

A quality of service assessment technique for large-scale management of multimedia flows^{*}

J. L. García-Dorado¹, J. Aracil¹, J. A. Hernández¹, S. Lopez-Buedo¹, J. E. López de Vergara¹, P. Reviriego², G. Huecas³, S. Pavón³, and J. Quemada³

¹ Universidad Autónoma de Madrid, Spain,
Contact email: jl.garcia@uam.es.

² Universidad Carlos III de Madrid, Spain.

³ Universidad Politécnica de Madrid, Spain.

Abstract. This paper presents the concept and preliminary experiments of a system for assessing on the Quality of Service of multimedia flows. The goal is to devise a mechanism that allows a service provider to take action whenever poor quality of service is detected in the delivery of multimedia flows. Such procedure is fully automatic since it is based on a goodness-of-fit test between source and destination packet inter-arrival histograms. If the null hypothesis of the test is accepted the flow is marked as in good standing, otherwise it is marked as anomalous and the network management system should take action in response. The proposed technique is analyzed in terms of hardware complexity and bandwidth consumption. The results show this technique is feasible and easily deployable at a minimum hardware and bandwidth expense.

1 Introduction and problem statement

Multimedia networks are very challenging to operate and manage. Indeed, the assurance of quality of service implies the surveillance of a very large number of flows, possibly thousands. This is due to the tremendous growth of multimedia applications over the past few years, which is expected to remain steady in the near future.

With such a large number of flows, it is essential to implement *automatic procedures for quality of service management* that actually save operators to take action manually. Ideally, flows would be marked as either “pass” or “fail” and such binary decision would trigger the appropriate corrective mechanisms automatically. In this paper, we provide a technique that tackles the former problem, since the latter is out of the scope of this paper.

Besides overall latency, quality in multimedia applications is primarily measured with two different metrics: data loss and jitter [1]. In our proposal, we focus on obtaining a single metric that actually captures *how similar incoming and outgoing multimedia flows are*. In other words, we wish to evaluate the amount

^{*} The authors would like to acknowledge the support of the Comunidad Autónoma de Madrid to this work, under project *e-Magerit (S-0505/TIC/000251)*.

of distortion suffered by a multimedia flow as it traverses a given path across the network. To do so, we propose to perform a goodness-of-fit test [2] between the interarrival time probability distributions of the multimedia flow at the ingress and egress points of the path under study. Note that such procedure allows us to merge the above-mentioned two metrics (i. e. data loss and jitter) into a single one, since it is clear that the packet interarrival time series will be distorted by data loss and jitter.

The advantages of the proposed technique are manifold: On one hand, there is no need for time synchronization between the monitored endpoints. The measurement devices at the path entry and exit points are only required to calculate histograms, yet remaining low system complexity, as shown in the next section. In addition, this technique is non-intrusive and consumes nearly no bandwidth. Unlike RTCP [3], there is no need to periodically report to the transmitting end about the observed jitter, which results in bandwidth consumption and it is difficult to implement in channels with limited outbound bandwidth from the client. Finally, our method is mathematically rigorous, since it is based on well-known and reliable statistical techniques.

2 System architecture

The system architecture is depicted in fig. 1. Let us consider a given ingress-egress pair of nodes. The ingress measurement node builds a histogram with the interarrival times of a given multimedia flow, and so does the egress measurement node. Each node transmits its histogram to a centralized network management node, which performs a goodness-of-fit test between them all. Such node thus marks flows as “pass” or “fail” for a given significance level.

Clearly, this approach paves the way for autonomic network management, where, as stated in [4], a multimedia system can self-optimize the traffic of the flows sent to users to achieve a given objective QoE (Quality of Experience) based on feedback information about that traffic. This capability is very important for Video-on-Demand (VoD) scenarios, where users pay for receiving multimedia contents.

Moreover, this information can also be used to perform admission control and let new users start new multimedia sessions. If the current set of flows over the network passes a given QoS test, then a new flow can be admitted. Otherwise, no new sessions will be allowed in order to avoid higher QoS degradation. Furthermore, it is possible to use this information for capacity planning of networks with such multimedia constraints.

Concerning the physical implementation of the proposed solution, histogram calculation is a low-complexity task that can be easily handled by any cheap FPGA-embedded microprocessor, such as MicroBlaze [5]. Even for small FPGAs, this processor occupies few resources of the device. For example, MicroBlaze uses less than 15% of XC3S1200E [6], and costs less than US\$10. For instance, if $\sqrt{N_s}$ bins are used, where N_s is the number of samples, the amount of memory required to store one such histogram would be $\sqrt{N_s} \cdot \log_2(N_s)$ bits. For a typical




Fig. 1. Architecture and testbed scenario.

10 minute session, using 500-byte packets and a high quality 1Mbps multimedia stream, the average histogram size would be 6.5 Kb. Current low-cost FPGAs have internal memories typically two orders of magnitude bigger, and they also feature glueless interface to conventional DRAMs. For example, a 32 MB DDR memory like MT46V16M16 [7] is currently priced under US\$5, thus making completely viable a FPGA-based low-cost solution capable of storing tens of thousands of such histograms.

3 Results and conclusions

The authors have used the ISABEL application as testbed for the performance evaluation of our proposed technique. ISABEL [8] is an advanced collaborative work application that supports video, audio and data to create collaborative sessions adapted to the users' needs.

In our performance evaluation, we set ISABEL to collect packet interarrival times at all intermediate nodes in the network. This way permits us to perform a quality-of-service test on a per-hop basis. With N_s measurements, we perform the standard χ^2 goodness-of-fit test, with $\sqrt{N_s}$ histogram bins. It is worth pointing out that the number of bins can be modified in order to focus on a particular range of interarrival times, or as a way to demand a higher fidelity between the matching histograms.

Specifically we consider the traffic generated by the source “cebolla” node (see fig. 1). Fig. 2 shows the cumulative distribution function (CDF) for the interarrival times between packets sent by the “cebolla” node to the “apollo” node, which are then forwarded by “apollo” to “lamata”, and finally transmitted from “lamata” to “toledo” node.

Table 1 analyzes the QoS degradation as the number of hops increases, for a particular flow under study. It turns out that “apollo” and “lamata” pass the χ^2 test for low values of α but the “toledo” node does not. Thus, we can

IV




Fig. 2. CDF of packet interarrival time (in milliseconds) in all the nodes.

Node	Significance level (α)	Null hypothesis (H_0)
apollo	0.01	✓
lamata	0.02	✓
toledo	0.05	X

Table 1. χ^2 test between the interarrival packet times.

conclude that “toledo” node is suffering anomalies and action should be taken if it repeatedly shows such behavior.

References

1. Claypool, M., Riedl, J.: The effects of high-speed networks on multimedia jitter. In: SCS Euromedia Conference. (1999)
2. Montgomery, D.C., Runger, G.C.: Applied Statistics and Probability for Engineers. John Wiley and Sons (1999)
3. Schulzrinne, H., Casner, S., Frederick, R., Jacobson, V.: RTP: A transport protocol for real-time applications. Internet proposed standard RFC 1889 (January 1996)
4. López, D., López de Vergara, J.E., González, F., Sánchez-Macián, A.: An OWL-S based architecture for self-optimizing multimedia over ip services. In: Proc. 1st IEEE Int. Workshop on Modelling Autonomic Communications Environments. (2006)
5. Xilinx Inc.: Microblaze Product Brief (2005)
6. Xilinx Inc.: Spartan-3E FPGA Family Data Sheet (2007)
7. Micron Technology Inc.: 256 Mb DDR SDRAM Datasheet (2006)
8. Quemada, J., et. al.: Isabel: an application for real time collaboration with a flexible floor control. In: Proc. Int. Conference on Collaborative Computing: Networking, Applications and Worksharing. (2005)