Accepted from Open Call

On the Performance of New Generation Satellite Broadband Internet Services

Alessio Botta and Antonio Pescapé

Abstract

In the context of Internet access technologies, satellite networks have traditionally been considered for specific purposes or as a backup technology for users not reached by traditional access networks, such as 3G, cable or ADSL. In recent years, however, new satellite technologies have been pushed to the market, reopening the debate on the possibilities of having high-performance satellite access networks. Therefore, the performance monitoring and analysis of such networks is of great interest and importance for both industry and academia.

In this article we present experiments we conducted for and results we obtained from the performance monitoring and measurement of two satellite broadband Internet access services. Moreover, we describe the simulator we built for the black-box analysis of the traffic shaping mechanism employed by the satellite operator.

We quantify the real performance achievable with these technologies. Our unique results show when and how new-generation Internet satellite services can be a promising way to provide broadband Internet connections to users, and shed light on the traffic shaping mechanism employed by the operator and the possibilities left for the users.

Introduction

Originally launched for long-distance telephony and for television broadcasting, communication satellites are more and more used for Internet access [2–6]. In its infancy, due to its poor performance, Internet access via satellite has traditionally been chosen by users — often in rural areas — not served by other access networks (e.g. 3G, cable, or ADSL), having special needs (e.g. bank communications) or as a backup and secure link. The first commercial services for residential satellite Internet access were monodirectional, requiring another technology (e.g. the telephone) for the uplink direction. The poor performance (few hundred kb/s) and the necessity of another Internet access strongly limited their spread. Later on, bidirectional commercial services were launched, but still their performance was poor and the costs high. In the following years, however, a great effort has been put on this technology and several improvements have been achieved. Among the most relevant, we cite the new TCP versions and improved TCP acceleration mechanisms, which highly increased the performance of TCP (and then of applications relying on it) over the satellite link [7–10] and the launch of satellites with a set of features specifically designed for Internet access.

For example, new satellites launched recently (e.g. in the 2011) have Internet connectivity as a primary target and promise high individual and aggregated performance, thanks to several improvements such as multi-spot illumination/frequency reuse, TCP accelerators, robust terrestrial network (based on MPLS), and so on. As a consequence, recent commercial services for Internet access via satellite promise tens of Mb/s user data rates and stable performance, and therefore reopen the debate on the possibility to use satellite networks for Internet access on a larger scale. The several advantages of Internet access via satellite (quick and easy installation and deployment of terminals, low environmental impact, etc.) are counterbalanced by a number of disadvantages, including high latency, necessity to employ middleboxes and accelerators, impact of the weather conditions, and so on. These aspects and the need for real and updated performance figures when referring to operational satellite networks were the main motivations of the work described in this article.

To this end we have set up a testbed and we have experimentally measured and analyzed the performance of two satellite broadband Internet access services: a first- and a second-generation bidirectional satellite connection. Also, we have built a simulator for the black-box analysis of the behavior of the Fair Access Policy, a traffic shaping mechanism used by the operator. Our results provide a fresh sketch of actual performance (in terms of throughput, delay, jitter, and loss) of both the first and latest generation Internet satellite services, and they show when and how they are a promising way to provide Internet connection to users. Also, the proposed simulator sheds light on the mechanisms employed by the operator for shaping user traffic.

This article is organized as follows. In the next section we introduce the technologies subject of this study. Following that we present the testbed we set up for it. We describe the results of the performance measurements in the next section and the results of the FAP analysis in the section following that. Finally we draw conclusions.

The authors are with University of Napoli Federico II.

This research has been partially funded by the MIUR projects PLATINO (PON01_01007), SMART HEALT (PON04a2_C), and SIRIO (PON01_02425).

Preliminary results within this framework have been presented in [1].

 0163-6804/14/$25.00 © 2014 IEEE

Broadband Satellite IP Networks

In this section we briefly review the two satellite broadband Internet access services that are the subject of our study: the first-generation satellite network, or FGS for short, and the second-generation satellite network, or SGS for short. Then we describe the mechanism to control the volume of traffic allowed to the users employed in both FGS and SGS.

First and Second Generation Satellite Networks

FGS — This service was deployed in 2007 and has a number of users of the order of ten/hundred thousand. The connection is bidirectional, and the satellite has a single spot covering Europe and some countries in the Mediterranean area. The network operates through wireless links working in the Ku and Ka frequency bands, using adaptive coding and modulation on the forward link and rate adaptable, multi-carrier TDMA on the return link. The operator advertises about 3.6 Mb/s of maximum downlink throughput and 384 → 512 kb/s of maximum uplink throughput per user. Three different profiles are offered to the users, mainly differing in terms of maximum volume of traffic allowed (see subsection below).

SGS — This service is based on a new geostationary satellite deployed recently (mid 2011) by the same operator of FGS. The satellite operates in the Ka band and it is equipped with highly directional antennas, with spot size of the order of a few hundred kilometers. Thanks to this feature the satellite is able to cover a large part of Europe using about a hundred spots. The concept of spatial reuse of the frequencies (as done in the cellular network) has been used, reaching an aggregated satellite throughput of the order of a hundred Gb/s. The satellite link uses adaptive coding and modulation on the forward link and automatic power control and rate adaptation on the return link. The commercial service for bidirectional Internet access, launched in late 2011, comprises different profiles, with advertised maximum throughput in the range of 6 → 10 Mb/s in the downlink and in the range of 1 → 4 Mb/s in the uplink. The different profiles are also characterized by a different volume of traffic allowed to the users. As the satellites are geostationary, the minimum one-way delay (i.e. time to go from the earth to the satellite, and back) of these two access technologies is on the order of 250 ms.

Both FGS and SGS employ Performance Enhancement Proxies (PEPs) to speed-up TCP traffic [10]. PEPs can employ different techniques to enhance performance. Very commonly, PEPs split the TCP connection established by and toward the host connected via satellite into two or more parts. This allows the client to use a standard TCP and the PEPs to use a special transport protocol (e.g. a custom TCP version) specifically designed for the satellite link. In this way it is possible to strongly increase the performance of the end-to-end connection. For both FGS and SGS the PEP is located inside the satellite modem, that is, in the LAN of the user, so that the TCP connection is closed very near to the user and a new TCP connection (with a special TCP version) is opened just before the satellite link. In the following sections we analyze the impact of the PEP on performance measurements, using both FGS and SGS. Further studies on the PEP are also reported in [11].

For both FGS and SGS the PEP is located inside the satellite modem, that is, in the LAN of the user, so that the TCP connection is closed very near to the user and a new TCP connection (with a special TCP version) is opened just before the satellite link.

The Fair Access Policy (FAP)

To carefully share the aggregated satellite bandwidth among the users, both FGS and SGS limit the maximum traffic volume allowed to the users. The technique used by the operator is called Fair Access Policy (FAP). It operates as follows. It periodically checks the volume of traffic produced by the user in different sliding time windows (one hour, four hours, one day, etc.). If one of these volumes exceeds a threshold (that is larger for larger time windows), the maximum allowed throughput is limited (the limitation is more sever for larger time windows. For example, the maximum allowed throughput is about 250 kb/s if the user exceeds the 1–h threshold; it is about 150 kb/s if the user exceeds the 4–h threshold, and so on). The limitation is removed when the volume in the sliding time window becomes smaller than the threshold. Below we present a methodology and a simulator we devised for the black-box analysis of the FAP behavior, useful to study other similar approaches.

The Testbed

As shown in Fig. 1, the testbed is composed of Linux servers connected to the Internet through our University network as well as through FGS or SGS. FGS and SGS are based on different technologies. For the former the testbed is equipped with offset Gregorian parabolic antennas with diameter of the main dish of 65 cm, connected to the satellite modems located in our server room, connected to the FGS measurement clients. For SGS there are front-feed parabolic antennas with diameter of 77 cm, connected to the satellite modems located in our server room, connected to the SGS measurement clients. The testbed comprises clients that are both subjected to the FAP and not (i.e. with no limits on the volume of traffic). The non-limited clients have been used to measure the performance of the satellite connections, while the FAP-limited clients have been used to test the shaping mechanism. We believe this testbed has unique characteristics that makes it possible both to deeply test these connections, and experiment the same conditions of the standard users.

As illustrated in Fig. 1, we generated traffic through the connections under test, from a host using this connection to another host in our University. We used D-ITG [12] for traffic generation and performance measurement, with different traffic profiles (VBR and CBR), rates (from about 5 kb/s to 12 Mb/s), and protocols (TCP and UDP), measuring the throughput, delay (both one way and round trip), jitter, and loss. The experimental campaigns were performed between February 2010 and July 2012, involving a few thousands of minutes-long measurement experiments and collecting about 1 TByte of measure-

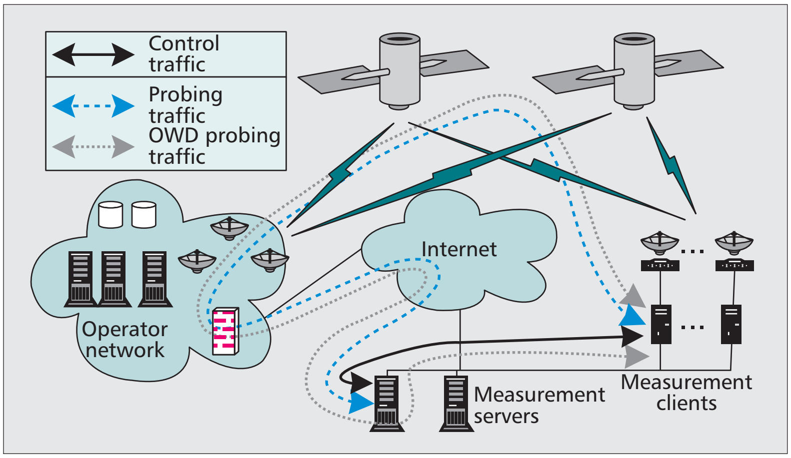

Figure 1. Testbed used for the experiments.

ment data. Traces collected during the measurement campaigns are available on demand.[1]

On the Performance of IP Satellite Services

Here we present the most interesting results obtained:

- The throughput of FGS to draw a baseline.
- The impact of the PEP (see above) on measurement and monitoring procedures and results.
- The throughput and one way delay of SGS, underlying also the PEP impact here.
- The impact of the TCP version on the throughput of SGS, and performance results obtained with video, game, and voice applications.
- A comparison between SGS and FGS using real applications and users.

Performance of FGS

Throughput — Figure 2a reports the throughput obtained in the uplink and the downlink on FGS with both UDP and TCP as a function of the packet size (each point represents the average of 100 measurements conducted in different days and daytimes, and the standard deviation is reported with the vertical segment). For the downlink we report the results obtained with a packet rate of 1000 pps, while for the uplink we report those with a packet rate of 100 pps. With these packet rates and starting from some packet sizes, we saturated the downlink and the uplink of FGS respectively, showing the maximum throughput. Figure 2a shows that FGS is able to achieve about 300 kb/s of throughput in the uplink with both TCP and UDP. The saturation is obtained with packet sizes larger than 512 Bytes. Before that, the throughput measured equals the expected one. In the downlink, **FGS is able to achieve about 3.5 Mb/s with UDP and about 2.5 Mb/s with TCP, with the latter also having larger variance**. This is due to the sensitivity of TCP to competing traffic: the high number of users of FGS, and consequently the high volume of traffic competing for resources cause the TCP throughput to be lower than that of UDP on average and also more variable. This important result will be considered in the analysis of the FAP below.

Figure 2b shows the throughput behavior during the day. The plot is obtained measuring the maximum throughput in both the uplink and the downlink with TCP every half hour for two entire weeks, and then grouping and averaging the results for each hour of the day (i.e. the result related to each of the hours reported on the x axis is obtained averaging the results for that particular hour in all 14 days). The right plot of Fig. 2b shows that **the throughput in the downlink is characterized by a very strong diurnal pattern**, with average values that range from 1.6 Mb/s during the day to 3.2 Mb/s during the night. The left plot of Fig. 2b shows that this behavior is less evident in the uplink. The main reason is that FGS uses multi-carrier TDMA in the uplink, making it possible to isolate the uplink channels of the users. Moreover, the aggregate volume of all downlink traffic is much larger than the uplink one, and this can also cause congestion in the terrestrial network of the operator.

PEP Impact — As already said, an important aspect to consider is the presence of the PEP [10], which can severely impact the performance measurements and results. Figure 2c shows the RTT estimated with three different approaches: ICMP Ping, TCP on port 80 (*syn-ack time*, TCP80 in the following), and TCP on port 81 (*syn-reset time*, TCP81 in the following). The effect of the PEP is notable: **the RTT with TCP80 is on the order of 10^0 ms, while with the other two approaches it is on the order of $10^2/10^3$ ms**. The reason is that the PEP operates only on TCP traffic on port 80, and the RTT estimated with TCP80 only regards the network path between the user and the PEP (i.e. the LAN connecting the measurement client and the satellite modem). This result can have important consequences on delay measurements: if the tools and methodologies are not chosen and operated carefully, the measures can be severely impacted and strongly misleading. This aspect has been investigated in the following (see One Way Delay) and in our previous paper [11].

Performance of SGS

Throughput — Figure 3a shows the results in the uplink and downlink of SGS, using the same format as in Fig. 2a. In this case, we report the results obtained with a packet rate of 1000 pps also in the uplink because at 100 pps the SGS uplink was not saturated. Figure 3a shows that **SGS is able to achieve much higher and more stable performance than FGS**. The uplink throughput is about 3.9 Mb/s and the downlink throughput is about 9.5 Mb/s with both TCP and UDP. The better performance is mainly due to two causes:

- SGS employs state-of-the-art TCP performance accelerators that help TCP traffic cope with the satellite link.
- Being very new, SGS has many fewer users, and the network, also thanks to the high aggregated throughput available, is far from being congested.

We measured the throughput behavior during the day also for SGS, and we verified that the throughput is stable during all the hours, as also witnessed by the small variance in Fig. 3.

[1] *http://traffic.comics.unina.it/Traces*

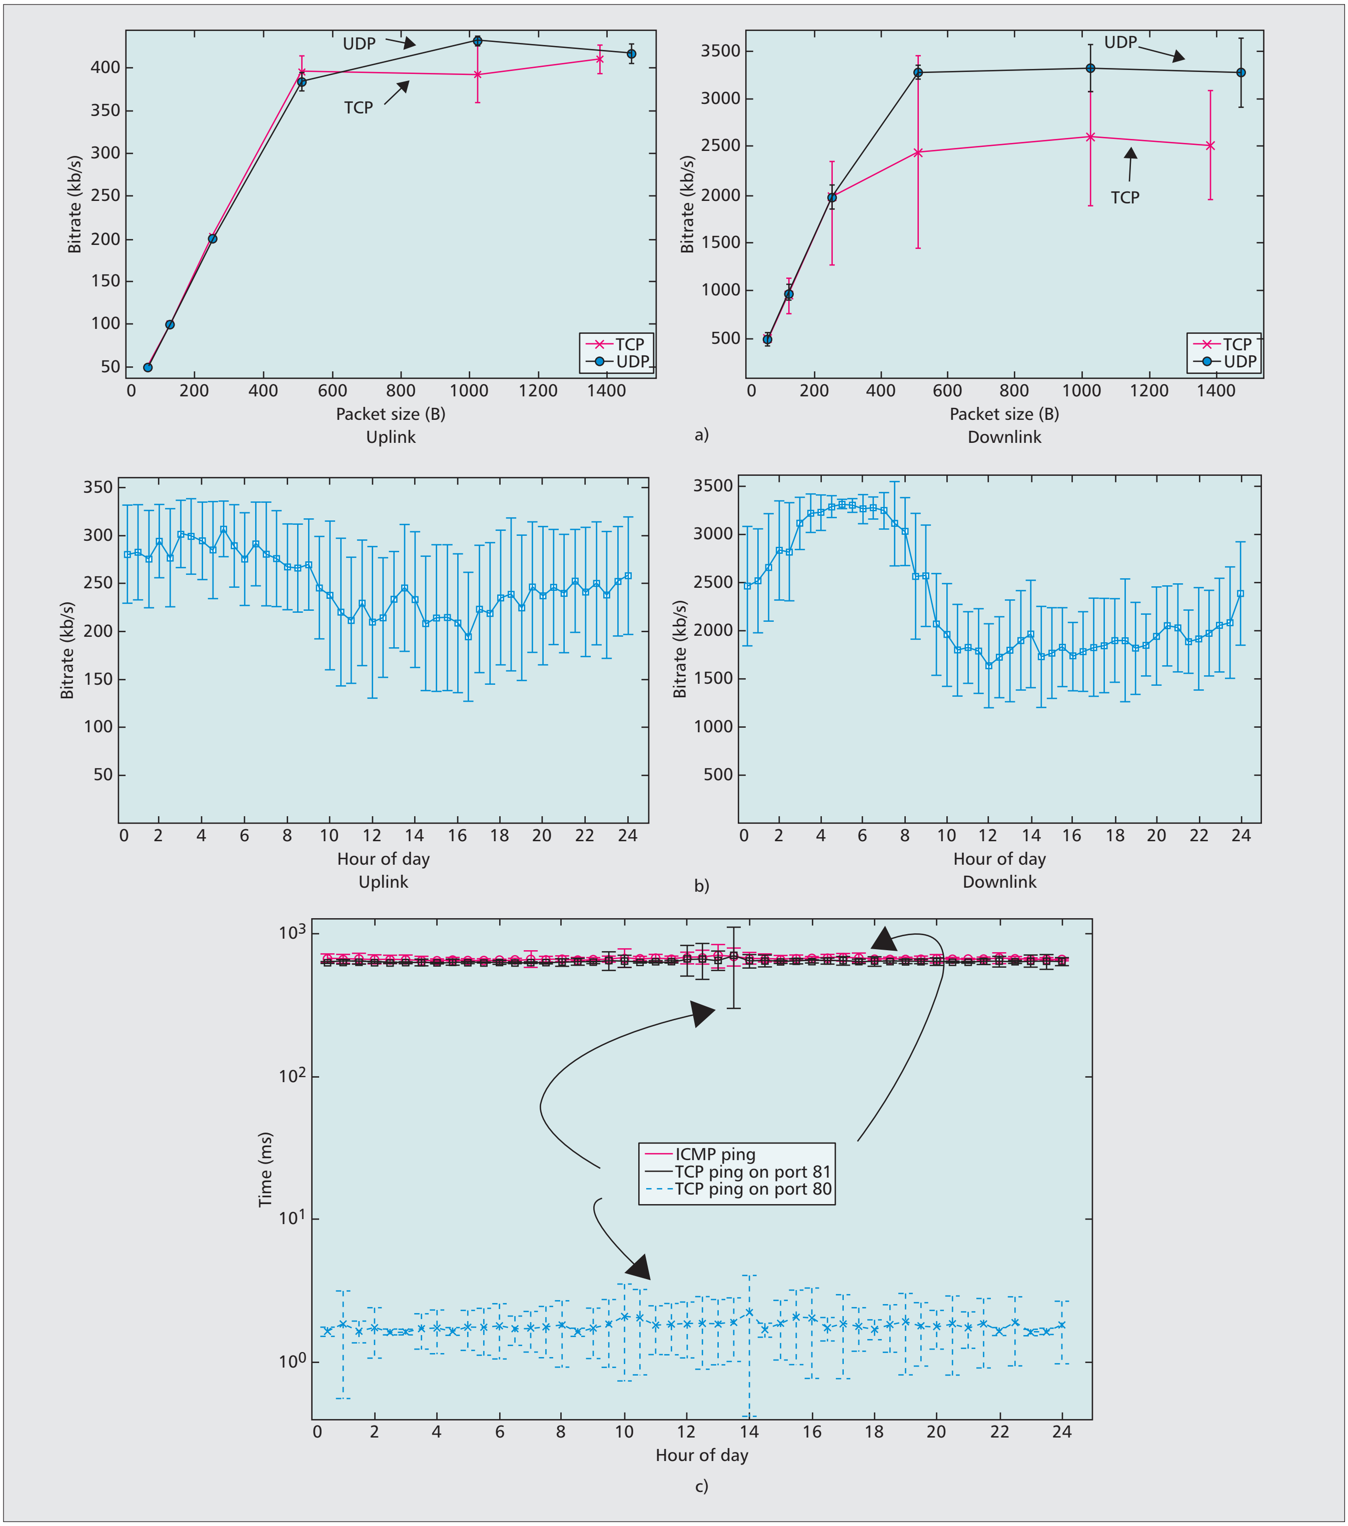

Figure 2. Performance of FGS.

One Way Delay — An important performance parameter for long-distance wireless networks is the latency, often evaluated through the round trip time (RTT) to overcome the clock synchronization issues. The RTT, however, is influenced by both uplink and downlink directions, which can be an issue on asymmetrical network connections such as satellite connections. To measure the One Way Delay (OWD) overcoming the clock synchronization issues, we set up the testbed to receive the packets on the same hosts that generated them, as reported in Fig. 1 (see OWD probing traffic). The measurement clients are both senders and receivers of measurement traffic, which is routed through the satellite network thanks to a host acting as a NAT (i.e. one of the measurement servers in Fig. 1). We experimentally verified that the impact of such operation is negligible with respect to the OWD of the link under test.

The left plot of Fig. 3b shows the cumulative distribution function of the OWD samples collected with TCP and UDP, using a packet rate of 100 pps and three different packet sizes. Counter intuitively, **with TCP the OWD decreases with the packet size: the larger the packet size, the smaller the OWD**. This behavior is due to a buffer in the satellite network that operates on TCP traffic (very likely being the PEP), working on bytes rather than on packets, which justifies the larger OWD values with smaller packet sizes. This is also confirmed by the fact that with UDP we did not observe this behavior, and the OWD is always very concentrated around 0.3 s. Such buffer causes the OWD with TCP and small packet sizes to increase up to 0.6 s. This high RTT does not cause throughput degradation in Fig. 3a because the link is still not saturated with packet sizes smaller than 512 Bytes and TCP is able to achieve the expected throughput even with this RTT. It is worth recalling that we already saw some possibly negative consequences of the PEP on delay measurements, RTT in that case. In conclusion, the PEP can be a serious issue when performing network measurements, and these results (see also [11]) can

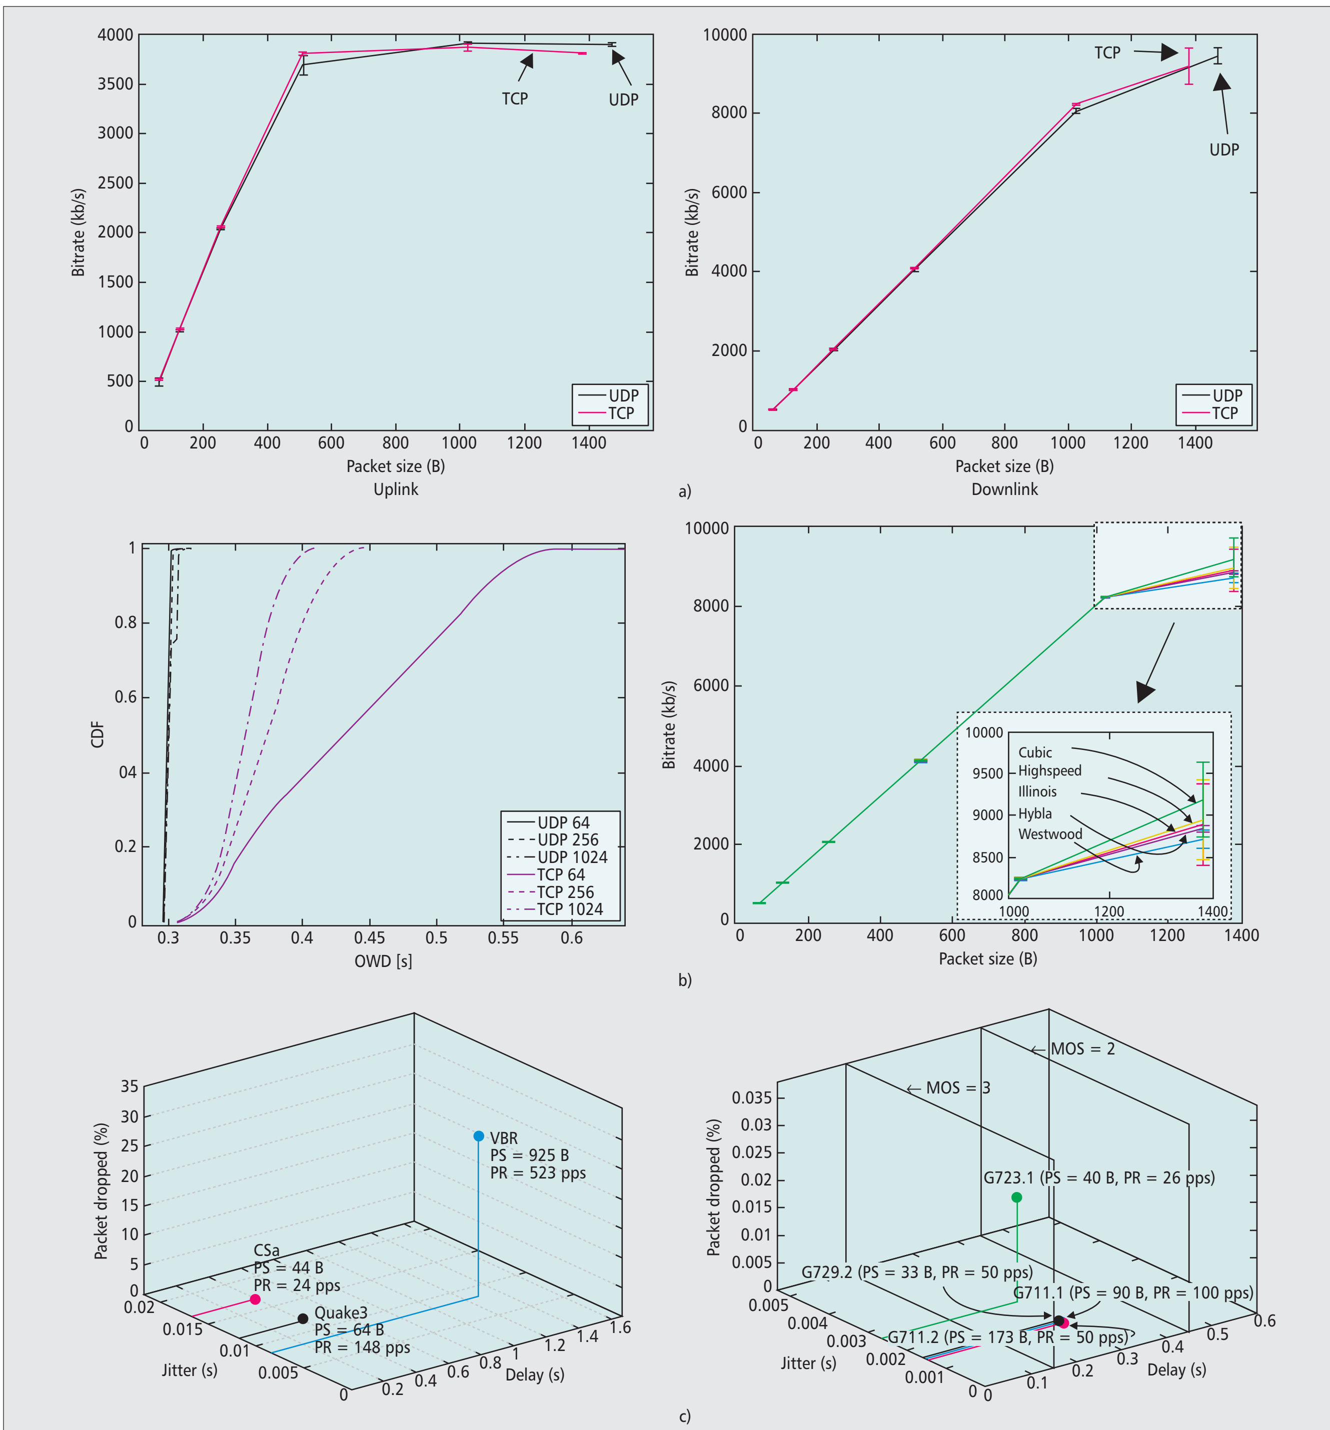

Figure 3. Performance of SGS.

help guide the experimenter to use the proper tools and methodologies.

Impact of TCP Versions — SGS is characterized by high bandwidth-delay product, and there are various TCP versions optimized for these kinds of networks. Moreover, the PEP used by SGS terminates the TCP connections of the clients just before the satellite network and opens new TCP connections, using an optimized TCP version. Furthermore, Internet users have different TCP versions that may interact in different ways with the PEP. Driven by these considerations, we tested the performance of different TCP versions on SGS to understand the impact of the various congestion control algorithms on throughput, jitter, and delay. In particular we used the TCP versions supported by the latest Linux kernel that are more indicated for this connection (*westwood*,[2] *illinois*,[3] *high speed*,[4] *hybla*,[5] and *cubic*[6]). We generated CBR and VBR traffic in the uplink and the downlink with different packet rates (PR) and packet sizes (PS), and we measured the throughput, jitter, and OWD. For space constraints, we only report the throughput (average and standard deviation over 100 experiments) obtained with the considered TCP versions, shown in the right plot of Fig. 3b; minor differences are only notable in saturation (see the zoom in right plot of Fig. 3b). **This means that the PEP is actually able to decouple the satellite connection from the LAN of the client, so that its TCP version does not impact the performance. Among the minor differences, we can note that cubic, the default TCP version of Linux, obtains on average the best performance.**

Video, Game, and Voice Performance — We also tested the performance achieved by SGS with other applications. In particular, we tested if and how video, game, and voice applications can properly work on this network. Using D-ITG, we generated traffic emulating these applications and we measured the OWD, jitter, and losses. The left plot of Fig. 3c shows the average values (over 100 measurements) of these parameters with video and game traffic in the uplink. For video traffic (VBR in the figure), we used a measurement methodology based on the model presented in [13]: PR equal to 720 pkt/s and random PS with Normal distribution. With this traffic on SGS we measured about 8 ms of jitter, 1.3s of delay, and 25 percent of packet loss, which may severely impact the quality of the video and therefore the Quality of Experience (QoE). For game traffic, we used the parameters from the models presented in [14]: bimodal PS and Student inter-packet times (IPT) for Counter Strike (CSA in the figure) and Normal PS and Exponential PR for Quake3. These applications obtained better performance than the video: the jitter is about 16 ms for Counter Strike and 12 ms for Quake3, the delay is about 400 ms, and the packet loss is 0 for both games. This behavior is mainly due to the smaller bitrate and packet rate of game traffic with respect to video traffic. The performance of voice applications in the downlink are reported in the right plot of Fig. 3c. For this application we tested different codecs (G711.1, G711.2, G723.1, G729.2), according to the VoIP models of D-ITG. In the tri-dimensional plot we also report the planes related to the Mean Opinion Score (MOS) using a simple model[7] based on the ITU E-model. As shown, with all the codecs we obtain a MOS between 2 and 3. However, with codec 723.1 we observe a small packet loss of about 0.2 percent. Therefore, **as for video, game, and voice, the QoS of satellite IP services seems to be still far from the expectations**. In the following section we will see what happens when real applications and real users are involved in the tests.

Performance of Satellite IP Services with Real Applications and Real Users: FGS vs SGS

In this section we present the results of measurement experiments conducted with the aim to compare SGS and FGS and to test if real users are satisfied when using real applications on these access networks. In particular, Table 1 shows the MOS obtained using Skype with voice only or voice and video. The test was conducted with 70 users and with different kinds of access connections (ADSL, corporate networks, etc.) toward FGS and SGS. The users were asked to make a two-minute voice only, and voice+video call toward the same user connected through the satellite networks and then to assign a quality score between 1 and 5. Table 1 shows that the two satellite access networks provide a very high MOS (average values of 4.25 and 4.35), and therefore QoE, when making voice calls through Skype. The standard deviation shows also that the MOS is consistent among the users. SGS shows a slightly higher MOS than FGS. However, when making video calls through Skype, the experience is much less appreciated by the users, with the average MOS going down to 2.55 with FGS and 2.72 with SGS. The small standard deviation value, again, shows that these values are consistent among the users. Thus, **our performance measurements highlight how multimedia applications (especially when using video too) are still hardly supported by new generation satellite services**.

Black-Box Analysis of the FAP

As introduced above, both FGS and SGS employ a bitrate limiter (called FAP), which starts operating when the user exceeds certain thresholds on volume of traffic produced. In this section we report results of our black-box analysis conducted to: test how the FAP calculates the volumes of user traffic; test how the FAP limits the bitrate allowed to the user; and explore possible ways for the user to best use the available traffic volume.

To answer the first two questions we performed a six-month experimental campaign using a measurement client with the FAP enabled. The left plot of Fig. 4 reports the volume of traffic generated with D-ITG and the volume of traffic seen by the FAP (accessible through the OSS of the satellite operator). As shown, the tests revealed that the FAP operates on 15-minute sliding time windows (updated every 15 minutes) and has a 15-minute delay with respect to when the actual volume has been produced (i.e. if x MB have been produced up to time t_i, such volume will be seen by the FAP at time t_i + 15min). This means that the user can

SGS is characterized by high bandwidth-delay product, and there are various TCP versions optimized for these kinds of networks. Moreover, the PEP used by SGS terminates the TCP connections of the clients just before the satellite network and opens new TCP connections, using an optimized TCP version.

[2] *http://c3lab.poliba.it/index.php/Westwood:Linux*

[3] *http://www.cs.fsu.edu/~baker/devices/lxr/http/source/linux/net/ipv4/tcpillinois.c*

[4] *http://www.icir.org/floyd/hstcp.html*

[5] *http://hybla.deis.unibo.it*

[6] *http://research.csc.ncsu.edu/netsrv/?q=content/bic-and-cubic*

[7] *TeKtronix, Common VoIP Service Quality Thresholds, http://www.tek.com/document/poster/common-voip-service-quality-thresholds.*

	Audio		Video	
	Average	Standard deviation	Average	Standard deviation
FGS	4.25	0.84	2.55	0.64
SGS	4.35	0.74	2.72	0.67

Table 1. MOS with Skype on FGS and SGS.

actually exceed the traffic volume for 15 minutes before being capped by the FAP, and the FAP requires 15 minutes of additional time before removing the cap from the user. Other important considerations have also been drawn from these tests. First, the FAP considers the volume of traffic at the network level (i.e. it looks at the size of the layer-2 payload). Second, and more important, **the cap introduced by the FAP operates as a drop-tail queue with packet granularity**. This causes bursty losses, which can severely impact the performance of TCP [15] of capped users.

Having characterized the behavior of the FAP, we developed a simulator in which we implemented a FAP to investigate the possible ways to interact with it. For example, the simulator allows us to understand how a user can make the best of the available traffic volume (i.e. to find an answer to the third question), to understand how the operators can set the thresholds in the most convenient way, and so on. The simulator is written in Matlab and allows us to reproduce the behavior of a user that has to download/upload a certain volume of traffic and that of the FAP that limits the throughput in case of threshold violation. Adopting the point of view of the user, we then performed a set of experiments aimed at exploring how the user can best use its available volume of traffic. For studying such interesting case, we use values consistent with those of a typical contract from the satellite operator and we refer to a particular working real example. We aimed at answering questions such as the following: given that the user has a maximum allowed volume of 15 GB of traffic per month before being severely capped, how can he/she download[8] such volume in the shortest time?

The right plot of Fig. 4 reports the time required to download 15 GB on FGS using three different approaches:

1 Standard-free, which consists of downloading without stopping until the 15GB are finished, and it implies that the download rate will vary, depending on the caps imposed by the FAP when exceeding the thresholds.

2 Standard-block, which consists of downloading only when not capped (i.e. the user stops when the cap is enforced and resumes the download when the cap is removed), and it implies that the download rate will always be equal to the maximum (i.e. the available bandwidth).

3 Advanced, which consists of downloading, in each 15-minutes time slot, a volume of bytes that allows the user not to be capped in the current and the following time slot.

We performed this analysis instrumenting the simulator for considering different values of the maximum download rate allowed by the network when not capped (i.e. the available bandwidth of the connection). The reason is that we verified that the available bandwidth of FGS is often below the nominal bandwidth of 3.6 Mb/s. As shown in the right plot of Fig. 4, the *advanced* approach makes it possible to reduce the download time to about 55 hours, while the minimum time is about 170 with the *standard-block* approach and 125 hours with the *standard-free* approach. However, if the network does not provide more than 2.75 Mb/s of available bandwidth, the *standard-block* approach makes it possible to obtain minimum download times. Summarizing, **smart download approaches can help cope with the FAP if the available bandwidth is higher than 2.75 Mb/s**. This result is clearly depending on the specific values of the FAP configuration parameters (mainly being the cap thresholds and rates) chosen by the operator. Changing these parameters may result in changing the best way to download a specific volume of Bytes or, in more general terms, may affect more or less the experience of the users. Moreover, as the conditions of the network are variable, dynamic parameter values (e.g. depending on the network status) may provide improved network conditions and user experience. In conclusion, the simulator allows the user and the operator to understand how to properly interact with and operate the FAP.

Conclusion

During the past few years we have set up a testbed and we have conducted a measurement campaign also comprising real users. To the best of our knowledge this is the first article providing quantitative results of a comprehensive performance measurement of new generation satellite Internet access comprising: real operational satellite networks; real users; both synthetic (but realistic) and real traffic. Also, for the first time in the literature our black-box analysis of the FAP sheds light on its behavior. In particular, we showed how new-generation satellite networks have much higher performance with respect to the previous generation, reaching throughput values up to about 4 Mb/s in the uplink and 10 Mb/s in the downlink. Thus, the integration of new-generation satellite Internet services with high speed wireless backbones [16] can be a way to connect people and communities, especially in rural areas. We saw how the PEP impacts measurement results, especially in terms of latency. We observed that real applications can obtain good performance on SGS, even though multimedia applications, especially with video, are still hardly supported. Finally, we saw that intelligent approaches can help users achieve smaller download times when subjected to the FAP. Thanks to this work we have real and updated figures regarding the performance of new-generation satellite access technologies, also quantifying the performance limitations with isochronous applications.

References

[1] A. Botta and A. Pescapé, "New Generation Satellite Broadband Internet Services: Should ADSL and 3G Worry?" TMA 2013, co-located with IEEE INFOCOM 2013, April 2013.

[8] *The analysis can easily be extended to the case of upload or mixed download/upload.*

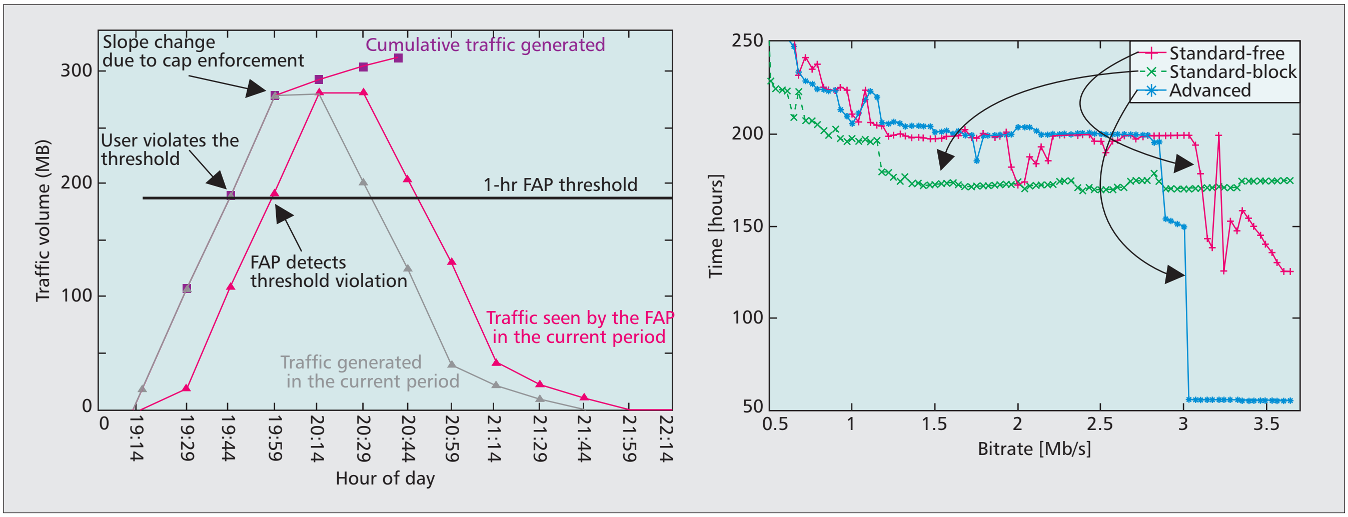

Figure 4. Black-box analysis of the FAP: capping behavior (left) and download time for 15 GB data versus the available bandwidth (right).

[2] P. Chitre and F. Yegenoglu, "Next-Generation Satellite Networks: Architectures and Implementations," *IEEE Commun. Mag.*, vol. 37, no. 3, Mar. 1999, pp. 30–36.

[3] C. Metz, "IP-Over-Satellite: Internet Connectivity Blasts Off," *IEEE Internet Computing*, vol. 4, no. 4, Jul./Aug. 2000, pp. 84–89.

[4] Y. Hu and Li V.O.K., "Satellite-Based Internet: A Tutorial," *IEEE Commun. Mag.*, vol. 39, no. 3, Mar. 2001, pp. 154–62.

[5] T. Taleb, Y. Hadjadj-Aoul, and T. Ahmed, T, "Challenges, Opportunities, and Solutions for Converged Satellite and Terrestrial Networks," *IEEE Wireless Commun.*, vol. 18, no. 1, Feb. 2011, pp. 46–52.

[6] Y. F. Hu *et al.*, "Broadband Satellite Multimedia," *IET Communications*, vol. 4, no. 13, Sep. 2010, pp. 1519–31.

[7] C. Partridge and T. J. Shepard, "TCP/IP Performance over Satellite Links," *IEEE Network*, vol. 11, no. 5, Sept./Oct. 1997, pp. 44–49.

[8] C. Barakat, E. Altman, and W. Dabbous, "On TCP Performance in a Heterogeneous Network: A Survey," *IEEE Commun. Mag.*, vol. 38, no. 1, Jan 2000, pp. 40–46.

[9] I. F. Akyildiz, G. Morabito, and S. Palazzo, "TCP-Peach: A New Congestion Control Scheme for Satellite IP Networks," *IEEE/ACM Trans. Net.*, vol. 9, no. 31, June 2001, pp. 307–2.

[10] C. Caini, R. Firrincieli, and D. Lacamera, "PEPsal: A Performance Enhancing Proxy for TCP Satellite Connections," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 22, no. 8, Aug. 2007, pp. B-9–16.

[11] A. Botta and A. Pescapé, "Monitoring and Measuring Wireless Network Performance in the Presence of Middleboxes," *Wireless On-Demand Network Systems and Services (WONS)*, 2011.

[12] A. Botta, A. Dainotti, and A. Pescapé, "A Tool for the Generation of Realistic Network Workload for Emerging Networking Scenarios," *Computer Networks* (Elsevier), vol. 56, no.15, Oct. 2012, pp. 3531–47.

[13] W. Garrett and W. Willinger, "Analysis, Modeling and Generation of Self-Similar VBR Video Traffic," *ACM SIGCOMM CCR.*, vol. 24, no. 4, 1994, pp. 269–80.

[14] W. C. Feng *et al.*, "A Traffic Characterization of Popular On-Line Games," *IEEE/ACM Trans. Net.*, vol. 13, no. 3, 2005, pp. 488–500.

[15] E. Altman, K. Avrachenkov, and C. Barakat, "TCP in Presence of Bursty Losses," *Performance Evaluation*, vol.42-23, Sep. 2000, pp. 129–47.

[16] A. Botta *et al.*, "Experimental Evaluation and Characterization of the Magnets Wireless Backbone," *ACM WINTECH 2006*, Los Angeles (USA), pp. 26–33.

Biographies

ALESSIO BOTTA (a.botta@unina.it) is a postdoc in the Department of Computer Engineering and Systems of the University of Napoli Federico II (Italy). He graduated in Telecommunications Engineering (M.S.) and obtained the Ph.D. in Computer Engineering and Systems, both at University of Napoli Federico II. His research interests are in the area of networking and, in particular, in the area of network performance measurement and improvement, with focus on wireless and heterogeneous systems. He has coauthored more than 40 international journal (*IEEE Communications Magazine*, *IEEE Transactions on Parallel and Distributed Systems*, *Elsevier Computer Networks*, etc.) and conference (IEEE Globecom, IEEE ICC, IEEE ISCC, etc.) publications. He has served and serves several technical program committees of several international conferences (IEEE Globecom, IEEE ICC, etc.) and he acts as a reviewer for different international conferences (IEEE Infocom, etc.) and journals (*IEEE Transactions on Mobile Computing*, *IEEE Network*, *IEEE Transactions on Vehicular Technology*, etc.) in the area of networking. In 2010 he was awarded with the best local paper award at IEEE ISCC 2010. He has served and serves as an independent reviewer of research and implementation project proposals for the Romanian government.

ANTONIO PESCAPÉ [M'00, SM'09] (pescape@unina.it) is an Assistant Professor in the Department of Computer Engineering and Systems at the University of Napoli Federico II (Italy) and Honorary Visiting Senior Research Fellow at the School of Computing, Informatics and Media of the University of Bradford (UK). He received the M.S. Laurea Degree in Computer Engineering and the Ph.D. in Computer Engineering and Systems, both at the University of Napoli Federico II. He teaches courses in Computer Networks, Computer Architectures, Programming, and Multimedia, and he has also supervised and graduated more than 100 among BS, MS, and Ph.D. students. His research interests are in the networking field with a focus on Internet Monitoring, Measurements and Management and on Network Security. He has co-authored over 130 journal (*IEEE ACM Transaction on Networking*, *Communications of the ACM*, *IEEE Communications Magazine*, *JSAC*, *IEEE Wireless Communications*, *IEEE Networks*, etc.) and conference (SIGCOMM, Conext, IMC, PAM, Globecom, ICC, etc.) publications, and he is co-author of several patents pending. He has served and serves as workshop and conference Chair and on more than 90 technical program committees of IEEE and ACM conferences, including IEEE ICC (NGN symposium). He serves as an Editorial Board Member of the *Journal of Network and Computer Applications* and has served as an Editorial Board Member of *IEEE Communications Surveys and Tutorials* (2008-2011), and was guest editor for the special issue of *Computer Networks* on "Traffic classification and its applications to modern networks". For his research activities he has received several awards. He is a Senior Member of the IEEE, and has served and serves as an independent reviewer/evaluator of research and implementation projects and project proposals co-funded by the EU Commission, Swedish government, several Italian local governments, Italian Ministry for University and Research (MIUR), and Italian Ministry of Economic Development (MISE).