

Latency Reduction in Home Access Gateways with Shortest Queue First

Nabil Benameur
Orange Labs, France Telecom
nabil.benameur@orange.com

Fabrice Guillemin
Orange Labs, France Telecom
fabrice.guillemin@orange.com

Luca Muscariello
Orange Labs, France Telecom
luca.muscariello@orange.com

ABSTRACT

To overcome quality of experience degradation due to the saturation of the uplink of ADSL lines, we propose in this paper the Shortest Queue First (SQF) queuing discipline to sort packets out of the buffer associated with the uplink of a home gateway. SQF is a flow aware queuing discipline which consists of serving those flows creating the less backlog in the buffer. We show experimental results proving that SQF outperforms other flow based queuing disciplines such as Deficit Round Robin.

1. INTRODUCTION

The saturation of the uplink of ADSL lines is a major source of dissatisfaction for residential customers, especially those with rather small bit rates (say, from 2 to 12 Mbit/s downlink and up to 1 Mbit/s uplink). From measurements made in the Orange ADSL network in France, about 5 % of users saturate their uplink persistently. This may appear at first glance as a small proportion of the global population but these customers are the most active (sometimes geeks) and are very inclined to complaint against poor quality.

Uplink saturation is due to applications giving rise to massive uploads of data onto the network, like peer-to-peer and personal cloud. The major issue with the saturation of the uplink is that packets, which would need to rapidly traverse the home gate, are queued in big buffers and are too much delayed, causing latency for the associated applications. This is typically the case for ACK packets corresponding to data downloads by end users. Those packets are delayed, which results in most cases in low web surfing. The same occurs for VoIP calls in the Internet channel with or without video (e.g., Skype, Apple FaceTime, Google Talk). One way of remedying this situation would consist of giving head-of-line priority to such flows, leading however to a complex flow identification and management system in the home gateway, which should also be kept cheap for economical reasons.

Without explicit flow differentiation (for ACK flows, VoIP flows, etc.), a possible way to overcome the saturation of the uplink is to increase the uplink bandwidth. However applications are more and more greedy and even in the case of FTTH access, we have observed the saturation of the uplink (see for instance [9] for measurements from the Orange networks). As this may appear as a short term solution, it is not certain that users perceive an improvement of their quality of experience when web surfing or using Skype. A more reliable solution consists of performing flow based queuing, for instance Deficit Round Robin (DRR) [5, 7, 8].

More precisely, DRR aims at equally sharing the available bandwidth among active flows so that packets of each of them can traverse a big buffer and no flows get completely stalled. This algorithm, however, does not distinguish between flows creating big backlogs in a buffer and those flows with packets regularly spaced out (typically ACK and VoIP flows). Thus, DRR does not make any difference between bulk packet transfers and smooth flows. To overcome this shortcoming of DRR, we propose in this paper the Shortest Queue First (SQF) packet scheduling scheme [3, 2, 4, 1].

The basic postulate of SQF is that the smoothest flows are associated with time sensitive applications and it is essential to prioritize them in a buffer. A field trial carried out in France with real low bit rate customers has shown that this postulate is valid with regard to quality improvement.

2. SQF PRINCIPLES

SQF is an algorithm to sort packets of flows out of a buffer, typically the buffer associated with the uplink of a home gateway. A virtual queue is associated with a flow with backlogged packets in the buffer. In practical implementations, up to N flows can be associated with a virtual queue. If a packet of a new flow arrives while there are already N flows with backlogged packets in the buffer, the packet is virtually queued in a default queue. According to SQF, the virtual queue with the smallest amount of backlogged bytes is served (one packet in removed from this virtual queue). In the experimental results reported below, the value of N has been set to 128. All queues (including the default virtual queue) are served according the SQF algorithm. To manage possible overflow of the buffer, when an arriving packet causes buffer overflow, then packets of those queues with the largest amount of backlogged bytes are removed so as to make sufficient buffering space for the incoming packet. This discipline is known as Longest Queue Drop [8].

In practice the output buffer of a home gateway can be decomposed into several class based queues in order to implement Diffserv class based QoS. In general the queue with the lowest priority is for best effort flows. SQF, as well as DRR, is precisely intended to sort packets out of such a best effort queue.

3. EXPERIMENTAL RESULTS

The SQF algorithm has been implemented in a Livebox v2 (the home gateway by Orange) in software as a Linux qdisc. In this version, packets coming from the home towards the network are handled by the CPU of the Livebox (slow path), whose software is based on a 2.6.x Linux kernel.

We have run several experiments with and without SQF enabled in the Livebox for comparison, in a residential access with 12 Mbit/s downlink and 1 Mbit/s uplink. An Agilent tester has been used to run emulated traffic following different kind of workloads: web surfing, email transmission/reception, VoIP communications, peer-to-peer in background. Such transactions are randomly and repeatedly initiated so as to compute statistical values. To estimate the gain brought by SQF, we have computed the ratio of download times of web pages (see Table 1) with and without SQF. This table clearly shows a reduction of download times thanks to the implicit prioritization of ACK flows, being SQF active in the uplink channel. Additional experiments have shown that DRR is similar to FIFO queuing, especially for big Web pages. To estimate the impact of SQF

Web page		Download time with SQF
Small page	Google	Time divided by 3.7
Big pages	Orange	/3.1
	Alapage	/1.3
	snf	/2.9

Table 1: Experimental download times with and without SQF.

on voice calls, we have used the E-model [10] to automatically compute the Mean Opinion Score. Results are given in Figure 1. SQF has a discriminatory impact: Those calls which can be handled by a virtual queue have an excellent QoE while those calls with packets queued in the default queue have bad QoE. Other experiments have shown that the QoE of voice calls with DRR when the uplink is heavily saturated is always bad. Tests have also been made in the operational network for a sample of customers who have experienced a better QoE using SQF than in the case in which the uplink was upgraded to a higher rate.

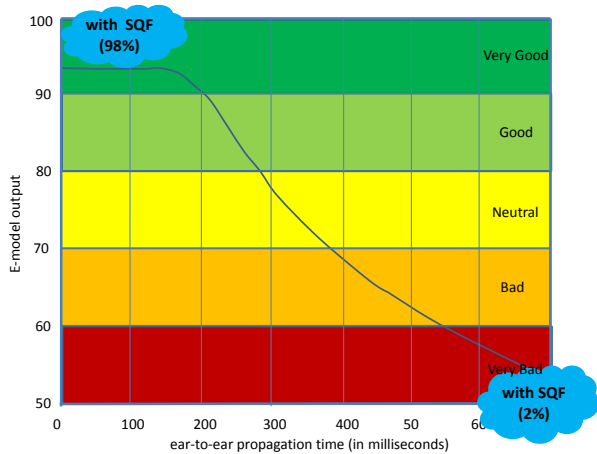


Figure 1: QoE of voice calls with SQF.

4. IMPLEMENTATION ISSUES

SQF has been implemented in the slow path of the Orange home gateway as a qdisc module in the Linux network stack. Such implementation is feasible for ADSL solutions not employing hardware processing offload used to implement fast path solutions. However, fast-path techniques tend to be

used in most of commercial chipsets for home gateways. In such a case, SQF should be implemented in the micro code of the chipset itself, which requires explicit agreements with the chipset company to develop and maintain such implementation.

5. CONCLUSION

Per-flow queuing is the most promising approach to reduce latency in the Internet at very low cost. In this paper, we have presented some performance gains, in terms of latency and download time reduction and QoE enhancement, by using the SQF per-flow scheduler in the uplink of an ADSL residential gateway. While the benefits of per-flow queuing are well known in the literature [5, 8] almost no relevant deployments have been made so far because of the misleading belief that per-flow queuing would not be scalable. However, it has been shown in the technical literature that per-flow queuing is feasible and scalable at all typical network utilization at any line rate (see [6] for instance).

Such mechanisms are not available today in QoS chipsets and software solutions, based on Linux, are the only available alternative. The evolution of current chipsets into distributed processing and fast path implementation breaks the development model based on a open programmable networking stack available in Linux. Therefore, lightweight mechanisms such as the SQF per-flow scheduler should be implemented in the micro code of chipsets.

6. REFERENCES

- [1] Bonald T. and Muscariello L. Shortest queue first: implicit service differentiation through per-flow scheduling. Demo at IEEE LCN 2009.
- [2] Carofiglio G. and Muscariello L. On the impact of TCP and per-flow scheduling on Internet performance (extended version). IEEE/ACM Transactions on Networking, April 2012.
- [3] Carofiglio G. and Muscariello L. On the impact of TCP and per-flow scheduling on Internet performance. In proc. of IEEE INFOCOM 2010.
- [4] Bonald T. and Muscariello L. and Ostallo N. Self-Prioritization of Audio and Video Traffic, In Proc. of IEEE ICC 2011.
- [5] Kortebi A.; Oueslati S. and Roberts J. Cross-protect: implicit service differentiation and admission control. In Proc. of IEEE HPSR 2004.
- [6] Kortebi A.; Muscariello L.; Oueslati S. and Roberts J. Evaluating the number of active flows in a scheduler realizing fair statistical bandwidth sharing. In Proc. of ACM SIGMETRICS 2005.
- [7] Shreedhar M. and Varghese G. Efficient fair queueing using deficit round robin. In Proc. of ACM SIGCOMM 1995.
- [8] Suter B.; Lakshman T.V.; Stiliadis D. and Choudhury, A.K. Design considerations for supporting tcp with per-flow queueing. In Proc. of IEEE INFOCOM 1998.
- [9] Vu Brugier G. Analysis of the impact of early fiber access deployment on residential Internet traffic. In Proc. of International Teletraffic Congress 2009.
- [10] ITU-T Recommendation G.107 : The E-Model, a computational model for use in transmission planning (2003)