

6 Multi-layer Recovery Strategy in Resilient Packet Ring over Intelligent Optical Transport Networks

This Chapter is devoted to describe the suggested resilience interworking strategies to be applied in an IP over Resilient Packet Ring over intelligent optical transport networks. Firstly, we discuss the RPR technology which is an emerging packet-based transport technology recently standardized by the IEEE. Specifically, we concentrate on the resilience mechanisms provided by RPR highlighting their advantages and drawbacks. Secondly, the recovery at the optical layer is discussed. Then, the DHOT approach, which is based on the interworking between the RPR and the optical layers is presented and evaluated. Finally, as a further contribution of this Thesis, an algorithm based on the use of the automatic switching of optical connections capabilities of the ASON networks to improve the bandwidth utilization of the RPR rings in case of failures or to respond to the client traffic fluctuations is presented and discussed.

6.1 Resilient Packet Ring technology

Resilient Packet Ring (RPR) is a new packet-based transport technology for ring-based metropolitan area networks. It includes a new MAC layer technology as well as powerful automatic recovery mechanisms. Indeed, RPR systems are seen as the successors to SONET/SDH ADM-based rings for the efficient delivery of IP-based data traffic achieving both the optimization of the bandwidth utilization. RPR technology has been recently standardized by the Institute of Electrical and Electronics Engineers (IEEE) as IEEE 802.17 RPR [97].

Many legacy 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 networks, however, as we discussed in Chapter 2, was designed and optimized for

point-to-point circuit-switched services such as voice services. On the other hand, for metropolitan environments, Ethernet technology may offer a simpler and cost-effective solution for the transport of the 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 fast protection mechanisms.

RPR technology fills this gap by acting as a multi-service transport protocol based on packets rather than circuits. While Resilient Packet Ring paradigm may provide performance-monitoring features similar to those of SONET/SDH, it maintains the advantages of the Ethernet technology such as low equipment cost, high bandwidth granularity and statistical multiplexing capability. The IEEE 802.17 RPR standard defines a set of protocols for detecting and initializing the shared ring configuration, recovery from failures and regulating the fair access to the shared medium.

For Carriers, RPR promises to deliver all the necessary end-user services, such as TDM voice, Virtual Private Networks, data and Internet access, at dramatically lower equipment, facility and operating costs. It is a very promising transport technology, since most of the major carriers and vendors (among others Cisco Systems, Luminous and Nortel Networks) have actively participated in the IEEE 802.17 standardization process and have shown much interest in the evolution of the standard [98]. In fact, the unique features of RPR were sufficiently interesting to trigger many pre-standard installations by important players in the telecommunications market (e.g., Cisco Systems, Nortel Networks, Sprint, Luminous, Bell Canada, Worldcom and SUNET). The first major pre-IEEE 802.17 RPR standard deployments RPR technology-based networks introduced by Sprint in 1999 and Macedonia Telecom as well as China Telecom in 2001.

Ring topology based on RPR is also studied by ITU-T and a preliminary version of Recommendation on Multiple Services Ring (X.msr) is available [99]. Table 10 summarizes the most significant differences between X.msr and IEEE 802.17 RPR.

Feature	ITU-T (X.msr)	IEEE 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 (pre-planned bandwidth)	Fairness algorithm for unprovisioned traffic (Class C)
Multicast	Supported	Supported

Table 10: Significant differences between X.msr and IEEE 802.17 RPR

RPR systems are seen by most carriers as the inevitable successors to SONET/SDH Add Drop Multiplexer/Digital Cross Connects-based rings.

Thus, the introduction of RPR-based metropolitan networks is gaining importance and it represents a very promising networking solution to transport data traffic in a short/medium term. In fact, the IEEE 802.17 RPR standard has been approved in June 2004 and thus, network equipments standard-compliant will not be available before the beginning of 2005. As a consequence, the earliest deployment of IEEE 802.17 RPR networks will be in the timeframe of two or three years. As concerns different geographical areas' readiness to implement RPR technology, we believe that Asia (mainly China) seems to be in first position, followed by the United States of America (USA). In Europe, the prospects for RPR seem not to be particularly promising. This conclusion is based on the current deployment of IEEE 802.17 RPR pre-standard 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, and 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 USA than in Europe, simply because OPTera 3000 is ready for SONET and not for SDH. Indeed, it is considered able to provide increased revenues for carries to transport data traffic in metropolitan environments.

Summarizing, in comparison with the currently enabled technology in the metropolitan environment, which RPR pursues to substitute (i.e., SONET/SDH rings), the RPR technology exhibits the following advantages: 1) Overcome of the limitations that TDM/circuit-based architectures impose on data communications allowing direct connectivity without circuit provisioning; 2) Increase of the bandwidth efficiency by implementing the spatial reuse of bandwidth and the statistical multiplexing of packets; 3) Enabling service integration by supporting various traffic priorities and 4) Reduction costs and complexity by eliminating intermediate layers between the IP layer and the optical layer.

6.1.1 Fundamentals of RPR technology

RPR-based networks enable efficient transfer of data traffic as well as fast protection mechanisms. It is a standard which consists of a superset of features derived from various

proprietary solutions such as Cisco Systems' DPT [100] and Nortel Networks' Optera Packet Edge System [108].

Network Operators claim that the functionalities of RPR and its implementation in real commercial environments present many advantages, namely:

- 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),
- OAM and advanced traffic and bandwidth management,
- Support for unicast, multicast and broadcast data traffic.

It has been designed to operate over a variety of physical layers, including SONET/SDH, Gigabit Ethernet (IEEE 802.3ab), DWDM and dark fibre, and is expected to work over higher-speed physical layers. The minimum supported data rate is 155 Mb/s.

RPR networks are based on two symmetric counter-rotating rings (external and internal ring) that carry data and control information (Figure 47). The nodes/stations may send data on either of the two ringlets. In most cases, the shortest path to the destination is used. Therefore, the nodes use a topology discovery protocol to obtain a topology map of the ring, which is then used for the shortest path computation.

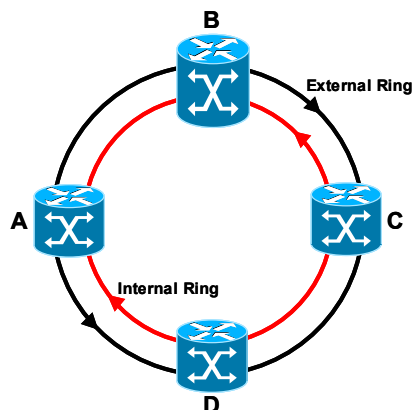


Figure 47:4-nodes Resilient Packet Ring network

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

The IEEE 802.17 RPR standard supports three types of services, namely Class A (high priority), Class B (medium priority) and Class C (low priority) [97]. The 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 on Figure 48) 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.

The Class B service is dedicated to near real-time applications that are less delay-sensitive but that still require some bandwidth guarantees. It provides guaranteed information transfer at the Committed Information 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 bandwidth for its Class B traffic.

The 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 divides the remaining bandwidth in proportion to administratively configured node weights.

The allocated rates for Class A and Class B services and node weights for Class C are configured in each station by a provisioning mechanism.

Each node has two MAC datapath entities, one for each ringlet (in the IEEE 802.17 RPR standard, the external ring is called Ringlet 0 while the internal ring is called Ringlet 1). Figure 48 presents an example of a three-node IEEE 802.17 RPR ring and a more detailed view of the MAC datapath entity. Specifically, it shows that a frame received from the 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 (in case of multicast traffic), 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 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.

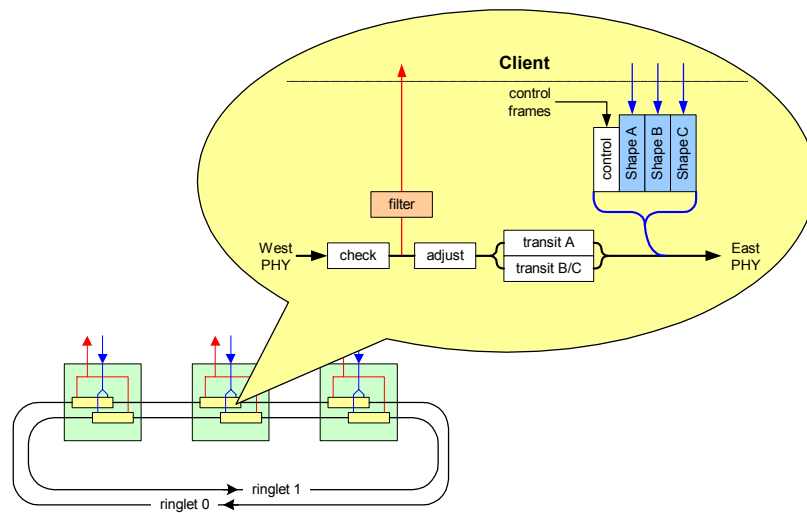


Figure 48: A three-node IEEE 802.17 RPR ring with a simplified structure of the MAC datapath entity

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 to avoid loss of frames in transit [97].

We carried out a simulation case study to evaluate the performance of the RPR MAC protocol. Specifically, a RPR ring composed by 5 nodes/stations was simulated. The traffic inserted by each node is uniformly distributed among the rest of the nodes. Specifically, 20% of the generated traffic by each node corresponds to Class A (i.e. high quality video traffic), other 20% to Class B (high priority IP traffic) and the rest to Class C (low priority IP traffic). Figure 49 depicts the throughput of the ring when increasing the offered load (ρ), being the offered load the ratio between the offered traffic and the maximum network throughput ($Throughput_{Max}$). Since the network is assumed error-free and no packets are lost, network throughput and offered traffic are equal until saturation. This occurs when the utilization factor of each link between nodes is equal to 1. It can be calculated that the $Throughput_{Max}$ is equal to $8 \cdot R_b / (n+1)$, where R_b is the RPR bit rate interface and n is the number

of nodes composing the ring. In Fig. 29, being $n = 5$ and $R_b = 2.5$ Gbps, the $Throughput_{Max}$ is equal to 16.67 Gbps.

In extreme load traffic conditions (i.e., the ring is saturated), the network is not able to support all the offered traffic. RPR reacts and thus the throughput of Class C decreases, while the throughput of the Class A and B still continue to increase.

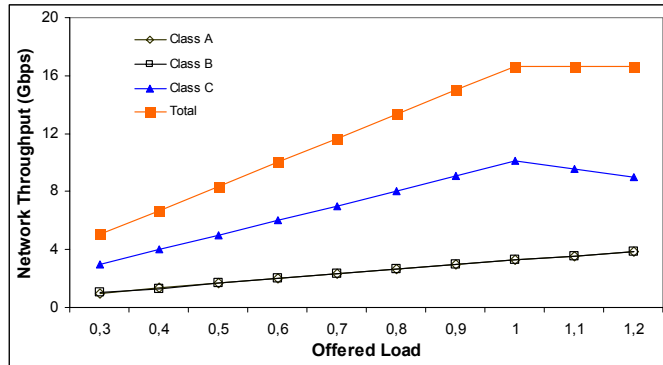


Figure 49: IEEE 802.17 RPR MAC: performance evaluation

If high priority traffic is used in RPR rings, the traffic must be shaped at ingress, and the service that uses this type of traffic must be carefully engineered. In fact, 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, 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. Specifically, each failure scenario has to be investigated in turn to determine whether a given load is handled properly.

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 a detailed awareness of congestion points in the whole ring but increases ring utilization and spatial reuse.

6.1.2 RPR resilience mechanisms

At the same time, RPR standard offers powerful protection methods, namely the *ring wrapping* and the *packet steering*. They are based on the ring wrap at the nodes surrounding the failure or on the packet steering by causing the source node to redirect packets [97]. Both of them have been designed to minimize the traffic losses in case of failures and aim to achieve recovery times of about 50 ms and no spare resources (capacity) are required to be pre-allocated [97]. RPR allows the

full ring bandwidth to be utilized under normal conditions and protects traffic in case of failure obviating the need for SDH/SONET-based protection.

There are few steps to recover from failure by RPR layer, which include the indication of a failure (or significant signal degradation) and the final wrapping or steering of the ring.

The RPR complete recovery time is the time required by the network to return to a steady state after a failure has occurred in the ring. Focusing specifically on the wrapping, the complete recovery time comprises the *response time*, the *topology discovery reconfiguration time*, and the *MAC protocol convergence time*. The *response time* includes the detection of the failure, the generation of protection messages and the node state transitions, and finally the ring wrap. Wrapping mechanism works as follows: if failure is detected (either equipment or link failure), packets directed towards failure direction are wrapped back in opposite direction. It is made possible because of internal node structure with dual homing (connection to external and internal ring). Figure 50 shows an example of how ring wrapping works. The example shows RPR ring composed of four nodes in which link failure (e.g. fibre cut) on the fibre from node A to B is considered. The interchanged control messages needed to run the ring wrapping are figured as $\{Request\ type, Source\ Address, Wrap\ Status, Path\ Indicator\}$. In absence of failure (Ring in idle), each node periodically sends control message (IDLE). When the failure occurs in the external ring between A and B, node B detects a signal fail (SF) on the external ring, (e. g. not receiving the periodic keep-alive message sent periodically from node A). Thus, it changes to wrapped state performing a wrap and transmits towards A on the internal ring (short path) the message $\{SF, B, W, S\}$ and on the external ring (long path) the message $\{SF, B, W, L\}$.




Figure 50: RPR wrapping protection

After receiving a protection request from node B on the short path node A changes to wrapped state. Then, it transmits towards B on the short path $\{0, A, W, S\}$ and on the long path $\{SF, B, W,$

L}. When the nodes C and D receive long path control packets (switch requests), they change from idle to pass-through mode (in each direction).

After the wrapping of the ring, the update of the topology nodes map is needed because RPR uses the shortest path for the communications between node pairs. To do this, each node runs a topology discovery algorithm [97].

Packet steering is based on the ability to choose the ringlet on which the data is sent. If the preferred path is unavailable due to failure, the other path will be used.

The implementation of the wrapping protection in the nodes is optional. Both protection modes may be mixed in a wrap-then-steer mode where the wrapping protection is activated first to avoid the loss of frames in transit; then 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 the wrapping protection.

At RPR layer, the failure detection may be carried out in two ways. The first is based on messages received from the physical layer, for example, Loss of Signal (LoS) from the SONET/SDH or the optical layer, and the second on periodical continuity checks within the IEEE 802.17 RPR layer.

6.1.3 Topology Discovery algorithm

The Topology Discover (TD) protocol is used for the network reconfiguration after the ring wrapping/packet steering. In normal RPR ring operation, both rings are utilized to carry traffic. After the wrapping of the ring, the available bandwidth is reduced, and low priority data traffic is reconverged fairly to the lowered bandwidth, which is accomplished dynamically by the fairness algorithm. In order to optimise the bandwidth utilization it is necessary to run the TD protocol because RPR technology supports the basic version of traffic forwarding based on number of spans, which means that the shortest path is chosen towards destination RPR node (Figure 51). Each RPR node performs this action by sending out special discovery packets on one or both rings. The originator of a topology discovery packet sets the egress ring identifier (internal or external ring) adding its own MAC address and length field. Such packet is sent hop-by-hop around the ring (however in nature it is a point-to-point packet). Each traversed node appends its MAC address, updates length field and forwards packet towards destination. After reaching again the originator of

this packet, topology discovery packet has all nodes MAC addresses in proper order and with relevant length field.




Figure 51: Topology map for node A before failure

When sending topology discovery packet over wrapped ring, the wrapped node indicates this situation and wraps the packet (i.e. sends it further along the reconfigured ring). On the way towards originator after passing wrapped node, MAC addresses are not added since it is travelling the same route in opposite order of nodes. Then, the topology map at each node is updated (Figure 52).




Figure 52: Topology map for Node A after the running of the TD algorithm

The topology map of the ring is changed after receiving two identical TD packets. This is done to avoid changes of topology in transient conditions. The delay introduced by the TD protocol is the time of traversing the whole ring (already being wrapped) plus necessary time to service packet inside the subsequent nodes.

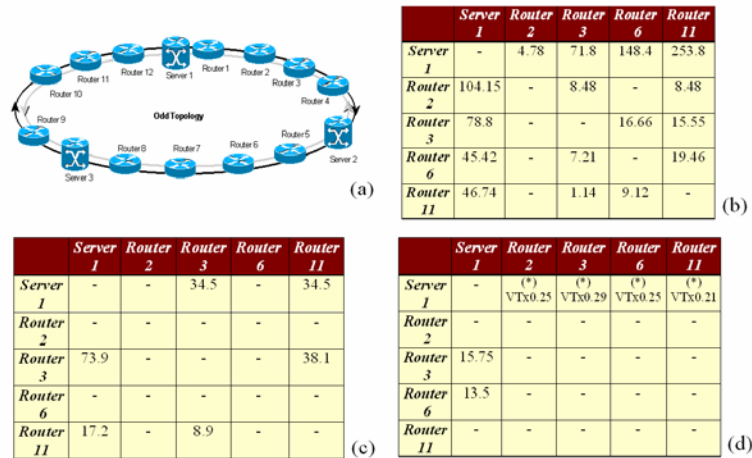
6.1.4 RPR resilience mechanism: Performance evaluation

We have carried out simulation case studies to assess the resilience features of RPR rings [102]. Specifically, a metropolitan IP/RPR network for the city of Milan was simulated using the OPNET

simulation tool [103]. 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 e-mail services) and streaming traffic with stringent delay requirements (telephone services and video streaming), which are the most common IP-based applications. The simulated network was an RPR ring connecting twelve IP routers, and three Internet servers, and we assumed that all of them use the wrapping protection mechanism.

The distance among nodes was set to 3 km, which results in the propagation time between nodes of 15 μ sec. We considered 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 e-mail services) sent as low priority (Class C) traffic.

The network consisted of a logical topology composed of three different segments, each one including four routers logically attached to one server. Each segment represents a geographical zone of the metropolitan environment. Moreover, we assumed traffic homogeneity in the three different segments and the same traffic matrices for all of these. As an example, Figure 53 includes the traffic matrices for one of these segments, namely the one composed of Server 1 (S1) and Routers 2, 3, 6 and 11.



(*) 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 53: (a) RPR network topology; Traffic matrix in Mb/s: (b) data traffic, (c) voice traffic, and (d) video traffic

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) [104], and we also used in

[105] and [106]. 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.

As concerns the traffic model, we used 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 per second for voice and video traffic sources. For data traffic, we considered the statistical distribution for the IP packet size given in [62], while for voice and video packets we used fixed packet sizes of 44 and 512 bytes, respectively. On this scenario, we simulated two different cases study, namely a fibre cut between two routers and a router (not server) failure. The simulated operation time was 200 ms for the case of the fibre cut and 300 ms for the case of the router failure. In both cases, it was assumed that the failure occurred at the instant $t = 70$ ms.

The aim of these simulations was to evaluate the impact of a failure, both on the mean end-to-end delay and on the network throughput.

Figure 54 depicts the mean end-to-end delay experienced by the Class A and Class C traffic before the failure (in this case the fibre 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 quite significant increase (about 50%). 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 the ring failure recovery.

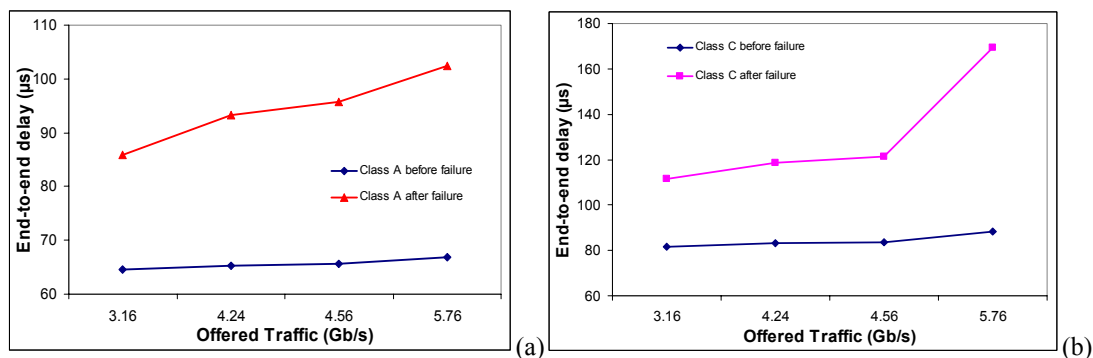


Figure 54: Impact of link failure on end-to-end delay for (a) high priority and (b) low priority traffic

Figure 55 depicts the throughput as a function of the simulated operation time in the case of the router failure (300 ms). Both plots of this Figure show that after node failure the throughput suddenly decreases and subsequently, after the complete recovery, converges towards 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 towards that router once they know that it has been excluded from the network. Figure 55 (a), which was obtained for the case of no video service, shows that the network throughput evolves towards a steady state after the stabilization of the MAC protocol, and the RPR ring continues to work efficiently. We estimated that, in this case, the time needed to return to stability (to the same network throughput value) is 70 ms. Figure 55 (b), which is obtained in case of higher load (including video traffic: VT per server = 0.43 Gbit/s), shows that, after node failure, the throughput decreases, and it takes longer time to reach a stable situation.




Figure 55: Network throughput evolution after a node failure: (a) no video traffic, (b) average video traffic generated by the servers is 0.43 Gbit/s

In the case of 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 (few ms), which is comparable to that of SONET/SDH networks. On the other hand, the complete recovery time (including the time required for the reconvergence of the RPR MAC protocol) depends on the ring size and on the actual traffic load, but if the traffic is well engineered the network, after the ring reconfiguration, reaches the stability and will not saturate.

Although fast and efficient, the RPR recovery mechanisms imply that the available bandwidth is reduced. The reduction factor strongly depends on the actual load and distribution of traffic. The next Section discusses this problem in detail.

6.1.5 Potential hazardous situation¹

The aim of this Section is to describe a situation in which a given traffic assignment leads to a significant degradation of network performance when failure occur. The presented example is valid both for 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 situations, e.g., a significant decrease in the bandwidth allocated to each Class C stream.

As an example, consider the situation in Figure 56, in which a RPR ring composed by N nodes is taken into consideration. Nodes in one part of the ring (here, on the right) send Class C traffic to their neighbours while the remaining nodes send class C traffic to a given hub node (depicted here as Node 3).




Figure 56: RPR ring with a worst-case traffic stream assignment

Let N_n be the number of nodes sending traffic to their neighbours and N_{to_3} the number of nodes sending traffic to node 3 (Figure 56, on the left).

$$N_n + N_{to_3} + 1 = N \quad (1)$$

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

$$B = N_n + 1 \quad (2)$$

The bandwidth unit used here is the full bandwidth of the RPR span, (e.g. 2.4 Gbit/s). Component 1 in the equation above is the result of all N_{to_3} streams in the left part of the ring sharing the bandwidth and therefore each stream reaches a steady state of $1/N_{to_3}$ units when

¹ This Section includes the work carried out by the University of Science and Technology (AGH), Krakow, Poland, as a contribution to reference [101]

sending traffic to node 3. Component N_n is the aggregated bandwidth of the streams associated with nodes sending traffic to their neighbours (on the right side of the ring). Obviously, the situation shown in Figure 56 is a simplification, since both optical rings are, in fact, used.

After the failure, the rings are wrapped and the traffic streams previously directed to node 3 must now travel in the opposite direction, thus interfering with traffic to neighbouring nodes. Due to the fairness feature for Class C traffic, each span of the ring shares its full capacity evenly among $(N_{to_3}+1)$ streams (as shown in Figure 56, where N_{to_3} equals 3), so the capacity of single streams equals $1/(N_{to_3}+1)$. The total bandwidth for RPR after failure is:

$$B = \frac{1}{N_{to_3} + 1} \cdot (N_n + N_{to_3}) \quad (3)$$

The total loss of traffic after failure is equal to (after maximizing it against N_n):

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

For $N = 63$ (in principle, N may be as high as 255), the loss of traffic is equal to 94% of the traffic sent before the RPR ring reconfiguration. This result is somewhat discouraging and leads to the following conclusions, namely: 1) RPR rings should be engineered to avoid such a traffic pattern (or similar) and 2) When network performance is an important issue it may be necessary to verify network operation in each of the assumed failure scenarios.

Apart from the theoretical case study carried out by AGH University, we carried out a simulation case study to evaluate the bandwidth reduction due to the failure. Specifically, we simulated a 5-nodes RPR network. The offered traffic (ρ) was set to 0.6. The simulated operation time was 120 ms and the failure (e.g. fibre cut connecting two underlying OXCs) occurs at the instant $t = 50$ ms. The wrapping method was used.




Figure 57: Effect of the RPR network reconfiguration: ring saturation

As it can be seen in Figure 57, once the RPR layer detects the failure and after the consequent reconfiguration of the ring, the available bandwidth is substantially reduced (in this case halved), so then the network is not able to carry the entire offered load (the RPR ring is saturated).

6.1.6 Summary of strengths and weakness of RPR technology

RPR technology has attracted great interest during the last 3 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 discuss 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 and, therefore, in a failureless state the resource utilization is high. However, in the case of 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 congestion problems have 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.

The 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 the 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.

Finally, 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 during the definition of the standard.

6.2 Resilience interworking strategy in RPR over intelligent optical networks

The increase of the number of wavelengths that can be multiplexed onto the same fibre (up to 160) each one carrying 2.5 or 10 Gb/s client signals implies that outages of the network infrastructure (e.g. fibre cut) can have serious consequences (economical as well as social) [26].

As before mentioned, in current network infrastructure recovery is carried out at SONET/SDH layer. Protection at the SDH layer, based for example on the Automatic Protection Switching (APS) protocol [107], although very robust (allowing network recovery within 50 ms) is not efficient from the network resources optimization point of view. In fact, it needs to pre-allocate spare network resources to be used for protection purposes and the bandwidth reserved for the backup paths is not used to carry traffic, increasing in such a way the required CAPEX.

On the other hand, the achieved advances in optical components as well as the introduction of intelligence (i.e. Control Plane) to the optical layer lead to the definition of recovery mechanisms directly in the server optical layer. The introduction of resilience mechanisms in the optical layer is very useful because the optical layer provides better management for certain kind of failures. As an example, let us suppose that a single optical fibre carries multiple wavelengths which correspond to some SONET/SDH streams. If recovery at SONET/SDH layer is used, a fibre cut therefore results in that all the streams are restored independently by this layer. As a consequence, the network management system is flooded with a large number of alarms generated by each of these independent entities. Contrarily, if the fibre cut is recovered at optical layer this operational inefficiency can be avoided.

Hence, failures such as fibres cut, or optical equipment damages can be more efficiently handled.

Both protection (1+1 and/or 1:1/1:N) and restoration can be used in the optical layer. The fault management is based on the detection of the failure, the notification of the detection of the failure, the failure localization, and finally on the recovery procedure. The latter can be based for example on dedicated path protection (Figure 58) and/or restoration.

To implement such fault management the GMPLS-based control plane can be used. Specifically, the optical nodes are equipped by OXC switch controllers connected through

signalling networks and generating the optical signalling. Different ways can be used to implement signalling between the controllers of the optical equipments.




Figure 58: Recovery at the optical layer, 1:1 dedicate path protection between node A and C

As described in Chapter 2, the way to implement signalling, the in-fibre in-band signalling, the in-fibre out-of-band signalling and the out-of-fibre out-of-band signalling methods can be used.

Specifically, Figure 59 shows an example in which the in-fibre out-of-band signalling is considered.

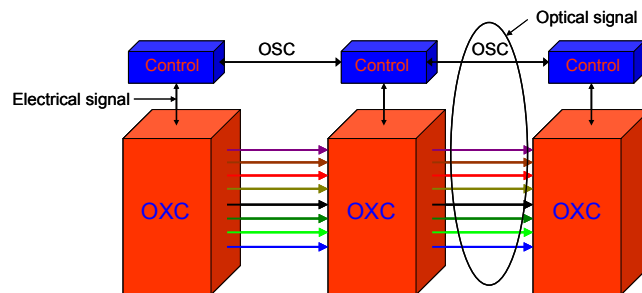


Figure 59: Control structure in the optical layer, in-fibre out-of-band signalling

To monitor the connectivity of the control channels, the Link Management Protocol (LMP) defined by the Internet Engineering Task Force (IETF) is used [108]. In such a context, the supervision of the control channels is based on a fast keep-alive procedure based on the interchange of *HELLO* messages between the optical Connection Controller (CC). Each node sends the *HELLO* messages periodically (each *HelloInterval*) to its neighbours. Nevertheless, the interchange of these messages could not provide data link failure detection since the failure of the control channels can be due to the failure of the transmission equipment of the signalling messages (i.e., control channel lasers). Thus, the data link failure detection is done combining the signalling based on the *HELLO* messages with the monitoring of the optical signals at the OXC interfaces/ports, in order to detect the Loss of Light (LOL). In such a scheme, when the optical layer detects the fault, first it signals to

the client layer that the failure has been detected and then it launches the predefined procedure for the management of the fault.

It has to be underlined that the signalling network can be used not just for the connectivity of the control channels but also for the management of the routing and signalling instances.

However, it worth noting that carrying recovery in the optical layer presents some drawbacks such as: 1) The optical layer is not aware of failures that occur at higher layers and 2) Link budget constraints limits the recovery capability because the length of the protection route or the number of nodes the protected traffic passes through may be physically constrained [3].

If we consider that RPR runs over intelligent optical transport networks, using a single-layer strategy, the only layer responsible for taking the recovery actions in case of a failure is the optical layer. This is due to the fact that it is better to recover from the failure at the optical layer since the RPR protection mechanisms imply the reduction of the available bandwidth.

The obvious advantage of this single-layer approach is that it does not require any interworking feature between RPR and optical transport layer. However, not all kinds of failures will be handled efficiently with this strategy, since the optical layer is not able to detect the failures occurred at the higher layers, such as the RPR line card failures or RPR site failures.

Hence, this Ph.D. Thesis proposes a multi-layer recovery strategy in order to efficiently coordinate the different mechanisms of each layer. It is to be used in an IP over RPR over OTN/ASON metropolitan network scenario.

When combining RPR over OTN/ASON, we are using technologies with similar reaction times but different features. While RPR recovers from failure by ring wrapping around failed span or by packet steering, the optical layer relies, for example, on dedicated resources to recover from failures (i.e. 1+1 and/or 1:1 dedicated protection).

If the *uncoordinated approach* is used to coordinate the resilience mechanisms, both RPR and optical layer resilience mechanisms act independently of each other. Nevertheless, due to the fact that resilience mechanisms detect the failure in similar time, it is very likely, that both layers will try to restore connectivity at the same time (Figure 60). This can lead to significant performance degradation for the layer above RPR (i.e. IP). If both RPR and OTN/ASON start recovery actions, independently of the procedure in the optical layer, once detected the failure, RPR will wrap its

ring, thus reducing available bandwidth depending upon the traffic pattern usage before and after the wrap. As a result, a failure at the optical level that could be efficiently managed by the optical layer moreover implies the reduction of the bandwidth at the client layer (RPR layer).

The problem of harmful interaction between RPR and underlying layer is only mentioned in [97]. There, RPR over SONET/SDH interworking scenario is considered and not specific conclusions are given. The document suggests avoiding such a case where RPR and SDH protection are used simultaneously and recommends either using a single layer protection (e.g. SDH APS) or implementing a hold-off timer (hereafter single hold-off timer approach).

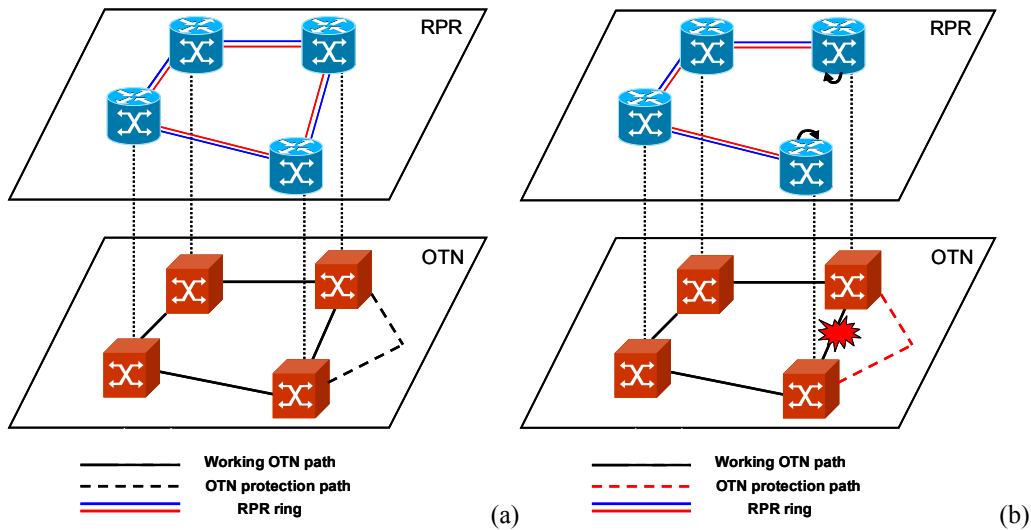


Figure 60: RPR/OTN scenario: a) basic arrangement, b) uncoordinated approach

Regarding the single hold-off timer, using this approach, the recovery action in OTN/ASON is launched immediately after a failure is detected, while the recovery in RPR is delayed for some time, needed to complete recovery tasks. If the optical layer manages to re-establish connectivity, there is nothing to do for the RPR protection and no action is thus taken at this level. If, on the other hand, the optical layer is unable to recover from a failure, after the hold-off timer has expired RPR will trigger its own protection, and will recover from the failure. The main advantage of this approach is its simplicity. Moreover, in case the failure is resolved in the optical layer, there is no need for RPR re-convergence (which slows down recovery process) and no bandwidth reduction occurs as it was in case of the uncoordinated recovery.

Nevertheless, the single hold-off timer (hereafter SHOT) approach presents a very important drawback, namely its length. If it is too short, the protection in RPR will be needlessly triggered before OTN/ASON has finished the recovery process. If it is too long and OTN/ASON cannot cope

with the failure, RPR will wait with its protection without any reason. Summarising, if the failure occurs at the optical layer, it works properly but its efficiency decreases very much when the failure occurs at the higher layers.

6.2.1 Double Hold-Off timer approach

This Thesis proposes a novel multi-layer resilience strategy, based on the interworking between RPR and the optical layer. It consists on implementing the double hold-off timer (hereafter DHOT) approach. RPR can detect a failure in different ways, depending on the used sources of information about failures: some of them are independent from other layers and one is the information signalled from the underlying layer (i.e. optical layer). RPR, detecting a failure (through signal fail (SF) signalling [97]), is able to distinguish between two cases: 1) When the optical layer has also detected the failure (it has occurred at optical layer) and 2) When the optical layer has not detected the failure (and RPR is the only layer that is able to do a successful recovery). In the latter case, the failure has occurred in the upper layers.

The suggested DHOT approach is based on dividing the entire single hold-off timer into two parts, namely the H1 (short) part and the H2 (long) one. The first part (H1) is activated after RPR detects the failure. It serves to give to the optical layer some time to detect the failure, to notify and signal it to the RPR layer. It has to be underlined that the detection failure at the optical level is strongly influenced by the optical components themselves and their management. Anyway, according to [109], it takes very few ms.

After the expiration of H1, if the optical layer has not detected the failure, RPR triggers its protection immediately. On the other hand, if the failure is signalled to RPR layer by the optical layer, the RPR layer waits during the H2 timer to give time for recovery in the optical layer. The DHOT approach is described in detail in the flow-chart of Figure 61.

The recovery procedure in the optical layer encompasses both fault localization and the recovery mechanism (i.e. dedicated path protection or restoration). Even in the case the optical layer detects the failure, it can be unable to solve it, for example due to the unavailability of resources in case of using restoration. Therefore, if after the expiration of H2 the failure is not recovered, then RPR protection mechanism is launched.

The required functionalities for the DHOT implementation have already been incorporated both in optical transport layer and in RPR. As stated above and according to [17] OTN is required to

signal to its client layer both signal degradation and signal failure while the RPR standard is able to accept such signals from the underlying layer [97].

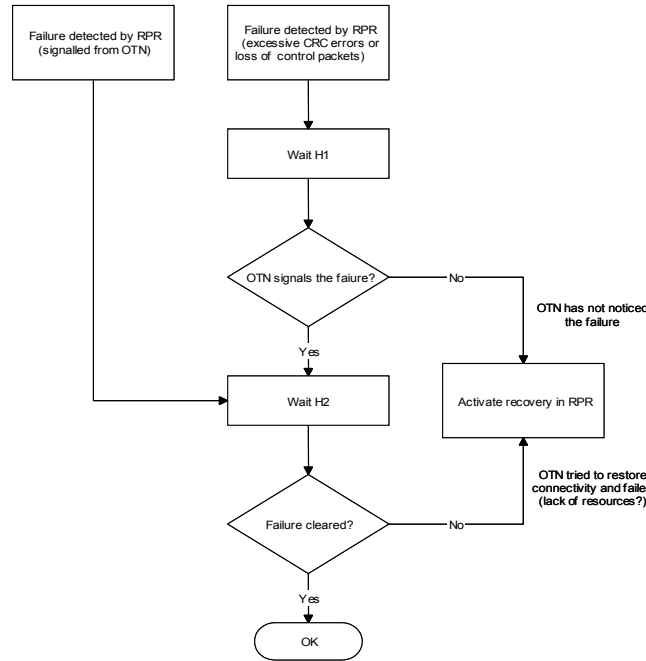


Figure 61: Coordinated approach: double hold-off timer (DHOT)

Figure 62 depicts the failure management both in the case the failure occurs at the RPR layer and in the case it occurs at the optical layer.

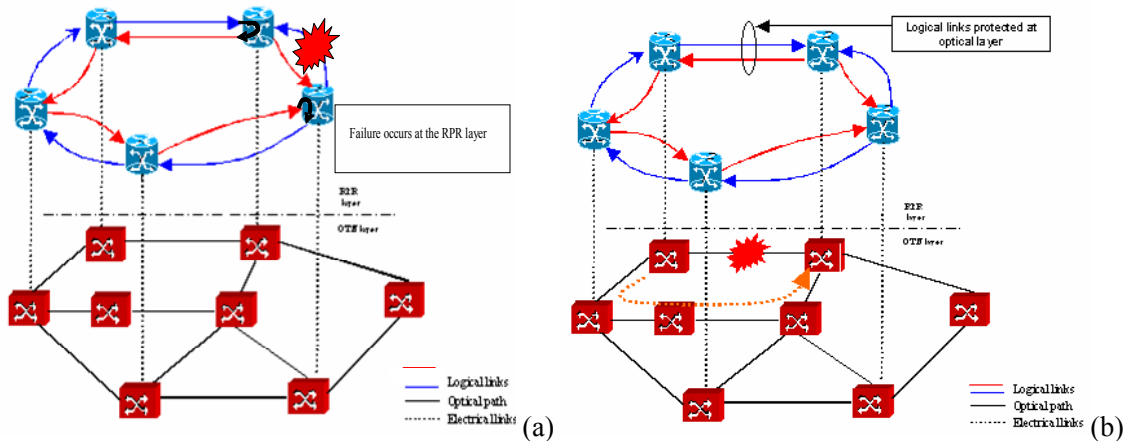


Figure 62: Failure management, a) at RPR layer and b) at the optical layer

The main advantage of the DHOT approach, when compared to the SHOT one is that the recovery time is much shorter when failure root is above the optical layer, allowing in this case minimizing the traffic lost due to the failure. Moreover, being based on the interworking between

the two layers, it is able to react to the failure more promptly and independently from the failure scenario optimizing, at the same time, the utilization of the network resources.

6.2.2 DHOT approach: Performance Evaluation

We carried out various simulation case studies in order to compare the SHOT and the DHOT approaches for different failure scenarios. The simulated scenario consists of 5 IP routers equipped with RPR cards, logically connected through a meshed optical transport networks composed by optical cross-connects (OXC). The optical nodes are connected through bi-directional optical paths (i.e. two physically disjoint optical fibres) and the 1:1 optical path protection was implemented. For the fault management in the optical layer, we implemented the GMPLS-based Link Management Protocol. In our simulation model, the fault detection is carried out through the implementation of the *HELLO* messages signalling between the optical nodes controllers combined with the Loss of Light (LoL) alarms from the OXCs. Specifically, an in-fibre out-of-band signalling approach has been implemented. To avoid to get the ring saturated after the RPR recovery process, the offered load (ρ) was set to 0.45. Class A represented the 20% of the offered traffic, the same for the Class B traffic and the rest represented Class C traffic. We also assumed that the traffic inserted in the ring by each node was uniformly distributed among the rest of stations/nodes. The simulated operation time was set to 120 ms and the bottom-up coordination approach was used. The failures occur at the instant $t = 50$ ms.

On one hand, two case studies were carried out. The first one concerned the case in which the failure at the optical level (e.g., cut of the fibre connecting two OXCs breaking the logical connection between two IP/RPR routers) and the second one concerned the case in which the failure occurs at RPR layer (e.g. failure of RPR card of one of the routers composing the ring). In both case studies, the H1 timer was set to 10 ms while the H2 timer was set to 30 ms.

Focusing in the first case study, Figure 63 shows the network throughput versus the simulated operation time. According to the defined DHOT approach, the RPR layer once detected the failure, instead of launching immediately the recovery action waits for H1 in order to leave to the optical layer the objective to recovery from the failure. In this case, since the failure root is in the optical layer, this signals the failure detection to the RPR layer and then it launches the optical path protection. Once the optical level has recovered the network from the failure (we implemented path protection recovery), the network throughput comes back again to the same value before the failure occurred. When the optical layer is able to handle the failure, the behaviour of the network

throughput is the same using both the DHOT and the SHOT approaches (In Figure 63 only the case of DHOT is plotted). We obtained that the time required returning to the steady condition, which includes the fast failure detection and the switch to the optical protection path, is about 12 ms.




Figure 63: Recovery from failure at optical level

We insist that the detection failure at the optical level is strongly influenced by the optical components themselves and by the way they are managed. Anyway, as stated before, it takes very few ms.

The second case study deals with the comparison between the SHOT and the DHOT when the failure occurs at the RPR layer or upper layers. The comparison of the network throughput in case of using the SHOT and in case of using the DHOT is depicted in Figure 64. The SHOT foresees that the RPR layer waits for the entire hold-off timer (i.e. H1+H2). Once the timer has expired and the failure has not been recovered by the optical layer, then RPR starts to recovery from failure. By using the DHOT approach, the RPR layer just waits for the first short timer (i.e. H1). If H1 expires and the optical layer has not signalled the failure detection, then the RPR starts immediately the recovery action (in this case, we implemented the wrapping mechanism).




Figure 64: SHOT vs. DHOT, failure at RPR level

In such a case, after H1, the network throughput comes back again to the value before the failure. We can estimate that the time required to the network throughput to come back to the same value before the failure in case of SHOT is about 45 ms while in case of using DHOT in the same conditions is about 20 ms.

We also carried a third simulation case study to compare the SHOT and the DHOT approaches, using the relative traffic losses. In this case, we assumed that the optical layer is unable to detect the failure occurrence and we considered various values for the H1 and H2 timers. We calculated the ratio (R) between the packets lost obtained with the SHOT and the DHOT approaches, that is $R = (\text{Packets Lost})_{\text{DHOT}} / (\text{Packets Lost})_{\text{SHOT}}$.

Table 11 and Table 12 depict the gain in terms of percentage of reduction of the traffic losses (i.e., $100 \cdot (1 - R)$) arising from the implementation of the DHOT approach with respect to the SHOT. For such comparison, both RPR protection mechanisms have been considered.

Specifically, the results showed in Table 11 indicate that the traffic losses reduction ranges from the 62% to 77%, according to the values of the H2 timer and the RPR protection mechanism.

H1 = 10 ms	<i>DHOT vs. SHOT: Traffic Lost reduction</i>	
H2 (ms)	<i>RPR wrapping</i>	<i>RPR steering</i>
20	62.5%	62.0%
30	71.5%	71.1%
35	74.5%	74.2%
40	77.0%	76.6%

Table 11: DHOT vs. SHOT: Packets lost

Table 12 depicted this gain when fixing the value for the H2 timer (30 ms) and considering different values for the H1 timer. The aim is to consider that the implemented failure detection at the optical switched strongly depend from the optical components. Specifically, the results indicate that the traffic losses reduction ranges from the 52% to 71%, according to the values of the H2 timer and the RPR protection mechanism.

H2 = 30 ms	<i>DHOT vs. SHOT: Traffic Lost reduction</i>	
H1 (ms)	<i>RPR wrapping</i>	<i>RPR steering</i>
10	71.5%	71.1%
15	63.9%	63.6%
20	57.5%	57.2%
25	52.7%	52.5%

Table 12: DHOT vs. SHOT: Packets lost

We carried out a fourth simulation case study in order to complete the comparison between the SHOT and the DHOT approaches when the failure occurs at the RPR layer. In particular, we

evaluated the recovery time that is the time required from the failure. This is the time required for the network reconfiguration after the failure and it encompasses the time required for the ring wrapping plus the time required by the TD algorithm. Table 13 reports the recovery times when the H1 timer is set to 10 ms and various values for H2 timer are considered. If the DHOT approach is used, the recovery time is given by the H1 timer plus the time required by RPR layer to recover from the failure (few ms). Contrarily, if the SHOT would be used, the recovery time, since the optical layer is not able to recover from the failure, is given by the total hold-off timer (H1+H2) plus the time required by the RPR mechanisms.

H1 = 10 ms	DHOT	SHOT
H2 (ms)	<i>RPR wrapping (ms)</i>	<i>RPR wrapping (ms)</i>
20	12.55	32.57
30	12.55	42.58
35	12.55	47.56
40	12.55	52.57

Table 13: DHOT vs. SHOT: Recovery Time

Specifically, it can be observed that DHOT performs much better than the SHOT approach. In fact, the time required by DHOT for recovery is about from the 23% to the 38% of the time which would be required by using the SHOT.

Table 14 reports the recovery time for different values of H1 while H2 is fixed to 30 ms. Also in this case, it can be observed how the recovery time is much lower using the DHOT than using the SHOT. Specifically, the recovery time required by using the DHOT is from the 29% to 48% of the recovery time required by using the SHOT.

H2 = 30 ms	DHOT	SHOT
H1 (ms)	<i>RPR wrapping (ms)</i>	<i>RPR wrapping (ms)</i>
10	12.55	42.58
15	17.57	47.56
20	22.57	52.60
25	27.56	57.55

Table 14: DHOT vs. SHOT: Recovery time

Finally, it worth noting that this percentage strongly depends on the actual traffic load, the failure scenario and the set of the double hold-off timers.

6.3 Resilience Interworking in RPR over ASON/GMPLS networks

As stated in the previous Sections, the protection at RPR layer implies “some” reduction of the available bandwidth at logical level.

Let us suppose that the traffic transported by the RPR ring and generated from the higher layers (i.e., IP/MPLS) is uniformly distributed among the ring nodes, hence the maximum offered traffic to avoid the ring saturation (after the ring reconfiguration) is $\rho = 0.5$. In fact, in this situation, the available bandwidth is halved after the wrapping/steering. If the offered traffic is higher than 0.5, the recovery action implies the saturation of the ring, basically blocking the low priority (Class C traffic). On the other hand, both the Class A and B have to be well-engineered in order to avoid packet losses due to the ring saturation.

On the other hand, as vastly discussed in Chapter 2 and 3, in an IP/MPLS over RPR environment, the client traffic offered to the ring is characterized by its fluctuations over time (e.g. on a daily time basis).

We propose here a procedure which aim consists in using the automatic switching of optical connections capabilities provided by the ASON networks to face with both the potential ring saturation in case of failures and the fluctuations over time of the traffic inserted in the RPR ring. By implementing this procedure, the available bandwidth of the ring is automatically increased/decreased when strictly required by the ring status.

Basically, this procedure is based on introducing at the IP/RPR routers a monitoring function in order to compute periodically (e.g., each Observation Window) the traffic being carried by the light paths connecting a couple of RPR nodes.

As illustrative example, let us suppose that IP/RPR routers are connected through light paths (e.g., permanent optical connections set up by the NMS). This is depicted in Figure 65 (a). If failure occurs and it has to be handled at the RPR layer by using, for example, the ring wrapping mechanism, the monitoring function is used to detect an overloading condition (i.e. ring saturation) as a consequence of the ring reconfiguration. When the over-loading condition is detected on the light path connecting two RPR nodes, then the IP/RPR router requests via UNI interface to the optical connection controller (CC) of the optical components (e.g., OXC) the dynamic establishment of a switched connection. If the GMPLS signalling is able to provide the switched connection, the two RPR nodes are connected through two light paths (Figure 65 (b)).

Then, by applying some TE rules (i.e., Load Balancing), the traffic to be transported is distributed between the two light paths.

On the other hand, when the monitoring function detect that the traffic between the two routers can be carried by only the permanent connection (under-utilization condition), then the router requests to GMPLS-based control plane the tear down of the switched connection.

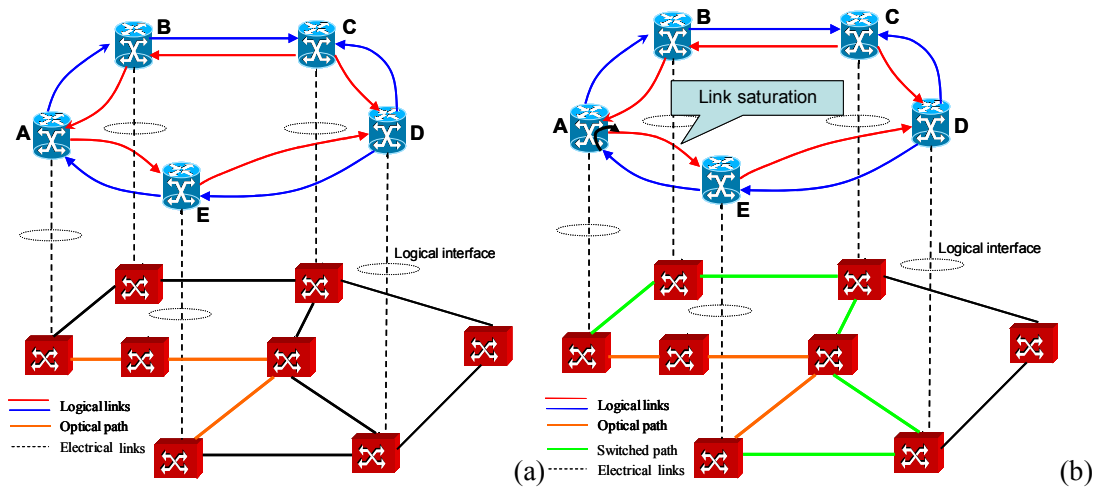


Figure 65: RPR over ASON interworking in case of failure, (a) Initially condition, (b) Avoiding ring saturation by requesting a switched connection

For the over-loading/under-utilization of the light paths, a threshold-based policy (the one defined for the TRIDENT procedure) is used, which means using a high threshold for the congestion detection and a low threshold to detect the under-utilization condition [110].

Of course, the tear down of the switched connection is also requested also when the failure has been physically solved and the ring logical bandwidth comes again to initial conditions.

Nevertheless, the implementation of such a mechanism implies the use of spare RPR cards to be used when saturation occurs

The aim of this interworking procedure is to keep limited the size of the transport networks. Indeed, when the failure has to be recovered at the RPR layer, the automatic switching capabilities offered by the ASON/GMPLS networks can be used to avoid to dedicate spare resources to protect at the optical level the logical links between routers. In general, it has to be considered that a transport network support different client networks and thus, avoiding to overdimension the transport network allows reducing the CAPEX for Network Operators.

The following Figure 66 illustrates the flow-chart of the suggested procedure.

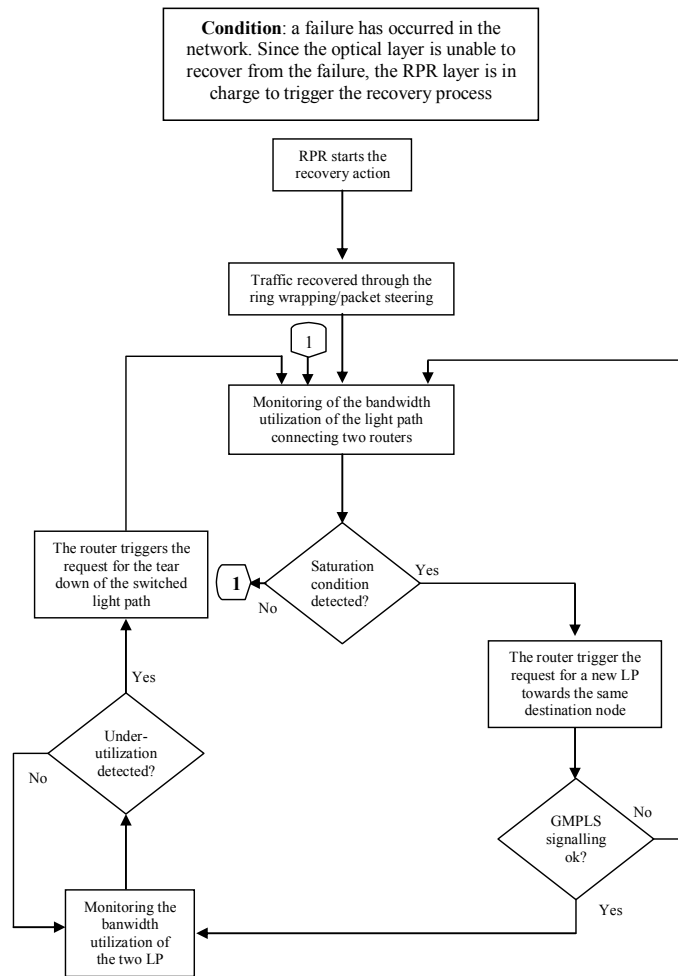


Figure 66: RPR over ASON interworking: flow-chart

