Queue Management and QoS Routing for Traffic Differentiation

Marilia Curado, Ricardo Veludo, Edmundo Monteiro
Laboratory of Communications and Telematics
CISUC/DEI
University of Coimbra
Pólo II, Pinhal de Marrocos, 3030-290 Coimbra, Portugal
Telephone: +351 239 790000 Fax: +351 239 701266

Abstract

This paper presents a simulation study of router mechanisms to provide differentiated levels of service to traffic with diverse performance requirements in IP networks. The paper focuses on queue management mechanisms and on Quality of Service routing. The performance of the Random Early detection dropper associated with the Weighted Round Robin scheduling discipline is compared with the Dynamic Degradation Distribution system. This system redistributes the resources among traffic classes according to the state of the route. Afterwards, the impact of Quality of Service routing in networks where there is class-based traffic differentiation is assessed. The results show that even though the queue management mechanisms actually deployed in commercial routers naturally support traffic differentiation and provide adequate levels of QoS, the Dynamic Degradation Distribution system is able to give better performance in situations of congestion. Moreover, the inclusion of Quality of Service routing capabilities clearly improves traffic performance and network utilization.

Keywords: Queue Management, Quality of Service Routing.

1 Introduction

Quality of Service (QoS) 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 Integrated and Differentiated Services architectures proposed by the Internet Engineering Task Force (IETF) [1, 2]. In the first approach, QoS management is associated with resource reservation and traffic is treated at the flow level in order to give guarantees to the requests made by the end-users. These characteristics pose some limitations in terms of scalability due to the amount of state that must be maintained in routers and to the complex resources management mechanisms associated with individual flow handling. The Differentiated Services approach was proposed to overcome these drawbacks. Traffic is aggregated in a limited number of classes with diverse QoS requirements, and the state maintained in each router is only per class instead of being per flow. Moreover, in the Differentiated Services framework, functions such as admission control and traffic classification are only made in border routers, leaving central routers with the only responsibility of performing expedition of traffic aggregates.

The deployment of Differentiated Services requires that the routers are able to give differentiated treatment to traffic with diverse QoS requirements. Scheduling and packet drop are the two main mechanisms deployed in routers to achieve multiple level QoS services. Scheduling algorithms decide the order by which packets in queues are processed and allow for the preferential treatment of some queues over other less important queues. Active drop algorithms control the discard of packets when the number of packets in queues increases. Although ordinary drop algorithms only discard packets when the queues become full, active drop algorithms start to discard packets when congestion rises, but before the queues are full. With this approach, the queues have enough space available to accommodate bursts and at the same time the average queue size remains small, contributing to a better performance of interactive applications, such as interactive audio and video sessions.

Weighted Round Robin (WRR) and Class Based Queuing (CBQ) are, among others, two types of scheduling algorithms used in existing routers [3, 4]. RED and its variants are the most widely used active packet drop mechanisms [5]. There are however other proposals to provide traffic differentiation among traffic classe. In particular, the Dynamic Degradation Distribution (D3) system developed at the University of Coimbra is a queue management mechanism that combines a scheduler and a dropper to give differentiated treatment to traffic with diverse requirements in terms of delay and loss sensitivity following the paradigm of the Differentiated Services [6]. In this paper the traffic performance achieved with both sets of mechanisms is done by simulation.

Even though the above mechanisms are able to provide traffic differentiation in each router, they can have limited performance when used in conjunction with traditional routing protocols. Current routing protocols used in the Internet, such as the Open Shortest Path First (OSPF) protocol lack characteristics for QoS provision to support emerging new services [7]. 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. The shortest path is selected, based on a single static metric that does not reflect the availability of resources and therefore congestion easily occurs on the shortest path, with the corresponding degradation of traffic performance. This scenario has motivated the development of QoS aware routing protocols, with the most significant developments on QoS routing aimed at communication systems where traffic differentiation is done per flow, as in the Integrated Services. Since the Differentiated Services framework does not explicitly incorporate QoS routing, it is interesting to evaluate the impact of providing class-based QoS routing for a service model that does class-based traffic differentiation. In this paper it is analyzed the impact that QoS routing can have in traffic performance in class-based networks, using the Laboratory of Communications and Telematics QoS Routing (LCT-QoSR) strategy developed at the University of Coimbra [8].

This paper is organized as follows: Section 2 gives an overview of the QoS mechanisms under evaluation; Section 3 presents the simulation test conditions and the analysis of the results; conclusions and issues to be addressed in future work are presented in Section 4.

2 Description of QoS mechanisms under evaluation

This section presents a brief description of the QoS mechanisms that are evaluated in this paper.

2.1 Active drop mechanism - RED

RED is an alternative to the traditional packet drop mechanism, tail drop, where a packet is dropped whenever the queue is full. This type of discard behavior brings several disadvantages to tail drop. Namely, tail drop is unable to accommodate bursty traffic since the queues are almost always full. Moreover, there is the possibility of lockout caused by the monopolization of the queue by just a few flows. Another problem pertains to the global synchronization among TCP flows due to the congestion control mechanisms used. RED overcomes these drawbacks by detecting early congestion and taking preventive measures. The RED algorithm has three different behaviors depending on the average queue size, av_{qs}, and two thresholds, min_{th} and max_{th}. If the average queue size is below the min_{th} threshold, no action is taken and if the average queue size is above the max_{th} threshold, all packets are dropped. When the average queue size is between the two thresholds, packets are marked with a discard probability that reflects the state of congestion of the queue. Although the advantages of using RED over passive drop mechanisms is clear, the fine tuning needed to achieve an interesting configuration of all the parameters involved has been pointed out as its main drawbacks [9].

2.2 Scheduling algorithm - WRR

Besides deciding on witch packets to drop, there is the need to decide from which queue to serve packets. WRR is the scheduling algorithm used in this paper. The WRR scheduling discipline distributes available bandwidth among the classes supported using a weighted round robin scheme.

2.3 Class-based traffic differentiation – D3

The main objective of D3 is to treat the traffic of each class at the router in such a manner that the performance degradation that may occur due to an increase in the traffic load will be mainly absorbed by the less sensitive classes, therefore protecting the most sensitive classes. The performance degradation in then dynamically distributed among classes according to their sensitivity to delay and loss. The redistribution of scheduling and queue length resources is based on a QoS metric composed of a delay and a loss congestion index. The indexes measure the impact that queuing delay and packet drop have on application performance.

Figure 1 illustrates the dynamics of the delay congestion index for three classes with different sensitivities to delay. In a situation when the load increases at a router, the delay of all the classes increases, but the scheduler will process packets from the class with higher sensitivity to delay in order to protect it. The objective of the scheduler is to process packets from different classes in such a way that all classes have the same delay congestion index, higher than the one with previous load conditions. This delay congestion index will correspond to different delay increases, being higher for low sensitive classes. A similar approach is used for the self-configuration of the packet drop mechanism.

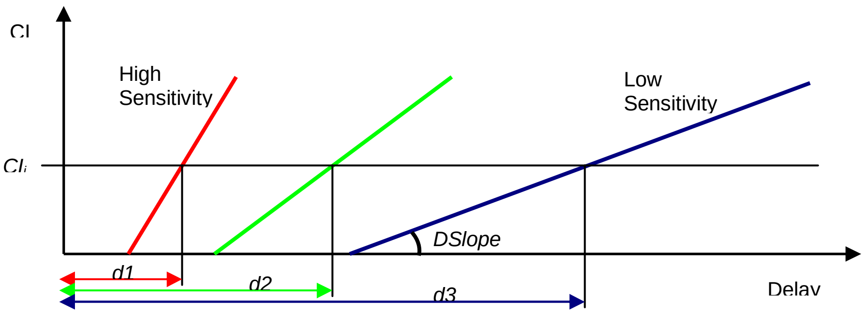

Figure 1 Graphical description of D3 behavior.

The main difference between the QoS mechanisms under evaluation, namely RED plus WRR and D3, is that D3 adapts the treatment given to traffic to the the level of congestion of the interface, while WRR performs scheduling based on the configured weights of each queue.

2.4 QoS routing – LCT-QoSR

The previously described mechanisms differentiate traffic in each router. QoS routing appears as the missing and complementary component to achieve QoS. The main objective of the LCT-QoSR strategy is to select paths suitable to each traffic class based on information about the state of the network. The LCT-QoSR scheme computes paths based on the delay and loss congestion indexes, with two major differences relatively to existing multi-constraint QoS routing proposals. Firstly, the path computation algorithm computes optimal QoS paths in relation to the two additive QoS metric in polynomial time, although most proposals use heuristics to handle algorithm complexity associated with multi-constrained QoS routing. Secondly, LCT-QoSR computes QoS aware paths for the traffic classes, without giving strict performance guarantees, just adding traffic dispersion over alternative paths, and thus needs a limited amount of state to be kept in routers.

The LCT-QoSR strategy extends the OSPF routing protocol to compute multiple paths for each destination, one per traffic class. The metrics concerning delay and loss are combined in a single weight function that takes into consideration class sensitivity to delay and loss. The weight of link *e* for class *i*, represented by $w_i(e)$, is computed by the function depicted in Equation 1, where *DcI(e)* is the delay congestion index, *LcI(e)* is the loss congestion index, the sensitivity of class *i* to delay is represented by δ_i and its sensitivity to loss is represented by σ_i.

$$w_i(e) = \delta_i DcI(e) + \sigma_i LcI(e) \qquad (1)$$

Based on this weight, a shortest path tree is computed for each traffic class using the Dijkstra algorithm as in OSPF.

3 Experimentation

This section describes the simulation settings used for the evaluation of the QoS queue management and QoS routing mechanisms presented above and discusses the results obtained.

3.1 Test conditions

The simulation study was done on the Network Simulator[1] using the mesh topology depicted in Figure 2.

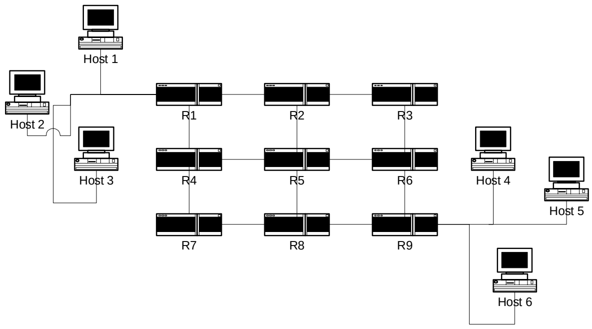

Figure 2 Topology used in the simulation study.

Traffic from three traffic classes was generated in order to create two different load conditions, one where the network is not congested and the other where the network is congested. The definition of the traffic classes took into consideration the requirements of real traffic in the Internet, namely, one class for best-effort traffic, one class for traffic sensitive to delay, and one class for traffic sensitive to loss, as shown in Table 1. Also, in order to cover several types of packets generated by various applications the experiments were done using three packet sizes, namely, 512, 1024 and 1460 Byte.

Type of traffic	Class
Best-effort	1
Delay sensitive	2
Loss sensitive	3

Table 1 Traffic classification.

The configuration of the parameters that control the behaviour of RED was done by experimentation and the values that showed the best performance were used. The values of the RED thresholds and the WRR weights of the queue associated with each traffic class are presented in Table 2. The configuration of class sensitivity to delay and loss used in D3 are also shown in Table 2.

Class	RED thresholds	WRR queue weights	D3 delay slope	D3 loss slope
1	50/70	2	30	30
2	40/60	7	60	30
3	20/70	5	30	60

Table 2 RED and D3 configuration parameters.

Two sets of experiments were done. On the first set, traffic performance under RED and WRR is compared with the performance obtained with D3. On the second set, the performance obtained with D3 using the OSPF routing protocol is compared with the performance when D3 is associated with the LCT-QoSR protocol. The results of the experiments are analysed according to the following aspects:

- Differentiation among traffic classes;

[1] http://www.isi.edu/nsnam/ns/

- Traffic performance for variable packet sizes;
- Comparison of the QoS mechanisms according to the performance metrics used.

3.2 Comparison of RED/WRR and D3

Figure 3 shows the comparison of the performance of traffic from the three classes considered when traffic differentiation is done using RED/WRR and D3 and the network is not congested. The graphic on the right-hand side shows the average delay of packets with variable size, and the left-hand side depicts the corresponding loss rate. Traffic differentiation is correctly achieved with RED/WRR and D3. Specifically, the class with worst treatment is class 1, where best-effort traffic is classified. Moreover, class 2, the class with highest sensitivity to delay has the lowest delay and class 3, the class with highest sensitivity to loss has the lowest loss rate for all packet sizes. However, the differentiation among QoS sensitive classes according to both parameters is quite limited. This result is clearly due to the low level of load, since without congestion there isn't the need to differentiate among traffic classes by penalizing them according to the parameter they are less sensitive to. The results show that there isn't a direct relation between packet size and traffic performance when the network is not congested, except for the best-effort class. The comparison of the traffic performance using RED/WRR and D3 gives a clear advantage to D3 for the best-effort class. Nevertheless, RED/WRR slightly outperforms D3 for QoS sensitive classes.

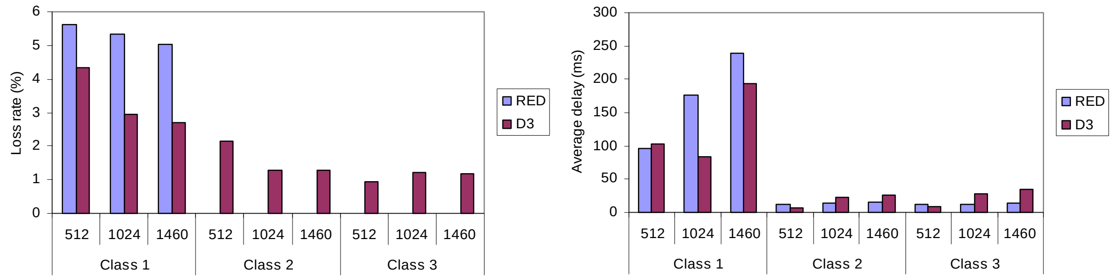

Figure 3 Comparison of RED and D3 without congestion.

When the network is congested the results change significantly from the situation when the network load is lighter, as depicted in Figure 4. The most relevant conclusion pertains to the fact that the D3 performance is better in terms of delay average for all traffic classes. Concerning loss rate, D3 also outperforms RED/WRR with the exception of class 2, the delay sensitive class. In this case, a lower loss rate is achieved for RED/WRR. This behavior is a side effect of the protection of the class against long delays.

Although packet size has a clear impact on the average delay felt by traffic from all classes it does not have a direct effect on the loss rate evaluation parameter. The influence of packet size on delay is due to the processing time associated with the transmission of larger chunks of data.

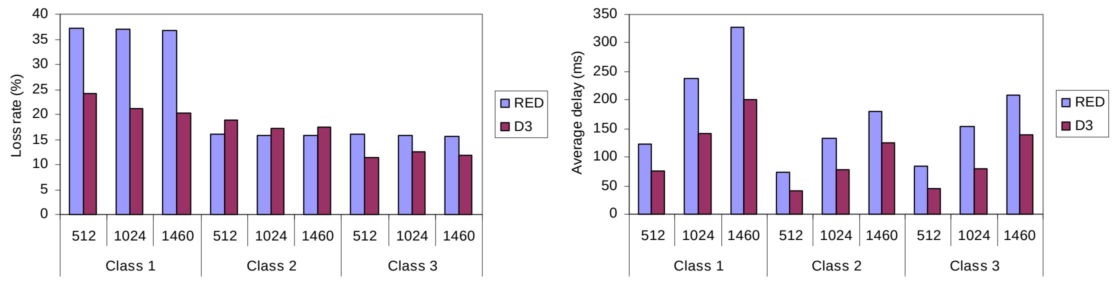

Figure 4 Comparison of RED and D3 with congestion.

The analysis of the results presented in this section shows that by redistributing resources among traffic classes, the D3 scheduling and dropping mechanisms outperform the combination of RED and WRR when the network is in a congestion state. These results agree with the ones presented by Christiansen *et al*. [10], where it is shown the poor performance of RED for web browsing. However, without congestion, the RED/WRR approach has generally better performance, except in terms of the loss rate of best-effort traffic. Since the need for QoS mechanisms is more important when there is congestion in order to achieve a better level of service, the results showed that D3 is able to improve traffic performance when compared to RED/WRR. Moreover, the configuration of D3 has shown to be more simple when compared to the well know complexity and parameter sensitivity of RED/WRR [9].

3.3 Interaction between Queue Management and QoS routing

The mechanisms used for resource management and traffic differentiation in each router are able to improve traffic performance. However, when traditional single shortest path routing protocols are used, it is not possible to take advantage of existing alternative paths. In this section the performance of D3 is compared in two situations, when using the widely used OSPF routing protocol and when using the LCT-QoSR strategy. Figure 5 illustrates the traffic performance obtained with the two situations mentioned when the network is not congested, for variable packet sizes. The results show that even though the network is not congested, there is a clear advantage in using QoS routing, specifically in terms of the delay experienced by all the traffic classes. With QoS routing it is possible to support more traffic than just a single link could otherwise carry.

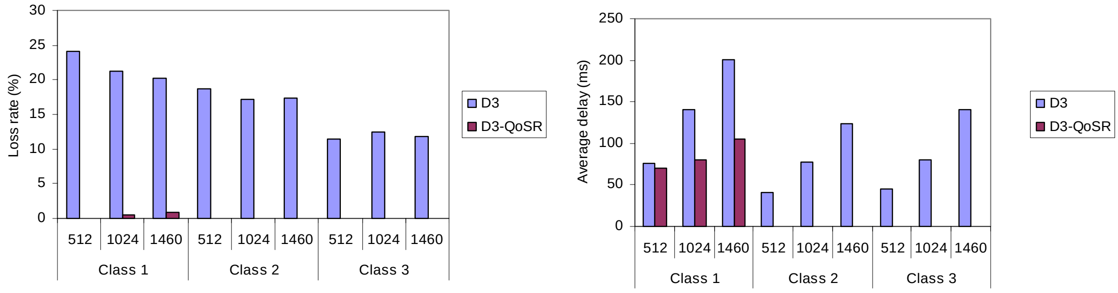

Figure 5 Impact of QoS routing in queue management without congestion.

The traffic performance of the different traffic classes when the LCT-QoSR strategy is used in a situation of congestion is depicted in Figure 6. The results show a clear improvement of using QoS routing, concerning both average delay and loss rate. Although expected, the added value of QoS routing is the distribution of traffic over the alternate paths originated at the source router. While the impact of the D3 system is at the interface level, QoSR affects resource management on all the router interfaces.

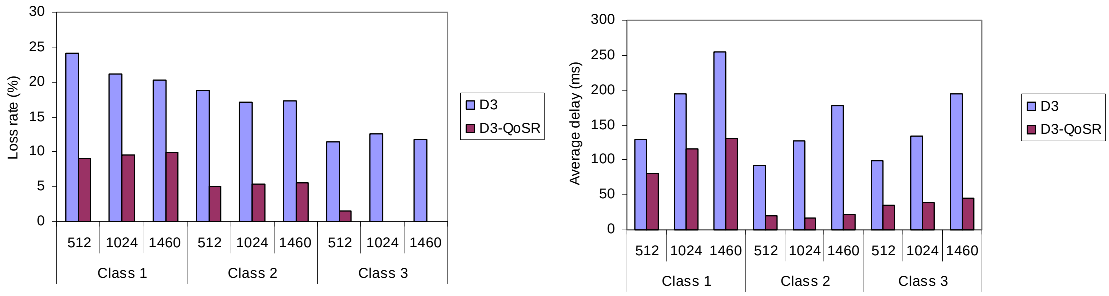

Figure 6 Impact of QoS routing in queue management with congestion.

Even though QoSR improves traffic performance, it has a communication cost due to the additional routing messages needed to distribute the state of the network to all the routers and a processing cost due to the frequency of application of the path computation algorithm. However, several mechanisms have been proposed and embedded in QoS routing protocols that are able to limit the impact of these factors [11, 12, 13]. Therefore, if the increase in routing overhead is bounded, and since traffic performance increases by using QoS routing, the extra cost associated with QoS routing should be supported by the network through adequate traffic engineering in order to provide a better service to end-users and to improve network utilization.

4 Conclusions and future work

In this paper a set of queue management and QoS routing mechanisms for class-based traffic differentiation were evaluated by simulation. The first group of queue management mechanisms included the widely deployed active packet drop algorithm RED and the scheduling algorithm WRR. The second group was the D3 system for resource redistribution among traffic classes through the adaptation of the scheduler and packet dropper to different levels of congestion. The results showed that both approaches are able to do traffic differentiation in variable network load conditions. Moreover, the combination RED/WRR has better performance when the network is not congested, while D3 shows better results when the network is congested. The evaluation of the impact of QoS routing used in association with queue management showed that it clearly improves traffic performance by using the available network resources instead of using just the shortest path as in traditional routing protocols.

In future work the evaluation on larger topologies and using more diverse traffic patterns will be addressed. Furthermore, the performance of the LCT-QoSR strategy will be compared with other class-based QoS routing proposals.

5 Acknowledgemnts

This work was supported in part by the Portuguese Ministry of Science and High Education (MCES), and by European Union FEDER under program POSI (Project QoS II and IPQoS) and under the PhD grant PRAXIS XXI/ BD/13723/97 and E-Next 6º FP NoE.

6 References

1. R. Braden, D. Clark, S. Shenker, “Integrated Services in the Internet Architecture: an Overview”, RFC 1633, Internet Engineering Task Force, June 1994.
2. S. Blake, D. Black, M. Carlson, E. Davies Nortel, W. Weiss, “An Architecture for Differentiated Services”, RFC 2475, Internet Engineering Task Force, December 1998.
3. S. Floyd, V. Jacobson, “Link-sharing and Resource Management Models for Packet Networks”, IEEE/ACM Transactions on Networking, Vol. 3 No. 4, pp. 365-386, August 1995.
4. W. Stallings, “High-Speed Networks and Internets – Performance and Quality of Service”, 2nd Edition, Prentice Hall, 2002.
5. S. Floyd and V. Jacobson, “Random Early Detection gateways for Congestion Avoidance”, IEEE/ACM Transactions on Networking, vol.1 n.4, pp. 397-413, August 1993.
6. G. Quadros, A. Alves, E. Monteiro, F. Boavida, “An Approach to Support Traffic Classes in IP Networks”, Proceedings of the first International Workshop on Quality of future Internet Services (QofIS 2000), Berlin, Germany, September 25-26, 2000.
7. J. Moy, “OSPF Version 2”, RFC 2328, Internet Engineering Task Force, April 1998.
8. 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, Boston, MA, USA, November 5-8, 2000.
9. M. May, J. Bolot, C. Diot, B. Lyles, “Reasons not to deploy RED”, Technical Report, June 1999.
10. M. Christiansen, K. Jeffay, D. Ott, F. Smith, “Tuning RED for Web Traffic”, ACM SIGCOMM, August 2000.
11. 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.
12 G. Apostolopoulos, R. Guerin, S. Kamat, and S. Tripathi. "Quality of Service Based Routing: A Performance Perspective", Proceedings of SIGCOMM'98, Vancouver, British Columbia, USA, September 1998.
13. M. Curado, O. Reis, J. Brito, G. Quadros, E. Monteiro, “Stability and scalability issues in Hop-By-Hop Class-based Routing”, Proceedings of the 2nd International Workshop on QoS in Multiservice IP Networks (QoS-IP 2003), Milano, Italy, February 24-26, 2003.