Measurements of Multicast Television over IP

Kashif Imran
Blekinge Institute of Technology
kaim05@student.bth.se
Marco Mellia, Michela Meo
Politecnico di Torino – Dipartimento di Elettronica
lastname@tlc.polito.it

Abstract— **In this paper, we present measurement results collected from real traces on the network FastWeb, an ISP provider that is the main broadband telecommunication company in Italy. The network relies on a fully IP architecture and delivers to the user services such as data, VoIP and IP television over a single broadband connection. Our measurements, that are based on a passive measurement technique, focus on IP TeleVision (IPTV) multicast, that consists of 83 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: CBR, 2-VBR (i.e., two typical bitrate values) 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, showing that the quality of IPTV offered by FastWeb is excellent.**

I. Introduction

The evolution of the Internet toward 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 technology that will trigger this revolution, definitively opening the path for convergence. Technology to support such kind of services is available since more than fifteen years, and standards are available since the mid of ’90s considering both signaling [1], [2] and transport protocols [3], as well as voice and video Codecs [4], [5]. Albeit networking technology has evolved, offering both users and Telecom providers high speed access and backbone networks, still today the revolution is far from being complete. Indeed, while the Internet has definitively been accepted as the only data communication network, the large majority of voice traffic is originated from circuit-oriented networks, and IP television is far from being a reality.

Traffic monitoring and characterization 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 [6]. Data traffic has hogged the majority of this effort, while the attention toward multimedia traffic measurements increased only recently [7], [8], [9], [10], [11], [12], [13]. However, most previously mentioned works focus on VoIP traffic, and rely on traffic characterization 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 environment. Considering video over IP measurements, the authors of [12] 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 [13], authors’ attention is completely devoted to the characterization of the behavior of users accessing a live or video-on-demand streaming. To the best of our knowledge, no measurement study is available in the literature that is based on purely passive monitoring of high-quality IPTV traffic from an operative network.

In this paper, we present the first extended set of measurement results collected via passive monitoring of high-quality IPTV traffic. Real traffic traces are collected from an ISP provider in Italy, called FastWeb [14], which is the main broadband telecommunication company in Italy, offering telecommunication services to more than 5 millions of families, with 1 million of subscribers (11% of market share). Thanks to its fully IP architecture, and the use of either Fiber to the Home (FTTH) or Digital Subscriber Line (xDSL) access, FastWeb has optimized the delivery of converged services, like data, VoIP, IPTV, over a single broadband connection.

Our measurements cover several indexes trying to offer a detailed characterization of network-centric indexes, such as bitrate, jitter, loss probability. Results show that the technology is mature enough to make the final convergence step, allowing for the integration of data and real-time services over the Internet.

The rest of the paper is organized as follows: the FastWeb architecture is first detailed in Sec. II, followed by the presentation of measurement methodology in Sec.III. Measurement results are reported in Sec. IV. Finally, Sec. V concludes the paper.

II. 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. Since then, FastWeb has become the main broadband telecommunication company in Italy. Thanks to

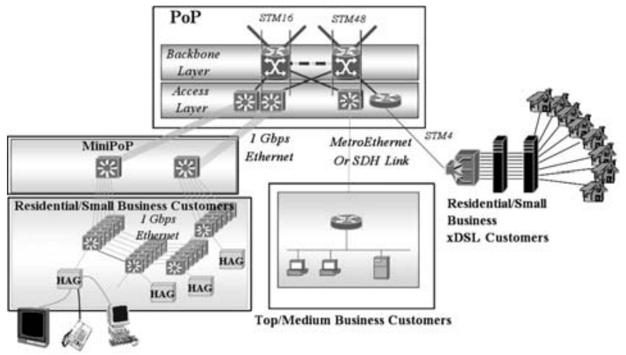

Fig. 1. The FastWeb infrastructure: FTTH and xDSL access, MiniPoP, PoP and backbone layers.

its fully IP architecture, and the use of either Fiber-To-The-Home (FTTH) or xDSL access technologies, FastWeb has optimized the delivery of converged services, like data, VoIP, 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 Fig. 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; while a modem port is used when xDSL access is offered. In the first case, L2 switches are interconnected by 1000Base-SX links forming a bidirectional ring. Rings are terminated at the so called MiniPoP by means of two L2 switches, configured as a spanning tree root to recover from 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 connect the backbone by means of Packet-Over-Sonet (POS) STM16 or STM48 links. In case of xDSL access, the HAG is connected to the traditional twisted pair phone cable terminated directly to a DSLAM. 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 Fig. 1; notice that no analog circuit is present even when using xDSL access. When FTTH access is adopted, customers are offered 10Mbps Half-Duplex Ethernet links, while, in case of xDSL access, customers are offered 512 or 1024kbps upstream and 6Mbps or 20Mbps downstream links. Finally, medium/top business customers are offered both MetroEthernet or SDH access by means of a router connected directly at the PoP layer.

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 12.400km 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 (720x576@25fps) 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 2.5Mbps up to 4Mbps, being either CBR or VBR. The video stream is then encrypted and encapsulated over UDP using a proprietary layer called FastWeb-VideoStation (FWVS) protocol that provides authentication mechanisms so that unauthorized receivers cannot correctly decode the MPEG stream. 1336B long packets are used by the videopump to avoid IP fragmentation problems (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. This corresponds to "broadcast" all TV channels to all network devices, accounting for an aggregate bitrate of about 280Mbps. Multicast packets are marked so that they are forwarded with high priority by network devices.

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. Thanks to the FWVS header, the VideoBox can correctly decode the encrypted MPEG stream if subscription is valid.

III. Measurement Methodology

In this section, we define the measurement methodology adopted to perform the traffic characterization, 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. We assume that the first bytes of the packet payload (up to the FWVS headers) are exposed to the analyzer.

All the developed algorithms have been implemented in Tstat [15]. Tstat is an IP networks 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 focused on

TABLE I

DISTRIBUTION OF THE FLOWS OVER THE CLASSES

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

data traffic, Tstat has been enhanced to monitor multimedia streams.

A. Performance Indexes

In the following, a brief description of the algorithms used to collect packet level measurements is given. Measurements are performed on the stream aggregate, and on a per-flow basis. Moreover, some measurements are taken packet by packet, while some others are averaged over time intervals with extension $\Delta T = 1s$. Both the evolution of a given performance index over time, and its probability density function (pdf) are tracked.

Given a video flow, the following indexes are monitored:

- *Average 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,

$$IPG_j(i) = t_i(j) - t_i(j-1), \qquad j = 2, \dots, N_i$$

  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 the interval i. Moreover, we compute,

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

  $IPG_j(i)$ is the IPG between consecutive packets, while $E[IPG](i)$ is the average IPG in time interval i.
- *Average Jitter, J*[1]: At each time interval i, the following measurement is taken,

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

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

$$Nlost(i),\ Ndup(i),\ Nlate(i),\ Nout(i)$$

To identify lost, duplicate, late and out-of sequence packets, i.e., to evaluate $Nlost(i)$, $Ndup(i)$, $Nlate(i)$, $Nout(i)$, a sliding window mechanism is adopted to record the observed packet sequence, i.e., the *sequence number* in the FWVS 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.

[1]The Jitter measurement is performed at the probe point and not at the receiver node. Being the probe very close to actual destinations, we neglect the missing contribution to the jitter.

TABLE II

AVERAGE 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]	%
CBR	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	-	-	-	-

IV. 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 MiniPoP L2-switches, that is configured to replicate all multicast traffic flowing through the links connecting the PoP backbone router. Tstat is directly run on the probe so that live traffic measurements are taken. An average load of 280Mbps has been processed for more than three weeks.

We first focus on the average 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 78kbps extension and we count the number of samples falling in each interval. By analyzing the behavior of the 83 flows, we identify three coarse classes of flows:

- **CBR** flows, whose value of B is almost constant, meaning that more than 90% of the measured samples fall in the bin containing the mean value

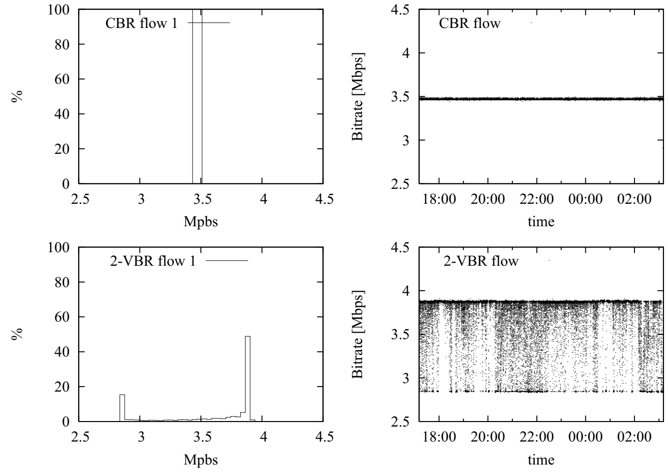

Fig. 2. Bitrate histogram and time evolution of CBR Flow 1 (top plots) and 2-VBR Flow 1 (bottom plots)

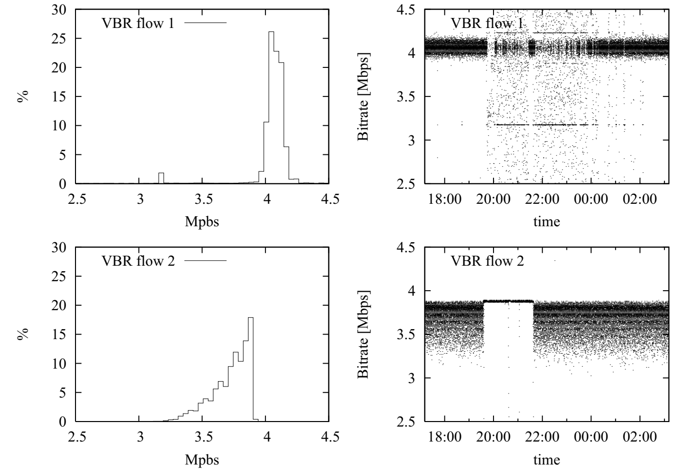

Fig. 3. Bitrate histogram and time evolution of VBR Flow 1 (top plots) and VBR Flow 2 (bottom plots)

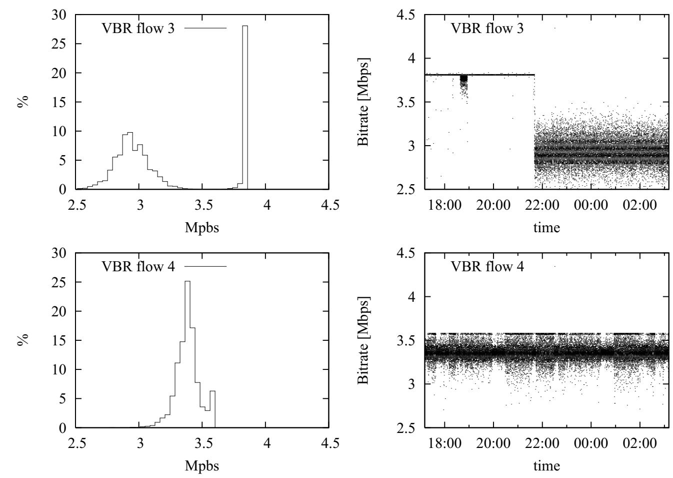

Fig. 4. Bitrate histogram and time evolution of VBR Flow 3(top plots) and VBR Flow 4 (bottom plots)

- **2-VBR** flows, whose bitrate takes two typical values: the highest peak is smaller than 90% and the two highest peaks, considered together, contain more than 60% of the samples
- **VBR** flows, where the bitrate is variable and it is not possible to identify one or two typical values for it.

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

In order to show some examples, the top part of Table II reports data about four sample CBR 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 shows also 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 extracted from the VBR class are considered, only detailing their mean and standard deviation. The data for the remaining flows are not reported in the table for the sake of brevity, but they are equivalent to the considered sample flows. Notice that, independently from the class, flow bitrate falls in the $[3, 4]$Mbps range for all flows except one. The standard deviation is really small with respect to the large bitrate.

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

Since the analysis of the possible behaviors of VBR flows requires more care, Figs. 3 and 4 show the histograms and time evolution of the bitrate for all the four sample VBR flows reported in Table II. The VBR flows are often characterized by the alternation of long periods with very different typical behavior; this possibly corresponds to the source changing encoder from time to time. In the case of Flow 1, for example, the bit rate is usually distributed around 4.1Mbps but from about 19:30 to 00:30 it varies much more significantly over a wider range of values, exhibiting two peaks at about 4.2Mbps and 3.2Mbps. The Flow 2 bitrate is usually quite variable during the considered interval, while from 19:30 to 21:30 it is basically constant. A similar behavior is exhibited by Flow 3, which moves from an initial period of almost constant bitrate to a second period with a very variable bitrate. Finally, Flow 4 is, on the contrary, quite variable in time for the whole duration of the considered interval, showing a more stationary behavior.

We now consider the distribution of the average jitter, $J(i)$ and the packet-by-packet IPG. In top left plot of Fig. 5 the

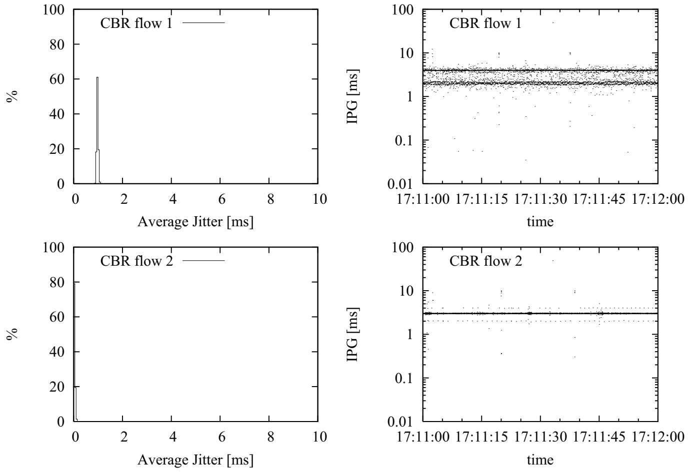

Fig. 5. Distribution of the average jitter and IPG time evolution of CBR Flow 1 (top plots) and CBR Flow 2 (bottom plots)

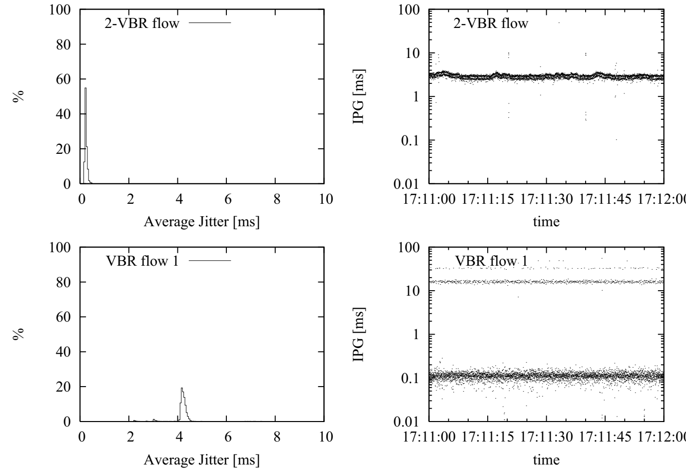

Fig. 7. Distribution of the average jitter and IPG time evolution of 2-VBR Flow 1 (top plots) and VBR Flow 1 (bottom plots)

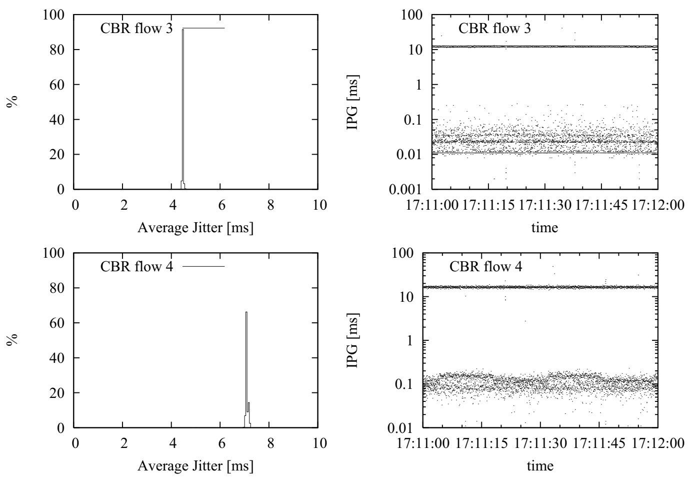

Fig. 6. Distribution of the average jitter and IPG time evolution of CBR Flow 3 (top plots) and CBR Flow 4 (bottom plots)

distribution of $J(i)$ for CBR Flow 1 shows that the average jitter is quite 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 behavior. This observation holds true for all flows. In order to have a deep understanding of the IPG between consecutive packets we report on the top right plot of the same figure the time evolution of the inter-packet gap, $IPG_j(i)$. Observe that the time scale on the x-axis refers to 1 minute only, while logarithmic scale is used on the y-axis. It is clear from the time evolution that the IPG mass is concentrated around two average values, respectively at 2ms and 4ms. Therefore, the average jitter over 1s, $J(i)$, is almost constant around 1ms. Flow 1 has the characteristic of a constant average bitrate, but quite bursty packet arrival process, in which packets are generated every 2ms or 4ms.

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

CBR Flow 3 (top plots of Fig. 6) exhibits constant and large jitter (about 90% of $J(i)$ take the 4.2ms 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 0.01ms, 0.02ms and 10.2ms). This leads to a source with constant average bitrate, but high burstiness is present at the packet level, since IPG is either very small (smaller than 0.02ms, i.e., back-to-back packets) or very large (larger than 10ms, i.e., long periods of silence are present). Similar to CBR Flow 3 is the behavior of CBR Flow 4 (bottom plots of Fig. 6): the large average jitter (about 7ms) is due to IPG taking either very small values (smaller than 0.2ms) or very large values (about 18ms).

From this analysis, as well as from the study of the other CBR flows that are not reported here for the sake of brevity, we can conclude that within the CBR class of flows, there is no typical behavior of jitter and IPG: some flows have a significant burstiness, some others have an almost constant IPG.

Fig. 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 CBR flow cases. Notice that, considering jitter and IPG of 2-VBR and VBR flows, there is no typical behavior that can be associated to the class (CBR, 2-VBR, VBR). Moreover, within each class, different flows can have quite different behavior of both jitter and IPG.

Finally, considering $Nlost(i)$, $Ndup(i)$, $Nlate(i)$, $Nout(i)$, 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 is able to deliver all packet to the destination, without even introducing out-of-sequence delivery of packets. This proves that the quality of service offered to IPTV traffic in the FastWeb network is excellent.

V. Conclusions

To the best of our knowledge, for the first time in this paper we presented an extensive measurement campaign focusing on high-quality IPTV traffic characterization. We investigated the characteristics of real-time, high-bitrate MPEG video sources transmitted using multicast IP in a commercial network, trying to highlight similarities and differences among different video channels. Measurement results on the FastWeb backbone show that IPTV and network technologies are mature to be deployed by large ISPs, opening the path to the convergence toward a single multi-service network.

Moreover, in the FastWeb solution, measurement results highlighted that the quality of service of IPTV is excellent: no packet loss was ever experienced during the whole measurement campaign, which lasted for more than 3 weeks, and jitter suffered by flows has been always very small, introducing no artifacts at the receiver.

Finally, all the algorithms and tools used to obtain the results presented in this paper are made available to the research community via open-source licensing, which we hope will allow other researchers and network operators to contribute to the understanding of multimedia transmission over the Internet.

Acknowledgments

This work was funded by the Italian Ministry of University, Education, and Research (MIUR) under the PRIN MIMOSA. We would like to thank FastWeb for allowing us to monitor some of their backbone links, and for kindly answering to all our questions.

References

[1] ITU-T Recommandation H.323, "Packet-based multimedia communications systems", `http://www.itu.org`, February 1998.

[2] J. Rosenberg and oth., "RFC 3261, SIP: Session Initiation Protocol", `http://rfc.net/rfc3261.html`, July 2003.

[3] H. Schulzrinne and oth., "RFC 3550, RTP: A Transport Protocol for Real-Time Applications", `http://rfc.net/rfc3550.html`, June 2002.

[4] R.V. Cox, P. Kroon, "Low bit-rate speech coders for multimedia communication", *IEEE Communications Magazine*, vol. 34, no. 12, pp. 34-41, December 1996.

[5] D. Le Gall, "MPEG: a video compression standard for multimedia applications" *Communications of the ACM*, vol. 34, no. 4, pp. 46-58, April 1991.

[6] R. Caceres, P. Danzig, S. Jamin, and D. Mitzel, "Characteristics of Wide-Area TCP/IP Conversations", *ACM SIGCOMM*, Zurich, SW, September 1991.

[7] A.P. Markopoulou, F.A. Tobagi, M.J. Karam, "Assessment of VoIP Quality over Internet Backbones", *IEEE Infocom*, New York, NY, June 2002.

[8] W. Jiang, K. Koguchi, H. Schulzrinne, "QoS Evaluation of VoIP End-points", *ACM SIGCOMM*, Zurich, SW, September 1991.

[9] P. Calyam, M. Sridharan, W. Mandrawa, P. Schopis, "Performance Measurement and Analysis of H.323 Traffic", *Passive and Active Measurement Workshop (PAM)*, April 2004.

[10] A.E. Conway, "A passive method for monitoring voice-over-IP call quality with ITU-T objective speech quality measurement methods", *IEEE ICC*, New York, NY, April 2002.

[11] R.G. Cole, J. Rosenbluth, "Voice over IP performance monitoring", *Journal on Computer Communication Review*, vol. 31, no. 2, pp. 9-24, April 2001.

[12] D. Loguinov, H. Radha, "Measurement study of low-bitrate internet video streaming Internet" *1st ACM SIGCOMM Workshop on Internet Measurement*, November 2001, San Francisco, CA.

[13] E. Veloso, V. Almeida, W. Meira, A. Bestavros, S. Jin, "A hierarchical characterization of a live streaming media workload", *2nd ACM SIGCOMM Workshop on Internet Measurement*, November 2002, Marseille, FR.

[14] "FastWeb Company Information", `http://company.fastweb.it`, 2006.

[15] M. Mellia, D. Rossi, "TCP Statistic and Analysis Tool", `http://tstat.tlc.polito.it`, 2006.