


Provided for non-commercial research and education use.  
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

## Computer Networks

journal homepage: [www.elsevier.com/locate/comnet](http://www.elsevier.com/locate/comnet)

## Impact of adverse network conditions on P2P-TV systems: Experimental evidence

Eugenio Alessandria, Massimo Gallo, Emilio Leonardi, Marco Mellia, Michela Meo\*

Politecnico di Torino, Dipartimento di Elettronica, Corso Duca degli Abruzzi 24, 10129 Torino, Italy

### ARTICLE INFO

#### Article history:

Received 13 July 2010

Received in revised form 11 November 2010

Accepted 3 February 2011

Available online 24 February 2011

Responsible Editor: I.F. Akyildiz

#### Keywords:

Peer-to-peer

Video streaming measurement

### ABSTRACT

In this paper we define an experimental set-up to analyze the behavior of commercial P2P-TV applications under adverse network conditions. Our goal is to reveal the ability of different P2P-TV applications to adapt to evolving network conditions, such as delay, loss, available capacity, and presence of background traffic and to check whether such systems implement some form of congestion control. We apply our methodology to four popular commercial P2P-TV applications: PPLive, SOPCast, TVAnts and TVUPlayer. Our results show that all the considered applications are in general capable to cope with packet losses and to react to congestion arising in the network core. Indeed, all applications keep trying to download data by avoiding bad paths and carefully selecting good peers. However, when the bottleneck affects all peers, e.g., it is at the access link, the applications become rather aggressive, and potentially harmful for both other applications and the network.

We then observe the interference between TCP and P2P-TV traffic. As expected, P2P-TV applications do not perform TCP-friendly congestion control, causing in some cases problems to TCP traffic and to P2P-TV performance itself. Finally, we also verify that the applications are fair towards clients sharing the same access link, even when congestion arises.

© 2011 Elsevier B.V. All rights reserved.

### 1. Introduction and motivations

A new class of peer-to-peer (P2P) systems providing real-time video streaming over the Internet is fast emerging and gaining popularity. Several commercial P2P streaming systems such as PPLive, SOPCast, TVAnts and TVUPlayer, just to mention the most popular ones, are already available on Internet.

P2P-TV systems may contribute to revolution the TV broadcasting by allowing ubiquitous access to a practically unlimited number of channels. This represents an important step forward in the direction of the Anything–Anyone–Anywhere–Anytime ubiquitous communications of the future Internet [1].

The adoption of a P2P paradigm reduces the network costs, pushing complexity from the network to the users, while helping to relieve the bandwidth cost burden at the servers. Although from the users' as well as from the server points of view this class of P2P applications has useful and interesting characteristics, the network operators' have serious concerns about the capability of the Internet to support large scale P2P-TV systems (mainly due to the potential high bandwidth requirements, large number of involved users, and the intrinsic inelasticity of video traffic). These concerns are confirmed by some news reports, see for instance [2]. It seems, therefore, rather urgent to have a better understanding of the potential impacts that these applications may entail on the underlying transport network. Since the most widely deployed commercial systems follow a closed and proprietary design, only an experimental and black box characterization of traffic injected by such systems is in general possible; we emphasize, indeed, that approaches requiring to partially reverse

\* Corresponding author. Tel.: +39 011 090 4167; fax: +39 011 090 4099.

E-mail addresses: [leonardi@tlc.polito.it](mailto:leonardi@tlc.polito.it) (E. Leonardi), [mellia@tlc.polito.it](mailto:mellia@tlc.polito.it) (M. Mellia), [michela@tlc.polito.it](mailto:michela@tlc.polito.it), [michela@polito.it](mailto:michela@polito.it) (M. Meo).

engineering complex P2P-TV systems are viable at much larger cost, and only in a few cases. Moreover, to develop new architectures and algorithms that improve the “network friendliness” of such applications [3,4], it is necessary to understand how current applications react to different network conditions and scenarios. Do they implement any congestion control algorithm? How do they react to packet drop? What is the impact of increased end-to-end delay?

In this paper we propose a testing methodology and test-bed experiments to assess how these applications react to different network conditions, like available bandwidth, loss probability, delay; both network load, and user perceived quality of service, should then be measured. Applying our methodology, we test and compare four popular P2P-TV applications, namely PPLive, SOPCast, TVAnts and TVUPlayer. All selected applications adopt the “mesh-based” approach [1], in which peers form a generic overlay topology to exchange chunks of data.

A number of lessons can be learnt from the results shown in the paper.

- All applications are effective in trying to overcome network impairments; for example, all applications avoid low quality paths by carefully selecting peers to download from, the applications can provide an almost unaffected quality even in presence of severe packet losses (up to 10–20%).
- The applications act as if the congestion always affects a subset of paths, and they keep on looking for good paths. Thus, when the access link is congested so that all paths are affected, the applications aggressively download data trying to receive the video stream. This might be quite harmful for both the network and other applications or users sharing the same access link.
- All applications implement a memory based algorithm that tracks good and bad neighbor peers, while no change is observed in the mechanisms to create the neighbor set.
- The coexistence of P2P-TV applications with elastic background traffic, such as the one carried by TCP connections, is critical in some cases.
- The applications tend to fairly share the bandwidth among peers connected through the same access link. However, if this link is congested and the available bandwidth is scarce, the fairness might translate into a useless waste of resources and no one enjoying the service.

The paper is organized as follows. We start by summarizing the related work in Section 2. Then, the measurement setup and methodology are defined in Section 3. Section 4 presents the results in which network impairment affects the incoming traffic from all peers, in Section 5 scenarios in which “good” and “bad” peers coexist are analyzed to investigate the ability of the applications to correctly handle them; scenarios in which the P2P-TV applications are sharing the access bandwidth with TCP connections are instead analyzed in Section 6 and the case of two clients sharing the same access link is made in Section 7. Finally, Section 8 summarizes our findings.

## 2. Related work

To the best of our knowledge, this is the first experimental work on P2P-TV systems exploring how such systems react to different network conditions. In a previous paper, we performed a similar characterization considering Skype [5], in which the focus was on the voice traffic sent/received by a Skype client.

Considering more general experimental results about P2P-TV systems, the research community has given a lot of attention to understand application internals [6–16].

A few works [6–8], relying on the implementation of an active crawler, focus on a single system. These approaches face the daunting task of partially reversing the engine of the P2P-TV 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. In particular, Hei et al. [6] investigate PPLive, whereas Li et al. [7] focus on the commercial re-engineer of Coolstreaming, and Wu et al. [8] consider UUSEE. All these papers mainly provide a big picture of the considered system, focusing on metrics such as the number of users, their geographical distribution, the session duration, the packet size distribution. None of the above mentioned papers considers the particular aspects we are interested into, i.e., the way the system reacts to evolving network conditions.

Other works, such as [9,10], instead, study specific aspects of a P2P streaming system. For instance, Vu et al. [9] give some preliminary results on the node degrees of popular versus unpopular channels in PPLive. Authors in [10] investigate the stability of nodes in PPLive, and devise schemes to identify the most stable nodes.

Quality of service is of concern in [11,12]. Authors in [11] 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 [12] focus on similar metrics but exploit logs made available from an (unspecified) commercial P2P streaming system.

Authors in [13] analyze and compare PPLive and SOPCast investigating the time evolution of different metrics, like transmitted/received bytes, number of parents and children, etc. In [14], on the contrary, a comparative evaluation of four commercial systems (namely PPLive, PPStream, SOPCast and TVAnts) is proposed showing 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 [15] PPLive, SOPCast and TVAnts systems are analyzed. A systematic exploration of the mechanisms driving the peer-selection in the different systems is performed. At last, in [16] a simple experimental analysis of PPLive and Joost is presented to evaluate the characteristics of both data distribution and signaling process for the overlay network discovery and maintenance.

A preliminary version of this paper has been presented at Infocom 2009 [17]. In this paper, we extend the analysis of [17] by investigating the transmitted traffic, the coexistence of P2P-TV and TCP traffic, and fairness towards clients sharing the same access link. Some results of the

previous version are omitted here, to leave room for the new results.

### 3. Methodology

The aim of this work is to study how P2P-TV applications react to different network scenarios. Given that all successful P2P-TV applications follow a proprietary and closed design, we have to follow a “black-box” approach. We therefore setup a testbed, in which clients running the Application Under Test (AUT) are connected to a Linux router, which is connected to the Internet via our Fast-Ethernet based campus LAN. The router itself is then used to enforce particular network conditions. In particular, we used the Linux Network Emulation functionality `netem` coupled with the Token Bucket Filter `TBF`. This allows us to emulate the properties of wide area networks, controlling available bandwidth, delay, loss, duplication and re-ordering of packets routed through the router. Other PCs are connected to the testbed to inject background traffic.

Note that, since we run real on-field experiments, our control on the experimental set-up is limited to the network interface of the router. This implies that the global network conditions are unknown and that possible effects due to congestion, loss, delay inside the Internet are superposed to the effects “artificially” introduced at the router under our control.

Two packet level traces are collected at the router: the first one logs all packets sent/received by the network interface that connects the router to the Internet; the second one logs all packets sent/received by the network interface that connects the PC running the AUT. Packet level traces are then post-processed to obtain the desired measurements. In this paper, we report results considering the *average received bitrate* measured in small time windows (set to 1 min); the received bitrate is denoted by  $r(t)$ , with  $t$  the time at the end of the measurement window, and it is evaluated at the application layer, i.e., neglecting transport, network and data-link overheads. Similarly, in some cases we report the *average transmitted bitrate*  $s(t)$ , intended as the bitrate transmitted by the peer interface;  $s(t)$  is also measured in 1 min long intervals. The number  $n(t)$  of peers that exchanged packets with the AUT during a time interval, i.e., the *number of active peers*, is evaluated as well; for  $n(t)$  we use shorter time windows of 5 s.

Finally, the PC running the AUT is used to capture the video stream that is received and to dump it on a file by means of a video grabber utility. To evaluate the video quality of the received stream, we cannot apply any standard reference-based technique, since they all rely on the comparison of the received and original video (being impossible to get the latter one). All selected applications generate 378 kbps streams encoded using the Microsoft VC-1 encoder; a typical bitrate of 450–500 kbps is received by the AUT, so that 100–150 kbps of additional overhead is required by the applications to successfully deliver the stream (not including transport, network and data-link headers). Since the codec relies on proprietary design, it is difficult to evaluate the quality of the received stream.

We are therefore forced to estimate the stream quality by simply counting the number of errors a decoder has to deal with when decoding the stream. In particular, we decoded each file using `ffmpeg` utility which reports a detailed list of corrupted video I-frames. Those are major impairments that will affect the video quality for several frames, i.e., up to when a good I-frame is received (usually several seconds later). Similarly, the audio stream decoding errors are evaluated as reported by `ffmpeg`. In the following, we report therefore the number of corrupted I-frames and audio blocks as quality index. While this allows only a qualitative evaluation of the stream quality, it allows us to fairly compare different applications.

#### 3.1. Scenario definition

The parameters we consider in this paper are the following:

- $c$ : Capacity limit.
- $l$ : Packet loss.
- $d$ : Delay.

The set  $L(t) = \{c(t), l(t), d(t)\}$  specifies the state of the controlled link during each instant of the experiment – we restrict to the cases in which only one of the three above parameters is evolving with time and we denote with  $p(\cdot)$  its *profile* over time.

As profile  $p(\cdot)$  we select a step function, with initial value  $p_0$ , increments  $I$ , and step duration  $\Delta T$ , so that

$$p(t) = p_0 + I \sum_{n=1} H(t - n\Delta T) \quad (1)$$

in which  $H(t)$  is the Heaviside step function

$$H(t) = \begin{cases} 0 & t < 0, \\ 1 & t \geq 0. \end{cases} \quad (2)$$

If  $I$  is positive,  $p(\cdot)$  corresponds to an *increasing* profile; negative values of  $I$ , on the contrary, generate *decreasing* profiles. A null increment,  $I=0$ , finally, leads to a *constant* profile.

The impairment defined by a scenario can affect all sent/received packets, so that *global* impairment is imposed, or only a subset of sent/received packets, so that *per peer* impairment is imposed; for example the scenario can affect a single peer, a subnet, an Autonomous System, or any generic subset of IP addresses.

#### 3.2. Considered general setup

We performed several experiments considering different scenarios and profiles, during various time periods and with clients tuned on different TV channels. We collected a total of more than 300 h of experiments. In this paper, we report a subset of the most representative results. We report the results of selected experiments and not some statistical results, because the experiments involve a number of variables that we cannot control, nor know, such as the number of users/peers watching the same channel, the distribution and characteristics of the peers,

the network conditions over the many involved paths. These uncontrollable (and unknown) variables make it impossible to repeat exactly the same experiment to get a statistical result. However, the fundamental behaviors described and analyzed in the paper were observed in all the experiments we performed. An extended set of results is available in [18].

For what concerns the scenarios, we focus on cases in which the download link is controlled, while the upload link capacity is limited to  $c_u(t) = 200$  kbps, unless otherwise specified. This indeed allows us to evaluate the application behavior when the peer has not enough capacity to act as an “amplifier”, i.e., to serve many peers; this is the typical condition of ADSL users.

### 4. Global impairment

#### 4.1. Effect of available capacity

In the first set of experiments, the download available capacity  $c(t)$  is imposed. Results are shown in Fig. 1 and are organized in the following way. The two largest plots in the left part of the figure report the bitrate  $r(t)$  versus time for profiles with either decreasing or increasing capacity limit (on the top and bottom plots, respectively). The eight small plots on the right part of the same figure depict the number of corrupted audio and video frames for the same experiments; each plot refers to a specific application.

Let us start by considering the decreasing profile. Every  $\Delta T = 5$  min, the available bandwidth is decreased by a  $l = -50$  kbps, starting from an initial value of  $c_0 = 800$  kbps. The average bitrate evaluated using 60 s time intervals is reported for all applications on left top plot of the figure. The experiment lasted 1 h, after which the available capacity was set back to 800 kbps. All applications have similar behavior: the bitrate remains basically constant for all the time the available capacity is larger than the data rate,  $c(t) > r(t)$ . When the capacity bottleneck kicks in, all the applications react by increasing the download data rate. Consider, for example, TVAnts, which exhibits the largest value of the bitrate. The normal data rate is about 600 kbps; when the capacity limit reaches 650 kbps, the receiver starts suffering the bottleneck (due to traffic burstiness), and it reacts by requesting other peers to send more traffic; the download rate becomes larger than 800 kbps. As the capacity limits  $c(t)$  decreases, the received rate decreases too, being always about twice the available capacity, i.e., the offered load to the congested link is about 2,  $r(t)/c(t) \approx 2$ . Other applications show similar behavior, with smaller values of the offered load in congested conditions; in particular, TVUPlayer exhibits the smallest overload factor  $r(t)/c(t) \approx 1.3$ .

Looking at the last 10 min of the experiment, when the capacity returns to high values, an unexpected, strange behavior is observed. Indeed, since  $c(t)$  is larger than the normally required capacity,  $r(t)$  should take again the typical values that can be observed when traffic is not bottle-

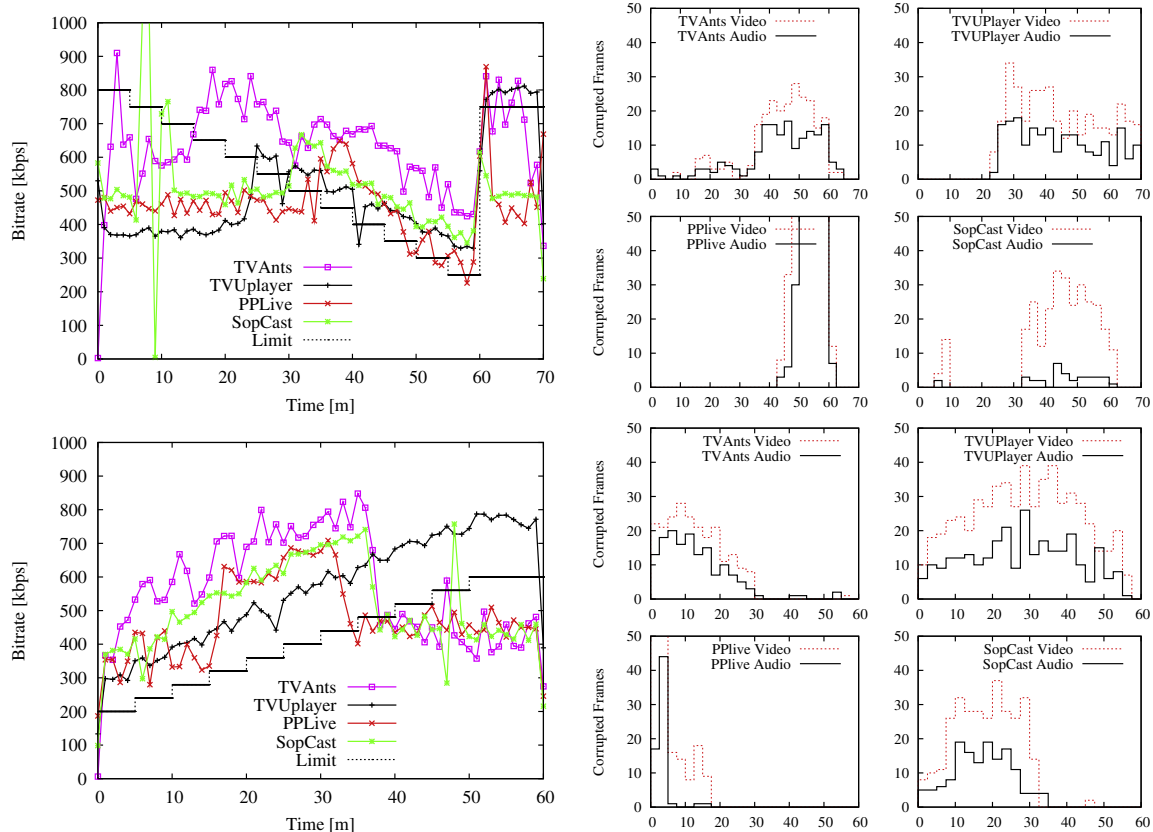


Fig. 1. Left plots: received bitrate for decreasing (top) and increasing (bottom) available bandwidth. Right plots: percentage of corrupted audio/video frames, one plot for each application.

necked. While this is true for PPLive and SopCast, both TVAnts and TVUPlayer keep on receiving at a rate which is about twice as large as the normal receiver rate. This maintains the bottleneck offered load higher than 1, so that audio/video impairment is observed up to the end of the experiment. Indeed, looking at the number of corrupted frames reported in the top right part of Fig. 1, audio/video impairment starts to show up as soon as the bottleneck kicks in, and it does not always disappear when the bottleneck capacity is set back to 800 kbps. TVAnts shows the longest period during which corrupted frames are observed, while, PPLive can cope with downlink capacity as small as 400 kbps without any audio/video error.

Results for the case of an increasing capacity profile are reported in bottom part of Fig. 1. The AUT is started now in scarce bandwidth conditions ( $b_0 = 200$  kbps), and  $I = 40$  kbps increments are applied every  $\Delta T = 5$  min. All applications react to the adverse condition by trying to download much more data than the available capacity; also in this case  $r(t)/c(t)$  varies from 1.3 (for TVUPlayer) to 2 (for TVAnts). Only when the available capacity is large enough to sustain the minimal download rate, all applications but TVUPlayer decrease  $r(t)$  to their typical values. This is reflected by the disappearance of audio/video impairment, as shown by the right plots. Again, TVUPlayer suffers major impairment, even when the bandwidth grows to large values.

We can conclude that P2P-TV applications do not correctly perform congestion control, in scenarios in which

peer access links get congested. They all try and react to limited access capacity by increasing the redundancy (by FEC or ARQ mechanisms) and, thus, the download rate. This may be potentially harmful for both the network and other applications sharing the congested link. Note that a single congested link may also be present when the unique peering link between a stub ISP and the rest of the Internet is congested. If this happens, P2P-TV applications may react as in the previously presented cases, causing further network problems and congestion.

#### 4.2. Effect of loss probability

The second set of results we report aims at investigating the impact of loss probability on the AUT. Organized in a similar way as the previous figure, Fig. 2 shows the receiver rate for increasing (top plots) and decreasing (bottom plot) packet loss probability profiles. The right y-axis of the larger plots reports the percentage of losses,  $l(t)$ , that varies in time steps of  $\Delta T = 5$  min, with loss increment  $I = 5\%$  ( $l_0 = 0\%$ ). In this case also, all the applications react to increasing packet losses by increasing the bitrate  $r(t)$ . By doubling its received data rate for  $l(t) > 35\%$ , TVUPlayer is the most aggressive application, while PPLive shows the smallest increase. Looking at the corresponding number of audio/video corrupted frames, it is impressive to observe that all applications achieve very good video quality as far as  $l(t) < 25\%$ . In particular, it is worth noticing that SopCast can cope with 25% packet loss probability with only

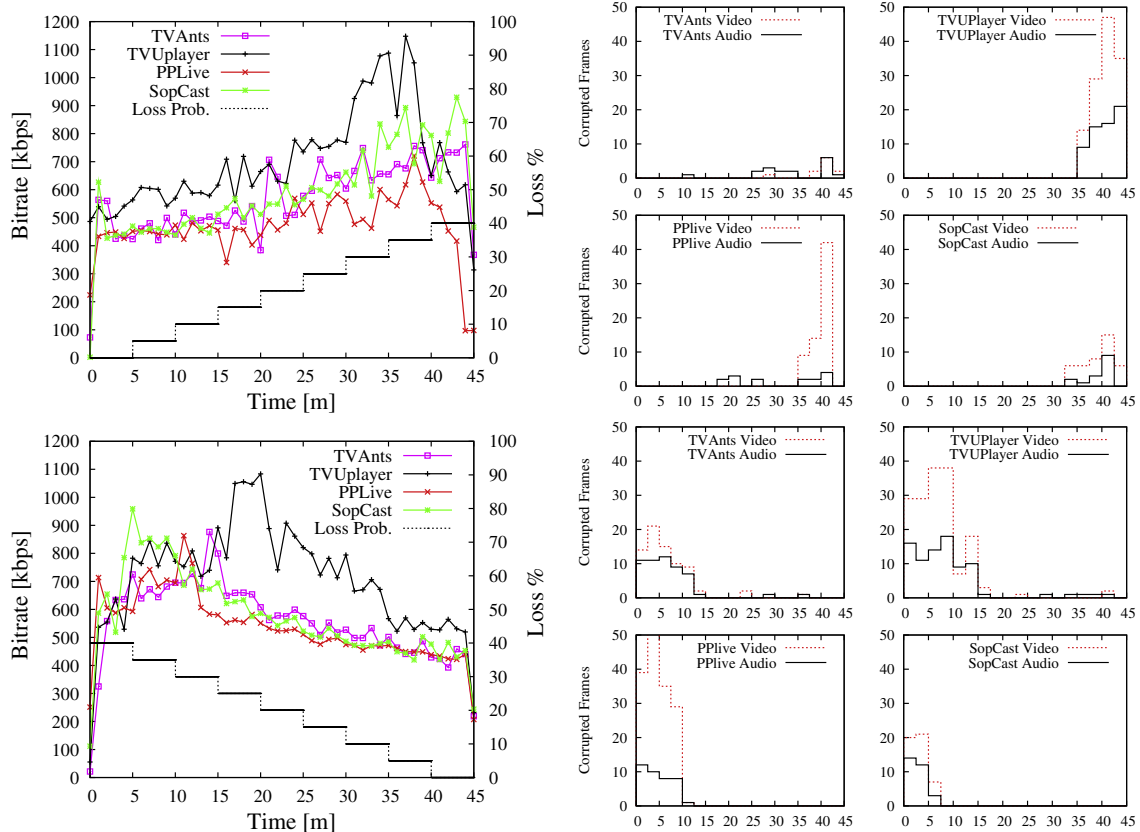


Fig. 2. Left plots: received bitrate for increasing (top) and decreasing (bottom) loss probability. Right plots: percentage of corrupted audio/video frames, one plot for each application.

about 100 kbps of additional data rate. TVUPlayer exhibits similar performance, but at a much higher cost that accounts to up to 600 kbps of additional data rate.

Similar observations can be drawn by looking at the decreasing packet loss probability scenario reported in bottom plots of Fig. 2. In this scenario,  $I = -5\%$ ,  $i_0 = 40\%$ ,  $\Delta T = 5$  min.

These results allow us to conclude that all applications react to packet losses by trying to recover them, using some kind of ARQ mechanism that causes an increase of the received traffic. While this is very efficient in repairing the audio/video stream, it comes at the expense of an offered load that can be as large as twice the rate in normal conditions. This definitively confirms that P2P-TV applications do not perform, in general, any congestion control.

### 4.3. Effect of delay

We now consider the effect of increasing and decreasing delay profiles. Fig. 3 reports the results for the received bitrate of the increasing (left plot,  $I = 200$  ms,  $d_0 = 0$  ms,  $\Delta T = 5$  minutes) and decreasing (right plot,  $I = -200$  ms,  $d_0 = 2000$  ms,  $\Delta T = 5$  min) profiles; plots about the number of corrupted frames are not reported for the sake of brevity.

Results about the increasing delay case show that the applications can manage quite well slow variations of the delay; they can stand up to 1.5 s of additional delay without any significant variation of the received bitrate (and any audio/video error). The applications start suffering the delay when it reaches almost 2 s, which is quite large; PPLive and TVUPlayer seem the most delay sensitive applications.

Interestingly, the applications suffer more for large values of the delay at the start up (see right plot Fig. 3). It is probably difficult for the applications to successfully create the neighbor list. Indeed, since the additional delay applies also to packets carrying signaling information, signaling dialogs are probably hardly completed with large values of the delay. In the decreasing profile, delay has to decrease below 1.2 s to allow the applications to start receiving the video stream. Again, PPLive seems the most sensitive

application: additional delay should be smaller than 1 s to allow it to work.

### 4.4. Number of active peers

Finally, upper plots of Fig. 4 report the number of active peers,  $n(t)$ , for the previously described scenarios with increasing and decreasing capacity limit and loss probability.

The network conditions have no impact on the behavior of  $n(t)$ , which repeats regularly during the whole experiment. On the contrary, different experiments show different absolute values of  $n(t)$ ; indeed, the absolute values change with channel popularity and time of the day so as to reflect the population of available peers. PPLive, that is an extremely popular application, has always the largest number of active peers.

The same observations can be made by considering the evolution of the number of total contacted peers, which is independent from the network conditions; the associated plots are reported in the right plots of Fig. 4. Notice also that the periodic peaks clearly visible in most of the applications are due to periodic keep-alive messages used to exchange signaling information.

These results allow us to conclude that the internal algorithms each application implements to discover, create and maintain the overlay, are insensitive and do not adapt to network conditions; network conditions influence only the video stream distribution mechanisms.

## 5. Per peer impairment

### 5.1. Effect of available capacity

We now investigate the capability of the AUT to cope with scenarios in which only a subset of peers is affected by network impairment, so that “good” and “bad” peers coexist. The goal is to verify if the AUT can identify the set of “good” peers to download from. In particular, for the results reported here, the imposed network impairment affects only peers having an odd IP address. The rationale behind this choice is to have the peer population


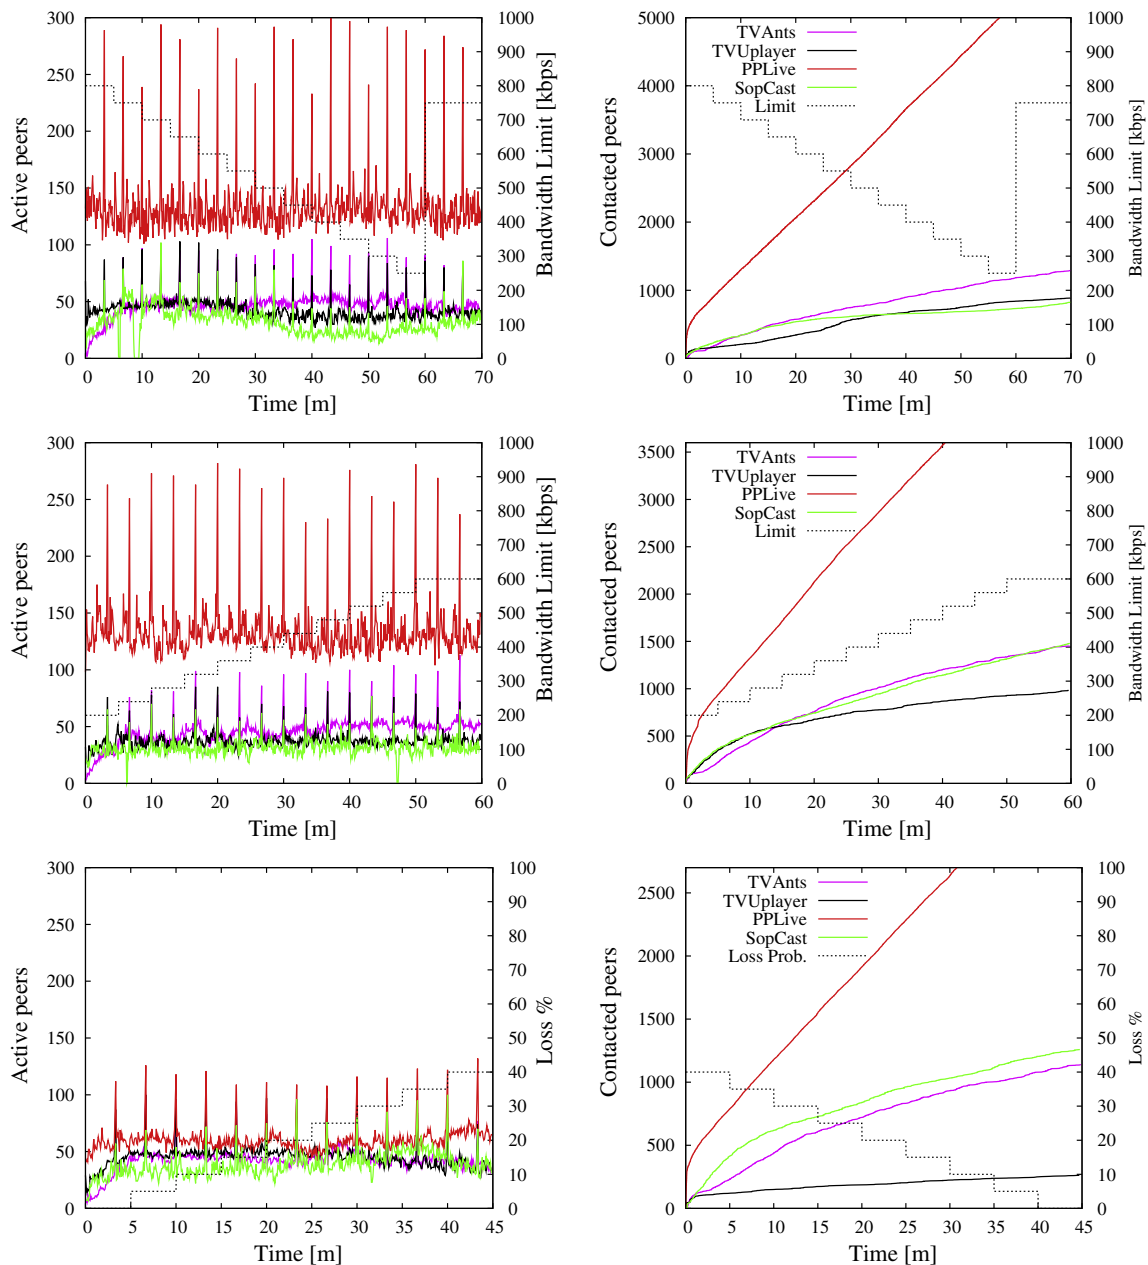


Fig. 3. Received bitrate for increasing (left plot) and decreasing (right plot) delay.



**Fig. 4.** Number of active peers (left) and total contacted peers (right) for decreasing or increasing available bandwidth and packet loss probability (from top to bottom).

split into two equally large subsets: odd peers, affected by network impairment, and even peers, not affected.

The plots of Fig. 5 report results considering a decreasing capacity limit profile. In particular, the profile  $c(t)$ , that is imposed to odd peers only, starts from  $c_0 = 400$  kbps, and every  $\Delta T = 5$  min a further bandwidth decrease of  $I = -25$  kbps is applied. After 60 min, the available bandwidth is again set to the initial value. Each plot reports, for a given application, the bitrate received from even and odd peers and the total received bitrate. The imposed profile is also given for completeness. Again, all applications exhibit similar behavior: during the initial phase there is no preference in receiving data from even or odd peers: they equally contribute to the total download rate. As soon as the bandwidth limit kicks in, reducing the

performance of odd peers, the applications preferentially download data from even peers. The preference is stronger for SopCast (rightmost plots) for which even peers contribute to 80–90% of download rate. TVAnts, on the contrary, adapts less than the other applications to these network conditions.

Fig. 6 reports results for an increasing bandwidth profile applied to odd peers only; the profile has parameters:  $c_0 = 125$  kbps,  $I = 25$  kbps,  $\Delta T = 5$  min. Similar considerations hold: All applications quickly identify the adverse capacity constraints affecting odd peers, so that even peers provide larger contribution to the total download bitrate. In particular, TVUPlayer (top right plot) has a very accurate control mechanism that allows it to quickly identify the changing network conditions. PPLive and SopCast also

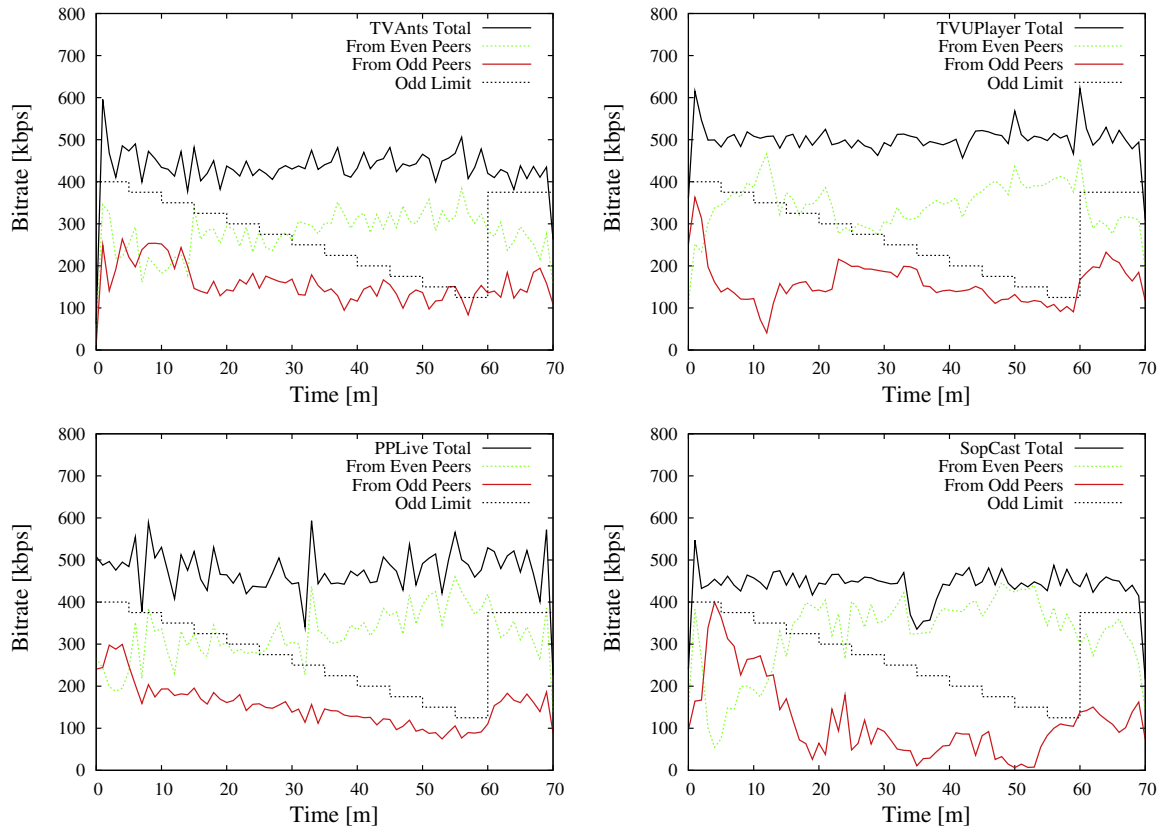


Fig. 5. Received bitrate for decreasing capacity limit. Odd peers only are affected by the impairment.

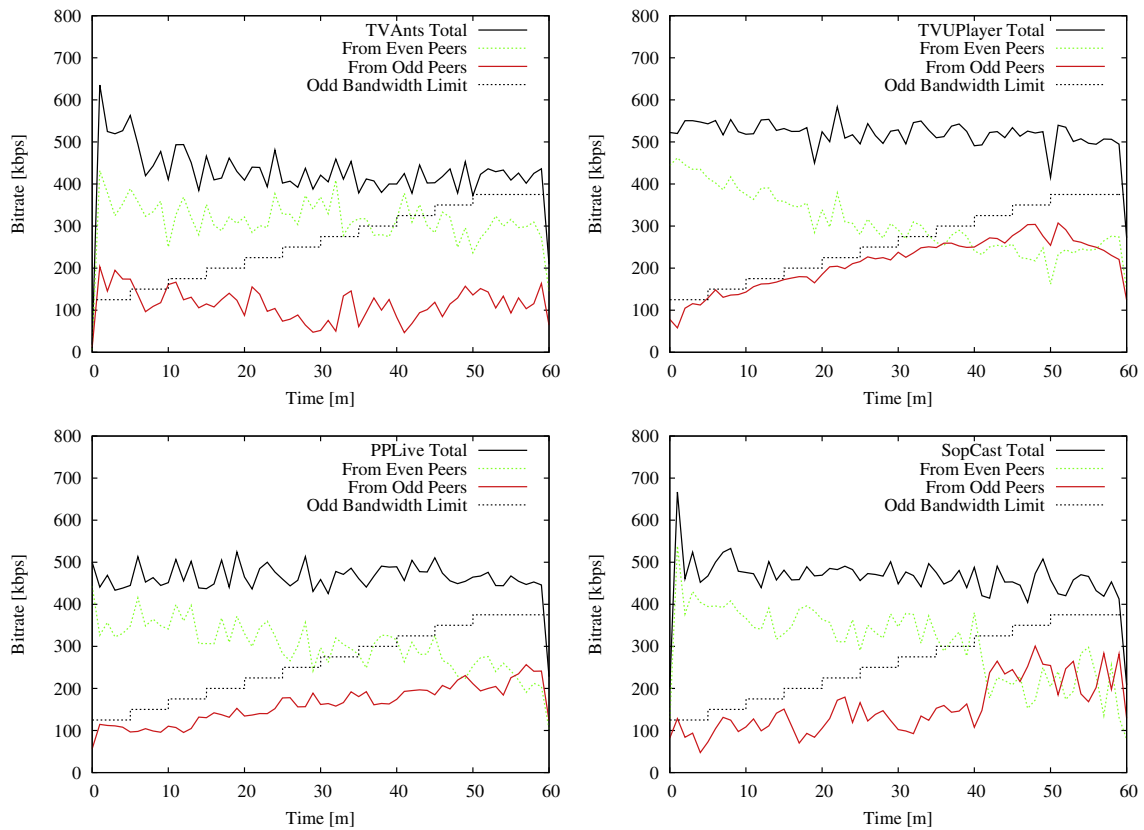


Fig. 6. Received bitrate for increasing capacity limit. Odd peers only are affected by the impairment.

exploit the additional bandwidth of odd peers that becomes gradually available, but a longer transient phase is required. Finally, TVAnts ignores the additional bandwidth, since about 70% of traffic is received from even nodes during the whole experiment.

In all cases, all the applications receive the minimum required amount of data that guarantees them to decode the audio/video streams without suffering any error.

### 5.2. Effect of loss probability

Figs. 7 and 8 report results considering increasing and decreasing profiles of  $l(t)$ , respectively. Let us start by considering increasing loss probability; plots on top row of the figure refer to a profile with  $l_0 = 0\%$ ,  $I = 5\%$ ,  $\Delta T = 5$  min. In this case, different reactions are observed. TVAnts has a strong preference to receive from even peers starting from  $l(t) \geq 5\%$ . TVUPlayer has an on/off behavior, so that no preference is shown up to  $l(t) = 10\%$ , and, then, 95% of data is received from even nodes only. PPLive is ignoring the loss impairment up to  $l(t) \geq 20\%$ , after which data are preferentially received from even nodes (but still 20% of traffic is received from odd peers when  $l(t) = 40\%$ ). Finally, SopCast shows a more irregular and uncontrolled behavior which causes a preference toward odd peers until  $l(t) = 15\%$ , after which about 90% of data is received from even nodes only.

Consider now Fig. 8, which reports results for a decreasing profile of  $l(t)$  (with  $l_0 = 40\%$ ,  $I = -5\%$ ,  $\Delta T = 5$  min). Since

all applications start in very unfavorable conditions for odd nodes, most of the traffic is received from even nodes. In particular, TVUPlayer constantly receives 98% of traffic from even nodes only, even when  $l(t)$  becomes small. Similarly, TVAnts and SopCast exhibit a very stable preference during the whole experiment duration, with TVAnts trying to received 15–20% of traffic from impaired peers. PPLive, on the contrary, keeps on receiving 20% of traffic from odd nodes, percentage that goes up to 50% when  $l(t) \leq 10\%$ . This confirms that PPLive is capable of coping with high packet loss rates (as already noticed in Fig. 2), hinting to an effective FEC algorithm.

### 5.3. Other results

We performed other similar tests, targeting with impairment: a particular peer, IP subnetworks, and Autonomous Systems. All the experiments showed consistent results, so that preference is given to good peers. We, thus, conclude that all applications implement a per-peer preference mechanism that is used to select the subset of good performing peers. While internal algorithms are unknown, the presented results suggest that the applications are using different algorithms. Due to space constraint, we do not report the figures referring to other scenarios we tested, and refer the reader to [17,18] for more details.

Considering other possible impairment, experimental results shows that:

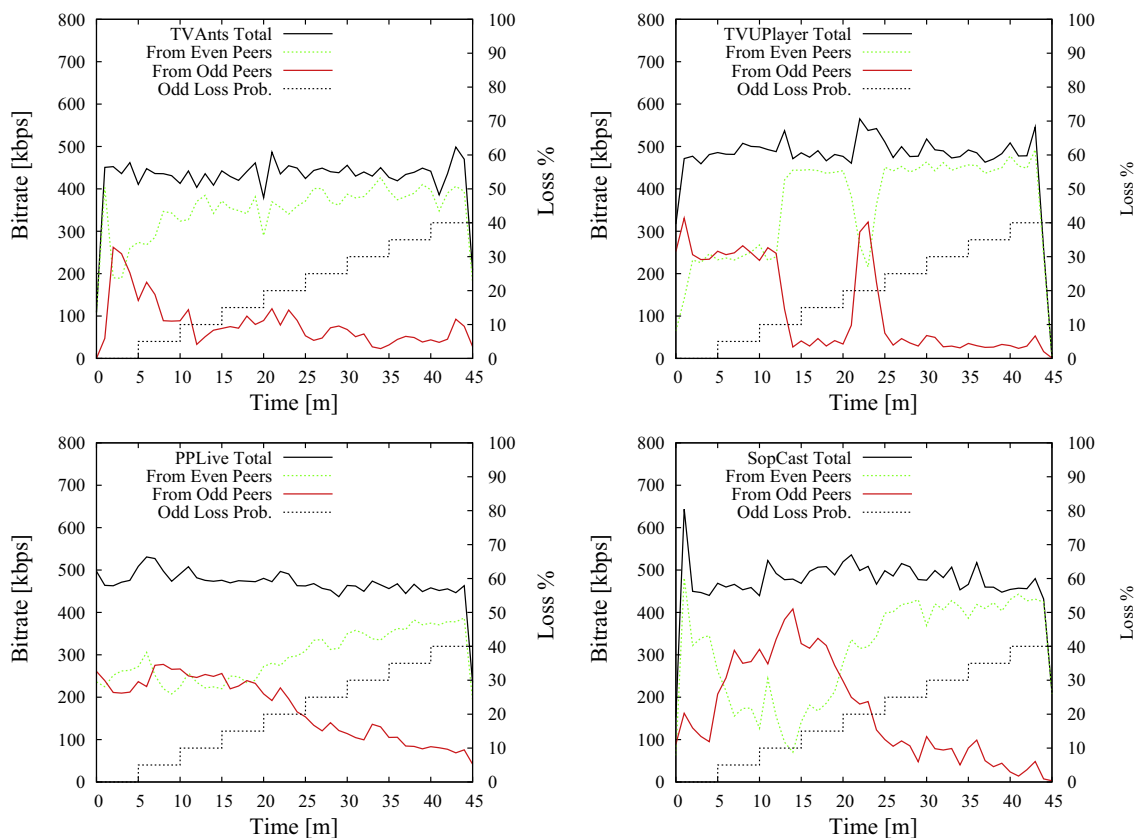


Fig. 7. Received bitrate for increasing loss probability. Odd peers only are affected by the impairment.

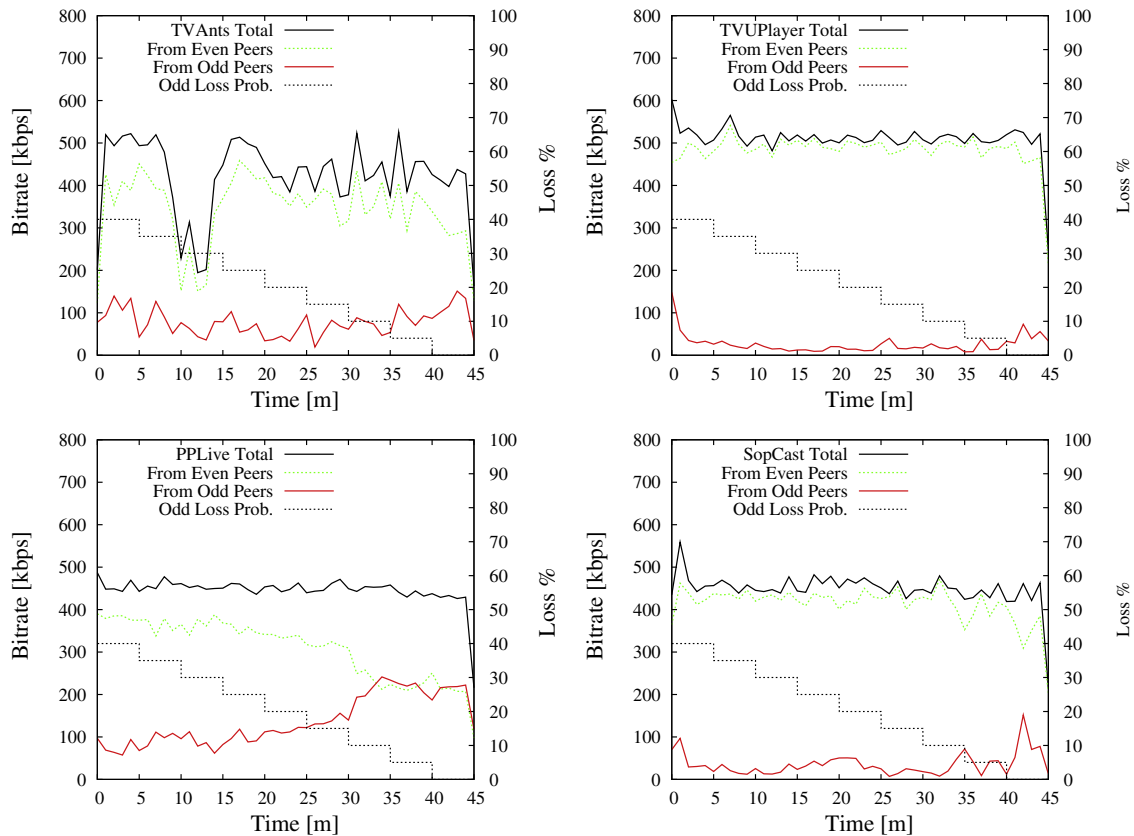


Fig. 8. Received bitrate for decreasing loss probability. Odd peers only are affected by the impairment.

- *Delay preference*: All applications are very sensitive to additional delay, so that content was almost exclusively retrieved from good peers as soon as the additional delay increased.
- *Hop distance*: When we artificially decreased the IP TTL value to artificially inflate the path hop distance, no application showed appreciable bias. This clearly indicates that hop-count information is not exploited by the peer selection algorithm.
- *Number of active peers*: Results show that the AUT keeps contacting odd peers, since  $n(t)$  is not correlated with  $c(t)$ ,  $l(t)$  or  $d(t)$ . This hints to control algorithms that react to different network scenarios by carefully selecting the good peers to exchange data with. However, signaling is exchanged with all peers (including “bad” peers) independently from the instantaneous end-to-end network quality. We also verified that all AUTs keep exchanging data with and probing “bad” peers even during very unfavorable conditions. In these cases only few, small packets are sent and (possibly) received. This suggests that only signaling information is exchanged between any two peers that are experiencing adverse network conditions, but the bad peers are not dropped in favor of the good ones.

5.4. Impact on the upstream traffic

In this section, we consider the impact of downlink impairment on uplink traffic. As before, we consider an ADSL-like scenario in which the peer has upstream

capacity equal to  $c_u = 200$  kbps, i.e., the peer contribution to video distribution is low.

Fig. 9 shows the total bitrate in the upstream for the considered applications when the downlink undergoes the same bandwidth decreasing profile previously presented. When the impairment on the downlink becomes severe, the peer reduces the number of chunk transmissions, that is, it reduces its contribution to the video content distribution. The transmitted information on the uplink is then basically limited to the signaling information.

To further investigate this phenomenon, we separately consider signaling and data information. Similar to what




Fig. 9. Transmitted bitrate for decreasing bandwidth limitation.

is done in [6,14,16], signaling and data information is distinguished based on the packet size: packets whose size is smaller than 300 B are assumed to carry signaling information, larger packets carry data. Fig. 10 reports the size of the observed packets for the four applications; green and black markers refer to signaling and data packets, respectively. Information packets are transmitted in the first part of the experiment, when the downlink bandwidth limitation is not severe, so that the peer can retransmit video chunks to its neighbors. When downlink limitation becomes so strong that the peer is perceived by its neighbors as poor performing, the peer reduces its contribution to video distribution (in the figure, few data packets are transmitted when severe downlink limitation applies). On the contrary, the bandwidth limitation profile marginally affects signaling information, whose pattern is basically unchanged for the whole duration of the experiment. Indeed, the signaling information that is needed by the peer to guarantee the correct reception of the video is continuously transmitted, regardless the peer performance. The same observations can be derived by focusing on the signaling and data information bitrate, reported in Fig. 11. For all the applications, signaling information requires more or less the same amount of bandwidth during the whole experiment, while data information diffusion depends on the network conditions (i.e., downlink bandwidth limitation). It is interesting that all the applications require roughly the same amount of signaling information,

between 50 and 100 kbps, Sopcast being the application with the lowest signaling overhead.

### 6. Sharing the access bandwidth with TCP flows

In this section we analyze the effects of possible access bandwidth competition between persistent downstream TCP flows and P2P-TV applications. We consider a static setup in which the downstream access capacity is constant and equal to 2 Mbit/s, while the upstream bandwidth is, as usual, limited to 200 kbit/s. The P2P-TV application is running from  $t = 0$ , while a total of 5 TCP flows are started, one every  $\Delta T = 300$  s. In particular, a PC connected on the same subnet of the P2P-TV client starts downloading a 5 GB long file from a Web server connected immediately after the Linux router.

Fig. 12 shows the results for the four considered applications, reporting both the average download and upload bitrate for the P2P-TV application, and the total TCP throughput. First observe that the received bit rate  $r(t)$  is mainly insensitive to the presence of concurrent traffic for all the four applications. This behavior is expected in light of the fact that P2P-TV applications are non-elastic. More surprisingly, TCP connections fail to efficiently exploit the available bandwidth. In the case of TVAnts and PPLive, in particular, the TCP connections are severely impacted by the presence of P2P-TV traffic, so that the available downlink bandwidth cannot be successfully exploited.

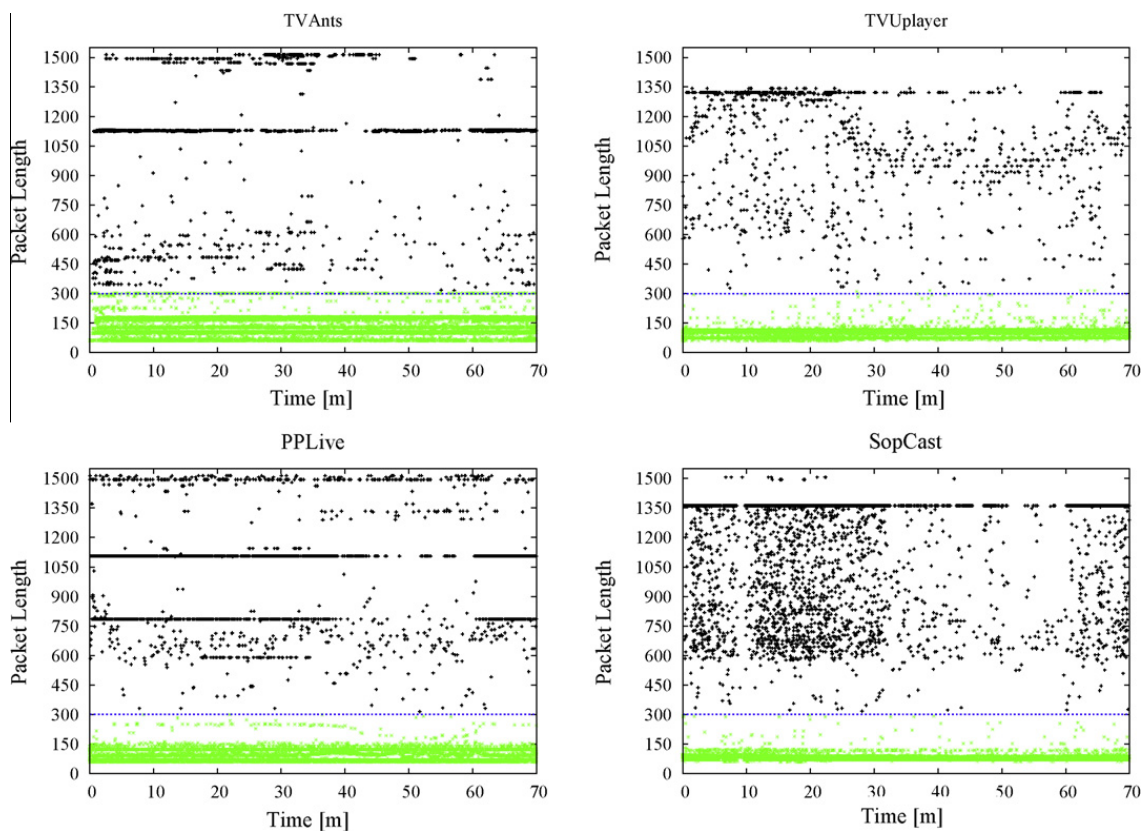


Fig. 10. Packet size for decreasing bandwidth: signaling (green markers) and data information (black markers). The solid line at 300 B represents the threshold to distinguish signaling and data information. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 11. Transmitted bitrate for decreasing bandwidth: signaling and data information.

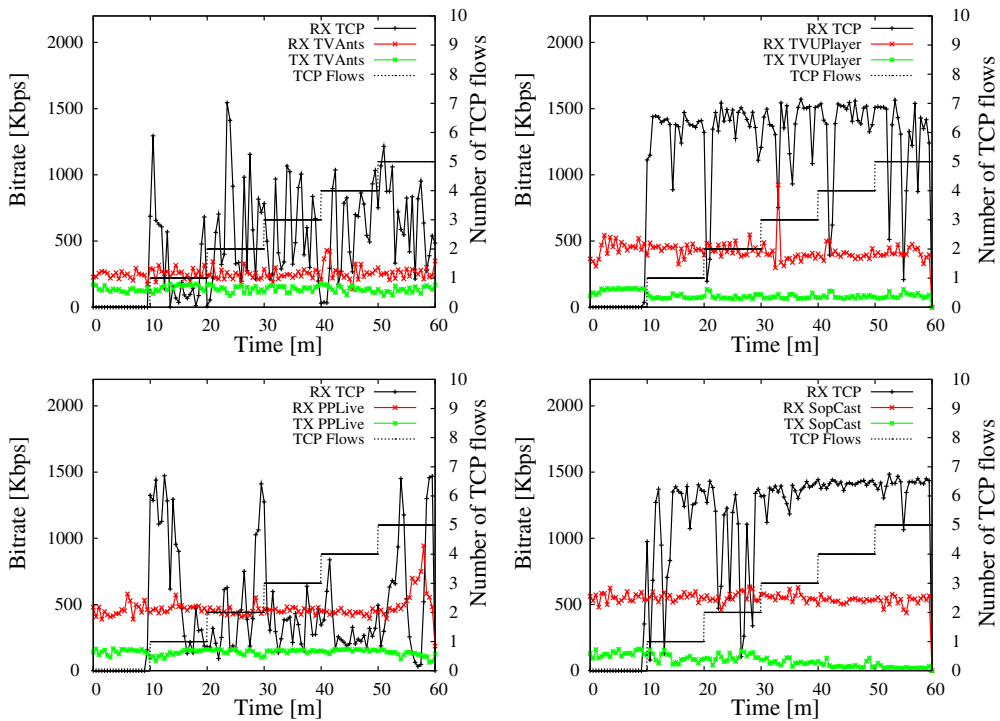


Fig. 12. Received and transmitted bitrate in presence of TCP background traffic. Access link bandwidth parameters:  $c(t) = 2 \text{ Mbit}/2$  and  $c_u(t) = 200 \text{ kbit}/s$ .

This phenomenon is explained by observing the upstream channel, where congestion arises. Indeed, the P2P-TV application keeps uploading traffic, so that the

narrow 200 kbit/s channel becomes congested, impairing TCP ACKs sent by receivers. Congestion on the upstream channel therefore induces ACK losses, preventing TCP from

effectively exploiting available capacity on the downlink path. Notice indeed that the average loss probability in the link is about 10%.

To better understand the interaction between TCP flows and P2P-TV systems, in Fig. 13 we also consider a second scenario in which the upload bandwidth has been increased to 10 Mbit/s. The figure shows that TCP flows succeed in efficiently exploiting the available bandwidth, since no congestion is present on the backward path. However, this time the P2P-TV application suffers for the presence of TCP traffic; note indeed, that the upload bit rate  $s(t)$  significantly decreases at about  $t = 300$  s, when the first TCP flow starts. In this case, congestion arises in the downstream link. Indeed, since TCP flows use all the available bandwidth, some packet loss and significant queuing delay occur. This makes the peer appear less performing to other peers, that might decide not to select it to download the content. While losses do not significantly affect the video quality (being the application capable of coping with packet loss and additional delay, as seen in the previous section), the large delay prevents the peer from effectively redistributing the chunks it gets.

As a conclusion, the coexistence of P2P-TV and background traffic can be rather critical. On the one hand, peers are still capable of correctly receiving the stream, but they cannot significantly contribute to its re-distribution. As a result upload bandwidth of peers, which is a very precious resource for P2P-TV applications, is wasted. On the other hand, not being elastic, P2P-TV traffic can prevent TCP from properly working, especially in presence of scarce bandwidth, as it is typical of today ADSL setups.

### 7. Clients sharing the access link

Finally, in this section, we consider the interaction between two clients of the same application sharing the access link. We perform experiments in which two clients are started at the same time and tuned on the same channel. We then apply a downlink bandwidth limitation profile as the one reported in dashed line of the plots of Fig. 14: the bandwidth is progressively reduced to 400 kbps and then progressively increased again. In the plots of the figure the red and green lines report the bitrate received by the two clients; the aggregate bitrate is reported also, for completeness. A substantial fairness is provided to users, even under severe bandwidth limitation. However, in these conditions, the video quality is severely impaired (as reported in Fig. 15 by video and audio impairments for the case of TVUPlayer). Both the clients undergo the same kind of quality degradation. While at a first sight fairness among clients might seem a positive characteristic of the application, we observe that when the bandwidth is insufficient to support both clients, no one receives an acceptable service. To this extend, in very adverse conditions, bandwidth fairness is not improving the overall service provided to users and implies a waste of bandwidth. The same conclusions can be drawn from the results obtained by starting the clients at different times.

Finally, we did not consider the performance achieved by clients of different applications sharing the same bottleneck link, since we cannot, in this case, provide a fair comparison. Indeed, when two clients run different applications, they belong to overlay networks that are

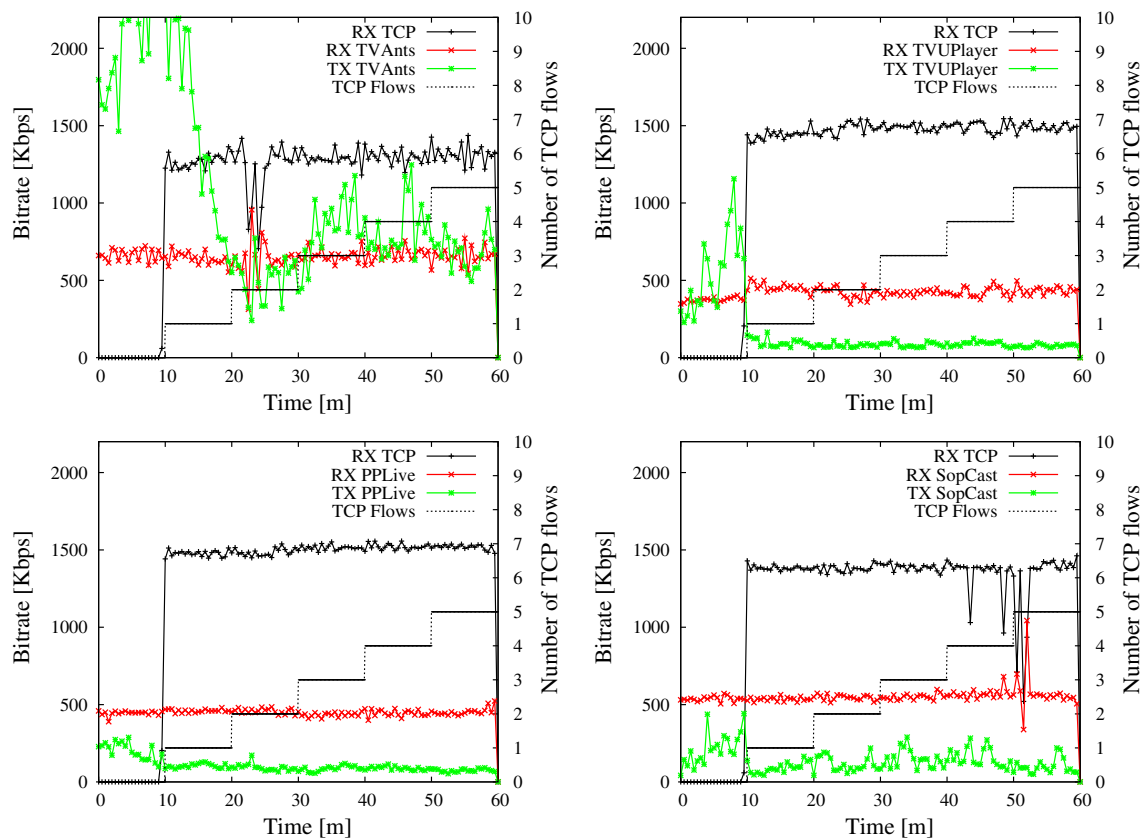


Fig. 13. Received and transmitted bitrate in presence of TCP background traffic. Access link bandwidth parameters:  $c(t) = 2$  Mbit/s and  $c_u(t) = 10$  Mbit/s.

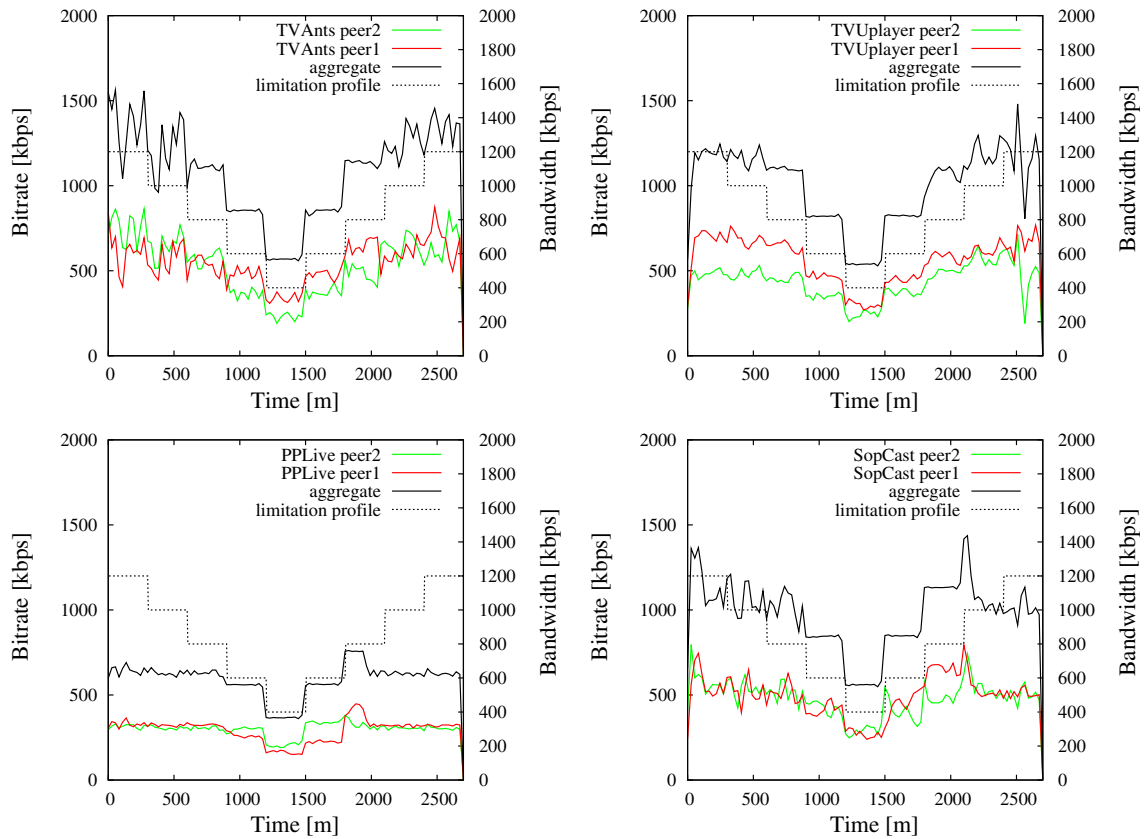


Fig. 14. Received bitrate for two clients concurrently running the same application and sharing the access link, with downstream bandwidth limitation.

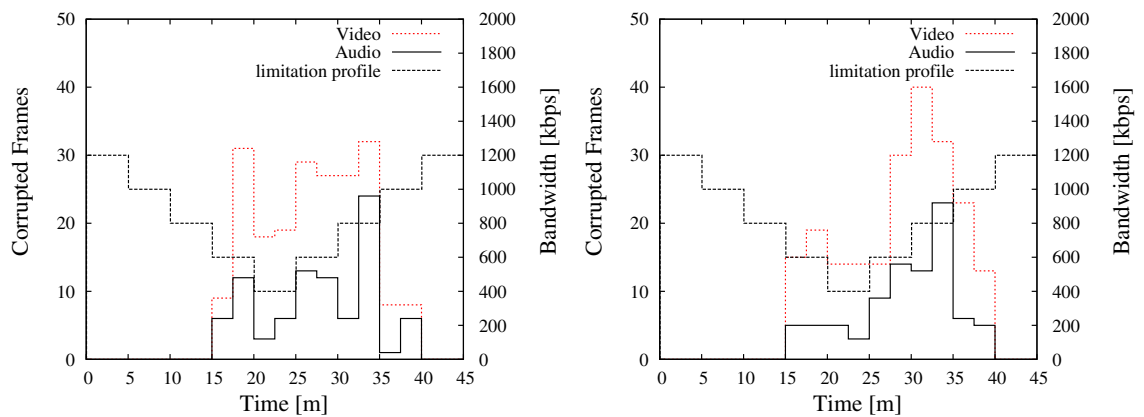


Fig. 15. Video and audio impairment for the two clients in the case of TVUPlayer and downstream bandwidth limitation.

possibly very different: based on the channel popularity and on the content and language of the program, the number, location, and characteristics of the peers in the overlay can be quite different; this makes a fair comparison impossible.

### 8. Conclusions

In this paper, we propose an experimental methodology to investigate the behavior of P2P-TV applications under adverse network conditions. Since most of the successful P2P-TV applications rely on a closed and proprietary design, it is indeed important to understand if these applica-

tions implement algorithms to cope with different and variable network scenarios. In particular, available bandwidth, delay, packet loss probability and presence of background traffic are the most important impairment today applications face in the Internet. We therefore explored how P2P-TV applications react to those parameters, by setting up real test-bed experiments.

We applied this methodology to four P2P-TV applications, namely, PPLive, SopCast, TVAnts and TVUPlayer. By observing the received bitrate and the number of contacted peers, we have shown that all applications effectively react to impairment caused by: (i) lack of bandwidth, (ii) packet loss probability, or (iii) large delay. Applications indeed

successfully select the subset of peers that offer the best performance, disregarding peers on impaired paths. However, in case the bottleneck affects all peers, e.g., it is at the access link, their behavior results rather aggressive, and potentially harmful for both other applications and the network. Interestingly, the control algorithm preferentially operates by selecting the active peers among the neighbors on the overlay, but it does not affect the neighborhood selection, i.e., the overlay topology discovery and setup. The analysis of the upstream traffic has revealed that the peer tries to contribute to the video distribution as far as the network conditions allow it; when the network conditions are bad, the peer acts as a receiver only.

Even if all applications show similar behaviors, some differences arise: TVUPlayer is the fastest and most prompt to react to changing conditions, but sometimes its control algorithm overreacts to dynamic situations; TVAnts, on the contrary, shows a less controlled behavior, which causes the highest overload when resources are scarce, and forces the client to keep downloading from impaired peers.

We have also analyzed scenarios in which P2P-TV applications share the access bandwidth with long lived TCP connections and we have investigated their interaction. In some cases the presence of inelastic P2P-TV traffic may have a negative effect on TCP preventing connections from fully exploiting the available capacity, while in other cases it is the P2P-TV application that suffers the presence of TCP traffic, becoming unable to redistribute the video stream. At last, we have considered clients sharing the same access link and verified that the applications are fair towards these clients even under adverse conditions.

## Acknowledgement

This work was funded by the European Commission under the 7th Framework Programme STREP Project “NAPA-WINE” (Network Aware Peer-to-Peer Application under Wise Network).

## References

- [1] J. Liu, S.G. Rao, B. Li, H. Zhang, Opportunities and challenges of peer-to-peer Internet video broadcast, *Proceedings of the IEEE* 96 (1) (2008) 11–24.
- [2] A. Murray-Watson, Internet groups warn BBC over iPlayer plans, *The Independent*, 12 August 2007.
- [3] E. Leonardi, M. Mellia, A. Horvath, L. Muscariello, S. Niccolini, D. Rossi, Building a cooperative P2P-TV application over a wise network: the approach of the European FP-7 STREP NAPA-WINE, *IEEE Communications Magazine* 46 (2008) 20–211.
- [4] H. Xie, Y. Yang, A. Krishnamurthy, Y. Liu, A. Silberschatz, P4P: provider portal for applications, in: *ACM Sigcomm 2008*, Seattle, WA, August 2008.
- [5] D. Bonfiglio, M. Mellia, M. Meo, N. Ritacca, D. Rossi, Tracking down Skype traffic, in: *IEEE Infocom*, Phoenix, AZ, April 2008.
- [6] X. Hei, C. Liang, J. Liang, Y. Liu, K.W. Ross, A measurement study of a large-scale P2P IPTV system, *IEEE Transactions on Multimedia* 9 (8) (2007) 1672–1687.
- [7] B. Li, Y. Qu, Y. Keung, S. Xie, C. Lin, J. Liu, X. Zhang, Inside the new coolstreaming: principles, measurements and performance implications, in: *IEEE INFOCOM'08*, Phoenix, AZ, April 2008.
- [8] C. Wu, B. Li, S. Zhao, Multi-channel live P2P streaming: refocusing on servers, in: *IEEE INFOCOM'08*, Phoenix, AZ, April 2008.
- [9] L. Vu, I. Gupta, J. Liang, K. Nahrstedt, Measurement of a large-scale overlay for multimedia streaming, in: *Proceedings of the 16th International Symposium on High Performance Distributed Computing*, Monterey, CA, June 2007.
- [10] F. Wang, J. Liu, Y. Xiong, Stable peers existence importance and application in peer-to-peer live video streaming, in: *IEEE Infocom'08*, Phoenix, AZ, April 2008.
- [11] X. Hei, Y. Liu, K.W. Ross, Inferring network-wide quality in P2P live streaming systems, *IEEE JSAC* 25 (9) (2007) 1640–1654 (special issue on P2P Streaming).
- [12] S. Agarwal, J.P. Singh, A. Mavlanckar, P. Baccichet, B. Girod, Performance and quality-of-service analysis of a live P2P video multicast session on the Internet, in: *IEEE IwQoS Enschede*, NL, June 2008.
- [13] S. Ali, A. Mathur, H. Zhang, Measurements of commercial peer-to-peer live video streaming, in: *Proceedings of Workshop on Recent Advances in Peer-to-Peer Streaming* Waterloo, ON, August 2006.
- [14] T. Silverston, O. Fourmaux, Measuring P2P IPTV systems, in: *ACM NOSSDAV'07*, Urbana-Champaign, IL, June 2007.
- [15] A. Horvath, M. Telek, D. Rossi, P. Veglia, D. Ciullo, M.A. Garcia, E. Leonardi, M. Mellia, Network awareness of P2P live streaming applications, in: *IEEE Hot-P2P 2009*, Rome, May 2009.
- [16] D. Ciullo, M. Mellia, M. Meo, E. Leonardi, Understanding P2P-TV systems through real measurements, in: *IEEE GLOBECOM 2008*, New Orleans, FL, 30 November 2008.
- [17] E. Alessandria, M. Gallo, E. Leonardi, M. Mellia, M. Meo, P2P-TV systems under adverse network conditions: a measurement study, in: *IEEE Infocom 2009*, Rio de Janeiro, Brazil, April 2009.
- [18] E. Alessandria, M. Gallo, P2P-TV systems measurements, Master Thesis, Politecnico di Torino, 2008.



**Eugenio Alessandria** is a network engineer working at France Telecom (Paris) since June 2010. He worked in research domain from begin 2009 until May 2010; his main search interest was Peer to Peer Streaming. He received the First and Second level degree in Master of science in Computer and Communication networks from Politecnico di Torino in 2006 and 2008.



**Massimo Gallo** is a Ph.D. student working at Orange Labs with the partnership of the Telecom ParisTech University since 2009. He received the First and Second level degree in Master of science in Computer and Communication networks from the Politecnico di Torino in 2006 and 2008. His main research interests are Information Centric Networks and Peer to Peer.



**Emilio Leonardi** received a Dr. Ing degree in Electronics Engineering in 1991 and a Ph.D. in Telecommunications Engineering in 1995 both from Politecnico di Torino. He is currently an Associate Professor at the Dipartimento di Elettronica of Politecnico di Torino. In 1995, he visited the Computer Science Department of the University of California, Los Angeles (UCLA), in summer 1999 he joined the High Speed Networks Research Group, at Bell Laboratories/Lucent Technologies, Holmdel (NJ); in summer 2001, the Electrical Engineering Department of the Stanford University and finally in summer 2003, the IP Group at Sprint, Advanced Technologies Laboratories, Burlingame CA. He is the scientific coordinator of the European 7th FP STREP

project “NAPA-WINE” on P2P streaming applications, involving 11 European research institutions, operators and manufacturers. His research interests are in the field of: performance evaluation of wireless networks, P2P systems, packet switching.



**Marco Mellia** received his Ph.D. degree in Telecommunications Engineering in 2001 from Politecnico di Torino. In 1999, he was with the CS Department at Carnegie Mellon University, Pittsburgh (PA) and since April 2001 he is with EE Department of Politecnico di Torino as Assistant Professor. He has co-authored 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.



**Michela Meo** received the Laurea degree in Electronics Engineering in 1993, and the Ph.D. degree in Electronic and Telecommunications Engineering in 1997, both from the Politecnico di Torino, Italy. Since November 1999, she is an Assistant Professor at Politecnico di Torino. She co-authored more than 120 papers, about 40 of which are in international journals. She edited six special issues of international journals, including ACM Monet, Performance Evaluation, Journal and Computer Networks. She was program co-chair of

two editions of ACM MSWiM, general chair of another edition of ACM MSWiM, program co-chair of IEEE QoS-IP, IEEE MoVeNet 2007, IEEE ISCC 2009 and she was in the program committee of about 50 international conferences, including Sigmetrics, Infocom, ICC and Globecom. Her research interests are in the field of performance evaluation and modeling, traffic classification and characterization, P2P, green networking.