Computer Networks 62 (2014) 43–54

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

# Achieving destination differentiation in ingress aggregated fairness for resilient packet rings by weighted destination based fair dropping

Mete Yilmaz [a,*], Nirwan Ansari [b]

[a] *Edge Routing Business Unit, Cisco, CA 95134, USA*
[b] *Department of Electrical and Computer Engineering, New Jersey Institute of Technology, NJ 07102, USA*

A R T I C L E   I N F O

*Article history:*
Received 29 July 2012
Received in revised form 16 December 2013
Accepted 6 January 2014
Available online 10 January 2014

*Keywords:*
MAN
WAN
Ring networks
Spatial reuse
Fairness
Active queue management

A B S T R A C T

The IEEE 802.17 is a standardized ring topology network architecture, called the Resilient Packet Ring (RPR), to be used mainly in metropolitan and wide area networks. This paper introduces destination differentiation in ingress aggregated fairness for RPR and focuses on the RPR MAC client implementation of the IEEE 802.17 RPR MAC in the aggressive mode of operation. It also introduces an enhanced active queue management scheme for ring networks that achieves destination differentiation as well as higher overall utilization of the ring bandwidth with simpler and less expensive implementation than the generic implementation provided in the standard. The enhanced scheme introduced in this paper provides performance comparable to the per destination queuing implementation, which is the best achievable performance, while providing weighted destination based fairness as well as weighted ingress aggregated fairness. In addition, the proposed scheme has been demonstrated via extensive simulations to provide improved stability and fairness with respect to different packet arrival rates as compared to earlier algorithms.



## 1. Introduction

The IEEE 802.17 Resilient Packet Ring (RPR) is a standardized network architecture based on ring topology deliberated by the IEEE LAN/MAN Standards Committee [1]. The ring is established by using point-to-point bidirectional connections between stations in RPR. The RPR network has two counter-rotating rings called ringlets. The RPR protocol utilizes both ringlets to overcome connection disturbances, and hence provides resilience. In addition, the RPR protocol provides a fairness mechanism among stations in terms of sharing the overall ring bandwidth applied to the best effort traffic.

The ring network is simpler to operate and administer than a complex mesh or an irregular network while virtually providing connection from each station to every other station [2]. SONET based ring networks have been already deployed by service providers in MANs or WANs. However, in these networks, many SONET rings consist of a dual-ring configuration in which one of the rings is used as the back-up ring and remains unused during normal operation to be utilized only in the case of a failure of the primary ring [2]. The static bandwidth allocation and network monitoring requirements increase the total cost of a SONET network. While plain Ethernet does not require static allocation and provides cost advantages, it cannot provide desired features such as fairness and auto-restoration [2]. In addition, the RPR standard is a media independent standard which can be used on top of SONET or Ethernet network to obtain the above advantages.

* Corresponding author. Tel.: +1 4088530767.
*E-mail addresses:* myilmaz@cisco.com (M. Yilmaz), nirwan.ansari@njit.edu (N. Ansari).

1389-1286/$ - see front matter © 2014 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.bjp.2014.01.003

The fairness aspects of RPR have been investigated in depth in light of interesting scenarios as those described in [3–5]. Improvements for the current fairness algorithm of the IEEE 802.17 have also been proposed as reported in [3,4,6–8]. In addition, in our earlier work we introduced weighted fairness and its use cases in [9,10]. The basic algorithm relies on identifying a single congestion point on the ring. This decreases the ring utilization and spatial reuse in certain traffic scenarios. The standard provides means to keep track of all the congested downstream nodes and their fair rates. A generic way to utilize this mechanism fully is to implement 255 separate queues (one queue per destination, including the multicast queue) within the MAC client. Supporting multiple destination queues at the MAC layer is not efficient and requires additional memory or complex shared memory queuing implementations. In addition, this may not be even possible to implement on the already deployed hardware. This paper presents a simpler and less expensive MAC client implementation that does not require separate queues while providing ring utilization comparable to the generic implementation. Our previous work [11] showed that it is possible to utilize an active queue management scheme at the MAC client level. However, the implementation that will be presented in this article achieves destination differentiation at the source stations by allowing different weights per destination while achieving RPR ingress aggregated fairness. In addition, the proposed implementation provides a higher degree of stability under different packet arrival rates and patterns. This is achieved by penalizing the flows which exceed their fair shares more than the others.

The rest of this paper is organized as follows. In Section 2, the RPR operation and fairness algorithm will be discussed with emphasis on MAC client implementations. Section 3 will discuss the MAC client implementation utilizing our proposed active queue management scheme based on the approximate fair dropping algorithm [12]. In Section 4, behavior of the current (aggressive) fairness algorithm with MAC client implementations of single queue, multiple queues, and two other algorithms with active queue management of a single queue will be illustrated through simulations. Section 5 provides simulation results for the modified scenario that provides destination differentiation. In Section 6, comparison of a large ring hub scenario will be provided. Finally, the conclusion will be drawn in Section 7.

In order to demonstrate different operational modes, some performance figures of merits are included and discussed. The scenarios have been executed on the RPR simulator model developed at Cisco during the RPR standardization process. The simulation model is implemented in Opnet. The suggested modification has been incorporated into the simulator model and its behavior is verified through simulations.

## 2. RPR operation and fairness

The operation of the RPR protocol has been discussed in detail in [2]. In this paper, a brief overview based on an example will be provided.

An RPR network can have a maximum of 255 stations. Each station on the ring removes a frame from the ringlet if that frame is destined for that station (except multicast frames, which are removed by the sender). This facilitates the spatial reuse property of an RPR network since the bandwidth will not be consumed by that frame around the ringlet.

Fig. 1 shows an example scenario. Stations *4* and *3* are transmitting to Station *2* on the inner ringlet, while Stations *2* and *1* are transmitting to Station *7*. In addition Station *2* has traffic destined to Station *1*. Each time a frame is received at Station *2* from the ring, the frame will be checked to see if that frame is destined to the station itself or not. The frame will be removed from the ring if it is destined to the station (in this case, any frame from Station *3* or *4*). In parallel, Station *1* will receive frames from Station *2*, and Station *7* will be able to receive the frames from Stations *1* and *2* without being impacted by the traffic generated at Stations *3* and *4*.

The objective of the fairness algorithm is to distribute the unallocated bandwidth around the ring among stations in a fair manner. In the case of Fig. 1 (assuming that all the stations have equal weights), Stations *3* and *4* will get an equal amount of the link bandwidth on the link between Stations *2* and *3*, while Stations *1* and *2* will get an equal amount of the link bandwidth on the link between Stations *1* and *7*. More detailed explanation of the standardized fairness algorithms can be found in [1,2].

In short, the fairness algorithm provides a fair sharing of the link bandwidth according to the weights of the stations when there is more traffic than that can be transmitted through that link. RPR fairness is based on ingress aggregation. This fairness is referred to as "Ring Ingress Aggregated with Spatial Reuse" (RIAS) fairness in an earlier article [3]. This definition follows the same methodology used in [13] for max–min flow control. The RIAS fairness definition, however, does not include the station weights in the generalized formula while this is included in the IEEE 802.17 standard in the calculation of the estimated fair rate of a node. In addition, the RIAS definition assumes equal sharing of the bandwidth among flows originating from the same station, while the standard does not require that. In our previous works [9,10], we have expanded this definition with the inclusion of source station weights. In this section, a more general definition will be provided with the inclusion of the destination station weights along with

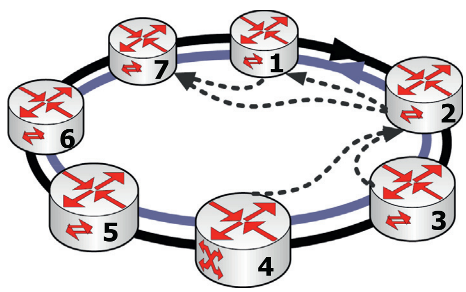


**Fig. 1.** Destination stripping and spatial reuse illustrated on the inner ring (ringlet 1).