CLASS-PINNING: A NOVEL APPROACH FOR HOP-BY-HOP QOS 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
University of Coimbra
Pólo II, Pinhal de Marrocos, 3030-290 Coimbra
Portugal

Abstract

The need for Quality of Service (QoS) on the Internet is increasing everyday due to the wide dissemination of new communication services with more stringent demands from the network. There have been several proposals for QoS provisioning, particularly, the Integrated Services and Differentiated Services frameworks developed on the Internet Engineering Task Force (IETF).

As the deployment of Differentiated Services on the Internet becomes a reality, it urges for its combination with QoS aware routing protocols. Among other proposals, at the Laboratory of Communications and Telematics of the University of Coimbra is being developed an intra-domain QoS routing protocol (UC-QoSR).

In this paper a mechanism of *class-pinning* to integrate the UC-QoSR strategy is proposed and evaluated. This mechanism addresses the stability problem associated with congestion based routing by controlling the instant when a traffic class shifts to a new path. A traffic class is moved to a new path only if it is considerably better than the current path. Results show that with the use of class-pinning it is possible to maintain a adequate path for each traffic class and to avoid the instability due to frequent path shifts.

Keywords: QoS routing, Differentiated Services, Stability

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 IEFT [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, it is selected the shortest path, 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 [3]. 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 protocol of the University of Coimbra (UC-QoSR) was conceived to fulfill this purpose.

The UC-QoSR protocol 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 First (OSPF) protocol in order to select paths appropriate for all traffic classes [4, 5].

A prototype of UC-QoSR was implemented over the GateD platform, running on the FreeBSD operating system [6]. This prototype was used to evaluate the routing strategy according to communication and processing overhead. The measurements obtained showed that the overhead introduced was affordable by the communication system [7] and traffic differentiation was achieved.

However some instability was noticed when the network was congested or in the presence of bursty traffic.

To overcome the instability problem described above a *class-pinning* mechanism was conceived to integrate the UC-QoSR. This mechanism is described and evaluated in this paper. The rest of the paper is organized as follows: Section 2 describes the UC-QoSR strategy and the mechanism of class-pinning; test conditions and analysis of results are presented in Section 3; the main conclusions and issues to be addressed in future work are presented in Section 4.

2 UC-QoSR Strategy

In this section the main characteristics of the routing strategy UC-QoSR are described. The mechanism of class-pinning conceived to maintain stability is presented in detail.

2.1 Characteristics of UC-QoSR

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

1. A QoS metric that represents the availability of resources in the network;
2. Traffic class requirements in terms of QoS parameters;
3. 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 [8]. The congestion indexes are distributed to all routers in the domain through modified OSPF routing messages (Link State Advertisements – 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, as shown in Table 1. The strategy is obviously configurable to accommodate traffic with other combinations of delay and loss sensitivity.

Class	**Sensitivity to delay**	**Sensitivity to loss**
1	Low	Low
2	High	Medium
3	High	Low
4	Medium	High

Table 1. Traffic classes in UC-QoSR.

The path computation algorithm uses a cost function that combines both congestion indexes, weighted according to delay and loss sensitivity of each traffic class. 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.

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 kept at a minimum, because it is used a mechanism of hysteresis that controls the frequency of distribution of routing information. This mechanism limits both communication and processing overhead.

This approach also avoids the number of path shifts that may occur due to bursts. The path selection decision is class-based and thus, if there are enough available paths on the network, each traffic class follows an individual path, achieving load balancing. Combined with these procedures, UC-QoSR uses a mechanism named class-pinning, that controls the path shifting frequency of all traffic classes.

2.2 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 [9, 10, 11].

In this paper 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.

3 Evaluation of Class-pinning

In this section the experimentation done to evaluate the class-pinning mechanism of UC-QoSR is presented and analyzed.

3.1 Test conditions

The testbed used for the evaluation of the class-pining mechanism 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[1].

The routers are PCs with 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 routing table with paths for all traffic classes and makes packet forwarding decisions based on destination IP address and Differentiated Services Code Point (DSCP) [12].

The interfaces between endpoints and routers are configured at 100 Mbps. Interfaces between routers are configured at 10 Mbps to introduce bottlenecks.

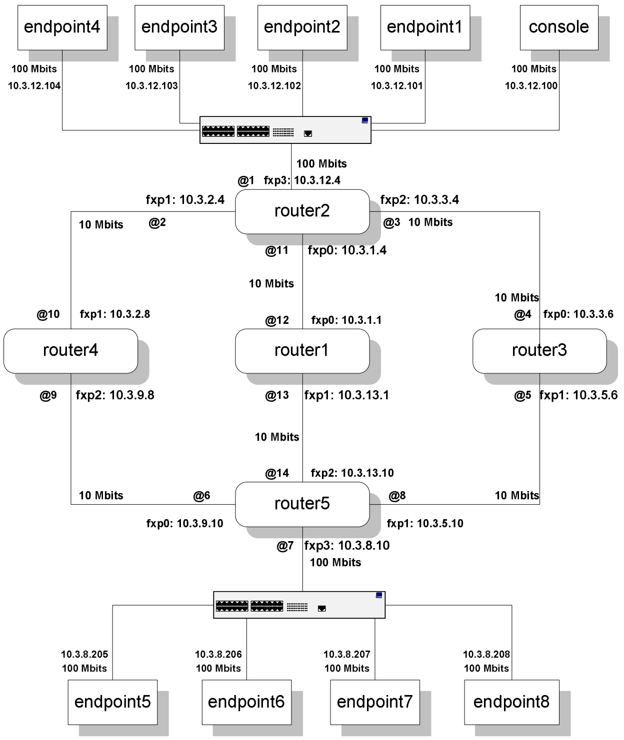

Figure 1. Experiments testbed.

Traffic flows of all the active classes were generated to create a total load of 36 Mbps, in order to induce a high level of congestion in the network. The traffic mix used is described in Table 2. The traffic of all classes is UDP.

Classes	Number of flows	Throughput of each flow (Mbps)
1 and 2	8	4,5
1, 2 and 4	12	3
1, 2, 3 and 4	16	2,25

Table 2. Traffic combinations.

The evaluation of the class-pinning mechanism was done at protocol and traffic performance levels.

The examination of protocol dynamics was done in all routers using the OSPF-Monitor tool included in GateD. The parameters evaluated were the following:

1. Number of routing messages issued (Router-LSA);
2. Number of routing messages received (Router and Network LSA);
3. Number of times the Shortest Path First (SPF) algorithm is applied;
4. Number of path shifts.

---

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

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.

3.2 Result analysis

Figure 2 shows an overall view of the dynamics of the UC-QoSR strategy in the experiments (routers are identified from R1 to R5).

The number of SPF computations is the same in all routers. This is due to the nature of the link state protocol used. However, router 2 issues a higher number of R-LSA. This event stems from the fact that this is the router where traffic enters the network, suffering a higher degree of load changes. Since router 2 issues more R-LSA, the other routers receive a higher number of LSA as it is depicted in Figure 2. On the other hand, router 2 receives a relatively small number of LSA.

A higher number of path shifts also occurs in router 2. This happens for the reason described above and also because this router has more interfaces than the other routers, and thus more alternative paths.

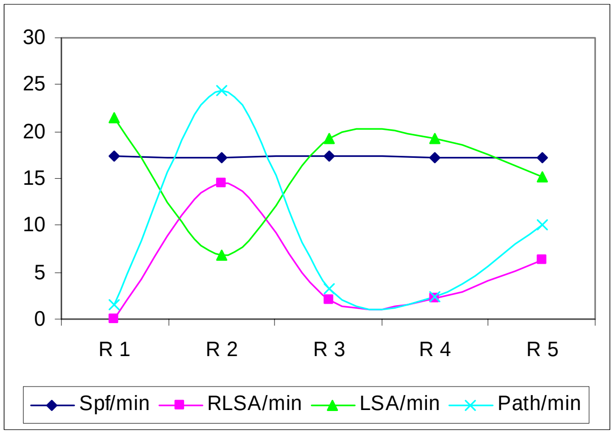

Figure 2. Comparison of routers dynamic parameters.

Next figures show the results concerning traffic performance. The values of the Degree of Significance parameter used by the class-pinning mechanism are 0% (the mechanism is not active), 25% and 50%. Flows with an odd number belong to class 1 and flows with an even number belong to class 2.

The evaluation of the class-pinning mechanism shows that there is a significant improvement in traffic performance when the mechanism is activated with a DS value of 25%. However, when DS is increased to 50% the impact is not noticeable. In some circumstances it even degrades performance. This is due to the fact that the pinning is too strong preventing the use of more interesting paths. These results clearly show the need to compromise between stability and adaptation to network dynamics.

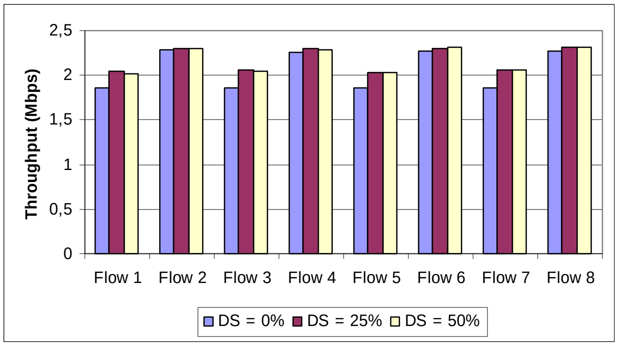

Figure 3. Throughput of all traffic flows.

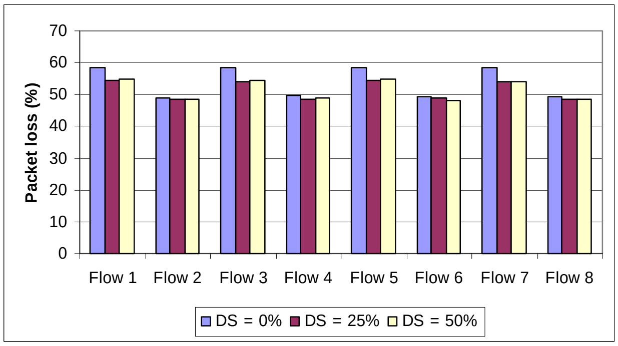

Figure 4. Packet loss percentage of all traffic flows.

In figures 5 and 6 the throughput and packet loss rate felt by traffic of classes 1 and 2 are shown, for different values of the DS parameter.

Two important conclusions can be drawn from these graphics. The first conclusion is that traffic differentiation is correctly made, since class 2 achieves higher throughput and suffers lower packet loss rate than class 1. The second conclusion is that the activation of the class-pining mechanism significantly improves traffic performance, as was observed in the previous graphics. Figure 5 also shows that, with a DS of 25%, the full capacity of the network interface cards is used (configured at 10 Mbps).

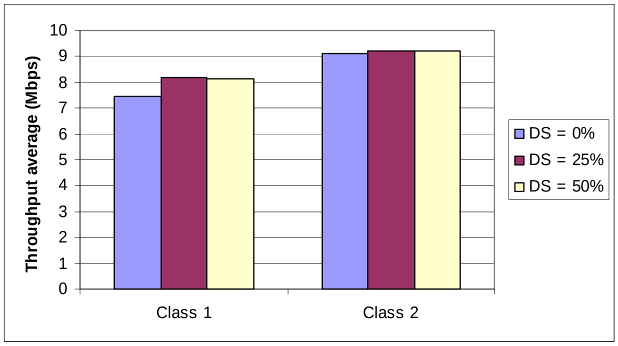

Figure 5. Throughput average of classes 1 and 2.

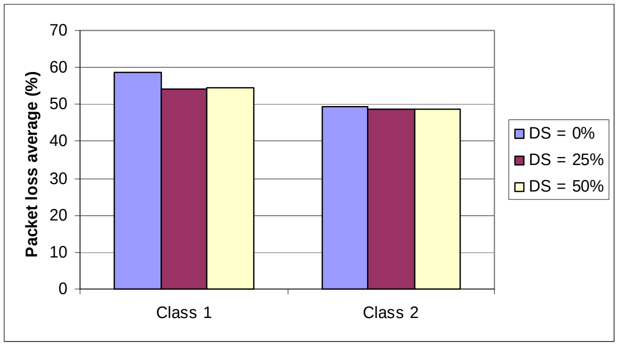

Figure 6. Loss average of classes 1 and 2.

Figure 7 shows the throughput average in the presence of four traffic classes. In this set of tests 16 flows of 2,25 Mbps each were generated to attain a total load of 36 Mbps. It can be observed that the activation of the class-pinning mechanism, with a DS of 25%, slightly contributes to performance improvement. However, when DS is 50%, the performance is worst than without the activation of the mechanism. This fact is originated by an excessive pinning that prevents the usage of alternative paths.

Figures 7 and 8 show that class 1 is better treated than the other traffic classes. It was observed that this class uses only one path and the others tend to share paths among them. This is an aspect that must be analyzed in more depth. Eventhough this undesired behavior, results show that the other classes are differentiated according to their delay and loss sensitivity.

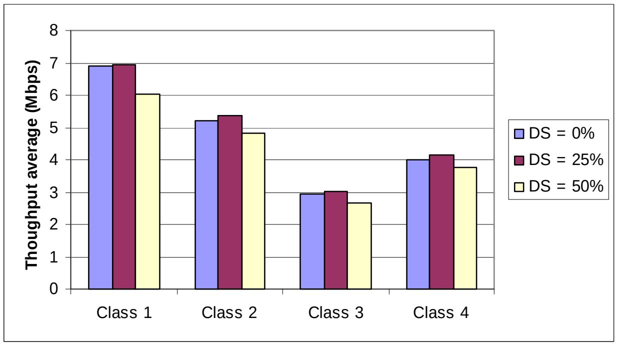

Figure 7. Throughput average of classes 1 to 4.

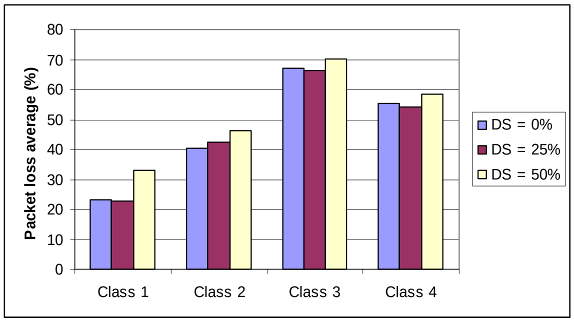

Figure 8. Loss average of classes 1 to 4.

In other test sets not included here, better performance was achieved for situations where congestion was not as strong as in the testes presented in this paper. However, even with this level of load, the proposed mechanism demonstrates very promising results.

4 Conclusions and Future Work

At the University of Coimbra a QoS routing strategy (UC-QoSR) is being 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.

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 or in the presence of bursty traffic. To overcome the instability problem described a *class-pinning* mechanism was conceived to integrate the UC-QoSR.

The focus of this paper was the evaluation of the class-pinning mechanism. The results showed that the use of the class-pining mechanism significantly improves traffic performance. However, for strong pinning, the impact in performance is not noticeable. In some circumstances it even degrades performance. This is due to the fact that the pinning mechanism prevents the use of more interesting paths.

The fine tuning of the class-pinning mechanism considering other network topologies and different traffic patterns is currently under study.

The full evaluation of the UC-QoSR strategy in wider scenarios is being carried out by simulation. These experiments will allow the scalability evaluation of the proposed strategy.

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, Network Working Group, June 1994.

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

[3] 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.

[4] J. Moy, OSPF Version 2, *Request For Comments 2328*, Internet Engineering Task Force, Network Working Group, April 1998.

[5] R. Guérin, S. Kamat, A. Orda, T. Przygienda, D. Williams, QoS Routing Mechanisms and OSPF Extensions, *Request For Comments* 2676, Internet Engineering Task Force, August 1999.

[6] 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.

[7] 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.

[8] 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.

[9] A. Shaikh , A. Varma , Lampros Kalampoukas , Rohit Dube, Routing stability in congested networks, *ACM SIGCOMM Computer Communication Review, Proceedings of the conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, August 2000.

[10] S. Vutukury, J.J. Garcia-Luna-Aceves, A Simple Approximation to Minimum-Delay Routing, *Proceedings of ACM SIGCOMM'99*, Cambridge, MA, USA, September 1999.

[11] Z. Wang, J. Crowcroft, Shortest Path First with Emergency Exits, *Proceedings of ACM SIGCOMM'90*, Philadelphia, USA, September 1990.

[12] D. Grossman, New Terminology and Clarifications for Diffserv, *Request For Comments 3260*, Internet Engineering Task Force, Network Working Group, April 2002.