Network Awareness of P2P Live Streaming Applications: A Measurement Study

Original

Network Awareness of P2P Live Streaming Applications: A Measurement Study / Ciullo, Delia; Garcia, M. A.; Horvath, A.; Leonardi, Emilio; Mellia, Marco; Rossi, D.; Telek, M.; Veglia, P.. - In: IEEE TRANSACTIONS ON MULTIMEDIA. - ISSN 1520-9210. - STAMPA. - 12:1(2010), pp. 54-63. [10.1109/TMM.2009.2036231]

Availability:

This version is available at: 11583/2298242 since:

Publisher:

IEEE

Published

DOI:10.1109/TMM.2009.2036231

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)

Network Awareness of P2P Live Streaming Applications: A Measurement Study

Delia Ciullo, Maria Antonietta Garcia, Akos Horvath, Emilio Leonardi, *Senior Member, IEEE*, Marco Mellia, *Senior Member, IEEE*, Dario Rossi, *Member, IEEE*, Miklos Telek, and Paolo Veglia

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, due to their closed and proprietary design. In this paper, we present large scale 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. We first define a general framework to quantify which network layer parameters leverage application choices, i.e., what parameters mainly drive the peer selection and data exchange. We then apply the methodology to a large dataset, collected during a number of experiments where we deployed about 40 peers in several European countries.

From analysis of the dataset, we observe that TVAnts and PPLive exhibit a mild preference to exchange data among peers in the same autonomous system the peer belongs to, while this clustering effect is less intense in SopCast. However, no preference versus country, subnet or hop count is shown. Therefore, 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.

Index Terms—Locality awareness, multimedia streaming, neighbor selection, overlay, peer-to-peer networks.

I. INTRODUCTION AND MOTIVATIONS

PEER-TO-PEER Live Streaming (P2P-TV) systems are candidates for becoming the next Internet killer applications as testified by the growing success of commercial systems such as PPLive, SopCast and TVAnts. They allow to “watch television” over the Internet, granting to anyone to become a content provider by limiting the infrastructure costs, while giving the chance to break broadcasting constraint so that anyone can watch any content anywhere, at anytime.

Manuscript received April 06, 2009; revised September 07, 2009. First published November 13, 2009; current version published December 16, 2009. This work was supported by the European Commission under the FP7 STREP Project “Network-Aware P2P-TV Application over Wise Networks” (NAPA-WINE). The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Jiangchuan (JC) Liu.

D. Ciullo, M. A. Garcia, E. Leonardi, and M. Mellia are with the Politecnico di Torino, Torino 10129, Italy (e-mail: ciullo@tlc.polito.it; antonieta@tlc.polito.it; leonardi@tlc.polito.it; mellia@tlc.polito.it).

A. Horvath and M. Telek are with the Budapest University of Technology and Economics, Budapest H-1521, Hungary (e-mail: horvakos@hit.bme.hu; telek@hit.bme.hu).

D. Rossi and P. Veglia are with the École Nationale Supérieure des Télécommunications, Paris 75014, France (e-mail: dario.rossi@enst.fr; paolo.veglia@enst.fr).

Digital Object Identifier 10.1109/TMM.2009.2036231

P2P-TV systems have already attracted an audience up to several millions of users [1] and drawn the attention of Internet Service Providers (ISPs). In particular, ISPs are worried by the impact that P2P-TV traffic can have over the network infrastructure. Indeed, while from the application point of view it is perfectly legitimate to exchange content with any peer worldwide, from the network perspective, it is much more efficient if peers download (and upload) chunks from “close” peers, e.g., peers in the same subnet, autonomous system (AS) or country.

Considering the most successful P2P-TV applications, 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 ISPs 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). Therefore a systematical analysis is needed to understand the impact that current P2P-TV services may have on the Internet. This is precisely one of the goals of our recently funded project called “Network-Aware P2P-TV Application over Wise Networks” (NAPA-WINE) [2]–[4]. We carried out this work within the confines of this project. This work aims at assessing of level of “network awareness” embedded in the currently deployed systems, i.e., which network property influences the decision taken by a P2P-TV application. Are peers selected at random? Is the traffic confined within the same AS the peer belongs to? Does a peer preferentially download traffic from nearby nodes?

To answer all the above questions, we define a general methodology first. We then run some large testbed experiments during which each P2P-TV system was considered. More than 40 peers spreading over four different countries were instructed to watch the same TV stream at the same time. By applying the proposed methodology, we highlight which parameters affect the peer selection and data exchange policies. We conclude that only TVAnts and PPLive exhibit a mild preference to exchange data among peers in the same AS. At the same time, no preference versus country, subnet or hop count is shown by any system. Despite the content is available from peers on the same LAN, about 82% of the video chunks are fetched from peers outside the LAN considering TVAnts. Percentages grows to 90% for PPLive and 98% for SopCast, respectively. Moreover, only 32% of the content is fetched from peers inside the AS where TVAnts peers are. Even worse, PPLive and SopCast peers receive the large majority of traffic from outside the AS (87% and 96%, respectively). The presented results

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			ASx	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	-	-

underline the need for the development of newer and network friendlier P2P-TV systems, an interesting topic deserving future investigation. To this extent, the principal goal of the NAPA-WINE project is to design a novel P2P-TV system that *explicitly* optimizes ISP resource utilization. According to the NAPA-WINE vision, peers should download/upload the stream from/to nearby peers, they should minimize the path length, and in general they should leverage information about the network status. According to the results presented in this paper, very little network awareness is embedded in current P2P-TV applications.

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. The second important difference lies on the *scale* of the experiment, which in our case involves more than 40 vantage points scattered across European countries and it is representative of very different network setups.

The rest of this work is organized as follows. First, we detail the measurement setup in Section II, where we also present a preliminary quantitative description of the performed experiments. We then introduce the methodology for the analysis of the experimental data in Section III, introducing the metric that we will use to assess P2P systems network awareness. Experimental results are reported in Section IV, while we devote Section V to a spatial and temporal analysis of the peer selection process. Section VI overviews related work, and finally, Section VII concludes the paper.

II. EXPERIMENTAL SETUP

The results of this paper are based on a large testbed we setup, whose main features are summarized in Table I. Partners of the NAPA-WINE project took part in the experiments by running P2P-TV clients on PCs connected either to their 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 seven different industrial/academic sites, and seven home PCs. PCs are distributed over four countries, and connected to six different ASs, while home PCs are connected to seven other ASs. Therefore, the setup is representative of a significant number of different network environments. In the following, we refer to the set of PCs used during the experiment as “NAPA-WINE peers”.

In P2P-TV systems, hosts running the application (called peers) form an overlay topology by setting up virtual links over which they transmit and receive information. A source peer is responsible to inject the video stream, by chopping it into segments (called chunks) of few kilobytes, which are then sent to a sub set of its neighboring peers (called neighbors). Each peer can contribute to the chunk diffusion process, by retransmitting them to its neighbors following a swarming like behavior, as in file sharing P2P systems like BitTorrent. The major differences between P2P-TV systems and traditional P2P file sharing applications are 1) that the source is generating the stream in real time, 2) that data must be received by peers at almost constant rate, and 3) that chunks must arrive almost in sequence so that they can be quickly played at the receiver.

We considered three different applications, namely PPLive, SopCast and TVAnts, and we performed several one-hour-long experiments during April 2008, where peers were watching the same channel at the same time. Packet-level traces were collected and later analyzed. Since P2P-TV applications are mostly popular in Asian countries, we tuned each application to CCTV-1 channel during China peak hours [5]. In all cases, the nominal video stream rate was 384 kbps, Windows Media 9 Encoder was used, and the video quality perceived by users was very similar across systems. Results reported in this paper refer to more than 120 h of experiments, corresponding to more than 140 M packets. Collected traces are also made available to the research community upon request.

Let us first give some preliminary definition. It has been previously observed that P2P-TV peers exchange packets of typical length, i.e., very short packets carrying signaling information, and much longer packets carrying video information [5]. Let $N_v(i, j)$ be the number of packets sent from peer i to peer j whose size is equal to the typical video packet length. To distinguish between peers that exchanged mainly signaling information, and peers that exchanged actual video content too, we say that Peer i is a “TX (transmitting) contributing” peer for j if $N_v(i, j)$ is larger than threshold $M = 5$, i.e., peer i transmitted at least five video packets to peer j . At the same time, j is a “RX (receiving) contributing” peer for i , i.e., j received at least five video packets from i . We verified that this heuristic gives accurate and conservative results for classifying the contributing peers. Results are also consistent with results of the heuristic presented in [5] in which only PPLive was analyzed.

We start by giving some preliminary insights from the collected data. Both the average value over all the NAPA-WINE peers and the coefficient of variation [i.e., $cv = \sqrt{\sigma^2(X)/E^2(X)}$, where $\sigma(X)$ and $E(X)$ are the standard deviation and average of X , respectively] are reported. Table II presents the following simple metrics which are evaluated considering all NAPA-WINE peers: 1) receiving data rate, 2) transmitting data rate, 3) number of contacted peers (i.e., the number of peers that successfully exchanged at least one packet), 4) number of RX contributing peers, 5) number of TX contributing peers, and 6) percentage of peers that have never replied to any message.

The first and the second rows show the average inbound and outbound data rate, including both video and signaling traffic. As we can expect, on the reception side, no significant differ-

TABLE II
MEAN AND COEFFICIENT OF VARIATION OF SOME STATISTICS COLLECTED
DURING ALL THE EXPERIMENTS

	PPLive		SopCast		TVAnts	
	Mean	<i>cv</i>	Mean	<i>cv</i>	Mean	<i>cv</i>
Avg RX rate [kbps]	549	0.22	482	0.05	415	0.04
Avg TX rate [kbps]	3150	1.02	262	1.04	396	0.65
N. contacted peers	22263	0.53	740	0.29	222	0.19
N. RX contrib peers	382	0.50	139	0.43	54	0.31
N. TX contrib peers	958	0.73	137	0.54	66	0.54
% non-resp. peers	27	0.87	25	0.73	30	0.31

ence can be observed among the different applications, as testified by the small coefficients of variation. This is due to the fact that the dominant component of the received traffic is constituted by the video content, whose average rate is the same for all the peers and applications (recall that all the considered applications adopt the same streaming encoding technique). For PPLive dataset, the higher *cv* and average values suggest that the receiver rate can be higher than the stream rate. This is due to the large incoming traffic that high-bandwidth peers receive, i.e., to the signaling messages they have to handle which are sent by the peers receiving the uploaded video content.

More interestingly, the transmission rates are significantly different. Indeed, the transmission data rates are strongly dependent on the specific mechanism adopted by each system to distribute the video content. First, the transmission data rate is largely correlated to the upload bandwidth of peers; all the applications successfully exploit high bandwidth peers, demanding to some of them a significant contribution. To confirm this, we investigated the upload rate of each NAPA-WINE peers and we found that high-bandwidth NAPA-WINE peers are acting as “amplifiers”, i.e., they upload much more than what they download. Instead, peers that are connected by CABLE/DSL show much smaller upload rate. On this regard, we observe that PPLive may be significantly demanding, so that high bandwidth peers push their average transmission data rate to more than 10 Mbit/s, with short time peaks reaching 30 Mbit/s. The high coefficient of variation is also due to difference among the upload capacity of the peers.

Huge differences among the systems arise considering the number of contacted peers, which is on the order of tens of thousands for PPLive, up to one thousand for SopCast, and in the order of few hundreds for TVAnts. We believe that the huge difference in the number of contacted peers is mainly due to the algorithms used to discover and to maintain the overlay, on which we will come back later in Section V.

The fourth and the fifth rows report the number of RX contributing peers and TX contributing peers, respectively. If we compare the number of contacted peers against the number of contributing peers, we can see that TVAnts exploits the contacted peers at the highest degree, i.e., one fourth of the contacted peers are also contributing peers. Considering SopCast, about one fifth of the contacted peers are used to download some content. Finally, in the PPLive case, the number of contacted peers is more than 50 times higher than the number of contributing peers. Focusing on the number of TX contributing peers (i.e., peers to which some content has been transmitted to), we observe that peers with low upload bandwidth serve fewer peers than peers with high bandwidth. This fact is expected since

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

low upload bandwidth peers have limited capability to upload data to other peers. For PPLive, the significantly larger number of both contacted and contributing peers explains the higher upload rate.

The last row shows the fraction of the peers which did not reply at all, i.e., failing peers. All of the systems show a very high failing ratio (25%–30%). This hints to little optimization on the P2P-TV control plane, so that outdated information is still distributed among peers. Moreover, we noticed that in all experiments, NAPA-WINE peers tried to contact peers with private IP address, with some peers performing address scan of whole subnetworks.

Fig. 1 shows the geographical distribution of the number of contacted peers, the amount of received and transmitted bytes, labeled #, RX and TX respectively. China (CN) and countries in which experiments were performed are shown, with the rest of the countries labeled with a star. Percentages are expressed over the total number of observed peers, which amounts to 181 729 for PPLive, 4057 for SopCast and 550 for TVAnts. As expected, China is the predominant country, though it is easy to gather that a non negligible fraction of the data is exchanged within the countries of NAPA-WINE peers: this hints to a bias in the peer selection, which will be more rigorously investigated in the following sections. The large difference in geographical distribution of contacted peers shows that different algorithms are used by the different applications. In particular, TVAnts seems to adopt a “smart” choice in selecting peers. Indeed, among the 550 total peers, 154 are located in Europe and 229 in China. Considering the watched channel and time of the experiment, the popularity of the application should be much higher in China than in Europe, so that we can conclude that the observed peers are only a biased and small fraction of the total active population. On the contrary, PPLive adopts a less smart peer discovery policy, so that the total number of contacted peers is more than 50 times higher than TVAnts or SopCast. In this case, 748 peers are located in Europe only (180.000 in China).

III. FRAMEWORK FOR PEER SELECTION 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. We define a flexible framework that allows us not only to inspect the level of “awareness” of a P2P system with respect to the underlying network, but also to assess whether peers behave fairly with respect one to another, i.e., if the peers are incentivized to the mutual data exchange. In particular, we consider:

- **AS** (p): the autonomous system where peer p is located;
- **CC** (p): the country, which peer p belongs to;
- **NET** (p): the subnetwork, which peer p belongs to;
- **HOP** (p, e): the IP hop-distance between peers p and e ;
- **SYM** (p, e): the symmetry of byte-wise data exchanges between peers p and e .

A. Framework Definition

Let $p \in \mathcal{W}$ denote a peer that belongs to the NAPA-WINE set $\mathcal{W}$. Let $\mathcal{P}(p)$ denote the set of *contributing* peers, p exchanges data with. That is, $\mathcal{P}(p)$ is composed by the peers to which p transmitted or/and from which p received some video information. Let $\mathcal{U}(p)$ denote the subset of peers to which p is uploading video content, and $\mathcal{D}(p)$ the subset from which p is downloading video from. $\mathcal{U}(p)$ and $\mathcal{D}(p)$ are two (non-disjoint) subsets of $\mathcal{P}(p)$, and $\mathcal{U}(p) \cup \mathcal{D}(p) = \mathcal{P}(p)$.

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 with $X(p, e) \in X$ the observed value of $X(\cdot)$ for the pair (p, e) . We partition $\mathcal{P}(p)$ into two classes based on $X(p, e)$, such that one class should intuitively be preferred from the application (e.g., good versus bad peers). More formally, we partition the support X into two disjoint sets: the 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 the ease of notation, let $\mathbf{1}_P(p, e)$ be the identity function which takes the value of 1 if $X(p, e) \in X_P$ 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 peer $p \in \mathcal{W}$, and let us define

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

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

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

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

where U and D subscripts are used to indicate the upload and download traffic, respectively. $Peer_{U|P}(p)$ counts 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 the 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 $\mathcal{W}$ of NAPA-WINE peers, 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 definitions hold for $Peer_{U|\bar{P}}$ and $Byte_{U|\bar{P}}$. 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-TV systems, since they are neither sensitive to the unit of measure, nor to the actual value of X .

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

B. Preferential Partitions

As preferential classes, we consider the following:

- **AS** : $\mathbf{1}_P(p, e) = 1$ if and only if $AS(p) = AS(e)$, i.e., both peers are located in the same autonomous system;¹
- **CC** : $\mathbf{1}_P(p, e) = 1$ if and only if $CC(p) = CC(e)$, i.e., both peers are located in the same country;
- **NET** : $\mathbf{1}_P(p, e) = 1$ if and only if $HOP(e, p) = 0$, i.e., peers belong to the same subnet;
- **HOP** : $\mathbf{1}_P(p, e) = 1$ if and only if $HOP(e, p) < \text{median}[HOP]$, i.e., the number of hops between e and p is smaller than the median distance among all peers;
- **SYM** : $\mathbf{1}_P(p, e) = 1$ if and only if $1/2 < B(e, p)/B(p, e) < 2$, i.e., the amount of data received (sent) is at most twice the amount of data sent (received).

While for AS, CC and NET the preferential set choice is straightforward, the HOP and SYM 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. We use the median of the distribution as threshold to define two subsets. Since the actual HOP median ranges from 18–20 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.

¹CC and AS have been determined by querying the “whois” database.

TABLE III
NAPA-WINE INDUCED BIAS

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

In case of incentive mechanism, we classify a data exchange as “symmetric” when the amount of data received is at most *twice* the amount of data sent, and vice versa. We point out that while this only enforces a loosely symmetrical relationship, we verified that the results are not very sensitive to these threshold choice (see Section IV-D).

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 peers *induced a bias* during the experiments. Recall that among NAPA-WINE peers there are several high-bandwidth peers, located in Europe only. Furthermore, all peers within the same institution are in the same LAN, and AS. This possibly represents an uncommon population subset.

A quantification of the induced bias is given in Table III. It reports the percentage of 1) NAPA-WINE peers over all peers observed during each experiment, and 2) bytes exchanged among NAPA-WINE peers over all exchanged bytes. Results are reported considering contributors only, or all peers. As first important remark, 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 less than 1% of the contributing peers. Similarly, they are about 10% and 30% of peers for SopCast and TVAnts respectively, but they contribute to 18% and 56% of exchanged bytes. We stress that by restricting the analysis to the set of peers other than NAPA-WINE, it will be possible to highlight and quantify which properties of the NAPA-WINE peers cause such a strong bias. To solve the issue concerning the induced bias, we introduce the set $\mathcal{P}'(p) \subset \mathcal{P}(p)$. Subset $\mathcal{P}'(p)$ is constituted by the peers in $\mathcal{P}(p)$ excluding the NAPA-WINE peers, formally $\mathcal{P}'(p) = \mathcal{P}(p) \setminus \mathcal{W}$. We evaluate the preference metrics also over the filtered set, getting P'_D, P'_U, B'_D, B'_U , accordingly. Intuitively, restricting the observation to $\mathcal{P}'$ is equivalent to consider peers not involved in the experiment. 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 peers. In case the bias is still evident, this means that the preference was *not* artificially induced by NAPA-WINE peers.

Another problem concerns the fact that it exists a *correlation* between the considered metrics: for example, peers within the same subnetwork ($\text{NET} = 1$) traverse paths of zero hop ($\text{HOP} = 0$), belong to the same AS and 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 peers mainly, since they form “clouds” of high-bandwidth PCs within the same LAN, CC, and AS. Considering the set $\mathcal{P}'$, where the correlation related to the locality among peers

is smaller, it will be possible to identify which metric has the highest impact.

All the observed parameters can be evaluated considering separately the download and upload direction of traffic, e.g., we can observe from (to) which countries the NAPA-WINE peers prefer to download (upload) the content. Notice that, for HOP metric, we can only directly measure $\text{HOP}(e, p)$, but not $\text{HOP}(p, e)$ which can be in general different from $\text{HOP}(e, p)$ due to Internet path asymmetry. However, we point out that the adoption of a coarse-granularity should minimize this issue. Indeed, it is likely that $\text{HOP}(e, p) \in \text{HOP}_P$, then $\text{HOP}(p, e) \in \text{HOP}_P$ as well, i.e., it is unlikely that the reverse path $\text{HOP}(p, e)$ is short when the direct path $\text{HOP}(e, p)$ is long. Finally, note that to compute the SYM metric it is necessary to compare the amount of transmitted and received data between any pair of peers.

IV. EXPERIMENTAL RESULTS

Empirical evaluation of PPLive, SopCast and TVAnts network-awareness is reported in Table 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. Table IV details results referring to the whole contributor set (P_U, P_D, B_U, B_D) and to the contributor set excluding the NAPA-WINE peers (P'_U, P'_D, B'_U, B'_D).

A. AS and Country Awareness

We first 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 ten times larger than a peer preference. The same factor holds including NAPA-WINE peers (which then do not bias the results). Similarly, for TVAnts, $B'_D = 7.6\%$ 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. Also, notice that 0.04% of *all* peers are in the same AS of NAPA-WINE peers in case of PPLive, and 3.6% in case of TVAnts. Still, as 1.3% of the *contributing* peers are located in the same AS for PPLive, and 13.5% for TVAnts, we can conclude that PPLive exhibits a stronger preference for peers within the same AS than TVAnts.

Looking at the downloaded traffic with respect to the peer CC, we notice 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 CC, we can conclude 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.

B. NET Awareness

We now evaluate the potential preference to exchange traffic with peers in the same subnet (NET). The set of peers in the

TABLE IV
 NETWORK AWARENESS AS PEER-WISE AND BYTE-WISE BIAS

Net	App	Download						Upload					
		Non-Napa			All Contributors			Non-Napa			All Contributors		
		B'_D	%	P'_D	%	B_D	%	P_D	%	B'_U	%	P'_U	%
AS	PPLive	6.5	0.6	12.8	1.3	0.8	0.2	1.8	0.5	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	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	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	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	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	14.3	3.1	37.7	12.5
NET	PPLive	-	-	9.9	0.8	-	-	1.4	0.3	-	-	1.4	0.3
	SopCast	-	-	2.0	2.6	-	-	3.5	2.6	-	-	3.5	2.6
	TVAnts	-	-	18.1	6.7	-	-	18.1	5.4	-	-	18.1	5.4
HOP	PPLive	42.2	41.1	51.4	42.4	30.4	40.4	31.7	41.0	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	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	57.8	53.0	78.9	67.2
SYM	PPLive	3.3	4.8	4.3	5.0	-	-	-	-	-	-	-	-
	SopCast	6.7	13.0	7.8	14.2	-	-	-	-	-	-	-	-
	TVAnts	12.4	10.9	20.0	14.3	-	-	-	-	-	-	-	-

Fig. 2. CDF of the number of peers and the amount of received data versus the number of hops.

same subnet includes only NAPA-WINE peers, i.e., $\mathcal{P}' = \emptyset$. 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. This points out that peers in the same autonomous system but not in the same NET are equally preferred as the peers in the same NET (and in the same AS). Therefore, the AS preference is stronger than the NET preference. Notice also that the AS locality is overall quite marginal, so that the majority of the traffic is still coming from other ASs. As such, there is large margin to improve the network friendliness of P2P-TV applications.

C. HOP Awareness

We also investigate the HOP count 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.

To further testify this finding, Fig. 2 reports the cumulative distribution function (CDF) of contacted peers (solid line) and of the received bytes (dashed line) versus the distance between peers in hop count, not including the NAPA-WINE peers. TVAnts only shows a slight commitment to the closest

peers, while SopCast and PPLive seem to ignore peer distance considering the hop number.

D. SYM Incentive Mechanism

Considering P2P file sharing applications, incentives mechanisms have been successfully introduced to improve system performance. For example, BitTorrent clients play a tit-for-tat game with other peers, so that the more a peer sends to a neighbor, the more it will receive from it. This enforces a sort of symmetry between the amount of bytes sent and received by peers.

We explore whether there exists some incentive mechanism that enforces symmetry in P2P-TV systems as well. Results are reported in Table IV: Even if we arbitrarily report SYM under the download section of the table, we recall that it is a metric that requires to compare the amount of traffic exchanged in both directions (upload and download) between two peers. Considering non NAPA-WINE peers, it emerges that only a small percentage (from 5% considering PPLive to 13% considering SopCast) of the links are symmetrical. Moreover, the amount of data exchanged between these peers is not predominant (less than 12%). This suggests that P2P-TV systems do not enforce any tit-for-tat like mechanism. Indeed, being the download rate constrained by the actual video rate, these systems are engineered in such a way that peers with limited upload capacity can receive the video stream anyway, even if they are not able to redistribute it.

This is highlighted in Fig. 3, which reports the amount of transmitted versus received bytes considering contributing peers. Intuitively, if a tit-for-tat like incentive mechanism were implemented, then a strong correlation should be observed

Fig. 3. Scatter plot of RX versus TX bytes, guidelines represent the tit-for-tat boundaries $\text{TX}/\text{RX} = 2$ and $\text{RX}/\text{TX} = 2$.

so that points accumulate along the $y = x$ diagonal. Log/log scale is used to better represent results. The area between the $\text{TX}/\text{RX} = 2$ and $\text{TX}/\text{RX} = 1/2$ lines corresponds to symmetrical exchanges as previously defined. Looking at Fig. 3, it can be seen that the wide majority of points fall outside this area, as already reported in Table IV. Only in the SopCast case, a cloud of points lies in the symmetry strip, though such points correspond to moderate amount of data (i.e., few thousand Bytes). Considering PPLive, we observe that a lot of points accumulate along the $y = 10x$ line, corresponding to peers that mostly download data from the NAPA-WINE peers.² The dense points accumulating around $y = 10^4$ and $x = 10^4$ are also a consequence of a private mechanism of the application. Summarizing, no evidence of a symmetric tit-for-tat like incentive emerges for any system.

To summarize, we have shown that the three applications behave differently, and by means of inference on passive measurements, we have empirically quantified these differences. While our results point out the lack of network awareness of such systems, the picture is far from being complete: for instance, the different behaviors are a direct consequence of specific, proprietary and therefore unknown mechanisms adopted by such systems. However, we point out that by pure black-box measurement is unfortunately impossible to understand what are the specific algorithms implemented, as well as the parameters adopted by each system.

V. DYNAMICS OF CONTACTED PEERS

In this section we supplement the analysis of peer selection, by inspecting its dynamics from both a temporal and a spatial point of view as well.

A. Temporal Analysis

To better understand the peer selection process, Fig. 4 plots the dynamics of the contacted peers versus time. One arbitrary NAPA-WINE peer is represented in each figure, since they are qualitatively all similar. For PPLive the behavior of both high-bandwidth and DSL nodes are reported. The continuous line reports the total number of contacted peers versus time, while the squared dots show the arrival of contributing peers, whose departure is shown by the crosses in the same line. In this context, the arrival and the departure of a peer is identified by the time of the first and last observed packet from it, respectively.

²The factor 10 can be explained by considering that PPLive protocol uses some acknowledgment message to possibly confirm the reception of data packets. The size of ACKs is about one tenth of the size of data packets.

Positive y-axis reports the remote peers that were contacted by the NAPA-WINE peer first, while negative y-axis reports the peers that were the initiator of the connection toward the NAPA-WINE peer. For PPLive, the evolution of the number of contacted peers is reported in Fig. 5, since it is much larger than other quantities.

Both TVAnts and SopCast limit this rate as soon as a good set of contributing peers is obtained (after about 250 s and 500 s, respectively). On the contrary, PPLive has a stronger greedy behavior, essentially contacting new peers at an almost constant rate. These different overlay exploration algorithms clearly drive the total number of contacted peers, which is much higher for PPLive.

As already observed in Table II, the number of contributing peers is limited to few tens for TVAnts and SopCast. In addition, the set of contributing peers is rather stable along the whole experiment duration, i.e., the contributing peer contact time lasts several tens of minutes. In the case of PPLive on the contrary, the number of contributing peers is much higher (several hundreds, up to 1000) and it exhibits a higher degree of variability. This can be explained considering the fact that the number of possibly good candidates is higher, and the peer selection policy continuously tries to improve performance by testing new peers. No major difference is shown between DSL or high bandwidth nodes considering the number of contributing peers that are contacted by the NAPA-WINE peer. On the contrary, the number of peers that initiated a connection to the high-bandwidth peer is larger than the one that initiated a connection to the DSL node. That is, the high-bandwidth peer gets more requests. This suggests that the information about the peer upload capacity is made available to other peers.

B. Spatial Analysis

Finally, we complete our analysis of the peer selection process by considering the spatial properties that can be inferred by exploiting our large number of measurement points. Fig. 6 shows the CDF of the common contributing peers, i.e., the probability that a contributing peer is seen by N different NAPA-WINE peers (on the x-axis).

For example, Fig. 6 shows that for PPLive, there are 50% of the peers can be observed by only 1 NAPA-WINE user and 70% of the peers can be observed by either 1 or 2 NAPA-WINE users. For SopCast, about 30% of peers are seen by only one NAPA-WINE peer, and 40% as seen by two NAPA-WINE users. These percentages reduces to less than 15% for both cases. Similarly, considering the probability that a peer has been contacted by

Fig. 4. Arrival and departure process of the all and video contributing peers.

Fig. 5. Arrival and departure process of the all peers for PPLive.

Fig. 6. CDF of common contributing peers.

at least 20 NAPA-WINE peers, we notice that for PPLive, the probability is close to one, meaning that a negligible set of peers have been contacted by more than 20 NAPA-WINE hosts; for SopCast, this probability is 0.8, meaning that there are about 20% of the peers that exchanged data with more than 20 different NAPA-WINE hosts; for TVAnts, there are 50% of the peers that are contacted by more than 20 different NAPA-WINE hosts.

In case a random independent and identically distributed (i.i.d.) selection is performed, the common peers CDF follows a Binomial distribution. On the contrary, in case of a correlated choice (i.e., when certain peers are preferred to other peers), a different trend is expected; for example, a more linear CDF

would suggest that peers prefer to contact the same subset of peers.

We assume that the number of contributing peers that exchanged data with NAPA-WINE peers during the experiment is a small fraction of all available peers. This assumption is supported by Fig. 1, in which TVAnts population is largely biased, suggesting a “smart” choice by peers. Furthermore, Fig. 4 shows that both TVAnts and Sopcast use a peer discovery mechanisms which is very greedy during the first part of peer life, after which the peer discovery rate slows down.

Fig. 6 clearly shows that for SopCast and TVAnts experiments the selection of peers from which to download is performed according to some algorithm that tends to correlate peer choice. For example, consider the probability that no more than 20 NAPA-WINE peers select the same contributing peer. For PPLive, is almost one, it is about 0.8 for SopCast, and it is about 0.5 for TVAnts. Indeed, for TVAnts, there are some peers that have been selected as contributing peers by most of the NAPA-WINE peers.

VI. RELATED WORK

A fairly large number of public P2P architectures and algorithms for the support of video distribution over the Internet has been proposed in the last three-four years within the scientific community [6]–[10]. Despite their clear merit, the above systems have only gained limited popularity—especially in comparison with commercial systems. The latter systems have indeed attracted a larger audience, up to several millions of users. The fact that such commercial systems follow a closed and proprietary design has motivated further research [5], [11]–[18], aimed at understanding these systems through on-field measurements.

Some works focus on single system, as [5], [11], and [12]. Such works, which exploit partial reverse engineering techniques and typically rely on active crawling methodologies, face the daunting task of understanding and implementing part of the system under analysis. As a consequence, this methodology is limited by the ability to break closed and proprietary systems, and we believe that they can be hardly extended to characterize all the possible P2P-TV applications. For example, [5] investigates PPLive, whereas [11] focuses on the commercial re-engineer of Coolstreaming, and [12] considers UUSee. All these papers focus on complementary metrics with respect to our work.

Other works, such as [13]–[16] instead study specific aspects of a P2P streaming system. For instance, [13] gives some preliminary results on the node degrees of popular versus unpopular channels in PPLive, while [14] instead investigates the stability of PPLive peers. Quality of service is of concern in [15] and [16]. Authors in [15] exploit an analysis of PPLive buffer maps, collected through protocol reverse engineering, to infer QoS metrics such as network-wide playback continuity, startup latency, playback lags among peers, and chunk propagation timing. Authors in [16] focus on similar metrics but instead exploit logs made available from an unspecified commercial P2P streaming system.

Despite all the above valuable work, to date, very few measurement studies compare different systems, such as [17]–[19], which are closest to our work. Authors in [17] analyze and compare PPLive and SopCast, investigating the time evolution of different metrics, like transmitted/received bytes, number of parents and children, etc. Authors in [18] present instead a comparative evaluation of PPLive, PPStream, SOPCast and TVAnts. Analysis is carried on in terms 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. In [19], authors set-up an *active* testbed to investigate the congestion control algorithms of different P2P-TV applications. Using active probes,

authors enforce artificial bandwidth limitations, packet loss and delay, and examine P2P-TV reaction to adverse network conditions. However, not all metrics potentially exploited by the overlay for neighbor selection and chunk scheduling can be artificially enforced—as, for instance, it is the case of the geographical and AS location of the contributing peers.

Our work differs from [17]–[19] in several aspects. The first is the aim, as our work focuses on a systematical exploration of the metrics, if any, that drive the peer-selection in the different systems. Second, we consider different aspects related to the overlay setup and download policies which are complementary to those addressed in [17]–[19]. An important last difference lies on the scale of the testbed, which in our case involves multiple vantage points scattered across European countries and it is representative of very different network setups.

This paper is a development of our previous work in [20]. We extended [20] with the investigation on the symmetry of the traffic and with new analysis of the dynamics of the peer selection process, from temporal and spatial point of view.

VII. 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 only TVAnts and PPLive exhibit a mild preference 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, neither the use of incentive mechanism emerge from any of the system under observation.

We believe that a much higher level of “network-awareness” has to be embedded in P2P-TV systems to better exploit and optimize the ISP resource utilization. In the context of the NAPA-WINE project, we are currently investigating how to reach this goal, e.g., to improve traffic localization, seeking shorter paths, exploiting topology knowledge, etc.

REFERENCES

- [1] G. Huang, “Experiences with PPLive,” in *Proc. Sigcomm 2007 Peer-to-Peer Streaming and IP-TV Workshop (P2P-TV)*, Kyoto, Japan, Aug. 2007.
- [2] Network-AwareP-TV Application Over Wise Networks. [Online]. Available: <http://www.napa-wine.eu>.
- [3] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz, “P4P: Provider portal for applications,” in *Proc. ACM SIGCOMM*, Seattle, WA, Aug. 2008.
- [4] Application-Layer Traffic Optimization (alto), IETF Working Group. [Online]. Available: <http://www.ietf.org/html.charters/alto-charter.html>.
- [5] X. Hei, C. Liang, J. Liang, Y. Liu, and K. W. Ross, “A measurement study of a large-scale P2P IPTV system,” *IEEE Trans. Multimedia*, vol. 9, no. 8, pp. 1672–1687, Dec. 2007.
- [6] C. Yang-Hua, S. G. Rao, S. Seshan, and H. Zhang, “Enabling conferencing applications on the internet using an overlay multicast architecture,” in *Proc. ACM SIGCOMM*, San Diego, CA, Aug. 2001.
- [7] S. Banerjee, B. Bhattacharjee, and C. Kommareddy, “Scalable application layer multicast,” in *Proc. ACM SIGCOMM*, Pittsburgh, PA, 2002.
- [8] V. N. Padmanabhan, H. J. Wang, P. A. Chou, and K. Sripanidkulchai, “Distributing streaming media content using cooperative networking,” in *Proc. ACM NOSSDAV'02*, Miami, FL, May 2002.
- [9] M. Castro, P. Druschel, A. M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh, “Splitstream: High-bandwidth multicast in a cooperative environment,” in *Proc. ACM SOSP'03*, Lake Bolton, NY, Oct. 2003.

- [10] X. Zhang, J. Liu, B. Li, and T. S. P. Yum, "DONet/Coolstreaming: A data-driven overlay network for peer-to-peer live media streaming," in *Proc. IEEE INFOCOM 2005*, Miami, FL, Mar. 2005.
- [11] B. Li, Y. Qu, Y. Keung, S. Xie, C. Lin, J. Liu, and X. Zhang, "Inside the new Coolstreaming: Principles, measurements and performance implications," in *Proc. IEEE INFOCOM'08*, Phoenix, AZ, Apr. 2008.
- [12] C. Wu, B. Li, and S. Zhao, "Multi-Channel live P2P streaming: Refocusing on servers," in *Proc. IEEE INFOCOM'08*, Phoenix, AZ, Apr. 2008.
- [13] L. Vu, I. Gupta, J. Liang, and K. Nahrstedt, "Measurement of a large-scale overlay for multimedia streaming," in *Proc. 16th Int. Symp. High Performance Distributed Computing*, Monterey, CA, Jun. 2007.
- [14] F. Wang, J. Liu, and Y. Xiong, "Stable peers: Existence, importance, and application in peer-to-peer live video streaming," in *Proc. IEEE Infocom'08*, Phoenix, AZ, Apr. 2008.
- [15] X. Hei, Y. Liu, and K. W. Ross, "Inferring network-wide quality in P2P live streaming systems," *IEEE J. Select. Areas Commun.*, Special Issue on P2P Streaming, vol. 25, no. 9, pp. 1640–1654, Dec. 2007.
- [16] 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," in *Proc. IEEE IwQoS'08*, Enschede, The Netherlands, Jun. 2008.
- [17] S. Ali, A. Mathur, and H. Zhang, "Measurements of commercial peer-to-peer live video streaming," in *Proc. Workshop Recent Advances in Peer-to-Peer Streaming*, Waterloo, ON, Canada, Aug. 2006.
- [18] T. Silverston and O. Fourmaux, "Measuring P2P IPTV systems," in *Proc. ACM NOSSDAV'07*, Urbana-Champaign, IL, Jun. 2007.
- [19] E. Alessandria, M. Gallo, E. Leonardi, M. Mellia, and M. Meo, "P2P-TV systems under adverse network conditions: A measurement study," in *Proc. IEEE Infocom'09*, Rio de Janeiro, Brazil, Apr. 2009.
- [20] M. Mellia, A. Horvath, M. Telek, D. Rossi, P. Veglia, D. Ciullo, M. A. Garcia da Rocha Neta, and E. Leonardi, "Network awareness of P2P live streaming applications," in *Proc. IEEE IPDPS'09, HotP2P Workshop*, Rome, Italy, May 2009.

Delia Ciullo received the degree in telecommunications engineering from Politecnico di Torino, Torino, Italy, in 2007. Since January 2008, she has been with pursuing the Ph.D. degree in the Electronic Department of Politecnico di Torino.

Her research interests are in the fields of energy-efficient network design, scaling properties in wireless networks, and traffic measurement in peer-to-peer systems.

Maria Antonietta Garcia received the B.Sc. degree in computer science from Catholic University of Pernambuco, Recife, Brazil, in 2005. In 2007, she has started pursuing the Ph.D. degree in electronic engineering and telecommunications at Politecnico di Torino, Torino, Italy.

Between 2004 and 2007, she worked as a Systems Engineer at C.E.S.A.R.—Recife Center for Advanced Studies and Systems. Her research interests are in the field of: peer-to-peer TV, modeling and analysis of peer-to-peer networks, and network

simulations.

Akos Horvath received the M.Sc. degree in computer science from Budapest University of Technology and Economics, Budapest, Hungary, in 2007.

He is currently working as a Research Fellow at the Department of Telecommunications at the Budapest University of Technology and Economics. His research interests include performance evaluation of peer-to-peer networks and telecommunication networks.

Emilio Leonardi (SM'09) received the Ph.D. degree in telecommunications engineering in 1995 from Politecnico di Torino, Torino, Italy.

He is currently an Associate Professor at Politecnico di Torino. In 1995, he visited the Computer Science Department of the University of California, Los Angeles; in 1999, the High Speed Networks Research Group, at Bell Labs, Holmdel, NJ; in summer 2001, the Electrical Engineering Department of Stanford University, Stanford, CA; and in summer 2003, the IP Group at Sprint Lab, Burlingame, CA.

His research interests are in the field of performance evaluation of wireless networks, P2P systems, and packet switching. He is the Scientific Coordinator of the EU 7-th FP STREP project "NAPA-WINE".

Marco Mellia (SM'08) received the Ph.D. degree in telecommunications engineering in 2001 from Politecnico di Torino, Torino, Italy.

In 1999, he was with the Computer Science Department at Carnegie Mellon University, Pittsburgh, PA, and since April 2001, he has been with the Electrical Engineering Department of Politecnico di Torino. He has coauthored over 140 papers published in international journals and conferences, and he participated in the program committees of several conferences like IEEE Infocom and ACM Sigcomm.

His research interests are in the fields of traffic measurement, P2P applications, and energy aware network design.

Dario Rossi (M'02) received the Ph.D. degree from Politecnico di Torino, Torino, Italy, in 2005.

During 2003–2004, he was with the Computer Science Department at the University of California, Berkeley. Since October 2006, he has been an Associate Professor at Telecom ParisTech, Paris, France. He has coauthored over 40 papers in leading conferences and journals. He participated in the program committees of several conferences like IEEE ICC, IPCCC, and Globecom. His research interests include P2P networks, Internet traffic measurement,

and sensor and vehicular networks.

Miklos Telek received the M.Sc. degree in electrical engineering from the Technical University of Budapest, Budapest, Hungary, in 1987 and the C.Sc. and D.Sc. degrees from the Hungarian Academy of Science, Budapest, in 1995 and 2004, respectively.

Since 1990, he has been with the Department of Telecommunications of the Technical University of Budapest, where he is a full Professor now. His current research interest includes stochastic performance modeling and analysis of computer and communication systems.

Paolo Veglia received the B.S. and M.S. degrees in computer science engineering from Politecnico di Torino, Torino, Italy, in 2005 and 2008, respectively. Since 2008, he has been pursuing the Ph.D. degree in the "Network, Mobility and Security" group at Telecom-ParisTech, Paris, France. His research interests include peer-to-peer, with particular attention to live video distribution, network measurements, and high speed packet processing.