PERFORMANCE OF AUDIO/VIDEO SERVICES ON CONSTRAINED VARIABLE USER ACCESS LINES

M. Vilas, X. G. Pañeda, D. Melendi, R. Garcia and V. Garcia
Computer Science Department, University of Oviedo, Campues Viesques sn, Xixón/Gijón, Spain

Keywords: User access line, audio/video services, constrained, variable, heavy loaded.

Abstract: Nowadays, it is more and more common for the same access line to be shared among different services and even among different users. This change in home users' behaviour, that has given rise to resource consumption close to the maximum available in user access lines, is mainly due to the increase in subscriber access capabilities which have taken place in the last few years. At the same time, contracts fulfilled by customers and network operators only provide guarantees for a reduced percentage of the maximum download/upload capacity of the line. In this paper, a study of the effects on streaming services caused by variations on the access line and by the traffic of other services is carried out. One of the main conclusions of the paper is that the delivery rate of *UDP* streaming sessions is mainly guided by the quality of the contents and does not consider the congestion in the network. For this reason, a method for delivery rate estimation for *UDP* streaming sessions is presented.

1 INTRODUCTION

The technological evolution of user access lines which has occurred in the last ten years has led to deep changes in user behavior. The number of broadband access lines has increased and audio/video streaming services, digital newspapers, virtual offices, email, chats, p2p, etc are services which are progressively more and more common.

In the contracts fulfilled by customers and network operators, only a reduced percentage of the maximum download/upload capacity of the line is guaranteed, sometimes around 10% of the maximum. Thanks to this margin, network operators can reach a balance between investment in resources on access networks and the level of service provided to end users. Depending on the number of users active in each access segment the capacity of these access lines can vary.

Another parameter that can affect user experience is the congestion of the access line. It is common to find users that are sharing files using p2p, while downloading *PDF* files using http and meanwhile accessing an audio/video service looking for the cinema trailers of next week's premieres.

For these reasons, it is very interesting to analyze the effects of bandwidth restrictions on streaming technology, the effects of other types of user traffic over streaming services and vice versa. Using this information, content managers and network operators can clearly establish the minimum service levels. In this way, questions such as how much quality can be offered to a streaming session in the most loaded hours of the day and how the quality of streaming sessions is affected can be answered.

In this paper an analysis of the effects on streaming services of changes in the delivery rate of access lines of cable network users is performed. This study is carried out considering another traffic sharing the bandwidth of the line causing resource utilization in the user access line close to 100%.

The rest of the paper is organized as follows. In section 2, previous work in the same field is analyzed. After that, the test-bed for performance evaluation is described. In Section 4 the effect of user access line variations with constrained bandwidth over streaming services is described. Based on the results of previous sections, an estimator of available bit rate that improves the performance on variable congested user access lines is presented in section 5. Finally, Conclusions and Future Work are presented.

Vilas M., G. Pañeda X., Melendi D., Garcia R. and Garcia V. (2007).
PERFORMANCE OF AUDIO/VIDEO SERVICES ON CONSTRAINED VARIABLE USER ACCESS LINES.
In *Proceedings of the Second International Conference on Signal Processing and Multimedia Applications*, pages 22-27
DOI: 10.5220/0002134900220027

2 PREVIOUS WORK

The analysis of the traffic generated by some of the most common streaming platforms, *Real Networks* and *Windows Media*, has been treated in Li, Claypool & Kinicki (2002) and Kuang & Carey (2002). These papers analyzed packet inter arrival times, packet sizes, and the generated packet rate at different levels.

The co-existence of streaming traffic and other types of applications has been intensively researched previously. In Chung & Claypool (2006), the authors deploy a study of the fairness of *Real Networks* streaming flows delivery rate consumption when they have to share resources with other *TCP* flows. The main conclusion of this study is that *Real Networks* only present a *TCP* friendly behavior when the encoding quality was less than the fair share of the capacity. In Boyden, Mahanti & Williamsom (2005), the authors highlight that the non-*TCP* friendly nature of *Real Networks* streaming was increased when contents are delivered using *Turboplay* (RealNetworks, n.d.). In Doshi & Cao (2003), authors present a detailed study on the mutual effects of different types of traffic (*TCP* and *UDP* flows) over extremely low bandwidth *WAN* links (128Kbps). In spite of the interesting results, bitrates of nowadays user access lines and backbone links are several times higher than the bitrates considered in this work.

The design of protocols to deliver streaming contents, maintaining fairness with other types of services, is another interesting research field with a large number of different approaches like Wu, Claypool & Kinicki (2005), Song, Chung & Shin (2002), Handley, Floyd, Pahdye & Widmer (2003) or Balk, Gerla & Sanadidi (2003).

3 EXPERIMENTAL SETUP

With the goal of analyzing the influence of user access lines variations, the particular architecture of the service or the elements of the core network are not representative. For this reason we have focused the analysis on the emulation of the user access line.

To emulate different network conditions and user access lines, we have combined *Linux* queuing management (Iperf, n.d.) and a kernel module that emulates packet loss and jitter, called *NetEM* (TC, n.d.) as can be seen in Figure 1. We have used as reference qualities for Downstream/Upstream (*US/DS*) channels those offered by the Asturian network operator *Telecable*: 640Kbps/128Kbps, 2Mbps/320Kbps and 4Mbps/640Kbps.

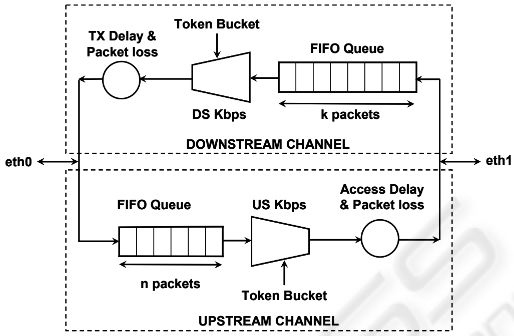

Figure 1: Model of the user access line implemented using a linux router, 2FastEthernet cards and *Netem*.

To restrict the maximum bit rate that a user can use to download/upload data, we have defined two Token Buckets applied in the interfaces of the router. This method of controlling bandwidth is used, for example, in cable networks (Laskshminarayanan and Padmanabhan, 2003).

Usually, network operators do not apply special *QoS* policies for different types of user traffic, except for *VoIP* telephony traffic. For this reason we have taken the decision of deactivate the default *Linux Advanced Queuing Management* (*AQM*). Substituting *AQM* and including a *FIFO* (First In - First Out) queue management. The size of the buffers is set to a maximum latency of 1024ms, the default on Cisco *CMTS* (Cable Modem Termination System) (Kennedy and Atov, 2006).

We have also included *NetEm* module in the *DS/US* to emulate some types of network delays like Request-and-Grant control (Zhenglin, Chongyang, 2002). To know if with the current configuration of network operator *Telecable*, the effect of Request-and-Grant delays is significant we have tested a real user connection. During the most loaded hours of the day, *RTT* (Round Trip Times) between user and a server installed close to the *CMTS*, was nearly constant and equal to 69msec±3msec. Only 0.02% of the *RTT* measurements present higher values.

Providing that the target of the study is the influence of user access line constraints and variations, we have deployed in the testbed a simplified version of the network and the different services that share the resources of the access line (Figure 2).

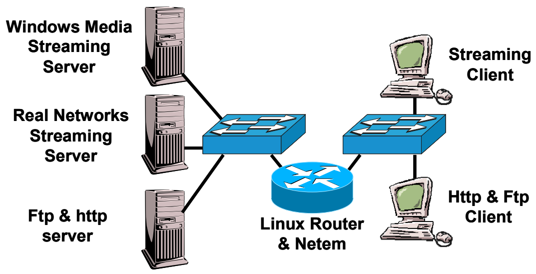

Figure 2: Testbed for the evaluation of streaming services performance on constrained variable user access lines.

4 EXPERIMENTAL RESULTS

4.1 Single Streaming Session

We have tested the behavior of streaming sessions when the channel capacity progressively decreases from the maximum capacity of 4Mbps to 350Kbps, and then, after reaching the minimum value of 350Kbps the quality was recovered upon reaching 4Mbps again. Intermediate values were 3.5Mbps, 3Mbps, 2.5Mbps, 2Mbps, 1.5Mbps, 1Mbps, 750Kbps, 500Kbps and 350Kbps. Contents were produced using multi-rate encoding, with qualities of 700, 450, 225, 100, 50Kbps.

In the case of *Windows Media* delivery, the *TCP* streaming connection consumes as much resources as is possible (Figure 3). Unfortunately, using windows streaming platform, each change in the quality is performed stopping content reproduction with the selected quality and starting a little later with the new one. This can be due to the *TCP API* that hides most network congestion indicators (Chung and Claypool, 2006).

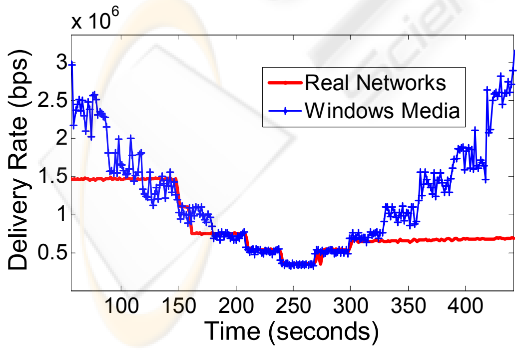

Figure 3: Delivery rate for Windows Media and Real Networks when access line conditions change.

We have found that *Real Networks* streaming sessions tend to react quickly to decrease in user access lines. In spite of the use of *Turboplay*, the delivery rate is not always equal to the available bit rate of the line, for example, after setting 1Mbps as the delivery rate, the delivery rate of *Real Networks* was set to 700Kbps and the selected quality was 450Kbps. Once the channel conditions improved, the quality of the content was not recovered at the same speed. On most occasions, quality was maintained at one of the lowest qualities until the end of the session. When an adjustment of quality occurs, this is done without stopping the playback of contents, transparently for the client.

4.2 Streaming Session and FTP

To perform an analysis of the mutual influence between *TCP* and *UDP* streaming traffic, we perform some tests downloading a file using *FTP* and after that, a streaming session is started. The bit rate of the downstream channel progressively decreased from 4Mbps to 200Kbps.

Results for *Real Networks* streaming are shown in Figure 4. As can be seen, two different periods show an aggressive behavior: the first 20 seconds of the session and when channel capacity goes under 300Kbps. After that, there is a sudden decrease down to 800Kbps. In the middle of these two values, the bit rate achieved by the *FTP* session is almost double that achieved by the streaming session.

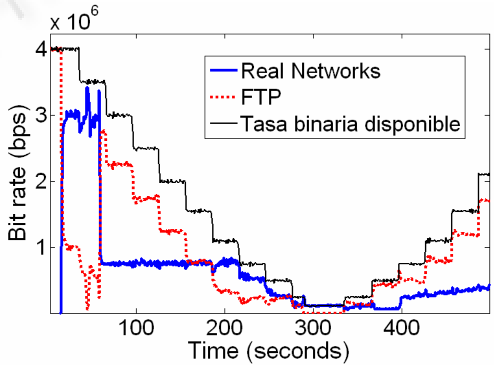

Figure 4: Real Networks streaming session.

Initial delivery rates used by *Real Networks* clients are not selected based on any measurement of the actual situation of the line. This initial value relies on the analysis of attainable delivery rate performed during the installation phase. At the same time, despite the fact that *FTP* obtains higher bandwidths than the streaming session the adaptation of streaming session delivery rate is performed on extremely negative conditions. The process of adjust the delivery rate seems to be guided by the quality of the playback and the encoding qualities of the stream.

In the case of *Windows Media* streaming (Figure 5) it can be seen that *Windows Media* streaming does not share bandwidth with *FTP* for bit rates in the access line higher than 2Mbps. When channel conditions go below 300Kbps, the streaming session slowly gets starved of bandwidth. This is possibly due to a failure of the quality adjustment on *Windows Media* platform that, as is stated previously, reacts slowly to network congestion. Approximately 30% of *Windows Media* streaming sessions have achieved a maximum delivery rate of 500Kbps in presence of *FTP* competing traffic.

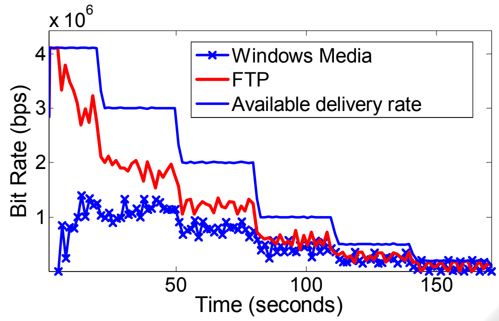

Figure 5: Windows Media streaming session and FTP delivery rate when bit rate decreases down to 350Kbps.

5 DELIVERY RATE ESTIMATION FOR UDP STREAMING

Once it is clear that *Real Networks* streaming sessions react slowly to improvements in the delivery rate and that they do not consider the congestion caused in the network, more advanced techniques are needed to adapt the consumption of streaming flows to the real conditions of the channel.

With this goal, we have analyzed the delivery rate, packet loss, and inter arrival time of *UDP* flows under different network conditions. Figure 6 shows the *CDF* function of inter-arrival times for a 1Mbps *UDP* over different access lines conditions. Results are presented with and without simultaneous *FTP* session. As can be seen, as the *UDP* flow consumption grows variation of inter arrival time gets bigger. Same results are obtained maintaining the quality of the access line and varying the delivery rate of the flow.

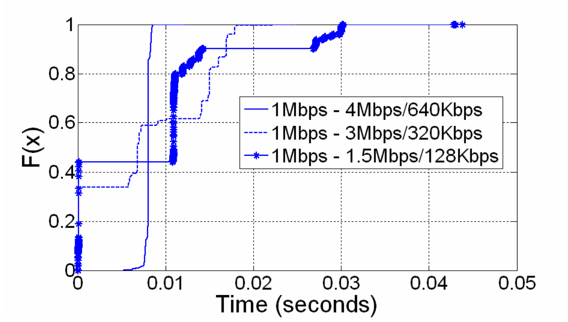

Figure 6: CDF of packet inter arrival time for different access line conditions.

Based on this analysis, it is possible to add an extra source of information to *UDP* streaming in order to detect and react to variations in the user access line quality. *VTP* (Balk, Gerla & Sanadidi (2003) estimates the correct value for the delivery rate based on the actual delivery rate measured on the receiver's side plus a smoothing function. The approach presented in this paper is based on the values of the variation of inter arrival times in the receiver and inter departure times in the sender. *BART* algorithm (Ekelin, Nilsson, Hartikainen, Johnsson, Mangs, Melander and Bjorkman, 2006), has previously used this type of estimation in order to evaluate end to end available bandwidth on a network path. Unlike *BART*, that generates extra congestion on the network in order to estimate available bandwidth in real time, the approach presented in this paper relies on the data packets generated by the streaming sessions.

Adding to sent packets the departure time from the server, it is possible to obtain the difference between departure times of consecutive packets. At the same time it is possible to register arrival times and calculate the inter arrival time. In this way, it is possible to compare these two values and take a decision about channel conditions. Using as a measure of the congestion of the line the inter-packet strain used in *BART*, it is possible to analyze channel conditions and react correctly. Inter-packet strain (ε) is defined as:

$$1+\varepsilon_i = {t_i}/{t_i^*}$$

Being t_i the inter arrival time of consecutive packets and t_i^* the inter departure time of the same pair of packets. The formula shown before has been modified from the original source in order to avoid sudden increases due to packet compression; two consecutive packets that arrive extremely close to the destiny. With inter-arrival times close to zero, *BART* estimator presents great oscillations. For this reason we have interchanged the position of t_i and t_i^*.

From empirical observation (Figure 7) it is possible to conclude that the average of inter-packet strain increases as the delivery rate of the *UDP* flow gets close to the available delivery rate. Whereas, if the delivery rate increases over 80% of the delivery rate of the access line, average value decreases as long as the interference caused by the *FTP* session decreases. Also, the variance of inter-packet strain increases as the delivery rate of the *UDP* flow gets close to the available delivery rate.

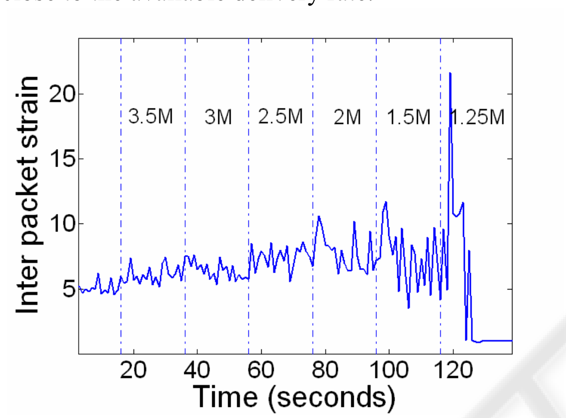

Figure 7: Inter-packet strain for a 1Mbps UDP flow and different DS rates (Mbps).

5.1 Experimental Results

We have tested the behavior of inter-packet strain as an estimator of channel congestion. We have deployed a simplified version of a streaming server and client that establish a *TCP* control connection. After a short negotiation, the server starts sending *UDP* packets at a predefined rate to the client. These packets include the departure time. The client estimates channel conditions based on inter-packet strain and takes a decision about the suitability of a delivery rate adaptation. Delivery rate adaptation messages are sent from the client to the server using the *TCP* control connection in the same fashion as *Real Networks* streaming.

In this first approach only fixed packet sizes are considered and the thresholds to cause a delivery rate adaptation were determined experimentally during a set of trials. To avoid the effects of temporary drops in network conditions, the mean value and the variance of packet strain is calculated every 0.2 seconds. Possible delivery rate values were selected at 200Kbps steps, ranging from 200Kbps to 2Mbps, simulating a multiple-bit rate file. The value of the delivery rate, used as the base in the estimations, is obtained from the size of the packets divided by the difference between the departure times of two consecutive packets from the server.

Results, when available delivery rate varies, are shown in Figure 8. Using inter-packet strain in heavy loaded user access lines we have achieved values for the delivery rate that accomplish these basic properties:

- During phases of channel stability no adaptations are generated.
- When congestion increases, the delivery rate is adjusted, once or more, until the estimator shows no congestion. These adaptations are non erratic, following the same pattern (increase or decrease) as the conditions of the access line.
- The estimator behaves in a conservative way, maintaining *UDP* flow consumption to values close to one half of the attainable delivery rate or lower.

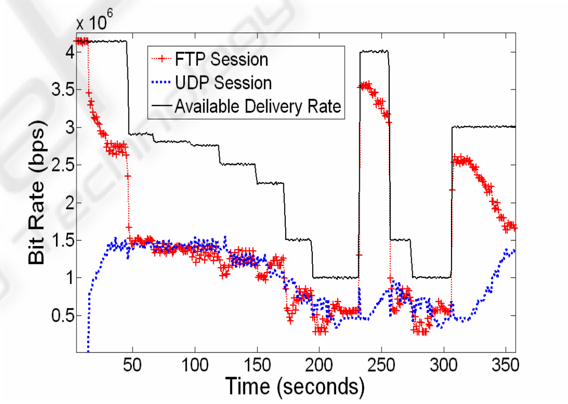

Figure 8: UDP streaming based on packet strain delivery rate estimation.

At the same time, it can be seen in Figure 8 that the *FTP* session never gets starved of bandwidth. Also, reaction to the increase of capacity is much faster than that obtained using *Real Networks* and follows the evolution of the access line, sharing the available delivery rate with the *FTP* session.

6 CONCLUSIONS

Previous studies have shown that *UDP* streaming is a bandwidth intensive application that does not perform a fair share of resources with *TCP* connections. As is described in this paper, the

Turboplay feature presents strong unfriendly behavior with other flows at the beginning of the session, especially when estimations of user access line delivery rate are aborted. Adaptability of *Real Networks* streaming sessions seems to be completely driven by the set of encoding qualities and the quality of the playback, without extra measurements of channel congestion.

To solve these problems, some techniques, like inter-packet strain measurement, could be useful to detect congestion and react so reducing it. In this way, the interference between streaming sessions and other services could be minimized. Combining session quality measurements (like percentage of packet loss and client buffer size) with congestion estimators is possible to improve the behavior of *UDP* streams.

7 FUTURE WORK

Streaming platforms typically use variable packet sizes and, some of them, even variable bit rates. It is necessary to study the reaction of inter-packet strain in the presence of variable packet sizes and variable inter arrival times. Also, the reaction of inter-packet strain when resources are shared with another type of services has to be carefully evaluated.

It is also necessary to perform an in-depth study of the suitability of inter-packet strain estimator when the packets generated in the server have to travel through several routers, interacting with a variable number of flows.

Finally, the integration of inter-packet strain on a real streaming platform is a very interesting field of study.

ACKNOWLEDGEMENTS

This work was supported in part by Telecable, La Nueva España and Asturies.com within the projects of NuevaMedia, Telemedia, ModelMedia and MediaXXI and the Spanish National Research Program within the project INTEGRAMEDIA (TSI2004-00979).

REFERENCES

Li, M., Claypool, M., Kinicki, R. (2002), *MediaPlayer versus Real Player – A Comparison of Network Turbulence*, Paper presented at ACM SIGCOMM, Marseille, France, 2002.

Kuang, T., Carey, W. (2002), *A measurement study of Real Media Streaming Traffic*, Paper presented at ITCOM, Boston, Massachusetts USA, 2002

Chung, J., Claypool, M. (2006), Empirical Evaluation of the Congestion Responsiveness of Real Player Video Streams, *Kluwer Multimedia Tools and Applications*, Volume 31, Number 2, November 2006.

Real Networks, (n.d), Retrieved March 12th, 2007, from

Boyden, S., Mahanti, A., Williamsom, C. (2005), *Characterizing the Behaviour of RealVideo Streams*, Paper presented at SCS, Philadelphia, USA, 2005.

Doshi, R., Cao, P. (2003), Streaming traffic fairness over low bandwidth WAN links, Paper presented at WIAPP 2003. San Jose, CA, USA, 2003.

Wu, H., Claypool, M., Kinicki, R. (2005), Adjusting Forward Error Correction with Temporal Scaling for TCP-Friendly Streaming MPEG, *ACM Transactions on Multimedia Computing, Communications and Applications*, Vol 1, Issue 4, November 2005.

Song. B., Chung, K., Shin, Y. (2002), SRTP: TCP-Friendly Congestion Control for Multimedia Streaming, *Paper presented at ICOIN*, Korea, 2002.

Handley, M., Floyd, S., Pahdye, J., Widmer, J., TCP Friendly Rate Control (TFRC): Protocol Specification, RFC 3448, January 2003.

Iperf: Testing the limits of your Network. TCP/UDP Bandwidth Management Tool, (n.d) Retrieved March 12th, 2007 from http://dast.nlanr.net/Projects/Iperf/

TC. Linux QoS Control Tool, (n.d), Retrieved March 12th, 2007 from http://www.rns-nis.co.yu/~mps/

NetEM. Network Emulation (n.d), Retrieved March 12th, 2007 from http://linux-net.osdl.org/index.php/Netem

Laskshminarayanan, K., Padmanabhan, V. (2003), Some findings on the Network Performance of Broadband Hosts, Paper presented at ACM conference on Internet measurement, Florida, USA, 2003.

Kennedy D., and Atov, I., (2006) Implementing a Testbed for the Evaluation of FAST TCP in DOCSIS-based Access Networks, Technical Report 060119A, Swinburne University of Technology, 2006.

Zhenglin, L., Chongyang, X., (2002), An Analytical Model for the Performance of the DOCSIS CATV Network, *The Computer Journal*, Vol 45, No. 3, 2002.

Ekelin, S., Nilsson, M., Hartikainen, E., Johnsson, A., Mangs, J.-E., Melander, B., Bjorkman, M., (2006), Real-Time Measurement of End-to-End Available Bandwidth using Kalman Filtering, 10th IEEE/IFIP NOM 2006, Vancouver, Canada.

Balk, A., Gerla, M., Sanadidi, M., (2003) Adaptive MPEG-4 Video Streaming with Bandwidth Estimation, *Lecture Notes In Computer Science*; Vol. 2601, 2003.