every 5 min. Interestingly, the percentage of peers that change channel is higher at the beginning, when possibly users are seeking for a good swarm/channel to follow the event they are interested in. The churning percentage of users leaving the system is small, but it suddenly increases at the event end when more than 50% of users leave the system at the same time.

These results highlight the “human-factor” implication when designing a P2P-TV system. Both flash crowd and sudden peer departures are not negligible, so that algorithms must explicitly deal with them.

V. OVERLAY GRAPH PROPERTIES

We now turn our attention to the properties of the overlay graph. As well known, a generic P2P system implements mechanisms to discover peers so that signaling information is continuously exchanged to find new peers. Most of P2P-TV systems (SopCast included) implement a peer discovery algorithm based on gossiping protocols [20]. In SopCast a continuous discovery process is carried out by peers that look for new peers at a practically constant rate. In the overlay

(a) Event A

(b) Event B

Fig. 4. Percentage of peers changing channel and leaving the system during two soccer matches.

Fig. 5. Number of peers contacted every 60s by internal peers.

graph, an edge is created every time a peer exchanges information with another peer. To assess what are the properties of the SopCast overlay graph, we observe the peer discovery process.

A. Peer discovery rate

We start by observing the peer discovery rate, i.e., the number of neighbors a peer contacts in a period of time ΔT . The discovery rate depends on the rate with which the peer is contacted by or contacts other nodes. Fig. 5 shows an example of measurement for peers in swarm 14. We separately plot the average discovery rate of peers in the TP and in the IT-FTTH data sets. We choose $\Delta T = 60s$. Let us first focus on the TP peers. The discovery rate is pretty much constant over time, so that about 170 peers are contacted every 60s. Moreover, the standard deviation of the measurement computed among the different internal peers is very small. This highlights that all peers in the TP data set perform the same kind of discovery process. In contrast, the discovery rate of the IT-FTTH peer is much larger than the TP peer rate. It also shows a significant variation during peer lifetime. Similar results are observed when comparing low and high upload capacity peers: the formers exhibit smaller discovery rate than the latter. This suggests that SopCast implements some algorithms to exploit the upload bandwidth of high capacity peers.

To corroborate this intuition we performed several test-bed experiments in our Campus network. We started two peers

Fig. 6. External peer spatial distribution.

that watch the same channel at the same time. Both peers are connected to the Internet via the same router which limits the upload rate to 256kb/s and 32Mb/s, respectively. We then observe the evolution over time of the discovery rate on the two peers. Results are identical to Fig. 5: the peer with higher upload capacity exhibits a much higher peer discovery rate. Therefore, we conclude that the discovery rate depends on the peer upload capacity.

B. Bias in the discovery process

Let us now investigate if there is any preference based on the peer geo-position, distance or any other property. Fig. 6 shows the geographical breakdown of nodes contacted by different internal peers during their whole lifetime. Left and right groups of bars refer to swarm 11 and 14, respectively. For each swarm, we report the breakdown for i) two internal peers selected at random from the PL dataset, ii) all peers in the PL dataset, and iii) one peer in the HU and IT-FTTH dataset. Results show that there is no statistically significant difference, even if peers are located in different countries (PL or HU for swarm 11, PL or IT for swarm 14) or are connected through different access technologies (PL-ADSL vs IT-FTTH for swarm 14). However, the breakdown changes in the two events. This suggests that the SopCast peer discovery mechanism is not driven by any preference related to any peer property but it reflects only the natural distribution of users around the world. That is, the peer discovery mechanism follows a random process in which the probability of contacting (or being contacted) by a peer is independent from other peers. This is a very robust choice which allows SopCast to deal with the high churning rate we have seen in the previous section.

C. Swarm size estimation

Leveraging on the SopCast peer discovery process, in the following we exploit a simple model to estimate the swarm size. Let us consider independent observation periods of duration ΔT . We monitor internal peers watching the same channel from the same network, i.e., identical and independent peers. The discovery process they perform can be modeled as a random walk at constant rate where each peer discovers a random set of external peers every ΔT . Given any two internal

Fig. 7. Common peer distribution as observed during a $\Delta T = 360$ s time interval.

Fig. 8. Estimated overlay graph size and total number of discovered peers. $\Delta T = 360$ s.

peers a and b , the amount of common external peers $C(a, b)$ they discover follows a Poisson distribution with mean value determined by the degree $|P(a)|$ and $|P(b)|$ of a and b :

$$C(a, b) = \frac{|P(a)||P(b)|}{N} \quad (2)$$

Given a swarm and an observation time ΔT , we measure $C(a, b)$, $P(a)$, $P(b)$ for each possible pair of internal peers a, b . From (2) it is possible to estimate the total number of peers in the swarm, let this estimation be denoted by $\hat{N}(a, b)$. Since there are several internal peer pairs, we compute the average and standard deviation among the estimations $\hat{N} = E[\hat{N}(a, b)]$ and $\sigma_N = std[\hat{N}(a, b)]$. $\hat{N}$ is the estimated swarm size, i.e., the number of vertexes in the overlay graph.

The sampling time ΔT plays a key role, since it defines the result of the discovery process: on the one hand, ΔT should be large enough to allow a correct estimation of $C(a, b)$; on the other hand, ΔT should be small to minimize the impact of peers churning and channel swapping. We performed sensitivity analysis on the impact of ΔT , finding that a good trade off is obtained for ΔT larger than 2 min and shorter than 10 min. In the following, we choose $\Delta T = 360$ s. Fig. 7 reports the CDF of $C(a, b)$ for all the pairs a, b of internal peers in a swarm during a single observation period $\Delta T = 360$ s. $C(a, b)$ closely follows a Poisson distribution with mean value 36, confirming that (2) offers a good approximation.

Fig. 8 depicts the estimated average and standard deviation of the estimated swarm size. We compare it against the amount of unique peers M discovered by all internal peers during the same observation time ΔT . We consider swarm 6 which was previously reported in bottom plot of Fig. 1. $\hat{N}$ quickly grows at the beginning when the flash crowd phenomenon starts. During the event, $\hat{N}$ is then stable since the swarm population remains constant. Finally, at the event end we can observe the abrupt departure of peers. This confirms the impact of user habits on P2P-TV usage we already noticed

in Fig. 1. Comparing $\hat{N}$ with M , we observe that the latter does not provide a good estimation of the swarm size during the initial transient. In regime situation, M is comparable with $\hat{N}$ because more than 60 internal peers are present. The aggregated discovery process they perform allows to practically find all peers in the swarm in 360 s only. However, when the number of internal peers is small, M provides a lower bound to N .

VI. DISTRIBUTION GRAPH PROPERTIES

Moving to the characterization of the distribution graph, we investigate if there is any preference in choosing from which contributing peer video content is downloaded. We first investigate if any preference is given to peers within the same Autonomous System. This is a timely topic being investigated for both P2P and P2P-TV systems [17], [21]; indeed, preference to peers in the same AS would allow ISP to reduce traffic peering costs.

A. Preference to peers in the same AS

For each channel, we partition the set of all external peers according to the AS they belong to. By focusing on a particular AS, we then compare the fraction of external peers belonging to the AS with the fraction of data transmitted/received from internal peers. A difference in these fractions is interpreted as an indication of some form of preference given to internal peers. We show results for the analysis made on some prominent ASes.

The results of Fig. 9(a) are derived focusing on peers that belong to the same TP AS we are monitoring. Swarms are sorted in decreasing values of the fraction of peers in the TP AS, as in Tab. II. Consider transmitted data: internal peers are likely to transmit a big portion of data to peers which are located in the same AS. For example, for Swarm 0, more than 50% of transmitted traffic goes to only the 32% of peers. Considering instead received data, we observe the opposite: in Swarm 0, 32% of peers can only provide less than 18% of traffic to internal peers. Indeed, customers of the TP ISP are offered ADSL lines and have small upload capacity. Since low-capacity peers can only upload a fraction of the traffic they download, the rest of the traffic has to come from other ASes, in which high upload capacity peers are present.

B. Preference to peers in the same country

To quantify if there is a preference at the country level, we consider another AS located in Poland, named the PL2 AS, whose customers are offered ADSL access lines. Results are reported in Fig. 9(b). Comparing the fraction of peers in PL2 and the fraction of traffic transmitted to PL2, we notice that this time the two curves are clearly similar. This shows that peers in TP send an amount of traffic to peers in PL2 which is proportional to the number of contacted PL2 peers, i.e., a peer in PL2 is selected with the same probability of peers in any other AS. Considering the amount of traffic received from PL2 peers, we observe that it follows exactly the same trend seen in Fig. 9(a). This is expected, since PL2 customers have

the same upload capacity as TP customers, so that they can provide only a fraction of the traffic. Repeating the analysis considering other countries ADSL providers, similar results are obtained. These results suggest that there is no preference based on the country the peers belong to.

C. Preference to peers with large capacity

Figs. 9(c) and 9(d) refer to results considering two ASes with a large fraction of high upload capacity peers. The first one is an educational Polish AS - PL3, the second one is a commercial AS in Russia - RU1. It is clear that when these high capacity peers are present, internal peers download from them a large fraction of traffic. For example, in Swarm 1, peers from PL3 contribute to more than 10% of downloaded traffic despite they represent only 0.3% of external peers. For the swarms on the left, i.e., popular swarm channels in Poland, it is easy to find peers with high capacity in PL3 AS. On the contrary, for rightmost swarms that are unpopular channels in Poland, traffic is received from any AS, provided that peers there have large upload capacity. This is the case of RU1.

We can, thus, conclude that SopCast peers tend to transmit traffic to peers in the same AS they belong to. But if bandwidth availability in the AS is not sufficient, peers fetch the content from any high bandwidth peer, whatever AS it belongs to.

D. Characterization of peer upload capacity

The uplink capacity of peers is a key characteristic to understand the feasibility and scalability of P2P-TV services. Indeed, the total available upload capacity offered by peers must be equal or larger to the total download capacity dictated by the video rate and the number of peers in a swarm. Today, ADSL peers lack of upload capacity to sustain the video content diffusion, and high upload capacity peers are required to supply it. What is then the access capacity of contributing peers in SopCast overlay? To the best of our knowledge, no previous work has performed this characterization.

To estimate path capacity we rely on the so called *Probe Gap Model* (PGM) [22]. The intuition is to observe the minimum Inter-Packet Gap (IPG) of packets received by a given node, which, in case of packets sent back-to-back at the source, provides an estimation of the available capacity C along the path. Several assumptions must be verified to obtain reliable measurements: i) packets must be sent back-to-back by the transmitter; ii) measurement must be repeated several times to estimate the minimum IPG; iii) time-stamping at the receiver must be very accurate, and packets must not be artificially delayed at the receiver by some buffering. To this extent, large packets are preferred, so that the packet transmission time is large (a 40 Bytes long packet lasts only 320 ns on a 1 Gb/s link).

For each swarm, we then estimate the upload capacity of all external contributing peers from which at least 10 packets larger than 1300B were received. Those are packets containing video data, sent in bursts by the transmitter. Given the preference we have seen to download content from high capacity peers, we expect that the estimation can be performed

Fig. 9. TP Aggregated AS Contribution to Swarms

only for a limited subset of contacted peers, with a bias to include higher capacity peers. Fig. 10 reports the estimated capacity versus normalized peer ID; peers are ordered in decreasing estimated capacity. Log scale is used on the y-axis. The aggregated distribution is also shown. For all swarms, it emerges that about 10% of peers have 1 Gb/s upload capacity,

Fig. 10. Contributing peer capacity for every swarm; the thickest line refers to the aggregate statistic.

while [60,80]% of them have a capacity smaller than 2 Mb/s, with more than 50% of peers having less than 1 Mb/s capacity. This reflects the intuition that nowadays P2P-TV systems rely on high capacity peers, which act as amplifiers, and help in re-distributing the video stream to several peers. By trying to identify the heavy contributing peers, we verified that the majority of them are actually hosts at University campuses around the world.

VII. CONCLUSIONS AND IMPLICATION FOR TRAFFIC LOCALIZATION

In this paper, we have presented a methodology to investigate the behavior of P2P-TV applications and, in particular, to study the characteristics of the fundamental graphs created by these applications: the social network of the users, the overlay network of the peers, the distribution graph through which video content is delivered to users. We apply the methodology to SopCast traffic passively collected in operative networks of some ISPs in Europe.

Results about users' habits show that, today in Europe, the users tend to run the P2P-TV applications as a cheap alternative to traditional television broadcasting systems, in correspondence of special events (such as sport events). The user behavior is therefore extremely bursty, with relevant flash crowd phenomena and sudden peer departures.

Through the proposed methodology, we could observe that the SopCast peer discovery process that creates the overlay network is based on random mechanisms that do not exploit any relevant information about the peer characteristics, so that the overlay results biased only by the cultural preferences of the users. The distribution mechanism, instead, tends to favor the choice of peers in the same AS. However, whenever the upload bandwidth provided by peers in the same AS is not enough to distribute the video content, high upload capacity peers are selected from other ASs as distributing peers.

This variety of results and observations show that the proposed methodology can be effectively used to investigate, through passive traffic analysis, the effect of internal mechanisms of P2P applications, even when their design is proprietary and unknown.

REFERENCES

- [1] E.K. Lua, J. Crowcroft, M. Pias, and R. Sharma. A survey and comparison of peer-to-peer overlay network schemes. *IEEE Communications Surveys and Tutorials*, 7:72–93, 2005.
- [2] M. Ripeanu, I. Foster, and A. Iamnitchi. Mapping the gnutella network: Properties of large-scale peer-to-peer systems and implications for system design. *Arxiv preprint cs/0209028*, 2002.
- [3] S. Saroiu, K.P. Gummadi, and S.D. Gribble. Measuring and analyzing the characteristics of Napster and Gnutella hosts. *Multimedia systems*, 9(2):170–184, 2003.
- [4] J. Liang, R. Kumar, and K.W. Ross. The kazaa overlay: A measurement study. *Computer Networks Journal (Elsevier)*, 49(6), 2005.
- [5] D. Stutzbach and R. Rejaie. Understanding churn in peer-to-peer networks. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*, pages 189–202. ACM, 2006.
- [6] D. Stutzbach, R. Rejaie, and S. Sen. Characterizing unstructured overlay topologies in modern P2P file-sharing systems. *Networking, IEEE/ACM Transactions on*, 16(2):267–280, 2008.
- [7] S. Spoto, R. Gaeta, M. Grangetto, and M. Sereno. Analysis of PPLive through active and passive measurements. In *Parallel & Distributed Processing, 2009. IPDPS 2009. IEEE International Symposium on*, pages 1–7. IEEE, 2009.
- [8] D. Stutzbach, R. Rejaie, N. Duffield, S. Sen, and W. Willinger. On unbiased sampling for unstructured peer-to-peer networks. *IEEE/ACM Transactions on Networking (TON)*, 17(2):377–390, 2009.
- [9] B. Yang and H. Garcia-Molina. Efficient search in peer-to-peer networks. In *Proceedings of the International Conference on Distributed Computing Systems (ICDCS)*. Citeseer, 2002.
- [10] S. Guha, N. Daswani, and R. Jain. An experimental study of the skype peer-to-peer voip system. In *Proc. of IPTPS*, volume 6. Citeseer, 2006.
- [11] Y. Yu, D. Liu, J. Li, and C. Shen. Traffic identification and overlay measurement of Skype. In *Computational Intelligence and Security, 2006 International Conference on*, volume 2, pages 1043–1048. IEEE, 2007.
- [12] Dario Rossi, Marco Mellia, and Michela Meo. Understanding skype signaling. *Computer Networks*, 53(2):130 – 140, 2009. QoS Aspects in Next-Generation Networks.
- [13] B. Li, S. Xie, Y. Qu, G.Y. Keung, C. Lin, J. Liu, and X. Zhang. Inside the new coolstreaming: Principles, measurements and performance implications. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 1031–1039. IEEE, 2008.
- [14] C. Wu, B. Li, and S. Zhao. Multi-channel live P2P streaming: Refocusing on servers. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 1355–1363. IEEE, 2008.
- [15] X. Hei, C. Liang, J. Liang, Y. Liu, and K.W. Ross. A measurement study of a large-scale P2P IPTV system. *Multimedia, IEEE Transactions on*, 9(8):1672–1687, 2007.
- [16] A. Kermarrec, E. Le Merrer, Y. Liu, and G. Simon. Surfing Peer-to-Peer IPTV: Distributed Channel Switching. *Euro-Par 2009 Parallel Processing*, pages 574–586, 2009.
- [17] E. Leonardi, M. Mellia, A. Horvath, L. Muscariello, S. Niccolini, and D. Rossi. Building a cooperative P2P-TV application over a wise network: the approach of the European FP-7 strep NAPA-WINE. *Communications Magazine, IEEE*, 46(4):20–22, 2008.
- [18] M. Mellia, R. Lo Cigno, and F. Neri. Measuring IP and TCP behavior on edge nodes with Tstat. *Computer Networks*, 47(1):1–21, 2005.
- [19] N. Cascarano, F. Risso, A. Este, F. Gringoli, L. Salgarelli, A. Finamore, and M. Mellia. Comparing P2PTV Traffic Classifiers. In *Communications (ICC), 2010 IEEE International Conference on*, pages 1–6. IEEE, 2010.
- [20] D. Ciullo, M.A. Garcia, A. Horvath, E. Leonardi, M. Mellia, D. Rossi, M. Telek, and P. Veglia. Network Awareness of P2P Live Streaming Applications: A Measurement Study. *Multimedia, IEEE Transactions on*, 12(1):54–63, 2009.
- [21] H. Xie, Y.R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz. P4P: Provider portal for (P2P) applications. In *Proc. of ACM SIGCOMM*, 2008.
- [22] L. Lao, C. Dovrolis, and MY Sanadidi. The probe gap model can underestimate the available bandwidth of multihop paths. *ACM SIGCOMM Computer Communication Review*, 36(5):29–34, 2006.