Stability and Scalability Issues in Hop-by-Hop Class-Based Routing

Marília Curado, Orlando Reis, João Brito, Gonçalo Quadros, Edmundo Monteiro

marilia@dei.uc.pt, {oreis, jbrito}@student.dei.uc.pt, {quadros, edmundo}@dei.uc.pt
Laboratory of Communications and Telematics
CISUC/DEI
University of Coimbra
Pólo II, Pinhal de Marrocos, 3030-290 Coimbra
Portugal

Abstract. An intra-domain Quality of Service (QoS) routing protocol for the Differentiated Services framework is being developed at the University of Coimbra (UC-QoSR). The main contribution of this paper is the evaluation of the scalability and stability characteristics of the protocol on an experimental test-bed. The control of protocol overhead is achieved through a hybrid approach of metrics quantification and threshold based diffusion of routing messages. The mechanisms to avoid instability are: (i) a class-pinning mechanism to control instability due to frequent path shifts; (ii) the classification of routing messages in the class of highest priority to avoid the loss of accuracy of routing information. The results show that a hop-by-hop, link-state routing protocol, like Open Shortest Path First, can be extended to efficiently support class-based QoS traffic differentiation. The evaluation shows that scalability and stability under high loads and a large number of flows is achieved on the UC-QoSR strategy.

1 Introduction

Quality of Service plays a major role in the deployment of communication system for applications with special traffic requirements, such as video-conferencing or Internet telephony. The need to support these types of traffic has motivated the communication research community to develop new approaches. Some of this work resulted in the Differentiated and Integrated Services architectures proposed by the Internet Engineering Task Force (IETF) [1, 2].

Current routing protocols used in the Internet lack characteristics for QoS provision to support emerging new services. All traffic between two endpoints is forwarded on the same path, even if there are other alternative paths with more interesting properties for the requirements of a specific flow or traffic class. Usually, the shortest path is selected, based on a single static metric, that does not reflect the availability of resources. In these situations, congestion easily occurs on the shortest path, with the corresponding degradation of traffic performance, despite the underutilization of network resources on alternative paths. This scenario has motivated the development of QoS aware routing protocols.

The most significant developments on QoS routing are aimed at communication systems where traffic differentiation is done per flow, as in the Integrated Services [1]. The Differentiated Services framework does not explicitly incorporate QoS routing. It is, thus, essential to develop QoS routing protocols for networks where traffic differentiation is done per class. The Quality of Service Routing strategy of the University of Coimbra (UC-QoSR) was conceived to fulfill this purpose.

The UC-QoSR strategy selects the best path for each traffic class based on information about the congestion state of the network. This strategy extends the Open Shortest Path (OSPF) routing protocol [3] in order to select paths appropriate for all traffic classes as described in [4, 5].

A prototype of UC-QoSR was implemented over the GateD[1] platform, running on the FreeBSD operating system [4]. The description of the mechanisms introduced to allow for scalability and to avoid instability and the evaluation of its robustness are the main objectives of the present paper. The rest of the paper is organized as follows: Section 2 presents some related work; Section 3 describes the UC-QoSR strategy; test conditions and analysis of results are presented in Section 4; the main conclusions and issues to be addressed in future work are presented in Section 5.

2 Related Work

The issues concerning stability of congestion based routing have been addressed by several researchers. This problem becomes even more important in protocols for QoS routing.

The advertisement of quantified metrics, instead of the advertisement of instantaneous values, is a common approach to avoid the instability of dynamic routing protocols. The quantification can be done using a simple average [6], or using hysteresis mechanisms and thresholds [7].

Another methodology to avoid routing oscillations is to use load-balancing techniques, allowing for the utilization of multiple-paths from a source towards the same destination. A simple approach of load balancing is to use alternate paths when congestion rises, as in the algorithm Shortest Path First with Emergency Exits (SPF-EE) [8]. This strategy prevents the excessive congestion of the current path because it deviates traffic to an alternate path when congestion starts to rise, and thus avoids routing oscillations. As an alternative to shortest path algorithms, algorithms that provide multiple paths of unequal cost to the same destination were proposed by Vutukury and Garcia-Luna-Aceves [9]. The algorithm proposed by these authors finds near-optimal multiple paths for the same destination based on a delay metric.

Even tough the proposals described above permit load balancing and avoid routing oscillations, they do not take into consideration the requirements of the different types of traffic. This problem has been addressed by several proposals within the connec-

---

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

tion oriented context. Nahrstedt and Chen conceived a combination of routing and scheduling algorithms to address the coexistence of QoS and best-effort traffic flows. In their approach, traffic with QoS guarantees is deviated from paths congested with best-effort traffic in order to guarantee the QoS requirements of QoS flows and to avoid resource starvation of best-effort flows. Another routing strategy that addresses inter-class resource sharing was proposed by Ma and Steenkiste [10, 11]. Their strategy comprises two algorithms: one to route best-effort traffic and the other to route QoS traffic. The routing decisions are based on a metric that enables dynamic bandwidth sharing between traffic classes, particularly, sending QoS traffic through links that are less-congested with best-effort traffic. Although these proposals achieve active load balancing, they use source routing algorithms that do not reflect the actual hop-by-hop Internet routing paradigm and thus are not able to adapt dynamically to changes in network.

Routing stability can be achieved using some other mechanisms, like route pinning and doing load sensitive routing at the flow level [12]. However these approaches are not suitable for a situation where routing is done hop-by-hop and there is not connection establishment phase.

The development of hop-by-hop QoS routing protocols for communication systems where traffic differentiation is made per class has been the subject of recent studies. Van Mieghem *et al.* evaluated the impact of using an exact QoS routing algorithm in the hop-by-hop routing context [13]. These authors showed that such an algorithm is loop-free, however it may not find the exact solution. To solve this problem, the use of active networking is proposed. An algorithm called Enhanced Bandwidth-inversion Shortest-Path (EBSP) has been proposed for hop-by-hop QoS routing in Differentiated Services networks [14]. This algorithm is based on a widest-shortest path algorithm that takes into account the hop count. The hop count is included in the cost function in order to avoid oscillations due to the increased number of flows sent over the widest-path. Although the algorithm EBSP selects the best path for Premium-class traffic, it does not consider other traffic classes.

QoS routing introduces additional burden in the network, pertaining to the processing overhead due to more complex and frequent computations and the increased routing protocol overhead. The trade-off between the cost of QoS routing and its performance was evaluated in some works [15, 16]. The results included in these references are applicable to systems where there is a flow establishment phase. Furthermore, the above references do not address QoS routing scalability in terms of number of flows and packet sizes, treating only the different rates of arrival of flows and bandwidth requirements.

Despite the relevant QoS issues addressed, the proposals for QoS routing analyzed lack the analysis of the applicability to a class-based framework and are only evaluated theoretically or by simulation. The use of a prototype approach limits the dimension of the test-bed, however it introduces processing and communication systems dynamics, being closest to a real situation.

3 UC-QoSR Strategy

In this section the main characteristics of the routing strategy UC-QoSR are briefly described. A more detailed description can be found in previous publications of the authors [4, 5]. The mechanisms used to allow for scalability and to control instability, which consist the main objectives of the present work, are presented in detail.

3.1 UC-QoSR System Model

The UC-QoSR strategy was designed for intra-domain hop-by-hop QoS routing in networks where traffic differentiation follows the class paradigm. This strategy is composed of three main components, as follows:

a) A QoS metric that represents the availability of resources in the network;
b) Traffic class requirements in terms of QoS parameters;
c) A path computation algorithm to calculate the most suitable path for each traffic class, according to the dynamic state of the network expressed by a QoS metric.

The availability of resources in the network is measured through a QoS metric that represents the congestion state of routers interfaces. This metric consists of two congestion indexes, one relative to packet delay (*DcI*) and other to packet loss (*LcI*). These indexes evaluate the impact that delay and loss at the router, will have on application performance [17]. The congestion indexes are distributed to all routers in the domain through modified OSPF routing messages (Router Link State Advertisements – R-LSA).

The UC-QoSR strategy was conceived for communication systems where traffic characterization is based on class sensitivity to delay and loss. Currently, four classes are considered with different delay and loss sensitivities.

The problem of QoS routing when using two additive or multiplicative metrics, or one additive and one multiplicative metrics is a NP-complete problem [18, 19]. Thus, since the congestion indexes are additive metrics, the selection of a path that minimizes both congestion indexes is a NP-complete problem. However, due to their nature, the indexes represent comparable measures, and can be combined in a single metric without loss of information from aggregation of different kinds of units.

The path computation algorithm uses a cost function that combines both congestion indexes, weighted according to delay (δ) and loss sensitivity (λ) of each traffic class. The cost of link l for class i ($c_{l,i}$) results from the combination of the loss congestion index of link l (l_l) and the delay congestion index of link l (l_l), according to:

$$c_{l,i} = \delta_i d_l + \lambda_i l_l \, . \qquad \textbf{(1)}$$

The merging of the congestion indexes origins a value that represents the congestion state of the interface, as it is perceived by traffic belonging to each class. The Dijkstra algorithm is then used to compute the shortest path tree for each traffic class. The UC-QoSR strategy remains fully compatible with original OSPF because the path selection algorithm is not altered, and because the OSPF configured costs are also advertised in R-LSAs. It is thus possible to establish adjacencies among routers running UC-QoSR and OSPF.

3.2 Mechanisms for scalability

QoS routing protocols must contribute to a significant improvement in traffic performance and network resource usage to compensate for the burden they introduce on the network. This overhead is twofold, comprising an increase in the communication load due to routing traffic and a raise in the processing capacity of routers caused by the frequency of path computations. In UC-QoSR, these overheads are controlled by a policy that controls the emission of link state updates. This policy combines metrics quantification and threshold based diffusion. A similar approach was followed by Apostolopoulos *et al.* but in flow establishment context [15].

The quantification rule is a moving average of the congestion indexes, with a variable window size (*N*). The congestion indexes are monitored every second (the lowest time granularity provided by GateD) and the samples are taken continuously. In Equation 2, $MA_d(k)$ is the moving average of *N* values of the delay congestion indexes at sample *k*. This function is used to filter the peaks of the QoS metric.

$$MA_d(k) = \sum_{i=k-N}^{k} \frac{d(i)}{N}. \quad \textbf{(2)}$$

The filtered values, resulting from the application of Equation 2, are then presented to the diffusion control module. In this module, the new value is compared with the one that was previously advertised, and will be diffused only if it significantly different. The decision to issue the advertisements is controlled by the value of a defined threshold.

Besides the link state update policy described above, in UC-QoSR, OSPF was modified, in order to control even further the protocol overhead and thus increase the possibility of scalability. In original OSPF, the routing messages denominated Network-LSA (N-LSA) identify the routers connected to the network and its diffusion occurs wherever R-LSAs are issued. In the UC-QoSR strategy, the emission of N-LSAs has been detached from the emission of R-LSAs, because R-LSAs are issued at a higher rate than in OSPF and the information transported in N-LSAs does not change at such a rate. Thus, in the UC-QoSR, the emission of N-LSA remains periodic and dependent on router connectivity, while the emission of R-LSA is controlled through the threshold of the diffusion module. This strategy allows for a significant reduction of routing messages in the network.

The policy to control protocol overhead described above contributes also to avoid the number of path shifts that may occur in the network. Combined with these procedures, the UC-QoSR strategy uses a mechanism named class-pinning, that controls the path shifting frequency of all traffic classes.

3.3 Mechanism of class-pinning

The main role of QoS routing is to dynamically select paths based on information about the state of the network. Therefore, they enable the avoidance of congested paths, contributing to the improvement of application performance. However, the dynamic selection of paths may cause routing instability and network oscillatory behavior. This will naturally degrade application performance. In face of this scenario it is necessary to achieve a compromise between the desired adaptability of the protocol and the unwanted instability [20, 21].

In this work a mechanism of class-pinning to avoid instability is proposed. This mechanism addresses the stability problem described above by controlling the instant when a traffic class shifts to a new path.

When the state of the network changes (due to an event like the start of a new flow or a traffic burst) routing messages are sent to all routers, and new paths are computed. After the calculation, traffic will shift to the less congested paths, leaving the paths currently used. The next time this process occurs, traffic will eventually go back to the original path, and, thus, instability happens.

With the class-pinning mechanism, new paths are computed upon the arrival of routing messages. However, they will be used only if they are significantly better than the path that is currently used by that class. The *Degree of Significance* (DS) parameter is used to support the pinning decision. This parameter establishes the threshold for path shift from the old to the new path.

When the routing information about the state of the network is outdated, bad routing decisions can be made, and congestion may rise. Routing information may be outdated because it is not distributed with enough frequency or because it does not arrive due to congestion. The first cause can be avoided with the appropriate tuning of threshold used in the diffusing module. In order to avoid the delay or loss of routing information, the UC-QoSR routing messages are classified in the class with higher priority. The importance of the priority treatment of routing messages was shown by Shaikh *et al.* [20].

4 Experimentation

In this section the experimentation made to evaluate the stability and scalability of UC-QoSR are presented and its results are analyzed.

4.1 Test Conditions

The test-bed used for the experiments presented in this section is depicted in Figure 1. The *endpoints* 1 to 4 are traffic sources and *endpoints* 5 to 8 are traffic destinations. Each endpoint only generates or receives traffic of a single class to avoid the influence of endpoint processing on traffic patterns. Traffic was generated and measured with the traffic analysis tool Chariot from NetIQ[2].

The routers are PCs with the FreeBSD operating system. The kernel is modified, at the IP level, to include the delay and loss metric modules and to schedule and drop packets according to class sensitivity to these parameters. The monitoring of the delay and loss congestion indexes is needed for the routing decision. The kernel is also modified to interact with the UC-QoSR protocol embedded in GateD. It keeps the

[2] <http://www.netiq.com>

routing table with paths for all traffic classes and makes packet forwarding decisions based on destination IP address and Differentiated Services Code Point (DSCP) [2].

The interfaces between endpoints and routers are configured at 100 Mbps. Interfaces between routers are configured at 10 Mbps to introduce bottlenecks. In the results presented, the moving average window size is 80 samples and the threshold that controls the diffusion of R-LSAs is 30%. These values resulted from the tuning that was done by extensive experimentation with combinations of configurations [5].

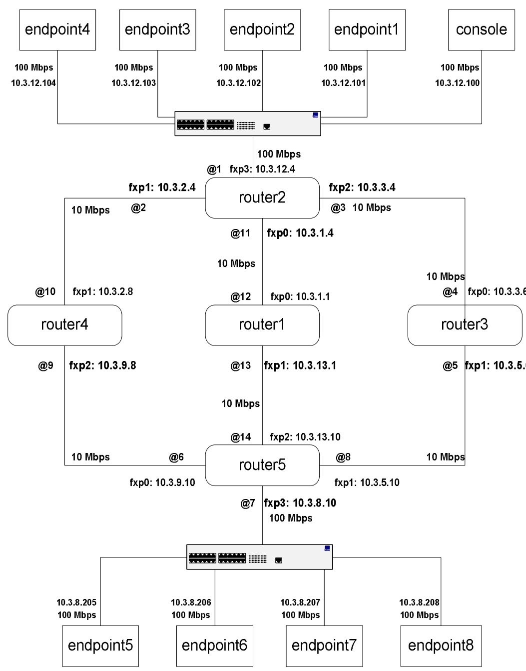

Fig. 1. Experimental test-bed

The evaluation of the UC-QoSR stability and scalability capabilities was done at protocol behavior and traffic performance levels. At the protocol level, the indicators used to measure the degree of stability and scalability of the UC-QoSR strategy are the parameters that represent the dynamics of the protocol, as follows:

a) Number of routing messages issued (Router-LSA);
b) Number of routing messages received (Router and Network[3] LSA);
c) Number of times the Shortest Path First (SPF) algorithm is applied;
d) Number of path shifts.

---

[3] The measure used is the total number of LSAs received, that is, the number of Router and Network LSAs received.

Each experiment was carried out for five minutes and was repeated ten times. The results present the averaged values of all tests. The inspection of protocol dynamics was done in all routers using the OSPF-Monitor tool included in GateD. The evaluation of traffic performance was made according to throughput and loss rate of all active traffic classes. These values were measured by the application Chariot. The plotted results have a degree of confidence of 95%.

4.2 Scalability evaluation

The evaluation of the scalability of the UC-QoSR strategy was done considering different combinations of traffic characteristics namely, number of traffic classes, levels of traffic load, number of flows and packet sizes. The combinations that were used are depicted in the following tables.

Table 2 shows the load distribution that was used in the experiments with 2, 3 and 4 traffic classes to attain the level of total load used in each experiment. The traffic of all classes is UDP to avoid the influence of TCP flow control in the results.

Table 1. Test conditions for stability evaluation

Set of tests	Parameter	Values
A	Number of classes	1, 2, 3, 4
	Level of total load	5, 10, 15, 20, 25, 30, 35, 40 Mbps
B	Number of flows	4, 8, 12, 16, 20
	Packet size	64, 128, 256, 512, 1024, 1460 byte

Table 2. Load distribution per different traffic classes

Total load (Mbps)	Load of each class (Mbps)		
	2 Classes	3 Classes	4 Classes
5	2,5	1,6	1,25
10	5	3,3	2,5
15	7,5	5	3,75
20	10	6,6	5
25	12,5	8,3	6,25
30	15	10	7,5
35	17,5	11,6	8,75
40	20	13,3	10

The results of the evaluation of the impact of the number of classes and level of load on the scalability of the UC-QoSR strategy are presented in Figures 2 and 3. The results shown are measured at *router 2*, since this is the router where exists the most significant bottleneck. These figures show that the load on the network only influences the indicators used for scalability evaluation for loads over 20 Mbps. This stems from the fact that congestion starts to occur in the routers while they adapt the paths to the traffic that is being generated. In the case where there are two traffic classes, each class is generated at 10 Mbps (Table 2). When four traffic classes are generated, each at 5 Mbps, two of those classes must share one path, since there are only three alternative paths in the test-bed. This setting clearly creates congestion, since all the capacity of the interfaces is used. However, the impact of higher loads is not significant, showing that the mechanisms used in UC-QoSR allow for scalability, controlling the rate of routing information diffusion and path calculation.

The application of the smoothing moving average filter to the QoS metric and the utilization of the threshold to control the emission of routing messages limit the number of R-LSA that the router emits. The control over the emission rate of R-LSA will naturally influence the other protocol indicators, the number of SPF and the number of path changes.

The comparison of figures 2 and 3 also shows that the impact of the number of classes is not significant in protocol dynamics. There is only a slight increase in the number of LSA received by *router 2* and, consequently, on the number of SPF calculated. This is, in fact, what was expected, since with four traffic classes all alternative paths are used. In these circumstances, all routers experience traffic changes, issuing the adequate routing messages to distribute their state.

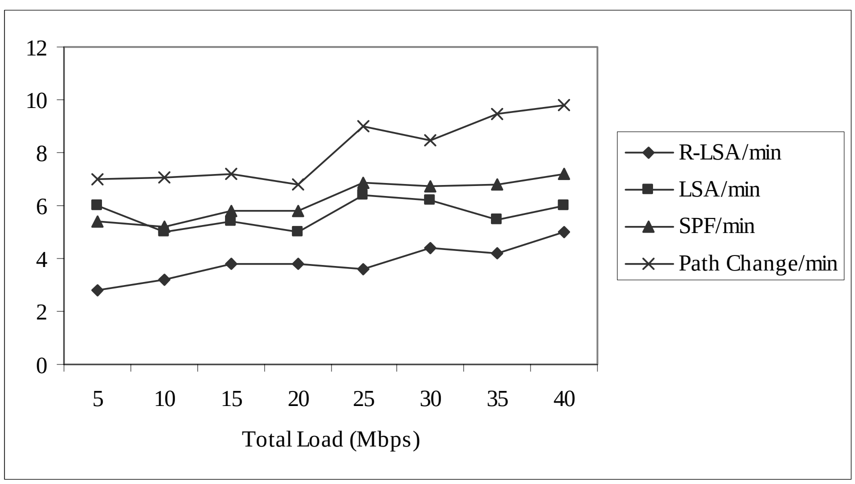

Fig. 2. Measures of protocol dynamics for different levels of traffic load (2 traffic classes)

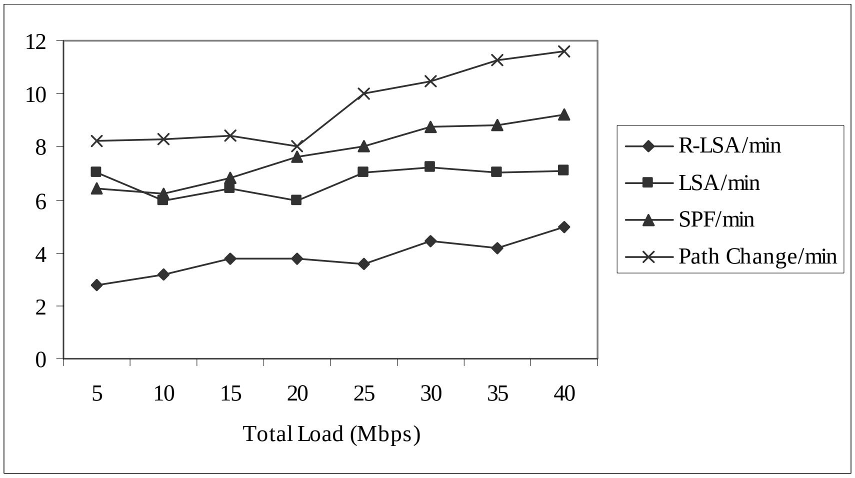

Fig. 3. Measures of protocol dynamics for different levels of traffic load (4 traffic classes)

In the evaluation of the scalability of the UC-QoSR strategy according to the number of flows and packet size, were generated four traffic classes and the traffic load was kept constant at 20 Mbps. The results are shown in Figures 4 and 5. Figure 4 shows that the indicators of protocol dynamics increase with the number of flows.

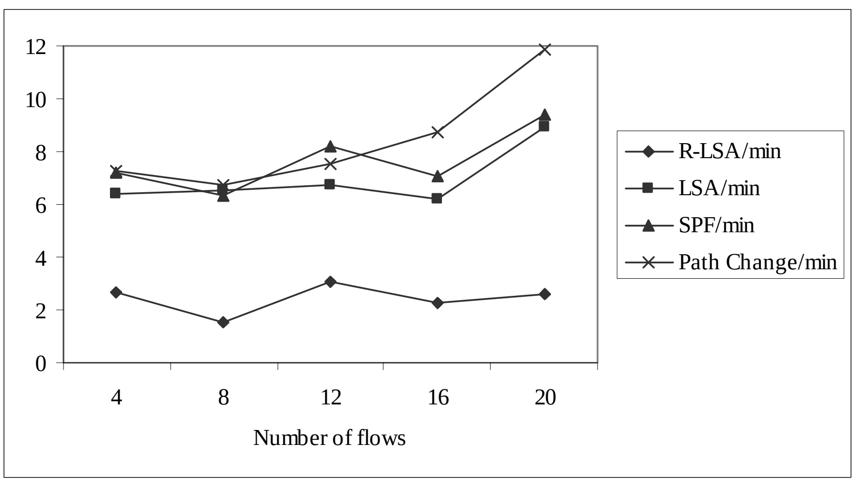

Fig. 4. Measures of protocol dynamics for packets size of 256 byte

The number of LSA received is higher (about 15%), and thus the number of SPF also increases. The raise in the number of path shifts is due to the adaptation process of the flows in each class. For packets of 512 byte and over, the relationship between the number of flows and the indicators of protocol dynamics does not have the same nature (Figure 5). In this case, there is only an increase in the LSA received and the number of SPF. The number of path shifts is not affected by the number of flows. These results show that the UC-QoSR strategy has good scaling capabilities under the presence of an increasing number of flows, unless packets are too small. However, this problem can be overcome if faster processors are used.

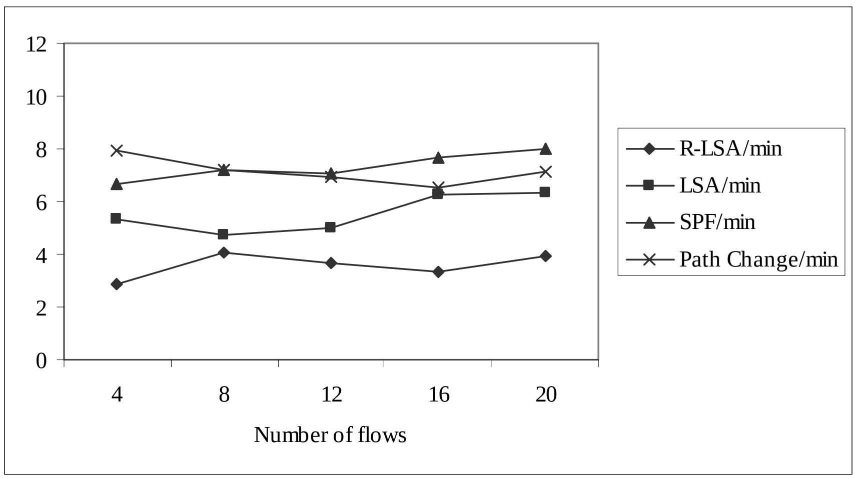

Fig. 5. Measures of protocol dynamics for packets size of 1460 byte

Another conclusion from the observation of the results above is that the UC-QoSR strategy achieves load balancing over alternate paths. This capability is achieved through class-based routing and avoids sudden traffic fluctuations and the corresponding routing oscillations. However, this has shown insufficient to avoid oscillations in the presence of high levels of congestion, leading to the proposal of class-pinning included in the next sub-section.

4.3 Stability evaluation

In the evaluation of the stability of the UC-QoSR strategy four traffic classes were generated. In order to evaluate the degree of stability achieved by the class-pinning mechanism, the *Degree of Significance* parameter was varied from 0% to 50%. The indicator used to evaluate stability at the protocol level was the number of path shifts as in [22].

Figure 6 shows the results of the experiments concerning stability evaluation under different traffic loads. The level of the load introduced in the network has a direct impact in the number of path shifts. However, this impact is reduced for DS values of 20% and 30%, showing the effectiveness of the class-pinning mechanism. It is interesting to remark that the use of strong class-pinning, under higher load, introduces instability in the network, showing results similar to those obtained when the mechanism is not active. This fact is originated by the exaggerated postponement of path change. When the difference in path cost is sufficient to shift traffic, it is high and all traffic classes tend to shift at the same time to the same path, generating instability.

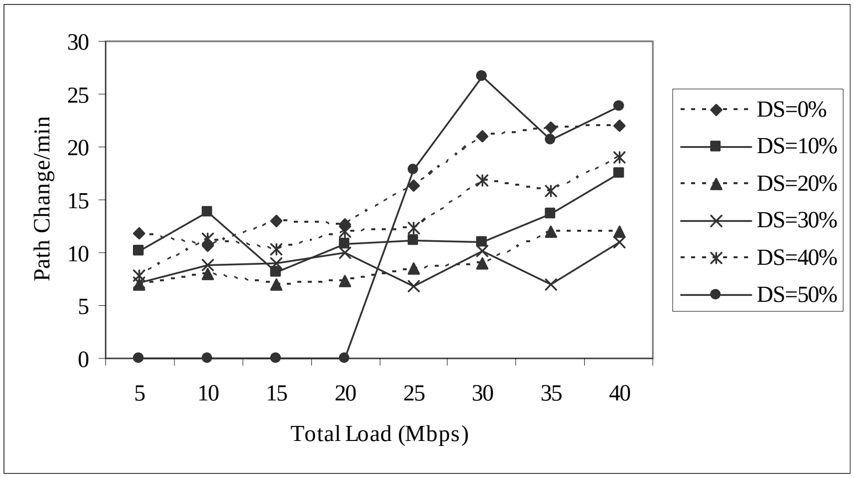

Fig. 6. Evaluation of the class-pinning mechanism

The throughput of traffic classes under two different values of DS is deployed in Figure 7. Traffic of classes 1 and 2 was generated at 6 Mbps and traffic of classes 3 and 4 at 4 Mbps.

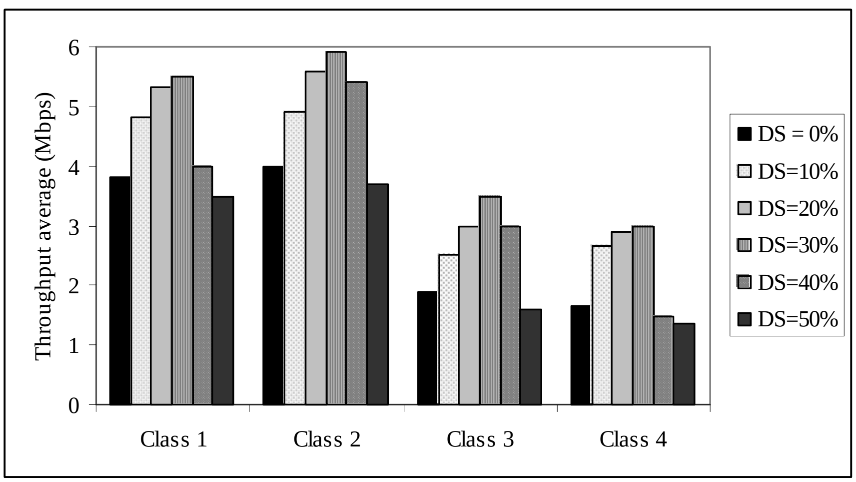

Fig. 7. Traffic throughput with different levels of class-pinning

This figure shows that the activation of the class-pinning mechanism contributes to increase traffic performance while maximizing resource utilization (throughput increases about 50% when is used a DS of 30%).

5 Conclusions and Future Work

At the University of Coimbra a QoS routing strategy (UC-QoSR) was developed to support class-based traffic differentiation. The paths suitable for traffic classes are computed based on a QoS metric that evaluates the state of the network in terms of delay and losses. The proposed mechanisms were implemented as an extension to OSPF routing protocol on the GateD platform.

Previous experiments with the UC-QoSR showed that the overhead introduced was affordable by the communication system and traffic differentiation was achieved. However some instability was noticed when the network was congested. To overcome this instability, a set of mechanism where conceived to integrate the UC-QoSR. Besides instability, scalability is other important issue in QoS routing. The focus of this paper was the proposal and evaluation of stability and scalability mechanisms to integrate the UC-QoSR strategy.

The results showed that the UC-QoSR strategy with the proposed mechanisms scales well under heavy loads. The impact of the number of flows is only significant when packets are small. The experiments also showed that the use of the class-pining mechanism significantly improves stability and traffic performance. However, for strong pinning, the behavior is worst due to the difference between paths.

Acknowledgements

This work was partially supported by the Portuguese Ministry of Science and Technology (MCT), under program POSI (Project QoS II and IPQoS) and under the PhD grant PRAXIS XXI/ BD/13723/97.

References

[1] R. Braden, D. Clark, S. Shenker, “Integrated Services in the Internet Architecture: an Overview”, Request For Comments 1633, Internet Engineering Task Force, June 1994.

[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] J. Moy, “OSPF Version 2”, Internet Engineering Task Force, Request For Comments 2328, April 1998.

[4] M. Oliveira, J. Brito, B. Melo, G. Quadros, E. Monteiro, “Quality of Service Routing in the Differentiated Services Framework”, Proceedings of SPIE’s International Symposium on Voice, Video, and Data Communications (Internet III: Quality of Service and Future Directions), Boston, Massachusetts, USA, November 5-8, 2000.

[5] M. Oliveira, J. Brito, B. Melo, G. Quadros, E. Monteiro, “Evaluation of a Quality of Service Routing Strategy for the Differentiated Services Framework”, Proceedings of the 2001 International Conference on Internet Computing (IC'2001), Monte Carlo Resort, Las Vegas, Nevada, USA, June 25-28, 2001.
[6] A. Khanna, J. Zinky, “The Revised ARPANET Routing Metric”, Proceedings of SIGCOMM’89, Austin, Texas, September 19-22, 1989
[7] 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.
[8] Z. Wang, J. Crowcroft, “Shortest Path First with Emergency Exits”, Proceedings of SIGCOMM'90, Philadelphia, USA, September 1990.
[9] S. Vutukury and J.J. Garcia-Luna-Aceves, “A Simple Approximation to Minimum-Delay Routing”, Proceedings of ACM SIGCOMM’99, Harvard University Science Center, Cambridge, Massachusetts, USA, 31 August-3 September, 1999.
[10] K. Nahrstedt, S. Chen, “Coexistence of QoS and Best Effort Flows - Routing and Scheduling”, Proceedings of 10th IEEE Tyrrhenian International Workshop on Digital Communications: Multimedia Communications, Ischia, Italy, September, 1998.
[11] Q. Ma, P. Steenkiste, “Supporting Dynamic Inter-Class Resource Sharing: A Multi-Class QoS Routing Algorithm”, Proceedings of IEEE INFOCOM’99, New York, USA, March 1999.
[12] A. Shaikh, J. Rexford, K. Shin, “Load-Sensitive Routing of Long-Lived IP Flows”, Proceedings of ACM SIGCOMM’99, Harvard University Science Center, Cambridge, Massachusetts, USA, August 31-September 3, 1999.
[13] P. Van Mieghem, H. De Neve and F. Kuipers, "Hop-by-hop Quality of Service Routing", Computer Networks, vol. 37. No 3-4, pp. 407-423, 2000.
[14] J. Wang, K. Nahrstedt, “Hop-by-Hop Routing Algorithms for Premium-class Traffic in DiffServ Networks”, Proceedings of IEEE INFOCOM 2002, New York, NY, June, 2002.
[15] G. Apostolopoulos, R. Guerin, S. Kamat, and S. Tripathi. "Quality of Service Based Routing: A Performance Perspective", Proceedings of SIGCOMM'98, Vancouver, British Columbia, September 1998.
[16] B. Lekovic, P. Van Mieghem, “Link State Update Policies for Quality of Service Routing”, Proceedings of Eighth IEEE Symposium on Communications and Vehicular Technology in the Benelux (SCVT2001), Delft, The Netherlands, October 18, 2001.
[17] G. Quadros, A. Alves, E. Monteiro, F. Boavida, “An Approach to Support Traffic Classes in IP Networks”, Proceedings of QoFIS’2000 – The First International Workshop on Quality of future Internet Services, Berlin, Germany, September 25-26, 2000.
[18] Z. Wang, J. Crowcroft, “Quality of Service Routing for Supporting Multimedia Applications”, IEEEJSAC, , vol. 14. No 7, pp. 1228-1234, September 1996.
[19] H. De Neve, P. Van Mieghem, “TAMCRA: A Tunable Accuracy Multiple Constraints Routing Algorithm”, Computer Communications, 2002, Vol. 23, pp. 667-679.
[20] A. Shaikh, A. Varma, L. Kalampoukas and R. Dube, “Routing stability in congested networks” Proceedings of ACM SIGCOMM’00, August 28-Sepetmber 1, Grand Hotel, Stockholm, Sweden, 2000.
[21] A. Basu, J. Riecke, “Stability Issues in OSPF Routing”, Proceedings of ACM SIGCOMM’01, San Diego, California, USA, August 27-31, 2001.