

Network Awareness of P2P Live Streaming Applications

Original

Network Awareness of P2P Live Streaming Applications / Ciullo, Delia; GARCIA DA ROCHA NETA, MARIA ANTONIETA; Horvat, A.; Leonardi, Emilio; Mellia, Marco; Rossi, D.; Telek, M.; Veglia, P.. - STAMPA. - (2009), pp. 1-7. (Intervento presentato al convegno Parallel & Distributed Processing, 2009. IPDPS 2009. IEEE International Symposium on tenutosi a Rome, Italy nel May 29, 2009) [10.1109/IPDPS.2009.5160953].

Availability:

This version is available at: 11583/1936662 since:

Publisher:

IEEE

Published

DOI:10.1109/IPDPS.2009.5160953

Terms of use:

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)

Network Awareness of P2P Live Streaming Applications

Delia Ciullo*, Maria Antonietta Garcia*, Ákos Horváth[†], Emilio Leonardi*, Marco Mellia*, Dario Rossi[‡],
Miklós Telek[†] and Paolo Veglia[‡]

*Politecnico di Torino, lastname@tlc.polito.it

[†]Budapest University of Technology and Economics, lastname@hit.bme.hu

[‡]TELECOM-ParisTech, firstname.lastname@enst.fr

Abstract—Early P2P-TV systems have already attracted millions of users, and many new commercial solutions are entering this market. Little information is however available about how these systems work. In this paper we present large scale sets of experiments to compare three of the most successful P2P-TV systems, namely PPLive, SopCast and TVAnts.

Our goal is to assess what level of “network awareness” has been embedded in the applications, i.e., what parameters mainly drive the peer selection and data exchange. By using a general framework that can be extended to other systems and metrics, we show that all applications largely base their choices on the peer bandwidth, i.e., they prefer high-bandwidth users, which is rather intuitive. Moreover, TVAnts and PPLive exhibits also a preference to exchange data among peers in the same Autonomous System the peer belongs to.

However, no evidence about preference versus peers in the same subnet or that are closer to the considered peer emerges. We believe that next-generation P2P live streaming applications definitively need to improve the level of network-awareness, so to better localize the traffic in the network and thus increase their network-friendliness as well.

I. INTRODUCTION AND MOTIVATIONS

P2P-TV systems are candidates for becoming the next Internet killer applications as testified by the growing success of commercial systems such as PPLive [1], SopCast [2] and TVAnts [3], which have already attracted an audience up to several millions of users and drawn the attention of Telecom operators and Service providers.

However, little information is available about the internal algorithms and protocols used by these applications, which are proprietary and closed. Therefore, the very same potentialities of P2P-TV systems constitute a worry for network carriers since the traffic they generate may grow without control, causing a degradation of the quality of service perceived by Internet users or even the network collapse (beside the consequent failure of the P2P-TV service itself). This has motivated further research [4], [5], [6], [7], [8], [9], [10], [11], [12], aimed at understanding these systems through on-field measurements. Most work, though, focuses on the study of a single P2P-TV system [4], [5], [6], [7], [8], [9], [10]. For instance, active crawling methodologies are used to investigate PPLive [4], CoolStreaming [5] and UUSee [6]. The downside of this approach is that it relies on heavy reverse engineering, hardly extendable to characterize all the possible P2P-TV applications. Other works instead focus on

very specific aspects of P2P streaming systems: e.g., node degree of popular versus unpopular channels [7] and node stability [8], while quality of service is of concern in [9], [10].

While all the above are valuable works, the risk is that observation gathered from a single system cannot be generalized. However, to date, very few measurement studies exist that compare different systems [11], [12], which are closer to our work. Considering PPLive and SopCast, [11] limitedly focuses on the temporal evolution of different metrics (e.g., like transmitted/received bytes, number of parents and children, etc). Authors in [12] instead compare PPLive, PPStream, SopCast and TVAnts, by means of flow-level scatter plots of mean packet size versus flow duration and data rate of the top-10 contributors versus the overall download rate.

Therefore, despite the above works pose a first important milestone, a more systematical analysis is needed to provide a deeper understanding of the impact that a large deployment of general P2P-TV services may have on the Internet. This is precisely one of the goals of the recently funded Project called “Network-Aware P2P-TV Application over Wise Networks” (NAPA-WINE) [13]. In this paper, we aim at providing an assessment of level of “network awareness” embedded in the currently deployed systems; that is the capacity of the P2P application to discovery some properties of the underlying network and to exploit them to optimize its decisions. To determine if a P2P-TV application is “network-aware” is equivalent to answer the questions: does it randomly select peers? Or does it preferentially look for high-bandwidth peers? Is the traffic confined within the same Autonomous System the peer belongs to? Does it preferentially download traffic from nearby nodes? Hence we defined a methodology, and focusing on SopCast, PPLive and TVAnts, we inferred their level of network awareness from the characteristics of the traffic they generate. We believe our work to be novel in two main aspects. The first is the aim, as we focus on a systematic exploration of the metrics, if any, that drive the P2P streaming in different systems. A second important difference lies on the scale of the testbed, which in our case involves more than 40 vantage points scattered across European countries and it is representative of very different network setups. Finally, the presented results underline the current need for the development of new and network friendly P2P-TV systems, an interesting topic deserving future research.

TABLE I
SUMMARY OF THE HOSTS, SITES, COUNTRIES (CC), AUTONOMOUS
SYSTEMS (AS) AND ACCESS TYPES OF THE PEERS INVOLVED IN THE
EXPERIMENTS.

Host	Site	CC	AS	Access	Nat	FW
1-4	BME	HU	AS1	high-bw	-	-
5			ASx	DSL 6/0.512	-	-
1-9	PolITO	IT	AS2	high-bw	-	-
10			ASx	DSL 4/0.384	-	-
11-12			ASx	DSL 8/0.384	Y	-
1-4	MT	HU	AS3	high-bw	-	-
1-3	FFT	FR	AS5	high-bw	-	-
1-4	ENST	FR	AS4	high-bw	-	Y
5			ASx	DSL 22/1.8	Y	-
1-5	UniTN	IT	AS2	high-bw	-	-
6-7				high-bw	Y	-
8			ASx	DSL 2.5/0.384	Y	Y
1-8	WUT	PL	AS6	high-bw	-	-
9			ASx	CATV 6/0.512	-	-

II. EXPERIMENTAL SETUP

The results of this paper are based on a large testbed we setup, whose main features are summarized in Tab. I. Partners took part in the experiments by running P2P-TV clients on PCs connected either to the institution LAN, or to home networks having cable/DSL access. In more detail, the setup involved a total of 44 peers, including 37 PCs from 7 different industrial/academic sites, and 7 home PCs. Probes are distributed over four countries, and connected to 6 different Autonomous Systems, while home PCs are connected to 7 other ASs and ISPs. Therefore, the setup is representative of a significant number of different network environments.

We considered three different applications, namely PPLive, SopCast and TVAnts and we performed several 1-hour long experiments during April 2008, where partners watched the same channel at the same time and collected packet-level traces. Since P2P-TV application are mostly popular in Asian countries, we tuned each application to CCTV-1 channel during China peak hours [4]. In all cases, the nominal stream rate was 384kbps, Windows Media 9 Encoder was used, and the video quality perceived by partners was not remarkably different across systems. Results reported in this paper refer to more than 120 hours of experiments, corresponding to more than 140.000.000 collected packets. Collected traces are also made available to the research community from the NAPA-WINE website upon request.

A short summary of the experiments is given in Tab. II, which reports the mean and maximum values, as seen by NAPA-WINE peers, of i) the stream rates (in upload and download directions), ii) the number of peers and iii) the number of *contributing* peers for the different applications. By contributing peers, we denote peers with whom some video segment has been exchanged, either in upload (TX) or in download (RX), and that are identified according to the heuristic in [14], which we verified to give accurate and conservative results. A *significant heterogeneity* across systems emerges from the data: for instance, despite the received stream rate is similar across systems (PPLive one being larger in reason of a larger signaling overhead tied to the number




Fig. 1. Geographical breakdown of the number of peers, transmitted and received bytes.

of peers contacted), the number of contributing peers varies widely from an average of 58 peers in the TVAnts and 391 for PPLive. Also, notice that PPLive potentially uses a significant amount of peers' resources, as the average (maximum) upload rate reaches about 3 Mbps (12 Mbps) for peers in our testbed.

To complete a brief overview of the experiments, Figure 1 shows the geographical distribution of the number of contacted peers (#) and the amount of received (RX) and transmitted (TX) bytes. The labels on the bars refer to China (CN) and the four countries in which experiments were performed, with the rest of the Countries labeled "*". Percentages are expressed over the total number of observed peers, which amounts to 4057 for SopCast, 550 for TVAnts, and 181729 for PPLive. As expected, China is the predominant country, though it is easy to gather that a non negligible fraction of the data is exchanged within European countries: this hints to the existence of a bias in the peer selection, which we will dig more rigorously in the following section.

III. PEER SELECTION: METHODOLOGY AND ANALYSIS

As previously stated, our aim is to develop a rigorous framework to unveil the "network-awareness" exhibited by P2P-TV applications, i.e., which network parameters current P2P-TV systems take into account when distributing the stream. Among all the possible properties that can be used, we have to pick those which: (a) can be measured quite simply from the application's point of view, (b) can be drawn off-line from packet traces. For example, an application can not discover path costs unless this information is given by the AS; on the contrary, it is straightforward to actively measure RTT between two end-points but it is very hard to infer it passively. Taking into account these requirements, as network properties, we consider:

- **BW**: the peer access capacity
- **AS**: the Autonomous System a peer is located in
- **CC**: the Country a peer belongs to
- **NET**: the subnetwork a peer belongs to

TABLE II

SUMMARY OF EXPERIMENTS: MEAN AND MAXIMUM VALUES OF STREAM RATES, OVERALL NUMBER OF PEERS AND NUMBER OF CONTRIBUTING PEERS FOR THE DIFFERENT APPLICATIONS

App.	Stream RX [kbps]		Stream TX [kbps]		All peers		Contrib. RX		Contrib. TX	
	Mean	Max	Mean	Max	Mean	Max	Mean	Max	Mean	Max
PPLive	552	934	3384	11818	23101	39797	391	841	1025	2570
SopCast	449	542	293	1070	776	1233	139	229	152	243
TVAnts	419	478	464	1001	229	270	58	90	75	118

- **HOP**: the distance measured in hop number between two peers

In the following, we propose a general framework to highlight the possible preferences.

A. Framework Definition

Let $p \in \mathcal{W}$ denote a peer that belongs to the NAPA-WINE probe set \mathcal{W} . Let $\mathcal{P}(p)$ denote set of peers p exchange data with: with respect to our previous terminology, in the following we will in other words restrain our attention to the set of “contributors”. Considering contributing peers, let $\mathcal{U}(p)$ denote the subset of peers which p is uploading video content, and $\mathcal{D}(p)$ the subset which p is downloading from. $\mathcal{U}(p) \cap \mathcal{D}(p)$ represents the set of peers that are both downloading/uploading video from/to p .

Let $e \in \mathcal{P}(p)$ be an arbitrary peer that exchanges traffic with p . Denote by $B(p, e)$ the amount of bytes transmitted from p to e , so that $B(e, p)$ represents the amount of bytes received by p from e .

Consider now a generic network parameter $X(\cdot)$, and denote by $X(p, e) \in X$ the observed value of $X(\cdot)$ for peer pair (p, e) . X is the support of metric $X(\cdot)$. We partition $\mathcal{P}(p)$ into two classes based on $X(p, e)$, such that one class should intuitively be preferred from the P2P application (e.g., high-bandwidth vs low-bandwidth peers). More formally, we partition the support X into two disjoint sets: a preferred set X_P and its complement $X_{\bar{P}}$, such that $X_P \cup X_{\bar{P}} = X$ and $X_P \cap X_{\bar{P}} = \emptyset$.

For ease of notation, let $\mathbf{1}_P(p, e)$ be an identity function which is equal to 1 whether $X(p, e) \in X_P$ and 0 otherwise; similarly, $\mathbf{1}_{\bar{P}}(p, e) = 1 - \mathbf{1}_P(p, e)$. Without loss of generality, let us focus on the upload traffic of a NAPA-WINE probe $p \in \mathcal{W}$, and further define¹:

$$Peer_{U|P}(p) = \sum_{e \in \mathcal{U}(p)} \mathbf{1}_P(p, e) \quad (1)$$

$$Byte_{U|P}(p) = \sum_{e \in \mathcal{U}(p)} \mathbf{1}_P(p, e) \cdot B(p, e) \quad (2)$$

$$Peer_{U|\bar{P}}(p) = \sum_{e \in \mathcal{U}(p)} (1 - \mathbf{1}_P(p, e)) \quad (3)$$

$$Byte_{U|\bar{P}}(p) = \sum_{e \in \mathcal{U}(p)} (1 - \mathbf{1}_P(p, e)) \cdot B(p, e) \quad (4)$$

The function $Peer_{U|P}(p)$ states the number of peers of which p is a contributor and which belongs to the preferential partition X_P . Similarly, $Byte_{U|P}(p)$ represents the total amount of

bytes uploaded from peer p to peers in its preferential partition X_P . Conversely, $Peer_{U|\bar{P}}(p)$ and $Byte_{U|\bar{P}}(p)$ represent the number of peers and bytes to which p is uploading despite they belong to the non-preferential partition $X_{\bar{P}}$. Considering now the whole set of NAPA-WINE probes, we define the total amount of peers and bytes as:

$$Peer_{U|P} = \sum_{p \in \mathcal{W}} Peer_{U|P}(p) \quad (5)$$

$$Byte_{U|P} = \sum_{p \in \mathcal{W}} Byte_{U|P}(p) \quad (6)$$

Similar definition holds for $Peer_{U|\bar{P}}$ and $Byte_{U|\bar{P}}$; notice that a peer e may be counted more than once, e.g., if it exchanges traffic with more than one NAPA-WINE peer.

Finally, we define the peer and byte preference as:

$$P_U = 100 \frac{Peer_{U|P}}{Peer_{U|P} + Peer_{U|\bar{P}}} \quad (7)$$

$$B_U = 100 \frac{Byte_{U|P}}{Byte_{U|P} + Byte_{U|\bar{P}}} \quad (8)$$

Intuitively, P_U expresses the chance that the peer selection mechanism favors the discovery and data exchange among peers belonging to the preferred partition X_P . Similarly, B_U quantifies the chance that any given byte is uploaded to peers belonging to the X_P class. Clearly, the greater P_U and B_U are, the greater the bias with respect to the preferential partition of metric X is. The advantage of using these simple metrics is that they allow a *direct and compact* comparison of different network properties and P2P systems, since they are neither sensitive to the unit of measure nor to the actual magnitude of the X metric.

Downlink metrics P_D and B_D can be defined by simply considering $e \in \mathcal{D}(p)$ in the previous derivation.

B. Preferential Partitions

As preferential classes, we consider the following:

- **BW**: $\mathbf{1}_P(e, p) = 1 \Leftrightarrow BW(e, p) > 10Mbps \Leftrightarrow \min IPG(e, b) < 1ms$, i.e., peer e is a high-bandwidth peer, as it can be inferred from the minimum inter packet-gap (IPG);
- **AS**: $\mathbf{1}_P(p, e) = 1 \Leftrightarrow AS(p) = AS(e)$, i.e., both peers are located in the same Autonomous System;
- **CC**: $\mathbf{1}_P(p, e) = 1 \Leftrightarrow CC(p) = CC(e)$, i.e., both peers are located in the same Country;
- **NET**: $\mathbf{1}_P(e, p) = 1 \Leftrightarrow HOP(e, p) = 0$, i.e., peers belongs to the same subnet;

¹Subscript U or D denotes the upload and download traffic respectively.

- **HOP:** $1_P(p, e) = 1 \Leftrightarrow HOP(e, p) < \text{median}[HOP]$, i.e., the number of hops between e and p is smaller than the median distance.

While for most properties the preferential set choice is straightforward, the BW and HOP cases require additional discussion. Considering HOP metric first, the hop count $HOP(e, p)$ has been evaluated as 128 minus the TTL of received packets, since 128 is the default TTL considering Windows O.S. As threshold to define two classes, we use the median of the distance distribution as threshold. Since the actual HOP median ranges from 18 to 20 depending on the application, we use a fixed threshold of 19 hops for all applications. This means that, approximately 50% of the peers falls in the preferential class, which includes the shorter paths.

Considering BW, we infer whether a peer e has an high-bandwidth path to p considering the minimum inter packet-gap (IPG) of the packet it sends to p . Since we are considering contributor traffic, a significant number of video chunks are sent by the transmitter. Being a chunk built of several packets, the source transmit them in a burst, so that they are sent as a train of packets. They can be then considered as several “packet-pairs”, that can be used to infer the bottleneck capacity. By measuring the minimum IPG, it is possible to easily classify a peer as a high- or low-bandwidth peer, using 1 ms threshold, which corresponds to the transmission time of a 1250 bytes packet over a 10Mbps link. Evidence of this is available in [14].

C. Preliminary Analysis and Issues

Given the black box approach based on passive measurement, several issues could undermine the significance of the results unless carefully dealt with. The first issue is that the NAPA-WINE probes *self-induce a bias* during the experiments. Recall that among NAPA-WINE peers there are several high-bandwidth peers, located in Europe only, that belongs to the same LAN within single institutions. This possibly represents an uncommon population subset. To overcome this limitation, we have to properly handle the self-induced bias by conditioning the observation set accordingly. A quantification of the self-induced bias is given in Tab. III. It reports the percentage of peers and bytes exchanged among NAPA-WINE peers, considering contributors only, or all peers. A first important remark holds: NAPA-WINE peers clearly prefer to exchange data among them. For example, considering contributors in the PPLive experiment, NAPA-WINE peers contribute to more than 3.5% of exchanged data, even if they represent only 1% of the contacted peers. Similarly, they are 10% and 30% of observed peers considering SopCast and TVAnts respectively, but they contribute to 18% and 56% of exchanged bytes. We stress that by considering the set of peers other than NAPA-WINE, it will be possible to highlight and quantify which properties of the NAPA-WINE peers causes such a strong bias. Thus, to solve the issue concerning the self-induced bias, we explicitly filter the contributor set $\mathcal{P}'(p) = \mathcal{P}(p) \setminus \mathcal{W}$ in the above formulation, over which to evaluate P'_D, P'_U, B'_D, B'_U accordingly. Intuitively, restricting

App	Contributors		All-peers	
	Peer%	Bytes%	Peer%	Bytes%
PPLive	0.95	3.54	0.10	3.33
SopCast	10.25	17.71	4.60	19.45
TVAnts	29.82	56.31	15.56	56.06

TABLE III
NAPA-WINE SELF-INDUCED BIAS

the observation to \mathcal{P}' is equivalent to consider peers not involved in the experiment, i.e., to get rid of NAPA-WINE probe bias. For example, we expect that a preference versus a metric noticed in the full contributor set should be noticeable also in the set deprived of NAPA-WINE probes. In case the bias is still evident, then the preference was *not* artificially self-induced by NAPA-WINE peers.

Another issue concerns the fact that it exist a *correlation* between the considered metrics: for example, peers within the same subnetwork (NET=1) traverse zero hop (HOP=0) paths and belong to the same Autonomous System (AS) and Country (CC) as well. It may be therefore difficult to properly isolate the impact of each metric. At the same time, this correlation is likely to hold for the NAPA-WINE probes mainly, since they forms “clouds” of high-bandwidth PCs within the same LAN, CC, and AS. Considering the set \mathcal{P}' , it will be possible to identify which metric is having the highest impact, being the correlation smaller.

Finally, the *directionality* of the network property under consideration must be carefully handled. Indeed, we only dispose of information available at single vantage points in the network, collected at either the information source or sink. In particular, considering HOP metric, we can only directly measure $HOP(e, p)$, but not $HOP(p, e)$ which can be in general different from $HOP(e, p)$ due to Internet path asymmetry. However, we point out that the adoption of a coarse-granularity set should minimize the directionality issue. Indeed, it is likely that $HOP(e, p) \in HOP_P \Rightarrow HOP(p, e) \in HOP_P$ as well, i.e., it is unlikely that the reverse path $HOP(p, e)$ is short when the direct path $HOP(e, p)$ is long. Similarly, access bandwidth BW of a non NAPA-WINE peer e can be inferred only considering the uplink direction of peer e , thus only provided that e is one of p ’s contributors. However, in our experiments, the $\mathcal{U}(p)$ and $\mathcal{D}(p)$ sets are typically disjoint, which significantly limits the set of peers of which we are able to assess the access capacity: therefore, in order to gather conservative results, in the following we will limitedly consider the downlink direction for the BW metric.

IV. EXPERIMENTAL RESULTS

Empirical evaluation of PPLive, SopCast and TVAnts network-awareness is reported in Tab. IV. Specifically, we report, for both upload (U) and download (D) directions, the peer-wise (P) and byte-wise (B) preference metrics for each of the different network properties early considered. Tab. IV details results considering to the whole contributor

Net	App	Download				Upload			
		Non-Napa		All Contributors		Non-Napa		All Contributors	
		B'_D %	P'_D %	B_D %	P_D %	B'_U %	P'_U %	B_U %	P_U %
BW	PPLive	95.9	85.9	95.6	86.1	-	-	-	-
	SopCast	98.2	83.3	98.5	85.3	-	-	-	-
	TVAnts	96.5	83.2	98.2	89.6	-	-	-	-
AS	PPLive	6.5	0.6	12.8	1.3	0.8	0.2	1.8	0.5
	SopCast	0.6	0.7	3.5	3.9	1.7	0.7	6.4	3.9
	TVAnts	7.3	3.3	32.0	13.5	11.6	1.8	30.1	9.6
CC	PPLive	6.5	0.6	13.1	1.4	1.1	0.3	2.1	0.6
	SopCast	0.6	0.8	4.0	4.4	1.7	0.8	7.2	4.4
	TVAnts	7.6	4.0	37.9	16.3	14.3	3.1	37.7	12.5
NET	PPLive	-	-	9.9	0.8	-	-	1.4	0.3
	SopCast	-	-	2.0	2.6	-	-	3.5	2.6
	TVAnts	-	-	18.1	6.7	-	-	18.1	5.4
HOP	PPLive	42.2	41.1	51.4	42.4	30.4	40.4	31.7	41.0
	SopCast	29.0	40.7	37.9	48.0	45.9	43.0	56.9	49.8
	TVAnts	62.1	55.0	81.1	71.9	57.8	53.0	78.9	67.2

TABLE IV
NETWORK AWARENESS AS PEERWISE AND BYTEWISE BIAS

set (P_U, P_D, B_U, B_D) or to the contributor set excluding the NAPA-WINE probes (P'_U, P'_D, B'_U, B'_D) .

A. BW Awareness

As previously stated, in this case we are able to confidently infer the access capacity of peers only provided that they are uploading video content to a NAPA-WINE peer: therefore, we are forced to limitedly consider the downlink directions. From Tab. IV it can be clearly seen that, for all applications, a very strong preference for high-bandwidth peers is shown. Indeed, high-bandwidth peers represent 83% – 86% of the contributors, from which 96% – 98% of the traffic is received. This clearly shows that all applications are i) very efficient in pinpointing high-bandwidth peers, that ii) are then preferentially exploited to download the stream. The NAPA-WINE peers add little bias, so that percentages do not change much by excluding them from the statistics.

Not surprisingly, we can conclude that BW-awareness is definitively embedded in all P2P-TV applications.

B. AS and Country Awareness

We turn our attention to location awareness by considering the AS and CC metrics. Considering download direction, it can be seen that SopCast is unaware of AS location. Indeed, P_D is almost equal to B_D , which suggests that peers in the same AS are not preferentially selected to download data from. On the contrary, both PPLive and TVAnts show higher AS-awareness. Considering non-NAPA-WINE contributors, a PPLive peer downloads from $P'_D=0.6\%$ of peers $B'_D=6.5\%$ of traffic, i.e., there is a byte preference 10 times larger than a peer preference. The same factor holds including NAPA-WINE peers (which then do not bias the results). Similarly for TVAnts, in which $B'_D=7.6\%$ of the bytes are downloaded from $P'_D=3.3\%$ of the non-NAPA-WINE contributors, i.e., a B'_D/P'_D ratio equal to 2. Recalling that the total number of peers observed in the TVAnts experiment is two order of magnitude smaller than the one involved in the PPLive

experiment, we can conclude that TVAnts is also much more efficient in discovering peers within the same AS (13.5% on average) than PPLive (1.3% on average).

Looking at the downloaded traffic with respect to the Country geolocation of peers, we observe that almost the same percentages are observed as in the AS preference case. Since two peers in the same AS are also located within the same Country, we can state that no country preference is shown, i.e., the CC preference is due to the AS preference. Finally, considering the upload directions, similar conclusions can be drawn.

To better explore the issue related to peer locality, Figure 2 shows the average amount of traffic transferred from a high bandwidth NAPA-WINE peer belonging to AS- i to a high bandwidth NAPA-WINE peer within AS- j , for all the AS pairs. The intra-AS traffic is enlightened in black.

At a first look, only the PPLive-Popular experiment clearly suggests that the system favors intra-AS traffic over inter-AS-traffic. However, we notice that most of the intra-AS traffic is in this case local traffic (hop count equal to zero). To a better look, also TVAnts presents some bias in favor intra-AS traffic. Indeed, the ratio between the average amount traffic exchanged among intra-AS peers (reported in black) versus inter-AS peers (reported in gray) R is equal to 1.93, i.e., about twice the traffic is exchanged among peers within the same AS compared to peers in other ASs. Moreover the contribution provided by non local traffic is significant in this case. Neither SopCast nor PPLive show such bias, being $R = 0.2$ and $R = 0.98$ respectively.

Thus, we conclude that both SopCast and PPLive do not tend to favor traffic exchange within the same AS (excluding the traffic exchanged among peers in the same SubNet). However, intra high-bandwidth NAPA-WINE traffic is significant. This strong bias exhibited by intra high-bandwidth NAPA-WINE peers confirms that both applications tend to only favor downloading traffic from high-bandwidth peers.

In conclusion it turns out that for all the systems, the




Fig. 2. Average amount of exchanged data among the ASs involved in the experiments.

peer upload bandwidth seems to be the dominant metric that drives the selection of the peers from which downloading. However PPLive and especially TVAnts exploits also some form of locality while in SopCast the choices seem completely independent on the peers location.

C. NET Awareness

We now evaluate the potential preference to exchange traffic with peers in the same subnet (NET). The set of peers in the same subnet includes only NAPA-WINE peers, so that $\mathcal{P}' = \emptyset$. Results show that also in this case, PPLive and TVAnts only exhibit NET awareness, for both upload and download directions. Indeed, about 10% and 18% of the bytes are received from about 1% and 7% of hosts which are in the same subnet respectively. Conversely, SopCast does not show any evidence of subnet awareness. However, the NET preference can be also enforced by the AS preference. Looking at the ratio between P over B for the AS and NET preferences, we observe that they are very similar, which underlines that the NET preference is due to the AS preference and thus is not explicitly enforced.

D. HOP Awareness

Finally, we investigate the IP distance preference. In this case, no particular evidence of preference toward shorter paths is underlined. Indeed, looking at the non-NAPA-WINE peers, almost no difference emerges comparing P' and B' . Only

TVAnts shows a small preference to download from closer nodes. Considering the complete set \mathcal{P} , the self-induced bias of NAPA-WINE peers shows up, artificially highlighting a HOP preference, which is instead due to BW and AS preference. We can conclude that no HOP awareness emerges.

V. CONCLUSIONS

In this paper we have proposed a methodology to highlight which metric is exploited by P2P-TV applications to optimize the video delivery. Considering three popular P2P-TV applications, namely PPLive, SopCast and TVAnts, we have shown that a clear preference to exploit high-bandwidth peers emerges in all analyzed systems. Additionally, TVAnts and PPLive prefer to exchange data among peers in the same Autonomous System. However, no evidence of preference versus peers in the same subnet, or having a shorter path, emerges from any of the system under observation.

Results therefore suggests that future P2P-TV applications could improve the level of “network-awareness”, by better localizing the traffic the network has to carry, seeking shorter paths, exploiting topology knowledge, etc.

ACKNOWLEDGMENTS

This work has been funded by the European Commission under the FP7 STREP Project “Network-Aware P2P-TV Application over Wise Networks” (NAPA-WINE). We would like

to thank all partners that took part in the experiment for the valuable contribution to collect the data set.

REFERENCES

- [1] PPLive, <http://www.pplive.com>.
- [2] SOPCast, <http://www.sopcast.com>.
- [3] TVAnts, <http://www.tvants.com>.
- [4] 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*, Vol.9, No.8, pp.1672-1687, Dec. 2007.
- [5] B. Li, Y. Qu, Y. Keung, S. Xie, C. Lin, J. Liu, X. Zhang, "Inside the New Coolstreaming: Principles, Measurements and Performance Implications", *IEEE INFOCOM'08*, Phoenix, AZ, Apr. 2008.
- [6] C. Wu, B. Li, S. Zhao, "Multi-channel Live P2P Streaming: Refocusing on Servers" *IEEE INFOCOM'08*, Phoenix, AZ, Apr. 2008.
- [7] L. Vu, I. Gupta, J. Liang, K. Nahrstedt, "Measurement of a large-scale overlay for multimedia streaming" *Proc. of the 16th International Symposium on High Performance Distributed Computing*, Monterey, CA, June 2007.
- [8] F. Wang, J. Liu, Y. Xiong, "Stable Peers: Existence, Importance, and Application in Peer-to-Peer Live Video Streaming" *IEEE Infocom'08*, Phoenix, AZ, Apr. 2008.
- [9] X. Hei, Y. Liu, K.W. Ross, "Inferring Network-Wide Quality in P2P Live Streaming Systems," *IEEE JSAC*, special issue on P2P Streaming, Vol.25, No.9, pp.1640-1654, Dec. 2007.
- [10] S. Agarwal, J. P. Singh, A. Mavlankar, P. Baccichet, and B. Girod, "Performance and Quality-of-Service Analysis of a Live P2P Video Multicast Session on the Internet," *IEEE IwQoS'08*, Enschede, NL, June 2008.
- [11] S.Ali, A.Mathur, H.Zhang, "Measurements of Commercial Peer-To-Peer Live Video Streaming," *In Proc. of Workshop on Recent Advances in Peer-to-Peer Streaming*, Waterloo, ON, Aug. 2006.
- [12] T. Silverston, O. Fourmaux, "Measuring P2P IPTV Systems," *ACM NOSSDAV'07*, Urbana-Champaign, IL, June 2007.
- [13] "Network-Aware P2P-TV Application over Wise Networks ", <http://www.napa-wine.eu>.
- [14] D.Ciullo, M.A. Garcia, A. Horvath, P. Veglia, "Characterization of P2P-TV traffic," NAPA-WINE technical Report WP3.1n1. <http://www.napa-wine.eu/>