Inter-ring Traffic Management in Bridged Resilient Packet Rings: Global Fairness and Buffer Overflow Prevention

Pisai Setthawong ***Surat Tanterdtid***

Department of Telecommunications Science, Faculty of Science and Technology, Assumption University, Thailand

Summary
Resilient Packet Ring (RPR) is a dual-ring network, also known as the IEEE 802.17 Standard. As with other IEEE 802 networks, multiple RPR networks can be bridged together to form a bridged network when necessary. However, further research is necessary on additional issues that arise from bridging RPR networks. In this paper, we place emphasize on two of these issues; that is, fairness for inter-ring traffic and buffer overflow prevention at the bridges. We propose a global fairness reference model as the benchmark for global fairness of inter-ring traffic. We then propose a global fairness algorithm for implementing such a global fairness. We also propose how we can use the same algorithm to minimize the occurrence of buffer overflow at the bridges. Simulations were performed to evaluate the algorithms using various network scenarios and traffic patterns. The results show that the proposed algorithm can achieve the intended results of ensuring global fairness and minimizing buffer overflow in bridged RPR networks.

Key words:
Resilient Packet Ring, Bridging, Inter-ring Traffic, Fairness Definition, Fairness Algorithm, Buffer Overflow

1. Introduction

Resilient Packet Ring (RPR) is a packet-based high-speed dual-ring network that can be used for implementing local area networks (LAN) and metropolitan area networks (MAN). It support rates up to many gigabits per second, as well as increased bandwidth availability, up to an average four-fold increase over the link capacity, through spatial reuse and the use of both ringlets for data transmission. RPR has been approved as the IEEE 802.17 Standard since June 2004 [1].

RPR is scalable by bridging multiple RPR rings together to form a bridged network, as is possible with other IEEE 802 networks. However, the IEEE 802.17 Standard requires that all packets whose source and destination nodes are on different rings must be flooded on all the connected rings in the bridged network. This results in a significant drop in bandwidth efficiency. Most of the research on bridged RPR networks, including that being discussed within the IEEE 802.17b Working Group, has so far focused on how to avoid this flooding of packets.

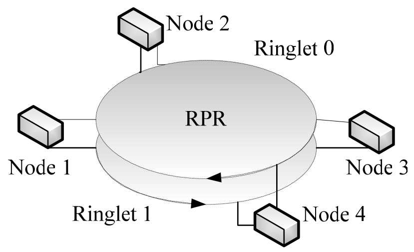

Fig. 1 RPR Network with 4 nodes. Both ringlets are used for data transmission. The packets in each ringlet flow in the opposite direction. Each packet is transmitted on the ringlet with the shortest path to its destination.

While this issue is critical, in this paper, we will raise other issues regarding bridged RPR networks that are also equally worth considering. As with any bridging solution, the bottleneck for inter-ring traffic, or traffic whose source and destination nodes are on different rings, is at the bridges that inter-connect the rings together. Without efficient bandwidth management, it is possible for the packet queues at the bridge to overflow and cause packets to be dropped in transit. In addition, there is no guarantee of fairness for the inter-ring packets. As a result, some nodes might use an unfair share of the available inter-ring traffic bandwidth at the expense of other nodes.

The remaining of this paper will explore further in detail the reasons why the issues stated in the previous paragraph are important and hence should be topics for further research. In addition, solutions to these issues are presented, supported by simulation results.

2. Background

In order to better comprehend the issues that are the focus in this paper, it is first necessary to understand the concept of fairness algorithm and bridging in RPR.

2.1 Fairness Algorithm

Because of spatial reuse in RPR, the destination node is responsible for removing the packets from the ring, instead of the source node. This frees sections of the ring up for other nodes to utilize at the same time. The benefit of this is that the available bandwidth in the ring exceeds the link capacity.

However, spatial reuse causes unfairness whenever congestion occurs on a link. Congestion occurs when a node attempts to transmit a packet on to the ring, but there are currently packets in transit through the node, being transmitted from the upstream nodes. In this case, priority is given to the packets in transit and the node has to queue its packet, in order to prevent packets from being dropped in transit. The result is that the upstream nodes will always have priority over its downstream nodes and its packets can thus monopolize the ring.

To prevent this from occurring, the RPR standard specifies two fairness algorithms - namely, the aggressive fairness algorithm and the conservative fairness algorithm [1]. Both algorithms are capable of ensuring fairness by limiting the transmit rate of the upstream source nodes. As a result, this will make bandwidth available to the downstream nodes as well and all the nodes will get to transmit their fair share.

Note that the terminology *fair* is used, instead of *equal*. The reason is because there are many definitions of fairness and each ring may determine its own definition. For this paper, we will consider only equal fairness, where all nodes are assumed equal, although the results may be extended to other definitions of fairness, such as weighted fairness.

These fairness algorithms have been shown to be effective in ensuring fairness within a single ring for intra-ring traffic [3] [4]. Other fairness algorithms [5] [6] [7], have also been proposed to enhance upon the standard fairness algorithms. However, there has not yet been any research into their effectiveness in ensuring fairness across rings for inter-ring traffic. As a result, hereafter, we will refer to these fairness algorithms as *local fairness algorithm* to differentiate them from the fairness algorithm that we will propose later in this paper.

2.2 Bridging

Bridges are special nodes with multiple interfaces that inter-connect between multiple RPR networks [8]. The bridges are responsible for forwarding packets between the connected rings. The Spatially Aware Sublayer (SAS) is the part of the MAC Client layer that is responsible for determining which packets are to be forwarded out of which interfaces. Although the current RPR standard specifies that packets must be flooded on all the interfaces if the destination is not on the connected ring, research on increasing the efficiency of bridging is centered on creating a more efficient forwarding algorithm for the SAS [9] [10]. The proposals in this paper will also be implemented within the SAS.

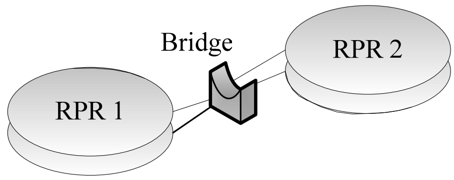

Fig. 2 The bridge connects multiple RPR networks together.

By bridging RPR ring networks, traffic can be divided into *intra-ring traffic*, where the source and destination nodes are on the same ring, and *inter-ring traffic*, where the source and destination nodes are on different rings. Also, we will use the term *ringlet* to refer to each of the dual rings in a RPR network and the term *ring* to refer to each RPR network in a bridged RPR network.

3. Issues with Bridged RPR Networks

Although previous research on bridging has focused on increasing the efficiency in transmitting inter-ring traffic, in the current paper, we would like to emphasis that there are other important issues that should also be focused on; namely, buffer overflow and global fairness.

3.1 Buffer Overflow

Buffer overflow occurs when the sum of all the inbound traffic to be forwarded out of the same outbound interface is larger than the available bandwidth of that outbound interface. In this case, packets must be buffered at the outbound interface to be transmitted later. However, if this condition continues, the buffer will continually fill up until it is full. At this point, it is no longer possible to queue additional packets and packets must then be dropped. This condition is known as *buffer overflow*.

Although buffer overflow is a common concern when bridging any IEEE 802 networks, the possibility of this occurring in bridged RPR networks is more likely than with other IEEE 802 networks.

The first reason is because both ringlets in a RPR network are used for transmission. As a result, it is highly possible that both ringlets will be used to transmit packets from different source nodes to the bridge. This means that the rate of inbound traffic to the bridge may be as high as twice the link capacity. If all of the inbound traffic is to be forwarded out of the same outbound interface, the available bandwidth at that outbound interface may not be sufficient to accommodate all the inbound traffic and buffer overflow will occur. In the worst case, where the rate of inbound traffic to be forwarded out of the same outbound interface is twice the link capacity, the rate of packets dropped at the bridge due to buffer overflow will be equal to the link capacity.

Another reason for the higher possibility of buffer overflow is the use of the local fairness algorithm in the RPR network. All the nodes in a ring must adhere to the local fairness algorithm, in order to ensure fairness in the ring, including the bridges that are connected to that ring. As a result, the available bandwidth for the outbound interfaces of a bridge is limited by the local fair rate of the bridge on that interface. The less the available bandwidth is, the greater the possibility of buffer overflow occurring.

Because of these 2 reasons, there is a greater possibility of buffer overflow occurring in bridged RPR networks when compared to other IEEE 802 networks. However, the current RPR standard does not specify any mechanism for reducing the possibility, or even preventing, buffer overflow from occurring. We would like to emphasize the need for such a mechanism and will introduce one later in this paper.

3.2 Global Fairness

Another issue that is of concern with inter-ring traffic is the fairness among the nodes in transmitting inter-ring packets. With limited bandwidth at the bridges, it is necessary to ensure that the available bandwidth for transmitting inter-ring traffic is fairly shared among the nodes in the bridged RPR network. This is to prevent certain nodes from monopolizing the available bandwidth. To differentiate from local fairness in single RPR network, fairness for inter-ring traffic will be referred to as *global fairness*.

Without global fairness, the nodes will be able to transmit as many inter-ring packets as allowed by the local fairness algorithm. As long as the bridges can accommodate the traffic and buffer overflow does not occur, unfairness will not occur. However, once the rate of traffic becomes greater and buffer overflow starts to occur, the bridge will start to drop packets. Depending on the packets from which source nodes are dropped, the available bandwidth for that source node is determined. Since this randomly occurs whenever the buffer is full, the available bandwidth for each node varies by time and is thus non-deterministic. In general, the larger the number of packets transmitted by a node, the larger the number of packets will be forwarded by the bridge (as well as the larger the number of drop packets). In fact, the available bandwidth for a node is proportional to the transmit rate of that node. In this way, unfairness will occur among the nodes.

Buffer overflow is inter-related with global fairness. With global fairness, the available bandwidth for inter-ring traffic is shared fairly among the nodes. If the nodes limit their transmit rate for inter-ring traffic to the global fair rate, then all the available bandwidth will be used and no more. The rate of inbound traffic at the bridges will not exceed the available bandwidth of the outbound interfaces and thus buffer overflow can be avoided. Because of this, both issues of buffer overflow and global fairness should be considered at the same time.

The remaining of this paper will present a model for defining global fairness, a global fairness algorithm for implementing global fairness and an analysis of how the proposed global fairness algorithm can prevent buffer overflow.

4. Global Fairness Reference Model

First of all, we need to define a global fairness reference model in order to analytically determine global fairness. Through this reference model, we can calculate the global fair rate for each node. We based our work on the accepted definition of local fairness, known as the Ring Ingress Aggregated with Spatial Reuse (RIAS) reference model [6], by making the same assumptions:

- Fairness is considered at the node-link level. Regardless of whether a source node transmits packets to a single destination or to multiple destinations on a particular link, the node will be treated the same and count as a single node-link

- Only equal fairness will be considered. The model can be extended later to include other types of fairness definitions

The global fairness reference model is based on the observation that the bridge is the key to ensuring fairness for inter-ring traffic. The reason is because the bridge is the first node that can identify when unfairness occurs. Therefore, the reference model focuses on the bridge and how the bridge should handle packets that it forwards in order to maintain global fairness. This is outlined in the two rules below:

Rule 1: The rate of traffic forwarded out of an outbound interface onto a ringlet at the bridge has to be less than or equal to the local fair rate of the bridge on that ringlet.

Rule 2: The rate of traffic forwarded out of an outbound interface onto a ringlet at the bridge must be fairly shared by all nodes whose packets are to be forwarded on to that ringlet. That is, a *global fair rate* will be calculated for that outbound interface and the nodes must ensure that they do not transmit beyond the global fair rate. In doing so, all the nodes can transmit their fair share of inter-ring traffic.

The calculation of the global fair rate depends on the number of nodes that are causing unfairness. We use the following notations:

Bw = available bandwidth at the outbound interface
g_n = current calculated global fair rate for iteration n
tr_i = transmit rate of node i on the outbound interface
N_n = number of nodes that transmit above the current calculated global fair rate ($tr_i > g_n$)

The global fair rate is calculated using the following steps:

Step 1: Initialize $g_0 = 0$

Step 2: $g_{n+1} = \left(Bw - \sum_{\forall i \in (tr_i < g_n)} tr_i\right) / N_n$

Step 3: Repeat step 2 until $g_{n+1} = g_n$

The calculated global fair rate is the target rate that the nodes should not exceed when transmitting inter-ring packets through that bridge out of that particular outbound interface.

As an example, we consider the bridged RPR network shown in Fig.2. Assume that there are 2 nodes in RPR ring RPR1 transmitting to the same destination node in RPR ring RPR2. One node would like to transmit at 1/3 the link capacity and another node would like to transmit at 5/6 of the link capacity. The calculation of the global fair rate, based on the global fairness reference model, will result in the global fair rate of 2/3 of the link capacity (g_0=0, g_1=1/2, g_2=2/3 and g_3= g_2 = 2/3).

5. Global Fairness Algorithm

Although the global fairness reference model can be used to calculate the global fair rate, it is only a conceptual model and cannot be applied directly in a bridged RPR network. In order to do this, we need to define a *global fairness algorithm*.

The global fairness algorithm defines the mechanism that the nodes and bridges must follow, in order to achieve global fairness. How well the proposed global fairness algorithm achieves global fairness can be measured by comparing the results of using the algorithm with the calculated global fairness using the global fairness reference model. In addition, the global fairness algorithm must be compatible with the local fairness algorithms. Both global fairness and local fairness must be achievable at the same time.

In this paper, we propose a global fairness algorithm that can be used to achieve global fairness. There are 3 key points in the proposed algorithm: 1) unfairness indicator, 2) global fair rate calculation, and 3) transmission of global fairness packets.

5.1 Unfairness Indicator

Normally, the global fairness algorithm is not invoked; it is only invoked when unfairness is identified for inter-ring traffic at the bridge. We propose two different indicators of unfairness: 1) percentage of buffer filled at outbound interface and 2) receiving of local fairness packets.

When congestion occurs at the bridge for an outbound interface, this means that the rate of inbound traffic to be forwarded out of an outbound interface is larger than the available bandwidth of that outbound interface. As a result, the buffer at the outbound interface will continue to be filled until eventually it is full. Packets will start to be dropped and unfairness will occur. As a result, we can use the percentage of buffer that is filled as an indicator of congestion and likelihood of unfairness occurring; once it is filled beyond a certain threshold level, λ, the bridge will invoke the global fairness algorithm. The threshold level can be expressed as a percentage of the total buffer size, such as 80%.

Another unfairness indicator is when the bridge receives a local fairness packet from the local fairness algorithm on one of its interfaces, with a smaller local fair rate than the current local fair rate for that interface. The reason is because the global fair rate is dependent on the available bandwidth, or the local fair rate. The global fair rate must then be recalculated using the new local fair rate. If the new local fair rate is less than the previous value, there is a possibility that unfairness may occur again because the available bandwidth is reduced. However, if the local fair rate is greater than the previous value, there is no need to recalculate the global fair rate. After the source nodes have not received a global fairness packet for a period of time, and then they will gradually increase their global fair rate automatically, as part of the global fairness algorithm.

5.2 Global Fair Rate Calculation

Once the bridge identifies that unfairness may occur and that the global fairness algorithm should be invoked for a particular outbound interface, the bridge must then calculate the global fair rate of that interface. In order to do this, each bridge needs to maintain the transmit rate of packets, tr_{yz}, from source node y forwarded by the bridge out of outbound interface z.

From this detail, we can calculate the global fair rate using the same recursive method as that used in the global fairness reference model, where l_z is the local fair rate of the outbound interface z:

Step 1: Initialize $g_0 = 0$

Step 2: $g_{n+1} = \left(l_z - \sum_{\forall i \in (tr_{iz} < g_n)} tr_{iz}\right) / N_n$

Step 3: Repeat step 2 until $g_{n+1} = g_n$

5.3 Transmission of Global Fairness Packets

Once the global fair rate has been calculated, the relevant source nodes that are causing the unfairness must be informed through the use of the *global fairness packet*.

Not all source nodes that transmit through the bridge out of outbound interface z are sent a global fairness packet. Only the nodes that transmit over a certain percentage, Δ, of the global fair rate are sent a global fairness packet directly. This is to reduce the control packet overhead, compared to sending the global fairness packet to all the source nodes. A fairly large Δ percentage is recommended, such as 97%, but it should be less than 100%. The reason the value 100% is not recommended is to compensate for the fact that the transmit rate may not be measured accurately. This may result in nodes that are causing unfairness not being sent the global fairness packet because their measured transmit rate is slightly lower than the actual transmit rate. Another reason is that the nodes may gradually increase their transmit rate if they have more packets to transmit than limited by the current global fair rate and they have not received a new global fairness packet within a specified interval. By informing the nodes before their transmit rate exceeds the global fair rate, the nodes can limit their transmit rate beforehand and unfairness can be avoided before it occurs.

Since a node may transmit to multiple destinations and hence its packets may transit through different bridges, a node may receive global fairness packets from different bridges. The node must maintain the global fair rates for each bridge. The node will update the global fair rate for a bridge whenever it receives a new global fairness packet from that bridge or if the previous value timed out after a specified interval. However, how the node uses this information to limit its transmit rate depends on how much it knows about the bridged RPR network topology. With the current RPR standard, the node does not know the next-hop bridge to the destination and thus the transmit rate for all inter-ring packets combined must not exceed the smallest global fair rate received. However, for the other proposed algorithm for bridged RPR networks [9] [10], the node will know the next-hop bridge to the destination. The node can then limit the transmit rate of inter-ring packets by next-hop bridge, using the smallest global fair rate received from that next-hop bridge.

Once a global fair rate has timed out and there are no other limiting global fair rates, the node can then increase its transmit rate gradually. The reason is because an increase in the node's transmit rate may no longer cause congestion and unfairness, due to the changing nature of traffic in the bridged RPR network. However, if the increased transmit rate does cause congestion and the start of unfairness, the global fairness packet will be transmitted to limit the transmit rate again.

6. Buffer Overflow Prevention

In addition to ensuring global fairness in bridged RPR networks, the global fairness algorithm can also be extended to prevent buffer overflow at the bridges. Buffer overflow occurs because the rate of traffic to be forwarded out of the outbound interface is larger than the available bandwidth and there is no available mechanism for the bridge to inform the source nodes that are transmitting to reduce their transmit rate.

By using the global fairness algorithm, the bridge now have a mechanism to inform the source nodes to reduce their transmit rate before buffer overflow occurs. The suggested rate to reduce the transmit rate to is actually the global fair rate. If the nodes limit their transmit rate to the global fair rate, the buffer will no longer fill up and congestion will not occur. How effective the global fairness algorithm is in preventing buffer overflow depends on λ, or the percentage of buffer filled, before the global fairness algorithm is invoked. Obviously the larger the value, the less time there is to inform the nodes before buffer overflow occurs. However, the smaller the value is, the greater the possibility of prematurely triggering the global fairness algorithm unnecessarily.

Another factor affecting buffer overflow is the size of the buffer at the outbound interface. The smaller the buffer, the less packets can be queued and the easier buffer overflow can occur. There is a minimum buffer size such that it is not practical to prevent buffer overflow from occurring at all. However, too large a buffer will increase memory requirements of the interface and increase its cost unnecessarily.

We can calculate the minimum buffer size requirement and the recommended value of λ through some mathematical analysis of the bridged RPR network. First, we consider the simplest case: a bridged RPR network with 2 RPR rings, as shown in Fig. 2. If the bridge invokes the global fairness algorithm, we need to determine the size of the buffer that is necessary to queue the packets that transit through the bridge until the source nodes receive the global fairness packet and limit their transmit rate.

The maximum ring size for a RPR ring is 2000 km. Since the bridge will choose the ringlet with the shortest path to the source node to transmit the global fairness packet on, the packet will not traverse more than half the length of the ring, or 1000 km. This is equal to a propagation delay of about 5 ms through optic fiber. The number of bytes that transit through the bridge during this delay depends on the transmission rate *R* of the link and can be calculated as

$$2 \times 0.005 \times R / 8 \qquad (1)$$

The reason for the multiplier of 2 is because there are packets already on the ring when the global fairness packet is transmitted, in addition to the packets that will be transmitted before the source node receives the global fairness packet. This value gives the size of the buffer that must be available to queue packets when the bridge invokes the global fairness algorithm. For example, if OC-48 link (2.4 Gbps) is used, the required empty buffer size is 3.0 megabytes.

The value calculated by Eq.1 is actually larger than required because at the same time, packets are also removed from the buffer to be transmitted out of the outbound interface. The larger the available bandwidth on the outbound interface is, the larger the rate that packets are removed from the buffer and the less the required buffer size. However, since the available bandwidth varies depending on the local fair rate of the outbound interface, we consider the bound on buffer size by assuming that there is no available bandwidth.

We can then calculate the actual buffer size by dividing the empty buffer size calculated from Eq.1 with the targeted value of λ.

$$\text{Buffer size} = \frac{2 \times 0.005 \times R / 8}{\lambda} \qquad (2)$$

For example, if OC-48 link is used and λ = 90% is targeted, then the actual buffer size is 30 megabytes.

For a more complicated bridged RPR network, the empty buffer size requirement will increase. Given *x* as the maximum number of RPR rings a bridge is connected to and *y* as the span of the bridged RPR network, then the actual buffer size can be calculated as

$$\text{Buffer size} = \frac{2 \times (x-1) \times (y-1) \times 0.005 \times R / 8}{\lambda} \qquad (3)$$

7. Ensuring Fairness during Buffer Overflow

Careful consideration of the 2 factors, the λ value and the buffer size, will be able to prevent buffer overflow. However, it may not be realistic to choose the values that prevent buffer overflow from occurring 100% of the time. For example, the buffer size to prevent buffer overflow may be significant large such that it is not cost-effective. Because of this, we need to consider how we can guarantee fairness as much as possible even when buffer overflow is occurring.

Normally, the bridge uses tail drop to drop packets when the buffer is full. That is, if a packet arrives and the buffer is full, then that packet is dropped. However, this leads to packets of random source nodes to be dropped. In fact, the larger the number of packets transmitted by a node through the bridge, the larger the number of packets of that node will be forwarded out of the outbound interface. This obviously leads to unfairness.

In order to ensure fairness even in the case of buffer overflow, we propose a different drop algorithm. When the buffer is filled to a certain percentage of the buffer size, λ', the proposed drop algorithm will be invoked. Note that λ' must be larger than λ but less than 100% to be effective. The proposed drop algorithm will check the source node of each packet that arrives at the buffer. If the transmit rate of the source node already exceeds the global fair rate, then the packet is dropped, even though the buffer is not yet completely filled. The rate of traffic to be forwarded out of the outbound interface will thus decrease. Also, this will prevent any source node from transmitting more than the global fair rate and causing unfairness; no one node can monopolize the buffer. The remaining empty space in the buffer is reserved for the source nodes that have not yet transmitted its fair share of packets. By using the proposed drop algorithm, these source nodes are more likely to be able to transmit their fair share of the available bandwidth on the outbound interface.

8. Simulation Results

We will use simulations in order to evaluate our proposed algorithms for global fairness and buffer overflow prevention. We have developed a RPR simulation model using the OMNeT++ simulation library [14] that supports bridged RPR networks. We compare the results of the simulation with and without the proposed global fairness algorithm.

In the simulations, each RPR ring consists of 7 nodes. The link capacity is 2.4Gbps (OC-48) and there is a 0.1 ms link delay between neighboring nodes. A uniform packet size of 1536 bytes (24 bytes header and 1512 bytes data) is used. All nodes that transmit have an unlimited number of packets to transmit, unless stated otherwise. The various global fairness algorithm parameter values are λ=80%, λ'=95% and Δ =90%. The aggressive fairness algorithm is used as the local fairness algorithm.

We consider two different networks in our simulations. The first network consists of a bridged RPR network with 2 RPR rings connected together through a single bridge. The second network consists of a bridged RPR network with 4 RPR rings.

8.1 Bridged RPR Network with 2 RPR Rings

We consider a simple bridged RPR network with 2 RPR rings, as shown in Fig.3. The buffer size for each outbound interface at the bridge is 1.6MB.

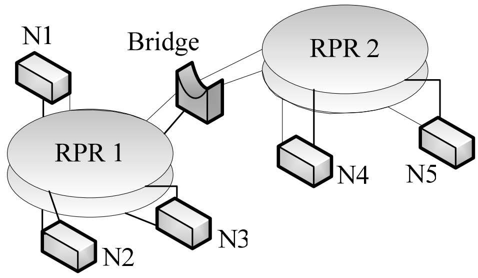

Fig. 3 Bridged RPR network with 2 RPR rings. N5 is the destination node for all packets.

First, we do not consider the effect of the local fairness algorithm in ring RPR2 by assuming that only nodes N1, N2 and N3 are transmitting to node N5. Using the global fairness reference model, we can calculate the global fair rate to be 0.8 Gbps, or 1/3 of the link capacity. The results of the simulations are shown in Fig.4, Fig.5, and Fig.6.

In Fig.4, we can observe that without the proposed global fairness algorithm, node N1 transmits at nearly the link capacity, or 2.4 Gbps. Nodes N2 and N3 share the link capacity through the local fairness algorithm in ring RPR1 and so transmit at 1.2 Gbps. In contrast, with the global fairness algorithm, all the nodes transmit at 0.8 Gbps. This means that the global fair rate was calculated correctly by the bridge and its value distributed to all the relevant source nodes. The source nodes then limit their transmit rate to the global fair rate.

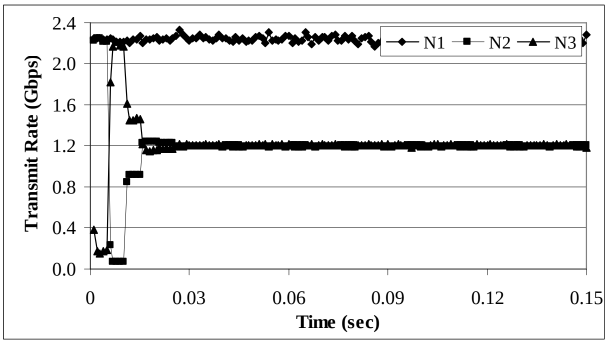

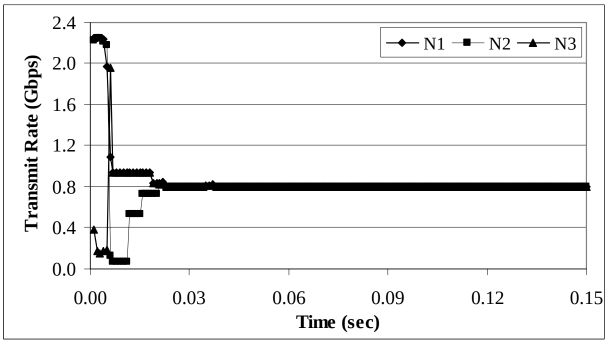

Fig. 4 Transmit rate of nodes N1, N2 and N3. The top figure is without the proposed global fairness algorithm and the bottom figure is with.

In Fig.5, we can observe the rate of packets dropped at the bridge on the outbound interface to node N5. Without the global fairness algorithm, packets are dropped randomly, depending on whether the buffer was full at the time of the packet arrival. Buffer overflow occurs because the rate of traffic arriving at the bridge is twice the available bandwidth of the outbound interface. It is worth noting that more packets from node N1 are dropped, since there are twice as many packets originating from node N1 when compared to nodes N2 and N3. In contrast, with the global fairness algorithm, buffer overflow is avoided and no packets were dropped. The reason is because all the 8 nodes limit their transmit rate to the global fair rate and the buffer size of the outbound interface is sufficiently large.

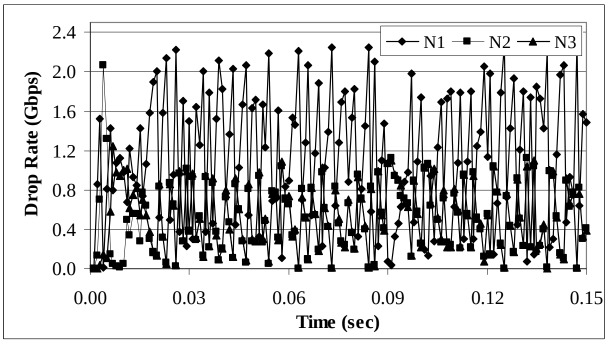

Fig. 5 Drop rate of packets from nodes N1, N2 and N3 at the buffer of the bridge's outbound interface. The figure is without the proposed global fairness algorithm; there are no packets dropped with.

In Fig.6, we can observe that rate of packets arriving at the destination node N5 successfully. Without the global fairness algorithm, the receive rate varies non-deterministically. This is because of the random nature in which packets are dropped at the bridge. However, if the rate is averaged over a long period of time, it can be observed that the receive rate of node N1 is about 1.2 Gbps, while the receive rate of nodes N2 and N3 is about 0.6 Gbps. This is proportional to the transmit rate of those nodes; the higher the transmit rate, the higher the receive rate. Unfairness results because the available bandwidth is not fairly shared among the 3 source nodes. In contrast, with the global fairness algorithm, the receive rate is the same as the transmit rate of 0.8 Gbps for each of the 3 source nodes.

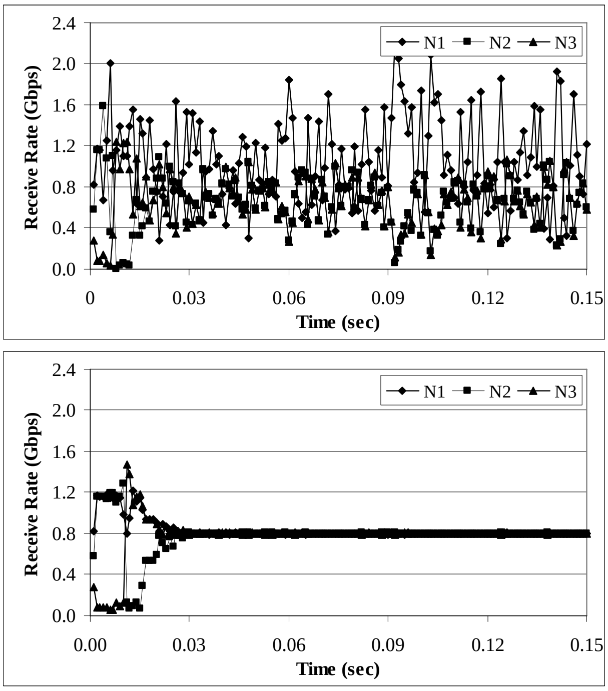

Fig. 6 Receive rate of packets from nodes N1, N2 and N3 at destination node N5. The top figure is without the proposed global fairness algorithm and the bottom figure is with.

In order to further validate the proposed global fairness algorithm, we consider the scenario where node N4 is also transmitting to node N5, in addition to nodes N1, N2 and N3 in the previous scenario. The effect of node N4 transmitting is that the link capacity in ring RPR2 will need to be shared fairly between node N4 and the bridge. That is, the available bandwidth at the outbound interface of the bridge is not the link capacity, but reduced to half the link capacity, or 1.2 Gbps. The global fair rate is thus calculated to be 0.4 Gbps. Without the global fairness algorithm, nodes N1, N2 and N3 are not aware of the reduced available bandwidth at the bridge and so they will transmit at the same rate is in the previous scenario. However, the rate of drop packets at the bridge will be even higher, twice the previous amount, and the receive rate will be half the previous amount.

This scenario is even more undesirable due to the wasted bandwidth utilization in ring RPR1 for the packets that get dropped at the bridge. With the global fairness algorithm, the source nodes limit their transmit rate to the global fair rate, 0.4 Gbps, calculated by the bridge. Buffer overflow is avoided and no packets were dropped at the bridge. The receive rate at the destination node N5 is shown in Fig.7. We can observe that the rate of packets from the nodes N1, N2 and N3 reaching destination node N5 is the same as the global fair rate.

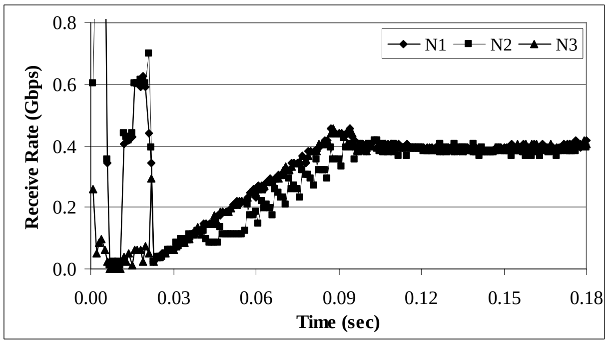

Fig. 7 Receive rate of packets from nodes N1, N2 and N3 when node N4 is transmitting to node N5.

8.2 Bridged RPR Network with 4 RPR Rings

Next, we consider a more complicated bridged RPR network with 4 RPR rings, as shown in Fig.8. The buffer size for each outbound interface at the bridge is 3.2MB. Packets from nodes N1, N2 and N6 arrive at bridge B3 for node N5 on the same inbound interface via ringlet 0 in ring RPR3. Packets from node N3 arrive at bridge B3 for node N5 on a different inbound interface via ringlet 1. Using this network, we also will consider more complicated traffic patterns.

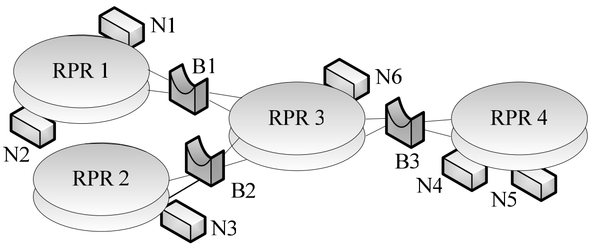

Fig. 8 Bridged RPR network with 4 RPR rings. N5 is the destination node for all packets.

First of all, we consider the scenario where nodes N1, N2, N3 and N4 transmit to node N5. This is similar to the scenario in the previous section; however, we want to evaluate whether the global fairness algorithm will also work well across multiple bridged RPR networks. The result with the global fairness algorithm for the receive rate at destination node N5 is shown in Fig.9. The results are similar to that in Fig.7, confirming that the global fairness algorithm can also be used in larger bridged RPR networks. It can be observed that at around 0.03s, the receive rate for node N3 suddenly drops to nearly 0. The reason is because the source nodes receive the global fairness packet at different times – node N3 receives the packet first and nodes N1 and N2 receive the packet later at the same time. As a result, in the interval between when node N3 and the other source nodes receive the global fairness packet, almost all of the packets reaching the bridge are from nodes N1 and N2. Combined with the small available bandwidth at the bridge at that time, it takes some time before all those packets are forwarded out of the bridge.

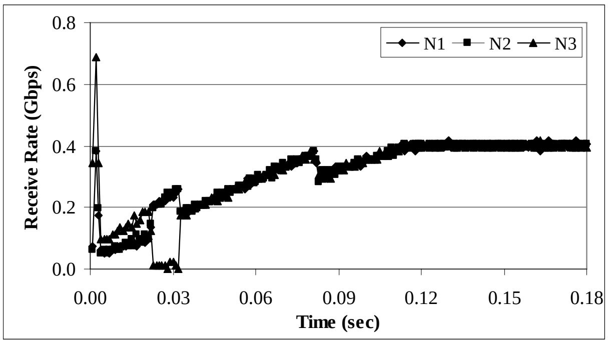

Fig. 9 Receive rate of packets from nodes N1, N2 and N3 at node N5 with the global fairness algorithm.

Next, we consider the effect when some nodes do not have an unlimited number of packets to transmit. This is done by fixing the transmit rate of node N3 to 0.2 Gbps. This is half the value of the global fair rate in the previous scenario. This means that the nodes that will cause unfairness in the network are nodes N1 and N2 only. The global fair rate for this scenario, as calculated using the global fairness reference model, is 0.5 Gbps. The result with the global fairness algorithm for the receive rate at destination node N5 is shown in Fig.10. All the packets from node N3 are able to reach the destination, since the node's transmit rate is below the global fair rate. The other nodes N1 and N2 use 0.5 Gbps of the bandwidth, which is equal to the global fair rate. Thus, the global fairness algorithm can ensure global fairness, even when some of the nodes do not have an unlimited number of packets to transmit.

Last, we consider the effect when some of the source nodes themselves are on the path of the inter-ring traffic. For this scenario, we consider nodes N1, N2, N4 and N6 transmitting to destination node N5.

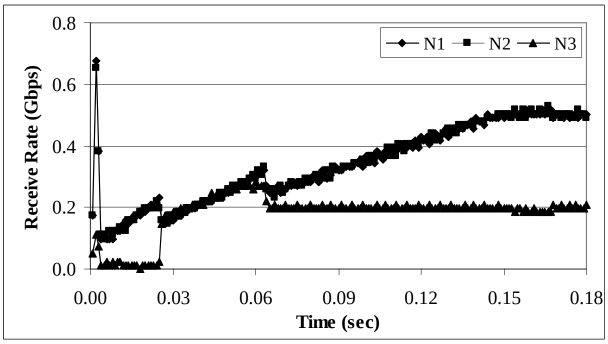

Fig. 10 Receive rate of packets from nodes N1, N2 and N3 at node N5 with the global fairness algorithm. Node N3 transmits at a lower rate, 0.2Gbps, than nodes N1 and N2.

Packets from nodes N1 and N2 will have to transit through node N6 on the path to node N5. This is interesting because node N6 itself has its own packets to transmit. The local fairness algorithm will be triggered in both rings RPR3 and RPR4. The global fair rate for this scenario is 0.4 Gbps.

In Fig.11, the results for the drop rate of packets at the bridge B3, without the global fairness algorithm, is shown. First of all, the local fairness algorithm will be triggered in ring RPR3, between node N6 and bridge B1. Each will be given an equal share of the link capacity, or 1.2 Gbps.

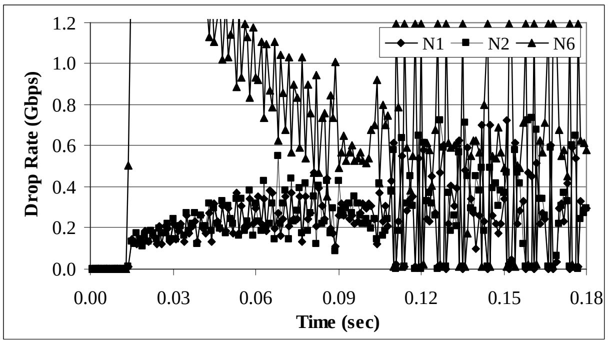

Fig. 11 Drop rate of packets fro nodes N1, N2 and N6 at the bridge B3 without the global fairness algorithm. The packets from nodes N1 and N2 transit through node N6.

Therefore, even though nodes N1 and N2 transmit at the link capacity, three quarters of the packets will be dropped at bridge B1 and only 1.2 Gbps will be forwarded on to ring RPR3. Next, the available bandwidth at bridge B3 is also only 1.2 Gbps because of node N4 transmitting to node N5. Therefore, the packets from nodes N1, N2 and N6 will also be dropped at bridge B3 to reduce the traffic forwarded on to ring RPR4 to 1.2 Gbps. Since more packets from node N6 reach bridge B3, the drop rate of that node is larger than the other two nodes N1 and N2. This case is similar to that shown in the top figure of Fig.5. Unfairness occurs in the receive rate at node N5, with more packets from N6 reaching the destination than the other 2 nodes.

In Fig.12, we show the results for the same scenario, but with the proposed global fairness algorithm. No packets were dropped at the bridge and so only the receive rate at node N5 is shown. It can be observed that the receive rate of all the source nodes are the same at 0.4 Gbps. Global fairness is ensured.

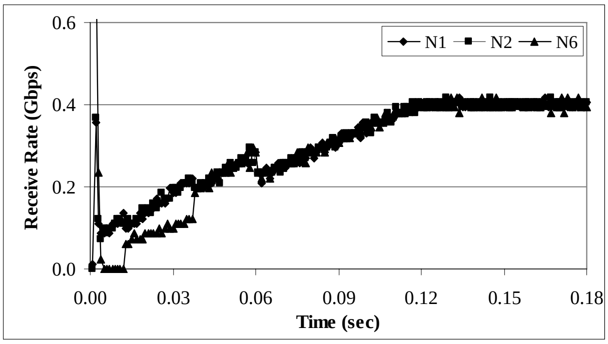

Fig. 12 Receive rate of packets from nodes N1, N2 and N6 and node N5 with the global fairness algorithm. The packets from nodes N1 and N2 transit through node N6.

8.3 Effects of Parameters for Buffer Overflow

In the previous two scenarios, we have chosen the buffer size sufficiently large such that buffer overflow does not occur with the global fairness algorithm. In this subsection, we consider the effect of the parameters of the global fairness algorithm on buffer overflow, using the network scenario shown in Fig.3 with nodes N1 and N3 transmitting to node N5.

First, we consider the effect of the buffer size by fixing the value of λ to 80% and observing the drop rate with varying buffer sizes. Using Eq.2, we can calculate the buffer size to totally prevent buffer overflow must be larger than 600,000 bytes. However, since the rate at which the packets leave the outbound interface is close to the link capacity, it is sufficient for the buffer size to be slightly larger than half of that size. The result is shown in Fig.13, with the buffer size in bytes.

From the figure, we can observe that the smaller the buffer size, the higher the drop rate. The buffer is not sufficiently large to queue the packets before the global fairness packet reaches the source nodes. There are only some packets dropped when the buffer size is 306,200 bytes; when the buffer size is 409,600 bytes, packets are no longer dropped. This corresponds with our calculation.

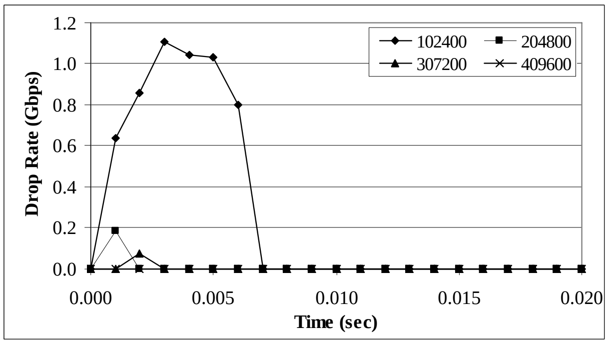

Fig. 13 Total drop rate of packets at the bridge with λ = 80% for various buffer sizes. The buffer size is given in bytes.

Next, we consider the effect of the value of λ by fixing the buffer size to 204,800 bytes and observing the drop rate with varying values of λ. Using Eq.2 and considering that the packets leave the outbound interface at the bridge close to the link capacity, we can calculate the minimum value of λ to totally prevent buffer overflow to be 70%. The result is shown in Fig.14.

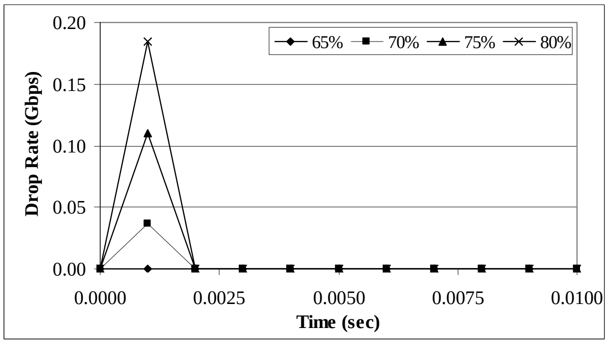

Fig. 14 Total drop rate of packets at the bridge with buffer size of 204,800 bytes for various values of λ.

From the figure, we can observe that the smaller the value of λ, the smaller the drop rate. The smaller the value of λ, the earlier the source nodes are informed of the global fair rate and thus the less the buffer size required to queue packets at the bridge. From the results, it is worth noting that when λ is 70%, there are still some packets being dropped; only when λ is 65% are packets no longer dropped. The reason for this is because the value of λ' is set to 95%. According to the proposed drop algorithm in Section 7, packets are prematurely dropped at the bridge because the source nodes have transmitted more than the global fair rate through the bridge, even though the buffer is still not filled. Therefore, the value of (100% - λ') needs to be subtracted from the calculated value of λ to find the actual minimum value of λ to use.

Another parameter to consider for buffer overflow is the value λ'. For this, we use the network scenario shown in Fig.8, with nodes N3 and N6 transmitting to node N5. We fix the value of λ = 80% and the buffer size to 51,200 bytes. In this case, the buffer size is not sufficient and packets will be dropped at the bridge. The receive rate at destination node N5 for λ' = 95%, 100% are shown in Fig.15 for each source node. Note that λ' = 100% is equivalent to not using the proposed drop algorithm.

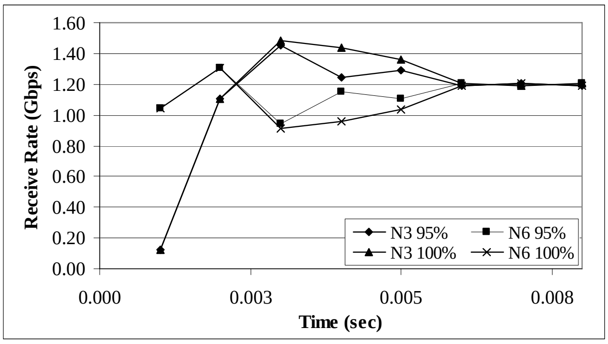

Fig. 15 Receive rate of packets from nodes N3 and N6 at node N5 for various values of λ'.

During the interval 0.001s to 0.006s, packets are dropped at the bridge. When λ' = 95%, the receive rate for packets from nodes N3 and N6 are close to the global fair rate, when compared to when λ' = 100%. That is, during the interval when packets are dropped, using the proposed drop algorithm allows for a greater degree of global fairness.

9. Conclusion

In this paper, we have emphasized the importance of considering global fairness and buffer overflow prevention in bridged RPR networks. Although the current RPR standard specifies local fairness algorithms for ensuring fairness within a single RPR ring, fairness for inter-ring traffic in bridged RPR networks should also be considered, in order to ensure that the available bandwidth for transmitting inter-traffic is fairly shared among all the nodes. In addition, it is important to prevent, or at least minimize, buffer overflow at the bridges; otherwise, packets will be dropped in transit at the bridges, resulting in inefficient network utilization.

As a result, we propose the global fairness reference model in order to define fairness for inter-ring traffic. Using the reference model, we can calculate the global fair rate at the bridge for each outbound interface. If the nodes do not transmit beyond the global fair rate, global fairness can be maintained. Next, we propose a global fairness algorithm for implementing global fairness according to the global fairness reference model. Whenever unfairness is detected, the algorithm is triggered in order to inform the source nodes of the global fair rate, through the use of the global fairness packet. In addition, we propose that the global fairness algorithm can also be used to prevent buffer overflow at the bridges. The equation to calculate the buffer size in order to prevent buffer overflow is presented. In order to evaluate our proposed global fairness algorithm, simulations were performed for various bridged RPR networks and traffic patterns. The results show that the proposed global fairness algorithm ensures global fairness and minimize buffer overflow successfully.

References

[1] IEEE Standard 802.17, *Resilient packet ring (RPR) access method and physical layer specifications*, IEEE Std., 2004.

[2] I. Cidon and Y. Ofek, "MetaRing – A Full-Duplex Ring with Fairness and Spatial Reuse," *IEEE Transactions on Communications*, vol. 41, no. 1, pp. 110–119, 1993.

[3] S. G. F. Davik, M. Yilmaz and N. Uzun, "IEEE 802.17 Resilient Packet Ring Tutorial," *IEEE Communications Magazine*, vol. 42, no. 3, pp. 112–118, 2004.

[4] V. G. P. Yuan and E. Knightly, "The IEEE 802.17 Media Access Protocol for High-speed Metropolitan-Area Resilient Packet Rings," *IEEE Network*, vol. 18, no. 3, pp. 8–15, 2004.

[5] S. Gjessing and F. Davik, "Performance Evaluation of Back-Pressure Fairness in RPR," in *Proceedings of the 7th European Conference on Networks & Optical Communications*, Germany, 2002, pp. 293–300.

[6] V.Gambiroza, P.Yuan, L. Balzano, Y. Liu, S. Sheafor, and E. Knightly, "Design, Analysis, and Implementation of DVSR: A Fair, High Performance Protocol for Packet Rings," *IEEE/ACM Transactions on Networking*, vol. 12, no. 1, pp. 85–102, 2004.

[7] C. L. L. Tan, Y. Yang and N. Xiong, "PID-RPR: A High Performance Bandwidth Allocation Approach for RPR Networks," *IEICE Transactions on Communications*, vol. E88-B, no. 7, pp. 2872–2878, 2005.

[8] IEEE Standard 802.1D, *IEEE 802.1D-2004, Media Access Control (MAC) Bridges*, IEEE Std., 2004.

[9] A. Kvalbein, S. Gjessing, and F. Davik, "Performance Evaluation of an Enhanced Bridging Algorithm in RPR Networks," in *Proceedings 3rd International Conference on Networking ICN'04*, vol. II, French Caribbean, 2004, pp. 760–767.

[10] P. Setthawong and S. Tanterdtid, "Efficient Traffic Management for Bridged Resilient Packet Rings Using Topology Discovery and Spanning Tree Algorithm," in *Proceedings IASTED International Conference on Networks and Communication Systems NCS'05*, 2005, pp. 392–397.

[11] F. Yuan and C. Huang, "Performance analysis of resilient packet rings with single transit buffer," in *Proceedings of IEEE ICT'2002*, Beijing, China, 2002.

[12] C. R. W.Dongmei, K.K. Ramakhrisnan and A.Smiljanic, "Congestion Control in Resilient Packet Ring," in *Proceedings of the 12th IEEE International Conference on Network Protocols ICNP'04*, 2004, pp. 108–117.

[13] F. Davik, A. Kvalbein, and S. Gjessing, "An analytical bound for convergence of the resilient packet ring aggressive mode fairness algorithm," in *Proceedings of the 40th IEEE International Conference on Communications, Seoul, Korea, May 16-20*, 2005.

[14] A. Vargas, "The OMNeT++ Discrete Event Simulation System," in *Proceedings of the European Simulation Multiconference ESM'01*, Czech Republic, 2001.

Pisai Setthawong is a Ph.D. candidate at the Department of Telecommunications Science, Assumption University, Thailand. He received his B.S. (1996) and M.S. (1998) degrees from the University of Tokyo, Japan in Information and Communication Engineering and Electronics Engineering, respectively. His research interests include protocols for high-speed networks, traffic engineering, and quality of service.

Surat Tanterdtid received his B.Eng, M.Eng, and D.Eng degrees in Electrical Engineering from Chulakongkorn University, Thailand in 1993, 1996, and 2000 respectively. He is currently a special instructor at the Faculty of Science and Technology, Assumption University, Thailand. He is a correspondence member of the CIGRE SC D2 Study Committee. His research interests include MPLS, IPv6 and RPR.