

The Utility of Characterizing Packet Loss as a Function of Packet Size in Commercial Routers

Jose Saldana, Julián Fernández-Navajas, José Ruiz-Mas, Eduardo Viruete Navarro, Luis Casadesus
 Communication Technologies Group (GTC) – Aragon Institute of Engineering Research (I3A)
 Dpt. IEC. Ada Byron Building. CPS Univ. Zaragoza. 50018 Zaragoza, Spain
 e-mail: {jsaldana, navajas, jruiz, eviruete, luis.casadesus}@unizar.es

Abstract—This work presents a test called “*sizogram*” which characterizes the packet loss vs. packet size behavior of a router. This graph can help us to take some decisions about traffic, e.g. multiplexing or not, the number of samples or frames to be included in a packet, etc. In addition, it can be useful to have a better idea of the internal architecture of the router. Simulations have been carried out using different buffer implementations and sizes, and the results show the usefulness of this test.

Keywords—router buffer, packet loss, packet size

I. INTRODUCTION

The Internet was designed as a best-effort network, and for many years routers were optimized for managing big packets [1], which are typically used by services like e-mail or web browsing. Nowadays, the network infrastructure has to deal with real-time interactive services, which typically use smaller packets, as information has to be fragmented in order to be delivered in a short time.

Packet loss is one of the main causes of quality impairment in IP networks: in wireless environments, packets can be corrupted by the radio link, but in wired networks the main cause of packet loss is dropping in router queues [2]. Services based on TCP can use retransmission in order to recover lost packets, but real-time applications, which normally use UDP, can be seriously affected, as their interactivity makes it impractical to wait for the retransmission of a packet.

There is a relationship between packet size and loss probability: in wireless networks, as the transmission time increases, the probability of packet corruption grows; this effect was explored in [3]. The relationship for drop-tail buffers in wired networks was analytically studied in [4]. But commercial routers present different buffer implementations, which, although do not significantly modify their behaviour regarding delay and jitter, can affect the results for packet loss. For example, some routers measure their buffers in number of packets, while others do it in bytes or defining a queuing time limit [5]. Other problem is that, in addition to bandwidth limit, the processing capacity of the router establishes an upper bound in packets per second (pps) [1], [6].

Some studies [2], [7], [8] have characterized the behaviour of the public Internet in terms of packet loss. This characterization can be useful in order to take some decisions, e.g. if we multiplex a number of VoIP flows, we can reduce bandwidth, but at the cost of increasing packet size, and this will have an influence on packet loss, depending on the

implementation and buffer size of our router [9]. The number of samples per packet also affects packet size. In [7] a study was conducted, exploring how the number of video frames included in each packet could modify packet loss patterns.

In this work we present a test called “*sizogram*”, which characterizes the packet loss vs. packet size behaviour of a router. In addition to helping us make decisions on the packet size, it can also give us an idea of the router implementation, as the internal architecture of many commercial routers is not easily available [10]. Its interest is not in the absolute values but on the shape of the graphs. Although we will only conduct it to characterize a router, as future work we expect it will also be useful so as to characterize real Internet links.

II. TESTS AND RESULTS

For each packet size, which corresponds to a point of the graph, two Poisson traffics are sent to the same buffer: a *test* traffic with a fixed packet size, and a *background* traffic which is 100% of the link capacity, in order to saturate the buffer and thus producing packet loss. The distribution of *background* traffic includes small, medium and big packets: 50% of the packets are of 40 bytes, 10% are of 576 bytes and 40% of 1,500 bytes [11]. Traffic amounts of *test* traffic are at least two orders of magnitude smaller than *background* ones so as not to modify the size distribution of *background* traffic. We will not depict delay or jitter, because, as expected, their graphs do not significantly vary with *test* traffic packet size.

In the concrete tests we present, the capacity of the link is 10 Mbps, and the *test* traffic is 100 kbps. The duration of each test, which ranges from 200 to 3,000 seconds, has been calculated in order to have 25,000 packets of *test* traffic. The tests have been conducted using Matlab simulation.

Two buffer implementations have been tested: one measures the buffer size in bytes (*byte-sized*) and the other measures it in packets (*packet-sized*). Buffer sizes have been selected as *tiny*, as they are the most suitable for real-time applications [10]. They use drop-tail queues: if there is no place for the entire packet at the queue, it is discarded.

Fig. 1 shows the *sizogram* for a *byte-sized* buffer ranging from 10 to 100 kB. The graphs are monotonically increasing, as the bigger the packet, the bigger the probability of having no place at the queue. The last point (1,500 bytes) presents a small decrease because in this case there is a coincidence between *test* and *background* packets of 1,500 bytes.




Figure 1. Sizogram for *byte-sized* buffer of 10kB, 20kB, 50kB and 100kB

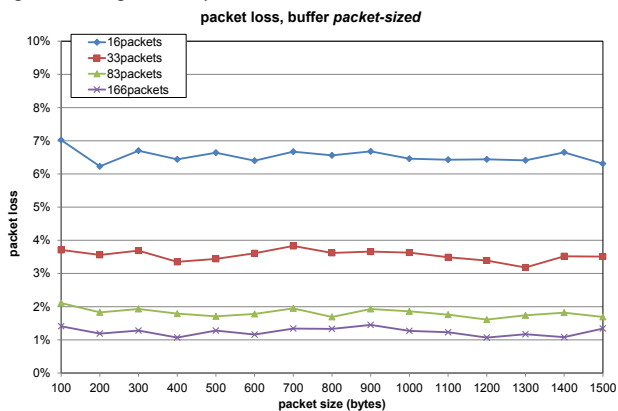


Figure 2. Sizogram for *packet-sized* buffer of 16, 33, 83 and 166 packets

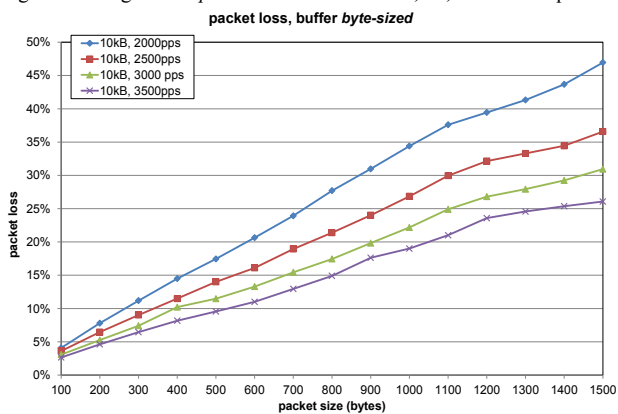


Figure 3. Sizogram for *byte-sized* buffer of 10kB with pps limitations




Figure 4. Sizogram for *byte-sized* buffer of 10kB, with processing time

Fig. 2 has been obtained for a *packet-sized* buffer which size ranges from 16 to 166 packets. The *sizogram* shows that packet loss is independent of packet size, which is logical as the buffer only considers whole packets despite their sizes. In case of having this behaviour, we should use big packets when possible by merging smaller ones or increasing the number of samples or video frames they include.

Fig. 3 shows the *sizogram* for a *10 kB* router with a limitation in terms of pps, when the number of packets per second ranges from 2,000 to 3,500 (note that the scale of the figure is different). It can be seen that this parameter seriously increases packet loss, even if the pps limitation of the router is significantly higher than the pps of the offered traffic (1,970 pps). Fig. 4 shows the same test when different packet processing times are added. In this case, the increase of packet loss is not severe. Nevertheless, the shape of the curves is still the same.

As future work, we plan to perform this test in commercial routers, and also using different Linux *tc* (Traffic Control) queuing disciplines, which implementation details are publicly available. Other background traffic models will also be tested.

ACKNOWLEDGMENT

This work has been partially financed by CPUFLIPI Project (MICINN TIN2010-17298), MBACToIP Project, of Aragon I+D Agency and Ibercaja Obra Social, Project of Catedra Telefonica, Univ. of Zaragoza.

REFERENCES

- [1] W. Feng, F. Chang, W. Feng, J. Walpole: "A Traffic Characterization of Popular On-Line Games". IEEE/ACM Trans. Netw., pp. 488-500, 2005.
- [2] D. X. Wei, P. Cao, S. H. Low, "Packet Loss Burstiness: Measurements and Implications for Distributed Applications," Parallel and Distributed Processing Symposium, International, p. 222, 2007.
- [3] J. Korhonen, Y. Wang, "Effect of packet size on loss rate and delay in wireless links," Wireless Communications and Networking Conference, 2005 IEEE , vol.3, no., pp. 1608- 1613 Vol. 3, 13-17 March 2005.
- [4] G. Dan, V. Fodor, G. Karlsson, "On the effects of the packet size distribution on FEC performance," Computer Networks, Volume 50, Issue 8, Selected Papers from the 3rd International Workshop on QoS in Multiservice IP Networks (QoS-IP 2005), pp 1104-1129, June 2006.
- [5] J. Sommers, P. Barford, A. Greenberg, W. Willinger, "An SLA perspective on the router buffer sizing problem," SIGMETRICS Perform. Eval. Rev. 35, 4, pp 40-51, March 2008.
- [6] J. Yu, I. Al-Ajarmeh. "Call Admission Control and Traffic Engineering of VoIP", In Proc. Second International Conference on Digital Telecommunications, IEEE ICDDT 2007.
- [7] J. M. Boyce, R. D. Gaglianella, "Packet loss effects on MPEG video sent over the public Internet," in Proc Sixth ACM MULTIMEDIA '98. ACM, New York, NY, USA, 181-190, 1998.
- [8] M. S. Borella, D. Swider, S. Uludag, G. B. Brewster, "Internet Packet Loss: Measurement and Implications for End-to-End QoS," Parallel Processing Workshops, International Conference on, p. 3, International Conference on Parallel Processing Workshops (ICPPW'98), 1998.
- [9] J. Saldana, J. Murillo, J. Fernández-Navajas, J. Ruiz-Mas, E. Viruete, J. I. Aznar, "Evaluation of Multiplexing and Buffer Policies Influence on VoIP Conversation Quality", Proc. CCNC 2011-3rd IEEE Workshop on Digital Entertainment, Networked Virtual Environments and Creative Technology (DENVECT), pp 1147-1151, Las Vegas. Jan 2011.
- [10] A. Vishwanath, V. Sivaraman, and M. Thottan. "Perspectives on router buffer sizing: recent results and open problems," SIGCOMM Comput. Commun. Rev. 39, 2, 34-39, March 2009.
- [11] Cooperative Association for Internet Data Analysis "NASA Ames Internet Exchange Packet Length Distributions".