Evaluation of a Quality of Service Routing Strategy for the Differentiated Services Framework

Marília Oliveira, João Brito, Gonçalo Quadros, Edmundo Monteiro
Laboratory of Communication and Telematics <http://lct.dei.uc.pt>
Center for Informatics and Systems of the University of Coimbra
Pólo II, Pinhal de Marrocos, 3030 Coimbra, Portugal

Abstract

In this paper we present the evaluation of the University of Coimbra Quality of Service routing (UC-QoSR) strategy for the Differentiated Services framework. This strategy enables routers to dynamically select paths suitable for each traffic class considered and to perform the respective forwarding. UC-QoSR was implemented as an extension to the OSPF routing protocol included on the GateD routing platform, running on the FreeBSD operating system. We present the results concerning the evaluation of the processing overhead due to the more frequent application of the path computation algorithm and to an increased complexity of the forwarding process. The determination of the factors that influence convergence time is also subject of this analysis. The results obtained for UC-QoSR are compared to the results for OSPF. We show the effectiveness of the mechanisms used in UC-QoSR to make QoS routing scalable in terms of processing cost and responsive as to convergence time.

Keywords: Quality of Service routing, Differentiated Services.

1. Introduction

The majority of Quality of Service routing proposals use source routing and on-demand path computation. These proposals are mainly intended for connection-oriented networks, where there is resource reservation, as in the Integrated Services architecture [1]. We developed a Quality of Service routing strategy that has the ability to embed a Differentiated Services network [2] with QoS routing capabilities. This strategy extends the hop-by-hop routing model, with path pre-computation of the actual Internet, to provide for the dynamic computation of paths suitable for different traffic classes and for the corresponding forwarding mechanisms. The main goals behind the conception of this strategy were simplicity and to use actual proven protocols, widely used and keeping modifications as small as possible.

It is well known that Quality of Service routing introduces a non-negligible overhead in the network [3]. This overhead is twofold, including communication and processing components. Communication overhead is due to more frequent updates that are necessary in order for the routers to maintain the information about the state of the network Processing overhead is influenced by the complexity of the path computation algorithm, the frequency of its application and complexity of the forwarding process.

In this paper we present the evaluation of the processing overhead introduced by UC-QoSR when compared to OSPF. The paper is organised as follows: in Section 2 we describe the UC-QoSR strategy; the experimental evaluation is described in Section 3; the major conclusions and some guidelines for future work are presented in Section 4.

2. Description of UC-QoSR

At the Laboratory of Communication and Telematics of the University of Coimbra (LCT-UC) we are developing a new service model for the Internet, where traffic differentiation is done per class, following the approach of the Differentiated Services. The differentiation of classes is achieved at each router by a scheduler and dropper that implement the Per Hop Behavior Dynamic Degradation Distribution (PHB D3) [4], conceived at LCT. In the service model proposed, we consider four traffic classes. Traffic characterization is based on delay and loss sensitivity of traffic generated by multimedia applications [5], as in Table 1.

Table 1- Classes of traffic of the LCT service model.

Class	Delay sensitivity	Loss sensitivity	Application
1	Medium	High	Video Training
2	High	Low	Video real -time
3	High	Medium	Internet telephony
4	Low	Low	*Best-effort*

We use the QoS metric developed at LCT-UC [6] to measure the impact of the degradation of QoS characteristics in application performance. This metric is composed of a delay congestion index (*DcI*) and a loss congestion index (*LcI*). These indexes represent the impact that delay and loss at the router will have on application performance degradation. These congestion indexes represent the state of each interface of the routers in the domain.

The objectives of UC-QoSR are to distribute the congestion indexes of the interfaces of each router to other routers in the network and, based on this information, to select paths adequate to each class of service. These goals are achieved by the inclusion of QoS mechanisms in OSPF. The congestion indexes are distributed in OSPF update messages, and paths are computed using OSPF path selection algorithm, that is, the *Dijkstra* algorithm. For each class, the cost of each link in the network is obtained by the combination of the congestion indexes of the corresponding router interface weighted according to the sensitivity of the class to delay and loss. Due to its nature, the indexes represent comparable measures, and there is no loss of information from aggregation of different kinds of units. The weights allow for different costs in each class, according to Equation 1.

$$Cost_i = WL_i * LcI + WD_i * DcI \quad (1)$$

Here, $Cost_i$ is the cost of class i, WL_i and WD_i are the weights of the loss and delay congestion indexes for class i, respectively.

The *Dijkstra* algorithm is used five times, once to obtain the paths that minimize the cost for each class and also to obtain the path that minimizes OSPF original metric. With this approach, it is possible for routers running UC-QoSR to forward packets that are not compliant with the codepoints of the LCT-UC service model.

We have implemented the UC-QoSR strategy as an extension to Open Shortest Path First (OSPF) [7], using the GateD[1] platform, running on the FreeBSD operating system. Some of the extensions included in OSPF were inspired in the QoS routing mechanisms described in [8].

3. Performance evaluation of the UC-QoSR strategy

QoS routing mechanisms introduce communication and processing overhead in the network that must be evaluated.

The communication cost inherent to UC-QoSR was evaluated on the prototype developed. The results showed that it is feasible to deploy this strategy because the mechanisms to limit the communication overhead are effective controlling the overhead introduced [9]. The idea behind the control of routing traffic is the definition of a compromise between the actuality of the information maintained at each router and the communication burden due to an increased number of update messages.

Processing overhead is affected by a number of factors including the path selection criteria and the trigger for path selection computation. In UC-QoSR, we minimise complexity because, instead of using algorithms that would take into consideration both congestion indexes, we combine them into one single metric and use the Dijkstra algorithm to select paths. The processing overhead introduced is due to the more frequent computation of paths and to the fact that each time paths are computed, the algorithm is applied five times, one for each traffic class and one for compatibility with OSPF. The type of trigger for path computation (periodic, on request, after receiving each routing update) can, in certain circumstances, reduce processing overhead. However, in UC-QoSR, paths are computed after receiving each routing update. This fact does not affect negatively the processing performance, because

[1] <http://www.gated.org>

the thresholds used for triggering updates control communication overhead. When they are issued it is necessary to update the state of the routers and thus paths must me recomputed.

3.1. Experiments conditions

In this section we present the analysis of the tests concerning the evaluation of the processing cost associated with UC-QoSR. This evaluation was conducted on a testbed with six highly interconnected routers running the UC-QoSR prototype in order to maximise the number of alternate paths (Figure 1). Each router is a Celeron at 466 MHz, with 64 Mbyte of RAM, and running the FreeBSD operating system.

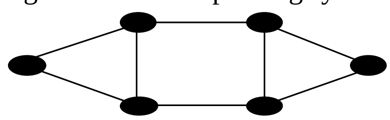

Figure 1- Experimental testbed.

Each test was repeated 20 times. The values presented represent the average of the values measured. Each test for the overall GateD performance evaluation had the duration of 10 minutes and was done for situations of smooth and bursty loads. The performance measures evaluated are described bellow.

The evaluation of the processing overhead introduced by UC-QoSR is determined by comparing the processing cost (number of clock cycles) associated with UC-QoSR and OSPF, running respectively on a modified FreeBSD kernel and on an original kernel.

The evaluation of local convergence time includes GateD notification of the event, the time spent in path computation and the installation of the new paths in the kernel. This evaluation is done after metrics variation and after link failure.

Forwarding cost corresponds to the processing time of each packet at the IP layer. This is measured in terms of forward processing time and routing lookup time. Forwarding processing time with a cache hit corresponds to the time of processing of a packet at the IP layer if the next-hop for its destination and class is on the forwarding cache. Forwarding processing time with a cache miss corresponds to the time of processing of a packet at the IP layer if the next-hop for its destination and class is not on the forwarding cache.

Routing table lookup time is the time needed to search the routing table for the adequate next-hop for the packet being processed.

Besides the evaluation of the parameters described above, we also conducted an overall GateD performance evaluation, comparing original and UC-QoSR systems.

3.2. Presentation of results

The results concerning the evaluation of the processing overhead are presented in Table 2, both for the original version of GateD and for the UC-QoSR strategy. It can be seen that for the test conditions described, namely for the topology used in the testbed, the processing overhead of UC-QoSR is 84%. This is an interesting result, since in UC-QoSR the shortest path calculation is applied five times in each execution of the path computation algorithm. This observation stems from the fact that the sequential execution of the same computation takes advantage of the cache.

Table 2- Processing overhead evaluation.

	Original GateD	UC-QoSR GateD
Average processing time (μs)	162	298

The evaluation of local convergence time after metrics variation is depicted in Table 3 and it is well bellow one second.

Table 3- Local convergence time after metrics variation.

	UC-QoSR GateD
Average convergence time (μs)	587995

The evaluation of local convergence times after link failure, with link going down and going up are shown in Table 4. It can be seen that UC-QoSR takes longer to react to these changes, but not exceeding the double of original GateD.

Table 4- Local convergence time after link failure.

	Original GateD (clock cycles)	UC-QoSR (clock cycles)
Link going down	1915	3916
Link going up	2005	2566

The evaluation of forwarding cost is depicted in Figure 2, when there is a cache hit or miss. The results show the relevance of using a forwarding cache, since a cache miss can correspond to an increase in as much as 200% of the forwarding time.

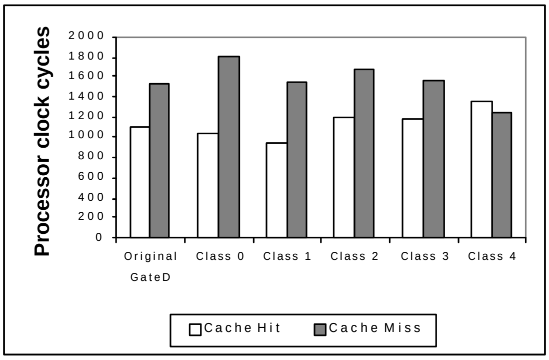

Figure 2- Forwarding processing time evaluation.

Routing table lookup times are presented in Table 5. These results show that the search for a next-hop in the UC-QoSR strategy is more expensive than the same search on the base kernel.

Table 5- Routing table lookup time.

	Average (Clock cycles)
Original GateD and kernel	372
UC-QoSR model - Class 0	586
UC-QoSR model - Class 1	500
UC-QoSR model - Class 2	496
UC-QoSR model - Class 3	625
UC-QoSR model - Class 4	566

The overall GateD performance evaluation is presented in Figure 3. During these experiments, the original GateD did 2 path calculations and Uc-QoSR did 300. The processing time observed in both cases does not significantly reflect this difference. This allows for the conclusion that UC-QoSR does not impose too much processing effort on the routers.

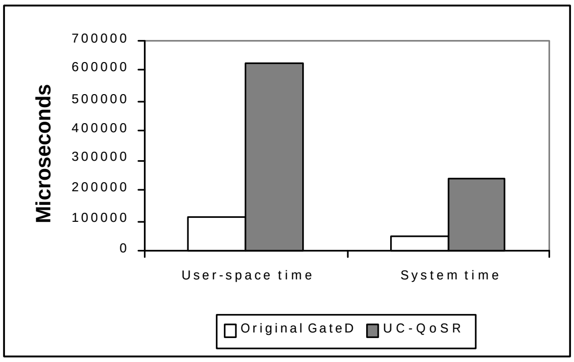

Figure 3- Overall GateD performance evaluation.

4. Conclusions and Future Work

We presented the evaluation of the UC-QoSR strategy on a prototype developed on the GateD platform, running on the FreeBSD operating system. This strategy was evaluated upon processing and forwarding cost, local convergence time, and globally. The results showed that the mechanisms introduced to compute QoS aware paths do not put an excessive burden on the routers, when compared to the characteristics of OSPF.

In future work we will perform the impact of the UC-QoSR strategy on traffic performance.

Acknowledgments

This work was partially supported by the Portuguese Ministry of Science and Technology, under program PRAXIS XXI (Projects QoS II, IPQoS, PhD grant BD/13723/97) and by POSI - Programa Operacional Sociedade de Informação of Portuguese Fundação para a Ciência e Tecnologia and European Union FEDER

References

[1] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, IETF, June 1994.

[2] S. Blake, D. Black, M. Carlson, E. Davies Nortel, W. Weiss, "An Architecture for Differentiated Services", IETF, RFC 2475, December 1998.

[3] G. Apostolopoulos, R. Guérin, S. Kamat, S. Tripathi, "Quality of service Based Routing: A Performance Perspective", ACM SIGCOMM'98, Aug 31-Sept 4, 1998, Vancouver, BC, Canada.

[4] G. Quadros, A. Alves, E. Monteiro, F. Boavida, "An Approach to Suport Traffic Classes in IP Networks", QoFIS'2000, Berlin, Germany, Sept 25-26, 2000.

[5] J. Gollick, "Catching the Multimedia Wave", Data Communications, March 1999.

[6] G. Quadros, E. Monteiro, F. Boavida, "A QoS Metric for Packet Networks", SPIE's symposium on QoS Issues Related to Internet, Boston, MA, USA, Nov 2-5, 1998.

[7] J. Moy, "OSPF Version 2", IETF, RFC 2328, April 1998.

[8] R. Guérin, S. Kamat, A. Orda, T. Przygienda, D. Williams, "QoS Routing Mechanisms and OSPF Extensions", IETF, RFC 2676, Aug 1999.

[9] M. Oliveira, J. Brito, B. Melo, G. Quadros, E. Monteiro, "Quality of Service Routing in the Differentiated Services Framework", SPIE's Symposium on QoS and Future Directions, Boston, Massachusetts, USA, Nov 5-8, 2000.