

Multicasting in IEEE 802.17 resilient packet ring

Martin Maier

*Institut National de la Recherche Scientifique (INRS), Montréal, QC, H5A 1K6, Canada
maier@emt.inrs.ca*

Michael Scheutzow

Department of Mathematics, Technical University Berlin, 10623 Berlin, Germany

Martin Herzog

*Institut National de la Recherche Scientifique (INRS), Montréal, QC, H5A 1K6, Canada
herzog@emt.inrs.ca*

Martin Reisslein

*Department of Electrical Engineering, Arizona State University, Tempe,
Arizona 85287-5706*

Received June 26, 2006; revised August 31, 2006; accepted September 11, 2006; published October 19, 2006 (Doc. ID 72330)

By using destination stripping and shortest-path routing, the new standard IEEE 802.17 resilient packet ring (RPR) allows for *spatial reuse*, resulting in a significantly increased capacity compared to source-stripping legacy ring networks. We show that for multicast traffic the performance of RPR reduces to that of legacy ring networks that do not support spatial reuse. We propose a bandwidth-efficient and cost-sensitive multicast approach for RPR networks that exploits RPR's built-in *topology discovery* and supplementary time-to-live field to enable spatial reuse for multicast traffic. By means of analysis and simulation we investigate the proposed multicast approach in terms of transmission, multicast, and reception capacities as well as throughput efficiency for different numbers of nodes and multicast fanout under various traffic scenarios. Our findings show that for multicast traffic, the transmission capacity and multicast capacity of RPR are increased significantly, in particular for small to medium multicast group sizes, and a multicast throughput efficiency of 100% is achieved. Besides the performance gain, the proposed multicast approach is able to maintain RPR's simplicity. © 2006 Optical Society of America

OCIS codes: 060.4250, 000.4430.

1. Introduction

Optical packet-switched (OPS) networks have been attracting a great amount of attention over the past few years [1]. OPS networks can be basically categorized into first-generation and second-generation optical networks [2]. In first-generation OPS networks, each intermediate node performs optical-electrical-optical (OEO) conversion, while in second-generation OPS networks, all-optical (OOO) node structures allow in-transit traffic to optically bypass intermediate nodes. Significant progress has been made in the design of all-optical architectures for optical packet switches [3,4]. All-optical packet-switched networks are a promising solution for future high-performance networks, but they currently face technological feasibility issues, e.g., the lack of optical random access memory (RAM). At present, first-generation OPS networks provide a more practical solution to build OPS networks in a cost-effective manner. In this paper, we concentrate on a recently standardized state-of-the-art first-generation OPS ring network. More precisely, we investigate the multicast performance of the recently approved standard Institute of Electrical and Electronics Engineers (IEEE) 802.17 resilient packet ring (RPR) and show how its shortcomings can be overcome by using RPR's built-in mechanisms without sacrificing the simplicity of RPR, which was one of its key design goals [5].

The new standard IEEE 802.17 RPR aims at combining synchronous optical

network/synchronous digital network (SONET/SDH) carrier-class functionalities of high availability, reliability, and profitable time-division multiplexing (TDM) service support and Ethernet's high bandwidth utilization, low equipment cost, and simplicity [6–8]. One of RPR's main performance-enhancing techniques is *spatial reuse*. In RPR, packets are removed from the ring by the destination node instead of the source node. As a result, nodes downstream from the destination node are able to use the bandwidth to transmit data. This so-called destination-stripping technique enables nodes in different ring segments to transmit simultaneously, resulting in spatial reuse and an increased bandwidth utilization compared to source-stripping legacy ring networks where the source node takes its sent packets after one ring round-trip propagation delay from the ring, e.g., the IEEE 802.5 token ring and the American National Standards Institute (ANSI) fiber-distributed data interface (FDDI). For instance, under uniform *unicast* traffic, where each of the N network nodes generates the same amount of traffic and a given data packet is destined to any of the remaining $(N-1)$ nodes with equal probability $1/(N-1)$, the mean hop distance equals approximately $N/4$ hops by using shortest-path routing, whereby one hop denotes the distance between two adjacent nodes. This translates into a spatial reuse factor of 4 on each fiber ring, i.e., on each fiber ring, up to four nodes can send data simultaneously, as opposed to a single transmission in source-stripping ring networks.

Beside unicast (point-to-point) traffic, multidestination (point-to-multipoint) traffic is expected to account for a significant portion of the load on RPR networks due to emerging applications such as videoconferences, multimedia stream distribution, telemedicine, and distributed games. Multidestination traffic is sent by means of *multicasting*, where a single multicast packet is received by multiple destination nodes. In this paper, we investigate multicasting in RPR as defined in IEEE 802.17. As we see below, for multicast traffic the performance of RPR reduces to that of legacy ring networks that do not provide spatial reuse and allow for only a single (multicast) transmission on either ring. As a result, for multicast traffic the performance of RPR is deteriorated significantly. To overcome this shortcoming, we propose and examine an approach that makes use of RPR's built-in topology discovery protocol and improves the multicast performance of RPR dramatically.

Note that the focus of this work is on RPR networks with a ring topology. For more detailed information on the performance evaluation of topologically augmented RPR networks, including fairness and survivability issues, we refer the interested reader to Refs. [9,10].

The remainder of the paper is organized as follows. In the following subsection we review related work. Section 2 provides a brief overview of RPR. In Section 3, we outline the shortcomings of multicasting as defined in IEEE 802.17 and describe a more bandwidth-efficient multicast solution, which is subsequently analyzed in Section 4. Numerical results are presented in Section 5. Section 6 concludes the paper.

1.A. Related Work

Recently, the study and improvement of the performance of RPR networks have received a considerable amount of attention. Most of the previous work on RPR focused on its fairness control. It was shown in Ref. [7] that, under unbalanced and constant-rate traffic inputs, the spatial reuse in RPR is decreased due to severe and permanent oscillations. These oscillations that span nearly the entire range of the link capacity not only hinder spatial reuse but also decrease throughput as well as increase delay jitter. Novel fairness algorithms have been proposed that are able to mitigate the oscillations and achieve nearly complete spatial reuse [11–13]. The stability and convergence of the aggressive-mode (AM) fairness algorithm of RPR was examined in Ref. [14]. An enhanced version of the conservative-mode (CM) fairness algorithm of RPR was proposed and investigated in Ref. [15].

Apart from fairness, other aspects of RPR networks were examined recently. The access delay of RPR networks in steady state was analyzed in Refs. [16,17]. A new scheduling scheme, which helps improve the quality of service for high-priority traffic in RPR networks with a single transit queue, was proposed in Ref. [18]. The wrapping and steering protection techniques of RPR were evaluated with respect to service disruption, packet reordering, and packet loss in Ref. [19]. And finally, for a wavelength-division-multiplexing (WDM) upgrade of current single-channel RPR networks, we refer the interested reader to Ref. [20].

We note that the multicast capacity of a simplified multicast transmission strategy in a tunable transmitter (TT) and fixed-tuned receiver (FR) WDM ring network is formally analyzed in Ref. [21]. The simplified multicast transmission strategy analyzed in Ref. [21] serves a multicast packet by sending *one* packet copy per wavelength channel around the ring to the last multicast destination. In contrast, in this paper we consider the IEEE 802.17 RPR network and propose and formally analyze a multicast transmission strategy that exploits RPR's time-to-live (TTL) field and generally transmits *two* packet copies in both ring directions so as to minimize the network resources required for the multicast. The analysis in Section 4 is fundamentally different from the analysis in Ref. [21] in that it finds the largest gap among all the multicast destination nodes and the sending node along the ring perimeter (whereby the largest gap could lie anywhere along the ring perimeter) so as to minimize the distances traveled by the two multicast copies, whereas the analysis in Ref. [21] only considers the gaps right next to the sending node.

2. Resilient Packet Ring

In this section, we briefly highlight the salient features of RPR. For a more detailed description of RPR we refer the interested reader to Refs. [5–8].

RPR is an optical dual-fiber bidirectional ring network that interconnects up to $N = 255$ nodes. Each fiber ring carries a single wavelength channel. Destination stripping in conjunction with shortest path routing is deployed to improve the spatial reuse of bandwidth. Each node is equipped with two fixed-tuned transmitters (FTs) and two fixed-tuned receivers (FRs), one for each fiber ring. Each node performs optical-electrical-optical (OEO) conversion. Each node has separate (electrical) transit and station queues for each ring. Specifically, for each ring a node has one or two transit queues for in-transit traffic, one transmission queue for locally generated data packets, one reception queue for packets destined for the local node, and one add_MAC (media access control) queue that stores locally generated control packets.

The RPR *topology discovery* provides a reliable and accurate means for all network nodes to discover the topology (number and ordering) of the nodes on the ring network and any changes to that topology due to link or node failures and to added or disconnected nodes. This is achieved by collecting information about all nodes and links via the topology discovery protocol. To this end, each node broadcasts topology and protection (TP) control packets on both rings by means of source stripping. TP control packets are sent periodically and when triggered by protection-state changes. Each node uses received TP control packets to build and update its *topology database*. Among others, the topology database is used by the MAC control entity to compute a complete topology image and to determine the shortest path. Note that the topology database enables each node to calculate the number of hops to all remaining nodes on both rings.

3. Multicasting in Resilient Packet Ring

Next, let us consider multicasting in RPR. First, we review multicasting as defined in the IEEE standard 802.17 and outline the shortcomings. We then propose a more bandwidth-efficient multicast approach for RPR that supports spatial reuse for multicast traffic, as opposed to IEEE 802.17. In the following, we use the terms frame and packet interchangeably.

3.A. Multicasting in IEEE 802.17

According to Ref. [5], a multicast packet is transmitted to a set of nodes whose multicast group membership is identified by a group MAC address carried within the destination address field of the multicast packet. A multicast packet is sent in one of two ways: (i) *unidirectional flooded* or (ii) *bidirectional flooded*. To transmit data packets over an RPR ring, each data packet of RPR MAC clients is supplemented with additional information by the RPR MAC control entity at the RPR MAC layer, among others with a 1-byte TTL field. A flooded packet traverses a sequence of nodes and is stripped from the ring based on the expiration of the TTL field. The TTL field expires when it reaches zero, which indicates that the packet has passed through the intended number of nodes.

In unidirectional flooded multicast transmission, the multicast packet is sent to *all* other nodes on either the clockwise or counterclockwise ring. The multicast packet is taken from the ring by either the source node based on a matching source MAC address or an expired TTL field, whichever comes first, or the node just short of the source node based on an expired TTL field. In bidirectional flooded multicast transmission, the multicast packet is generally replicated and sent to *all* other nodes on both clockwise and counterclockwise rings. A bidirectional flooded multicast packet is stripped from the ring based on the expiration of the TTL field at an agreed upon span, called the *cleave point*. Both multicast packets are not allowed to overlap their delivery and therefore cannot be sent beyond the cleave point. The overlap is prevented by setting the TTL values in both multicast packets to the number of hops it takes from the source node to the cleave point, as determined by the number of hops to the nodes just short of the cleave point. Recall from Section 2 that the number of hops is provided by the topology database available at each node. Unless protection switching has occurred, the cleave point can be put on any link of the ring. Note that the cleave point can also be put immediately adjacent to the source node. In this case, a given multicast packet does not need to be replicated since it is sent only in one direction. For the sake of completeness, we note that in the event of protection switching the cleave point is the point at which the protection switch has taken place. Thus if the ring experiences protection switching, the location of the cleave point is fixed and cannot be chosen arbitrarily.

As detailed above, note that in IEEE 802.17, multicasting is realized by means of broadcasting, either via unidirectional or bidirectional flooded transmissions. That is, each multicast packet is sent to all nodes and thus unnecessarily consumes the bandwidth of $N-1$ or N links, which is equivalent to the bandwidth of one directional ring. As a consequence, for multicast traffic, only a *single* multicast transmission can take place on each directional ring, *preventing RPR from capitalizing on spatial reuse and thus reducing it to legacy ring networks* that do not support spatial reuse. In the following subsection, we propose a more bandwidth-efficient multicast approach for RPR that allows for spatial reuse by exploiting the built-in topology discovery protocol and supplementary TTL field of RPR.

3.B. Multicasting with Spatial Reuse

Note that due to the limited connectivity of ring topologies (each node is connected only to its two adjacent ring nodes), multicast transmissions in RPR are inherently less efficient than in topologies with a higher degree of connectivity, e.g., mesh or star networks. For instance, in Fig. 1 the source node has to send a multicast packet in the clockwise direction to two intermediate regular nodes that do not belong to the corresponding multicast group in order to reach both multicast destination nodes on the left-hand side of the figure via the shortest path in terms of hops. This inherent multicast inefficiency of the underlying ring network architecture of RPR cannot be avoided, unless RPR is augmented by additional (fiber) links, which is a network architecture proposed in Ref. [20]. Multicasting, which is the focus of this paper, is not considered in Ref. [20].

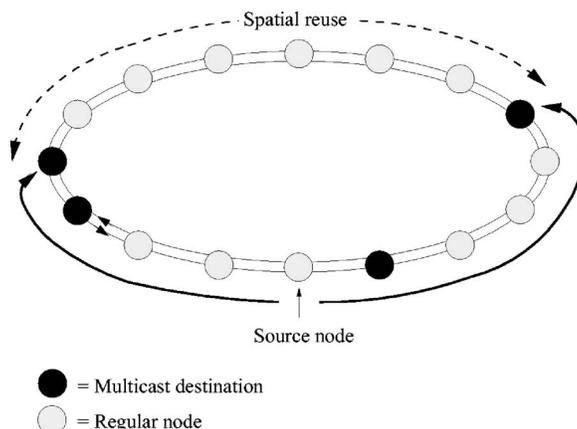


Fig. 1. Multicasting with spatial reuse for a multicast group of four destination nodes among $N=16$ nodes.

At the protocol level, however, multicast transmissions in RPR can be made significantly more efficient by exploiting RPR's built-in topology discovery protocol and supplementary TTL field without requiring any costly upgrades. Specifically, the TTL field of multicast packets can be set such that as few regular nodes as possible (i.e., nodes that do not belong to the respective multicast group) need to be visited. To minimize the amount of wasted bandwidth, the TTL values of both multicast packets have to be set such that they expire after the two multicast destination nodes on the left-hand side and the two multicast destination nodes on the right-hand side of Fig. 1 are reached. In doing so, the multicast packets are stripped from the ring by both final multicast destination nodes instead of being forwarded onward to the cleave point to all remaining regular nodes that do not belong to the respective multicast group. As a result, the multicast packets consume bandwidth only up to the final multicast destination node in either direction, giving rise to spatial reuse on the remaining part of the ring, as shown in Fig. 1.

To achieve multicasting with spatial reuse, a given source node has to execute the following three steps:

1. *Mapping*: The multicast group MAC address is mapped to the individual MAC addresses of all destination nodes belonging to the given multicast group. This mapping may be done by using information provided by the RPR MAC client that generates the given multicast packet.

2. *Localization*: By using its topology database, which gives the number and ordering of all nodes in the RPR network, the source node is able to determine the hop distance to the corresponding multicast destination nodes. The source node is thus able to determine the hop count to each individual multicast destination node.

3. *Routing*: Shortest-path routing is applied, i.e., the source node sends the multicast packet in one or both directions such that the total number of hops required to reach all respective multicast destination nodes is minimized. Note that shortest-path routing is achieved in that multicast packets must not traverse the "largest gap" between any two neighboring multicast destination nodes and the source node. More precisely, the largest gap is given in hops and may be between (i) any two neighboring destination nodes of the given multicast group, or (ii) the source node and the first multicast destination node in either direction. If the largest gap lies between any two neighboring destination nodes of a given multicast group, then the source node sends two multicast packets, one in each direction. Otherwise, if the source node borders on the largest gap, then the source node sends a single multicast packet in the opposite direction (i.e., away from the largest gap). In both cases, the source node sets the TTL value of each multicast packet to the hop distance of the multicast destination node or nodes bordering on the largest gap.

We note that the proposed bandwidth-efficient multicast approach for RPR networks maintains the simplicity of RPR. The proposed multicast approach makes use of RPR's built-in discovery protocol and supplementary TTL field without increasing the complexity of RPR network nodes. Specifically, setting the TTL field properly can be easily done by a source node since each RPR node has its own topology database that contains the hop count to all other RPR nodes. By simply copying the hop count into the TTL field, a given source node can set the TTL field of multicast packets properly. Similarly, each RPR node is able to check the current value in the TTL field of arriving packets. As discussed above in Subsection 3.A, each RPR node reads the TTL field of flooded packets in order to pull packets from the ring whose TTL field has expired. This reading capability at each RPR node can be used in our proposed multicast approach to strip multicast packets from the ring without requiring any modifications of the existing RPR mechanisms.

4. Analysis

In this section, we analyze the proposed multicasting approach with spatial reuse for RPR in terms of capacity, i.e., maximum achievable throughput. Besides multicast traffic, we also consider unicast and broadcast traffic. The performance measures of interest are multicast capacity, reception capacity, and transmission capacity (to be defined shortly).

4.A. Traffic Model

Throughout the analysis, we consider the following traffic model: Each generated packet is destined to F nodes, $1 \leq F \leq N-1$, where F is a random variable denoting the

fanout (number of destination nodes) of a given packet and N denotes the number of network nodes. The distribution of F is given by

$$\mu_l = P(F=l), \quad l = 1, 2, \dots, N-1, \quad (1)$$

where $0 \leq \mu_l \leq 1$ and $\sum_{l=1}^{N-1} \mu_l = 1$. Thus

- for unicast traffic we have $\mu_1 = 1$ and $\mu_l = 0, l = 2, 3, \dots, N-1$;
- for broadcast traffic we have $\mu_{N-1} = 1$ and $\mu_l = 0, l = 1, 2, \dots, N-2$; and
- for multicast traffic we have $\mu_1 = \mu_{N-1} = 0$ and $0 \leq \mu_l \leq 1, l = 2, 3, \dots, N-2$. We consider uniform traffic, i.e., (i) each node generates the same amount of traffic, (ii) a source node must not send any packets to itself, and (iii) the fanout set (set of destination nodes) \mathcal{F} for a given packet with given fanout F is drawn uniformly randomly from among the remaining $(N-1)$ nodes. Note that uniform traffic with any-to-any traffic between all attached nodes is typically found in metropolitan area core ring networks [22] that represent one of the major target applications of RPR. For an analytical and simulative investigation of the throughput-delay performance of RPR networks under uniform, hot-spot, symmetric, and asymmetric traffic demands, we refer the interested reader to Ref. [23].

4.B. Overview and Definition of Largest Gap

We first analyze the mean hop distance traveled by a multicast packet on the RPR ring, which in turn is used to evaluate the capacities. The key idea behind the analysis is to consider the set of nodes containing the source node of the multicast and all the multicast destination nodes. We initially evaluate for a given realization of the number of multicast destination nodes $F=l$ the expected length of the largest gap $g(l, N)$ (in number of hops) that lies between any two neighboring nodes in the considered set, as we express below more formally. The hop distance required to serve the multicast packet is then $N-g(l, N)$ hops. To see this, note that the multicast can be served (i) by sending one copy of the multicast packet in the clockwise direction from the source node to the destination node that borders to the largest gap, and (ii) by sending another copy of the multicast packet from the source node in the counterclockwise direction to the node bordering on the largest gap. If the largest gap borders on the source node, then only one copy of the multicast packet is transmitted, namely, in the direction opposite the largest gap.

In our analysis we suppose without loss of generality that the source node of the multicast is node N . This choice simplifies the notation in the analysis in that the node with index $n, n = 1, \dots, N-1$, has a hop distance of n hops in the clockwise direction from the source node. We introduce the following notation. Let $E_i, i = 1, \dots, l$, be a random variable denoting the ordered indices of the multicast destination nodes, which satisfy $1 \leq E_1 < E_2 < \dots < E_l \leq N-1$; we also define $E_0 := 0$ and $E_{l+1} := N$. Let G_i be a random variable denoting the length of the gap (in number of hops) between the $(i-1)$ th and i th destination node (in the clockwise direction), i.e., $G_i = E_i - E_{i-1}, i = 1, \dots, l+1$. For brevity we use the terminology “the gap has k hops” to express that “the gap has a length of k hops.” Note that $G_1 = E_1$ is the hop distance between the source node and the first destination node in the clockwise direction and that $G_{l+1} = N - E_l$ is the hop distance between the source node and the l th destination node in the counterclockwise direction. Note also that the random variables G_i are identically distributed, but not independent. Let $g(l, N)$ denote the expected length of the largest gap, i.e.,

$$g(l, N) := E[\max\{G_1, \dots, G_{l+1}\}]. \quad (2)$$

Our goal is to evaluate $g(l, N)$ for $1 \leq l \leq N-1$.

4.C. Evaluation of Expected Length of Largest Gap

Toward the evaluation of $g(l, N)$, we first examine the lengths of the gaps G_i in more detail. Without loss of generality, we examine the length of the gap G_1 between the source node and the first destination node of the multicast packet in the clockwise direction and denote $G := G_1$. Note that

$$P(G \geq k) = \frac{\binom{N-k}{l}}{\binom{N-1}{l}} \text{ for } k = 1, 2, \dots, N-l. \tag{3}$$

To see this, note that there are $\binom{N-1}{l}$ ways of selecting the l destination nodes from among the $N-1$ eligible nodes (note that the source node does not send to itself). The gap G_1 between the source node N and the first destination node in the clockwise direction (which has index E_1) has k or more hops if the index of the first destination node is equal to k or larger, i.e., if $E_1 \geq k$. For this inequality to be met, the l destination nodes must all have indices larger than or equal to k ; note that there are $(N-1) - (k-1) = N-k$ nodes with indices larger than or equal to k (and less than N). Hence there are $\binom{N-k}{l}$ ways of selecting the destination nodes such that the first destination node in the clockwise direction is k or more hops away from the source node. Note also that the expected length of the considered gap is

$$E[G] = \frac{N}{l+1}. \tag{4}$$

This can be seen formally by evaluating

$$E[G] = \sum_{k=1}^{\infty} P(G \geq k) \tag{5}$$

$$= \frac{\sum_{k=1}^{\infty} \binom{N-k}{l}}{\binom{N-1}{l}} \tag{6}$$

$$= \frac{\binom{N}{l+1}}{\binom{N-1}{l}} \tag{7}$$

$$= \frac{N}{l+1}, \tag{8}$$

or more intuitively, by noting that $E[G_1] + \dots + E[G_{l+1}] = N$ and that the G_i are identically distributed.

Before we analyze in general the expected length of the largest gap $g(l, N)$, we consider selected special cases. We first note that for unicast traffic, i.e., $l = 1$, we have

$$g(1, N) = \begin{cases} \frac{3}{4}N - \frac{1}{4} & \text{for } N \text{ odd} \\ \frac{3}{4}N^2 - \frac{N}{N-1} & \text{for } N \text{ even} \end{cases}. \tag{9}$$

For broadcast traffic, i.e., $l = N-1$, we have $g(N-1, N) = 1$. For $l = N-2$ we have $g(N-2, N) = 2$, and for $l = N-3$ we have $g(N-3, N) = 3\frac{2}{N-1} + 2\frac{N-3}{N-1} = \frac{2N}{N-1}$ (provided $N \geq 4$). For the other values of the number of multicast destinations l it becomes rather involved to write out explicit expressions for $g(l, N)$. We note that in general $g(l, N)$ is lower bounded by the lengths of the gaps when the destination nodes and the source node are equally spaced over the ring, i.e., $g(l, N) \geq \lceil N/(l+1) \rceil$.

We proceed to develop a recursion that allows for the fast numerical evaluation of $g(l, N)$ for all numbers of multicast destinations $l, l = 1, \dots, N-1$. We let $p_{l,N}(k)$ denote the probability that the considered gap G has k hops, i.e.,

$$p_{l,N}(k) := P(G = k) \tag{10}$$

$$= P(G \geq k) - P(G \geq k+1) \tag{11}$$

$$= \frac{\binom{N-k}{l}}{\binom{N-1}{l}} - \frac{\binom{N-k-1}{l}}{\binom{N-1}{l}} \tag{12}$$

$$= \frac{\binom{N-k-1}{l-1}}{\binom{N-1}{l}}, \quad (13)$$

whereby $\sum_{k=1}^{N-l} p_{l,N}(k) = 1$. Let G_{\max} be a random variable denoting the length of the largest gap, i.e., $G_{\max} = \max\{G_1, \dots, G_{l+1}\}$, and let $q_{l,N}(k) := P(G_{\max} = k)$, $k = 1, \dots, N-1$, denote its distribution. This distribution can be efficiently evaluated with the recursion

$$q_{l,N}(k) = p_{l,N}(k) \times \sum_{m=1}^k q_{l-1,N-k}(m) + \sum_{m=1}^{k-1} p_{l,N}(m) \times q_{l-1,N-m}(k), \quad (14)$$

for $k = 1, \dots, N-1$. Note that $q_{l,N}(k) = 0$ for $k < \lceil N/(l+1) \rceil$ since the largest gap can be no smaller than the gaps when the destination nodes and the source node are equally spaced over the ring. Also, $q_{l,N}(k) = 0$ for $k > N-l$, since the largest gap can have at most $N-l$ hops, which is achieved when the source node and the l destination nodes are direct neighbors on the ring. The recursion is initiated with

$$q_{0,N}(k) = \begin{cases} 1 & \text{for } k = N \\ 0 & \text{for } k < N \end{cases} \quad (15)$$

The reasoning behind the recursion is the following. The largest gap for the transmission to l multicast destinations on the ring with N nodes has k hops if one of two scenarios arises: (A) The gap between the source node and the first destination node in the clockwise direction has k hops (which occurs with probability $p_{l,N}(k)$), and none of the other gaps is larger than k hops. The event that none of the other gaps is larger than k hops is equivalent to the event that the largest gap for the transmission to $l-1$ destination nodes on a ring with $N-k$ nodes is no larger than k hops, which has probability $\sum_{m=1}^k q_{l-1,N-k}(m)$. (B) The first gap has less than k hops, and the largest among the other gaps has k hops. More specifically, suppose the first gap has m , $m = 1, \dots, k-1$ hops [which occurs with probability $p_{l,N}(m)$]. The event that the largest among the other gaps has k hops is then equivalent to the event that the largest gap for the transmission to $l-1$ destinations on a ring with $N-m$ nodes has k hops, which has probability $q_{l-1,N-m}(k)$. Considering the values m , $m = 1, \dots, k-1$, leads to the second summand in Eq. (14).

From the distribution $q_{l,N}(k)$, $k = 1, \dots, N-1$, of the length of the largest gap, we obtain the expected length of the largest gap as

$$g(l,N) = \sum_{k=1}^{N-1} k \times q_{l,N}(k). \quad (16)$$

Note that so far, we have considered a given realization of the number of multicast destinations l . Recall that in our traffic model the number of multicast destinations has distribution μ_l , $l = 1, \dots, N-1$. Hence we obtain the expected length of the largest gap for a given (arbitrary) multicast as $\sum_{l=1}^{N-1} \mu_l \times g(l,N)$ and the corresponding minimum expected hop distance required to serve a multicast as $N - \sum_{l=1}^{N-1} \mu_l \times g(l,N)$.

4.D. Evaluation of Multicast and Reception Capacities

The expected minimum hop distance required to serve a multicast can be used to evaluate the multicast capacity as follows. The multicast capacity C_M gives the maximum mean number of multicasts that can be transmitted simultaneously in the network. This capacity is in general governed by the ring segment (whereby a ring segment is the segment of the fiber between two successive nodes) with the largest utilization. In the uniform traffic scenario with equivalent traffic generation rates in the nodes as considered for the capacity evaluation, all ring segments experience the maximum utilization u_{\max} . This utilization is given for the bidirectional ring with one wavelength channel in each direction by

$$u_{\max} = \frac{N - \sum_{l=1}^{N-1} \mu_l \times g(l,N)}{2N}. \quad (17)$$

Thus we obtain the effective multicast capacity as

$$C_M = \frac{1}{u_{\max}} = \frac{2N}{N - \sum_{l=1}^{N-1} \mu_l \times g(l, N)}. \quad (18)$$

The corresponding reception capacity C_R gives the maximum mean number of simultaneously receiving multicast destination nodes in the network. It is given by $C_R = C_M \times E[F]$ with $E[F] = \sum_{l=1}^{N-1} l \times \mu_l$.

4.E. Evaluation of Transmission Capacity

We define the transmission capacity C_T as the smallest long-run average number of active transmitters required to achieve the multicast capacity. Note that either one packet copy transmission (in either the clockwise or counterclockwise direction) or two packet copy transmissions (in both the clockwise and counterclockwise directions) are required to reach all the multicast destinations with the smallest hop distance. In “tie” situations when the smallest hop distance can be achieved by either sending one or two packet copies, we count only one transmission toward the transmission capacity in order to save on transmission resources.

The main idea behind the evaluation of the transmission capacity is to evaluate the probability $P(1 \text{ copy})$ for the event that a given arbitrary multicast can be served by sending only one copy of the multicast packet. With this probability we then obtain the transmission capacity by noting that

$$C_T = [2 \times P(2 \text{ copies}) + P(1 \text{ copy})] \times C_M \quad (19)$$

$$= [2 - P(1 \text{ copy})] \times C_M, \quad (20)$$

where Eq. (20) follows by noting that the event that two packet copy transmissions are required is complementary to the event that one copy transmission is required.

Toward the evaluation of $P(1 \text{ copy})$ note that

$$P(1 \text{ copy}) = P(\text{source node borders on largest gap}) \quad (21)$$

$$= \sum_{l=1}^{N-1} \mu_l \times P(\text{source node borders on largest gap} | F = l). \quad (22)$$

For convenience we denote

$$\delta_l = P(\text{source node borders on largest gap} | F = l). \quad (23)$$

Let $\beta_{LR}(l)$ denote the probability for the event that (i) the gap between the source node and the smallest indexed multicast destination has $G_1=R$ hops, (ii) the gap between the source node and the highest indexed multicast destination has $G_{l+1}=L$ hops, and (iii) none of the other gaps has more than $\max(L, R)$ hops (i.e., $G_i \leq \max(L, R)$, $i=2, \dots, l$), given that the multicast has l destination nodes. With this definition we have

$$\delta_l = \sum_{L, R=1}^{N-1} \beta_{LR}(l), \quad (24)$$

and we note that for $l=1$ multicast destination, the source node always borders on the largest gap, i.e., $\delta_1=1$. Furthermore, for l , $2 \leq l \leq N-1$, we have

$$\beta_{LR}(l) = p_{l, N}(L) \times p_{l-1, N-L}(R) \times \sum_{k=1}^{\max(L, R)} q_{l-2, N-L-R}(k), \quad (25)$$

whereby $p(\cdot)$ is given by Eq. (13) and $q(\cdot)$ is given by Eq. (14). Inserting the expression for $p(\cdot)$ in Eq. (25) we obtain

$$\beta_{LR}(l) = \frac{\binom{N-L-R-1}{l-2}}{\binom{N-1}{l}} \times \sum_{k=1}^{\max(L, R)} q_{l-2, N-L-R}(k). \quad (26)$$

Some computations in these evaluations can be saved by noting that $\beta_{LR}(l) = \beta_{RL}(l)$.

5. Results

In this section, we present numerical results obtained from our analysis and simulation on the multicast performance of RPR. Specifically, in Subsection 5.A we present the numerical results from our analysis and examine the effect of the number of network nodes N and multicast fanout F on the multicast, transmission, and reception capacities for unicast, multicast, and broadcast traffic as well as different mixes of unicast and multicast traffic. In Subsection 5.B, we provide additional simulation results on the throughput efficiency of RPR. (For a delay performance analysis of RPR we refer the interested reader to Refs. [16,17].)

5.A. Capacity

Figures 2–4 depict the transmission, multicast, and reception capacities versus number of network nodes N , respectively. We vary N from 2 up to 255, which is the maximum number of supported nodes in RPR. In all three figures, we consider uniform unicast, broadcast, and multicast traffic. Clearly, for unicast traffic, we set the distribution of the fanout F to $\mu_1=1$ and $\mu_l=0, 2 \leq l \leq N-1$. For broadcast traffic, we set $\mu_{N-1}=1$ and $\mu_l=0, 1 \leq l \leq N-2$. For multicast traffic, we consider three different maximum multicast fanouts $F_{\max} \in \{1/4(N-1), 1/2(N-1), 3/4(N-1)\}$. More precisely, we let the fanout of a multicast packet be uniformly distributed over the interval $[2, F_{\max}]$, i.e., we set $\mu_1=0$ as well as $\mu_l=0$ for $F_{\max} < l \leq N-1$ and $\mu_l=1/(F_{\max}-1), 2 \leq l \leq F_{\max}$. Note that for now we consider unicast traffic and multicast traffic separately. Mixed unicast–multicast traffic scenarios are examined shortly.

Figure 2 shows the transmission capacity C_T versus number of nodes N . We observe that for $N=2$ the transmission capacity equals $C_T=4$ for both unicast and broadcast traffic. This is because on either ring both unicast and broadcast transmissions require a single hop and thus traverse half the ring circumference. Hence on each ring, two unicast/broadcast transmissions can take place simultaneously, resulting in a transmission capacity of $C_T=4$. For increasing N , the transmission capacity asymptotically approaches $C_T=8$ and $C_T=2$ for unicast and broadcast traffic, respectively. Clearly, under the assumption of uniform traffic and shortest-path routing a unicast packet has to traverse asymptotically a quarter of either ring on average. As a result, on each ring, four unicast transmissions can on average take place simultaneously, which translates into a transmission capacity of $C_T=8$. For broadcast traffic, each broadcast transmission consumes the bandwidth of almost an entire ring ($N-1$ links). Thus for broadcast traffic the dual-fiber ring architecture of RPR asymptotically provides a transmission capacity of $C_T=2$. Note that RPR as defined in IEEE 802.17 would provide the same transmission capacity $C_T=2$ also for multicast traffic due to unidirectional/bidirectional flooding and the lack of spatial reuse. The transmission capacity of RPR under multicast traffic can be increased dramatically by setting the TTL value of multicast packets properly, as proposed in Subsection 3.B, and thereby capitalizing on spatial reuse, as shown in Fig. 2 for different maximum multicast

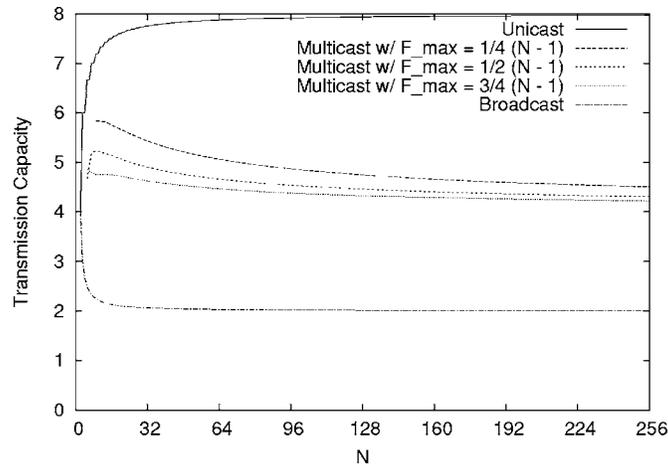


Fig. 2. Transmission capacity C_T versus number of nodes N for unicast, broadcast, and multicast traffic with different maximum multicast fanouts $F_{\max} \in \{1/4(N-1), 1/2(N-1), 3/4(N-1)\}$.

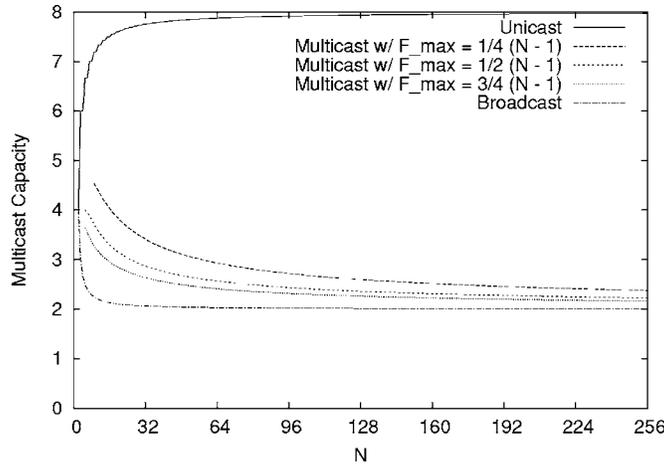


Fig. 3. Multicast capacity C_M versus number of nodes N for unicast, broadcast, and multicast traffic with different maximum multicast fanouts $F_{max} \in \{1/4(N-1), 1/2(N-1), 3/4(N-1)\}$.

fanout $F_{max} \in \{1/4(N-1), 1/2(N-1), 3/4(N-1)\}$. Note that the multicast curves do not start at $N=2$. This is because the fanout of multicast packets is assumed to be at least two, which implies that F_{max} and thus N have to be larger than two. We observe from Fig. 2 that for multicast traffic the transmission capacity C_T generally decreases with increasing F_{max} and increasing N . This is because with increasing F_{max} and N , multicast packets have to traverse more ring nodes in order to reach the intended destination nodes. As a result, more nodes have to forward multicast traffic and are thus prevented from sending locally generated packets, leading to a decreased C_T . Note that for all multicast curves the transmission capacity is larger than four. Hence, by using RPR's built-in topology discovery protocol and setting the TTL value properly, the transmission capacity of RPR is increased by more than 100% for multicast traffic, compared to RPR's current unidirectional/bidirectional flooding of multicast packets that achieves $C_T=2$.

Figure 3 shows the multicast capacity C_M versus number of nodes N . We observe that for unicast and broadcast traffic the multicast capacity C_M equals the transmission capacity C_T of Fig. 2. Clearly, for unicast traffic a single transmission is required to reach the corresponding destination node. Therefore, no packet replication is done and we have $C_M=C_T$. Broadcast packets can be sent either in only one or both directions, each approach providing the same smallest hop distance $(N-1)$. As described in Subsection 4.E, in such a tie situation we give priority to a single broadcast transmission without packet replication. Hence we have $C_M=C_T$ for broadcast traffic as well. As opposed to unicast and broadcast traffic, however, packet replication is generally

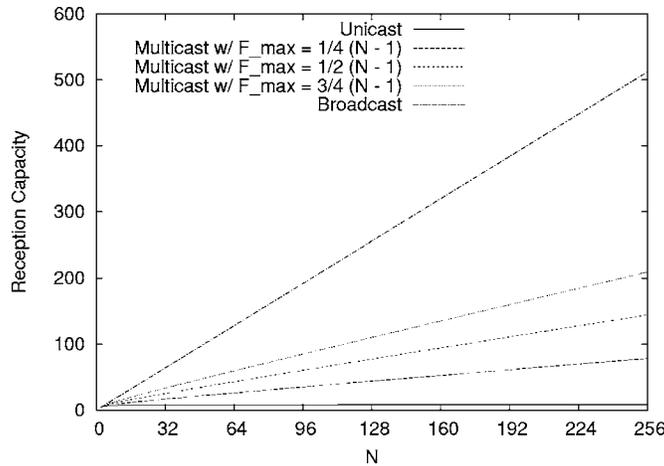


Fig. 4. Reception capacity C_R versus number of nodes N for unicast, broadcast, and multicast traffic with different maximum multicast fanouts $F_{max} \in \{1/4(N-1), 1/2(N-1), 3/4(N-1)\}$.

required for multicast traffic in order to minimize the hop distance. More precisely, for a medium to large number of nodes N , multicast packets generally have to be replicated and sent in both directions, which holds especially for an increasing maximum multicast fanout F_{\max} . As a consequence, for medium to large N , the multicast capacity is smaller than the transmission capacity for multicast traffic, as can be observed from Figs. 2 and 3. However, for small N , the multicast capacity is only slightly smaller than C_T . This indicates that for small N , many multicast packets do not need to be replicated and must be sent in only one direction. Similar to C_T in Fig. 2, we observe that C_M decreases with increasing F_{\max} and increasing N , since more nodes are prevented from transmitting locally generated traffic due to the increased traffic-forwarding burden. Note that, especially for small to medium values of N , the proposed multicasting approach clearly outperforms RPR's current flooding that corresponds to the broadcast curve in Fig. 3. For large N the multicast capacity for multicast traffic approaches that of broadcast traffic since multicast packets have to traverse almost all nodes in order to reach the intended multicast destination nodes.

Figure 4 shows the reception capacity C_R versus number of nodes N . For unicast traffic the reception capacity equals approximately $C_R=8$ independent of N , since each transmitted packet is received by a single destination node. As shown in Fig. 4, C_R increases with increasing N for multicast and broadcast traffic because each multicast and broadcast packet is destined to more nodes. Furthermore, we observe that for a given number of nodes N , the reception capacity increases with increasing multicast fanout. Again, this is because with increasing fanout, each multicast packet is received by more destination nodes. Under broadcast traffic each transmitted packet is received by all nodes except the source node, resulting in an asymptotic reception capacity of $C_R=2(N-1)$, where two source nodes send two broadcast packets, each destined for the remaining $(N-1)$ nodes on either directional ring.

So far, we have investigated unicast and multicast traffic separately for different multicast fanout. We next study the effect of mixed unicast–multicast traffic on the transmission, multicast, and reception capacities of RPR. In Figs. 5–7 we set the maximum multicast fanout to $F_{\max}=N-1$. Beside unicast, multicast, and broadcast traffic we consider the following unicast–multicast traffic ratios: 25% multicast/75% unicast, 50% multicast/50% unicast, and 75% multicast/25% unicast.

Figure 5 depicts the transmission capacity C_T versus number of nodes N for unicast, multicast, broadcast traffic, and different unicast–multicast traffic ratios. Again, for unicast and broadcast traffic the transmission capacity equals $C_T=4$ for $N=2$ and asymptotically approaches $C_T=8$ and $C_T=2$ for increasing N , respectively. Given that the maximum multicast fanout equals $F_{\max}=N-1$, we let the multicast curves in Fig. 5 start at $N=3$ to guarantee a multicast fanout of $F \geq 2$. For multicast and mixed unicast–multicast traffic we observe that for small N , the transmission capacity initially increases and then levels off with increasing N . Similar to unicast traffic, this is because for small N , unicast and multicast packets have to traverse relatively large portions of the ring perimeter such that only a few transmissions can take place

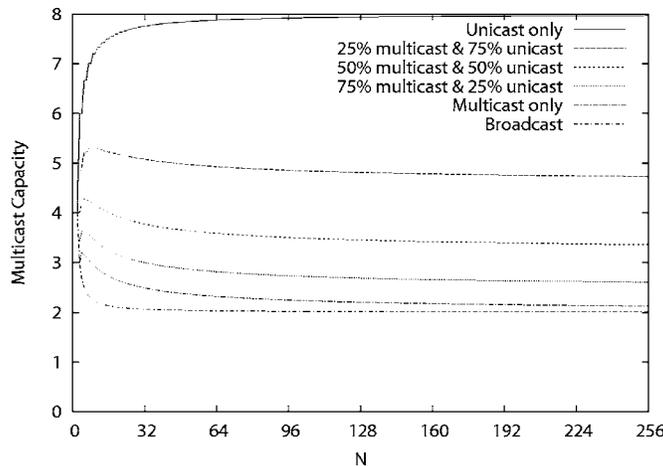


Fig. 5. Transmission capacity C_T versus number of nodes N for unicast, multicast, broadcast traffic, and different mixed unicast–multicast traffic ratios with maximum multicast fanout $F_{\max}=N-1$.

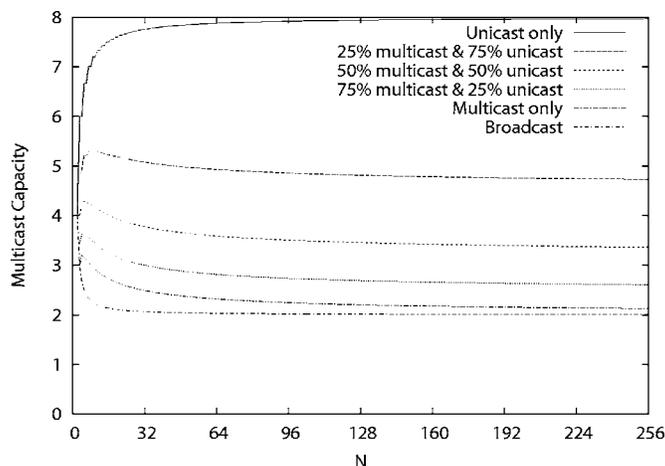


Fig. 6. Multicast capacity C_M versus number of nodes N for unicast, multicast, broadcast traffic, and different unicast-multicast traffic ratios with maximum multicast fanout $F_{\max}=N-1$.

simultaneously, resulting in a rather small C_T . By slightly increasing the number of nodes up to approximately $N=16$, the unicast and multicast transmissions traverse smaller ring segments, leading to more simultaneous transmissions and an increased C_T . By further increasing the number of nodes beyond approximately $N=16$, each multicast packet is destined to more nodes and thus consumes more ring bandwidth resources and diminishes spatial reuse. As a result, the transmission capacity slightly decreases with increasing N . As shown in Fig. 5, for a given number of nodes N the transmission capacity decreases if the fraction of multicast traffic increases. Clearly, this is because with an increasing fraction of multicast traffic, more nodes have to forward traffic, leading to an increased forwarding burden and a decreased transmission capacity.

Figure 6 depicts the multicast capacity C_M versus number of nodes N for unicast, multicast, broadcast traffic, and different unicast-multicast traffic ratios. Figure 6 illustrates that for mixed unicast-multicast traffic scenarios, which are typically found in operational networks, the transmission capacity C_M is significantly increased over a wide range of N compared to multicast-only traffic. For a mix of 25% multicast traffic and 75% unicast traffic, the multicast capacity equals approximately five, which is rather large given that we assume the maximum possible multicast fanout of $F_{\max}=N-1$.

Figure 7 shows the reception capacity C_R versus the number of nodes N for unicast, multicast, broadcast traffic, and different unicast-multicast traffic ratios. As shown in the figure, C_R increases with increasing number of nodes N and increasing fraction of

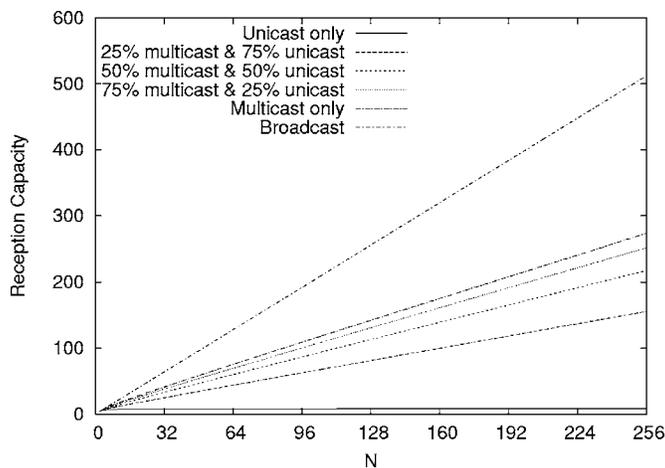


Fig. 7. Reception capacity C_R versus number of nodes N for unicast, multicast, broadcast traffic, and different unicast-multicast traffic ratios with maximum multicast fanout $F_{\max}=N-1$.

multicast traffic due to the larger number of receiving destination nodes. Note that for both scenarios 50% multicast/50% unicast and 75% multicast/25% unicast, C_R is larger than that for the multicast-only traffic of Fig. 4. This is due not only to the larger maximum multicast fanout of $F_{\max}=N-1$ but also to the increased multicast capacity C_M (see Figs. 3 and 6).

5.B. Throughput Efficiency

In this subsection, we investigate the mean aggregate transmitter throughput, multicast throughput, and receiver throughput of RPR by means of simulation, and we compare them to the above analyzed transmission, multicast, and reception capacity, respectively, in order to evaluate the respective throughput efficiency, which is defined as the ratio of throughput and capacity. The mean aggregate transmitter and receiver throughput is equal to the mean number of simultaneously transmitting and receiving nodes in steady state, respectively. The mean aggregate multicast throughput equals the mean number of simultaneous multicast transmissions in steady state.

As mentioned in Subsection 1.A, several improved fairness control protocols have been proposed for RPR. In our simulations, we adopt a fairness control protocol similar to that in Ref. [11], as explained in the following. To establish fair transmission rates, one fairness control packet circulates on each directional ring. Each fairness control packet consists of N fields that contain the fair rates of all links of the other directional ring. Each node monitors both fairness control packets and writes its locally computed fair link rates in the corresponding fields of the fairness control packets. To calculate the fair link rates, each node measures the number of bytes l_k arriving from node k , $1 \leq k \leq N$, including the node itself, during the time interval T between the previous and actual arrival of the control packet. Each node performs separate measurements for either directional ring using two separate time windows. The fair rate, F , of a given link is equal to the max-min fair share among all measured link rates l_k/T with respect to the link bandwidth available to fairness-eligible traffic. Each node limits the data rate of its $(N-1)$ ingress flows, one for each destination, by using token buckets whose refill rates are set to the current fair rates of the corresponding destinations, which are calculated as follows. Using the same time window T as in the calculation of the fair link rates above, each source node i , $1 \leq i \leq N$, counts the bytes ρ_{ij} sent to destination node j during the time window T . Thus there are two sets of $(N-1)$ byte counters, one for each time window. Each time a fairness control packet arrives, a given source node calculates the fair rates as follows. According to the Ring Ingress-Aggregated with Spatial Reuse (RIAS) fairness objective [11], the bandwidth available to a given source node on a certain link equals the fair rate F of this link that is shared among all its ingress flows crossing this link. Based on the measured ingress rates ρ_{ij}/T of these flows and the available link bandwidth F , the max-min fair share, f , is calculated for each crossed link. The refill rate of the corresponding token buckets is set to the minimum fair share f of these links.

In the subsequent simulations, we consider both uniform Poisson and self-similar traffic with Hurst parameter 0.75. The self-similar traffic is generated out of ON-OFF sources with Pareto distributed ON periods and exponentially distributed OFF periods. For Poisson traffic we provide 95% confidence intervals by using the batch mean method. We consider only best-effort (class C) traffic, which is fully fairness eligible in RPR. Furthermore, we consider the typical trimodal Internet protocol (IP) packet size distribution, i.e., 50% 40 byte packets, 30% 552 byte packets, and 20% 1500 byte packets. We examine a mixed traffic scenario of 75% unicast packets and 25% multicast packets, where the maximum fanout of multicast packets is set to $F_{\max}=N-1$. We concentrate on the single-queue mode of RPR, i.e., each node has a primary transit queue (PTQ) but no secondary transit queue (STQ). The network parameters are set to the following values: number of nodes $N=32$, ring circumference equals 100 km, size of the PTQ equals 6000 bytes, size of the transmit queue equals 15000 bytes, and each directional ring operates at a line rate of 2.5 Gbit/s. Each simulation was run for a total of 10^7 generated packets, including a warm-up phase of 10^6 packets.

Figures 8–10 depict the mean aggregate transmitter, multicast, and receiver throughputs versus mean arrival rate under Poisson and self-similar traffic (for Poisson traffic the confidence intervals are too small to be seen). The figures also show the transmission, multicast, and reception capacities for $N=32$ and a traffic mix of 75% unicast and 25% multicast traffic with a maximum multicast fanout $F_{\max}=N-1=31$.

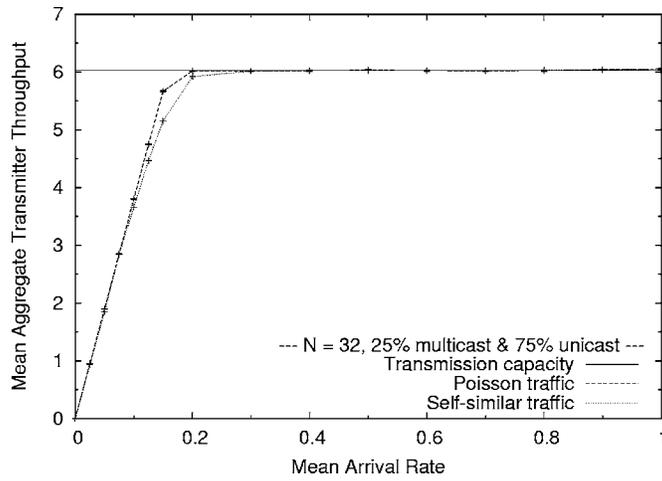


Fig. 8. Mean aggregate transmitter throughput versus mean arrival rate under uniform Poisson and self-similar traffic for 25% multicast traffic and 75% unicast traffic with $N=32$ and maximum multicast fanout $F_{\max}=N-1=31$.

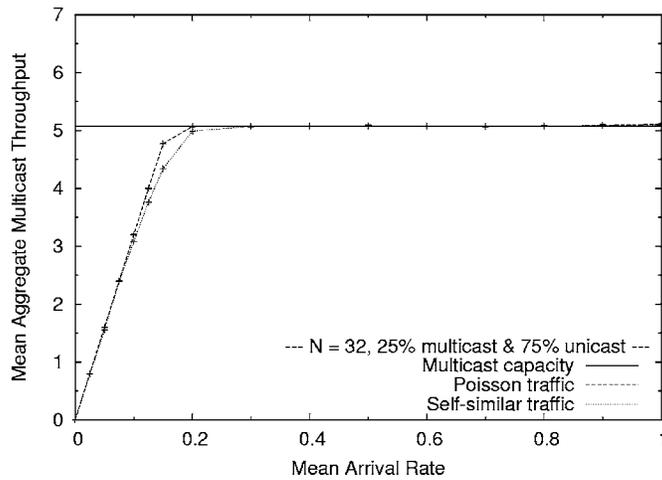


Fig. 9. Mean aggregate multicast throughput versus mean arrival rate under uniform Poisson and self-similar traffic for 25% multicast traffic and 75% unicast traffic with $N=32$ and maximum multicast fanout $F_{\max}=N-1=31$.

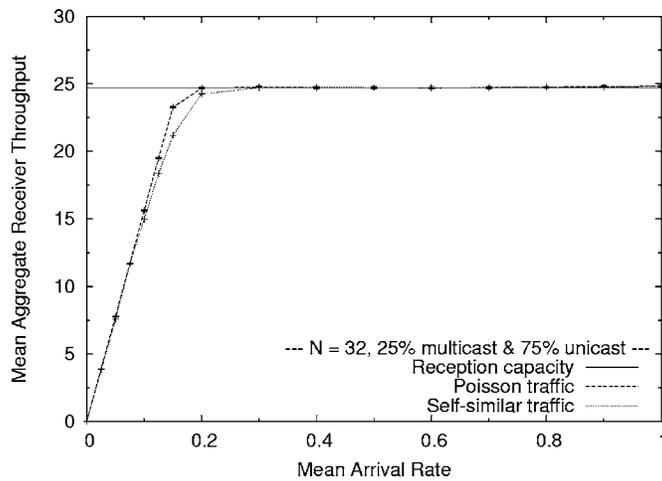


Fig. 10. Mean aggregate receiver throughput versus mean arrival rate under uniform Poisson and self-similar traffic for 25% multicast traffic and 75% unicast traffic with $N=32$ and maximum multicast fanout $F_{\max}=N-1=31$.

Apparently, the simulation results match very well with the analytical results in that, for an increasing mean arrival rate, the mean aggregate transmitter, multicast, and receiver throughputs approach the transmission, multicast, and reception capacities, respectively, for both Poisson and self-similar traffic. Thus by using the proposed multicast approach, RPR is able to achieve a multicast throughput equal to the improved multicast capacity, which translates into a multicast throughput efficiency of 100%. We note that for small to medium mean arrival rates, self-similar traffic achieves a slightly smaller throughput than Poisson traffic. This is because self-similar traffic is more bursty than Poisson traffic, resulting in fewer backlogged packets at each node's transmit queue. At medium to high mean arrival rates, the difference between Poisson and self-similar traffic disappears, since the transmit queues are constantly fully occupied. Clearly, at medium to high traffic loads, Poisson and self-similar traffic differ with respect to other performance metrics, e.g., delay. For a delay analysis of RPR we refer the interested reader to Refs. [16,17].

6. Conclusions

Resilient packet ring (RPR) networks as defined in IEEE 802.17 realize multicasting by means of unidirectional or bidirectional flooding. As a result, spatial reuse is prevented, and RPR's capacity reduces to that of legacy ring networks. By using RPR's topology discovery protocol and supplementary TTL (time-to-live) field, which is added to each packet by the RPR MAC control entity, multicast transmissions in general have to traverse fewer ring nodes, giving rise to spatial reuse. The proposed multicast approach does not require any costly upgrades of current RPR hardware and software and thus maintains the simplicity of RPR, which was one its key design goals, and also meets the requirements of cost-sensitive metropolitan area networks, which is the major deployment area of RPR. Furthermore, by means of analysis and simulation, we have shown that the proposed multicast approach is able to increase the transmission capacity of RPR by a factor of more than 2 and to achieve a multicast throughput efficiency of 100% and a significantly improved multicast capacity, especially for small to medium multicast group sizes.

Finally, we note that this work focused on multicasting in the standard IEEE 802.17, which was approved in June 2004. The standardization of further enhancements of IEEE 802.17 RPR is currently underway. Notably, the so-called spatially aware sublayer (SAS), IEEE 802.17b, provides spatial reuse via layer 2 bridging. The IEEE 802.17b draft is expected to be ratified in the second half of this year. SAS allows RPR nodes to learn about the MAC addresses of clients attached to RPR nodes. When a destination MAC address appears on the RPR ring for the first time, the corresponding packet is broadcast to all RPR nodes on the ring. In doing so, all RPR nodes learn about the destination MAC address and the corresponding RPR node to which the destination node is connected. Each RPR node uses this knowledge subsequently to send traffic only to the RPR node corresponding to the packet's destination. As a result, the broadcasting of packets to all RPR nodes on the ring is avoided, giving rise to spatial reuse. Note that SAS implies two types of traffic on the RPR ring network: broadcast traffic during the learning phase and unicast traffic afterward. Both types of traffic have been taken into account in our analysis and simulation and compared to multicast traffic.

References

1. B. Y. Yu, P. Toliver, R. J. Runser, K.-L. Deng, D. Zhou, I. Glesk, and P. R. Prucnal, "Packet-switched optical networks," *IEEE Micro* **18**(1), 28–38 (1998).
2. R. Ramaswami and K. N. Sivarajan, *Optical Networks—A Practical Perspective*, 2nd ed. (Morgan Kaufmann, 2001).
3. M. C. Chia, D. K. Hunter, I. Andonovic, P. Ball, I. Wright, S. P. Ferguson, K. M. Guild, and M. J. O'Mahony, "Packet loss and delay performance of feedback and feed-forward arrayed-waveguide gratings-based optical packet switches with WDM inputs-outputs," *J. Lightwave Technol.* **19**, 1241–1254 (2001).
4. K. Onohara, H. Sotobayashi, K. Kitayama, and W. Chujo, "Photonic time-slot and wavelength-grid interchange for 10-Gb/s packet switching," *IEEE Photon. Technol. Lett.* **13**, 1121–1123 (2001).
5. IEEE, "IEEE standard 802.17: resilient packet ring," (September 2004); <http://ieee802.org/17>.
6. F. Davik, M. Yilmaz, S. Gjessing, and N. Uzun, "IEEE 802.17 resilient packet ring tutorial,"

- IEEE Commun. Mag. **42**(3), 112–118 (2004).
7. P. Yuan, V. Gambiroza, and E. Knightly, "The IEEE 802.17 media access protocol for high-speed metropolitan-area resilient packet rings," *IEEE Netw.* **18**(3), 8–15 (2004).
 8. S. Spadaro, J. Solé-Pareta, D. Careglio, K. Wajda, and A. Szymański, "Positioning of the RPR standard in contemporary operator environments," *IEEE Netw.* **18**(2), 35–40 (2004).
 9. M. Herzog and M. Maier, "RINGOSTAR: an evolutionary performance-enhancing WDM upgrade of IEEE 802.17 resilient packet ring," *IEEE Commun. Mag.* **44**(2), S11–S17 (2006).
 10. M. Maier, M. Herzog, M. Scheutzow, and M. Reisslein, "PROTECTORATION: a fast and efficient multiple-failure recovery technique for resilient packet ring (RPR) using dark fiber," *J. Lightwave Technol., Special Issue on Optical Networks*, **23**, 2816–2838 (2005).
 11. 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 Trans. Netw.* **12**(1), 85–102 (2004).
 12. F. Alharbi and N. Ansari, "A novel fairness algorithm for resilient packet ring networks with low computational and hardware complexity," presented at the 13th IEEE Workshop on Local and Metropolitan Area Networks, Mill Valley, California, April 2004.
 13. Y. Robichaud, C. Huang, J. Yang, and H. Peng, "Access delay performance of resilient packet ring under bursty periodic class B traffic load," in *Proceedings of the IEEE International Conference on Communications* (IEEE, 2004), Vol. 2, pp. 1217–1221.
 14. F. Davik and S. Gjessing, "The stability of the resilient packet ring aggressive fairness algorithm," presented at the 13th IEEE Workshop on Local and Metropolitan Area Networks, Mill Valley, California, April 2004.
 15. D. Wang, K. K. Ramakrishnan, and C. Kalmanek, "Congestion control in resilient packet rings," in *Proceedings of the 12th IEEE International Conference on Network Protocols* (IEEE, 2004), pp. 108–117.
 16. C. Huang, H. Peng, F. Yuan, and J. Hawkins, "A steady state bound for resilient packet rings," in *Proceedings IEEE GLOBECOM* (IEEE, 2003), Vol.7, pp. 4054–4058.
 17. C. Huang, H. Peng, and F. Yuan, "A deterministic bound for the access delay of resilient packet rings," *IEEE Commun. Lett.* **9**(1), 87–89 (2005).
 18. J. Zhu, A. Matrawy, and I. Lambadaris, "A new scheduling scheme for resilient packet ring networks with single transit buffer," in *Proceedings IEEE GLOBECOM Workshops* (IEEE, 2004), pp. 276–280.
 19. A. Kvalbein and S. Gjessing, "Analysis and improved performance of RPR protection," in *Proceedings of the IEEE International Conference on Networks* (IEEE, 2004), Vol.1, pp. 119–124.
 20. M. Herzog, M. Maier, and A. Wolisz, "RINGOSTAR: an evolutionary AWG-based WDM upgrade of optical ring networks," *J. Lightwave Technol.* **23**, 1637–1651 (2005).
 21. M. Scheutzow, P. Seeling, M. Maier, and M. Reisslein, "Multicast capacity of packet-switched ring WDM networks," in *Proceedings IEEE INFOCOM* (IEEE, 2005), Vol. 1, pp. 706–717.
 22. N. Ghani, J.-Y. Pan, and X. Cheng, "Metropolitan optical networks," *Opt. Fiber Telecommun.* **IVB**, 329–403 (2002).
 23. M. Herzog, S. Adams, and M. Maier, "Proxy stripping: a performance-enhancing technique for optical metropolitan area ring networks," *J. Opt. Netw.* **4**(7), 400–431 (2005).