Quality of Service Routing in the Differentiated Services Framework

Marília Oliveira, Bruno Melo, Gonçalo Quadros, Edmundo Monteiro

marilia@dei.uc.pt, bmelo@student.dei.uc.pt, {quadros, edmundo}@dei.uc.pt

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

ABSTRACT

In this paper we present a quality of service routing strategy for networks where traffic differentiation follows the class-based paradigm, as in the Differentiated Services framework. This routing strategy is based on a metric of quality of service. This metric represents the impact that delay and losses verified at each router in the network have in application performance. Based on this metric, it is selected a path for each class according to the class sensitivity to delay and losses. The distribution of the metric is triggered by a relative criterion with two thresholds, and the values advertised are the moving average of the last values measured.

We present the results of tests concerning the evaluation of the criterion used, and show that it has higher sensitivity than a criterion with one threshold. We also evaluate the impact that the time scale used for computing the moving average has in the number of advertisements issued. Parameter tuning and evaluation of message overhead introduced by a more frequent distribution of routing information are also discussed.

Keywords: QoS routing, Differentiated Services

1. INTRODUCTION

In a communication system that aims at providing different levels of quality of service (QoS), it is advisable to use routing protocols that make decisions based simultaneously on the needs of traffic and the state of routers in the network, that is, QoS capable routing protocols. There are several research projects that aim at providing QoS routing capabilities to integrated services networks (networks where there is path establishment), performing routing at the flow level. For an overview and further references see [1].

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 model [2]. We use the QoS metric developed at LCT-UC [3] to measure the impact of the degradation of QoS characteristics in application performance.

The objectives of our routing proposal, University of Coimbra - Quality of Service Routing (UC-QoSR), are to distribute the metric of each router to other routers in the network and, based on this information, to select paths adequate to each class of service. Our routing proposal is based on the QoS routing mechanisms presented in [4]. We developed our strategy as an extension to the OSPF routing protocol [5] and use similar mechanisms to propagate the QoS metric. However, there are some differences, namely the type of metric used, the path selection algorithm and the identification of traffic characteristics.

The main characteristics of the UC-QoSR proposal are:

- Unicast QoS routing protocol.
- QoS routing for a network where traffic differentiation is done per class.

- Path selection according to a composed metric, representative of delay and loss.
- Load balancing capability.
- Compatibility with OSPF.

It is well known that the deployment of QoS routing introduces a non-negligible overhead in the network. The two major problems that must be accessed are an increased communication cost due to a higher number of routing message updates and processing cost due to the more frequent application of the path selection algorithm. In this paper we will evaluate de communication cost associated with UC-QoSR.

This paper is organized as follows: in Section 2 we present the framework of our routing proposal; in Section 3 we describe the main characteristics of UC-QoSR; Section 4 contains the description and analysis of the most significant evaluation results concerning the distribution of the QoS metric; the major conclusions and some guidelines for future work are presented in Section 5.

2. FRAMEWORK

In the LCT-UC service model we consider four traffic classes (Table 1). Traffic characterization is based on delay and loss sensitivity of traffic generated by multimedia applications [6].

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

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

The LCT-UC QoS metric comprises two congestion indexes, one relative to delay (*IcD*) and the other relative to loss (*IcL*). These indexes represent the impact that delay and losses experimented in each network element (e.g. router) have in the performance of applications. This approach to measuring quality of service yields values for the congestion indexes that are comparable, which is not possible with absolute measures of delay and loss. This characteristic is nuclear to our routing strategy, since allows the combination of the values of both congestion indexes in a cost function that represents the state of each link in the network.

3. ROUTING STRATEGY

The main objective of the UC-QoSR strategy is to select paths suitable to each class of traffic, based on the information provided by the congestion indexes. This section describes the strategy used for distribution of routing information, and the issues related to the path selection mechanism.

3.1 Routing information distribution

In the UC-QoSR routing strategy, the state of the interfaces of each router is represented by the delay and loss congestion indexes that constitute the QoS metric used. Information is distributed as an extension to the routing protocol OSPF, following an approach similar to [4]. We use the fields of OSPF routing messages that were devoted for type-of service routing. With this approach there is not the need for introducing new types of messages, keeping our strategy very similar to OSPF and thus allowing inter-operation among routers running both protocols.

The congestion indexes are monitored in a time scale of ms. In order to avoid routing oscillations, the values advertised are moving averages of the values registered, with a window size that is configurable. Instability could also be avoided if we used a periodic criterion. However, in this case the triggering of advertisements would be independent of the state of the router, and thus would not allow for the dynamic adjustment of the protocol.

The emission of advertisements is triggered by a relative criterion consisting of two thresholds (*T1*, *T2*) and a transition point (TP). With this kind of trigger, a new advertisement is issued only when the actual value of the congestion index is significantly different from the last value announced. The threshold used determines the level of significance; if the value being evaluated is below the transition point, *T1* is applied, otherwise *T2* is used. We define two thresholds to avoid the loss of sensitivity when the congestion indexes have higher values since, in these situations, the same absolute variation may not trigger an advertisement as would happen with smaller index values.

3.2 Path selection mechanism

Each router in the network advertises the delay and loss congestion indexes of its interfaces. The selection of a path that minimizes both indexes is a NP-complete problem [8]. In our purpose of keeping the UC-QoSR approach simple, we adopted the path selection algorithm used in OSPF, 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:

$$Cost_i = WL_i * IcL + WD_i * IcD \qquad (1)$$

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. Since costs are calculated locally, it is possible to attribute weights according to local policies regarding class treatment.

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.

According to Equation 1, a link with an elevated loss congestion index will have a cost that is higher for loss sensitive classes. Thus, loss sensitive traffic will avoid paths with high loss congestion indexes, and delay sensitive traffic will avoid paths with high delay congestion indexes. We pretend to deviate best-effort traffic from the path with lower delay congestion indexes in order to achieve an implicit reservation for classes of traffic that are more sensitive to delay. To achieve this goal, the weight of the loss congestion index for class 4 must be higher than the weight of the delay congestion index. Thus, this class will prefer a low loss path instead of a low delay path, leaving the low delay path free to other types of traffic, specifically, delay sensitive traffic.

4. EXPERIMENTATION

We have implemented the UC-QoSR strategy as an extension to OSPF, using the GateD platform [7]. The testbed used for evaluation of the metric distribution process is composed of three serially interconnected routers (R1, R2, and R3), as it is depicted in Figure 1. Each router is a Celeron at 450 MHz, with 64 Mbyte of RAM, and running the FreeBSD operating system.

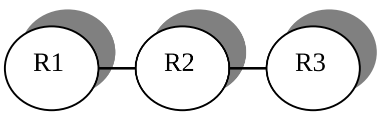

Figure 1- Testbed.

The main objectives of the tests are:

(1) Comparison of criteria sensitivity, with one and two thresholds.
(2) Tuning of parameters, namely the values of thresholds, and the time scale used for the calculation of values to be presented to the distribution criterion (moving average window).
(3) Evaluation of the overhead introduced by more frequent updates.

The values of the congestion indexes were generated according to a uniform distribution, whose limits represent the type of load in the network.

The conditions of the tests are presented in Table 2. Each test had the duration of 5 minutes.

We evaluated the number of Router Link State Advertisements (R-LSA), Link State Updates (LSU) and Acknowledgements (Ack).

Criteria	1 threshold	Values of 5, 10, 15, …, 100%
	2 thresholds	*T2* = ½ * *T1*; *TP* = 25, 50, and 75
Type of load	Low	*IcL* = 0; *IcD* = U(0,25)
	Medium	*IcL* = U(0,50); *IcD* = U(25,75)
	High	*IcL* = U(50,100); *IcD* = U(75,100)
	Random	*IcL* = U(0,100); *IcD* = U(0,100)
Time-scale of advertised values	Instantaneous	1 second (wind = 1)
	Averaged	Window size of 5, 10, 15, and 20

Table 2- Conditions of experimentation.

Figure 2 shows the number of R-LSA issued per minute, when it is used the criteria with one and two thresholds, for a random load. The criterion with two thresholds has higher sensitivity than the criterion with one threshold. The study of the influence of the value of the transition point showed that:

(1) When the load is low, it is always applied *T1*, for every TP considered.
(2) When the load is high, *T2* is applied most of the times, and the value of the TP has not significant impact in the number of R-LSA.
(3) For random and medium loads the criterion sensitivity is higher when TP is smaller, since thus *T2* is more frequently applied; however, the number of R-LSA issued is equivalent for transition points of 25 and 50.

The analysis done above shows that a TP of 50 is adequate for all load types.

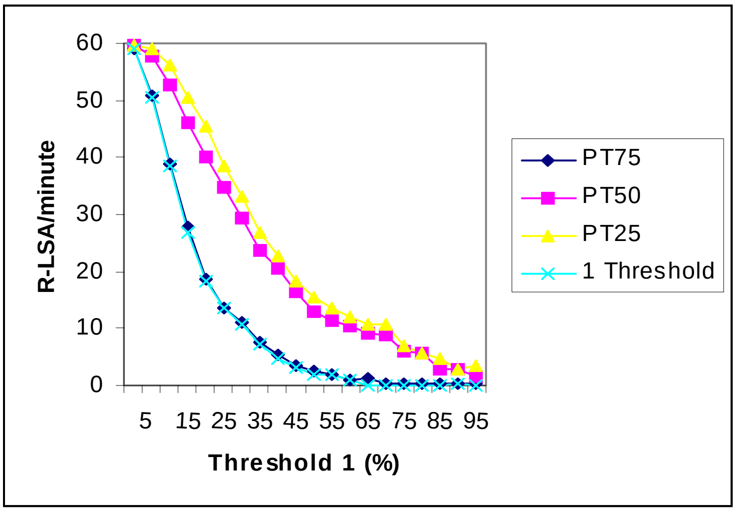

Figure 2- Criteria comparison for random load.

Figure 3 shows the number of R-LSA issued per minute when it is used the criterion with one threshold for high load, and the criterion with two thresholds for all loads. Even though this criterion shows an increased sensitivity, the number of R-LSA is smaller for high loads.

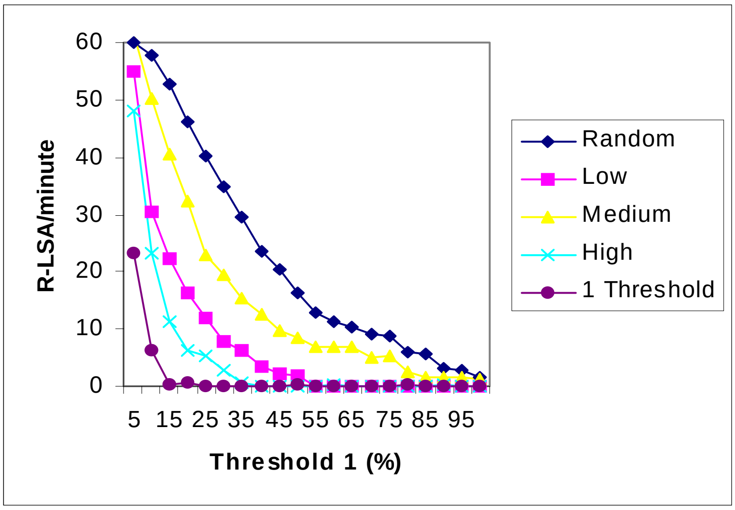

Figure 3- Load influence in message overhead.

We were able to equalize the number of R-LSA for high and low loads using smaller values for *T2* (*T2*=1/4**T1*), however at the expense of a higher number of R-LSA for medium and random loads. The results of the tests done are shown in Figure 4. In can be seen that this would introduce excessive communication overhead, so we decided to keep the threshold relation presented in Table 2.

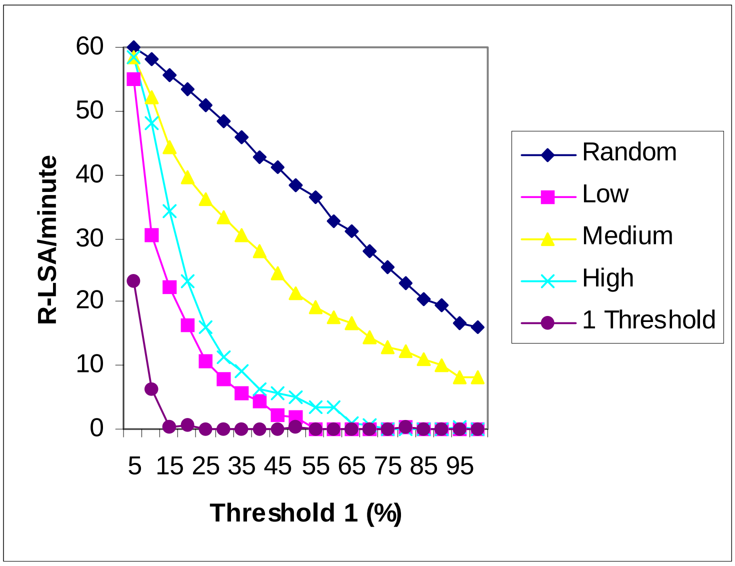

Figure 4- Load influence in message overhead (*T2*=1/4**T1*).

Figure 5 shows the relevance of the time scale used for the distribution of the metric, for medium loads. When the criterion is applied to the instantaneous values of the congestion indexes (wind = 1), the number of R-LSA is elevated, even for high thresholds. The utilization of a value that is the result of the moving average over a number of congestion index values allows for the reduction of routing overhead, specially for larger windows. For instance, using a window size of 20 can reduce routing overhead in about 40% relatively to a window size of 10.

In Table 3 are shown the threshold ranges and the respective periods of R-LSA emission for a window of size 20. These results show that *T1* should be in the range of 25-30% in order to control routing messages overhead. The period of emission of R-LSA for a *T1* of 30% is shown in parenthesis.

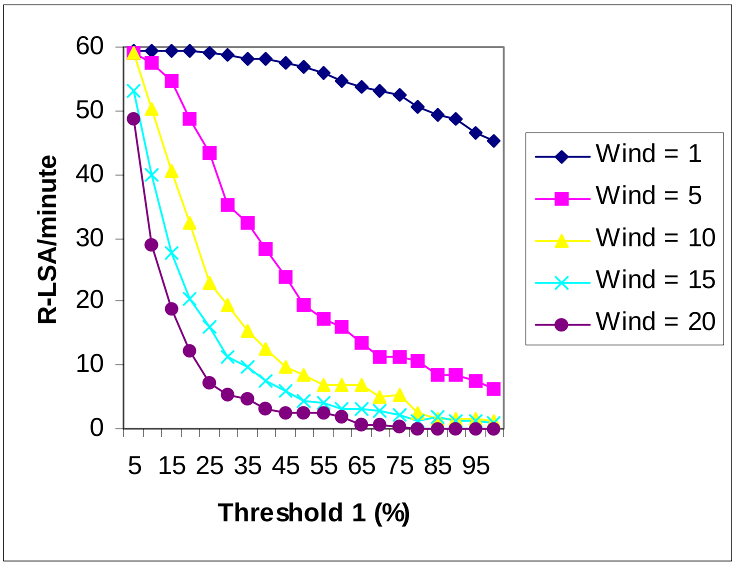

Figure 5-Time scale comparison for medium load.

	Random load	Low load	Medium load	High load
Threshold range	20-70% (30%)	20-35% (30%)	25-70% (30%)	15-30% (30%)
Period of R-LSA emission range	5-50s (10s)	12-60s (20s)	9-60s (12s)	15-60s (60s)

Table 3- Reference parameters.

We did some compatibility tests, with the middle router (R2) running OSPF and routers R1 and R3 running UC-QoSR. It was observed that both end routers retained the R-LSA corresponding to all routers of the testbed. We also noted that R2 had in its Link State Database the R-LSAs from R1 and R3. However, R2 only used the TOS 0 metric in the path selection algorithm, that is the metric that corresponds to original OSPF.

Since the flooding mechanism used in OSPF is a reliable mechanism it is important to evaluate the impact of confirmation messages [9]. The percentage of total bandwidth consumption due to routing traffic that is relative to confirmations is shown in Figure 6. It can be seen that the percentage of total bandwidth that is consumed by confirmations increases with thresholds. This is due to the fact that when there is a very frequent emission of advertisements, confirmation messages are grouped in the same acknowledgement packet. When advertisements are received less frequently, each acknowledgement is usually sent in a separate packet, thus consuming a higher bandwidth share of all routing protocol messages.

From the values presented in Table 3 and from Figure 6, we can see that, if it is used a threshold 1 of 30% (with a threshold 2 of 15%), confirmation messages are responsible for 12% of all the bandwidth consumed by routing messages. We believe that this is not a very high cost to pay for reliability and also for keeping changes as small as possible. The other advantage of keeping acknowledgments pertains to the fact that it allows inter-operability with OSPF.

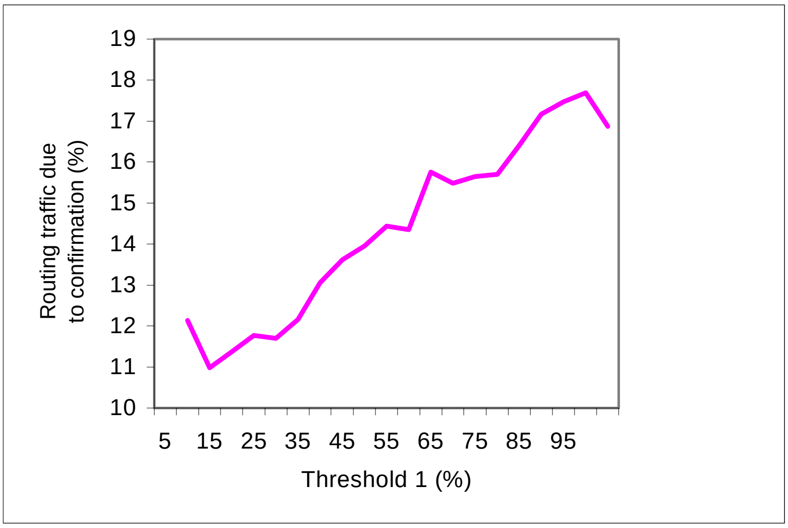

Figure 6- Impact of confirmations on routing traffic.

5. CONCLUSIONS AND FUTURE WORK

We propose a quality of service routing strategy for the LCT-UC service model, where traffic differentiation is done per class and the state of the network is represented by a QoS metric that represents the impact that loss and delay have in application performance. We compared two criteria for triggering updates, and showed that the criterion with two thresholds presents an increased sensitivity. We did tests to evaluate the impact of the time-scale of the values presented to the distribution criteria in the number of R-LSA issued. The results showed that using a window of size 20 in conjunction with a *T1* of 30% and a *T2* of 15% allow for periods between 10 and 60 second for the distribution of R-LSA, depending on the type of load. The evaluation of the percentage of routing bandwidth consumption due to acknowledgements also showed that confirmations can play a significant role in routing traffic bandwidth consumption, specially if high thresholds are used.

Ongoing and future work includes the implementation and test of the path selection algorithm. It is also our objective to adapt the path selection algorithm in order to select paths with maximally disjoint links [10], taking both congestion indexes into consideration. We believe this approach will contribute to a better traffic differentiation in the network.

ACKNOWLEDGEMENTS

This work was partially supported by the Portuguese Ministry of Science and Technology (MCT), under program PRAXIS XXI-PRAXIS/P/EEI/10168/1998 (Project QoS II – Quality of Service in Computer Communication Systems) and under the PhD grant PRAXIS XXI/ BD/13723/97.

REFERENCES

1 S. Chen, K. Nahrstedt, “An Overview of Quality-of-Service Routing for the Next Generation High-Speed Networks: Problems and Solutions”, IEEE Network Magazine, Special Issue on Transmission and Distribution of Digital Video, 1998.

2 S. Blake, D. Black, M. Carlson, E. Davies Nortel, W. Weiss, An Architecture for Differentiated Services, Internet Engineering Task Force, Request for Comments 2475, December 1998.

3 G. Quadros, E. Monteiro, F. Boavida, “A QoS Metric for Packet Networks”, Proceedings of SPIE's symposium on Voice, Video, and Data Communications conference on Quality of Service Issues Related to Internet, Boston, MA, USA, November 2-5, 1998.

4 R. Guérin, S. Kamat, A. Orda, T. Przygienda, D. Williams, QoS Routing Mechanisms and OSPF Extensions, Internet Engineering Task Force, Request for Comments 2676, August 1999.

5 J. Moy, OSPF Version 2, Internet Engineering Task Force, Request For Comments 2328, April 1998.

6 J. Gollick, “Catching the Multimedia Wave”, Data Communications, March 1999.

7 Merit GateD Consortium, <http://www.gated.org>.

8 Z. Wang , J. Crowcroft, “Quality of Service Routing for Supporting Multimedia Applications”, IEEEJSAC, September 1996.

9 G. Apostolopoulos, R. Guérin, S. Kamat, and S. Tripathi, “Quality of service Based Routing: A Performance Perspective”, ACM SIGCOMM’98, Aug 31- Sept 4, 1998, Vancouver, BC, Canada, pp 17-28.

10 N. Taft-Plotkin, B. Bellur, and R. Ogier, “Quality-of-Service Routing Using Maximally Disjoint Paths”, IWQoS’99, London, England, June 1999, pp. 119-128.