

Applying the DiffServ Model to a Resilient Packet Ring Network

Fredrik Davik^{1,2,3} and Stein Gjessing¹

¹ Simula Research Laboratory

² University of Oslo

³ Ericsson Research Norway

{bjornfd,steing}@simula.no

Keywords: Resilient Packet Ring, Differentiated Services, Per Hop Behavior.

1 Introduction

In this paper we introduce a formal specification of parts of the service differentiation mechanisms of the recent IEEE 802.17 Resilient Packet Ring (*RPR*) standard [1,2,3], and assess RPR's suitability for use in a DiffServ environment. We propose a simple mapping between RPR's traffic classes and three standardized DiffServ Per Hop Behavior groups. When using this mapping, we discuss, by use of an analytical example and simulation results, the behavior of the traffic assigned to each PHB group in terms of its throughput (all PHBs) and delay properties (for the EF PHB only). For the simulation part, we use our OPNET [4] implementation of the RPR standard. In the full length version of this paper, for the proposed mapping, we analyze conformance to the requirements specified for the three DiffServ PHBs [5].

2 Delay Guarantees and Rate Control in RPR

In an RPR node, the *Rate Control Block* imposes several per ringlet and per node rate constraints on local ingress (add)- and transit traffic. In this chapter, we introduce a set of invariants that expresses the effect of the RPR scheduling constraints, before illustrating this using an analytical example. We use the notation R_X to specify the rate constraint in effect for a particular RPR traffic class, where $X \in \{A, A0, A1, B, C\}$. *offered*(X) represents the amount of traffic of a particular traffic class, X , that a node or a set of nodes **want**(s) to transmit. *accepted*(X) refers to the corresponding amount of traffic that **will** be transmitted. R_L is the link rate (all links in an RPR network have the same rate).

For class A ($A0$ and $A1$) traffic, the per-node amount of traffic is constrained, regardless of the destination of the traffic transmitted, as shown in invariants 1 and 3-6. For class B and class C traffic, we use per link invariants as shown in invariants 7-10.

invariant 1 $R_R = \sum_{j \in \{nodes\}} R_{A0j} \leq R_L$

invariant 2 $R_U = R_L - R_R$

invariant 3 $\sum_{j \in \{nodes\}} (R_{A1_j} + R_{B_j}) \leq R_U$

invariant 4 $\forall nodes : offered(A) \leq R_{A0} \Rightarrow accepted(A0) = offered(A)$

invariant 5 $\forall nodes : R_{A0} < offered(A) \leq R_{A0} + R_{A1} \Rightarrow$
 $accepted(A0) = R_{A0} \wedge accepted(A1) = offered(A) - R_{A0}$

invariant 6 $\forall nodes : offered(A) > R_{A0} + R_{A1} \Rightarrow$
 $accepted(A0) = R_{A0} \wedge accepted(A1) = R_{A1}$

invariant 7 $\forall links : offered(B) + accepted(A1) \leq R_U \Rightarrow accepted(B) = offered(B)$

invariant 8 $\forall links : offered(B) + accepted(A1) > R_U \Rightarrow$
 $accepted(B) = R_U - accepted(A1)$

invariant 9 $\forall links : offered(C) + accepted(A1) + accepted(B) \leq R_U \Rightarrow$
 $accepted(C) = offered(C)$

invariant 10 $\forall links : offered(C) + accepted(A1) + accepted(B) > R_U \Rightarrow$
 $accepted(C) = R_U - accepted(A1) - accepted(B)$

The DiffServ Expedited Forwarding (EF) PHB is a DiffServ building block to be used for the provisioning of services that provides low loss, low delay and low jitter. Another DiffServ PHB, namely the Assured Forwarding (AF) PHB is a specification of a PHB group, that may contain up to four independent AF classes. Each class must be allocated a separate amount of forwarding resources. Within each AF class, an implementation must provide a minimum of two and a maximum of four different drop probabilities. The DiffServ default PHB is specified as a best-effort (BE) forwarding behavior. For the remainder of the paper, we refer to the default PHB as the BE PHB.

Given the detailed requirements associated to the EF, AF and BE PHBs discussed in detail in [5], and the properties of RPR's traffic classes A , B and C as defined by the invariants; we propose that the DiffServ EF, AF and BE PHBs are mapped to respectively RPR traffic classes A , B and C . In figure 1, we have shown the analytical resulting throughput when using this mapping and enforcing the invariants introduced above. The figure shows the analytical




Fig. 1: Theoretical Throughput for the EF, AF and BE PHBs on a single link, in an RPR network.

result, using the offered load as the free variable with a 55/27/18 ratio for respectively BE, AF and EF offered traffic, and plotting the corresponding values for accepted EF, AF and BE traffic as well as the aggregate of these values. Three points for offered load in this plot are of particular interest. As long as we are on the left side of point 1, the aggregate EF traffic can be increased further while maintaining invariant 4 or 5. At point 1, we have reached the load level where an additional load increase causes the premise of both invariants 4 and 5 to be false. Beyond this point, the maintenance of invariant 6 effectively prohibits the acceptance of additional EF traffic. On the left side of point 2, as long as invariant 9 is maintained, an increase in offered BE traffic leads to an equivalent increase in accepted BE traffic. At point 2, we have reached the load level where an additional load increase causes the invariant's premise to be false. Beyond this point, the maintenance of invariant 10 effectively prohibits the acceptance of more BE traffic. Finally, between points 2 and 3, as long as invariants 7 and 10 are maintained, an increase in offered AF traffic leads to an equivalent increase in accepted AF traffic, on the expense of an equivalent decrease in accepted BE traffic. At point 3, we have reached the load level where an additional load increase causes the premise of invariant 7 to be false. Beyond this point, the maintenance of invariants 8 and 10 effectively prohibits the acceptance of additional AF traffic as well as excluding BE traffic from the link.

3 Performance Evaluation by Simulations

We present results obtained for a 16 node ring, where the link delay and capacity is respectively $250 \mu s$ and 1Gbit/s for all links (i.e. a 800km metro ring). For each simulation, an offered load parameter is given. 101 load values are used, linearly distributed on a scale from 1 to 10 times the base load. For each load value, 16 simulations are executed with different sets of source destination pairs. The base load uses the same ratio between the PHBs as specified in section 2. All nodes transmit traffic of each PHB group using a Poisson distribution. Figure 2a plots the offered vs. accepted load for the EF, AF and BE PHBs, as well as the aggregate of all PHBs. Each point on the curves is the mean of the throughput values obtained from the 16 different sender-receiver simulations for that particular offered load. The vertical lines in the figure corresponds roughly to the points discussed for the analytical results in the previous section. When looking at the throughput-performance of the ring, the behavior is not as clear-cut as for the theoretical behavior, illustrated in figure 1. In a real (or simulated) RPR network, the individual links become congested at different points in time and at various levels of per-node offered load. Thus for a given sender-receiver pair setup, the effect of the invariants discussed in the previous section, will occur at different values of per node offered load for the various congested links on the ring. The effect of the invariants is additionally concealed as we take the mean of the values obtained from a number of different sender-receiver pair configurations. However, the general trend from figure 2a clearly resembles that of the analytical example. Figure 2b show the EF-traffic delay statistics broken

down on a per-hop level. The mean values are in the range $[0.6, 5.3]\mu s$, which is well within the transmission time of two 500B packets. The standard deviation is very low (less than transmission time of a 500B packet). As expected, the median, mean, standard deviation and max values all increase as the number of hops increase for the samples observed. But as shown, the jitter is very low and in the worst case within $32.9 \mu s$.

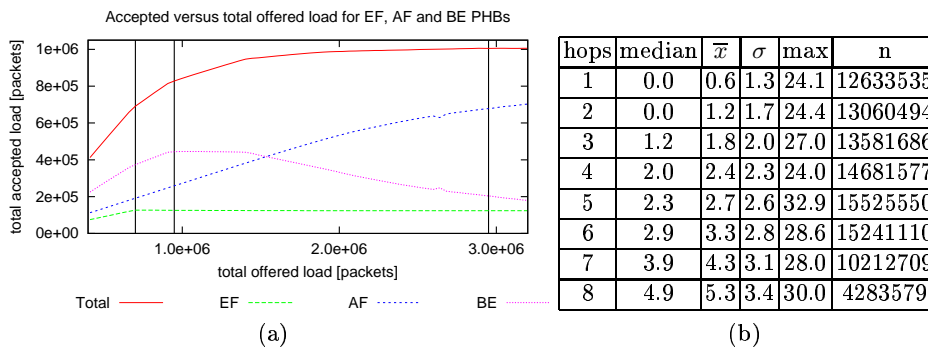


Fig. 2: a) Accepted vs. total offered load for EF, AF and BE PHBs. b) EF traffic delay statistics (in units of μs). Per link propagation and transmission delays are excluded. The min value is 0 for all hop counts.

4 Conclusion

In this short paper, we have proposed and evaluated, analytically and by simulations, a mapping between the three DiffServ PHBs: Expedited Forwarding (EF), Assured Forwarding (AF) and BE (Default) and the RPR traffic classes: *A*, *B* and *C*. Our results indicate, that the proposed mapping scheme appear reasonable and that conformance to DiffServ's PHB requirements appears to be possible.

References

1. IEEE Computer Society: IEEE Std 802.17-2004 (2004) 1
2. Davik, F., Yilmaz, M., Gjessing, S., Uzun, N.: IEEE 802.17 Resilient Packet Ring Tutorial. IEEE Commun. Mag. **42** (2004) 112–118 1
3. Gambiroza, V., Yuan, P., Balzano, L., Liu, Y., Sheafor, S., Knightly, E.: Design, analysis, and implementation of DVSR: a fair high-performance protocol for packet rings. IEEE/ACM Trans. Networking **12** (2004) 85–102 1
4. : (OPNET Modeler. <http://www.opnet.com>) 1
5. Davik, F., Gjessing, S.: Applying the DiffServ Model to a Resilient Packet Ring Network. Technical Report 2005-1, Simula Research Laboratory (2005) <http://www.simula.no/publication.php>. 1, 2