Non-Flooding Bridging Solutions for Resilient Packet Rings

Pisai Setthawong[1] and **Surat Tanterdtid**[2], Non-members

ABSTRACT

Resilient Packet Ring (RPR) is one of the IP-based technologies that have been proposed to replace SONET/SDH in metropolitan area networks because it is well-adapted to handle diverse traffic, including multimedia traffic, found in present-day networks. Additionally, the RPR network is used efficiently with nodes transmitting simultaneously as long as the paths of the packets do not overlap. However, when bridging RPR networks, packets are flooded on the bridged RPR network if the packet destination is on a remote network other than the source network. As a result, the network is used inefficiently, decreasing the available bandwidth for other traffic. In this paper, we propose enhanced topology discovery protocol and enhanced spanning tree protocol to prevent the flooding of packets on the bridged RPR network. Simulations were performed in order to evaluate the proposed solutions. The results show that the flooding of packets is successfully prevented and that the network is used more efficiently as compared to the RPR

Keywords: Resilient Packet Rings, IEEE 802.17 Standard, Bridging, Topology Discovery, Spanning Tree Protocol, Traffic Engineering

1. INTRODUCTION

The prevalent technology found in present-day metropolitan area networks (MANs) is based on SONET/SDH. However, SONET/SDH is a circuit-switching technology that is ill-adapted to handle the bursty nature of multimedia traffic increasingly found in present-day networks. IP-based networks, which benefit from statistical multiplexing, are more appropriate for handling multimedia traffic. Resilient Packet Ring (RPR), also known as the IEEE 802.17 Standard [3], is one such technology. RPR is a dual-ring network that defines a SONET/SDH reconciliation sublayer. This sublayer enables RPR to operate on the same physical layer as SONET/SDH and to utilize the existing SONET/SDH physical infrastructure.

There are other key benefits of using RPR instead of SONET/SDH. The former defines three service classes with different service guarantees for supporting various traffic types. Bandwidth can be reserved for time-sensitive traffic such as VoIP and videoconferencing. For best-effort traffic such as web and file transfer, fairness algorithms ensure that the nodes are dynamically allocated their fair share of the available bandwidth. The RPR network is also efficient. Both dual-rings are used for packet transmission. The packets are removed at the destination, allowing the nodes to transmit simultaneously as long as the paths of the packets do not overlap.

Although MANs can be implemented as a single large RPR network spanning hundreds of kilometers, most carrier networks consist of multiple small-sized MANs, or access MANs, interconnected by inter-area MANs. This network architecture can also be implemented using RPR. Each of the MANs is implemented as a single RPR network and the RPR networks are interconnected by bridges.

However, the bridged RPR network is inefficient. Packets are flooded on the bridged RPR network if the packet destination is on a remote network other than the source network. The flooding decreases the available bandwidth in the network for other traffic. Network efficiency, a key benefit of RPR, is no longer achievable. In order to achieve network efficiency in bridged RPR networks, this paper proposes two solutions to prevent the flooding of packets – enhanced topology discovery protocol and enhanced spanning tree protocol.

The remainder of this paper is organized as follows. Section 2 presents background information about RPR. Section 3 describes previous research on bridged RPR networks. In Section 4, we present the proposed solutions. The proposed solutions are then evaluated by means of discrete event simulations in Section 5. Finally, Section 6 concludes the paper.

2. BACKGROUND

The structure of an RPR network consisting of two unidirectional ringlets and four nodes is shown in Figure 1. Both ringlets are used for data transmission, with the traffic on one ringlet flowing in a clockwise direction (ringlet 0) and that on the other flowing in a counterclockwise direction (ringlet 1). In order to use these ringlets efficiently, the nodes transmit packets on the ringlet with the smaller hop count to the destination. For example, while node 1 transmits packets to node 2 on ringlet 0, it transmits packets to node 4

Manuscript received on June 21, 2007 ; revised on October 13, 2007.

[1,2] The authors are with department of Telecommunications Science, Assumption University, Thailand, E-mail:

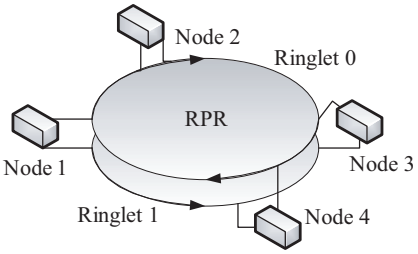

Fig.1: *An RPR network consisting of two unidirectional, counter-rotating ringlets and four nodes.*

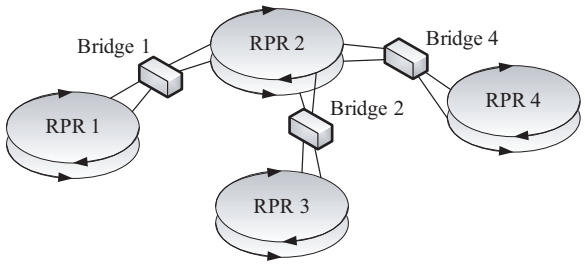

Fig.2: *A bridged RPR network consisting of four RPR networks interconnected by three bridges.*

on ringlet 1 [2]-[9].

2.1 Topology Discovery Protocol

In order to determine the ringlet to transmit the packet on, the node maintains a topology database. The topology database includes a list of nodes on the network and the hop count to those nodes on each ringlet.

The node maintains the topology database using the topology discovery protocol. The node broadcasts a topology discovery packet on both ringlets. When the other nodes on the network receive the topology discovery packet, they learn about the existence of the node that originated the packet. The hop count between the nodes is then calculated as the difference between the *ttlbase* and the *ttl* fields in the packet. The *ttlbase* field contains the original time-to-live value of the packet and the *ttl* field contains the current time-to-live value.

Once all the nodes on the network have broadcasted their topology discovery packets, the nodes will have a topology database of the network. The node will then be able to determine the ringlet to transmit the packet on with the smaller hop count to the destination. The node must continue to broadcast the topology discovery packet periodically. If the node fails to do this, the other nodes on the network will assume that the node has been detached and they will rebuild their topology database to exclude this node.

2.2 Bridging with RPR

Multiple RPR networks can be bridged together to form a bridged network, as shown in Figure 2.

The specialized node that interconnects the RPR networks is known as the bridge. The bridge forwards or copies the packets that transit it on one network to the other network(s) it is connected to, in order to ensure that the packets reach their destination.

There are many possible reasons for bridging multiple networks together. First, a single network can support a maximum of 255 nodes and is optimized for a maximum ring length of 2000 km. When a greater number of nodes or a longer ring length is required, the solution is to bridge multiple RPR networks together. Second, each network may be under different administrative domains but interconnections between the domains are still desired, such as between different faculties in the same university. Third, for carrier networks, the RPR networks can be used to implement access MANs that are bridged together to form inter-area MANs.

2.3 Spanning Tree Protocol

One of the major issues with bridging, not only in RPR but also in other IEEE 802 networks, is the risk of creating loops within the bridged network. If a loop exists, there is a possibility that a bridge will continue to copy packets to its connected networks without realizing that the packets have already been copied to those networks previously. The packets will continue to loop in the bridged network indefinitely and use bandwidth unnecessarily. In Figure 3, networks 1, 2, and 3 form such a loop.

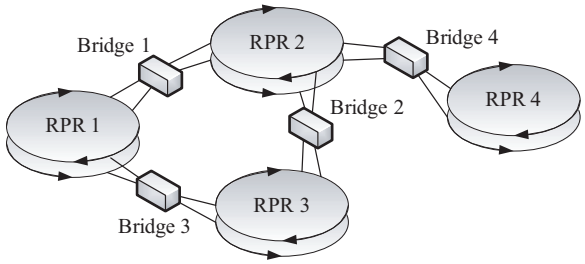

Fig.3: *Bridges 1, 2 and 3 form a loop consisting of RPR networks 1, 2, and 3.*

The loops may be created unintentionally, especially when the networks are bridged over multiple administrative domains. On the other hand, the loops may also be created intentionally, for example, to provide redundant paths in the case of link failure. When a link fails, an alternative path still exists between nodes.

In any case, if there are loops in the bridged network, they must be removed. For this purpose, the IEEE standard specifies the spanning tree protocol [**?**,]. The result of executing this protocol is that particular links to the bridges are identified to be deactivated, removing the loops in the bridged network. This protocol involves the following steps:

• Bridges periodically exchange configuration messages, also known as bridge protocol data units (BPDUs), among themselves. The BPDUs include: (1) the address of the bridge that created the BPDU, (2) the address of what the bridge has identified as the root bridge, and (3) the distance from the bridge to the root bridge.
• A bridge will initially identify itself as the root bridge and will create its own BPDUs. When a bridge receives a better BPDU, it will no longer consider itself the root bridge and will stop creating its own BPDUs. Instead, it will only forward the BPDUs it receives from the other bridges.
• A better BPDU is one that: (1) identifies a root bridge with a smaller MAC address, (2) identifies the same root bridge but with a closer distance to the root bridge, or (3) identifies the same root bridge with the same distance to the root bridge, but the bridge that generated the message has a smaller MAC address.
• A bridge will no longer forward BPDUs over a link if it determines that another bridge connected to the same network has received a better BPDU than the one it will forward.

After the protocol has converged, only one bridge will remain that still identifies itself as the root bridge. The path taken by the BPDUs created by this root bridge determines the links that remain activated; the other links are then deactivated. In this manner, all the loops in the network will be removed completely.

As an example, after executing the spanning tree protocol over the network shown in Figure 3, the resulting network is that shown in Figure 2. The links between bridge 3 and networks 1 and 3 are deactivated.

3. PREVIOUS RESEARCH ON BRIDGING RPR NETWORKS

Although bridging RPR networks can be done in the same manner as that with other IEEE 802 networks, it also presents its own unique issues. The focal point of these issues is the algorithm a bridge uses to decide whether to forward or copy packets from one of its connected networks to its other connected network(s).

The original IEEE 802.17 Working Group realized the complexities involved in solving these issues and opted for a simple solution in order to complete the work on the standard. Nevertheless, the need for further work was recognized; therefore, after the IEEE 802.17 Standard was approved, the IEEE 802.17b Working Group was formed to continue the work on solving the issues involved in bridging RPR networks. In this section, we present the IEEE 802.17 Standard for bridged RPR networks, as well as the results of other ongoing research.

3.1 Bridging in the IEEE 802.17 Standard

Bridging in the IEEE 802.17 Standard, hereafter referred to as standard bridging, makes the assumption that each node only knows the location of the nodes that are on the same network as itself. When the node wants to transmit a packet to a destination not on the source network, the node does not know which bridge it should transmit the packet to in order to reach the destination. Similarly, when a bridge receives a packet for a destination not on the network it is connected to, it will also not know the next bridge to forward the packet to in order to reach the destination.

In order to overcome this limitation, the IEEE 802.17 Standard proposes that if the packet destination is on a remote network, the packet is flooded on the bridged RPR network. The source node will set the flood bit on its packets so that it is flooded on the source network. The bridges connected to the source network will then copy and flood the packets on all their connected networks. Even when the packet reaches a bridge that is connected to the destination network, the bridge must flood the packet on the destination network. By flooding the packet on the bridged network, it is guaranteed that the packet reaches its destination. In Figure 4, the flooding of packets results in the packet being transmitted on network 4, although the network is not on the direct path between the source and destination nodes.

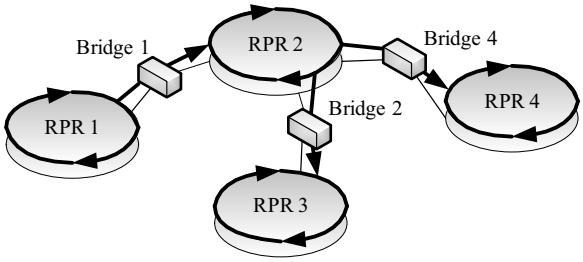

Fig.4: *Path taken by the packet transmitted from the source node on network 1 to the destination node on network 3.*

As can be observed from the example, standard bridging is inefficient since it uses unnecessary bandwidth when flooding the packet on the bridged network. The larger the size of the network, the greater is the inefficiency. This decreases the available bandwidth for other traffic in the network. As a result, while the algorithm is simple to implement, it does not use the network efficiently.

3.2 Enhanced Bridging

In order to improve upon standard bridging, it is necessary to add some learning capabilities to the bridges. For Ethernet (IEEE 802.3) networks, the concept of transparent bridging is used, where the bridges passively learn about the location of the

nodes by observing the packets that transit through it. Once the bridges learn about the location of the nodes, they will know which network they should forward the packets to and flooding is no longer necessary. This process is transparent to the non-bridge nodes.

Enhanced bridging [**?**,] uses the same concept of passive learning to prevent flooding in bridged RPR networks. In order to accomplish this, it makes use of the extended data packet format. The extended data packet format is similar to the standard data packet format, except for 2 additional fields – the extended source field and the extended destination field.

Initially, the node does not know the location of the packet destination if the destination is on a remote network and thus the node floods the packet on the source network. When the bridge connected to the network receives the packet, it will initially also not know about the location of the destination. It will then create a copy of the packet using the extended packet format by placing the actual source and destination addresses in the extended source and extended destination fields. In the source field, it will place its own address and in the destination field, it will place a bridge group address. The packet will then be flooded on all the bridge's other connected networks.

When the packet reaches another bridge on the next network, that bridge will remember that in order to reach the node in the extended source field, it needs to forward the packet to the bridge in the source field. In short, the bridges learn about how to reach a particular node by observing the extended source and source fields in the extended data packet. Eventually, the packet will reach its destination. The destination node will remember that in order to reach the node in the alternative source field, it needs to transmit the packet to the bridge indicated in the source field, rather than flood the network with the packet.

In this manner, the nodes and bridges will gradually learn about the network topology and flooding will no longer be necessary. Although enhanced bridging improves upon standard bridging, there are other issues with enhanced bridging. First, the traffic in the network has to be bidirectional for enhanced bridging to work. Second, all nodes need to maintain a forwarding database of the next-hop bridge to transmit the packet to in order to reach the destination. This increases the logic and memory requirements of the non-bridge nodes. Finally, the two pairs of extended address fields in the extended data packet format increases the packet size by 12 bytes. This increases the bandwidth overhead when transmitting packets.

4. PROPOSED SOLUTIONS

In order to improve upon standard bridging and enhanced bridging, this paper proposes two solutions for improving the network efficiency of bridged RPR networks and increasing the overall network bandwidth availability. The solutions make use of existing protocols, but makes minor enhancements to those protocols to achieve the desired improvements.

4.1 Enhanced Topology Discovery Protocol

First, our proposal makes enhancements to the topology discovery protocol by allowing the nodes to not only identify the nodes on their connected network, but also identify which of these nodes are bridges. The node can then choose the closest bridge as its default gateway bridge. When the packet destination is on a remote network, the node transmits the packet to the default gateway bridge instead of flooding the packet on the source network. When the bridge receives the packet, it will perform the following actions:

- If the bridge knows that it is the next-hop bridge toward the destination, it forwards the packet to the next network that is on the path to the destination.
- If the bridge knows that another bridge on the same network is the next-hop bridge, it lets the packet transit towards that bridge without flooding the packet on its other connected networks. That bridge will forward the packet toward the destination instead.
- If the bridge does not know the next-hop bridge, it copies and forwards the packet on all its connected networks. It also lets the packet transit so that the other bridges on the source network can also process the packet.
- If the bridge is the last bridge on the source network before the packet returns back to the source node, it can safely remove the packet from the network instead of the source node.

In this manner, when the source node is uncertain of the next-hop bridge for a packet, flooding of the packet at the source network can be avoided.

In the best-case scenario, the default gateway bridge is the next-hop bridge for the packet and the smallest hop count is achieved. In the worst-case scenario, the next-hop bridge is closer to the source node on the opposite ringlet that was used to transmit the packet to the default gateway bridge. In this case, the hop count to the actual next-hop bridge is not optimal. Note that this will only be an issue when there is more than one bridge in the source network. Regardless of the scenario, the enhanced topology discovery will still allow the network to be used more efficiently as compared to standard bridging.

In order to implement the enhanced topology discovery protocol, an additional bit flag is required to indicate whether the node is a bridge or not. We propose to include the flag in the attribute discovery (ATD) packet instead of the topology and protection (TP) packet. Note that both types of packet are defined by the topology discovery protocol. This is because the TP packet must be transmitted fre-

quently enough that a link failure can be detected and automatically resolved within a maximum of 50 ms. Adding the bit flag to the TP packet will lead to an increase in the bandwidth consumed by the TP packets.

On the other hand, the ATD packet is transmitted less frequently for less time-critical information on the node, such as station name or secondary MAC address. The ATD packet supports defining multiple attributes through its type-value format. A type can be defined for the bridge indicator bit flag using one of the available reserved types.

4.2 Enhanced Spanning Tree Protocol

For the bridges to learn about the bridged network topology, we propose the enhanced spanning tree protocol as a means for distributing topology information among the bridges. This can be viewed as similar to routing information being distributed among routers in a wide area network. The bridges learn the networks on which the nodes are connected and the next-hop bridge to forward the packet to in order to reach the nodes. Once the bridges have learnt the bridged network topology, there will no longer be any need to flood packets on the bridged network.

The enhanced spanning tree protocol is as follows. When the bridge creates a BPDU, it will append the list of nodes in its topology database and the next-hop bridge to reach each of the nodes to the BPDU. If the next-hop bridge is on a different network than that which the BPDU will be transmitted on, the current bridge address will be substituted as the next-hop bridge instead.

When the other bridges receive the BPDU, they will add the list of nodes in the BPDU into their topology database with the indicated next-hop bridge. The nodes whose next-hop bridge is the same as the bridge that received the BPDU are excluded. The bridge will also duplicate the BPDU to transmit on all its other connected networks, but include only the list of nodes it had just added to its topology database and indicate itself as the next-hop bridge for those nodes. Further, before allowing the original BPDU to transit onwards, the bridge will add the list of nodes in its topology database to the BPDU, but exclude the nodes with the next-hop bridge that is on the same network as that which the BPDU was received. The bridge will indicate itself as the next-hop bridge for those nodes it added to the BPDU.

In this manner, the list of nodes in the BPDU will grow as the BPDU traverses the network. When the BPDU finally returns to the bridge that created it, that bridge will observe the BPDU for additional nodes that have been added to the list since it was created. It will then add those nodes to its topology database with the indicated next-hop bridge. This process is necessary for the upstream bridges to learn about the nodes connected to the downstream bridges. In this context, the closer the bridge is to the root bridge, the more upstream the bridge is. Eventually, the bridges will learn about all the nodes in the bridged network and the next-hop bridge for each of those nodes.

In order to implement the enhanced spanning tree protocol, the BPDU format must be modified so that it can include a variable-length list of node MAC addresses in the packet payload. The list should be separated by the MAC addresses of the next-hop bridges of those nodes.

After the bridge has finished building its topology database, the bridge can now perform the following steps when a packet transits:

- If the destination is on the current network, the packet is allowed to transit.
- If the destination is on a different network, it checks the next-hop bridge for the destination.
  - If the next-hop bridge is reached on the current network, then the packet is allowed to transit.
  - If the next-hop bridge is on another connected network, the bridge forwards the packet on to the network where the next-hop bridge can be reached.
  - If the next-hop bridge cannot be found, this means that the enhanced spanning tree protocol has not yet converged and the packet is flooded on all the connected networks.

Once the enhanced spanning tree protocol has converged, there will no longer be any flooding of packets and the network can be used efficiently. It is interesting to note that the enhanced spanning tree protocol is executed only by the bridges; there is no additional information that needs to be maintained at the non-bridge nodes. Hence, transparent bridging can be achieved.

4.3 Enhanced Spanning Tree Protocol with Non-Bridge Nodes

As discussed in Section 4.1, there are some scenarios where the packets are not transmitted on the optimal path from the source node to the next-hop bridge since the node is only aware of the default gateway bridge. This can lead to bandwidth being used unnecessarily. However, this is unavoidable as long as bridging is transparent, since the node cannot differentiate the next-hop bridge by destination.

In cases where the amount of bandwidth being used unnecessarily becomes significant, it might be worthwhile to consider non-transparent bridging; that is, having the non-bridge nodes maintain a forwarding database of destination nodes and their next-hop bridges. This can be easily accomplished by having the non-bridge nodes observe the BPDUs exchanged between the bridges. In this manner, the nodes can also passively learn about the bridged network topology without having to be actively involved in the enhanced spanning tree protocol.

By thus making use of the enhanced spanning tree

protocol, the decision on whether bridging should be transparent or whether the non-bridge node should maintain a forwarding database can be decided on a node-by-node basis. If there is only one bridge on the network the node is connected to, there is no gain achieved from maintaining a forwarding database. On the other hand, the nodes that are connected to a network with multiple bridges might need to consider maintaining a forwarding database.

4.4 Bound on Completion of the Enhanced Spanning Tree Protocol

The enhanced spanning tree protocol takes $n + 2$ cycles to converge, before all the bridges in the bridged RPR network learn about all the nodes. A cycle refers to each time the root bridge creates and transmits a BPDU. n refers to the maximum number of networks between the root bridge and any other bridge in the bridged network.

The spanning tree protocol itself takes one cycle to determine the root bridge and the links to deactivate. Next, it takes at most n cycles for the topology information from all the bridges to reach the root bridge. In each of the cycles, the topology information is transmitted further to the next upstream network, one network per cycle, until it finally reaches the root bridge. Then, once the root bridge has the topology information from all the bridges, it takes one more cycle to distribute the information downstream to all the other bridges. Hence, it takes a total of $n+2$ cycles for all the bridges to build their topology database. The greater the number of networks in the bridged network, the longer it takes for the enhanced spanning tree protocol to converge.

5. SIMULATION RESULTS

In order to evaluate the proposed solutions, discrete-event simulations were performed. We implemented an RPR network simulator using the simulation tool OMNeT++ [8]. The SONET/SDH physical layer is simulated as a 2.4 Gbps (OC-48) link, with a propagation delay of 0.1 ms between adjacent nodes. The fairness algorithm is also implemented to ensure that the source nodes are allocated their fair share of the available bandwidth for transmitting packets.

The simulated bridged RPR network is shown in Figure 5. With this network, the proposed solutions can be evaluated under various conditions such as for various distances between the source and destination nodes and for different numbers of bridges on the source network. There are nine nodes on each network (not shown) and the bridges are connected such that the number of nodes between the bridges is evenly distributed within each network.

The source nodes have an unlimited number of packets to transmit, unless stated otherwise, and start to transmit the packets after the topology discovery and the spanning tree protocols have converged. The packets have a uniform size of 1536 bytes, the maximum size for a packet. The ATD packet and the BPDU are transmitted periodically every 1.0 s, the default interval according to the standards [3] [4].

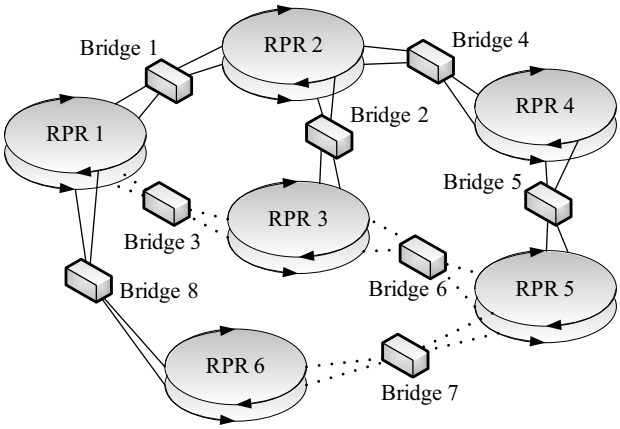

Fig.5: *Bridged RPR network used in the simulations; it consists of six RPR networks interconnected by eight bridges. The dotted lines represent links that are deactivated by the bridges.*

5.1 Evaluation of the Enhanced Spanning Tree Protocol

The simulation results show that for the bridged RPR network shown in Figure 5, the enhanced spanning tree protocol converges in five cycles. This is within our bound of $n + 2$ from Section 4.4 where $n = 3$ for bridge 1, the root bridge, and bridge 7.

Next, we evaluate the enhanced spanning tree protocol by observing the total hop count of the packet as it is transmitted from the source to the destination. If the packet is forwarded onto multiple networks, the hop count on each of the networks is also included in the total hop count. For the optimal case, the packet takes the path with the smallest hop count to the destination. The greater the total hop count, the less efficient is the network.

The simulation results for the source node on network 6 transmitting packets to the destination nodes on networks 1, 2, 4, and 5 are shown in Figure 6. The node on network 6 is chosen as the source node because network 6 has only one bridge. Thus, we can evaluate the enhanced spanning tree protocol independently of the enhanced topology discovery protocol.

The results obtained using the enhanced spanning tree protocol are the same as those derived for the optimal case. That is, flooding is successfully prevented and the network is used efficiently.

In contrast, the results obtained using standard bridging increasingly differ from those derived for the optimal case as the number of networks between

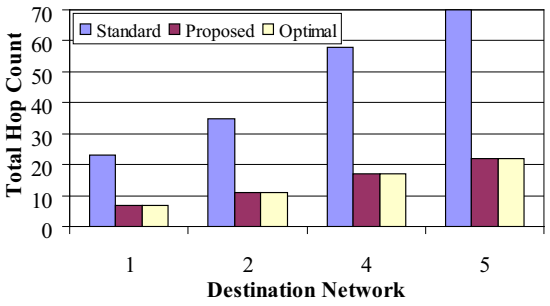

Fig.6: *Total hop count of packets transmitted between the source node on network 6 and the destination nodes on networks 1, 2, 4, and 5. The results are shown for standard bridging (Standard), the enhanced spanning tree protocol (Proposed), and the optimal case (Optimal).*

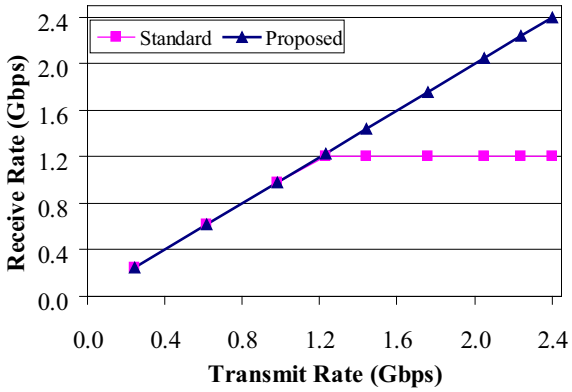

Fig.7: *Receive rate for standard bridging (Standard) and the enhanced spanning tree protocol (Proposed) at the destination node on network 5 as the transmit rate of the source node on network 6 is increased.*

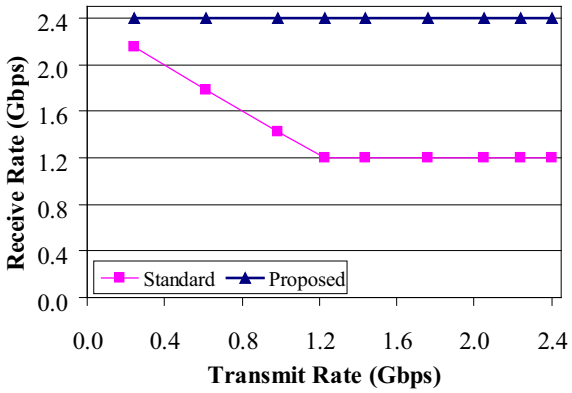

Fig.8: *Receive rate for standard bridging (Standard) and enhanced spanning tree protocol (Proposed) at the destination node on network 2 as the transmit rate of the source node on network 6 is increased.*

the source and the destination nodes increases. This is expected because the packets are flooded on the bridged RPR networks between the source and destination networks.

In addition to the total hop count, we also evaluate the enhanced spanning tree protocol by observing the bandwidth available for transmitting packets in the bridged RPR network. For the optimal case, the bandwidth available is derived from the packets taking the path with the smallest hop count to the destination. The lesser the available bandwidth, the lesser is the network efficiency.

We consider the scenario in which the source node on network 6 transmits to the destination on network 5. In addition, packets are transmitted between two nodes on network 2. The path taken by both traffic do not overlap for the optimal case.

The transmit rate of the source node on network 2 is kept constant at 2.4 Gbps while the transmit rate of the source node on network 6 is gradually increased from 0.2 to 2.4 Gbps. The receive rate at the destination nodes on network 5 and network 2 are shown in Figures 7 and 8, respectively.

In Figure 7, the receive rate increases with the transmit rate. However, once the transmit rate exceeds 1.2 Gbps, the receive rate for standard bridging is limited to 1.2 Gbps. The reason is that the packets from the source node on network 6 are flooded on network 2, thereby overlapping with the traffic from the source node on network 2. As a result, the fairness algorithm assigns half of the available bandwidth, or 1.2 Gbps, to each of the traffic. The bridge drops the packets it cannot forward onto network 2 and the receive rate remains constant at 1.2 Gbps. With the enhanced spanning tree protocol, the receive rate continues to increase with the transmit rate until it reaches 2.4 Gbps.

In Figure 8, the receive rate for the destination node on network 2 decreases as soon as the traffic is transmitted from the source node on network 6 with standard bridging. This is because the flooding of packets from network 6 reduces the bandwidth available for the traffic on network 2. This trend continues until the receive rate decreases to 1.2 Gbps, at which point the fairness algorithm ensures that the link is fairly shared between both traffic. On the other hand, the receive rate remains constant at 2.4 Gbps with the enhanced spanning tree protocol.

These two figures show that more bandwidth is available for transmitting packets with the enhanced spanning tree protocol as compared to standard bridging. In fact, the results from the enhanced spanning tree protocol are the same as those derived for the optimal case, showing that the network is used efficiently.

Next, we increase the amount of traffic with the destination on a different network than the source. Traffic is added between a source node on network 3 and a destination node on network 1 and between

another source node on network 3 and a destination node on network 4. The traffic also traverses network 2. Figures 9 and 10 show the receive rate of the traffic at the destination nodes as the transmit rate of the source nodes are increased.

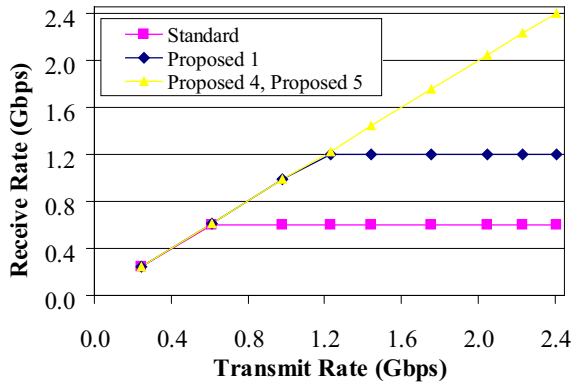

Fig.9: *Receive rate at the destination node on networks 1, 4, and 5 for standard bridging (Standard) and the enhanced spanning tree protocol (Proposed 1, Proposed 4, and Proposed 5, respectively).*

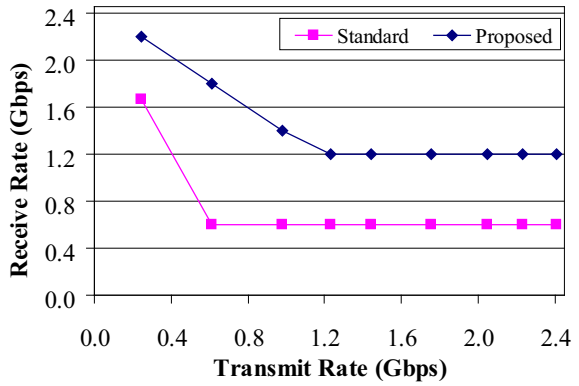

Fig.10: *Receive rate at the destination node on network 2 for standard bridging (Standard) and the enhanced spanning tree protocol (Proposed).*

With standard bridging, the receive rate of each traffic is limited to only 0.6 Gbps. This is because of the flooding of packets on network 2. As the amount of traffic flooding the network increases, the bandwidth available for each traffic decreases.

The results for the enhanced spanning tree protocol are the same as those derived for the optimal case. The traffic for the destination nodes on networks 4 and 5 do not overlap with any other traffic and the receive rate reaches 2.4 Gbps. However, the receive rate for the destination nodes on networks 1 and 2 is limited to 1.2 Gbps. The reason is that the traffic for the destination node on network 1 overlaps with the traffic for the destination node on network 2. Although the traffic takes the optimal path to the destination without flooding, there is still the possibility that it will overlap with other traffic. The probability that the traffic will overlap increases with the amount of traffic with the destination on a different network than the source. At the same time, the probability decreases with the enhanced spanning tree protocol. As a result, the network is used more efficiently with the enhanced spanning tree protocol as compared to standard bridging.

5.2 Evaluation of the Enhanced Topology Discovery Protocol

We evaluate the enhanced topology discovery protocol by observing the hop count of the packet on the source network and comparing the results with standard bridging and the optimal case. For the optimal case, the packet takes the path with the smallest hop count to the bridge on the source network closest to the destination. The greater the hop count, the less efficient is the network. The results for the available bandwidth are omitted because within an RPR network, the greater the hop count, the less the available bandwidth for other nodes to transmit packets.

The simulation results for the source nodes on networks 6, 1, and 2 are shown in Figure 11. In the simulations, all the source nodes transmit to the same destination node on network 5, and the average hop count is calculated for each source network. The average hop count is used in order to eliminate the dependency of the hop count on the location of the source node relative to the bridges on the network. Networks 6, 1, and 2 are chosen because there are 1, 2, and 3 bridges on these networks, respectively.

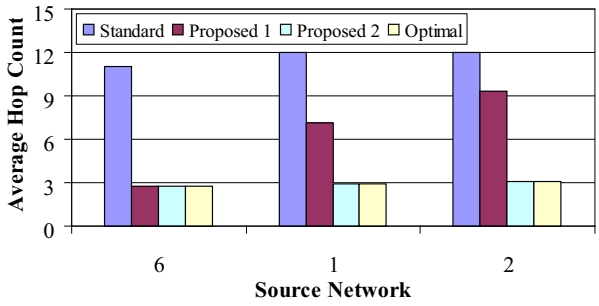

Fig.11: *Average hop count of packets on source networks 6, 1, and 2. The results are shown for standard bridging (Standard), enhanced topology discovery protocol with transparent bridging (Proposed 1), enhanced topology discovery protocol with non-transparent bridging (Proposed 2), and the optimal case (Optimal).*

The results show that standard bridging has the highest average hop count and that the average hop count is equal to the number of nodes on the source network. That is, the packet is flooded on the source network.

For the enhanced topology discovery protocol with transparent bridging, the packets are not flooded and the average hop count is less than that for standard bridging. However, when the number of bridges on the source network is more than one, the average hop count is greater than that derived for the optimal case. The reason is that the default gateway bridge is not always the closest bridge to the destination. For such cases, the hop count will be greater. The results also show that the average hop count increases with the number of bridges. The reason is that as the number of bridges increases, the probability that the default gateway bridge is the closest bridge to the destination is smaller.

For the enhanced topology discovery protocol with non-transparent bridging, the same results as those derived for the optimal case are achieved. This is because the source node is able to identify the closest bridge to the destination node to transmit the packet to. The results show that for networks with more than one bridge, non-transparent bridging achieves the same results as those derived for the optimal case. However, this increases the logic and memory requirements at the node for maintaining the forwarding database.

6. CONCLUSION

The RPR standard for bridging requires that packets are flooded on the bridged RPR network if the packet destination is on a remote network. The network is used inefficiently because of the flooding, and this reduces the bandwidth available for other traffic. In this paper, we propose two solutions – enhanced topology discovery protocol and enhanced spanning tree protocol – for preventing flooding in bridged RPR networks. The nodes use the enhanced topology discovery protocol to prevent flooding on the source network by identifying the default gateway bridge to transmit packets to if the packet destination is on a remote network. The bridges use the enhanced spanning tree protocol to prevent flooding in the bridged RPR network by identifying the bridge closest to the destination to forward the packet to.

The simulation results show that for the proposed solutions, flooding is successfully prevented and the network is used more efficiently as compared to standard bridging. The proposed solutions achieve the same results as those derived for the optimal case, except for source networks with more than one bridge. In this case, the enhanced topology discovery protocol with non-transparent bridging can be used to achieve the same results as those derived for the optimal case.

References

[1] I. Cidon and Y. Ofek. MetaRing, "A Full-Duplex Ring with Fairness and Spatial Reuse," *IEEE Transactions on Communications*, 41(1):110–119, 1993.

[2] F. Davik, M. Yilmaz, S. Gjessing, and N. Uzun, "IEEE 802.17 Resilient Packet Ring Tutorial," *IEEE Communications Magazine*, 42(3):112–118, 2004.

[3] IEEE Standard 802.17. Resilient packet ring (RPR) access method and physical layer specications, 2004.

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

[5] N. Jovanovic, D. Sorgic, T. Ji, and S. Song, "An overview of metropolitan and enterprise networks current and future," In *Proceedings of Canadian Conference on Electrical and Computer Engineer-ing*, pages 160–163, 2005.

[6] 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*, volume II, pages 760–767, French Caribbean, 2004.

[7] S. Spadaro, J. Sole-Pareta, D. Careglio, K. Wa- jda, and A. Szymanski, "Positioning of the RPR standard in contemporary operator environments," *IEEE Network*, 18(2):35–40, 2004.

[8] A. Vargas, "The OMNeT++ Discrete Event Simu- lation System," In *Proceedings of the European Sim- ulation Multiconference ESM'01*, Czech Republic, 2001.

[9] P. Yuan, V. Gambiroza, and E. Knightly, "The IEEE 802.17 Media Access Protocol for High-speed Metropolitan-Area Resilient Packet Rings," *IEEE Network*, 18(3):8–15, 2004.

Pisai Setthawong 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. He is currently pursuing his Ph.D. degree in Telecommunications Science at Assumption University, Thailand. His research interests include protocols for high-speed networks, traffic engineering, and QoS.

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.