

Communication Networks

Measurement of IPTV traffic from an operative network

Marco Mellia* and Michela Meo

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

SUMMARY

In this paper, we present measurement results of IP TeleVision (IPTV) multicast traffic collected from real of the FastWeb network. FastWeb is an ISP provider, which is the main broadband telecommunication company in Italy. The network relies on an IP architecture and, besides television channels, delivers to the user services such as data transfer and VoIP over a single broadband connection. The traffic we focus on consists of 83 high-quality digital TV channels encoded using different MPEG-2 encoders. The results show that, depending on the encoder and based on the bitrate, flows can be classified as being: 1-VBR, i.e. with one typical value of the mean bit rate, 2-VBR, i.e. two typical bitrate values are observed, and VBR. Measurement of the packet loss, jitter and inter-packet gap show that, independently from the class, packet generation process of the flows can have various degrees of burstiness. Despite the packet level burstiness, average jitter is limited to few milliseconds and no packet loss was ever observed in the backbone. Moreover, the evolution of inter-packet gap over time remains basically unchanged when observing a flow in different time periods. Our results prove that the FastWeb network is properly carrying IPTV traffic, and that today technology is mature to make TV over IP feasible and of good quality. Copyright © 2009 John Wiley & Sons, Ltd.

1. INTRODUCTION

The evolution of the Internet towards a universal communication network has been foreseen by both researchers and telecom providers. Voice over IP (VoIP) and IP TeleVision (IPTV) have long been indicated as the technologies that can trigger this revolution, definitively opening the path for convergence. Technology to support such kind of services is available since more than 15 years, and standards are available since the mid of 1990s considering signalling [1, 2] and transport protocols [3], as well as voice and video codecs [4, 5]. By offering both users and telecom providers high speed access and backbone networks, the networking technology is nowadays mature enough to complete the revolution from a data-centric to an universal communication network. New services are indeed being deployed; however, while VoIP is now successfully

adopted by users on a planetary scale, as the case of Skype proves, IPTV is still far from being reality, probably due to some kind of skepticism of content providers and a difficulty in being accepted by users.

In this paper, we show by measurement over real traces collected on the network of FastWeb [6], an ISP provider which is the main broadband telecommunication company in Italy, that IPTV can already be successfully provided and nowadays the technology allows the final convergence step, opening the door for the integration of data and real-time services over the Internet.

To the best of our knowledge, our results are the first extended set of measurement results collected via passive monitoring of high-quality IPTV traffic in the backbone network. FastWeb offers telecommunication services to more than 5 millions of users, with more than 1 million of subscribers (corresponding to the 11% of market share).

* Correspondence to: Marco Mellia, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129, Torino, Italy. E-mail: mellia@polito.it

*Received 2 February 2009
Revised 18 September 2009
Accepted 3 October 2009*

Due to its IP architecture, and the use of either Fiber to the Home (FTTH) or Digital Subscriber Line (xDSL) access, FastWeb has optimised the delivery of converged services, like data, VoIP and IPTV, over a single broadband connection. Our measurements focus on IPTV service and cover several network-centric indexes, such as bitrate, jitter, and loss probability. In addition, stochastic characterisation of both single and aggregate streams is provided; tests to assess whether Markovian assumptions hold and to evaluate Long-Range-Dependency (LRD) are performed.

We anticipate that results show that:

- different classes of streams can be identified based on the bitrate variability;
- no LRD is present in the measured processes, while periodicity is clearly visible;
- no packet loss was ever observed and jitter is almost constant, proving that the simple strict-priority mechanism implemented in the network is effective enough to provide IPTV quality of service guarantees;
- the inter-packet-gap is generally limited, but some burstiness occasionally appears and should be accounted for when dimensioning network element queues, and the receiver play-out buffer.

While this work mainly focuses on IPTV traffic characterisation, the presented results may also give interesting insights onto the source and network behaviour, these may be of interest to researchers that study next generation video encoders and streaming services. Indeed, we believe that the possibility for the research community to access the extended measurement set collected from real installments is a great opportunity to understand the current achievements and the still open challenges that need to be solved. For these motivations, part of the traces studied in this paper are made available and can be freely downloaded at [7].

In our study we could not collect data at the access network, meaning that we could not infer information about the actual quality of experience perceived by the users; the evaluation of indices such as zapping delay, or losses and delays introduced by the play-out buffer, could not be derived. However, our results show that, for what the backbone network is concerned, IPTV traffic can be carried in a proper way so that, with an accurate dimensioning and engineering of the access network, good quality can be provided to users.

The rest of the paper is organised as follows: after discussing related work in Section 2, the FastWeb architecture is described in Section 3; the presentation of measurement methodology follows in Section 4.

Measurement results are reported in Section 5. Finally, Section 7 concludes the paper.

2. RELATED WORK

Traffic monitoring and characterisation has always been seen as a key methodology to understand telecommunication technology and operation, and the complexity of the Internet has attracted many researchers to face traffic measurements since the pioneering times [8]. Data traffic has hogged the majority of this effort, while the attention towards multimedia traffic measurements increased only recently [9–16]. However, most previously mentioned works focus on VoIP traffic, and rely on traffic characterisation and measurement obtained from active probes, in which controlled sources, either PCs or traffic generators, are used to inject packets in a LAN or simple WAN environments. Considering video over IP measurements, the authors of Reference [14] focus on unicast streaming of low-quality video from a high-capacity server to dial-up clients; measurements about the path quality collected from the destination clients are presented. In Reference [15], authors' attention is completely devoted to the characterisation of the behaviour of users accessing a live or video-on-demand streaming. Finally, in Reference [16], authors describe design consideration and implementation of an IPTV monitor system for real time monitoring of the IPTV QoS. As claimed by the authors, the monitor must measure three ultimate metrics: IP cumulative jitter, packet loss and MPEG TS errors. In this paper, we consider a larger set of measurement indexes to better detail both the end-user perceived QoS, and the impact of IPTV traffic on the network.

Considering the characterisation of encoded video traffic, a plethora of papers is available, since from the beginning complex LRD properties were observed in MPEG traffic. Both source traffic characterisation [17, 18], and the its impact on network elements buffers has been deeply studied [19]. However, also in this case, only simple *ad hoc* test-bed experiments have been adopted, and no extensive measurements from real network has ever been presented.

This paper is an extension of our previous work [20], in which preliminary results were presented. In this paper, we investigate the stochastic characteristic of single and aggregate video streams, and we present a more detailed characterisation of the packet arrival process. A deeper

discussion on the implication of the impairment the receiver can suffer due to the observed packet arrival process burstiness is detailed. To the best of our knowledge, our work is the first in the literature to present a large set of experimental measurements collected by purely passive monitoring of high-quality IPTV traffic from a large operative network.

3. THE FASTWEB NETWORK

FastWeb was born in October 1999 with a revolutionary idea of delivering only Internet access to end users (consumers, SOHOs and large business customers) and then providing telecommunication services over IP. In October 2000, the service was opened to consumers and business customers, offering Internet access, VoIP telephony, IPTV and video-on-demand services over the same IP network infrastructure. Since then, FastWeb has become the main broadband telecommunication company in Italy. Its IP architecture and the use of either Fiber-To-The-Home (FTTH) or xDSL access technologies, allow FastWeb to optimise the delivery of converged services, like data, VoIP and IPTV, over a single broadband connection. In this section we briefly introduce the FastWeb architecture, describing the access network, the backbone network and finally the IPTV architecture.

As shown in Figure 1, a Metropolitan Area Network (MAN) Ethernet-based architecture is adopted in the last mile. Residential and small business customers are connected to a Home Access Gateway (HAG), which offers Ethernet ports to connect PCs and the VideoBox, as well as Plain Old Telephone Service (POTS) plugs to connect traditional phones. The HAG is essentially an Ethernet Switch, combined with a H.323 gateway to convert POTS analog input to VoIP transport. In case of FTTH access, a 10Base-F port is used to connect the HAG to a L2 switch installed in the basement; a modem port is used when xDSL access is offered. In the first case, L2 switches are interconnected by 1000Base-SX links forming bidirectional rings. Rings are terminated at the so called MiniPoP by means of two L2 switches, configured as a spanning tree root to manage possible faults. A trunk of several 1000Base-SX links connects each MiniPoP switch to a L2 switch in the PoP, in which two routers are used to reach the backbone by means of Packet-Over-Sonet (POS) STM16 or STM48 links. In case of xDSL access, the HAG is directly connected to a DSLAM by means of a traditional twisted pair phone cable. Then, either a STM4 or STM16 link is used to connect DSLAMs to the PoP by means of an additional router, as shown in the right part of Figure 1; notice that no analog circuit is present even when using xDSL access. When FTTH access is adopted, customers are offered 10 Mbps Half-Duplex Ethernet links, while, in the case of xDSL access, customers are offered 512 or 1024 kbps

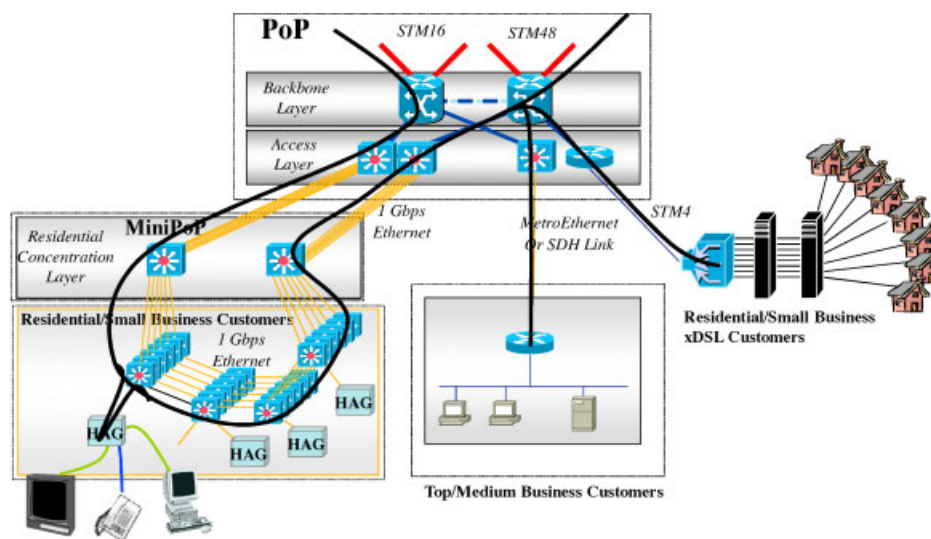


Figure 1. The FastWeb infrastructure: FTTH and xDSL access, MiniPoP, PoP and backbone layers. The black line represents the multicast video stream that spans through all network elements.

upstream links and 6 Mbps or 20 Mbps downstream links. Finally, medium/top business customers are offered either MetroEthernet or SDH access by means of a router directly connected at the PoP level.

Cities covered by the MAN access infrastructure are interconnected by means of a high-speed backbone based on IP-over-DWDM technology. The largest cities in Italy are directly connected by more than 12 400 km of optical fibers. In each city, one or more PoPs are present, while several MiniPoPs are installed so that each one collects traffic from up to 10 000 users.

Considering the services provided to customers, FastWeb offers traditional data access, telephony, video-on-demand and multicast streaming of digital TV channels. At the risk of being tedious, we recall that all services use IP at the network layer. The IPTV architecture, which is the object of the measurements in this paper, is based on both standard and proprietary protocols. In particular, at the time of measurement, 83 digital TV channels were broadcasted. Each TV channel is encoded using high-quality (720 × 576 @ 25 fps) MPEG-2 standard by a 'VideoPump'. VideoPumps are responsible for transcoding the original high-definition digital TV source into broadcasting quality MPEG-2 system stream. Different MPEG-2 encoders are used, resulting in stream bitrate ranging from 2.5 Mbps up to 4 Mbps. The video stream is then encrypted and encapsulated over UDP using a proprietary layer called FastWeb-VideoBox (FWVB) protocol that provides authentication mechanisms so that unauthorised receivers cannot correctly decode the MPEG stream. The FWVB header contains also an explicit sequence number field to allow the receiver to correctly manage missing packets or out of sequence events. 1336B long packets are used by the VideoPump to avoid IP fragmentation problems, so that different bitrates are obtained using variable inter-packet times. At the network layer, standard IP multicast is adopted to transport video streams through the FastWeb network, forming a multicast tree spanning all network routers and switches, as represented by the solid lines spanning all network element in Figure 1. This corresponds to 'broadcast' all TV channels to all network devices, accounting for an aggregate bitrate of about 280 Mbps. A strict priority policy is adopted, so that multicast and Video on Demand traffic is given higher priority than data packets, but lower priority than VoIP streams.

The 'VideoBox' is used to watch TV by home subscribers using traditional TVs connected via a SCART cable. The VideoBox is responsible for decoding the selected TV channel. To 'tune' in a given TV channel, the VideoBox performs a *join* action on the corresponding IP multicast

address. A *leave* action is performed to stop receiving a given channel, so that only one video stream can be selected at a time. Through the FWVB header, the VideoBox can correctly decode the encrypted MPEG stream if subscription is valid.

4. MEASUREMENT METHODOLOGY

In this section, we define the measurement methodology adopted for traffic characterisation, focusing on multimedia streams in particular. A monitoring probe is used to sniff packet headers from traffic flowing on a backbone link in which all multicast streams are present. The first bytes of the packet payload (up to the FWVB headers) are exposed to the analyser.

All the developed algorithms have been implemented in Tstat [21], which is an IP network monitoring and performance analysis tool developed by the Telecommunication Networks Group at Politecnico di Torino; Tstat is a free open-software tool. By passively observing traffic on a network link, Tstat computes a set of performance indexes at both the network (IP) and transport (TCP/UDP) layers. Originally designed for data traffic, Tstat has been enhanced to monitor multimedia streams.

4.1. Performance indexes

Measurements are performed on the stream aggregate, and on a per-flow basis. Some measurements are taken packet by packet, while some others are averaged over time intervals of duration $\Delta T = 1$ s.

Let $X(i)$ be a discrete-time stochastic process, obtained by observing packets from a video stream. To characterise $X(i)$, we focus on:

- The evolution of $X(i)$ versus time i .
- The probability density function (pdf) or Cumulative Distribution Function (CDF) estimated during a stationary period of observation of $X(i)$. Let $\mu = E[X]$ and $\sigma(X) = \sqrt{E[(X - \mu)^2]}$ be mean and standard deviation of $X(i)$, respectively.
- The Autocorrelation Function (ACF)

$$R(k) = \frac{E[(X(i) - \mu)(X(i - k) - \mu)]}{\sigma^2(X)}$$

and the corresponding Power Spectral Density $S(f)$, i.e. the discrete Fourier transform of $R(k)$.

- The estimation of the Hurst parameter of $X(i)$.

Several classes of discrete-time stochastic process can be distinguished, and a huge amount of work is present in the literature to classify $X(i)$. In the following, we will distinguish among: (i) *Markovian*, (ii) *Self-Similar (SS)—Long Range Dependent (LDR)* and (iii) *Periodic (PER)* processes. Markovian processes, of which Poisson processes are a subset, are well-known for their memoryless properties, so that $X(i+1)$ depends only on the current value of $X(i)$. On the contrary, a stationary process $X(i)$ is LRD if its ACF decays to zero so slowly that its integral does not converge, i.e. $\sum_k |R(k)| = \infty$. Intuitively, memory is built-in to the process because the dependence among widely separated values is significant, even across large time shifts. $X(i)$ is SS if $X(at) = a^H X(t)$, $a > 0$, where the equality refers to equality in distributions, a is a scaling factor, and the self-similarity parameter H , $0.5 \leq H < 1$, is called the Hurst exponent. Intuitively, self-similarity describes the phenomenon in which certain process properties are preserved irrespectively of scaling in time. If the process is uncorrelated, then $H = 0.5$. If $H > 0.5$, some correlation exists, being it higher for larger values of H . SS processes are LRD, but LRD processes can be non-SS. Finally, $X(i)$ is periodic if $S(f)$ shows some peaks, i.e. a large amount of the process power is related to a small frequency subset. Equivalently, if $R(k)$ is periodic, then some periodicity exists in $X(i)$.

In order to classify $X(i)$, we start by applying the Lewis-Robinson test [22] to verify if $X(i)$ is Markovian. If the test fails, we look at the Hurst parameter, H , to test if the process is Self-Similar. A large literature is available for the estimate of H . Among all the tools and methodologies, we selected the Abry-Veitch (AV) estimators [23]. Finally, by looking at the $R(k)$ and $S(f)$, we check for the presence of periodicity in $X(i)$.

Given a video flow, the following indexes are monitored:

- *Bitrate, B*: At each time interval i , the following measurement is taken,

$$B(i) = \frac{b_i}{\Delta T}$$

where b_i is the number of observed bits (measured at the IP layer, i.e. including the IP header and payload) during time interval i .

- *Inter-Packet Gap, IPG*: At each time interval i , the following measurement is taken,

$$\text{IPG}_j(i) = t_i(j) - t_i(j-1), \quad j = 2, \dots, N_i$$

where N_i is the number of observed packets during time interval i and $t_i(j)$ is the arrival time of the j -th packet of

interval i . Moreover, we compute,

$$E[\text{IPG}](i) = \frac{1}{N_i - 1} \sum_{j=2}^{N_i} \text{IPG}_j(i) = \frac{t_i(N_i) - t_i(1)}{N_i - 1}$$

$\text{IPG}_j(i)$ is the IPG between consecutive packets, while $E[\text{IPG}](i)$ is the average IPG in time interval i .

- *Average Jitter, J*: At each time interval i , the following measurement is taken,

$$J(i) = \frac{1}{N_i - 1} \sum_{j=2}^{N_i} |t_i(j) - t_i(j-1) - E[\text{IPG}](i)|$$

- *Number of lost, duplicate, late and out-of-sequence packets*:

$$N_{\text{lost}}(i), N_{\text{dup}}(i), N_{\text{late}}(i), N_{\text{out}}(i)$$

To identify lost, duplicate, late and out-of sequence packets, a sliding window mechanism is adopted to record the observed packet sequence, i.e. the *sequence number* in the FWVB header. The sliding window algorithm limits memory usage at the probe node and allows us to identify: (i) numbering gaps, (ii) duplicate sequence numbers and (iii) out-of-sequence delivery. In particular, a *lost packet* is identified if its sequence number has never been observed by the probe node when the sliding window moves on. A *duplicate packet* is identified every time a packet with an already recorded sequence number is observed. An *out-of-sequence packet* is detected if the sequence number of the observed packet is not the expected one. Finally, a *late packet* is identified if the sequence number of the observed packet is outside the sliding window boundaries, e.g. its sequence number is too small to be stored in the sliding window sequence number interval. We set the window size to 32 packets. From a preliminary analysis, we realised that a window size of 32 packets is large enough not to alter our measurement. Indeed, in all tested cases with larger values of the window size, we never had different measurement than with window size 32. Note that the actual VideoBox play-out buffer is larger than 32 packets to absorb jitter too.

5. MEASUREMENT RESULTS

In the following, we present results obtained by monitoring traffic at the MiniPoP level. A probe node based on high-end PCs running Linux has been installed in a PoP located in Turin. The probe is connected to one of the two MiniPoPL2-switches, that is configured to replicate all multicast traffic flowing through the links connecting the PoP backbone

router. The measurement probe is therefore very close to the users' VideoBox, since video packets arrive at the MiniPoP from the VideoPumps that are located in the FastWeb data center in Milan. The measurements presented in this paper were performed during the whole month of December 2006. Reported results refer to the first week of the measurement campaign, and they are representative of the typical system behaviour. Tstat is directly run on the probe machine so that live traffic measurements are taken. An average load of 280 Mbps has been processed.

We first focus on the bitrate per time interval, B . In order to study the distribution of B , we derive the histograms of the values $B(i)$ measured at time interval i by dividing the interval of variability of the bitrate, which is [1.6, 4.8] Mbps, into 41 bins of about 78 kbps extension and we count the number of samples falling in each interval (remind that the time interval is set to $\Delta T = 1$ s). By analysing the behaviour of the 83 flows, we identify three coarse classes of flows:

- **1-VBR** flows, whose value of B is almost constant around one typical value only, namely more than 90% of the measured samples falls in the bin containing the mean value;
- **2-VBR** flows, whose bitrate takes two typical values: the highest peak is smaller than 90% and the two highest peaks, considered together, account for more than 60% of the samples;
- **VBR** flows, for which the bitrate is variable and it is not possible to identify one or two typical values.

As shown in Table 1, the occurrence of 1-VBR flows is remarkable (almost 40%) and only slightly lower of the one of VBR flows. One third of these 1-VBR flows have 100% of the $B(i)$ instances in the peak.

In order to show some examples, the top part of Table 2 reports data about four sample 1-VBR flows and, in particular, it reports: mean, standard deviation, value of the peak and percentage of samples falling in the bin of the histogram containing the peak. The same table also shows data concerning four sample flows from the 2-VBR class: besides the previously mentioned data, details about the second peak are reported. Finally, four sample flows

Table 1. Distribution of the flows over the classes.

Class	No. of flows	% (over 83)
1-VBR	33	39.7
2-VBR	13	15.7
VBR	37	44.6

Table 2. Bitrate per time interval: mean, standard deviation and values of the peaks (when applicable, depending on the flow class).

Class	FID	Mean (kbps)	Std (kbps)	1st peak (kbps)	%	2nd peak (kbps)	%
1-VBR	1	3471	16	3471	100	—	—
	2	3571	16	3571	100	—	—
	3	3486	18	3485	95	—	—
	4	2040	91	2042	97	—	—
2-VBR	1	3587	399	3863	49	2836	15
	2	3162	100	3241	55	3043	44
	3	3701	400	4017	45	2990	15
	4	3626	374	3863	53	2836	11
VBR	1	4064	291	—	—	—	—
	2	3706	149	—	—	—	—
	3	3188	420	—	—	—	—
	4	3365	97	—	—	—	—

extracted from the VBR class are considered, only their mean and standard deviation are given. The data of the remaining flows are not reported in the table for the sake of brevity, but they are equivalent to the considered sample flows. Note that, independent of the class, flow bitrate falls in the [3, 4] Mbps range for all flows except one (the 1-VBR flow 4). The standard deviation is really small with respect to the large bitrate.

We now consider the evolution of the $B(i)$ sequence versus time (or, equivalently, versus i). As a sample case for 1-VBR and 2-VBR flows, we consider the flow number 1 of each of the two classes and plot in Figure 2 the histograms (on the left) and the first part of the sequences versus time (on the right); the 1-VBR case is shown on top of the figure, the 2-VBR case on bottom. While the behaviour of the 1-VBR Flow 1 is straightforward, it is interesting to notice that the two typical bitrate values of the 2-VBR Flow 1 alternate almost uniformly in time.

Since the analysis of the possible behaviours of VBR flows requires more care, Figures 3 and 4 show the histograms and time evolution of the bitrate for all the four sample VBR flows reported in Table 2. The VBR flows are often characterised by the alternation of long periods with very different typical behaviour; this might correspond to the source changing encoder states from time to time possibly due to a change in the programme being transmitted in a given channel or, even with the same encoder, to the characteristics of the TV stream that change within the same programme. In the case of VBR Flow 1, for example, the bit rate is usually distributed around 4.1 Mbps but from about 19:30 to 00:30 it varies much more significantly over a wider range of values, exhibiting two peaks at about 4.2 Mbps and 3.2 Mbps. VBR Flow 2

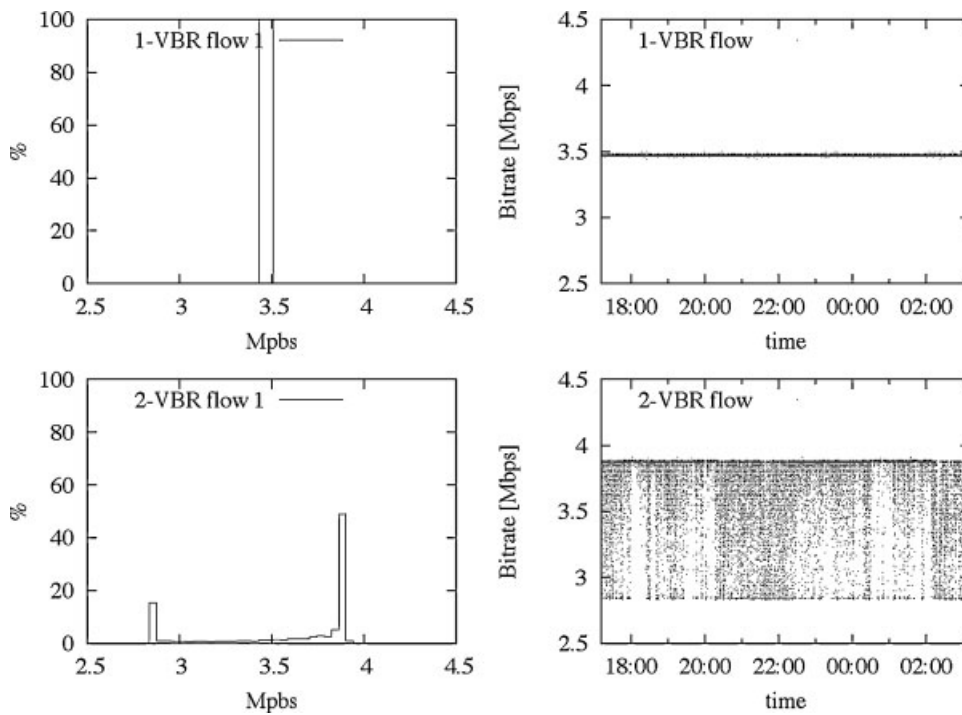


Figure 2. Bitrate histogram and time evolution of 1-VBR Flow 1 (top plots) and 2-VBR Flow 1 (bottom plots).




Figure 3. Bitrate histogram and time evolution of VBR Flow 1 (top plots) and VBR Flow 2 (bottom plots).

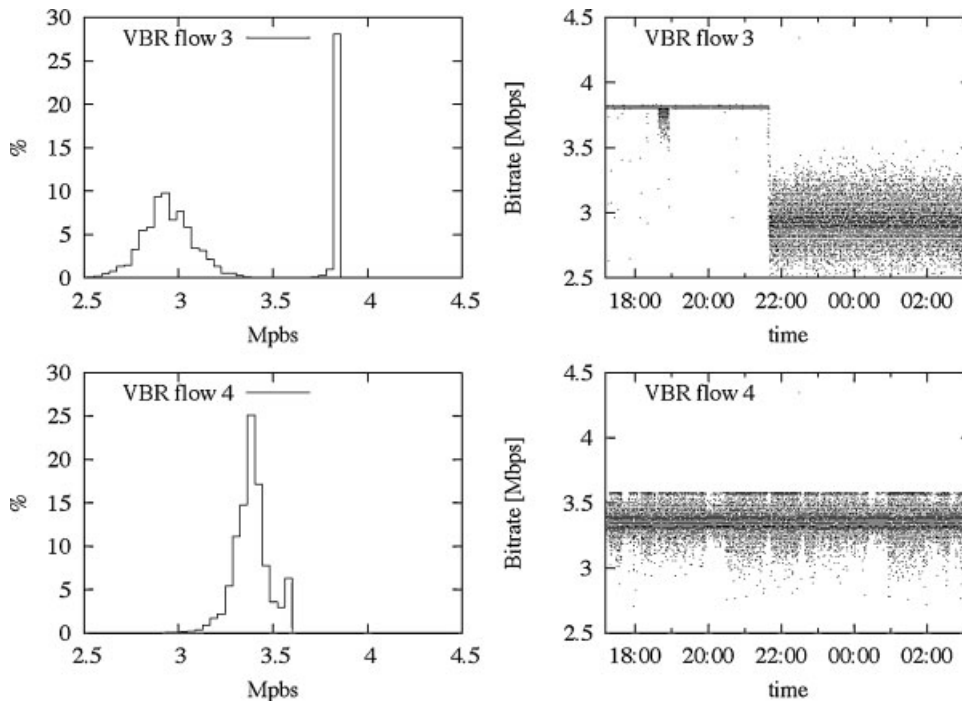


Figure 4. Bitrate histogram and time evolution of VBR Flow 3(top plots) and VBR Flow 4 (bottom plots).

bitrate is usually variable during the considered interval, while from 19:30 to 21:30 it is basically constant. A similar behaviour is exhibited by VBR Flow 3, which moves from an initial period of almost constant bitrate to a second period with a very variable bitrate. Finally, VBR Flow 4 is, on the contrary, quite variable in time for the whole duration of the considered interval, showing a more stationary behaviour.

We now consider the distribution of the average jitter, $J(i)$, and the packet-by-packet IPG. In top left plot of Figure 5 the distribution of $J(i)$ for 1-VBR Flow 1 shows that the average jitter is almost constant (in more than 60% of the time intervals the value of $J(i)$ is at about 1ms). This means that the average jitter in a $\Delta T = 1$ s time interval is stationary and there is no long-term trend in its behaviour. This holds for all flows. In order to have a deep understanding of the IPG between consecutive packets we report the time evolution of inter-packet gap, $IPG_j(i)$ on the top right plot of the same figure. Observe that the time scale on the x -axis refers to 1 minute only, while logarithmic scale is used on the y -axis. The IPG mass is concentrated around two values, respectively at 2 and 4 ms. Therefore, the average jitter, $J(i)$, is almost constant around 1 ms. 1-VBR Flow 1 has the characteristic of a constant average bitrate,

but quite bursty packet arrival process, in which packets are generated every 2 or 4 ms.

As reported on bottom plots of Figure 5, the behaviour of the average jitter and packet-by-packet IPG for 1-VBR Flow 2 is quite different. In this case, $IPG_j(i)$ is almost constant around 3 ms (right plot) and the corresponding jitter $J(i)$ is therefore practically negligible, as shown by its distribution whose mass is around 0.01 ms (left plot). Thus, 1-VBR Flow 2 is a very regular source: constant bit-rate, constant IPG (and therefore low-jitter) flow.

1-VBR Flow 3 (top plots of Figure 6) exhibits constant and large jitter (about 90% of $J(i)$ take the 4.2 ms value). By looking at $IPG_j(i)$ measurements reported on top right plot of the same figure, we observe that it assumes a small set of typical values (a few ‘horizontal lines’ are clearly visible in the IPG time evolution at 0.01, 0.02 and 10.2 ms). This leads to a source with constant bitrate, but high burstiness at the packet level, since IPG is either very small (smaller than 0.02 ms, i.e. back-to-back packets) or very large (larger than 10 ms, i.e. long periods of silence are present). Similar to 1-VBR Flow 3 is the behaviour of 1-VBR Flow 4 (bottom plots of Figure 6): the large average jitter (about 7 ms) is due to IPG taking either very small values (smaller than 0.2 ms) or very large values (about 18 ms).

MEASUREMENT OF IPTV TRAFFIC FROM AN OPERATIVE NETWORK

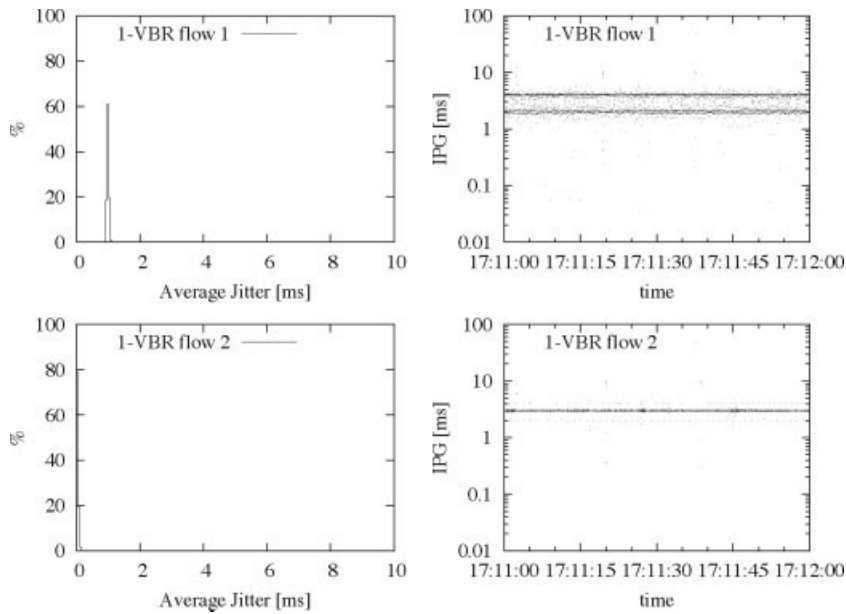


Figure 5. Distribution of the average jitter and IPG time evolution of 1-VBR Flow 1 (top plots) and 1-VBR Flow 2 (bottom plots).




Figure 6. Distribution of the average jitter and IPG time evolution of 1-VBR Flow 3 (top plots) and 1-VBR Flow 4 (bottom plots).




Figure 7. Distribution of the average jitter and IPG time evolution of 2-VBR Flow 1 (top plots) and VBR Flow 1 (bottom plots).

From this analysis, as well as from the study of the other 1-VBR flows that are not reported here for the sake of brevity, we can conclude that within the 1-VBR class of flows, there is no typical behaviour of jitter and IPG: some flows have a significant burstiness, some others have an almost constant IPG. This is possibly due to some encoder producing large blocks and those being segmented into smaller units by the IP packetisation engine.

Figure 7 reports the distribution of the average jitter and the time evolution of the IPG for Flow 1 of the 2-VBR class (top plots) and Flow 1 of the VBR class (bottom plots). Results are similar to the 1-VBR flow cases. Notice that, with respect to jitter and IPG, there is no typical behaviour that can be associated to the class (1-VBR, 2-VBR and VBR), and within each class, different flows can have quite different behaviour of both jitter and IPG.

To give the reader more insights about the packet arrival process, Figure 8 reports the pdf of the IPG measurements considering VBR Flow 2 (on the top plots) and VBR Flow 4 (on the bottom plots). Both linear scales (on the left) and logarithmic scales (on the right) are adopted. Few considerations hold. First, IPG is very regular, and the pdf support is narrow. However, by looking at the plot in logarithmic scale, it can be noted that there is a non-negligible amount of packets that arrive almost back-to-back, i.e. the IPG is in the order of tens of

microseconds. This bursty arrival can cause problems when packets arrive at the users' access switch from the high-speed backbone link. Packets have then to be forwarded to the low-speed user's access link so that the switch queue builds up quickly, and packet dropping may occur.

Similarly, there is a small percentage of packets that arrive after a quite long idle period, i.e. when the IPG is larger than 10 ms. In particular, by looking at the IPG pdf of VBR Flow 4, it can be noted that the IPG takes values larger than 50 ms. This large idle time can cause buffer under-run problems if the playout buffer at the user's VideoBox is not properly dimensioned. Notice that the playout buffer cannot be too large, due to delay constraints arising when watching a live event, or when 'tuning' to a particular channel.

In order to gauge the stationarity of the packet arrival process, Figure 9 shows the IPG distribution during different time periods. In particular, a 12 h long trace has been split into a series of 30 traces of 30 min each. For each trace, the IPG pdf has been evaluated, so as to observe possible changes of the IPG process over time. From bottom plot of Figure 9, it can be observed that the IPG process of VBR Flow 4 is extremely constant in time. Similarly, top plot reports the results for VBR Flow 2; it can be noted that the IPG process is extremely constant as well, but two different pdf classes can be identified that




Figure 8. Distribution of the IPG of VBR Flow 2 (top plots) and VBR Flow 4 (bottom plots). Linear scales on the left plots, logarithmic scales on the right.

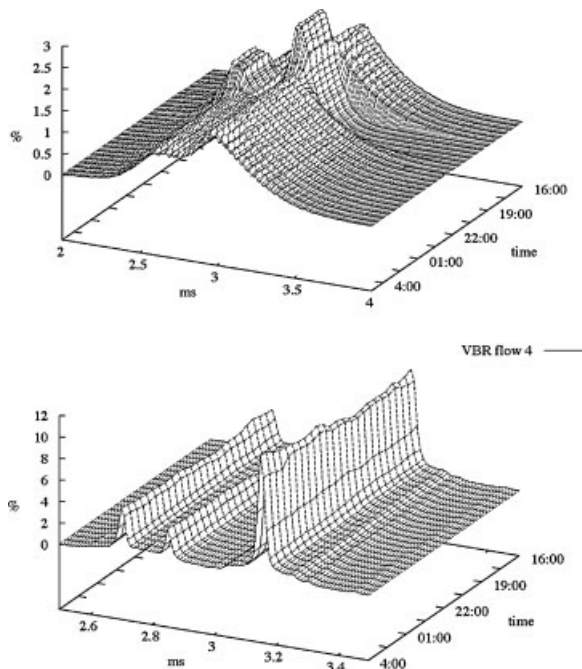


Figure 9. Time evolution of the IPG pdf of VBR Flow 2 (top plot) and VBR Flow 4 (bottom plot).

correspond to the two different source behaviours already observed in Figure 3. Given the regularity of the IPG pdf over time, it can be concluded that the video source emits a very regular stream of packets, and that the network introduces very small variations on packet arrival process. Similar conclusions can be drawn when considering other flows.

Finally, all flows never showed any lost, duplicate, late or out-of-sequence packet during the whole measurement campaign. This shows that, despite the possible large burstiness exhibited at the packet level, the network delivers all packets to the destination, without even introducing out-of-sequence delivery of packets.

Given the extreme regularity of the packet arrival process, we can conclude that the VideoPumps implement a very regular MPEG-2 encoding scheme, and that the network queues add negligible delay to in-transit packets. In particular, the strict priority mechanism adopted at the network layer that favours video packets is very effective in providing excellent quality of service to the video streams. This proves that the quality of service offered to IPTV traffic in the FastWeb network is excellent.

6. STOCHASTIC CHARACTERISATION

In order to characterise the packet arrival process of both the single and aggregate video stream, we performed the Lewis-Robinson test to verify if the Markovian assumption holds. In all cases, the test failed, showing that more complex models must be adopted. We then tested the LRD and periodic properties of different mixes of traffic. We considered several sequences $X(i)$, corresponding to different observation intervals, e.g. during night, day, week or weekend days. We considered different values of the length of $X(i)$ (from 15 min long to 12 h long traces), and estimated the corresponding Hurst exponent.

Results show that observed IPTV traffic cannot be modeled by LRD processes, but rather by periodic processes. Indeed, the estimated Hurst's parameter never exceeded 0.6, being typically in the range [0.5,0.57]

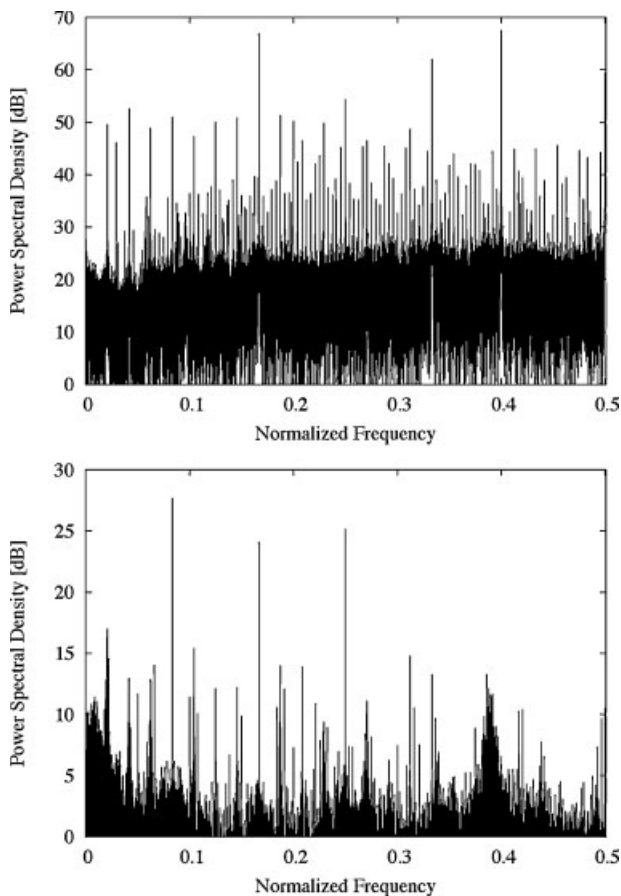


Figure 10. Power spectrum of the IPTV packet counting process. Aggregate stream on the top plot, single VBR stream on the bottom. $\Delta T = 10$ ms.

for all considered cases. This clashes with the general understanding that encoded video streams have LRD characteristics [17–19] and is probably due to the particular MPEG-2 encoders and configuration the VideoPumps adopt, that are almost constant and high bit-rate. The low variability of the packet sending rate and limited correlation over time of the packet arrival process are also due to the fact that the encoded video goes through various buffering stages.

To prove the periodic behaviour of IPTV traffic, Figure 10 shows $S(f)$ of the packet counting process $X(i)$, i.e. the number of IPTV packets arrived in a time window of 10 ms. $S(f)$ of the whole aggregate stream is reported on the top plot, while a single VBR flow is reported on the bottom plot. The aggregate traffic stream exhibits several periodical patterns; the most evident ones are at 20, 25, 30, 40, and 60 ms. They possibly correspond to the periodicity of PAL and NTSC video streams, and to the different Group of Picture (GOP) patterns inside individual streams. Indeed, from the single video stream plot, several periodicities can be noticed, namely at 40, 60, and 120 ms, and they are introduced by both the frame rate and the GOP pattern.

7. CONCLUSIONS

In this paper, we presented an extensive measurement campaign of high-quality IPTV traffic in a real innovative network in operation. We investigated the characteristics of real-time, high-bitrate MPEG video sources transmitted using multicast IP in a commercial network, enlightening similarities and differences among different video channels. Measurement results show that IPTV and network technologies are mature enough to be deployed by large ISPs, opening the path to the convergence towards a single multi-service network.

Measurement results showed that, in the FastWeb network, no IPTV packet loss was ever experienced during the whole measurement campaign, which lasted for more than 3 weeks, and jitter suffered by flows was always very small. This means that, if the access network (which we could not measure) is properly dimensioned and engineered, an excellent service can be provided to users.

No LRD properties have been observed considering both single and aggregate streams. While this clashes with previous studies, the particular setup adopted by FastWeb, i.e. almost constant and high-rate MPEG-2 sources, and constant packet size encapsulation at the network layer, reduces the variability of the MPEG-2 sources, limiting

large scale dependencies. Nonetheless, some small scale variability of the packet arrival process is still present, so that bursts of packets can be generated (causing possible network buffer overflow), and long idle periods can occur (causing possible playout buffer underrun).

Finally, the algorithms, tools and packet level traces used to obtain the results presented in this paper are made available to the research community; we hope they will allow other researchers and network operators to contribute to the understanding of multimedia transmission over the Internet.

ACKNOWLEDGEMENTS

This work was funded by the Italian Ministry of University, Education and Research (MIUR) under the PRIN MIMOSA. Authors thank FastWeb for allowing us to monitor some of their backbone links, and for kindly answering to all our questions. A special thank goes to Kashif Imran and Delia Ciullo for their contribution to this paper.

REFERENCES

- ITU-T Recommendation H.323. Packet-based multimedia communications systems. Available at: <http://www.itu.org>, February 1998.
- Rosenberg J, Schulzrinne G, Camarillo J, et al. RFC 3261, SIP: session initiation protocol. Available at: <http://rfc.net/rfc3261.html>, July 2003.
- Schulzrinne H, Casner S, Frederick R, Jacobson V. RFC 3550, RTP: a transport protocol for real-time applications. Available at: <http://rfc.net/rfc3550.html>, June 2002.
- Cox RV, Kroon P. Low bit-rate speech coders for multimedia communication. *IEEE Communications Magazine* 1996; **34**(12): 34–41.
- Le Gall D. MPEG: a video compression standard for multimedia applications. *Communications of the ACM* 1991; **34**(4): 46–58.
- FastWeb Company Information. Available at: <http://company.fastweb.it>, 2006.
- Mellia M. Tstat home page. Available at: <http://www.tstat.polito.it>
- Caceres R, Danzig P, Jamin S, Mitzel D. Characteristics of wide-area TCP/IP conversations. *ACM SIGCOMM*, Zurich, SW, September 1991.
- Markopoulou AP, Tobagi FA, Karam MJ. Assessment of VoIP quality over Internet backbones. *IEEE Infocom*, New York, June 2002.
- Jiang W, Koguchi K, Schulzrinne H. QoS evaluation of VoIP endpoints. *ACM SIGCOMM*, Zurich, SW, September 1991.
- Calyam P, Sridharan M, Mandrawa W, Schopis P. Performance measurement and analysis of H.323 traffic. *Passive and Active Measurement Workshop (PAM)*, April 2004.
- Conway AE. A passive method for monitoring voice-over-IP call quality with ITU-T objective speech quality measurement methods. *IEEE ICC*, New York, April 2002.
- Cole RG, Rosenbluth J. Voice over IP performance monitoring. *Journal on Computer Communication Review* 2001; **31**(2): 9–24.
- Loguinov D, Radha H. Measurement study of low-bitrate internet video streaming Internet. *1st ACM SIGCOMM Workshop on Internet Measurement*, San Francisco, CA, November 2001.
- Veloso E, Almeida V, Meira W, Bestavros A, Jin S. A hierarchical characterization of a live streaming media workload. *Second ACM SIGCOMM Workshop on Internet Measurement*, Marseille, FR, November 2002.
- Luo C, Sun J, Xiong H. Monitoring and troubleshooting in operational IP-TV system. *IEEE Transactions on Broadcasting* 2007; **53**(3): 711–718.
- Beran, J. Sherman, R. Taquq, M.S. Willinger, W. Long-range dependence in variable-bit-rate video traffic. *IEEE Transactions on Communications* 1995; **V.43**(N.234): 1566–1579.
- Krunz M, Tripathi SK. On the characterization of VBR MPEG streams. *ACM SIGMETRICS Performance Evaluation Review* 1997; **V.25**(N.1): 192–202.
- Rose O. Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems. *IEEE Conference on Local Computer Networks*, 1995; 397–406.
- Imran K, Mellia M, Meo M. Measurements of multicast television over IP. *Fifteenth IEEE LAN/MAN Workshop*, Princeton, NJ, June 2007.
- Mellia M, Lo Cigno R, Neri F. Measuring IP and TCP behavior on edge nodes with Tstat. *Computer Networks* 2005; **47**(1): 1–21.
- Ross SM. *Introduction to Probability Models*. Academic Press, Inc.: New York, 2006.
- Abry P, Veitch D. Wavelet analysis of long-range dependence traffic. *IEEE Transaction on Information Theory* 1998; **V.44**(N.1): 2–15.

AUTHORS' BIOGRAPHIES

Marco Mellia received his Ph.D. degree in telecommunications engineering in 2001 from Politecnico di Torino. From March to October 1999, he was with the Computer Science Department at Carnegie Mellon University, Pittsburgh, PA, as a Visiting Scholar. Since April 2001, he is with Electronics Department of Politecnico di Torino as Assistant Professor. He has co-authored over 140 papers published in international journals and presented in leading international conferences, all of them in the area of telecommunication networks, and he participated in the program committees of several conferences including IEEE Infocom and ACM Sigcomm. His research interests are in the fields of traffic measurement and modeling, 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.