



# Kent Academic Repository

Chennu, Srivas, Habel, Kai and Langer, Klaus-Dieter (2006) *Protected Ethernet Rings for Optical Access Networks*. In: 2006 ITG Symposium on Photonic Networks. VDE, pp. 29-36. ISBN 978-3-8007-2959-3.

## Downloaded from

<https://kar.kent.ac.uk/14493/> The University of Kent's Academic Repository KAR

## The version of record is available from

## This document version

UNSPECIFIED

## DOI for this version

## Licence for this version

UNSPECIFIED

## Additional information

## Versions of research works

### Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

### Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in *Title of Journal*, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

## Enquiries

If you have questions about this document contact [ResearchSupport@kent.ac.uk](mailto:ResearchSupport@kent.ac.uk). Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

# Protected Ethernet Rings for Optical Access Networks

Srivasa Chennu, Kai Habel, Klaus-Dieter Langer  
Fraunhofer Institute for Telecommunications, Heinrich-Hertz-Institut  
Einsteinufer 37, 10587 Berlin, Germany  
{chennu, habel, langer}@hhi.fhg.de

## Abstract

In this paper we propose a centralized link layer architecture for providing low latency fault recovery for optical access rings. This architecture exploits the naturally uneven breakdown of network management responsibilities between the components of an access ring. Important administrative operations like ring status checking, fault detection and recovery are aggregated at the HUB component located in the Central Office of the access network. Consequently, when compared with a standardized Rapid Spanning Tree Protocol implementation, the design of the proposed architecture benefits from a simplified link layer design of the Optical Network Unit, in addition to a significantly reduced fault recovery delay in the ring. We also present an Ethernet-based protocol that realizes our centralized protection model. The design principle of this protocol, responsible for the message passing required to react to topology changes in the network, is simple enough to allow quick reaction times, and to support QoS-aware prioritization of network traffic. The performance of the proposed architecture is evaluated using analytical and simulative means, and the performance aspects of the ring protocol relevant to network protection are compared with those provided by the Rapid Spanning Tree Protocol.

## 1 Introduction

Access networks form the ‘first mile’ of network infrastructure that connect Internet users to the high-capacity core networks forming the backbone of the Internet. These typically small-range networks merge traffic from end users into Metropolitan Area Networks that in turn pool traffic into larger internetworks. The lack of adequate bandwidth capacity in current access networks hinders the optimum leveraging of the large capacity of the backbones for providing broadband services to end users. Optical fiber technology, hitherto restricted to the backbone networks, holds the potential to relieve this bottleneck, and hence can deliver cost-effective, high-bandwidth Triple Play services to users.

For effective and reliable deployment of optical fibers in the access network, important design considerations need to be addressed. Pertinently, the large quantities of QoS-constrained data traffic generated by high-bandwidth services stress the need for survivability and fault protection to be incorporated in the design of the network. The speed and quality parameters specified for this resilience depend on the requirements of network applications that are to be supported. In this context, research in the IST MUSE project [1] has identified a wide range of QoS application classes in multi-service access networks [2]. Toward delineating a general framework for designing protected Ethernet architectures, the Metro Ethernet Forum defines a range of Restoration Time Categories suitable for a variety of network services [3].

Network protection can be incorporated at different layers in the network stack. In addition, protection facilities located in these layers can complement each other, to provide

a high level of network availability from the perspective of the end users. At the physical layer, a system for providing protection by means of redundant paths and hardware in Passive Optical Networks (PON) has been recommended by the ITU [4]. The research in [5] has demonstrated the feasibility of CWDM-based PON rings for constructing such last mile networks, as they inherently provide redundant paths for protecting end user connectivity in the event of a network fault. The focus of this paper is a centralized network management architecture and protocol that operates at the link layer in access rings. The so titled Fast Access Ring Protection Protocol (FARPP) is a simple, lightweight protocol that manages Ethernet-based traffic in the access ring. It exploits the functional simplicity therein, keeping the protocol overhead to a minimum. The physical path redundancy provided by the physical layer is used by FARPP to reconfigure the ring in response to ring faults. Owing to its lightweight design, it is intended to support low-latency path restoration in access rings, which can be tailored to satisfy the protection requirements for different service classes. Initial simulative evaluation of the protocol is focused on verifying its effectiveness in a realistic edge network.

This paper is organized as follows. Section 2 compares existing network protection schemes with regard to their advantages and drawbacks for deployment in access ring topologies. Section 3 begins by briefly describing the physical architecture of the optical access ring. It then elaborates on the link layer architecture and protection protocol that is the focus of this paper. Section 4 concentrates on the analytical and simulative evaluation of FARPP and presents results of the simulation effort. Finally, Section 5 concludes the paper, with a summary of further work.

## 2 Existing network protection schemes

There currently exist a variety of protection schemes designed for ring topology networks, based on different underlying technologies. It is pertinent to the aim of developing a link layer architecture for optical access networks, to study the functionality provided by these existing technologies. This effort serves the dual purpose of extracting useful design principles incorporated in these schemes that were applicable to the access network domain, and for identifying areas where important improvements could be made over them.

### 2.1 SDH/SONET technology

The Synchronous Digital Hierarchy (SDH) standard [6] is a well-known technology that provides protection for rings. It is a mature, standardized technology and currently finds wide deployment, as it is specially suited for ring topologies. SDH is a hierarchical technology suited for consolidating separate smaller traffic streams into larger ones, and inherently supports low restoration times, typically in the order of 50 ms. The associated drawback in using SDH is its inflexible scheme of bandwidth allocation, particularly unsuitable for bursty IP or multicast traffic. Furthermore, cost of SDH hardware is much higher than that of Ethernet switching hardware. Ring protection is provided by reserving 50% of available resources, used in the event of a fault. SDH does not balance traffic load on the ring, i.e. links closer to the HUB carry more traffic load than those further away. Lastly, since it is not based on statistical multiplexing, efficient bandwidth utilization requires the additional deployment of ATM at a higher layer.

### 2.2 Rapid Spanning Tree Protocol

The IEEE 802.1d Rapid Spanning Tree Protocol (RSTP) standard [7] is widely employed for constructing loop-free tree topologies from general mesh networks of Ethernet switches. Working in conjunction with Ethernet, it is a lower-cost solution than SDH. It is a standardized and well-understood mechanism, and is easily configurable and upgradeable. It is efficient in its use of ring bandwidth, and as opposed to SDH, requires no resource reservation for protection. But the Spanning Tree Protocol was designed to work on general mesh topologies, and does not exploit the simple configuration of access rings. This lack of optimization translates to relatively high restoration times in the order of 2-30 s, as fault notification proceeds upstream in a serialized manner from the original location.

### 2.3 Resilient Packet Ring

The IEEE 802.17 Resilient Packet Ring (RPR) standard [8, 9] from IEEE is a relatively new scheme currently in development, designed specially for application in the emerging Metro Ethernet domain. Its design is specifically optimized for the ring topologies deployed in this domain, and hence prescribes a maximal restoration time of 50 ms.

It is an Ethernet based, bandwidth efficient solution, suitable for bursty IP and multicast traffic. RPR also incorporates complex signaling mechanisms for fairness of bandwidth sharing amongst ring nodes, and is consequently is a difficult protocol to correctly implement. This means that hardware devices that implement the RPR standard are relatively new and relatively expensive. A complete implementation of RPR would also be an overkill for the simpler architecture envisioned for optical access rings.

### 2.4 Proprietary solutions

Many system vendors have developed specially tailored solutions for providing protection in Metro rings. Some of the most popular ones are listed below.

**Extreme Networks** offers switches equipped with *Ethernet Automatic Protection Switching* (EAPS) [10] technology.

**RAD Data Communications** offers switches equipped with their *Resilient Fast Ethernet Ring* (RFER) [11] technology.

**Telco Systems** offers *T-Metro* [12] network equipment for ring topologies.

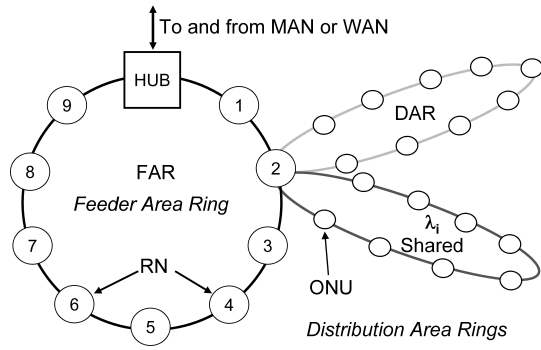
All of these proprietary solutions provide fast reconfiguration, typically less than 50 ms. They bundle together a large number of aggregated features at different layers of the network stack, and are primarily targeted at metro ring installations. The major disadvantage hindering the wider application of the technologies devised therein is the lack of open documentation and standardization. The concomitant lack of interoperability of the protection mechanisms used by the vendors of these systems is unattractive to network operators wishing to construct access rings using heterogeneous network hardware.

## 3 The optical access ring architecture

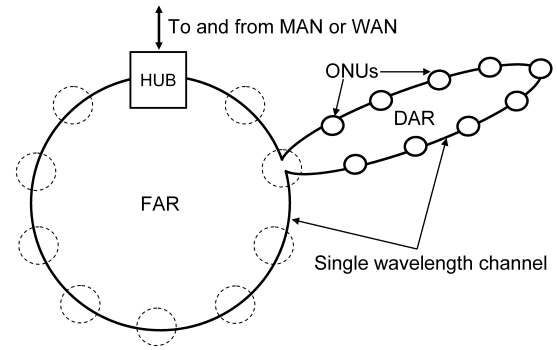
This section begins with a description of the underlying physical architecture of the optical access ring. The network entities at the link layer and FARPP are then discussed in detail.

### 3.1 The physical layer of the access ring

The design of the access network at the optical layer involves the definition of a suitable topology for the network, identification of channel multiplexing schemes, and the design of network components for optical transmission. Briefly summarized here is the prior research in [5] on protection-enabled optical access rings. Such ring structures inherently provide for network resilience, because each node along the ring can be reached by two disjoint paths. Hence protection at this layer is a matter of switching to the alternate path to a node when the first one fails during normal operation.



**Figure 1** The optical access ring at the physical layer



**Figure 2** A wavelength channel seen at the link layer

Figure 1 depicts the proposed access network architecture, consisting of two ring sub-structures: A single Feeder Area Ring (FAR), and many Distribution Area Rings (DAR) connected to it along its circumference.

The FAR is a bidirectional CWDM ring formed by one or more optical fibers. A CWDM optical fiber carries up to 18 wavelength channels, each of which supports either single fiber or dual fiber technologies. The FAR consists of a HUB component located at the Central Office (CO) of the access network, and Remote Nodes (RN) located out in the field. RNs are completely passive network elements, and consist of Optical Add Drop Multiplexers (OADM) that add or drop a specified wavelength channel onto a designated DAR.

The DAR is a bidirectional ring based on shared access of the wavelength channel provided to it by an OADM in the RN to which it is connected. A DAR connects one or more Optical Network Units (ONU), which are typically located at or close to the end user premises, depending on the extent of deployment of the optical fiber in the access network. ONUs are active network elements connected in a daisy-chain scheme, responsible for forwarding network traffic to and from end users. Data destined for an ONU is forwarded along by intermediate ONUs in the DAR till it reaches its destination. Similarly, data originating at an ONU is forwarded along toward the HUB.

### 3.2 The link layer of the access ring

At the link layer of the access network, only the active elements of the physical layer are visible. Consequently, the view from this layer looks like that depicted in Figure 2. The HUB and the ONUs are the only network entities having a link layer component. Further, a ‘ring’ at the link layer comprises the HUB, and the ONUs sharing a wavelength channel in a DAR. Consequently, the link layer protocol in the access ring operates independently in each DAR of the network. This protocol has been tailored to the optical access network introduced previously. In addition to the typical link layer tasks like data framing, error correction, etc. this architecture has been designed for fast

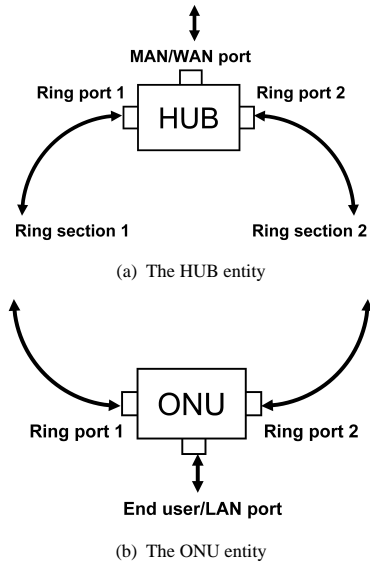
network reconfiguration and path recovery in the event of network faults. To accomplish this, the architecture exploits the naturally uneven breakup of responsibilities in the access network. By constructing a centralized network management model, it affords a simplified management policy, and reduced latency times for working around network faults. To this end, the protocol functionality is split up between two distinct entities elaborated on below.

#### 3.2.1 The HUB entity

The HUB entity of the access ring, a trusted component in the direct control of the network operator, is responsible for important network management tasks. It is responsible for least-cost forwarding of incoming and outgoing data in the access ring. It balances traffic load on the different links forming the ring, so as to generate an even traffic distribution across the ring. This avoids overloading some links more than others, keeps their utilization levels within acceptable limits, during both normal operation and fault recovery. The HUB also performs regular status checks to maintain a consistent view of current topology of the ring, and when a ring fault does occur, it reacts by updating its view of the topology and appropriately re-routing data around the fault.

The HUB and its ports are depicted in Figure 3(a). It is connected to the access ring via two ring ports, and to a metro or wide area network via another port. This link layer entity is considered to be the only ‘intelligent’ node in the ring, responsible for maintaining a consistent view of the current topology and reacting to changes in the same. Further, for the access network, it serves as a gateway to the metro or wide area network, forwarding and receiving data to and from the ONUs.

During normal operation i.e. when the ring is complete, the HUB uses its current view of the ring topology to optimize data routing in the ring. This optimization involves the choice of the least-cost path to reach an ONU, where the cost is a parameter defined on a per-link basis by the network operator. In the event of a ring fault, the HUB suitably routes traffic around the failure, so as to main-



**Figure 3** Protocol entities in the link layer architecture

tain connectivity to all available ONUs. When the fault is eventually repaired, it returns the ring back to its normal operational mode.

The HUB entity implements the ring management functionality required for routing optimization and reacting to ring faults. To do so, it splits the ONUs in the ring into two sections, one for each of its ring ports, such that the maximum path cost to an ONU in either section is minimized. During initialization of an ONU, the HUB receives JOIN messages on its ring ports. It then selects one of its ring ports for communicating with it and acknowledges the ONU's registration with a JOIN ACK message. The ONU's ID is then stored in an internal lookup table located at the HUB, along with the ring port chosen for the ONU. Consequently, frames coming in to the ring are selectively forwarded on one of the two ring ports, depending on the ring port registered for the destination ONU in the lookup table. During normal operation of the ring, the HUB regularly circulates HEALTH frames along both sections of the ring to ensure that the ring is closed. When a ring fault does occur, the lookup table is reconfigured so that the ONUs affected by the fault are shifted to the other section. This ensures that an alternate path is quickly restored, minimizing service disruption.

### 3.2.2 The ONU entity

As a consequence of the concentrating link layer complexity in the HUB, an ONU can be designed to be a simple, low-cost device deployed close to the end user, saddled with only the basic set of responsibilities. It forwards data destined for and originating from end users, and periodically reports its own operational status, and those of the links connecting it to the ring.

The ONU protocol entity, depicted in Figure 3(b), is connected to the access ring via two ring ports, and has an

additional port to interface with the end user or a local area network therein. ONUs are designed to be simple nodes, controlled by the HUB via administrative messages circulated on the ring. Each ONU has an identifier assigned uniquely within the ring. It is included in all frames sent out by the ONU, and is used by the HUB to register and address it. The key parameters that control traffic flow in and out of an ONU are its PRIMARY and SECONDARY ports. Under normal operating conditions, the PRIMARY port is the ring port on which an ONU receives data and administrative messages from the HUB, and sends out data to it. The SECONDARY port is the other ring port of the ONU. The ONU simply forwards data received on this port. These ports are set by the ONU on initialization, which is begun by the ONU sending out JOIN messages to the HUB on both its ring ports. The HUB collects these messages, and instructs the ONU to set its PRIMARY port to the one providing the least cost path to it, by sending out a JOIN ACK message on this least cost path.

During normal operation of the ring, the ONU receives and forwards the HEALTH messages it receives from the HUB. In addition, it regularly sends out its own HELLO messages on its PRIMARY port to maintain its membership in the access ring. This mechanism ensures that failed ONUs are eventually de-registered at the HUB. During ring faults, a subsection of the ONUs in the ring are no longer connected to the HUB via their PRIMARY port. On detection of the fault, they switch their data traffic to their SECONDARY port to work around the fault, and begin to send out HELLO messages on their SECONDARY ports. On receipt of this new HELLO messages the HUB reconfigures its lookup table to record the new forwarding path to the affected ONUs. When full connectivity is eventually restored in the ring, the ONUs return to using their PRIMARY port for communicating with the HUB. This simple mechanism for switching between ring paths ensures that the protocol overhead for restoring connectivity in the ring is kept as low as possible.

## 3.3 FARPP timers

The timers that control the operation of the HUB and the ONUs are crucial in maintaining synchronization between them. Described below the timers that run at the HUB and the ONU for this purpose.

### 3.3.1 At the HUB

**JOIN Receive Timer** specifies the amount of time the HUB waits for the second JOIN frame from an ONU after receiving the first one. If the HUB receives the second JOIN frame before the expiry of this timer, it computes the least-cost path and sends out a JOIN ACK frame to the ONU on the selected path. On the expiry of this timer, the HUB simply discards the first JOIN frame received from the ONU. The value of this timer should be at least equal to the time required for a JOIN frame to traverse the entire ring.

**HEALTH Send Timer** specifies the amount of time between two consecutive HEALTH frames sent out by the HUB. A low value for this timer provides for quicker reaction to ring faults.

**ONU Expiry Timer** specifies the amount of time an ONU's entry in the lookup table at the HUB ring ports is valid. If a HELLO frame from an ONU is not received by this time, its entry is removed from the tables. The value of this timer is set to be greater than or equal to the HELLO Send Timer of the ONU.

### 3.3.2 At the ONU

**JOIN Send Timer** specifies the amount of time an ONU waits after sending out a pair of JOIN frames on its ring ports. If the ONU does not receive a JOIN ACK from the HUB before the expiry of this timer, it sends out the JOIN frames again. The value of this timer should be greater than double the time required for a JOIN frame to traverse the entire ring.

**HELLO Send Timer** specifies the amount of time between two consecutive HELLO frames sent out by the ONU. The value of this timer can be flexibly configured.

**HEALTH Receive Timer** specifies the amount of time an ONU waits after receiving a HEALTH frame from the HUB. On the expiry of this timer, it assumes that a path to the HUB via its PRIMARY port is no longer available, and begins sending out HELLO frames from its SECONDARY port. This value of this timer is set to be greater or equal to the HEALTH Send Timer of the HUB.


### 3.4 FARPP timing diagram

The timing diagram in Figure 4 illustrates the flow of the protocol messages over time. Specifically, the messages exchanged between an ONU and the HUB during the four operational phases of FARPP, namely Initialization, Normal Operation, Link Down and Link Up.

### 3.5 FARPP protocol messages

The structure of the FARPP header is described in this section. This header is encapsulated in a standard Ethernet frame, and carries control information between the ONUs and the HUB. Summarized below are the important fields in this header along their functions.

**TYPE** - this field contains the type of the control frame. FARPP control frames are classified into DATA and ADMIN frames. DATA frames carry an application-specific payload. They contain DATA as the value of their TYPE field and correspond to a full-sized Ethernet frame of 1538 bytes (including an inter-frame gap of 12 bytes). ADMIN frames, on the other hand, carry protocol-related informa-



**Figure 4** Timing diagram depicting message exchange between an ONU and the HUB

tion. They correspond to a standard Ethernet header of 38 bytes (including an inter-frame gap of 12 bytes) with no payload. The TYPE field of an ADMIN frame can contain one of the following values, depending on the protocol-specific information it carries.

- JOIN - indicates a JOIN frame sent out by an ONU both its ring ports during initialization. The payload of a JOIN frame contains the accumulated path cost to the originator of the frame, and is updated by every ONU along the way to the HUB.
- JOIN\_ACK - indicates a JOIN ACK frame sent by the HUB to an ONU on the shortest path to it, to acknowledge an ONU's registration.
- REJOIN - indicates a REJOIN frame sent by the HUB to instruct an ONU to reinitialize itself.
- HEALTH - indicates a HEALTH frame circulated on the ring by the HUB.
- HELLO - indicates a HELLO frame sent by the ONU to maintain its registration at the HUB.

**SYSTEM\_ADDR** - contains the unique identifier of the source or destination of the frame, depending on whether an ONU or the HUB is the originator of the frame, respectively.

**PROTOCOL\_INFO** - contains protocol-related information used in conjunction with the control frame types listed above. This field could eventually be extended as additional functionality is added to FARPP.

### 3.6 Hardware based failure detection

In the case that detection of link failures is reliably provided by the optical layer, this functionality can be suitably employed at the link layer to provide an increased level of

robustness to network applications. Two methods to do so are discussed below.

### 3.6.1 Reduced detection latency

The amount of time required for detecting a failure at an ONU contributes significantly to PRT. If available, event-based failure detection could be used to complement the timer functionality used in FARPP to detect failures. On receiving a notification of link failure from the optical layer, an ONU could immediately begin broadcasting HELLO frames on its SECONDARY port, instead of waiting for the expiry of its HEALTH Receive Timer.

### 3.6.2 Frame reflection

Hardware based failure detection at an ONU would enable it to redirect data frames not intended for it back to the HUB, instead of simply forwarding them on a failed link. This ‘frame reflection’ technique, proposed for Metro rings as *Wrap Protection* [13], would appropriately handle frames mis-routed at the HUB during the transitional period ensuing immediately after the occurrence of a fault, and before the HUB has updated its lookup table to reflect the new topology. Such frames would hence be received back at the HUB, which could then suitably forward them according to the updated location of the recipient. These frames would consequently arrive at the intended destination instead of being lost, albeit out of order in the frame sequence.

## 4 Evaluation of FARPP

This section discusses a preliminary analytical and simulative performance evaluation of FARPP. As a first step toward the development of a detailed simulation and modeling of the protocol, a simple scenario usage scenario has been considered in this paper: a collection of ONUs in a DAR described in Section 3.1, which is allocated a optical wavelength channel providing a line rate of 1 Gbps. Each ONU in the DAR receives a Constant Bit Rate (CBR) stream of data frames constituting a High Definition Television (HDTV) video stream. The data rate of this stream is set to 12.304 Mbps, which translates to one complete Ethernet frame of 1538 bytes (including header and payload) transmitted every millisecond. Additionally, a uniformly random dither is added to the inter-frame transmission interval to vary the bit rate around the CBR value. In such a usage environment, one of the active links in the ring is disabled at a predetermined time, and the amount of time required for the restoration of the video stream to an affected user is measured.

### 4.1 Analysis of reconfiguration time

Figure 5 illustrates on a chronological time scale the events that occur during the process of reconfiguration as seen by an end user connected to an ONU in the DAR described above. The key parameter for evaluating the applicability

of FARPP is the Path Restoration Time (PRT) seen by the end user immediately after the occurrence of a fault. This measurable time is defined as the end-to-end time required for the protocol to react to a ring fault and suitably reconfigure itself to provide a valid path to each ONU. To facilitate an analysis of the same, the generic model for failure convergence schemes in telecommunication networks described in [14] has been applied to FARPP.

The end-to-end PRT as seen by an end user can be broken down into its individual components:

**Fault Detection Time (FDT)** is the amount of time that elapses between the actual occurrence of the fault and its detection at the ONU. If this detection is based on timers, the worst case value for this delay is equal to the HEALTH Receive Timer at the ONU.

**Notification Time (NT)** is the amount of time that elapses between the HELLO frame being sent out on the SECONDARY port of the ONU, and its reception at the HUB. This delay scales with the number of intermediate ONUs that the HELLO frame must transit, and the delay that it suffers in each of them. Consequently, the worst case value for this delay occurs when the link between the HUB and an immediately neighboring ONU along the ring fails. In this scenario, the HELLO frame must then traverse the entire length of the ring to reach the HUB.

**Recovery Operation Time (ROT)** is the amount of time that elapses starting from the HELLO frame being received by the HUB, which then updates its lookup table, till a DATA frame is eventually received by the ONU on the restored path. This delay scales with the number of intermediate ONUs that the DATA frame must transit. The delays it encounters on the way will be affected by the increased traffic load that the available links in the ring would have to carry during faults. The worst case value for this delay also occurs when the link between the HUB and an immediately neighboring ONU along the ring fails.

Summarizing the above, PRT can be expressed as:

$$PRT = FDT + NT + ROT. \quad (1)$$

### 4.2 FARPP simulation setup

A simulation of FARPP has been constructed using the Network Simulator [15]. A block diagram of the simulation setup is shown in Figure 6. It illustrates the customized modules designed for the FARPP simulation. Link modules (*LN*) represent the bidirectional links interconnecting the nodes in the access ring. Each ring node (*RN*) is an aggregate object composed of a pair of network interfaces, *eth0* and *eth1*. Each interface in turn contains a physical (*PHY*) and a medium access (*MAC*) layer, and an interface queue (*IFQ*) that prioritizes FARPP ADMIN frames over DATA frames. In addition, agents at the HUB




Figure 5 Breakdown of PRT into its components




Figure 6 Simulation setup for FARPP in NS-2

and the ONUs that model the FARPP protocol entities described previously have been developed.

In a given simulation run, the number of ONUs connected in a ring is specified by the parameter *ONUCount*. The values of the timer parameters that remain constant across all runs are listed in Table 1. It is to be noted that these values have not been specifically optimized, which is to be considered in future work.

In order to measure PRT in the simulation, a single HUB agent and *ONUCount* ONU agents are instantiated. Each of these agents is attached to a ring node, whose interfaces are connected together in series using links to form a DAR. In addition, *ONUCount* pairs of UDP agents and CBR applications attached to the HUB agent generate the HDTV streams destined for the ONUs. In conjunction, a sink is attached to each ONU agent, and is connected to its counterpart UDP agent at the HUB. The CBR applications feed their streams to the HUB agent, which forwards them toward the destination ONU agent. On reception of the stream at the ONU, the data is passed up to the sink. In order to simulate a worst-case traffic distribution in the ring, the costs associated with the links forming the ring are set such that, during normal operation, all the streams are forced to flow anti-clockwise along the ring. Consequently, the link connecting the last ONU to the HUB,  $LN(n:0)$ , is the most heavily loaded link, as it carries all the streams flowing through the ring. At intermediate points during a simulation run, the state of this link toggled on

Parameter	Location	Value
joinReceiveTimer	HUB	1.0 s
helloReceiveTimer	HUB	2.0 s
healthSendTimer	HUB	0.05 s
joinSendTimer	ONU	0.5 s
helloSendTimer	ONU	1.5 s
healthReceiveTimer	ONU	0.06 s

Table 1 Control parameters for FARPP simulation runs

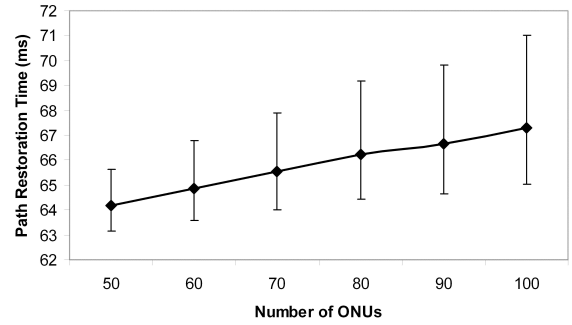


Figure 7 Scaling of PRT with number of ONUs

and off. Each time the link is switched off, a reconfiguration is triggered at the last ONU. Hence, the time it takes for the data stream destined to this ONU to be resumed after each disconnection represents the worst-case bound for PRT in this evaluation scenario.

### 4.3 Simulation results

Figure 7 depicts the PRT plotted against a linear increase in the number of ONUs (*ONUCount*) comprising the access ring. Each data point in the graph is generated by one simulation run, and represents the average PRT for the last ONU attached to the most heavily loaded link. The PRT is measured approximately 200 times during each run, and a mean value with a 95% confidence interval has been calculated and plotted from these individual measurements. As is deducible from the graph, though the utilization of the most heavily loaded link in the ring eventually reaches 100% with 100 ONUs, the PRT remains within bounds, discounting frame losses due to buffer overflows, which have not been considered for this simulation scenario. Fur-



ther, it can be seen that PRT is largely influenced by the HEALTH Receive Timer at the ONU. This timer has been set to a value of 60 ms for the simulation runs plotted in the graph (listed in Table 1) and corresponds to the FDT component in Equation 1. Since the combined contribution of NT and ROT to PRT is relatively small, the frequency of HEALTH packets sent out by the HUB and the associated timer values can be suitably configured to control the re-configuration latency of the access ring. Alternatively, if hardware-based failure detection discussed in Section 3.6 is available, the dominating contribution of FDT on PRT can be effectively eliminated.

The reconfiguration times measured in the FARPP simulation compare favorably to RSTP, which entails an increased processing overhead at each participating Ethernet bridge. RSTP control frames propagate information about the failure event progressively from the point of occurrence to neighboring nodes in the network. These nodes then react to the receipt of these topology change notifications to appropriately reconfigure their network ports. This distributed reconfiguration scheme eventually results in a per-hop delay in the order of 100 ms at each node, rendering it too slow to protect high-bandwidth data streams against disruption.

## 5 Conclusions and further work

This paper has outlined an optimized architecture for link layer protection of optical access rings. This architecture proposes a centralized network management model to exploit the salient features of access rings, and affords significant simplifications over existing schemes. In conjunction, FARPP, an Ethernet-based protocol that incorporates this design philosophy into the design of the HUB and ONU protocol entities has been modeled. This abstract model, which describes the message passing between the protocol entities at the link layer, has been used to evaluate the ability of the link layer architecture for providing low-latency reconfiguration. Although the evaluation is presently in a preliminary stage, a full-fledged model is in development, and is expected to provide us with concrete results that characterize the performance of FARPP under heavy network load conditions in a realistic access network environment.

In addition to a detailed evaluation model, the deployment of FARPP in an FTTP optical access network necessitates important enhancements in different aspects of its functionality. To mention a few, a comprehensive facility for notifying the network operator or a higher management layer of failures and repairs in the ring needs to be incorporated. When extended with this functionality, FARPP would be able to follow up its quick reaction to failures with a suitable report indicating the temporal and topological details of faults. Another essential feature is the incorporation of QoS-aware traffic prioritization in the access network, since the end-to-end QoS parameters agreed upon at the application level need to be honored as data flows

down the network stack. In conjunction, an effective and scalable admission control mechanism would be required for maintaining the load on the access network within the limits for which it is dimensioned. These features would in combination contribute toward ensuring better resiliency in the access network.

## Acknowledgments

This research is being supported by the IST-FP6 Integrated Project MUSE (Multi Service Access Everywhere) [1] under contract Nr. 507295. The authors would like to thank Prof. Dr. Ulrich Killat at the Hamburg University of Technology, and Jelena Grubor and Michael Schlosser at the Heinrich-Hertz-Institut for their guidance and discussions during the research leading up to this paper.

## 6 References

- [1] The IST MUSE project, [www.ist-muse.org](http://www.ist-muse.org), 2006
- [2] A. J. Elizondo Armengol and G. M. Gallizo Rueda: MUSE - Network Requirements for multi-service access, November 2004
- [3] Metro Ethernet Forum: Requirements and Framework for Ethernet Service Protection in Metro Ethernet Networks, February 2004
- [4] ITU-T Recommendation G.983.5: A broadband optical access system with enhanced survivability, January 2002
- [5] J. Grubor et al.: Feasibility of protected rings in optical access networks, ITG Fachbericht 182, Photonische Netze, pp. 155-162, Berlin, Offenbach: VDE Verlag, 2004
- [6] ITU-T Recommendation G.707: Network node interface for the synchronous digital hierarchy (SDH), December 2003
- [7] IEEE Std 802.1D-2004 Media Access Control (MAC) Bridges
- [8] IEEE Std. P802.17-2004 Resilient Packet Rings
- [9] The Resilient Packet Ring alliance: An introduction to Resilient Packet Ring technology, [www.rpralliance.org](http://www.rpralliance.org), October 2001
- [10] IETF RFC 3619: Extreme Networks Ethernet Automatic Protection Switching, October 2003
- [11] RAD Data Communications: Resilient Fast Ethernet Ring Technology, [www.rad.com](http://www.rad.com), November 2002
- [12] Telco T-Metro, Telco Systems, [www.telco.com](http://www.telco.com), 2005
- [13] Ziwen Lian et al.: Resilient Ethernet ring for metropolitan area networks, The Ninth International Conference on Communication Systems, 2004, pp. 316 - 320
- [14] De Greve, F. et al.: Rapidly Recovering Ethernet Networks for Delivering Broadband Services on the Train, The IEEE Conference on Local Computer Networks, 2005, pp. 294 - 302
- [15] The Network Simulator ns-2 version 2.29: <http://www.isi.edu/nsnam/ns>, 2006