WP233 (v1.0) October 19, 2005

IEEE 802.17, Resilient Packet Ring Networks Enabled by FPGAs

By: Nick Possley

Resilient Packet Ring (RPR) is a transport technology that has been available for several years. The standard IEEE 802.17 specifies the transport. This transport technology provides a wide range of services that are just now becoming very important to the industry for deployment of IP based multimedia services.

Demand for this technology is increasing, and Xilinx has invested in a supported reference design that utilizes the Virtex™-4 family of FPGAs to create a complete RPR MAC solution for equipment vendors. The solution includes hardware and software that provides full compliance to IEEE 802.17.

This paper describes RPR as a network, explains how the MAC operates to provide required network functionality, gives a high-level view of the Virtex-4 implementation of the MAC including device sizing, and covers a few system design use cases.

Background

Legacy communication systems are heavily deployed with ring topologies. DS1 and DS3s are often deployed in rings with ring switches. More recently, SONET/SDH is based on rings including the Bidirectional Line Switched Ring (BLSR) and the Uni-Directional Path Switched Ring (UPSR), the most commonly used ring architectures. These SONET/SDH rings already have fiber in place. Also, these ring architectures provide 50 ms protection switching in network failure scenarios in most cases. As a result, a certain level of network availability is achieved and expected.

One problem with SONET/SDH is that bandwidth is wasted in order to provide protection, and because it is channelized, no statistical multiplexing gain can be achieved between channels. This leads to network inefficiencies because unused bandwidth on one channel cannot be utilized by another channel.

Resilient Packet Ring (RPR) is a transport technology specified by IEEE specification 802.17. This technology provides the best of both worlds. It provides 50 ms protection switching for high network availability while having a packet-based transport that uses statistical multiplexing gain to better utilize all available bandwidth, including protection bandwidth. Further, RPR provides several levels of Quality of Service (QoS) guarantees, including QoS sufficient to support any type of TDM service transported over packets. Packets have become, by far, the most dominant traffic; RPR is therefore the most efficient transport technology for both packet and TDM traffic services going forward.

In addition, RPR is a ring topology, so it can easily be adapted to existing fiber plants that are deploying SONET/SDH. Also, service providers' familiarity with ring networks makes RPR a good solution for transitioning from a TDM-based network to a packet-based network.

Applications

RPR can be used in a variety of ways to distribute a variety of services. The most recent "buzz" about "triple play" services touts RPR as a perfect match for aggregation of service traffic in metropolitan areas.

Triple Play services include Voice, Video, and Data. There are nuances to each of these services. Voice services can include real-time voice traffic, e.g., phone call, or non real-time, streaming music distribution. Video services traditionally include streaming video distribution, e.g., Video On Demand. Video services can also include video teleconferencing, which is a real-time service. These differences in services require different levels of QoS from the transport mechanism in the network.

RPR works well to provide all of these services due to its defined service primitives in IEEE 802.17. These service primitives are:

- Class A Service - Provides an allocated, guaranteed data rate with low end-to-end delay and jitter bound. This class has precedence over all other classes.
- Class B Service - Provides an allocated, guaranteed data rate with bounded end-to-end delay and jitter for the allocated rate. This class also provides access to unallocated bandwidth that has no guaranteed data rate or bounded delay and jitter. The primitives referring to the different used bandwidths are Class B Committed Information Rate (classB-CIR) and Class B Excess Information Rate (classB-EIR). Class B takes precedence over class C.
- Class C Service - Provides a best effort delivery with no guaranteed data rate and no bounded delay or jitter.

Given this set of services, it is clear that all real-time services for Voice or Video will use Class A in order to maintain the lowest latency and guaranteed bandwidth. Streaming video or streaming voice distribution would use Class B. Data services would use

class C. Of course, these pairings all depend on what a service provider is offering. Some service providers might provide Class B data services for those customers willing to pay more. For consumer services, the breakdown described above is the most accurate because it provides the most efficient use of the available network bandwidth.

RPR can support multiple rates; 1G, 2.5G, and 10G are the standard rates. This allows for scaling of the technology to accommodate a growing network. 1G transport uses a Gigabit Ethernet physical layer. 2.5G transport uses a SONET OC48 physical layer. 10G uses a SONET OC-192 or 10GE physical layer.

Metro Networks

Figure 1 is a network diagram of an RPR metro network.

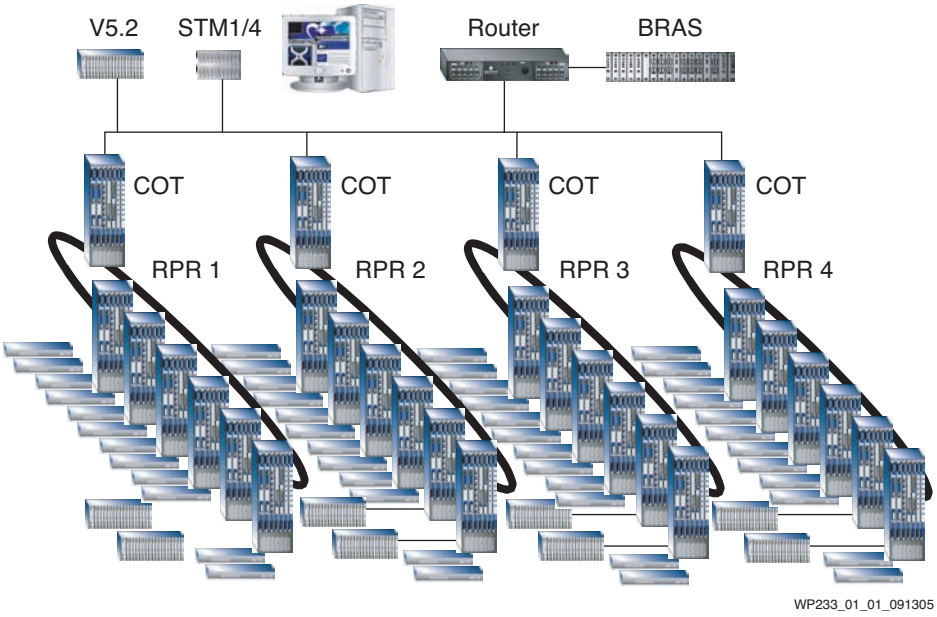

Figure 1: **RPR Metro Network**

This is an example network that illustrates how RPR is deployed and the types of equipment used for the deployment.

xPON and VDSL are the upcoming delivery mechanisms for triple play services. This places RPR in OLT equipment for PON and in DSLAMs for VDSL. Of course, many equipment vendors are producing one platform to do both, which is typically dubbed an MSPP (Multiservice Provisioning Platform).

In addition, there is a Central Office Terminal that must connect the metro network to the core or regional network. This terminal must also support RPR. The types of equipment in this category vary from edge routers to core switches.

The metro space is sometimes required to provide a mixture of TDM and packet services. RPR works well in these scenarios; it can easily be mapped into a VCAT tunnel inside of a SONET stream, which is a subset of the available bandwidth. TDM services can be mapped onto the extra bandwidth. This is probably most desirable when networks are being upgraded, and legacy services must be maintained.

TDM services can also be mapped onto RPR using a pseudo-wire-type approach similar to PWE-3. This allows for all services to be transported across an RPR ring.

Wireless Backhaul

Another interesting application for RPR is transport for wireless backhaul networks. RPR's service classes work well for guaranteeing the low latency required for real-time voice services. Figure 2 shows a potential network based on the UMTS network architecture and RPR.

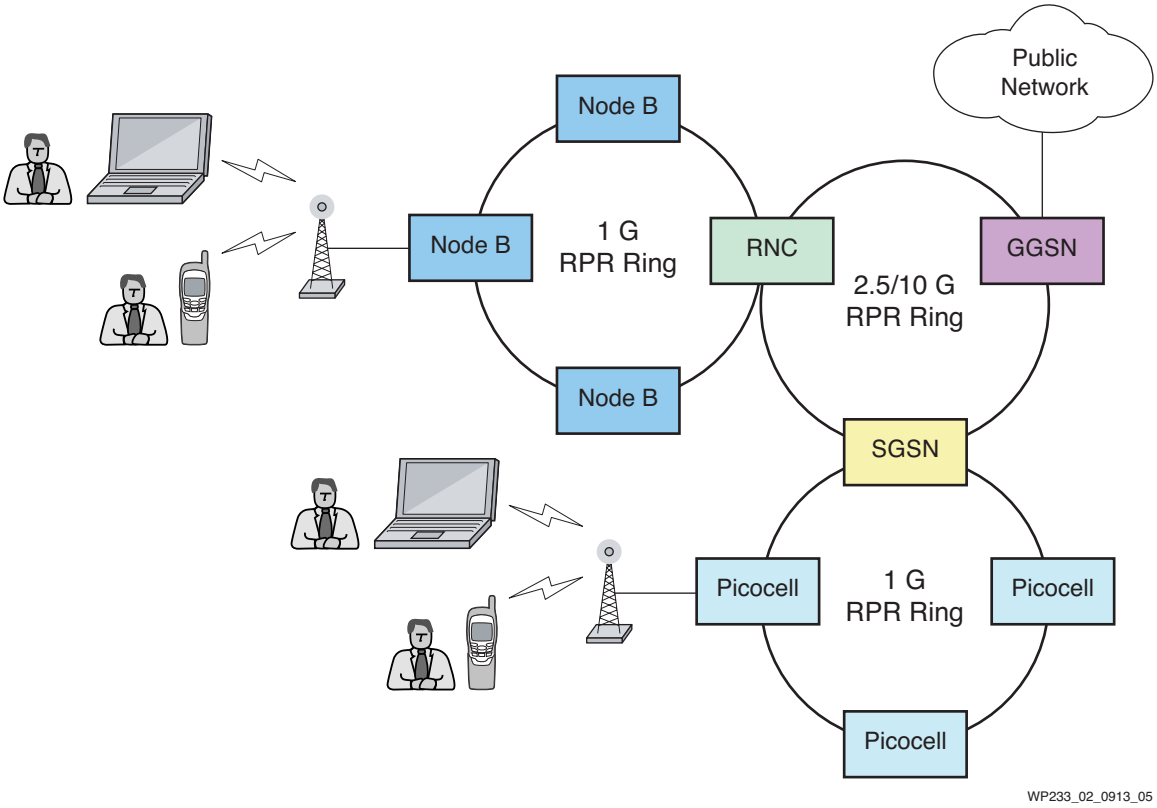

Figure 2: **RPR Wireless Backhaul Network**

RPR can remove much of the complexity in the Iu'x' interfaces because it incorporates the QoS transport service into the L2 transport instead of inside ATM PVCs that are bandwidth inefficient.

A common application scenario is MPLS over RPR. Using static LSP, MPLS can provide the secondary address for application ports that are transported over RPR.

Overall, these applications show that RPR has a wide market space and can play a key role in services moving forward.

One thing to note is the lack of RPR technology available in order to enable the deployment of the standard. FPGAs are an ideal device to support RPR if an efficient implementation that doesn't make the technology cost prohibitive can be achieved.

RPR MAC

The RPR MAC is the central component that enables RPR transport. This MAC governs all access to the ring. The MAC fits into the OSI reference model as shown in Figure 3.

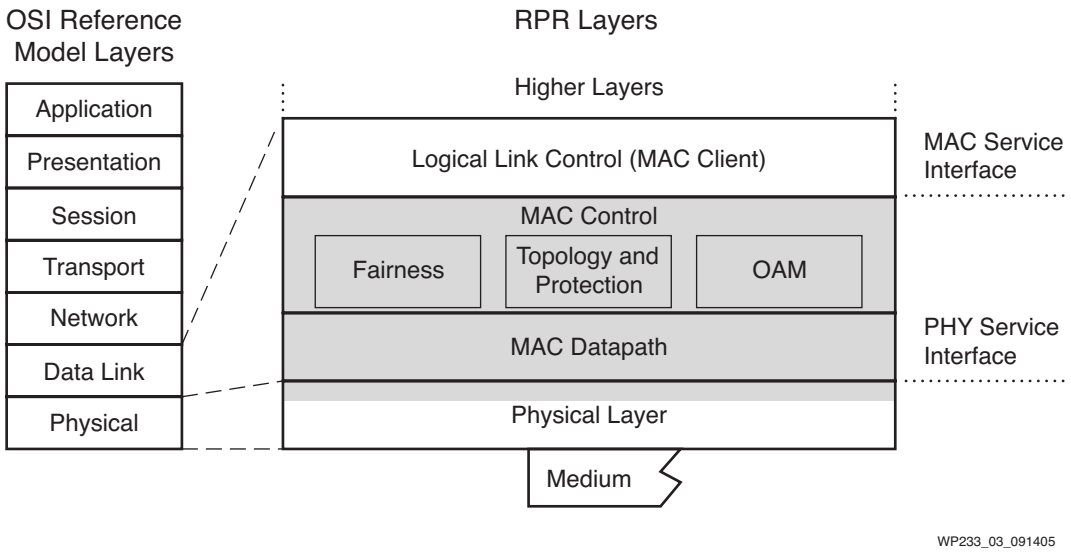

Figure 3: **RPR Relationship to the OSI Reference Model**

Overview

Functionally, the MAC has the following aspects:

- Network Protection
- Data Formatting and Transport
- QoS
- OAMP

Network Protection requires the MAC to track the topology of the network and respond to any changes via steering and wrapping of datapaths. Figure 4 shows the architecture for this functionality.

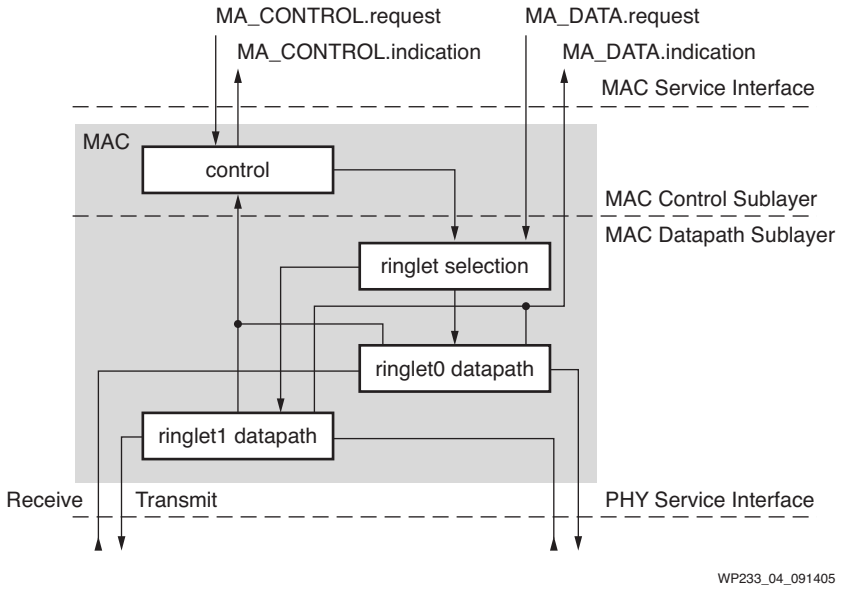

Figure 4: **RPR Protection Path Architecture**

Data formatting and transport requires the MAC to service client requests for transport by framing the data into the RPR packet format and resolving the hop count in the transmit direction. In the RX direction, the MAC is required to receive and pass along any packets not meant for the local station. For those packets addressed to the local station, the MAC parses, receives the data, and passes it to the client interface. Figure 5 shows the functionality for single and dual queue MAC architectures.

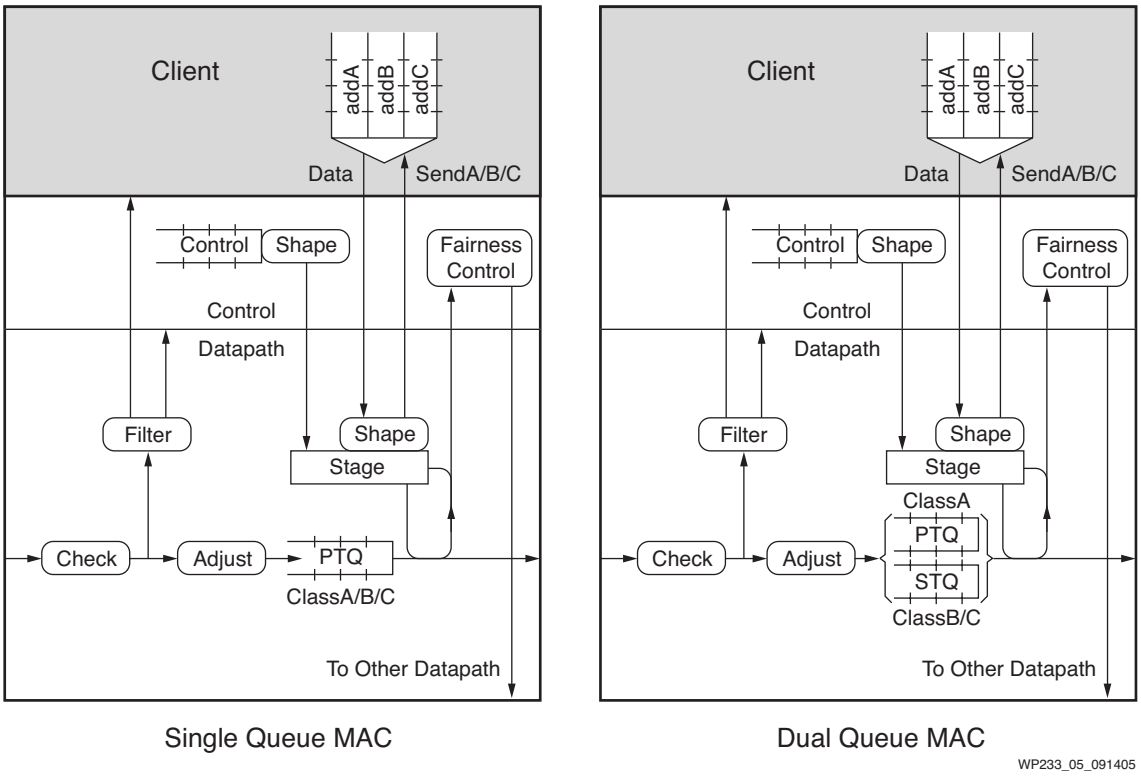

Figure 5: **RPR Datapath Architecture**

QoS requires the MAC to monitor the topology of the network and the resource utilization in the network and to perform traffic policing/shaping for the various levels of quality based on available resources. Dual queue versus single queue simply means whether or not a secondary set of queues is in the MAC. A dual queue MAC supports class B/C traffic classes; whereas, a single queue only supports class A (Figure 5). QoS requires a fairness algorithm to determine the utilization of network resources and to allocate the correct amount of bandwidth to the service classes. Figure 6 shows how the fairness algorithm allocates bandwidth.

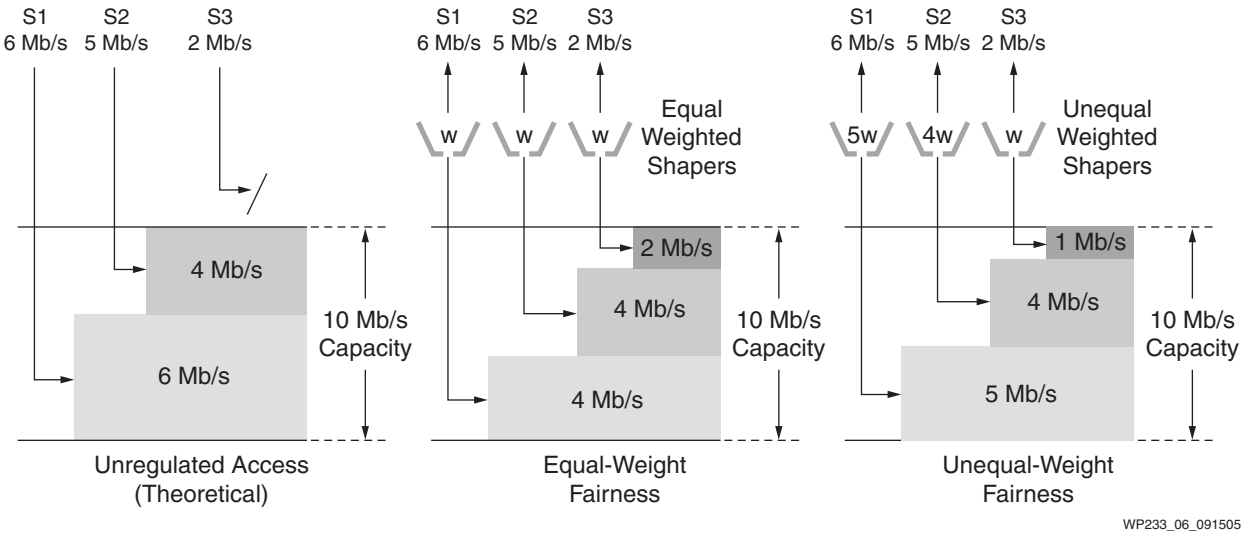

Figure 6: **RPR Fairness Example**

OAMP requires the MAC to gather statistics for each of the services and report the data back to the user via a standardized interface, e.g., MIB. In addition, the MAC is required to provide provisioning of all relevant parameters for each station.

For more details, please refer to the IEEE 802.17 latest specification.

Implementation

From the previous sections, it is obvious that the MAC is comprised of a complex set of functions. Given the complexity, it would make sense for this to be a large implementation, but Xilinx has not only implemented a most efficient and high performance design, the design is also half the price of any competing design. This implementation has been made possible by the Virtex-4 family of FPGAs. Through the use of the DSP48 slices and the speed improvements, the Virtex-4 FPGA provides unique ways to converge more functionality into less logic.

The implementation has a split between HW and SW in which all real time functions are implemented in hardware and everything else in software. This has resulted in a design that is only 11000 slices and 80 block RAMs for 1G and 2.5G. It will be slightly larger for 10G only to accommodate a wider memory interface. Table 1 shows the sizing.

Table 1: **RPR MAC FPGA Sizing**

Logic Block	Approx. Slice Count for Both Rings	Approx. Block RAMs	DCMs	DSP Slices
PHY Interface (SPI, GMII, XGMII)	250	4	2	–
RX Datapath	2000	20	1–2	–
Transmit Datapath	1200	18	–	–
Client Datapath	1000	8	1	–
Control Path	1000	6	1	–

Table 1: **RPR MAC FPGA Sizing *(Continued)***

Logic Block	Approx. Slice Count for Both Rings	Approx. Block RAMs	DCMs	DSP Slices
Fairness Block	3000	1	–	40
Topology Block	1800	6	–	–
Memory Management	600	5–20	1–2	–
Others (Self Test, Timer, etc.)	200	1	–	–
Total	11,000	72–87	6–8	40

The design requires at least 15,000 slices (actual logic may need 10,500 to 12,000 slices, and 20% head room must be kept for client interface customizations), 8 DCMs, RocketIO™ interfaces (either 3.125 Gb/s for XAUI or 10 Gb/s for XSBI), and 40–45 DSP slices are required for the design. Regarding IO calculations, the 10 Gb/s interface needs max IO signals, 200+ signals for memory interface, and 100+ signals for system. If the PHY interfaces are non-serial, IO requirements are even higher. If RocketIO is used as the interface to the PHY, a 350-I/O pin device is estimated for the design.

The DSP48 slices available in all Virtex-4 family devices allow the fairness block to be very small. This is a major contribution to the compact nature of the design.

In addition, the solution is customizable to a customer's exact application, without charging for any extra silicon. Figure 7, Figure 8, and Figure 9 show the block diagram from different aspects.

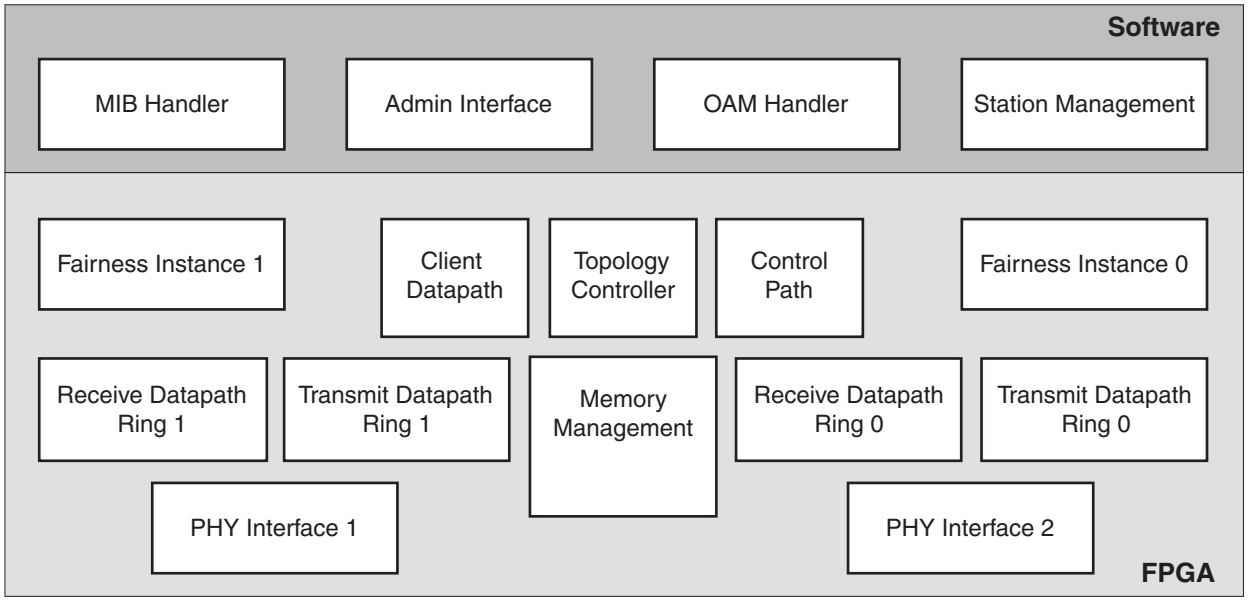

Figure 7: **RPR MAC FPGA HW/SW Split**

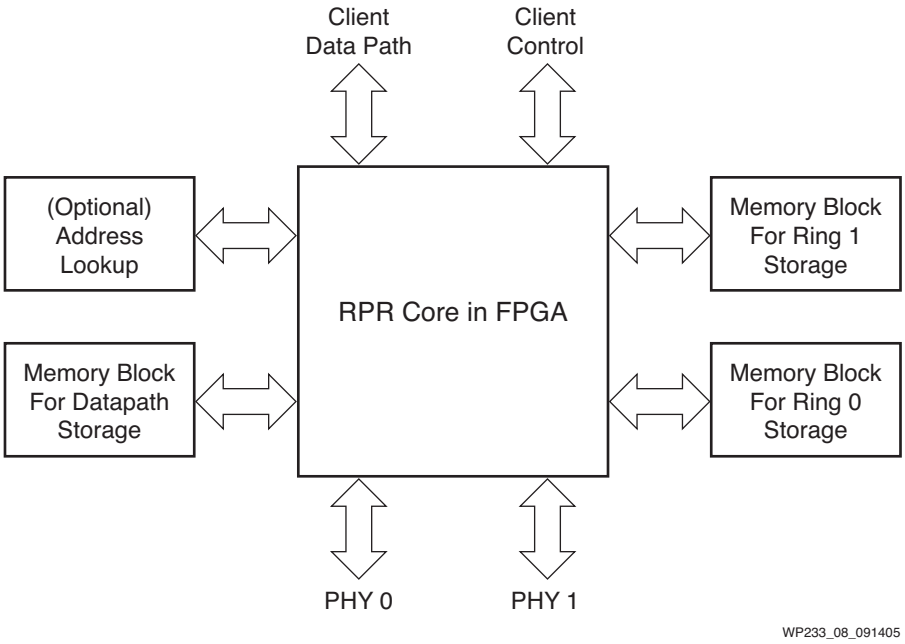

Figure 8: **RPR MAC FPGA Black Box HW View**

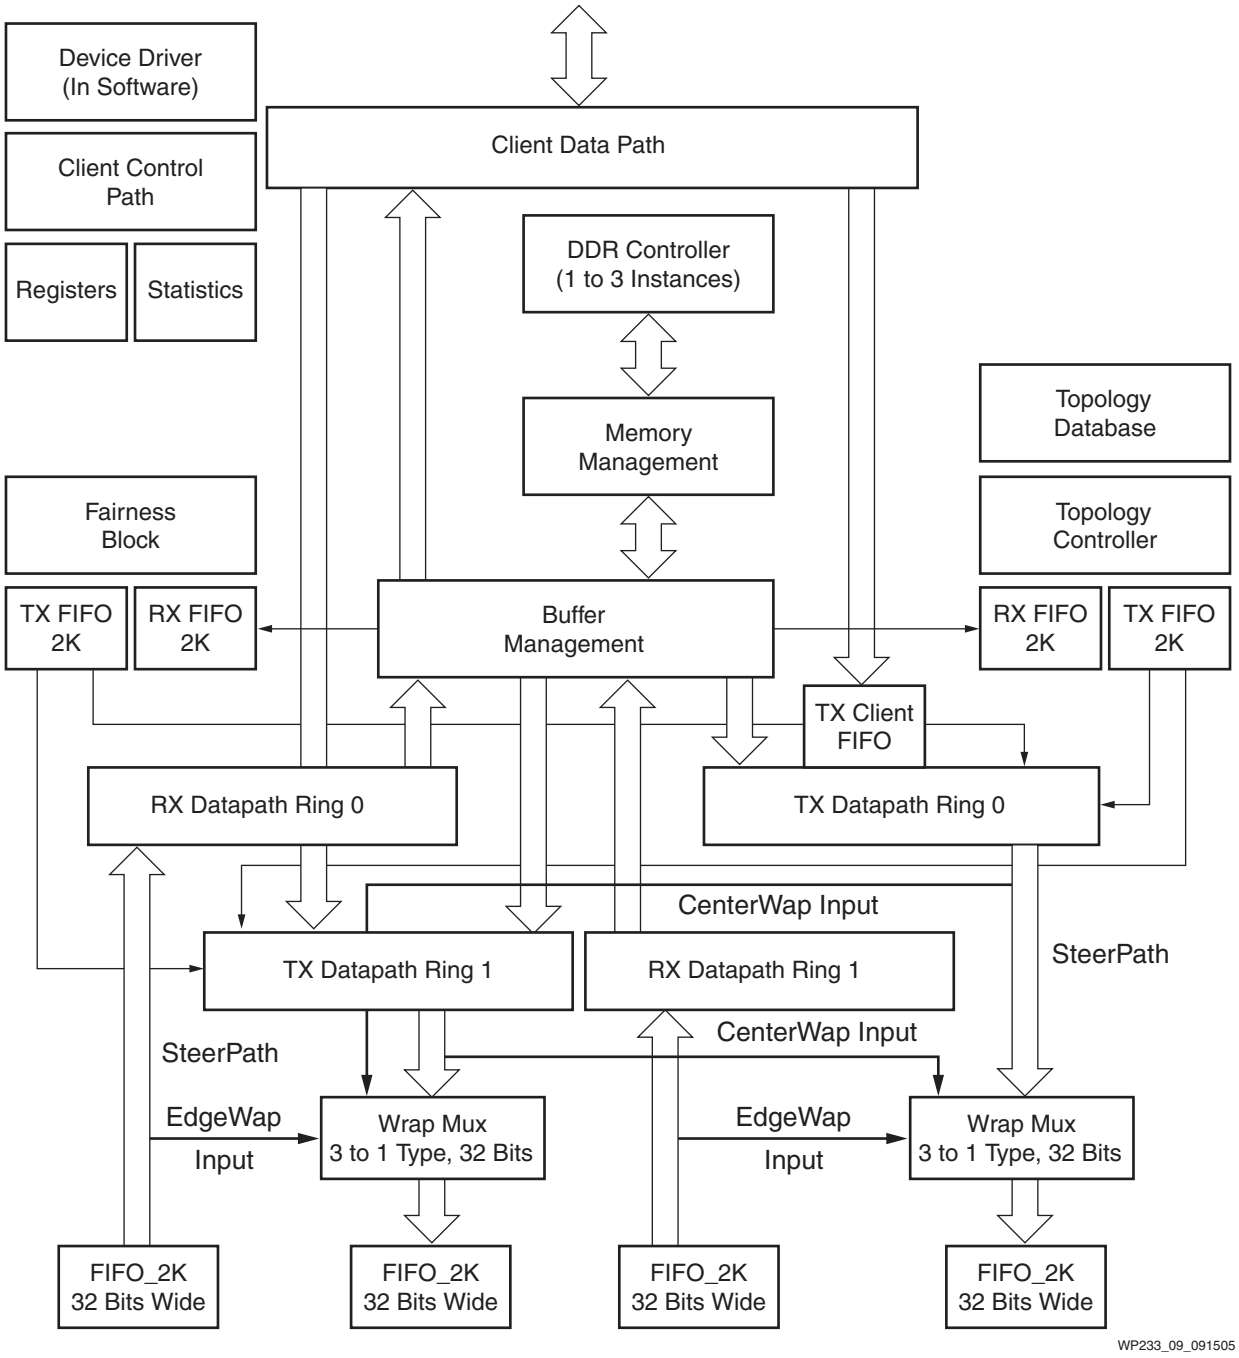

Figure 9: **RPR MAC FPGA Detailed Internal Architecture**

There are several different ways to view the core, and they each provide varying levels of detail.

Table 2 shows the customer implementation options.

Table 2: **RPR MAC FPGA Implementation Options**

Possibilities	1G Ethernet	10G Ethernet	OC48	OC192
PHY 0 and PHY 1	RocketIO, parallel (GMII)	XAUI or XSBI (RocketIO)	SPI 3.0	SPI 4.2
Memory	1 block for Ring 0, Ring 1, and datapath	1 block for Ring 0, 1block for Ring 1, and 1 block for the datapath (3 total)	1 Block for Rings and 1 for data Path	1 block for Ring 0, 1 block for Ring 1, and 1 block for the datapath (3 total)
Memory Bus Width	32	64	32	64
Client Datapath	LVDS, parallel, GbE	RocketIO, parallel (64 bits)	SPI 3.0, LVDS	SPI 4.2
Client Control Path	32-bit parallel, embedded CPU, merge with datapath	32-bit parallel, embedded CPU, merge with datapath	32-bit parallel, embedded CPU, merge with datapath	32-bit parallel, embedded CPU, merge with datapath

The Virtex-4 FX family provides the advantage of integrated PHYs for all interfaces except OC192. This allows for further integration of functionality inside the FPGA and, hence, less overall system cost.

For most customer systems, availability of the network is one of the most important factors to consider. RPR provides high availability through its network protection mechanisms. Most equipment vendors also require some equipment protection for higher availability and maintenance. The RPR core provides for equipment protection via a MATE interface and some software/hardware functionality. Figure 10 shows the MATE interface.

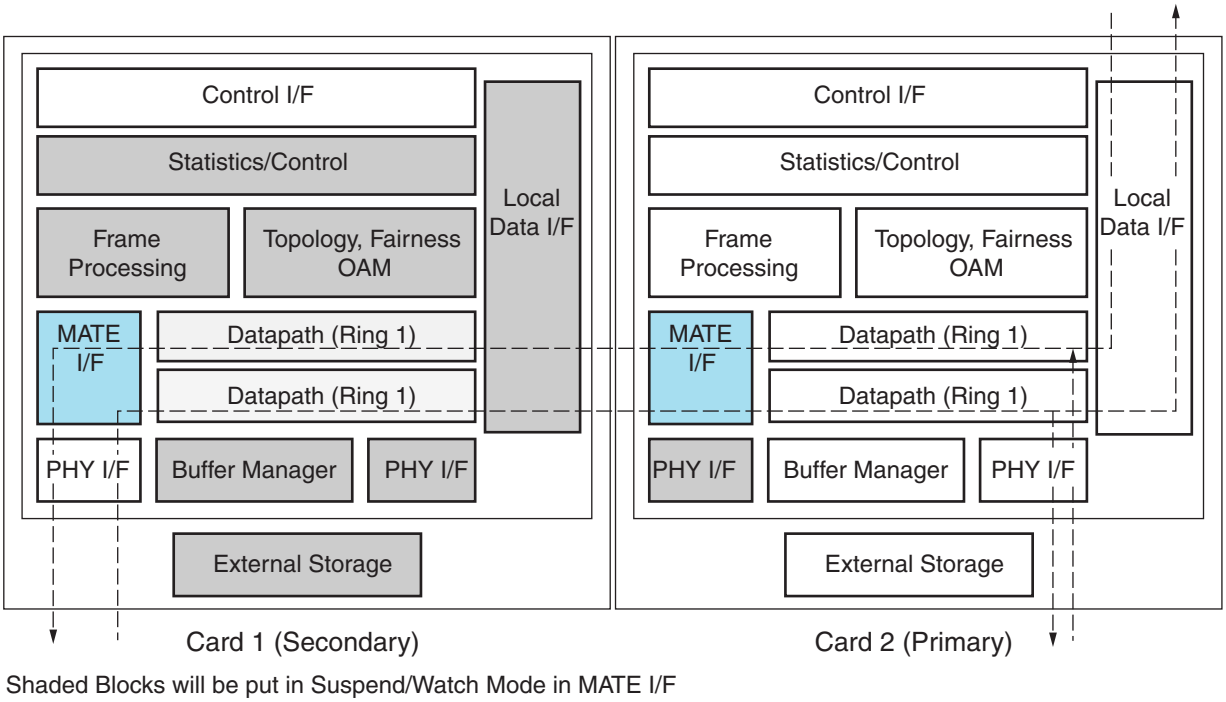

Figure 10: **RPR MAC FPGA Equipment Protection**

In Figure 10, there is a datapath interface and a control path interface between the MACs. The datapath interface allows for protection of PHY failures independent of MAC failures. In order to do this, the control of the MAC needs to remain on one controller. This requires a master/slave state machine between the two controllers. Also, if the whole MAC fails, then the functionality shown on a single card (see Fig. 10) is switched over.

The RPR MAC also has many options for pinouts. Table 3 shows a few scenarios.

Table 3: **RPR MAC FPGA Pinout Examples**

Functional Block	Memory	PHY	Datapath	Control and Other
GbE with RocketIO and RocketIO datapath interface (single card implementation). External CAM.	50	10	4–8	90
GbE parallel interface to the PHY (line side), GbE backplane interface (client side). External CAM.	50	48	24	90
OC48 with SPI3.0 interface to PHY and backplane as RocketIO 3.125G	100	160	4–8	90
10G with XAUI RocketIO interface to PHY, parallel to backplane.	260	18	140	90
OC48 with MATE interface, SPI3.0 PHY interface.	100	80 + 20 for MATE	8 for RocketIO	90

Testing

Of course, compliance testing to the IEEE 802.17 specification is the "real" test. Table 4 is the compliance matrix based on test cases.

The subclause number refers to IEEE 802.17 specification subclauses, and the name refers to Protocol Implementation Conformance Statement (PICS), outlined at the end of every section of the IEEE 802.17 specification.

Table 4: **RPR MAC Compliance Matrix**

Subclause	Name	Details	Support	Test Case Ref
Chapter 6				
6.3	FTS1	Generation of Strict Frames	Yes	RPR-0005
6.3	FTS2	Generation of Relaxed frames	Yes	RPR-0006
6.4.1	SP1	MA_Data Request	Yes	CDT-0002
6.4.2	SP2	MA_Data Indication	Yes	CDR-0001
6.4.3	SP3	MA_Control Request	Yes	Internal.
6.4.4	SP4	MA_Control Indication	Yes	Internal.
Chapter 7				
7.5.3	SM1	Idle Shaper	Yes	CCT-0004
7.5.4	SM2	MAC Control Shaper	Yes	CCT-0001
7.5.5	SM3	Class A Shaper	Yes	CCT-0002
7.5.6	SM4	Class B Shaper	Yes	CCT-0003
7.5.7	SM5	Fairness Eligible Shaper	Yes	TXP-0001, 0002
7.5.8	SM6	Downstream Shaper	Yes	TXP-0001, 0002
7.6.3.5	SM7	Receive Frame Edge	Yes	RXD-0014
	SM8	Receive Check	Yes	RXD-0003
	SM9	Receive Count	Yes	RXD-0023, RXD-0024
	SM10	Receive Strip	Yes	RXD-0027
	SM11	Receive Adjust	Yes	RXD-0030
	SM12	Receive Filter	Yes	RXD-0032
	SM13	Receive Filter Data Count	Yes	RXD-0023 - 27
7.6.3.12	SM14	Receive Filter Control Count	Yes	RXD-0023 - 27
7.6.4	SM15	Wrong Ringlet	Yes	RXD-0013
7.7.1	SM16	Ringlet Selection	Yes	CDT-0002
7.7.4	SM17	Stage Q selection	Yes	CDT-0005
7.7.5	SM18	Data Add Count	Yes	TXP-0004
	SM19	Control Add Count	Yes	TXP-0004
	SM20a	Single Q Transmit	Yes	TXP-0001
	SM20b	Dual Q Transmit	Yes	TXP-0002
	SM21	Transmit Count	Yes	TXP-0004
7.67.10	SM22	Transmit Route	Yes	TXP-0004
7.4	DP1	Transit Delay	Yes	RPR-0013 and 14
7.4.4	PM1	Steering	Yes	TXP-0003
7.4.4	PM2	Wrapping	Yes	TXP-0003
7.4.4	PM2a	Center Wrapping	Yes	TXP-0003

Table 4: **RPR MAC Compliance Matrix *(Continued)***

Subclause	Name	Details	Support	Test Case Ref
7.4.4	PM2b	Edge Wrapping	Yes	TXP-0003
7.4.3	PM3	Pass Through	Yes	TXP-0004
7.4.4	W1	Wrapped Station	Yes	TXP-0003
7.6.2	W2	Frame Fields for Wrapping	Yes	TXP-0003
7.6.2	W3	Frame Fields for Unwrap	Yes	TXP-0003
7.5	SH1	Crossing Below Low Limit	Yes	CCT-0001 to 3 and TXP-0001, 2
7.5	SH2	High Limit Value	Yes	CCT-0001 to 3 and TXP-0001, 2
7.5.1	SH3	Decrement Frequency	Yes	CCT-0001 to 3 and TXP-0001, 2
7.5.1	SH4	Accuracy	Yes	CCT-0001 to 3 and TXP-0001, 2
7.6.1	CC1	Steering Purge	Yes	TXP-0003
7.6.1.1	CC2	Queue Purge	Yes	TXP-0003
7.6.1	CC3	Wrapping Purge	Yes	TXP-0003
7.6.2	CC4	Wrapping Field Requirements	Yes	TXP-0003
7.6.2	CC5	Unwrapping Field Requirements	Yes	TXP-0003
7.7.3	TTL1	TTL Setting	Yes	CDT-0002
7.7.3	TTL2	TTL Base Setting	Yes	CDT-0002
7.7.7.1	FO1	Single Queue	Yes	TXP-0001
	FO2	Dual Q PTQ – FIFO Order	Yes	TXP-0002
	FO3	STQ FIFO Ordering	Yes	RPR-0008
	FO4	Dual Queue Cross Ordering	Yes	RPR-0008
Chapter 8				
8.3		1Gb/s PHY Reconciliation Sublayer	Yes	GMII-0001 to 7
		10Gb/s PHY Reconciliation Sublayer	Yes	XGMII-0001 to 7
8.4		SONET/SDH Reconciliation Sublayer	Yes	SPI3-0001 to 7
		GFP Reconciliation Sublayer	NO	
8.2.3.3		PHY_LINK_STATUS Indication	Yes	
8.3.1.1		PRS-1 Provides GMII Function	Yes	GMII-0001 to 7
		PRS-10 Functional Behaviour	Yes	XGMII-0001-7
8.4		SPI-3/4	Yes (SPI3.0)	SPI3-0001-7
		Matched Ringlet Rates	Yes	
		GFP Framing Compliance	No	
		HDLC Framing	No	
		HDLC Link Compliance	No	
		LAPS Framing	No	
Chapter 9: Frame Format				
9.2		Basic Data frames	Yes	CDT-0003
9.2		Extended Data Frame	Yes	CDT-0003
9.3		Control Frame	Yes	TPT-0005
9.4		Fairness Frame	Yes	RPR-0015, FCB-0010
9.5		Idle Frame	Yes	CCT-0004
9.6.2,		Fe Reserved	Yes	RPR-0015
9.6.6,		Parity Reserved	Yes	RPR-0015

Table 4: **RPR MAC Compliance Matrix *(Continued)***

Subclause	Name	Details	Support	Test Case Ref
9.7.5,		Res Reserved	Yes	RPR-0015
9.5.2.4		Idle Payload	Yes	CCT-0004
Chapter 10				
10.4.1		Rate Statistics	Yes	FCB-0002
10.4.1		Rate Statistics Maintenance Delay	Yes	FCB-0002
10.4.1		Policing Indicators	Yes	FCB-0002
10.4.2.4		Aging	Yes	FCB-0003
10.4.2.4		Low Pass Filtering	Yes	FCB-0003
10.4		Aging Interval	Yes	FCB-0003
10.4.3		Aggressive Rate Computation	Yes	FCB-0004
10.4.4		Conservative Rate Computation	Yes	FCB-0005
10.4.5		Rate Advertisement	Yes	FCB-0007
10.4.6		Rate Report	Yes	FCB-0007
10.4.7		Active Weights Computation	Yes	FCB-0008
10.4.8		Single Choke Fairness Frames	Yes	FCB-0009
10.4.8		Multi ChokeFairness Frames	Yes	FCB-0009
10.4.9		FDD Frame Transmit	Yes	FCB-0010
10.4.10		FRTT Computation	Yes	FCB-0011
10.4.3		Shaper-Based Admission	Yes	FCB-0003
10.4.3		Rate-Based Admission	Yes	FCB-0003
Chapter 11				
11.6.2	SM1	Receive Monitor State mc	Yes	TPB_0001, TPB_0004
11.6.3	SM2	Topology Control	Yes	TPB_0005 to TPB_0009
11.6.4		ParseTPframe	Yes	TPB_0010 to TPB_0017
11.6.5		ProtectionUpdate	Yes	TPB_0018 to TPB_0020
11.6.6		Topology Validation	Yes	TPV_0001 to TPV_0005
11.6.7		TransmitTPframe	Yes	TPT_0001 to TPT_0004
11.6.8		TransmitTCframe	Yes	TPT_0005 to TPT_0007
11.6.9		ReceiveTPframe	Yes	TPR_0001 to TPR_0003
11.6.10		ReceiveTCframe	Yes	TPR_0006
11.6.11		TransmitATDframe	Yes	TPT_0013 to TPT-0014
11.6.12		ReceiveATDframe	Yes	TPR_0007 to TPR_0008
11.6.13		Secondary Update	Yes	TPR_0009 to TPR_0013
11.6.14		TimingLRTTframe	Yes	TPR_0014, TPT_0009 to TPT_0012
11.3.1	375	TP Frame	Yes	TPR-0001- TPR-0003
11.3.2		TC Frame	Yes	TPR-0006
11.3.3		LRTT Request Frame	Yes	TPR_0014 to TPR-0017 and TPT-0009 to TPT-0012
11.3.4		LRTT Response Frame	Yes	TPR_0014 to TPR-0017 and TPT-0009 to TPT-0012
11.3.4		Tail Latency IN Field Setting	Yes	TPR_0014 to TPR-0017 and TPT-0009 to TPT-0012

Table 4: **RPR MAC Compliance Matrix *(Continued)***

Subclause	Name	Details	Support	Test Case Ref
11.3.4		Tail Latency OUT Field Setting	Yes	TPR_0014 to TPR-0017 and TPT-0009 to TPT-0012
11.3.5		ATD Frame	Yes	TPR_0008
11.3.5		ATT Advertising	Yes	TPR-0008
11.3.5		ATT Advertising, Default	Yes	TPR-0008
11.3.5		ATT Not Included	Yes	To be updated.
11.3.5		Type Dependent Length	Yes	TPR-0008
11.4.1	376	Weight ATT	Yes	TPR-0008
11.4.2		Station Bandwidth ATT	Yes	TPR-0008
11.4.2		Reserved Bandwidth Within Station Bandwidth ATT	Yes	TPR-0008
11.4.2		Excess Reserved Rate Defect	Yes	TPR-0008
11.4.3		Station Settings ATT	Yes	A software function.
11.4.4		Station Name ATT	Yes	A software function.
11.4.5		Station Management ATT	Yes	To be updated. It is a device driver function.
11.4.5		Managed Entity	Yes	A software function.
11.4.6		Station Interface Index ATT	Yes	TPR-0008
11.4.7		Secondary MAC ATT	Yes	TPR-0008
11.4.7		First 2ndary MAC Q Address Setting	Yes	TPR-0008
11.4.7		2nd 2ndary MAC Address Setting	Yes	TPR-0008
11.4.8		Organization Specific ATT	Yes	
11.1.6.3	PP1-7	Protection for Relaxed Data Frames	Yes	TPB-0008
11.1.6.3		Protection for Strict Frames	Yes	TPB-0008, RXD-0027
11.1.6.1		Steering Protection	Yes	TPB-0008, 0009
11.1.6.2		Wrapping Protection	Yes	TPB-0008, 0009
11.2.3		Revertive Operation	Yes	TPB-0008, 0009
11.2.3		Non-Revertive Operation	TBD	TPB-0008, 0009
11.6.5.2		Failure of Transit Path	Yes	TPB-0008, 0009
11.6.8		Topology Discovery in Maximum Size Ring	Yes	TPB-0010 to TPB-0014
11.6.6.1		Persistent Topology Inconsistency	Yes	TPB-0010 to TPB-0014
11.2.3		ATD Timer	Yes	TPT_0013
11.6.11		ATD Transmission (ATD2,3)	Yes	TPT_0013
11.6.11		Claiming Ownership of Resources Through ATTs	Yes	TPR_0011, TPR_0018, 19
11.6.11		Handling Race Conditions in ATD	Yes	To be updated.
11.6.12		ATD Reception	Yes	TPR_0007
11.6.12		Illegal ATD Discards	Yes	TPR_0007
11.6.14	LRTT1	LRTT Request Processing		TPT_0009
11.6.14.4		LRTT Response Time	Yes	TPT_0009
11.6.14.4		No Late LRTT Response	Yes	TPT_0016
11.6.14.4		LRTT Re-Measurement	Yes	TPR_0015
Chapter 12				
12.3.3	398	Organization Specific Frames	Yes	OAM-0006

Table 4: **RPR MAC Compliance Matrix *(Continued)***

Subclause	Name	Details	Support	Test Case Ref
12.6		Performance Monitoring	Yes	OAM-0010 - 0014
12.5.1		OAM Transmit State m/c	Yes	OAM-0003
		OAM Frame Receive stmc	Yes	OAM- 0004
		ECHO Frame Formats	Yes	OAM_0001, 0002
		Flush Frame Formats	Yes	OAM-0005
		Organization Specific Frame Format	Yes	OAM-0006
12.4.1		Echo Request/Response Formats	Yes	OAM-0001, 0002
12.4.1		Echo Request/Response Formats	Yes	OAM-0001, 0002
		Echo Flush Formats	Yes	OAM-0005
		Echo Flush Formats	Yes	OAM-0005
		Organization Specific Formats	Yes	OAM-0006
		Organization Specific Formats	Yes	OAM-0006
12.6.1.2		Scff Errors	Yes	OAM-0010
12.6.1.3		Errored Sec	Yes	OAM-0011
12.6.1.4		SES	Yes	OAM-0012
12.6.1.1		Unavailable Time	Yes	OAM-0013
12.6.1.2,3,4		PM5	Yes	OAM-0014
Chapter 13				
13.1	410	SME	Yes	
13.2		LME Primitives	Yes	LME_0001
13.2		LME Initialization Primitives	TBD	LME_0002
13.2		LME Initialization Attributes	TBD	LME_0002
13.3.1.1		Administrative Status Down (LME4, LME5, LME6)	TBD	LME-0003 to 5
13.3.2		Topology Discovery	Yes	
13.3.3		Performance and Accounting Measurements	Yes	To be added.
13.3.3.1		Intervals	Yes	

System Architecture Use Cases

Several different use cases can be explored. One use case is common for an IPDSLAM/GPON OLT in which multiple IP streams must be aggregated and sent through the ring. Another use case is an MSPP application where legacy TDM traffic is mixed with data traffic. These two use cases comprise the majority of applications.

Figure 11 shows the basic system architecture for 1G and/or 2.5G. This is a good starting point.

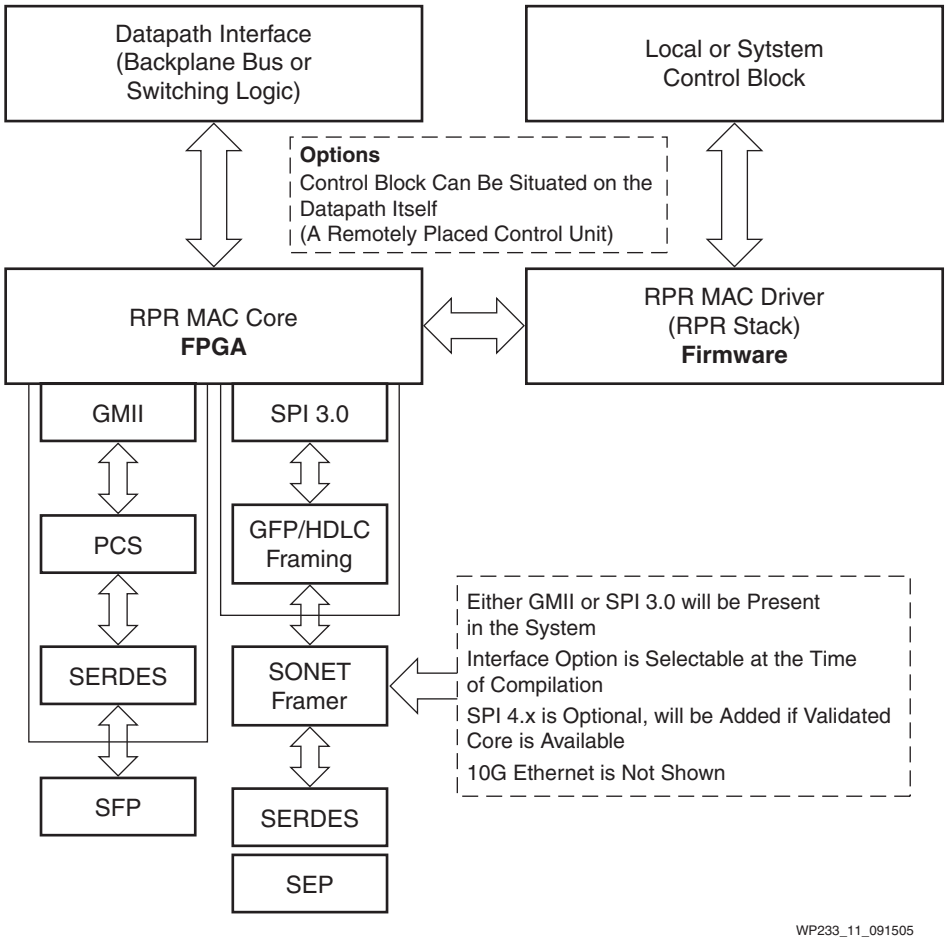

Figure 11: **System Architecture Baseline for 1G/2.5G RPR Interface**

Note that for the Sonet side, the framer could incorporate VCAT/LCAS to mix traffic.

IPDSLAM/GPON OLT

IPDSLAMs are rolling out from several vendors. GPON OLTs are starting to roll out as well. The application is the same; distribute and aggregate a bunch of IP streams over GPON or DSL. This is the access story from most tier 1 service providers. RPR is an ideal transport mechanism to distribute and aggregate this traffic. Figure 12 shows a trunk card implementation for this application.

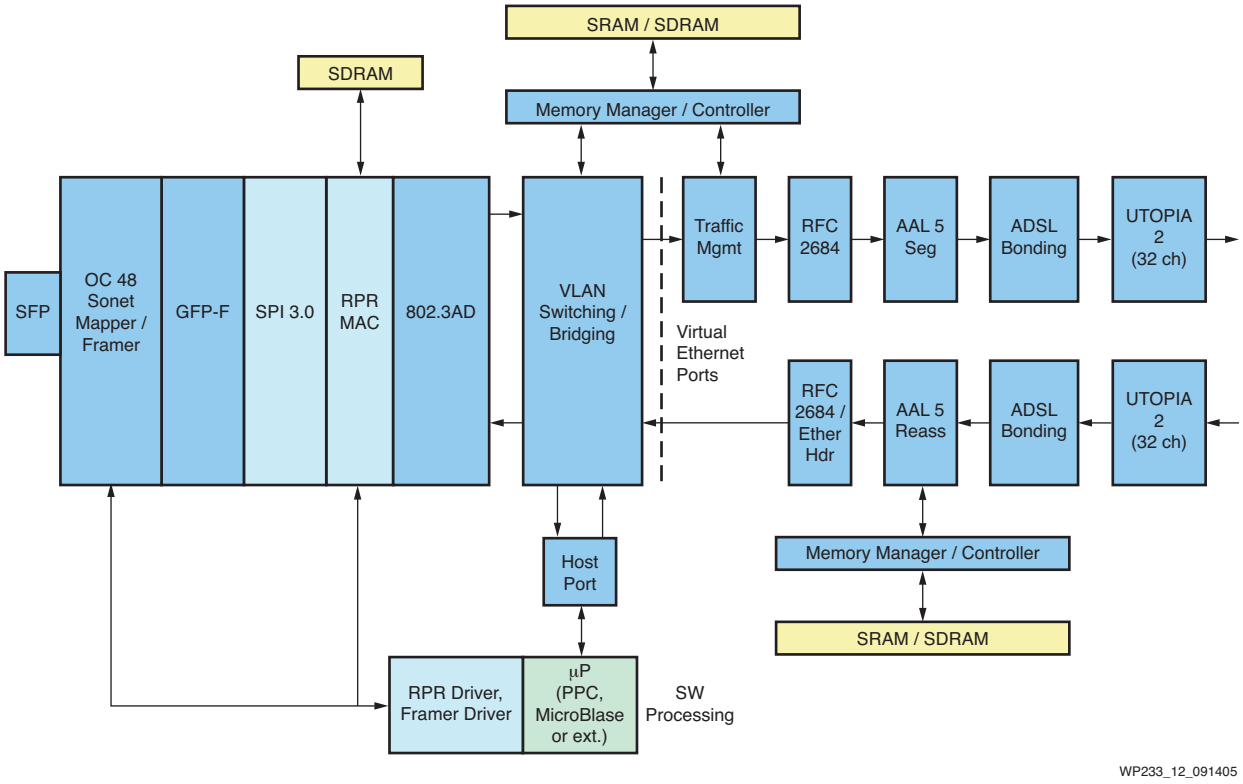

Figure 12: **IPDSLAM Trunk Card Architecture**

The RPR MAC, SPI3.0 interface and RPR driver are packaged as a part of the RPR solution. The rest is all IP blocks from other vendors. The traffic manager might not be needed if the granularity of QoS on the RPR MAC is enough. A GPON OLT trunk card would look very similar except the ATM interfaces would all be gone. IP can go directly to the OLT line cards.

MSPP Application

This application's main difference is that legacy TDM traffic must be preserved. A block diagram of a system is shown in Figure 13.

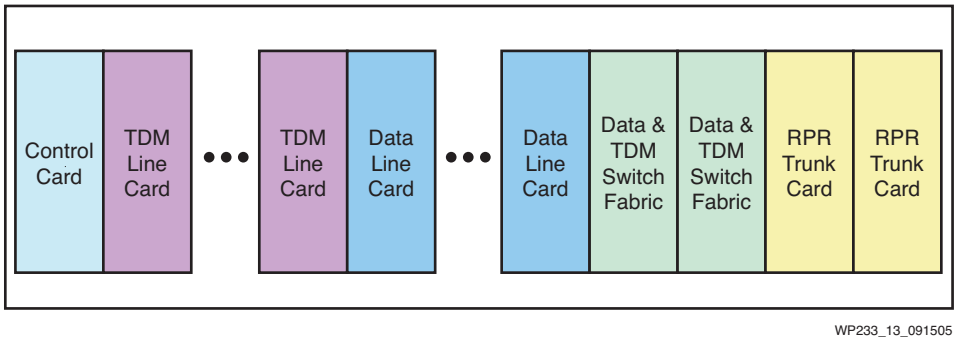

Figure 13: **MSPP**

Note that the dual switch cards and trunk cards are for redundancy, which is a common configuration. Figure 14 is a block diagram that illustrates the architecture and shows how the RPR MAC fits into the trunk card.

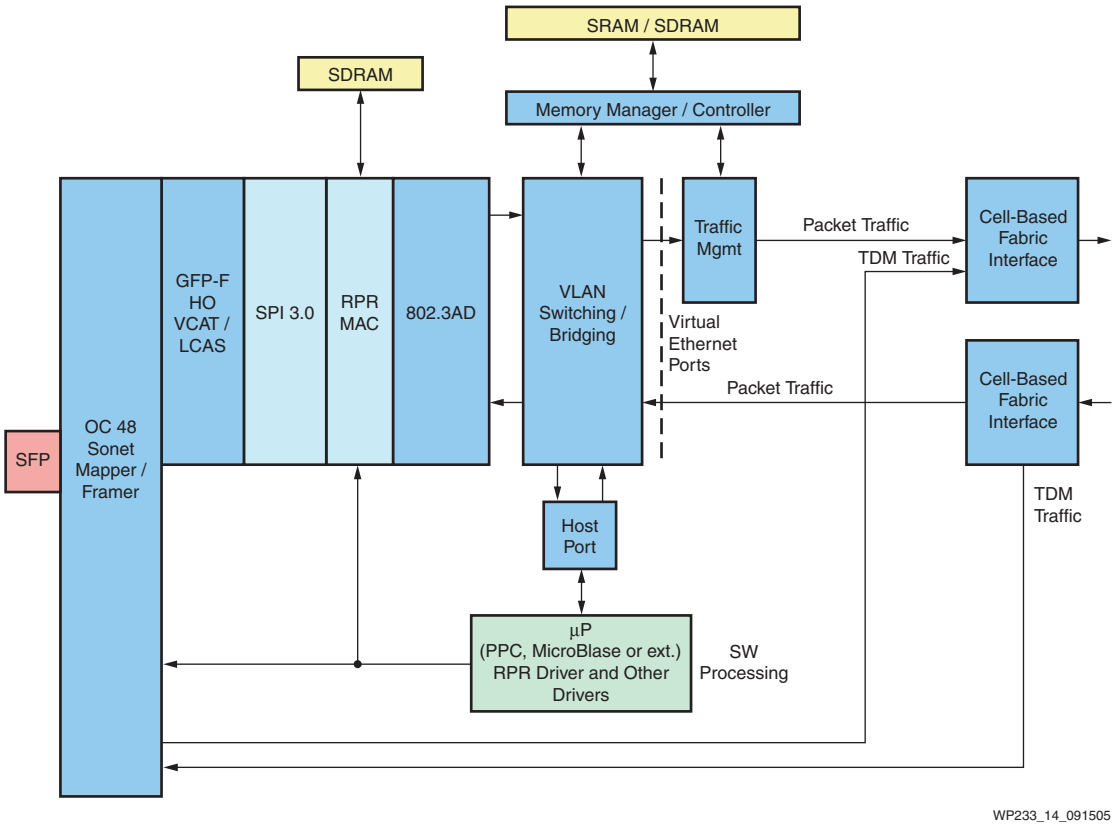

Figure 14: **MSPP RPR Trunk Card**

As shown in Figure 14, RPR is mapped into the Sonet frame at a sub-rate using a VCAT tunnel. The SPI3 interface is throttled back in order to accommodate whatever sub-rate is required. The remaining bandwidth in the OC48 can be used for any TDM traffic, or RPR can be assigned the entire bandwidth. This is a nice flexible option as network services evolve.

Conclusion

RPR is a very versatile transport technology that can enforce and provide the quality required for the new IP based services. Xilinx FPGAs in combination with the Xilinx RPR core is a powerful and cost effective solution that easily enables the development of the technology into next generation metropolitan access gear.

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
10/19/05	1.0	Initial Xilinx release.