QoS-aware Traffic Protection for Access Rings

Srivas Chennu[1,2], Kai Habel[1], Klaus-Dieter Langer[1]
[1] Fraunhofer Institute for Telecommunications, Heinrich-Hertz-Institut
Einsteinufer 37, 10587 Berlin, Germany
Tel: +49 30 31002 586, Fax: +49 30 31002 250, E-mail: chennu@hhi.fhg.de
[2] Hamburg University of Technology,
Schwarzenbergstrasse 95, 21073 Hamburg, Germany

In this paper we propose a centralized link layer architecture for providing low latency fault recovery for access rings. This architecture exploits the naturally uneven breakdown of network management responsibilities between the elements of the ring. Administrative operations like ring status checking, fault detection and recovery are aggregated at the HUB. The proposed architecture benefits from a simplified link layer design of the ONU, in addition to lower delays in recovering from faults in the ring. We 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 resulting from faults, is to be simple enough to allow quick reaction times. The architecture prioritizes the restoration of network services according to their QoS importance, to ensure that the quality guarantees promised to them are honored at all times. The fault recovery delay of the proposed architecture is evaluated by simulation. In conjunction, the delay and jitter characteristics experienced by different traffic flows during periods of normal operation and faults are measured and compared.

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. 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. The large quantities of QoS-constrained data traffic generated by high-bandwidth services stress the need for survivability and fault protection. 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 facilities located in different layers in the network stack can complement each other to provide a high level of network availability as seen from the perspective of the end users. At the physical layer,

Elasticity	Interactivity	Traffic Class	Service Priority	Application Example
Elastic	**Non-Interactive**	Background (BKG)	4 (Lowest)	File Download, P2P
	Interactive	Interactive (INT)	3	Web browsing, Telnet
Inelastic	**Non-interactive**	Streaming (STR)	2	Video on Demand
	Interactive	Conversational (CON)	1 (Highest)	VoIP, Gaming

Table 1: Traffic class hierarchy proposed by MUSE

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 first mile networks. An optimized centralized protection architecture and protocol operating at the link layer in access rings was introduced previously in [6], which also compared existing protection schemes for their applicability to access rings, and elaborated on the so titled Fast Access Ring Protection Protocol (FARPP), a lightweight Ethernet-based access ring protection protocol. This paper expands on this design by providing for lower recovery latency when supported by physical-layer fault detection, and QoS-aware prioritization of ring traffic during normal operation and faults. It also reports on a comprehensive simulation environment and results that demonstrate the performance of FARPP.

The paper is organized as follows. Section 2 describes the key aspects of the link layer protection architecture. Section 3 elaborates on the QoS-awareness that is incorporated therein. Section 4 focuses on a detailed simulative evaluation of the architecture and presents results. Section 5 concludes with a summary and an outlook on further work.

2 Link Layer Protection for Access Rings

This section describes a full-service link layer protection architecture for access rings. A brief introduction to relevant aspects of the QoS philosophy embodied in MUSE is followed by a discussion of the main network elements comprising the access ring, the breakup of responsibilities therein, and an overview of the protection mechanism in FARPP.

2.1 The MUSE QoS Architecture

The MUSE QoS architecture defines a comprehensive set of principles for providing service guarantees in broadband multi-service access networks. It considers these network domains to be based exclusively on Layer 2 functionality, with the goal being to leverage the available Layer 2 QoS mechanisms to offer and manage IP service guarantees [1]. User traffic is classified into a well-defined set of traffic classes, based on the QoS parameters that characterize it. It is the responsibility of the network to ensure appropriate service prioritization to traffic, based on the class that it belongs to. Table 1 summarizes the traffic classes proposed by MUSE. This proposal draws from recommendations made by the ITU and 3GPP standardization bodies. The listed traffic classes are characterized by two primary differentiators, Elasticity and Interactivity. Elasticity of a traffic flow is a measure of its tolerance to modifications to its original shape. Interactivity, on the other hand, reflects the delay constrained nature of a traffic flow. The four classes resulting from combinations of these differentiators provide a flexible yet relatively simple scheme that can be implemented at the link layer of access rings.

2.2 Functional Ring Elements

A representative access ring as seen at the link layer is illustrated in Figure 1. It consists of a HUB connected by bidirectional links to a series of ONUs in a daisy-chain interconnection scheme. The administrative structure of such access rings in the MUSE aggregation network lends itself to an uneven breakup of responsibilities between the two key network elements therein, the HUB and the ONU, introduced in the following sections.

2.2.1 The HUB

The HUB of the access ring, as depicted in the figure, is a trusted, 'intelligent' network element in the direct control of the network operator, and is responsible for important network management tasks. It is connected to the access ring via two ring ports, and to a metro or wide area network via another port. The HUB performs least-cost forwarding of incoming and outgoing data, aiming to balance traffic load over the links in the ring and maintain an even traffic distribution during normal operation. The selection of the least-cost path to reach an ONU is performed by calculating and comparing the cumulative costs of the two disjoint paths to it from the HUB. The cumulative cost is an aggregate of the costs of the individual links along a path, assigned on a per-link basis by the network operator.

The HUB maintains an updated view of the ring topology at all times. During normal operation i.e., when the ring is complete, the HUB uses its topology knowledge to optimize data routing in the ring. In the event of a ring fault, the HUB suitably routes traffic around the fault, so as to maintain connectivity to all available ONUs. When the fault is eventually repaired, it returns the ring back to its normal operational mode.

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 the ONU and acknowledges its registration with a JOIN ACK message. The ONU's ID is then stored in an internal lookup table located at the HUB. The ONU's entry in the table also stores the ring port chosen for it, the cumulative cost of the path chosen for the ONU, and the regularly updated status of the links connecting it to the ring. Thereafter, 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 sends out HEALTH frames on both its ring ports, the subsequent reception of which ensures 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 ring port on which they are still accessible, ensuring that an alternate path is quickly restored to minimize service disruption.

2.2.2 The ONU

As a consequence of concentrating administrative complexity in the HUB, the ONU can be designed to be a simple, low-cost device deployed close to the end user, saddled with only a 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 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. ONUs are controlled by the HUB

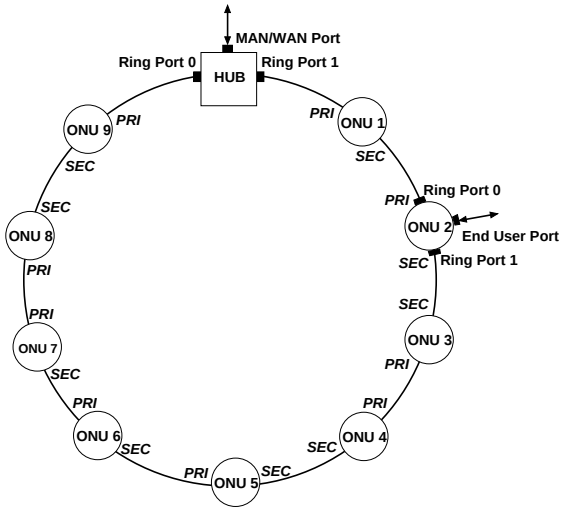

Figure 1: The link layer access ring

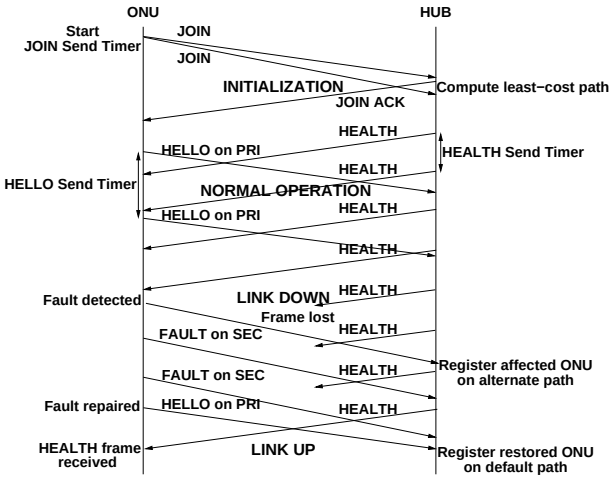

Figure 2: A sample FARPP timing diagram

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, indicated for each ONU in Figure 1. 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 sends out a JOIN ACK message on the least cost path, confirming the ONU's registration and instructing it to set its PRIMARY port to the one providing the least cost path. During normal operation of the ring, the ONU receives and forwards the HEALTH messages it receives on its ports. 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.

2.3 Event-Based Fault Detection and Recovery

The occurrence of faults in the ring is detected and acted upon by ONUs in the ring by one of two complementary mechanisms, one depending on hardware-based detection, and the other based on link layer timers. If the physical layer of the ring provides event-based notification of faults, ONUs neighboring the fault react immediately by sending out FAULT frames on their functioning ports, and switching to their SECONDARY ports if their PRIMARY ports have been affected. Downstream ONUs that receive and forward these FAULT frames to the HUB can then reconfigure themselves similarly, if their least cost paths are affected by the fault. Affected ONUs send out HELLO frames to the HUB on their SECONDARY ports. When the HUB eventually receives these FAULT and HELLO frames, it updates its lookup table to record the new path to the affected ONUs, and localizes the fault within the ring. When full connectivity is eventually restored, the ONUs return to using their PRIMARY ports for communicating with the HUB. A simplified timing diagram of this protection mechanism is shown in Figure 2, which lists the FARPP messages exchanged

between the HUB and an ONU during four phases of operation - Initialization, Normal Operation, Link Down and Link Up.

In the eventuality that the physical layer fails to correctly report faults, the occurrence of the same is detected by affected ONUs via expiration of timers that record the flow of HEALTH messages through the ring. During normal operation, an ONU expects to regularly receive HEALTH frames from the HUB on its PRIMARY port. After a configurable number of consecutive timer expiries that indicate that the expected HEALTH frames have not been received, the ONU assumes a fault on its default path to the HUB. It then reacts as before, by reconfiguring itself and sending out HELLO messages on its SECONDARY port, which eventually trigger a topology update at the HUB. This alternative timer-based mechanism, though slower than the event-based one, provides an independent and software-based method for fault recovery.

2.4 Frame Reflection

Consequent to the topology change resulting from a ring fault, ONUs neighboring the fault use their awareness of the failure to bounce mis-routed frames back to the HUB instead of forwarding them on a failed link. From the time of occurrence of a fault, till when the HUB updates its lookup table, it would forward incoming frames intended for ONUs affected by the fault over an incorrect ring port. To prevent the loss of such mis-routed frames, an ONU receiving a frame not intended for it checks the status of the outgoing link for the frame before forwarding it. If the link is down, it marks the frame as REFLECTED by tagging a field in the frame header, and returns it back to the HUB. The HUB then forwards the frame to the destination ONU, if it is still active, on the alternate path to it. Hence mis-routed frames eventually reach the intended recipient, though delayed and out of sequence. This reflection technique, adapted from a similar mechanism proposed for Metro rings [7], drastically reduces the number of frames lost during the transitional period following a ring fault. Further, this technique can be tailored to only reroute frames belonging to specific traffic classes that guarantee stringent QoS parameters. In such cases, traffic classified as Background (BKG) would be dropped instead of being reflected. Though this would result in a higher perceived loss rate for best effort traffic, this selective reflection would reduce overall load and maintain acceptable QoS for other traffic classes during fault recovery periods.

2.5 Fault Reporting

A methodology for localizing faults in the access ring is important for any subsequent maintenance and repair activity that needs to be undertaken. To this end, a flexible mechanism has been incorporated into the message passing in FARPP, which allows the HUB to infer the exact location of faults in the ring when event-based detection is available at the ONUs. All ONUs report to the HUB the status of both links connecting them to the ring. This information is piggybacked on the HELLO and FAULT messages sent out by the ONUs. The HUB collects the FAULT frames received from ONUs neighboring the fault to identify the failed link or ONU, and notify a suitably designated receiver of the changes in the topology.

2.6 FARPP Message Passing

The Fast Access Ring Protection Protocol (FARPP) implements the message passing scheme described so far, and provides fault resiliency for data traffic. Protocol messages are carried by specialized FARPP frames whose format borrows from that introduced in protection architectures for the Metro domain. Specifically, the extended Ethernet-based frame format described in [8] has been adapted to FARPP. FARPP frames have an additional header that contains protocol-related state information and identifies such frames as administrative traffic, which receive the highest priority of service at every node in the ring to ensure that they propagate topology change information with minimum delay under all circumstances. Timestamp information in the FARPP header, included by the originator of the message, allows the HUB to verify the currency of the topology change information contained therein.

3 QoS-Awareness

The HUB acts as the local QoS manager responsible for facilitating the negotiation of Service Level Agreements (SLA) for end users connected to the access ring managed by it. It configures QoS parameters for policing the SLA compliance of the users attached to the ONUs . During normal operation and faults, the ONU actively multiplexes SLA-compliant incoming and outgoing traffic to and from end users with traffic transiting through it. Incoming data flows are tagged with QoS-related information using a combination of VLAN IDs and priority bits based on the IEEE 802.1Q standard [9]. This QoS information encodes the originator of the flow, along with the MUSE traffic class to which the flow belongs.

Both the HUB and the ONU implement prioritized interface queues for outgoing data frames on their ring ports. A logical interface queue combines multiple physical queues, one for every traffic class that the access ring supports. A strict, non-preemptive Priority Queuing (PQ) policy dictates the order in which the queues get access to the outgoing link. Previous research in MUSE [10] has shown that, when combined with appropriate network dimensioning, admission control and traffic policing, a simpler priority queuing policy performs well when compared with more complex queuing policies like Weighted Fair Queuing (WFQ). Specifically, when dealing with QoS aggregates in the form of a small set of traffic classes being forwarded over a moderately loaded access ring, the lower implementation cost and complexity of priority queuing at the ONUs offsets the hard QoS guarantees that can be honored by WFQ variants.

In conjunction with the priority queuing discipline, the queuing of frames at the interface queues in the ring is customized so as to give preferential treatment to higher priority traffic classes during periods of faults. The priority queues are dimensioned such that, during normal operation of a moderately loaded ring, SLA conformant traffic flows experience only acceptable loss rates. During a fault in the ring, the utilization of some links might approach 100%, as they now would have to carry traffic rerouted around the fault. In such cases, the least priority Background (BKG) traffic is preferentially dropped in favor of traffic belonging to the higher priority classes, which could hence expect to experience tolerable loss rates and jitter even during the fault period.

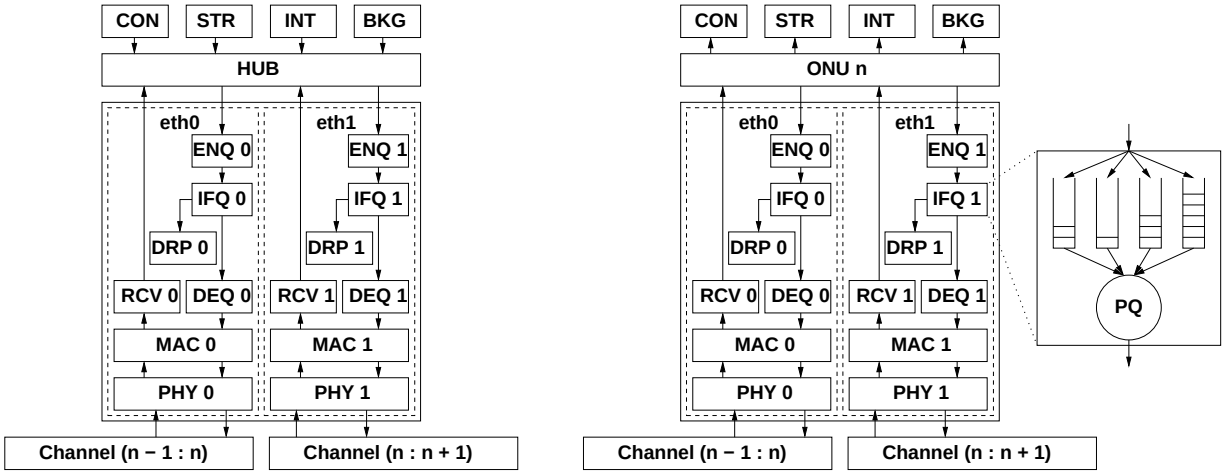

(a) The HUB

(b) The ONU

Figure 3: Block diagrams of the simulation stack for the HUB and the ONU

4 Simulation

4.1 Simulation Setup

In order to model and test the performance of the proposed link layer design, a full-service link layer access ring supporting event-based fault notification from the physical layer has been simulated. This simulation abstracts the key components making up the network stack in the ring nodes, and includes an implementation of the described QoS architecture.

The Network Simulator (NS-2) [11] has been used as the simulation environment. Modeling the behavior of FARPP in NS-2 required the extension of existing modeling support therein. NS-2 provides a large, extensible set of modular components, each of which models an aspect of network functionality in a node. These components can be assembled together to construct more complex network entities.

Figure 3 shows block diagrams depicting the components that simulate the HUB, the ONU, and the physical layer of the access ring. At the bottom, CHANNEL modules model the functionality of a bidirectional channel of specified bandwidth. A composite ring node combines the network interfaces (*eth0* and *eth1*) connecting a generic node to the ring. The network stack in each interface contains modules for modeling the Physical (PHY), Media Access Control (MAC) and QoS-aware queuing layers (IFQ). Attached above each ring node are link layer agents that realize the operation of FARPP at the HUB and the ONUs. Additionally, in order to track the flow of data in the ring, specialized trace modules selectively log the enqueue (ENQ), deque (DEQ), reception (RCV), and drop (DRP) events of frames. These are, when required, complemented by modules that monitor and collect aggregate traffic statistics about the traffic in the ring. At the top of the simulation stack are the traffic sources and sinks that create and absorb traffic belonging to the four MUSE traffic classes - CON, STR, INT and BKG.

4.2 Simulation Scenarios and Results

Different simulation scenarios have been configured and executed using the setup described above, to measure performance parameters of interest. In order to achieve reasonable simulation times, access rings with 100 Mbps link capacity have been simulated. Worst-case conditions for the desired parameters measured are ensured by dividing all of the bandwidth available in the most heavily loaded link amongst the ONUs sharing that link. This per-ONU bandwidth is then further subdivided amongst the four traffic sources, which generate traffic with the typical characteristics like average packet size, burstiness, etc. of the applications that they are modeling, with data rates sufficient to fill up the bandwidth allocated to them. During a simulation run, faults are simulated at desired points in the ring such that the total utilization of the most heavily loaded link reaches or exceeds 100%. In conjunction, the sizes of the queues in the ring have been appropriately dimensioned using estimates from queuing theory, to handle worst-case traffic loads.

4.2.1 Path Restoration Time Scaling

A key parameter measured with regard to evaluating resiliency is the Path Restoration Time (PRT). This parameter, studied in detail in [6], is the end-to-end time required for link layer to react to a ring fault and provide an alternate path to each affected ONU. In order to simulate the worst-case conditions for PRT, all data traffic is routed along one direction - clockwise or anti-clockwise - within a full-service access ring. The bandwidth of the most heavily loaded link is divided in the ratio of 10%, 20%, 40 and 30% amongst the traffic classes, respectively. In such a configuration, Figure 4 depicts the scaling of the PRT for the four traffic classes, measured at the ONU furthest from the HUB, and plotted against an increasing number of ONUs in the ring. As can be seen, the PRT values scale in a linear fashion with increase in ring size. In a fully loaded ring with 100 ONUs, they remain within 20 ms for the inelastic traffic classes (CON and STR) and within 60 ms for the elastic traffic classes (INT and BKG).

4.2.2 Queuing Delay

The QoS-aware queuing mechanism within the access ring ensures prioritized service and restoration of network services during faults. A worst-case scenario for this mechanism has been simulated using a ring with 50 ONUs divided equally into two disjoint groups based on the ring port of the HUB used to communicate with them. Traffic in the ring follows a 5% - 10% - 20% - 30% distribution amongst the traffic classes. Faults are simulated so as to force all the traffic in the ring along one direction, pushing the utilization of the most heavily loaded link above 100% at such times. Figure 5 compares the probability distributions of the queuing delays experienced by the different traffic classes at the most heavily loaded link, both during normal operation and during faults. It can be seen that the delay probabilities for the CON, STR and INT traffic classes remain well within the bounds stipulated by MUSE [12], and do not increase significantly during faults. The plateaus in the distributions for these classes arises from the non-preemptive queuing at the interface queues, because of which higher priority frames wait for any in-progress lower priority frames to complete transmission. Also evident are the much higher queuing delays suffered by the BKG traffic class during faults, frames belonging to which are delayed or dropped from the ring during periods of high load.

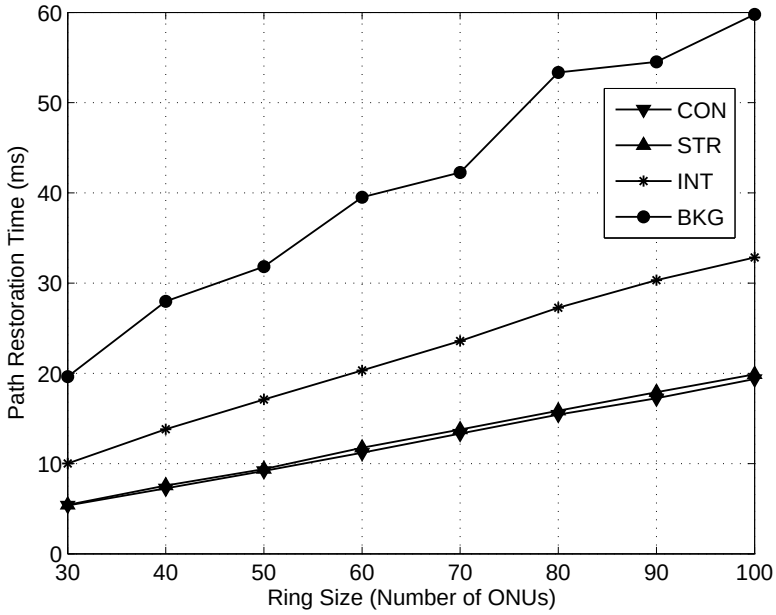

Figure 4: Scaling of Path Restoration Time against number of ONUs in a ring (Each data point has a tolerance within 6% at a 95% level of confidence)

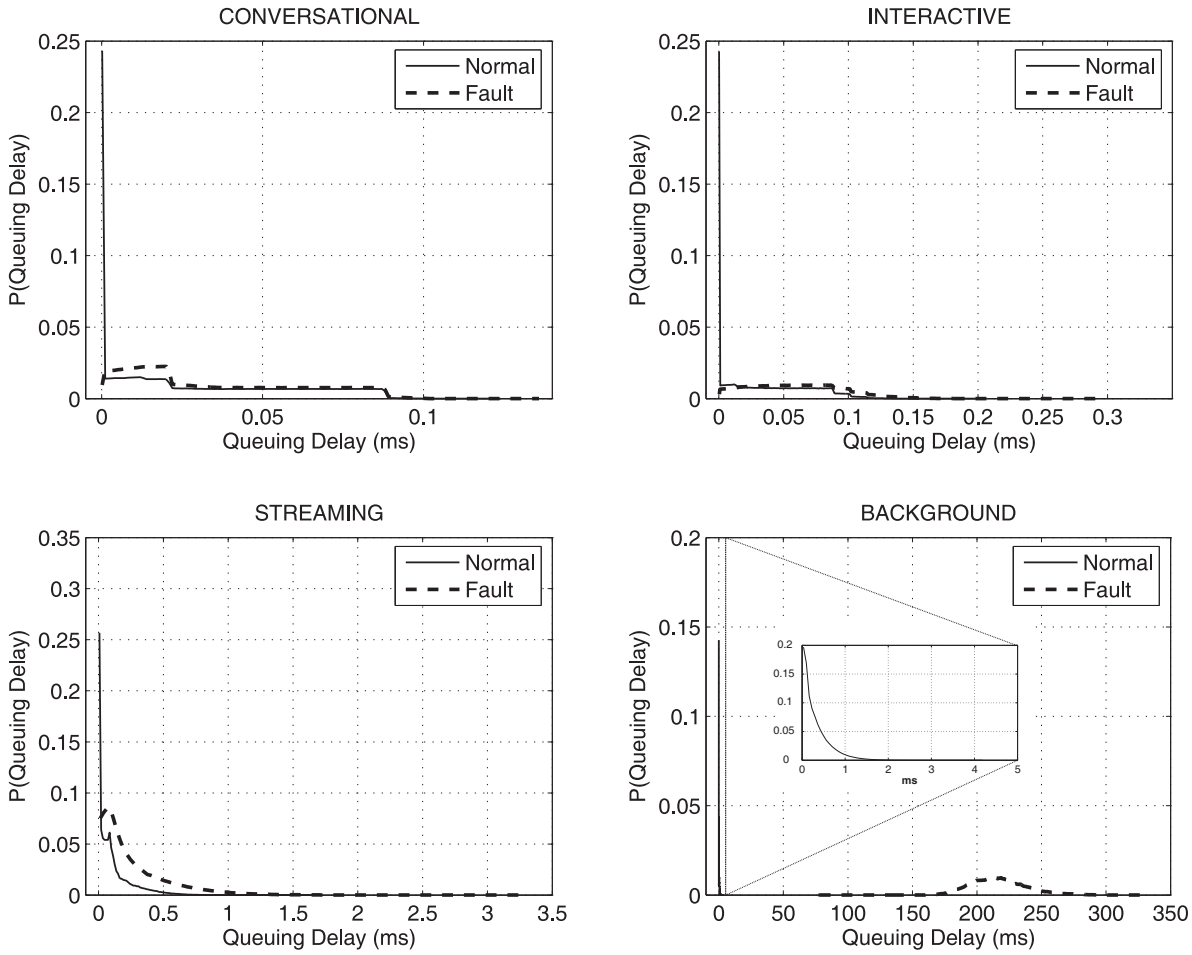

Figure 5: Probability distributions of queuing delays for the four MUSE traffic classes

5 Conclusions

This paper has outlined a QoS-aware link layer protection scheme optimized for access rings. It proposes a centralized network management model to exploit the salient features of such rings, and affords significant simplification over existing schemes. The division of labor among the two main ring elements, the HUB and ONU, have been detailed. By assigning primary ring management responsibilities to the HUB, and role of the relatively simple ONU in detecting and recovering from faults is minimized. Important functionality enhancements to the previously introduced FARPP messaging scheme that implements this architecture have been discussed. In order to verify the performance of FARPP, a detailed simulation of a full-service access ring has been implemented, supporting traffic prioritization based on the MUSE QoS classes. Relevant parameters for evaluating the protection capabilities of the proposed architecture, measured using simulation studies, demonstrate its ability to provide low-latency fault recovery, and QoS-aware management of traffic during normal operation and faults. Further work in the development of FARPP would involve detailed studies of the operational and dimensioning parameters of the protocol, essential for its effective deployment in real access ring environments.

Acknowledgment

This work was supported by the IST-FP6 Integrated Project MUSE (Multi Service Access Everywhere) [1] under contract nr. 026442. The authors would like to thank Prof. Dr. Ulrich Killat at the Hamburg University of Technology for the fruitful discussions and suggestions during the research leading up to this paper.

References

[1] The IST-MUSE Project, www.ist-muse.org

[2] Francois Fredricx: Network architecture and functional specifications for the multi-service access and edge, MUSE Project Public Deliverable D A2.2, January 19, 2005.

[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.: Protected Ring Network for Future Optical Access Domain, Proceedings of the 9th European Conference on Networks and Optical Communications, pp. 106-113, 2004

[6] S. Chennu et al.: Protected Ethernet Rings for Optical Access Networks, Proceedings of the 7th ITG Symposium on Photonic Networks, April 2006, ITG Fachbericht 193, pp 29 - 36

[7] Ziwen Lian et al.: Resilient Ethernet ring for metropolitan area networks, The Ninth International Conference on Communication Systems, 2004, pp. 316 - 320

[8] RFC 3619: Extreme Networks Ethernet Automatic Protection Switching (EAPS) Version 1

[9] IEEE Std 802.1Q-2003 Virtual Bridged Local Area Networks

[10] J. D. Angelopoulos et al.: Supporting Ethernet with QoS in a Local Access Multiplexer, Proceedings of the 10th European Conference on Networks and Optical Communications, pp 51 - 58, 2005

[11] The Network Simulator (NS-2), www.isi.edu/nsnam/ns

[12] A. J. Elizondo Armengol and G. M. Gallizo Rueda: MUSE - Network Requirements for multi-service access, MUSE Project Public Deliverable D A1.2, November 6, 2004