

Mobile Data Traffic Analysis: How do you Prefer Watching Videos?

Louis Plissonneau
Orange Labs, France
louis.plissonneau@orange-ftgroup.com

Guillaume Vu-Brugier
Orange Labs, France
guillaume.vubrugier@orange-ftgroup.com

Abstract—In this paper, we analyze an anonymized packet trace of mobile handset data traffic from a major European mobile Internet Service Provider (ISP). Our goal is to understand the effect of the differences between this traffic and the traffic in wired residential networks. This knowledge can be useful for adapting the network to specific mobile usage. We show that, compared to wired users, mobile users exchange data with fewer servers. We then evaluate the performance at flow level, and study the influence of the access type (2G or 3G). In particular, we quantify the improvement provided by a 3G access in terms of delays and throughput. To figure out the user experience, we focus on a popular usage: media streaming. We measure performance metrics for the two main content diffusion techniques: ISP-managed live TV and progressive download (PDL, *e.g.* YouTube or Deezer). We observe that live TV streams only suffer from moderate packet losses thanks to conservative encoding bitrates chosen according to the radio access type. As for progressive download, while audio streams generally cope well with mobile networks, only 3G users experience smooth video playback. Through the analysis of session durations, we determine the effect of the Radio Access Type on the user experience. We notably show that users react to impaired network conditions by aborting their streaming sessions. We also highlight usage differences between live TV and PDL.

I. INTRODUCTION

The increased capacity provided by recent 3G radio access technologies, combined with the rising computing power of the latest smartphones have resulted in a significant growth in mobile traffic. This growth is mainly fuelled by the popularity of large video content on mobiles, which has pushed mobile operators towards providing TV and Video on Demand on their networks. This joint increase in access capacity and user demand affects traffic patterns. As a result, a major concern for operators is to ensure that their mobile network provides a good level of performance when faced with these new traffic characteristics.

In this paper, we use passive measurements from a mobile ISP network to understand to what extent traffic patterns have changed, and how they compare to measurements collected on fixed networks. We first survey previous works in Sec. II before describing in Section III how data were collected and processed. We then present global traffic characteristics at user and flow level in Sec. IV. We show that the Web is still the most popular application, and that because most portals are specifically tailored for smartphones, mobile users often consult a limited number of servers compared to what is observed in fixed networks. We then quantify the performance

gain from 2G to 3G and show that 3G allows users to experience both higher TCP throughputs and shorter delays. In order to further evaluate the users' experience, we focus in Sec. VI on media streaming traffic. We study two types of content diffusion techniques: provider-managed live TV and progressive download (PDL, *e.g.* YouTube or Deezer). As encoding rates are chosen according to the Radio Access Type, we note that the quality of live TV streams (in terms of packet loss) is high, explaining the popularity of this service. Concerning PDL streams, we observe longer sessions as users watch several videos. We infer the number of playback interruptions and show that PDL streaming allows to efficiently deliver both video and audio content over a 3G access: for instance, 93% of audio streams experienced no interruption. On the contrary, 2G PDL video streams are often perturbed by interruptions. A comparison between announced content durations and effective session durations per access type reveals that users abort their streaming session when they encounter network perturbations.

Finally, we discuss the implications of these measurements for ISPs and content providers in the conclusion.

II. RELATED WORK

Previous studies have sought to evaluate the performance of commercial mobile data networks using data from real cellular networks. Their first objective was to measure the effect of a wireless environment on end-to-end TCP performance, starting with GPRS networks [1] and more recently on 3G networks [2]. Passive measurements on Mobilkom Austria's UMTS core network have for instance been extensively used to provide a measurement methodology and information about traffic composition on mobile networks [3], and to analyze capacity bottlenecks [4]. The way people use their mobile for Internet access has been studied by Trestian *et al.* who monitored more than 280,000 users in a 3G mobile network in order to assess the correlation between user's application interests and their location or movement patterns [5]. In [6], the authors try to sketch the most significant changes in wireless data traffic (including Wifi and cellular access) over ten years.

Our work differentiates from these in that we give a new insight into wireless traffic. Indeed, we extend our measurements of network performance to user-perceptible effects by evaluating how content streams behave in terms of playback

TABLE I: Data Overview

Date	November 22 nd , 2008	
Duration	1 H 39 min	
# of local clients	95,011	
2G users	50,062	
3G users	36,533	
TCP volume (GB)	up	5.6
	down	52.8
UDP volume (GB)	up	0.6
	down	16.4
2G volume (GB)	up	2.2
	down	18.1
3G volume (GB)	up	2.4
	down	38.7

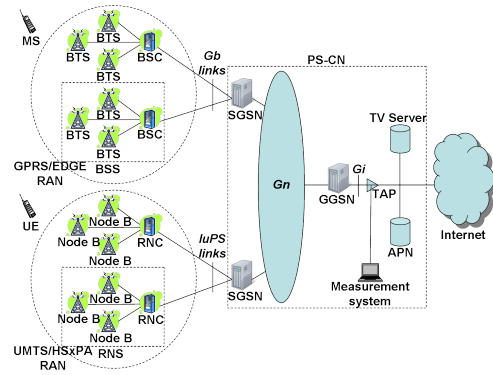


Fig. 1: Capture Schema

interruptions. This method can be used by ISPs and content providers to determine how contents should be adapted for delivery on mobile networks.

III. DATA COLLECTION AND ANONYMIZATION

We collected mobile data traffic from the network of a mobile ISP using a passive probe connected to an optical coupler. The monitored link was located between a GGSN¹ and the backbone network (see Fig. 1). We thus collected the data traffic generated by users located across the whole country. This traffic originated from around 39,000 cells and was exclusively generated by mobile handheld devices² (*i.e.* no modem or 3G USB stick).

In order to cope with the disk performance of the probe, we decided to store the traffic of a subset of the subscribers' IP addresses (one /17 range). As a result, the measurement device lost no packet during the traffic capture. However, due to the dynamic address assignment policy, we cannot guarantee that a client will stay in this address pool after a disconnection.

We parsed the RADIUS (Remote Authentication Dial In User Service) signaling messages in order to retrieve the associations between the clients and their successive connections. We can thus determine which IP address was assigned to the client at a specific time, and extract the following information:

- IMSI the International Mobile Subscriber Identity, which uniquely identifies a customer;
- RAT the Radio Access Type, which indicates if the device uses a 2G or 3G access at that moment.

As a first step, we **anonymized** this information by *hashing the IMSI* and by cutting packets just after their TCP or UDP header to ensure privacy. Some payload information was kept for audio/video traffic only in order to carry out more precise performance measurements (see Section VI).

Table I regroups relevant information about the trace. One can notice that:

- 3G access facilitates data usage: although there are 40 % more 2G clients, twice the traffic is on 3G access. Note

¹Gateway GPRS Support Node: interfaces the mobile network and the backbone.

²this traffic is using a different VLAN and can thus be easily identified

that some 3G/2G handovers were observed during our measurements. As a result, the sum of the 2G and 3G users counts is slightly greater than the number of local users.

- most of the traffic is downstream (received by the client): this reveals that no server is hosted in the mobile network, and peer-to-peer (P2P) softwares are not used;
- TCP traffic carries three times more bytes than UDP: this ratio is lower than in fixed networks because the ISP-managed live TV service uses UDP for transport³.

As we do not have access to the Radius tickets transmitted before and after the trace, we cannot determine the RAT and IMSI of some users (less than 10%).

The distribution of packet size further reflects the strong imbalance between upstream and downstream traffic:

- Upstream: most packets are TCP acknowledgments, or belong to small requests sent by mobile terminals to the servers. Their size is around 40 bytes.
- Downstream: content mainly flows in this direction. As a result, 40 % of the packets received by the mobile devices are full-sized (bigger than 1400 bytes).

IV. MOBILE VS. FIXED TRAFFIC CHARACTERISTICS

In this section, we look at global traffic characteristics to see how they compare with previous measurements carried out on wired networks. We first determine the dominant applications before describing how traffic is distributed among the hosts.

A. Application Breakdown

Application recognition is carried out by a custom deep packet inspection (DPI) software that has been calibrated with commercial DPI engines and proved to perform better than signature-based tools [7].

We decided to focus on the following application types as they account for most of the downstream volume:

- Web: including HTTP and HTTPS⁴;

³Another argument is that wired networks are dominated by P2P traffic, which makes heavy use of TCP.

⁴HTTPS can hide non-web traffic

TABLE II: Volume per application as a function of the RAT

Application volume in GBytes	Down				Up			
	2G		3G		2G		3G	
Web	12.2	67 %	17.8	46 %	1.8	83 %	1.7	71 %
Live TV streaming	1.1	6 %	10.2	26 %	6.5^{-5}	$\ll 1\%$	18.10^{-5}	$\ll 1\%$
PDL streaming	2.6	14 %	8.2	21 %	0.08	3 %	0.2	9 %
Other	2.2	12 %	2.5	6 %	0.3	12 %	0.5	18 %

- Live streaming: ISP-managed TV delivered using unicast UDP flows;
- Progressive download (PDL) streaming: transfer of multimedia content over HTTP, *e.g.* YouTube, iTunes or Deezer.

The “other” category aggregates all remaining application types (including mails and downloads) and unidentified traffic.

Table II shows that Web traffic is predominant both on 2G or 3G access.

Streaming is the second dominant application in terms of downstream volume, and is much more popular on 3G access. Indeed a minimum downstream throughput is required to watch TV in decent conditions (see Section VI-A). Note that live streaming generates almost no upstream traffic because it is carried through downstream UDP flows. Progressive download streaming generates less asymmetric traffic because of the upstream acknowledgments.

B. Traffic Distribution: Clients and Servers

In this part, we compare the concentration of traffic among clients and servers between our mobile trace and measurements from ADSL and FTTH (Fiber To The Home) residential access networks. Figures for fixed networks come from data collected in July 2008 [8]. The capture duration is similar for these three traces.

As revealed in Tab.III, the most active users carry almost the same fraction of downstream bytes in mobile and fixed networks. The absence of p2p in our trace only has a limited effect on this distribution. Indeed, p2p does not dominate the downstream traffic in residential wired networks; it only dominates upstream [8].

Another key observation is the way mobile data users concentrate on servers compared to the situation in fixed Internet traffic.

In order to study this, we computed server counts per client. We define this server count as the number of unique IP addresses accessed by a client over all of its sessions. Only flows bigger than 1 kB are taken into account in order to exclude invalid connection attempts which are common in

TABLE III: Contribution of downstream volume for clients and servers.

	Mobile	ADSL	FTTH
Top 1 % users	45 %	41 %	43 %
Top 20 servers	45 %	28 %	27 %

wired residential networks. Those connections typically occur in p2p traffic when peers attempt to contact sources that have left the p2p overlay, resulting in aborted 3-way handshakes with no payload transmitted. Still, the chosen threshold allows us to keep small web requests.

On the mobile trace, some older terminals access the Internet through a proxy. This device acts as a gateway to access other servers and has no equivalent in our fixed traces. For fair comparison with fixed networks, we thus have to consider it separately. We report access shares to this gateway for 2G and 3G in table IV(b).

Table IV(a) shows that the average number of servers accessed per subscriber is one order of magnitude lower on the mobile trace. This implies a greater concentration of mobile traffic on a few servers and is confirmed by Tab.III: the fraction of downstream volume created by the top 20 servers is higher on the mobile trace than on wireline traces. There are several reason that concur to this.

First, our mobile trace lacks p2p traffic. Such traffic involves lots of peers, especially for DHT-based softwares. Its absence thus tends to lower the figures reported for the mobile trace. To check if it is the only explanation, we applied the same process to Web traffic only (ports 80 and 443).

The “Web traffic” column shows almost the same significant difference in the concentration on servers. As a result, another immediate explanation for this gap lies in an intrinsically different user behavior when surfing the web: users would mainly connect to a limited number of services like social networks, webmails or news websites on their mobile, and browse lots of different websites on their computer.

A closer look at the dominant servers on the mobile trace

TABLE IV: Mobile clients access a lower number of servers (statistics on flows bigger than 1 kB)

(a) Average number of servers accessed per client, excluding the proxy.

	All traffic	Web traffic
ADSL	54.1	38.9
FTTH	71.3	43.6
Mobile (2G)	4.20	5.36
Mobile (3G)	5.35	7.16

(b) Clients accessing the proxy.

	Unique clients
2G access	17655 (35%)
3G access	9314 (26%)

V. GENERAL PERFORMANCE ANALYSIS

The objective of this section is to determine the performance of mobile flows. A first remark is that mobile users are very patient with the data they want to retrieve. Indeed we found only weak correlation between the throughput obtained and the volume transferred (compared to fixed networks). This observation holds for both 2G and 3G accesses: the correlation coefficient is 0.24 and 0.19 respectively. This shows that the user behavior concerning the load they put on their connection is almost the same regardless of the throughput. In order to better understand the reasons behind these throughputs, we refine the study with flow-level performance metrics: we focus on the round-trip time (RTT) to get information on delays and on the retransmission ratio as indicator of packet loss.

We measure the RTT of each TCP connection during the initial three-way handshake, and we only draw statistics on connections established by mobile devices. The TCP 3-way handshake can be described as follows: the mobile device first sends a SYN packet to a server which replies with a SYN-ACK packet, upon reception, the mobile device answers with an ACK packet. At the end of the process, the connection is fully established and ready for data transfers. In the rest of this section, we call:

- Wireless-side RTT: the time of the mobile loop, *i.e.* the delay between the SYN-ACK packet and the corresponding acknowledgment at the probe;
- Wired-side RTT: the time of the Internet loop, *i.e.* the delay between the initial SYN segment and the corresponding SYN-ACK response at the probe.

Figure 3(a) shows high wireless-side RTTs which are mainly due to retransmissions on the *air* interface because of interferences between frequencies. We measure much shorter RTTs for 3G than 2G, confirming previous studies [2]. As expected, no significant difference is observable on the Internet loop (not shown here); wired-side RTTs are around 50 ms, but some connections encounter backbone-side RTTs up to 200 ms. Flows exchanged with the proxy have very short wired-side RTTs (1 ms) as they stay in the ISP network.

For each TCP flow, we compute the retransmission ratio by dividing the number of packets with the same sequence number seen twice by the probe over the total number of packets for this flow. We count retransmissions in each direction, so that downstream retransmissions correspond to the wireless part and upstream ones correspond to the wired part. In Fig. 3(b), we observe clearly lower retransmission rates for 3G than 2G: only 8% of 3G flows have non-zero retransmissions whereas 16% of 2G flows have some retransmissions. Moreover, the overall retransmission ratio is higher for 2G than for 3G. The ratios on the wired part are identical for 2G and 3G, and thus omitted here.

The consequence of these shorter RTTs and lower loss rates is that TCP can really take advantage of the higher bandwidth provided to 3G clients, resulting in higher downstream peak rates.

Figure 4 shows the distribution of per client peak rates,





Fig. 2: Distribution of the number of servers accessed.


shows how specific they are to mobile networks: among the top 20 servers, we find (apart of the proxy) the Live TV and mail servers of the ISP, and portals operated by Apple and Google which are specifically tailored for smartphones. This contrasts with the situation in fixed residential networks where the top 20 servers are — apart of a few p2p sources — either content hosting servers or part of Content Distribution Networks and are not operated by the ISP.

Another possible cause behind the lower server counts in mobile web traffic is that we identify servers through their IP address. Content providers often do load balancing between different IPs to provide access to their contents. As mobile traffic represents a lower part of the Internet traffic for now, it may be that content providers provision less servers (hence IPs) for it, resulting in lower server counts on our mobile trace.

Figure 2 shows the distributions of the number of servers contacted per client. The gap between fixed and mobile traffic is clearly visible. Moreover, we observe that the higher the access capacity (*e.g.* 3G vs. 2G, and FTTH vs. ADSL), the higher the server count. In addition, isolating Web traffic tends to decrease the number of servers accessed in fixed trace while it increases this number for mobile users. This is typically due to p2p traffic (or its absence).



(a) Wireless-side RTTs



(b) Wireless Retransmission Ratio per Flow

Fig. 3: Performance metrics on the wireless part




Fig. 4: Peak-rate over 500 ms intervals.

evaluated over 500 ms intervals. We only consider for this figure the subscribers who transferred more than 10 kB of data in the trace.

We first observe that the use of 3G results in a marginal speedup in upstream peak rates. The main reason is that almost only TCP acknowledgments transit in this direction and that acknowledgment flows created by Web traffic rarely saturate the uplink.

The improvement is much more visible downstream as the peak rate experienced by clients is on average 2.7 times higher

TABLE V: Theoretical link capacity and expected throughput per RAT.

	Peak network capacity	Average user throughput
GPRS (2G)	115 kbps	30-40 kbps
EDGE (2G+)	473 kbps	100-130 kbps
UMTS (3G)	2 Mbps	220-320 kbps
HSDPA (3G+)	14 Mbps	550-1100 kbps

for those connected on 3G.

For comparison purposes, Tab. V summarizes the theoretical link capacities and usual throughputs for 2G and 3G systems, taken from the literature [9].

The measured improvement may seem limited when compared to the reported peak network capacities. However, it should be noted that the actual throughput can be far from the theoretical one for various reasons: imperfect transmissions conditions, distance between the mobile and the base station, overhead of the protocol stack, etc. In addition, we have to mention that most 2G access points were EDGE-enabled, whereas many 3G access points did not yet support High Speed Downlink Packet Access (HSDPA). Per client peak rates should increase as HSDPA is rolled-out.

VI. PERFORMANCE ANALYSIS OF MULTIMEDIA STREAMING

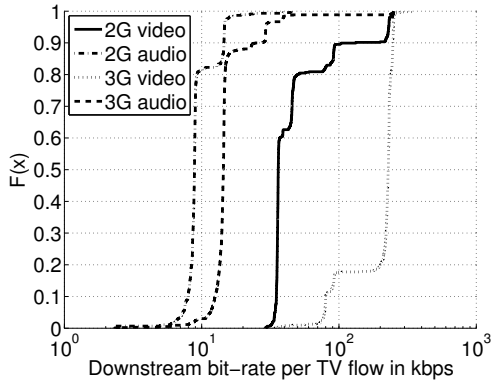
As previously shown in the paper, video streaming is among the most popular mobile applications. Analyzing its performance is thus of double importance as the required engineering and costs can vary widely depending on the chosen diffusion architecture. In this section, we distinguish carrier-based live TV and “over-the-top” media delivery through progressive download streams. Although in both cases a buffer is used to absorb jitter or losses, the induced usage is different: with live TV, users cannot pause and resume videos; with PDL like *e.g.* YouTube or Deezer, when the media stream is paused, the download continues and the user can resume playback later⁵.

A. UDP Streaming: Live TV

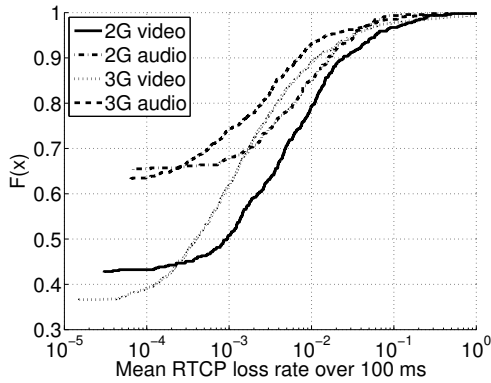
Mobile data users have access to TV programs delivered by dedicated servers. These servers stream contents using unicast UDP flows. TV streams are composed of one audio and one video stream. Each stream is composed of a RTP (Real-time Transport Protocol) data flow associated to a RTCP (RTP Control Protocol) control flow. The encoding rate of the TV flow is determined by the ISP according to the user’s radio access type. We have audio bit-rates at 8 or 12 kbps, and video bit-rates from 40 kbps up to 250 kbps for High Definition flows. These steps are clearly identifiable in Fig. 5(a) because there is no throughput adaptation in UDP, and because contents are streamed in real time.

We evaluate the user’s perceived quality of these TV flows using the loss rate extracted from the RTCP Receiver Reports. In our case, the mobile devices compute the loss ratio every

⁵In the extreme case, the media can be fully downloaded in advance, thus ensuring interruption-free playback



(a) Throughput of downstream live-TV flows per Radio Access Type and content type



(b) Loss rates of TV flows computed out of RTCP receiver reports

Fig. 5: UDP streaming characteristics

100 ms by relating the number of received packets to their RTP sequence numbers. In Fig. 5(b), most of the flows experience a low mean loss rate ($<1\%$). One can notice a difference between audio and video losses: video packets (that are bigger) are more likely to be lost. The loss rates for live TV are quite good, but losses are very bursty: some periods of 100 ms may be completely lost. Obviously, 3G networks are more capable of coping with video traffic, whereas audio streams always have a good quality whatever the radio access type. This indicates that bitrates are appropriately chosen.

B. TCP Streaming: Progressive Download

In this section, we focus on media streaming over TCP. In this so called “progressive download” (PDL) mechanism, a buffer allows the application to download a part of the media in advance before presenting it to the user, and also to absorb network perturbations. This traffic thus clearly differs from what is presented in the previous section.

To estimate the user-perceived performance of PDL multimedia streams, we parse HTTP headers and mp4/3gp/flv metadata to retrieve the content bitrate and detect the presence of sound and video tracks. In addition, we determine whether the user was connected on the 2G or 3G access network (only a handful of the observed transfers experienced handover

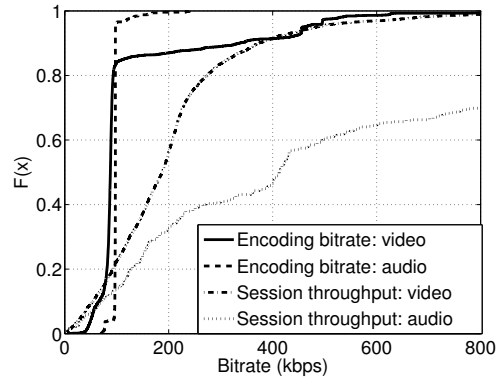


Fig. 6: Bitrates of progressive download media

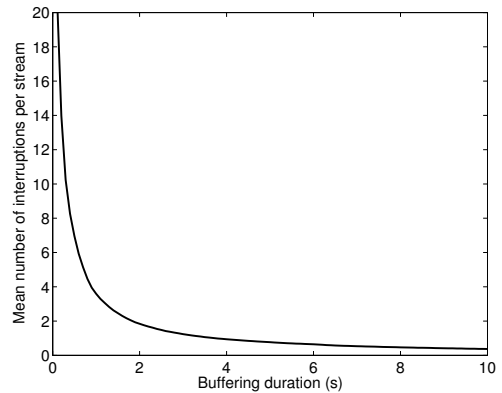


Fig. 7: Effect of the buffering duration on the number of playback interruptions

between the two). Table VI(a) shows that most of the identified streams carry video content (*i.e.* both the sound and video tracks were present), but that a significant number consists in purely audio content, here: music. We also note that 3G seems to encourage video content viewing.

The distribution of encoding bitrates depicted in Fig. 6 shows a strong mode around 96 kbps. The reason for audio streams is that we observed only one provider (`deezer.com`) which streams music tracks encoded at 96 kbps. Similarly, the low encoding bitrates for video content result from mobile-specific guidelines. However, the figure also reveals a huge difference between these encoding bitrates and the actual throughput on the network: for example, around 10% of all audio streams have a throughput greater than 2 Mbps. Actually, PDL streams are elastic traffic, and not true streaming. As a result, throughputs are mainly influenced by the capacity of the radio link and by the content provider’s traffic management policies. For instance, YouTube is known to enforce rate limitations [10] whereas Deezer clearly does not.

As we have access to the actual rate of the data stream, we can reconstruct the state of the playback buffer for each session. This allows us to detect the user-visible interruptions that occur when the network cannot keep up with the media bitrate. Before starting playback, we let the device buffer a volume (in




Fig. 8: Construction of streaming sessions

TABLE VI: Classification and performances of PDL streams

(a) Number of streams per RAT			
	2G	3G	
Audio streams	206	291	
Video streams	1112	2082	

(b) Fraction of streams with playback interruptions (%)			
	2G	3G	Overall
Audio streams	15.0	6.5	10.4
Video streams	32.0	16.4	22.3

bytes) that is equivalent to a fixed duration of the media⁶. This buffer size is calculated from the content encoding rate. Fig. 7 shows that the number of playback interruptions does not vary much for buffer sizes corresponding to more than 2-3 s. We thus summarize in Tab. VI(b) the results of this process for an initial buffering duration set to 3 s of content.

One can notice that the encoding bitrate allows these contents to be transmitted effectively on current mobile networks. Consequently, measurements show that most network perturbations are absorbed by the buffering: the network caused no playback interruption for more than 80% of the streams. Table VI(b) also illustrates how the increase in capacity provided by 3G access networks improves the user experience in this case: the fraction of interrupted transfers is almost halved.

C. Influence of the Radio Access Type on Session Durations

We first investigate how clients watch PDL programs. To that end, we reconstruct our *streaming sessions* as illustrated in Schema 8. Note that a user can play several media contents during one streaming session, and that interrupted sessions are also taken into account.

As the actual throughput can widely differ from the content encoding rate, the user is often still watching the (buffered)

video when the TCP flow ends. The fact that we observe no data transmission thus does not mean that the user is inactive. To overcome this problem, we extrapolate the instant the user stops watching based on the amount of data received and on the encoding rate. We then determine our streaming sessions by aggregating these viewing intervals and closing the session when the client was inactive for more than 30 s. We tried to vary this think time from 30 to 180 s and observed a smooth increase in the average session durations, *i.e.* no step that would represent a characteristic value.

These streaming sessions allow us to differentiate between streaming usages: a user interested in a single video may abort watching it if the quality is bad, whereas a user interested in watching any video may frequently “switch-over” to find more interesting or better quality content.


In Fig. 9(a) we represent the duration of the PDL contents as indicated in their content header. Note that the media requested by 2G and 3G users have identically distributed durations. Video streams have shorter duration overall. However, we notice that the distribution has a heavier tail than the one computed for audio streams (because the length of most songs is around 3 minutes).

If we now relate this to the actual session durations depicted in figure 9(b), we clearly observe smaller durations for 2G sessions. As the content themselves have similar durations, we explain this difference by the fact that 2G users experience more playback interruptions. They thus tend to react to the degraded quality by shortening their sessions.

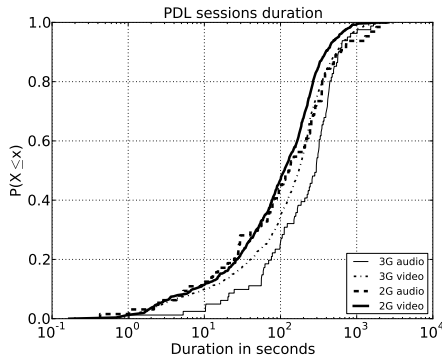
Finally, Fig. 10 confirms that the use of 2G also affects the duration of live TV sessions.

We also notice that most live TV sessions are shorter than the PDL ones shown on figure 9(b). As the loss rates are moderate (see Fig. 5(b)), we believe that the reason lies in a different user behavior. On the one hand, PDL streams come from sites that provide a huge choice of small user-generated content. Users are thus active and tend to fetch many related videos, leading to longer sessions. Live TV, on the other hand, is a more passive usage: fewer contents are provided, and switching over is not particularly encouraged.

⁶Note that this in no case restricts the maximum buffering duration



(a) Distribution of the duration of PDL streams (as specified in the content header)



(b) Distribution of the duration of PDL streaming sessions (as measured on the network)

Fig. 9: Comparison between announced and measured durations of PDL streams

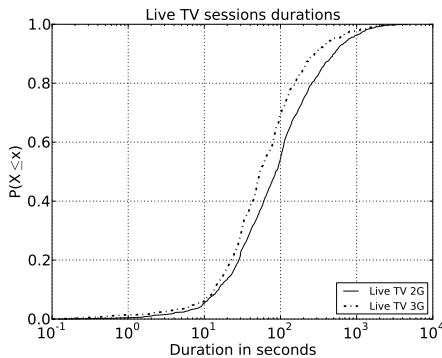


Fig. 10: Measured Live TV sessions duration

VII. CONCLUSION

In this article, we have analyzed a packet trace of mobile data traffic. We first show that, compared to wired access networks, mobile users usually retrieve data from a limited number of servers. Our results reveal that this phenomenon is due to a different user behavior when browsing the Web, combined with the absence of P2P.

We have also quantified a **2.7 throughput improve-**

ment factor with UMTS Radio Access Type compared to GPRS/EDGE. This gain is mainly visible downlink and is facilitated by the wireless RTT reduction at TCP level.

Focusing on the user experience when viewing multimedia content, we evaluate two content delivery techniques: live TV and progressive download. We show how their behavior differs and how the Radio Access Type influences their performances:

- **Live TV** is a popular service offered by the ISP whose bitrates are standardized and well adapted to the radio access. As the loss rate is low, the playback buffer should allow a good video quality apart of some loss bursts.
- **Progressive download** allows users to begin watching the content after a short buffering duration. In order to assess the user experience, we propose a method for estimating the number of playback interruptions from passive network measurements. We observe that 3G users generally enjoy an interruption-free playback as the actual network throughput is usually higher than the content encoding rate. However, most 2G users can conveniently retrieve audio streams only. We also show that PDL users react to bad playback quality by prematurely aborting their sessions.

The lessons of these measurements are that:

- HSDPA users could be provided higher encoding rates for videos;
- 2G users should mainly be proposed audio contents and shorter streams, if possible.

ACKNOWLEDGMENTS

We would like to thank Thierry Houdoin for his help in providing the data and the necessary tools and information for analysis, as well as for his feedback on drafts of this paper.

REFERENCES

- [1] P. Benko, G. Malicsko, and A. Veres, "A large-scale, passive analysis of end-to-end TCP performance over GPRS," in *INFOCOM 2004*, 2004.
- [2] P. Romirer-Maierhofer, F. Ricciato, A. D'Alconzo, R. Franzan, and W. Karner, "Network-Wide Measurements of TCP RTT in 3G," *Traffic Monitoring and Analysis*, pp. 17–25, 2009. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-01645-5_3
- [3] P. Svoboda, F. Ricciato, E. Hasenleithner, and R. Pilz, "Composition of GPRS/UMTS traffic: snapshots from a live network," in *IPS MoMe 2006, Salzburg*, 2006.
- [4] F. Ricciato, F. Vacirca, and P. Svoboda, "Diagnosis of capacity bottlenecks via passive monitoring in 3G networks: an empirical analysis," *Computer Networks*, vol. 51, no. 4, 2007.
- [5] S. Trestian, I. Ranjan, A. Kuzmanovic, and A. Nucci, "Measuring Serendipity: connecting people, locations and interest in a mobile 3G network," in *IMC'09*, 2009.
- [6] E. Halepovic, C. Williamson, and M. Ghaderi, "Wireless data traffic: a decade of change," *Network, IEEE*, vol. 23, no. 2, pp. 20–26, march 2009.
- [7] M. Pietrzyk, J.-L. Costeux, G. Urvoy-Keller, and T. En-Najjary, "Challenging Statistical Classification for Operational Usage: the ADSL Case," in *IMC'09*, 2009.
- [8] G. Vu-Brugier, "Analysis of the Impact of early Fiber Access Deployment on Residential Internet Traffic," in *ITC21*, 2009.
- [9] H.-H. Chen, *The Next Generation CDMA Technologies*, 1st ed. John Wiley & Sons, 2007.
- [10] L. Plissonneau, T. En-Najjary, and G. Urvoy-Keller, "Revisiting Web Traffic from a DSL Provider Perspective: the Case of YouTube," in *ITC SS 19*, 2008.