

July 2000

Performance analysis of QoS mechanisms in IP networks

D. Jia

University of Wollongong

E. Dutkiewicz

University of Wollongong, eryk@uow.edu.au

Joe F. Chicharo

University of Wollongong, chicharo@uow.edu.au

Follow this and additional works at: <https://ro.uow.edu.au/infopapers>



Part of the [Physical Sciences and Mathematics Commons](#)

Recommended Citation

Jia, D.; Dutkiewicz, E.; and Chicharo, Joe F.: Performance analysis of QoS mechanisms in IP networks 2000.

<https://ro.uow.edu.au/infopapers/215>

Performance analysis of QoS mechanisms in IP networks

Abstract

Integrated services IP networks are expected to provide a variety of services with differentiated QoS. This requires the implementation of mechanisms that can discriminate service classes in terms of QoS. The IETF has recently proposed a differentiated services (Diffserv) framework for provision of QoS. In this paper we analyse the performance of two Diffserv mechanisms: threshold dropping and priority scheduling in terms of packet loss and mean packet delay. A comparison of the two mechanisms is carried out with the requirement that both mechanisms provide the same level of packet loss for the preferred flow. This comparison extends the results reported in the literature for these two mechanisms. In particular, in this paper we determine the impact of buffer threshold and buffer size on packet loss and mean packet delay in these mechanisms.

Disciplines

Physical Sciences and Mathematics

Publication Details

This paper originally appeared as: Jia, D, Dutkiewicz, E & Chicharo, JF, Performance analysis of QoS mechanisms in IP networks, Proceedings. Fifth IEEE Symposium on Computers and Communications, 3-6 July 2000, 359-363. Copyright IEEE 2000.

Performance Analysis of QoS Mechanisms in IP Networks

Dix Jia, Eryk Dutkiewicz¹, Joe F. Chicharo
The Switch Network Research Centre
University of Wollongong, Australia
dix@snrc.uow.edu.au

Abstract

Integrated services IP networks are expected to provide a variety of services with differentiated QoS. This requires the implementation of mechanisms that can discriminate service classes in terms of QoS. The IETF has recently proposed a Differentiated Services (Diffserv) framework for provision of QoS. In this paper we analyse performance of two Diffserv mechanisms: Threshold Dropping and Priority Scheduling in terms of packet loss and mean packet delay. A comparison of the two mechanisms is carried out with the requirement that both mechanisms provide the same level of packet loss for the preferred flow. This comparison extends the results reported in the literature for these two mechanisms. In particular, in this paper we determine the impact of buffer threshold and buffer size on packet loss and mean packet delay in these mechanisms.

Keywords— Diffserv, QoS, Threshold Dropping, Priority Scheduling.

1. Introduction

Rapid growth of new applications and the need for differentiated Quality of Service (QoS) has increased the demand for better performance and flexibility of the Internet to support both existing and emerging applications. The current Internet offers best effort service to all users and is inadequate for those applications with more stringent QoS requirements. Differentiated Services (Diffserv) framework has been proposed by the IETF [6][7][8][9]. In Diffserv, packets are tagged with different priorities according to their service classes. Service differentiation is achieved when packets are processed and forwarded by Diffserv mechanisms according to

packets' priorities. Efficient support of different QoS services, however, may require the implementation of different QoS mechanisms in different parts of a network.

A number of QoS mechanisms have been proposed in literature including Threshold Dropping (TD) [8], Priority Scheduling (PS) [9], Random Early Detection (RED) [11], RED with In and Out profile packets (RIO) [3] and Weighted Fair Queuing (WFQ) [1][2][10]. TD and PS can be regarded as basic mechanisms from which the other mechanisms have been derived. Hence comparative performance of these two mechanisms in providing required QoS is an important issue. The results can be used to choose the appropriate mechanism to provide the required QoS for particular applications in the most efficient manner. The above mechanisms have been analysed in the literature to a certain extent. These include the analysis of RIO in [4] and WFQ in [1] and TD and PS in [5]. However, the important issue of how to engineer these mechanisms for optimal performance still needs to be tackled. In this paper we carry out a performance comparison of the TD and PS mechanisms with the aim of providing the same level of packet loss to the preferred flow. Our comparison allows us to determine resultant packet loss for the non-preferred flow and mean packet delay for both the preferred and non-preferred flows as a function of various parameters of the two mechanisms.

The paper is structured as follows. Section 2 briefly describes the operation of the TD and PS mechanisms. Section 3 presents a performance comparison of the mechanisms in terms of packet loss and mean packet delay. The impact of the threshold setting and buffer partitioning on the relative performance of the two mechanisms is also examined in this section. Section 4 concludes the paper.

2. Overview of TD and PS Mechanisms

¹ Now with Motorola ARC Sydney. This paper represents work done while still with the Switched Networks Research Centre

2.1. Threshold Dropping

A threshold dropping mechanism is depicted in Figure 1. Two arrival flows are considered: preferred flow and non-preferred flow. The preferred flow consists of packets which are tagged in profile (i.e. which do not violate their traffic contract) and the non-preferred flow consists of packets which are tagged out of profile. Preferred flow should receive preferential treatment with respect to the non-preferred flow. This is achieved in the TD mechanism by setting a threshold S . Non-preferred flow packets which arrive to the system when the queue length exceeds S are dropped. On the other hand preferred flow packets are only dropped when the queue length reaches the buffer size M .




Figure 1. Threshold dropping mechanism with two packet flows

Figures 2 and 3 show simulation results for the TD mechanism under various load and threshold conditions. These results were obtained assuming that preferred and non-preferred flows were Poisson with mean arrival rate λ_1 and λ_2 , respectively. Packet service time was assumed to be exponential. The mean packet delay is normalised with respect to service time. No flow control and packet re-transmission were considered




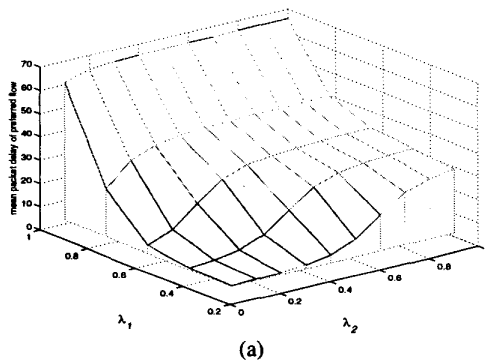


Figure 2. Loss and delay behaviors of TD mechanism under various load from both flows. (Buffer settings: $M=100$, $S=30$).

Figure 2 shows packet loss and mean packet delay as a function of λ_1 and λ_2 (normalised with respect to μ). In this figure the buffer size was set to $M = 100$ and the threshold was set to $S = 30$. As expected, increasing the load of the non-preferred flow has little effect on packet loss experienced by the preferred flow. The mean packet delays of both flows are bounded by their thresholds.

Figure 3 shows the impact of threshold S on packet loss and mean packet delay of the preferred and non-preferred flows. In this figure both flows had a fixed load of 0.7, the total buffer size was set to $M = 100$ and the threshold value S was varied from 10 to 90. Under the above conditions increasing the threshold value results in

little improvement in packet loss of the non-preferred flow. However, packet loss of the preferred flow increases sharply as the threshold is increased beyond approximately 40. Increasing the threshold leads to a linear increase in the mean packet delay for both flows.

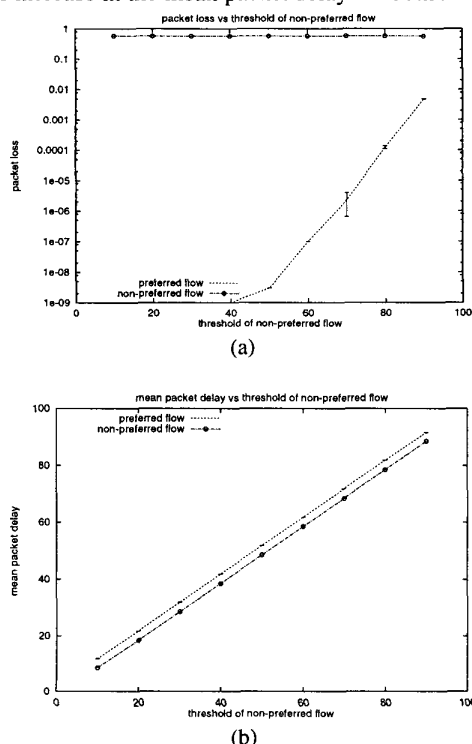


Figure 3. Impact of threshold of non-preferred flow on packet delay and loss

2.2. Priority Scheduling

A priority scheduling mechanism handling to packet flows is depicted in Figure 4. Packets belonging to the preferred flow receive non-preemptive priority over packets belonging to the non-preferred flow. Buffer sizes for the preferred and non-preferred flows are set to K and L , respectively.




Figure 4. Priority Finite Queues

Figure 5 shows simulation results for packet loss and mean packet delay experienced by the preferred and non-preferred flows in the PS mechanism as a function of the buffer size L allocated to the non-preferred flow. The total

buffer size ($K+L$) was set to 15 and preferred and non-preferred flows were Poisson with mean arrival rate λ_1 and λ_2 , respectively. Packet service time was assumed to be exponential. The mean packet delay is normalized with respect to service time. No flow control and packet retransmission were considered.

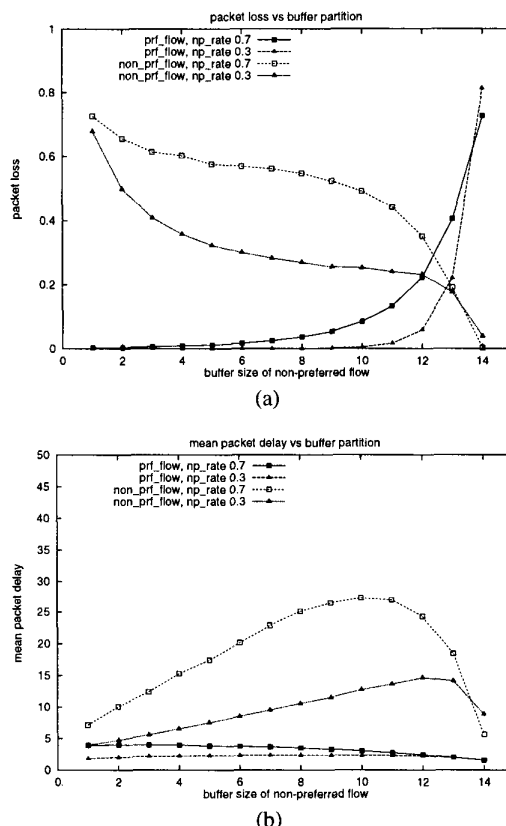


Figure 5. Packet loss and mean packet delay vs buffer partition for various of np_rate (λ_2). Normalized arrival rate of preferred flow (λ_1) is 0.7.

Figure 5 shows a clear trade-off between packet loss and mean packet delay when the buffer allocation is changed. Mean packet delay curves for non-preferred flow show interesting behavior when buffer space allocated to non-preferred traffic is varied. The mean packet delay for non-preferred flow is small when the buffer space allocation is either small (less than 2) or large (more than 12). This is because when the allocated buffer size is small, the mean delay is bounded by the small buffer size. When more buffer space is allocated to non-preferred flow, however, the buffer space left for preferred flow will be decreased due to the constant total buffer size. Under this scenario, packets from the non-preferred flow will spend less time waiting for the queue of the preferred flow to become empty. This behavior is due to the fact that we ignore

packet re-transmission in our simulation and only consider the mean delay of those packets which were not dropped from the queue.

3. Performance Comparison of TD and PS Mechanisms

In this section we present the results of a number of simulations carried out to obtain relative performance of the two mechanisms. We set the two mechanisms with the same total buffer space of 15 packets and the same link capacity (normalized to 1). As in earlier tests the preferred and non-preferred flows were modeled as Poisson processes. For given arrival rates of both flows, we varied the threshold S in the TD mechanism and the buffer size K in the PS mechanism until the same level of loss probability for the preferred flow was obtained from both mechanisms. We then compared the resulting packet loss of the non-preferred flow and the mean packet delay of both flows between these two mechanisms. The packet loss and mean packet delay results are shown in Figure 6 and Figure 7, respectively. The mean packet delay is normalized with respect to service time. Normalized arrival rate of non-preferred flow in both figures is 0.7.




Figure 6. Packet loss Comparison




Figure 7. Mean Packet Delay Comparison

The results of Figure 6 indicate that the TD mechanism has better performance in terms of packet loss for the non-

preferred flow when the load of the preferred flow is light. When the load is heavy the difference in packet loss between the two mechanisms is negligible. The results of Figure 7 indicate that as the load of the preferred flow changes, the PS mechanism provides a smaller mean delay to the preferred flow than does the TD mechanism. However, the TD mechanism results in a smaller mean delay for the non-preferred flow.

4. Conclusion

Threshold dropping (TD) and priority scheduling (PS) are two fundamental mechanisms that can provide the ability to discriminate between QoS of traffic classes in DiffServ. Our performance investigation of the TD mechanism indicated that changing the load of the non-preferred flow has a minimal effect on packet loss of the preferred flow. With a fixed total buffer size and the same arrival rate of both flows, there is a minimal improvement in loss for the non-preferred flow when its threshold is increased. The mean packet delays for both flows are bounded by their thresholds. A clear trade-off between packet loss and mean packet delay for the preferred and non-preferred flows is observed in the PS mechanism when the buffer allocation is changed. The PS mechanism has the advantage over the TD mechanism in providing a lower mean delay to the preferred flow when the two mechanisms are engineered so as to provide the same level of packet loss for the preferred flow. However, under the same scenario, the TD mechanism provides lower packet loss and mean packet delay to the non-preferred flow.

5. References

- [1] J.C.R. Bennett and H. Zhang, "Why WFQ Is Not Good Enough for Integrated Services Networks?", *Proceedings of NOSSDAV'96*, Apr. 1996.
- [2] J.C.R. Bennett and H. Zhang, "Hierarchical Packet Fair Queuing Algorithms," *Proceedings, SIGCOMM '96*, Palo Alto, California, August 1996.
- [3] Clark, D., Fang, W., "Explicit Allocation of Best Effort Packet Delivery Service", *Internet draft*, Sept. 1997
- [4] Martin May, J.C. Bolot, Alain J. Marie and C. Diot, "Simple Performance Models of Differentiated Services Schemes for the Internet", *Proceedings of INFOCOM'99*, New York, March 1999
- [5] Sambit Sahu, Don Towsley, Jim Kurose, "A Quantitative Study of Differentiated Services for Internet", *Umass CMPSCI Technical Report 99-09*.
- [6] S. Blake, D. Black, M. Davies, Z. Wang, W. Ewiss, "An Architecture for Differentiated Services", *RFC 2475*, December 1998.
- [7] Y. Bernet, J. Biner, S. Blake, M. Carlson, B. E. Carpenter, E. Davies, B. Ohlman, S. Keshav, D. Verma, Z. Wang, W. Weiss, "A Framework for Differentiated Services", *Internet Draft*, February 1999.

- [8] D.D. Clark, J. Wroclaski, "An Approach to Service Allocation in the Internet", *IETF Draft*, July 1997.
- [9] K. Nicholas, V. Jacobson, L. Zhang, "A Two-bit Differentiated Services Architecture for the Internet", *IETF Draft*, Nov. 1997.
- [10] Hui Zhang, "Service Disciplines for Guaranteed Performance Service in Packet-Switching Networks", *Proceedings of The IEEE*, Vol. 83, NO. 10, October 1995.
- [11] S. Floyd, V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance", *IEEE/ACM Trans. Networking*, 1(4):397-413, August 1993.