Positioning of the RPR Standard in Contemporary Operator Environments

Salvatore Spadaro, Josep Solé-Pareta, and Davide Careglio
Universitat Politècnica de Catalunya (UPC)
Krzysztof Wajda and Andrzej Szymański, AGH University of Science and Technology

Abstract

This article deals with the fundamentals and current standardization efforts for IEEE 802.17 Resilient Packet Ring. Its special resilience features make this technology robust against outages of the network infrastructure. The goals of this article are threefold. First, the fundamentals of RPR and the standardization process carried out under the auspices of IEEE and ITU are overviewed. Second, potentially hazardous situations involving traffic assignments are defined and illustrated. Finally, possible situations where the simplicity, enhanced throughput, and automatic resilience features of RPR may be advantageous for network operators are identified.

IEEE 802.17, Resilient Packet Ring (RPR), is a new technology for ring-based metropolitan area networks that enables efficient transfer of data traffic as well as fast protection mechanisms. It is a protocol that consists of a superset of features derived from various proprietary solutions such as Cisco Systems' Dynamic Packet Transport (DPT) and Nortel Networks' Optera. The work on defining a standard for RPR started in December 2000; the final version of the standard is scheduled for publication in 2004.

Operators claim that the functionality of RPR and its implementation in real commercial environments present many advantages:

- Advanced protection mechanisms
- Distributed control
- Interoperability with major transmission standards
- Scalability in speed and number of nodes
- Plug-and-play operation
- Performance monitoring capabilities
- Support for a limited number of priorities (two or three)
- Operations, administration, and management (OAM) and advanced traffic and bandwidth management
- Support for unicast, multicast, and broadcast data traffic

The unique features of RPR were sufficiently interesting to trigger many prestandard installations by important players in the telecommunications market (e.g., Sprint, Luminous, Bell Canada, WorldCom, and SUNET). The first major pre-IEEE 802.17 RPR standard deployments were DPT networks, introduced by Sprint in 1999 and Macedonia Telecom and China Telecom in 2001.

The remainder of this article is organized as follows. We overview IEEE 802.17 RPR technology. The behavior of RPR technology with realistic traffic streams is illustrated. A potentially dangerous situation in which there is a significant reduction in achievable throughput is described. We discuss the strengths and weaknesses of RPR. Finally, the state of standardization and RPR deployment is presented.

Fundamentals of RPR Technology

RPR technology, which will be standardized as IEEE 802.17 RPR [1], is based on two symmetric counter-rotating rings that carry data and control information. It is designed to operate over a variety of physical layers, including SONET/SDH, Gigabit Ethernet (IEEE 802.3ab), dense wavelength-division multiplexing (DWDM), and dark fiber, and is expected to work over higher-speed physical layers. The minimum supported data rate is 155 Mb/s.

An important feature of RPR technology is spatial reuse, which increases the overall aggregate bandwidth of the ring. Unicast frames are removed from the ring at their destination, which means they only occupy bandwidth on the links from source to destination. This is in contrast to earlier techniques, such as fiber distributed data interface (FDDI) and token ring, in which each frame had to traverse the whole ring, so spatial reuse could not be exploited.

The IEEE 802.17 RPR standard supports three types of services: class A (high-priority), class B (medium-priority), and class C (low-priority). Class A service is designed to support real-time applications that require a guaranteed bandwidth and low jitter. This service has absolute priority over the other types of services, and must be shaped at the ingress. A token bucket shaper (shape A in Fig. 1) is provided to ensure that the client traffic does not exceed the allocated rate. Each node/station advertises the amount of bandwidth it needs for its class A service. This allows calculating how much bandwidth is reserved for class A in the ring and how much is left unreserved for class B and C services. Traffic above the allocated rate is rejected.

Class B service is dedicated to near-real-time applications that are less delay-sensitive but still require some bandwidth guarantees. It provides guaranteed information transfer at the committed rate (CIR) and best effort transfer for excess traffic (beyond the committed rate). In contrast to class A, the bandwidth for class B CIR traffic is not statically allocated. In the presence of congestion, the node sends messages that throttle class C transmissions from other stations to leave

0890-8044/04/$20.00 © 2004 IEEE

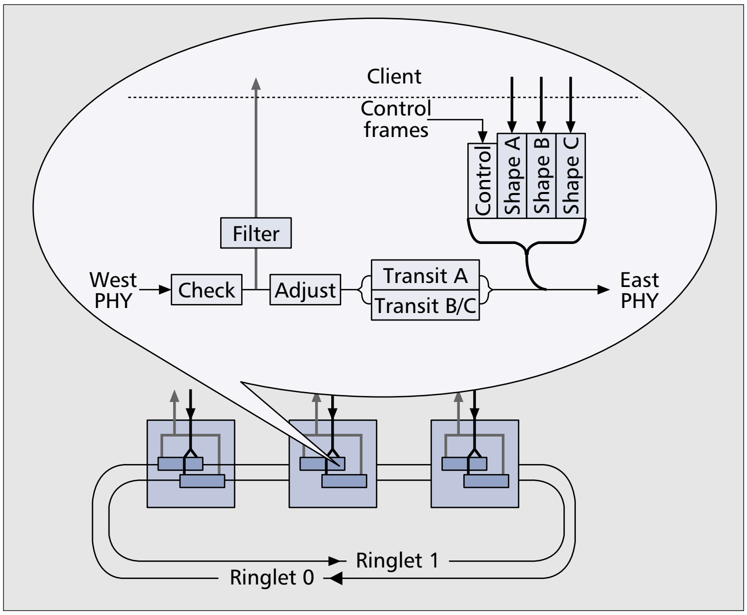

■ Figure 1. *A three-node IEEE 802.17 RPR ring with a simplified structure for the MAC datapath entity.*

bandwidth for its class B traffic. Class C service implements the best effort traffic class. This service is subject to weighted fairness mechanisms, which ensure that each station gets its fair share of the bandwidth available. The traffic is shaped by the IEEE 802.17 RPR medium access control (MAC), which uses a token bucket shaper. A fairness mechanism decides on the amount of bandwidth each station may currently use for its class C transmission. The calculation involves determining the amount of class A and B traffic present in the ring and dividing the remaining bandwidth in proportion to administratively configured node weights.

The allocated rates for class A and B services and node weights for class C are configured in each station by a provisioning mechanism. The provisioning mechanism, however, is outside the scope of the standard.

An IEEE 802.17 RPR node may use virtual output queuing to avoid head-of-line blocking for frames destined to nodes that are physically closer than the congestion point. This is called multi-choke implementation, which requires detailed awareness of congestion points in the whole ring but increases ring utilization and spatial reuse.

Figure 1 presents an example of a three-node IEEE 802.17 RPR ring and a more detailed view of the MAC datapath entity. Each node has two MAC datapath entities, one for each ringlet.

Figure 1 shows that a frame received from an IEEE 802.17 RPR ring is checked against bit errors and time-to-live expiration. Once this is performed, a filter module decides whether the frame should be copied to the client, passed to the control sublayer, or neither. The adjust function is responsible for stripping frames from the ring, adjusting frame fields (e.g., the time-to-live field), and placing the frame in the correct transit queue. The node described in this article has two transit queues, one for class A service and the other for classes B and C. An alternative implementation is characterized by a single transit queue.

More than one frame may be ready for transmission at a given moment. Two transit buffers, the control queue, and three queues corresponding to class A, B, and C local services may simultaneously hold a frame ready to be transmitted. A set of precedence rules is defined to maintain traffic priorities and avoid loss of frames in transit.

In a more general view, RPR nodes may send data on either of the two ringlets. In most cases, the shortest path to the destination is used. The nodes use a topology discovery protocol to obtain a topology map of the ring, which is then used for the shortest path computation.

The ringlets allow uninterrupted operation in the case of node or link failure. Two protection mechanisms may be used, steering and wrapping, both of which provide fast protection switching comparable with that of synchronous optical network/synchronous digital hierarchy (SONET/SDH) networks (50 ms). Neither of these require dedicated protection resources. Steering is based on the ability to choose the ringlet on which the data is sent. If the preferred path is unavailable due to a failure, the other path is used. All IEEE 802.17 RPR implementations must have steering capabilities. Wrapping protection, on the other hand, works as follows. If a failure condition is detected, the traffic going toward the failure is looped onto the opposite ringlet by the nodes adjacent to the failure. The implementation of wrapping protection in the nodes is optional. Both protection modes may be mixed in a wrap-then-steer mode where wrapping protection is activated first to avoid loss of frames in transit; following this, nodes switch to steering to improve ring utilization.

All the stations in the ring must use the same protection method; the default method in IEEE 802.17 RPR is steering. If, however, all the nodes support wrapping, the ring may be configured to use wrapping protection.

Failure detection may be carried out in two ways. The first is based on messages received from the physical layer, such as loss of signal (LOS) from the SONET/SDH layer, and the second on periodical continuity checks within the IEEE 802.17 RPR layer. When RPR is used over a physical media already equipped with protection mechanisms (e.g., an optical transport network, OTN), RPR supports a holdoff timer, which postpones RPR protection to avoid simultaneous recovery in both layers.

Although based on a similar concept, IEEE 802.17 RPR differs from DPT, one of the most popular of its predecessors, described in IETF RFC 2892 [2]. Its main differences are that it has three traffic classes instead of two, a shaping mechanism for high-priority traffic, optional station weights for low-priority traffic, steering as an additional protection mechanism, support for a broader range of physical media and line speeds, and support for bridging. IEEE 802.17 RPR also has a richer set of configurable parameters and OAM functions, such as a loopback (echo) request that may be used to check connectivity between two stations.

The process of RPR technology standardization had not yet been finalized when this article was written. Thus, it is possible that the functionality of the final IEEE 802.17 RPR standard will differ from the information presented here.

Ring topology based on RPR is also being studied by the International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) and a preliminary version of the Recommendation on “Multiple Services Ring” (X.msr) is available [3]. Table 1 summarizes the most significant differences between X.msr and IEEE 802.17 RPR.

Since the beginning of 2003 there has been ongoing cooperation between the ITU-T SG17 and IEEE 802.17 working groups, aimed at closing open items and determining the appropriate method for supporting X.msr within IEEE 802.17 RPR.

Nominal Traffic Assignments

We have carried out various studies to evaluate the behavior of RPR rings under different load conditions and different low and high-priority traffic streams [4, 5]. In [4] we aimed at verification and validation of the features of the Spatial Reuse

Feature	ITU-T (X.msr)	802.17 (RPR)
Topology	Two counter-rotating rings, max. 32 stations	*N* × dual counter-rotating rings, max. 256 stations
MAC address	Local with fixed addresses (4 octets) — possibly IP address	Globally unique MAC address (6 octets)
MAC transit	Unspecified buffer, 8 priorities	Single or dual buffers, 2 priorities
Protection	Wrapping	Wrapping and steering
Spatial reuse	Supported	Supported
Fairness	Not necessary (preplanned bandwidth)	Fairness algorithm for unprovisioned traffic
Multicast	Supported	Supported

■ Table 1. *Significant differences between X.msr and IEEE 802.17 RPR.*

Protocol fairness algorithm, the DPT equivalent of the IEEE 802.17 RPR MAC protocol; particular attention was paid to those features related to efficient use of bandwidth, fair access of different nodes to the ring, and support of priority traffic. In [5], a preliminary simulation case study was done to assess the resilience features of RPR rings.

In this section we extend and conclude the work done in [5]. A metropolitan IP/RPR network for the city of Milan is simulated. A similar scenario can be assumed for any major European city. We consider service classes for both elastic traffic (Web browsing, http-based services, and email services) and streaming traffic with stringent delay requirements (telephone services and video streaming), which are the most common IP-based applications. The simulated network is an RPR ring connecting 12 IP routers and three Internet servers (this network is shown in Fig. 2a), assuming that all of them use the wrapping protection mechanism.

The distance among nodes is set to 3 km, which results in a propagation time between nodes of 15 µs. We consider OC-48 (2.4 Gb/s) RPR node interfaces, with video and voice traffic sent as high-priority (class A) traffic, and data traffic (Web browsing, http, and email services) sent as low-priority (class C) traffic.

Finally, the network consists of a logical topology composed of three different segments, each including four routers logically attached to one server. Each segment represents a geographical zone of the metropolitan environment. Moreover, we assume traffic homogeneity in the three different segments and the same traffic matrices for all of these. As an example, Fig. 2 includes the traffic matrices for one of these segments, composed of server 1 and routers 2, 3, 6, and 11.

These traffic matrices were obtained from the estimation, carried out within the IST LION Project, of traffic flows in a realistic environment (the city of Milan) [5, 6]. The estimation took into account not only the characteristics of each service, but also the potential penetration (percentage of customers) for these kinds of services.

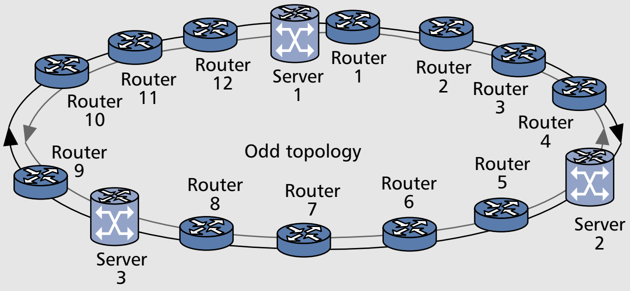

(a)

	Server 1	Router 2	Router 3	Router 6	Router 11
Server 1	-	4.78	71.8	148.4	253.8
Router 2	104.15	-	8.48	-	8.48
Router 3	78.8	-	-	16.66	15.55
Router 6	45.42	-	7.21	-	19.46
Router 11	46.74	-	1.14	9.12	-

(b)

	Server 1	Router 2	Router 3	Router 6	Router 11
Server 1	-	-	34.5	-	34.5
Router 2	-	-	-	-	-
Router 3	73.9	-	-	-	38.1
Router 6	-	-	-	-	-
Router 11	17.2	-	8.9	-	-

(c)

	Server 1	Router 2	Router 3	Router 6	Router 11
Server 1	-	(*) VTx0.25	(*) VTx0.29	(*) VTx0.25	(*) VTx0.21
Router 2	-	-	-	-	-
Router 3	15.75	-	-	-	-
Router 6	13.5	-	-	-	-
Router 11	-	-	-	-	-

(d)

(*) Concerning the video traffic (VT) generated by the servers, we considered four different cases:
VT = 0, VT = 0.33 Gb/s, VT = 0.43 Gb/s, and VT = 0.83 Gb/s

■ Figure 2. *a) RPR network topology; traffic matrix in megabits per second: b) data traffic; c) voice traffic; and d) video traffic.*

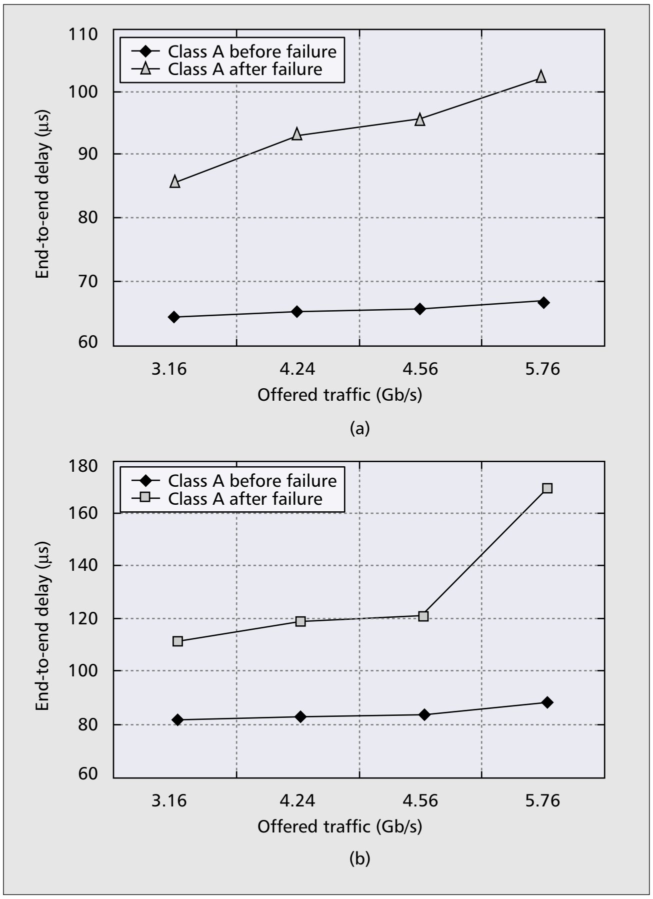

■ Figure 3. *Impact of link failure on end-to-end delay for a) high-priority; b) low-priority traffic.*

Concerning the traffic model, we use the ON-OFF model, with a burstiness (peak rate/average rate) of $b = 10$ and a mean burst length of $BL = 10$ packets for data traffic sources, and the Poisson model with a mean packet arrival intensity of λ packets/s for voice and video traffic sources. For data traffic, we consider the statistical distribution for the IP packet size given in [7], while for voice and video packets we used fixed packet sizes of 44 and 512 bytes, respectively.

On this scenario, we simulate two different case studies, a fiber cut between two routers and a router (not server) failure. The simulated operation time is 200 ms for the fiber cut and 300 ms for the router failure. In both cases the failure occurs at the instant $t = 70$ ms.

The aim of these simulations is to evaluate the impact of a failure on both the mean end-to-end delay and the network throughput. In particular, the *complete recovery time* is evaluated, which is the time required by the network to return to a steady state after a fiber cut or node failure. The complete recovery time comprises the *response time*, which includes detection of the failure, the generation of protection messages and node state transitions, and final ring wrap, the *topology discovery reconfiguration time*, and the *MAC protocol convergence time*.

Figure 3 depicts the mean end-to-end delay experienced by class A and C traffic before the failure (in this case the fiber cut) occurs and after the reconfiguration of the ring once the failure has occurred. The results show that the average end-to-end delay suffered a significant increase (about 50 percent). This is due to the fact that after the wrapping reconfiguration of the ring, the end-to-end path is longer for some traffic streams. It has to be noted that in the particular case of class A traffic, the end-to-end delay is below the typical requirements for voice and video services, even after ring failure recovery.

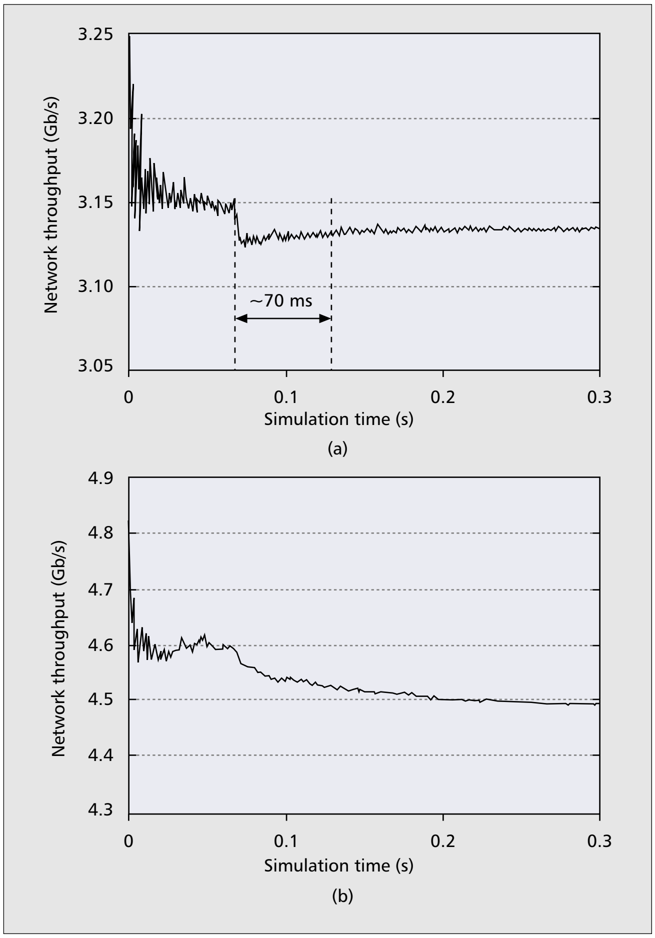

■ Figure 4. *Network throughput evolution after a node failure: a) no video traffic; b) average video traffic generated by the servers is 0.43 Gb/s.*

Figure 4 depicts the throughput as a function of the simulated operation time in a router failure (300 ms). Both plots of this figure show that after node failure the throughput suddenly decreases and subsequently, after *complete recovery*, converges toward a final throughput, which is lower than the throughput value before the failure. This is because a router, after it fails, no longer injects traffic into the ring, and the remaining nodes stop sending traffic toward that router once they know it has been excluded from the network.

Figure 4a, which is obtained for the case of no video service, shows that the network throughput evolves toward a steady state after stabilization of the MAC algorithm, and the RPR ring continues to work efficiently. We estimated that in this case the time needed to return to stability (complete recovery time) was 70 ms.

Figure 4b, obtained in higher load (including video traffic: VT/server = 0.43 Gb/s), shows that after node failure the throughput decreases, and it takes longer than 200 ms to reach a stable situation.

In failure the most important objective is maintaining network connectivity and minimization of packet losses. The results of the simulation experiments discussed above show that the RPR protection mechanisms have been optimized to do so. On one hand, traffic losses can only occur during the response time, which is comparable to that of SONET/SDH networks (a few milliseconds). On the other hand, complete recovery time depends on the ring size and actual traffic load, but if the traffic is well engineered the network reaches stability and will not saturate.

Potentially Hazardous Situations

The aim of this section is to describe a situation in which a given traffic assignment leads to significant degradation of network performance when a failure occurs. The example presented is valid for both steering and wrapping protection.

The consequence of the failure is that the routes traversed by frames switch from short to long ones. Additionally, the use of the fairness algorithm causes bandwidth to be shared between all active streams. This inevitably leads to potentially hazardous situation (e.g., a significant decrease in the bandwidth allocated to each class C stream).

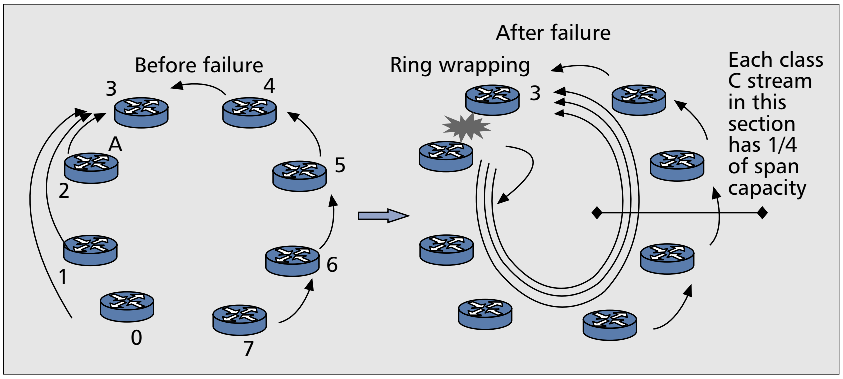
Before failure
3
4
A
5
2
6
1
0
7
After failure
Ring wrapping
3
Each class C stream in this section has 1/4 of span capacity

■ Figure 5. *An RPR ring with a worst-case traffic stream assignment.*

Consider the situation in Fig. 5, which is the result of detailed research. Nodes in one part of the ring (here, on the right) send class C traffic to their neighbors, while the remaining nodes send class C traffic to a given hub node (depicted here as node 3). This case seems to be important, since detailed analysis shows that for large rings the reduction in capacity after failure may be considerable (up to 94 percent) [8].

Let N_n be the number of nodes sending traffic to their neighbors and $N_{\text{to\_3}}$ the number of nodes sending traffic to node 3 (Fig. 5, on the left).

$$N_n + N_{\text{to\_3}} + 1 = N \qquad (1)$$

The total bandwidth (i.e., the maximum bandwidth available for all traffic streams in the case described) before failure is equal to

$$B = N_{\text{n}} + 1 \qquad (2)$$

The bandwidth unit used here is the full bandwidth of the RPR span (i.e., 2.4 Gb/s). Component 1 in the equation above is the result of all $N_{\text{to\_3}}$ streams in the left part of the ring that are sharing the bandwidth; therefore, each stream reaches a steady state of $1/N_{\text{to\_3}}$ units while it is sending traffic to node 3. Component N_n is the aggregated bandwidth of the streams associated with nodes that are sending traffic to their neighbors (on the right side of the ring). Obviously, the situation shown in Fig. 5 is a simplification, since both optical rings are in fact used.

After failure, the rings are wrapped, and the traffic streams previously directed to node 3 must now travel in the opposite direction, which thus interferes with traffic to the neighboring nodes. Due to the fairness feature for class C traffic, each span of the ring shares its full capacity evenly among ($N_{\text{to\_3}}$ + 1) streams (as shown in Fig. 5, where $N_{\text{to\_3}} = 3$), so the capacity of single streams equals $1/(N_{\text{to\_3}} + 1)$. The total bandwidth for RPR after failure is as follows:

$$B = \frac{1}{N_{to\_3} + 1} \cdot (N_n + N_{to\_3}) \qquad (3)$$

The total loss of traffic after failure (after being maximized against) is equal to

$$Loss = 1 - \frac{4(N-1)}{(N+1)^2} \qquad (4)$$

For $N = 63$ (in principle, N may be as high as 255), the loss of traffic is equal to 94 percent of the traffic sent before the RPR ring reconfiguration. This result is somewhat discouraging and leads to the conclusion that rings should be engineered to avoid such a traffic pattern (or similar ones); when network performance is an important issue, it may be necessary to verify network operation in each of the assumed failure scenarios.

Strengths and Weaknesses of RPR

RPR technology has attracted moderate interest in the last three years. It can be considered a *niche* technology. Important issues related to the use of RPR technology are discussed below, to point out its advantages and review its disadvantages.

The protection mechanisms implemented in RPR are fast: they aim to achieve recovery times of approximately 50 ms and to protect against any single failure in the ring. No bandwidth is dedicated for recovery purposes; therefore, in a failureless state resource utilization is high. However, in failure, the bandwidth available is substantially reduced. The reduction factor depends on the actual load and distribution of traffic.

If high-priority traffic is used in an RPR ring, the traffic must be shaped at ingress, and the service that uses this type of traffic must be carefully engineered. No mechanisms are provided to solve contention among high-priority traffic streams. If the high-priority traffic admitted exceeds the capacity of a given span, low-priority traffic is blocked. Thus, if problems are to be avoided, the amount of high-priority traffic injected into the ring must be controlled and limited by the higher layers, especially in the case of failure. We suggest that each failure scenario be investigated in turn to determine whether a given load is handled properly.

RPR would seem to be a wise choice for efficient and reliable transport of best effort traffic. It may be used to transport traffic with strict bandwidth and delay requirements, although in this case one would need to verify whether RPR would satisfy the necessary parameters for all conceivable traffic flow patterns. With regard to the use of different classes of traffic, RPR requires external measures to prevent congestion. These measures are not standardized or otherwise defined at present, so it is up to the user to provide them. However, it is possible that such measures will be defined as RPR technology matures and its use becomes widespread.

An important issue in modern telecommunications networks is interoperability among different layers. A new protocol should interwork smoothly with existing protocols. Interoperability with several physical layer techniques was explicitly considered during the standardization process of the IEEE 802.17 RPR. From the upper layer point of view RPR may be seen as a shared medium technology, and as such the problem was not widely studied.

State of the Standards and Deployment

Many metropolitan networks use a physical ring structure. It is a natural environment for the SONET/SDH networks that constitute the bulk of current metropolitan network infrastructure. SONET/SDH, however, was designed for point-to-point circuit-switched services (e.g., voice traffic), and its use for data traffic has several well-known disadvantages. Alternatively, Ethernet might offer a simpler and inexpensive solution for data traffic. However, because Ethernet is optimized for point-to-point or meshed topologies, its use of the available bandwidth is inefficient, and it does not take advantage of the ring topology in order to implement a fast protection mechanism [9].

RPR technology fills this gap by acting as a multiservice transport protocol based on packets rather than circuits. RPR systems are seen by many carriers as the inevitable successors to SONET/SDH add/drop multiplexer (ADM)-based rings.

RPR networks may provide performance-monitoring features similar to those of SDH and, at the same time, maintain Ethernet's advantages (e.g., low equipment cost, high bandwidth granularity and statistical multiplexing capability). Furthermore, the RPR MAC layer uses either packet over SONET (POS) (HDLC plus SDH framing) or GFP encapsulation, and may run on top of a physical SDH infrastructure or operate straight over the fiber (dark fiber or WDM).

For carriers, RPR promises to deliver all the necessary end-user services, such as time-division multiplexed (TDM) voice, virtual private network (VPN) data, and Internet access, at dramatically lower equipment, facility, and operating costs.

The key questions to be answered concern RPR's prospects in the evolution of transport networks and the expected penetration of RPR technology in metropolitan environments.

RPR seems to be a promising technology, since most of the major carriers have actively participated in the IEEE 802.17 standardization process and have shown much interest in the evolution of the standard. Furthermore, although RPR is still not standardized, at least three prestandard variations of RPR have already been put into live service in Asia, Canada, Europe, and the United States [10]. We believe the introduction of IEEE 802.17 RPR will take place in the medium term: the standard, which was forecast for March 2002, has been delayed and is expected no earlier than 2004. Thus, IEEE 802.17 RPR equipment will not be available before the beginning/end of 2004. As a consequence, the earliest deployment of IEEE 802.17 RPR networks will be in the timeframe of two or three years. Finally, regarding different geographical areas' readiness to implement RPR technology, Asia (mainly China) seems to be in first position, followed by the United States. In Europe, the prospects for RPR are not particularly promising. This conclusion is based on the current deployment of IEEE 802.17 RPR prestandard technology such as DPT-based products from Cisco Systems and OPTera Packet Edge Systems series 3000 from Nortel Networks. China is currently a good market for RPR products because there is not a great deal of SONET/SDH infrastructure installed, which thus opens the market to new and more efficient technologies (China Netcom already deployed Luminous' RPR-Based Metro Platform in several cities in 2002). Pre-RPR systems, such as OPTera-3000, are better positioned in the United States than in Europe, simply because OPTera 3000 is ready for SONET and not for SDH. The most important examples of pre-RPR deployments in Europe are DPT-based products. In Europe, Ethernet technology has a better chance than RPR.

Summary

RPR is a technology for networks of various sizes, from LAN to MAN or even WAN. Its ability to work over different kinds of underlying physical technologies makes it useful for both network operators, who may operate it over dark fibers or DWDM systems, and private users, who either own fiber infrastructure or lease SONET/SDH circuits.

Three features, resilience, simplicity, and efficiency, make this technology especially attractive. The RPR ring is able to maintain connectivity through equipment or fiber failure. To achieve this, neither dedicated bandwidth nor labor-intensive configuration is required. This, of course, comes at a cost, which is bandwidth reduction in a failure state. This must be taken into consideration whenever quality of service is important. Other interesting features are its ability to provide each node with a (weighted) fair share of bandwidth and to handle three traffic classes. The latter allows the implementation of a large number of services that require delay and bandwidth guarantees.

Acknowledgments

This work has been partially funded by the European Commission under the Information Society Technologies (IST) Program (LION Project IST-1999-11387) and the Spanish Ministry of Science and Technology under contract FEDER-TIC2002-04344-C02-02.

The authors would like to thank Prof. A. Jajszczyk for his valuable comments and suggestions during the preparation of this article. Finally, the authors wish to thank the anonymous reviewers for their precious suggestions.

References

[1] IEEE 802.17 Working Group, "Resilient Packet Ring Access Method and Physical Layer Specifications," Draft v. 3.0.

[2] D. Tsiang and G. Suwala, "The Cisco SRP MAC Layer Protocol," RFC 2892, Aug. 2000.

[3] ITU-T Draft Rec. X.msr, "Multiple Services Ring," Mar. 2002, temp. doc. 2053.

[4] J. Moyano *et al.*, "Performance Evaluation of the Spatial Reuse Protocol Fairness Algorithm (SRP-fa) Used in DPT Networks," *Proc. IEEE Int'l. Conf. Telecommun.*, Bucharest, Romania, June 2001, pp. 147–52.

[5] S. Spadaro *et al.*, "Assessment of Resilience Features for the DPT Rings," *Proc. Eurescom Summit*, Heidelberg, Germany, Oct. 2002, pp. 91–100.

[6] IST-1999-11387 LION Project, "Multilayer Resilient Network Planning and Evaluation: Preliminary Results," Del. 10, Jan. 2001.

[7] K. Thompson, G. J. Miller and R. Wilder, "Wide-area Internet Traffic Patterns and Characteristics," *IEEE Network*, vol. 11, no. 6, Nov./Dec. 1997, pp. 10–23.

[8] A. Jajszczyk, A. Szymański, and K. Wajda, "Issues of DPT Efficiency in Various Traffic Conditions," *Proc. 8th Polish Teletraffic Symp.*, Zakopane, Poland, Sept. 2001.

[9] Resilient Packet Ring Alliance, "An Introduction to Resilient Packet Ring Technology," white paper, Oct. 2001, http://www.rpralliance.org/articles/Whitepaper10.01.pdf.

[10] Resilient Packet Ring Alliance, "Positioning RPR in the Technology Universe," White paper, Dec. 2001, http://www.rpralliance.org/articles/Positioning_RPR.pdf.

Biographies

SALVATORE SPADARO (spadaro@tsc.upc.es) received an M.Sc. in electrical engineering from both the Universitat Politècnica de Catalunya (UPC) and Politecnico di Torino in 2000. He is an assistant professor at the Signal Theory and Communications Department of UPC and is also currently working toward a Ph.D. degree at the same university. He is a member of the Advanced Broadband Communications Center of UPC (http://www.ccaba.upc.es).

JOSEP SOLÉ-PARETA (pareta@ac.upc.es) was awarded his Master's degree in telecommunication engineering in 1984 and his Ph.D. in computer science in 1991, both from UPC. In 1984 he joined the Computer Architecture Department of UPC. Since 1992 he has been an associate professor with this department. He is co-founder and a member of the Advanced Broadband Communications Center of UPC (http://www.ccaba.upc.es). He is participating in NOBEL (IP project) and e-Photon/One (Network of Excellence) of the European VI Framework Program.

DAVIDE CAREGLIO (careglio@ac.upc.es) received an M.Sc. in electrical engineering from both UPC in 2000 and Politecnico di Torino in 2001. He is currently an assistant professor at the Department of Computer Architecture, UPC, and is also working toward a Ph.D. degree at the same university. He is a member of the Advanced Broadband Communications Center of UPC (http://www.ccaba.upc.es). He has recently been involved in the ACTS-SONATA, IST-LION, and IST-DAVID European projects.

KRZYSZTOF WAJDA (wajda@kt.agh.edu.pl) received his M.Sc. and Ph.D. in telecommunications from AGH University of Science and Technology, Kraków, Poland, in 1982 and 1990, respectively, and currently is an assistant professor with the Department of Telecommunications, AGH-UST. Since 1991 he has been involved in international projects: TEMPUS JEP N° 0971, COST 242, Copernicus ISMAN, ACTS BBL, IST LION, and many national projects.

ANDRZEJ SZYMAŃSKI (szymans@agh.edu.pl) received an M.Sc. degree in telecommunications in 1999 from AGH University of Science and Technology, Kraków and is currently working toward his Ph.D. degree in telecommunications. His field of study is network planning issues of reliable optical transport networks. Since 2000 he has been involved in the European IST-1999-11387 project LION, Layers Interworking in Optical Networks. Currently he is involved in the NOBEL project, Next-Generation Optical Network for Broadband in Europe, funded by the European Commission.