

Enabling OTN Switching over Packet/Cell Fabrics

White Paper

by Korby Mraze and Scott Wakelin

Issue No. 1: December 2011

PMC-Sierra, Inc.

Legal Information

Copyright

Copyright 2011 PMC-Sierra, Inc. All rights reserved.

The information in this document is proprietary and confidential to PMC-Sierra, Inc., and for its customers' internal use. In any event, no part of this document may be reproduced or redistributed in any form without the express written consent of PMC-Sierra, Inc.

PMC-2114338, Issue 1

Disclaimer

None of the information contained in this document constitutes an express or implied warranty by PMC-Sierra, Inc. as to the sufficiency, fitness or suitability for a particular purpose of any such information or the fitness, or suitability for a particular purpose, merchantability, performance, compatibility with other parts or systems, of any of the products of PMC-Sierra, Inc., or any portion thereof, referred to in this document. PMC-Sierra, Inc. expressly disclaims all representations and warranties of any kind regarding the contents or use of the information, including, but not limited to, express and implied warranties of accuracy, completeness, merchantability, fitness for a particular use, or non-infringement.

In no event will PMC-Sierra, Inc. be liable for any direct, indirect, special, incidental or consequential damages, including, but not limited to, lost profits, lost business or lost data resulting from any use of or reliance upon the information, whether or not PMC-Sierra, Inc. has been advised of the possibility of such damage.

Trademarks

For a complete list of PMC-Sierra's trademarks and registered trademarks, visit:
<http://www.pmc-sierra.com/legal/>

Other product and company names mentioned herein may be the trademarks of their respective owners.

Patents

The technology discussed in this document is protected by one or more patent grants.

Abstract

OTN switching delivers a cost-effective, scalable, low-latency switching layer to address the massive bandwidth scaling required to deal with packet traffic growth. Furthermore, it delivers this scale in such a way that it can transparently transport legacy services such as TDM concurrent with the packet traffic. Service providers are demanding converged packet/OTN boxes, often referred to as Packet Optical Transport Platforms (P-OTPs), which support high-capacity OTN switching as well as packet functions such as Ethernet statistical multiplexing and MPLS-TP switching. TDM traffic is traditionally switched over a TDM-optimized switch fabric, and packet is traditionally switched over a packet/cell fabric. Maintaining two fabrics is challenging in terms of cost, power, and operational complexity. This white paper discusses how OTN switching can be accomplished over a converged fabric which also natively supports packet switching.

About PMC-Sierra

PMC (Nasdaq:PMCS) is the semiconductor innovator transforming networks that connect, move and store digital content. Building on a track record of technology leadership, we are driving innovation across storage, optical and mobile networks. Our highly integrated solutions increase performance and enable next generation services to accelerate the network transformation. For more information, visit www.pmc-sierra.com.

About the Authors

Korby Mraze is a Manager in the Communications Products Division at PMC-Sierra and has responsibility for supporting customers with system-level implementation of OTN in next generation optical equipment, including P-OTPs. Mr. Mraze has been with PMC-Sierra for over 13 years and has supported virtually every product in PMC's communications portfolio.

Scott Wakelin is a Product Manager in the Communication Products Division at PMC-Sierra and has responsibility for OTN and optical transport solutions. Mr. Wakelin has over 12 years of experience delivering innovative semiconductor solutions that implement OTN, SONET/SDH, Ethernet, and PPP/Frame Relay to the networking industry. Mr. Wakelin holds a Master of Applied Science degree that was focused on network infrastructure and security.



Revision History

Issue No.	Issue Date	Details of Change
1	December 2011	Document Created

Table of Contents

Legal Information.....	2
Copyright.....	2
Disclaimer	2
Trademarks	2
Patents	2
Abstract.....	3
About PMC-Sierra	3
About the Authors	3
Revision History.....	4
1 Introduction	8
2 References.....	9
3 Requirements for OTN Switching over a Packet/Cell Fabric.....	10
3.1 Packet/Cell Fabrics	10
3.2 Criteria for OTN Switching	11
3.2.1 Scalability & Flexibility.....	11
3.2.2 Packet Delay Variation Tolerance	11
3.2.3 Tolerant to Packet Loss.....	12
3.2.4 Protection Switching.....	12
3.2.5 Timing Transparency.....	12
3.2.6 Jitter Performance.....	12
4 Enabling Technologies	13
4.1 Constant Bit Rate Interface (CBRI).....	13
4.1.1 ODUk Packet Format.....	14
4.1.2 Packet Delay Variation Compensation.....	15
4.2 OPSA	15
4.2.1 OPSA Building Blocks	16
4.2.2 Jitter Performance of OPSA.....	18
4.3 OTN Protection with CBRI and OPSA	19
4.3.1 OTN Line Card and OTN Fabric Card Protection	20
5 Conclusion	22

List of Figures

Figure 1 System Reference Model.....	10
Figure 2 Packet and Cell Fabric Requirements for OTN Switching	11
Figure 3 PMC-Sierra's OTN Packet Header Format.....	14
Figure 4 OTN Packet with Packet Header in Interlaken Protocol	15
Figure 5 Retaining OTN Rate Information Through a Fabric	16
Figure 6 Packet Maker Block	17
Figure 7 Packet Maker Packets	17
Figure 8 Packet Monitor Block	18
Figure 9 Reassembled ODUk and Packet Sizes	18
Figure 10 Generic System APS Topology.....	20
Figure 11 OTN Line Card Protection.....	20
Figure 12 OTN Fabric Card Protection.....	21



List of Tables

Table 1 CBRI ODUk Capacity	13
Table 2 Packet Header Usage	14
Table 3 G.8251 Jitter Generation Requirements	19
Table 4 G.8251 Jitter Transfer Requirements	19

1 Introduction

Today, carrier networks are anchored on SONET/SDH for both Layer 1 transport and switching, with G.709 Optical Transport Network (OTN) used primarily as a digital wrapper on long haul interfaces. However, as discussed in [1], rapid packet traffic growth calls for a transport mechanism capable of transporting packets efficiently. Efficient transport of packet traffic over SONET/SDH infrastructure requires the use of VCAT/LCAS, which is operationally complex and expensive to deploy.

OTN addresses the limitations of SONET/SDH by virtue of payload areas that are right-sized for GE, 10GE, 40GE or 100GE signals. OTN defines Optical channel Payload Units, OPUks, which are well suited for Gigabit Ethernet (k=0) to 100 Gigabit Ethernet (k=4), and via the OPUflex is adaptable to any packet traffic rate or to any other arbitrary client type, such as Fibre Channel. OTN switching delivers two primary benefits to service providers. First, it allows them to efficiently groom client traffic into higher rate containers such as OTU1, OTU2, OTU3, or OTU4 so that wavelengths are used efficiently. Second, it provides a mechanism to separate client and network optics, providing both capital cost and operational cost efficiency.

A recent Infonetics survey [2] revealed both a clear demand for electrical switching of OTN (74% of respondents) and a clear need for switching down to the ODU0 level, with 79% of respondents saying they plan to deploy ODU0 level switching in the next 2 to 3 years. Therefore, the question is no longer if OTN switching at the electrical level is required, but how it will be deployed. As discussed in [1], there are two leading architectures:

1. **Packet Optical Transport Platforms (P-OTP):** Combines OTN switching and packet switching in a single platform.
2. **Optical Transport Platforms (OTP):** Provides OTN switching only, with Carrier Ethernet services provided by a separate Carrier Ethernet Switch/Router (CESR) platform.

One area that has garnered particular interest from network equipment designers is the use of packet/cell fabrics for OTN switching. The availability of such fabrics, and the inherent packet switching capability make this a particularly appealing option if a method can be devised to switch OTN traffic within ITU performance specifications.

To this end, PMC-Sierra has developed two fundamental technologies which enable OTN traffic and its timing to be transmitted over a packet/cell fabric:

1. **Constant Bit Rate Interface (CBRI):** A standards based interface between an OTN processor and the fabric.
2. **OTN Phase Signaling Algorithm (OPSA):** A method to convey OTN timing across the fabric.

Most recently, PMC-Sierra has led an industry-wide initiative in the Optical Interworking Forum (OIF) to standardize the manner in which an OTN processor interfaces to a packet/cell fabric for OTN switching. As a result of this industry wide effort, the OIF has recently ratified the OTN over Packet Fabric Protocol Implementation Agreement which is based on PMC-Sierra's CBRI and OPSA technologies.

This white paper begins by examining the system level issues and requirements that must be met in order to switch OTN traffic over a packet or cell fabric. Following this, the two fundamental enabling technologies, CBRI and OPSA, are explored in detail.

2 References

1. PMC-2110752, PMC-Sierra, Enabling OTN Convergence White Paper, Issue 1
2. *OTN Deployment Strategies: Global Service Provider Survey*. Infonetics Research. Retrieved 22 March 2011 <http://www.infonetics.com>.
3. Interlaken Alliance, Interlaken Protocol Definition Version 1.2

3 Requirements for OTN Switching over a Packet/Cell Fabric

First and foremost, OTN switching in P-OTP and OTP platforms requires a centralized fabric in order to switch OTN traffic from one line card to another. The system designer has several fabric options, including:

1. Packet or cell based fabrics
2. TDM based OTN fabrics
3. Digital Cross Points
4. SONET/SDH fabrics


Though TDM-based OTN fabrics and crossbar fabrics present clear advantages over packet fabrics for OTN switching, such as extremely low latency, and SONET/SDH fabrics require complex adaptation and concatenation schemes in order to support either OTN or packet, the OIF made an effort to enable interoperable OTN switching line cards using off-the-shelf fabrics. This effort resulted in the OIF ratifying a set of requirements based around packet fabrics. Packet fabrics were chosen for their (a) availability and (b) ability to converge both packet and TDM services.

The remainder of this section explores a system reference model for OTN switching and the requirements that it imposes on a system.

3.1 Packet/Cell Fabrics

Figure 1 shows the system reference model for OTN switching over a packet/cell fabric. In this model, the OTN processor extracts the ODUk signals from received OTUk signals, packetizes them, and transfers them over the fabric via the FIC device. Similarly, client signals (Ethernet, SAN, Video and other CBR clients) are mapped into an ODUk by the OTN processor which then packetizes the ODUk, and transfers it to the fabric. In this manner, all of the intelligence for adapting the OTN traffic for switching purposes is held within the OTN processor.

Figure 1 System Reference Model

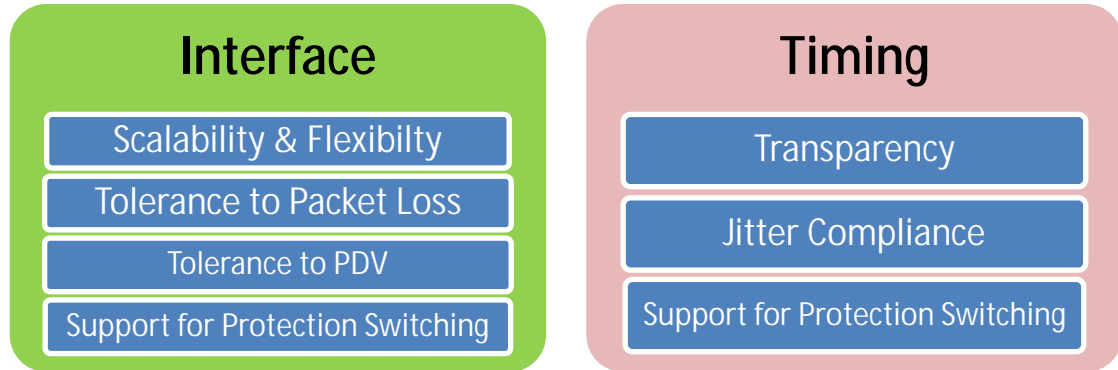


Packet/cell fabric architectures are centralized fabric topologies in which the fundamental unit of data transfer is a variable sized packet or fixed sized cell. The fabric typically operates asynchronously to the traffic that is switched over it.

As one would expect, a packet/cell fabric is meant for switching packets including Ethernet and MPLS. As such, packet fabrics are not inherently suitable for switching OTN without extra consideration. There are six basic requirements that must be met in order to switch OTN traffic over such a fabric. Figure 2 summarizes these requirements into two categories:

1. Requirements imposed on the interface between the OTN processor and the fabric
2. Requirements imposed on the timing characteristics of the system

Figure 2 Packet and Cell Fabric Requirements for OTN Switching



Let's explore these requirements in further detail.

3.2 Criteria for OTN Switching

This section explores the requirements for a mechanism that enables OTN switching over a packet/cell fabric.

3.2.1 Scalability & Flexibility

The interface between the OTN processor and the fabric must be scalable and flexible, including support for the following:

4. Line cards that scale from NxSFP to Nx100G port densities;
5. Hundreds of traffic flows per line card;
6. Payload rates from ODU0 (1.25G) to ODU4 (100G) on the same card, with adaptability to future rates including ODU5 (not yet defined); and
7. Arbitrary-size payloads, including any legal ODUFlex rate.

In ODUFlex implementations, the mechanism must allow for ODUFlex hitless resizing as defined by ITU-T G.HAO.

3.2.2 Packet Delay Variation Tolerance

Packet/cell fabrics are subject to varying delays due to congestion, scheduling changes and other factors. This packet delay variation (PDV) must not affect performance when switching OTN traffic through the packet/cell fabric. Appropriate buffering is required to emulate a fixed delay when switching OTN traffic from the OTN processor on an ingress line card, through the packet/cell fabric, to the OTN processor on the egress line card. Therefore, the interface between the OTN processor and the fabric (or Fabric Interface Chip) must support flow control.

3.2.3 Tolerant to Packet Loss

If a packet containing a portion of an ODUk frame were to be lost or dropped by the fabric, the lost data would result in a loss of frame event that could have network wide implications.

As a result, the mechanism for OTN switching over a packet/cell fabric must not only detect the occurrence of a lost packet, but take reparative measures such as replacing the lost data. Although the lost data will result in a burst of errors, the impact will be less severe than a loss of frame event.

3.2.4 Protection Switching

Protected OTN traffic switched over a packet/cell fabric must recover from failure conditions within 50ms. The packet stream start up mechanism following a failure must be done in a manner that ensures the 50ms switchover time can be met.

3.2.5 Timing Transparency

In SONET/SDH, the containers that transport the payload float within the overall frame structure. Pointer movements are used to adapt one timing domain to another. In SONET/SDH switching systems, the SONET/SDH signal is effectively retimed to synchronize it to the local fabric.

In contrast, OTN is an asynchronous protocol and has no such floating payload or system of pointer movements for retiming. In order to avoid buffer overflows or underflows, the egress line card must transmit an ODUk at the same rate at which it was received at the ingress line card. In order to accomplish this, the timing for an ODUk that is to be switched across a fabric must be transmitted along with the ODUk.

OTN is different from SONET/SDH in that the network need not be synchronous. Every ODUk circuit has a unique phase within the transport network. This simplifies the network since not all network elements need to be synchronized to one another. Furthermore, each ODUk circuit is rate-adapted to its constituent payload due to the fact that the OTN frame is created by simply wrapping around the constituent packet or TDM client.

The OTN frame structure is always a fractional multiple of the client being transported and therefore cannot have its payload synchronized to the fabric, as is done in a more traditional TDM-based SONET/SDH cross-connect. A new scheme is required to transport OTN over a packet/cell fabric transparently so that both its timing and its content are completely unaltered.

3.2.6 Jitter Performance

OTN client signals traversing a packet/cell fabric must meet the ITU-T Recommendation G.8251 ODCr and ODCp timing specifications. This must be done for the maximum number of network elements allowed by the G.8251 Hypothetical Reference Model. Any scheme for the segmentation of OTN traffic into packets and subsequent reassembly must adhere to the above specifications.

4 Enabling Technologies

Two fundamental technologies are required to enable an ODUk and its timing to be transmitted over a packet/cell fabric:

1. A standards-based interface between the OTN processor and the fabric.
2. A method to convey ODUk timing across the fabric.

A channelized scalable link-level protocol is needed from the OTN processor to the fabric that not only meets the requirements of the next generation of P-OTPs, but can scale to meet their needs in coming years. PMC-Sierra's implementation of this interface is the Constant Bit Rate Interface (CBRI) which is built on Interlaken. CBRI uses a packet header which is compliant to the OIF's OTN over Packet Fabric Protocol (OPFP).

In order to convey the timing information of OTN traffic over a packet/cell fabric, a method is required to encode the traffic rate within OTN packets without adding any overhead. PMC-Sierra's mechanism, the OTN Phase Signaling Algorithm (OPSA), is compliant to the OIF's OPFP.

This section introduces PMC-Sierra's CBRI and OPSA technologies and highlights how they meet the requirements in the previous section.

4.1 Constant Bit Rate Interface (CBRI)

As stated above, CBRI is based upon the Interlaken Protocol. The Interlaken Protocol Specification Version 1.2 is a multilane scalable packet interface that supports bandwidth transfers from 10 Gbps to 120 Gbps. Furthermore, Interlaken is channelized, allowing for 10s or 100s of ODUk's to be supported per line card. Flow control is also supported.

The Interlaken interface has standardized bandwidths that are composed of multiple SERDES lanes operating at a variety of rates. Table 1 shows the maximum number of ODUk streams that can be transferred over an interface of a given bandwidth. As an example, a 10 G Interlaken Interface can support 8 ODU0s, while a 50 G Interlaken interface can support 32 ODU0s. Combinations of various ODUk sizes are also supported so long as the overall bandwidth of the interface is not exceeded.


Table 1 CBRI ODUk Capacity

Interface Bandwidth	ODU0 Streams	ODU1 Streams	ODU2 Streams	ODU3 Streams	ODU4 Streams
10 Gbps	8	4	1	N/A	N/A
25 Gbps	16	8	2	N/A	N/A
50 Gbps	32	16	4	1	N/A
100 Gbps	80	40	10	2	N/A
120 Gbps	96	48	12	3	1

4.1.1 ODUk Packet Format

The OTN over Packet Fabric Protocol defines a packet header that is prepended to each OTN packet prior to being sent into the fabric. The mandatory section of the header provides bit fields used to perform such functions as packet delay variation compensation, dropped packet recovery, packet sequence numbers and client status. These bit fields are shown in Figure 3.

Figure 3 PMC-Sierra's OTN Packet Header Format



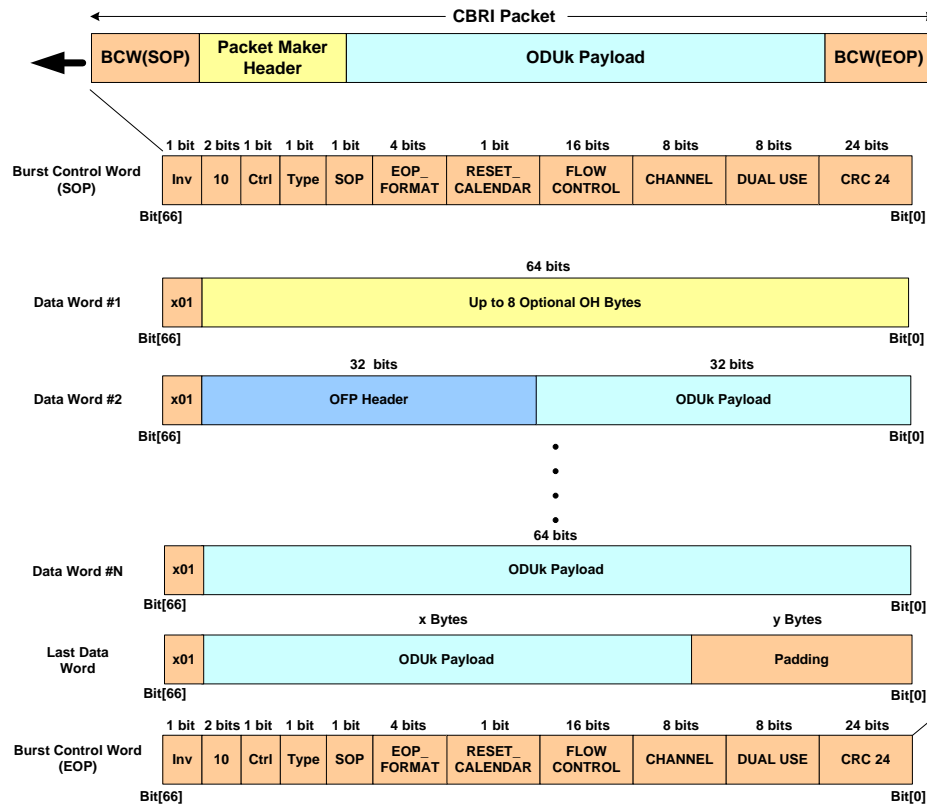
The fields of the OIF OTN over Packet Fabric Protocol packet header are described in Table 2 along with PMC's usage of each field.

Table 2 Packet Header Usage

Field	OIF Defined Usage	PMC-Sierra Usage
Timestamp	Used primarily to tune a packet delay variation (packet jitter) buffer on line cards in a packet/cell fabric	Same as OIF
RSV 1	Optionally combined with the existing sequence bits to generate a larger sequence number	Used by PMC devices to increase the sequence number range
SQ	Used to detect packet drops due to bit errors or congestion	Same as OIF
PPSI 1	Indicates the size of the previous packet to ensure the replacement packet of a dropped is packet sized correctly to avoid a change in frame alignment for the ODUk/ODUflex stream	Same as OIF
CSI	Optionally used to indicate the status of the client	Generates interrupts upon changes in state, used to declare signal fail indications
PPSI 2	Optionally indicates the size of the second previous packet	Not supported
P	A 1-bit odd parity calculation over the entire header	Same as OIF
User Specific and Fabric OH	Optional 0-12 byte field reserved for user proprietary data and fabric overhead	Up to an 8 byte user defined field is supported
Payload	ODUk/ODUflex packet of size $B_{nom}-1$, B_{nom} , or $B_{nom}+1$ bytes	Same as OIF

Figure 4 shows the entire OTN packet including standard and optional header along with the OTN payload inside an Interlaken formatted packet. The Interlaken protocol wraps the Burst Control Word Start of Packet and End of Packet around the OTN related data.

Figure 4 OTN Packet with Packet Header in Interlaken Protocol



4.1.2 Packet Delay Variation Compensation

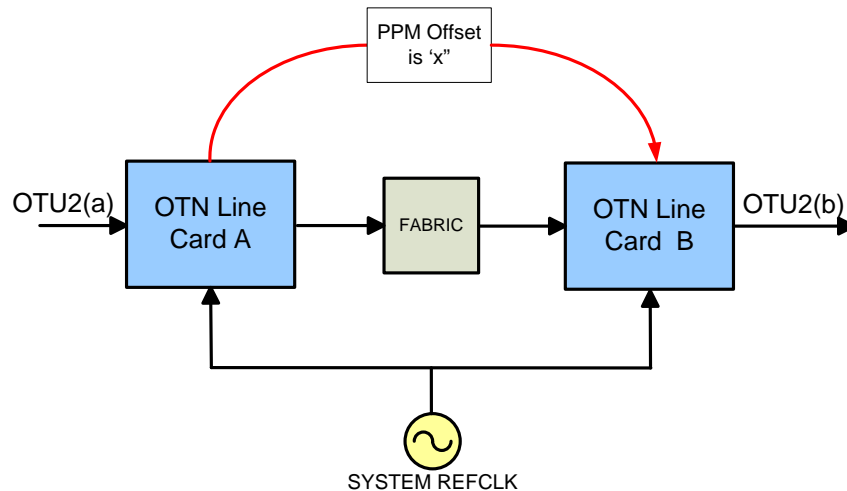
The packet delay variation compensation function of the packet header is achieved by monitoring the timestamp field. In order to use the timestamp field, all line cards in the node must share a common 8 kHz reference frame pulse and a common 311.04 MHz reference clock. The 8 kHz pulse is distributed to all line cards in the system to provide the ability to measure the “age” of the packets that it is receiving. This packet age is used to determine how old the packet is relative to the maximum age or maximum packet delay the fabric can have. The age information is then used as part of a flow control mechanism to tune a PDV buffer to the appropriate depth to ensure each packet reaches the packet reassembly block with the same delay.

4.2 OPSA

In the OIF’s OTN over Packet Fabric Protocol (OPFP), the OTN rate information is encoded by varying the size of the ODUk packets by ± 1 byte based upon a 311.04 MHz system reference clock. The packets are generated on the ingress line card at a nominal size (B), nominal minus one byte (B-1) or nominal plus one byte (B+1) over a period of T system reference clock cycles. To determine the rate of the ODUk stream, the size of the OPFP packets can be measured and then low pass filtered at the egress line card.

By encoding this rate information into the packet data itself, the packets are transferred over the packet/cell fabric and then reassembled and transmitted out at the original clock rate when the OTN traffic leaves the fabric. PMC-Sierra's OTN processors generate the B+/-1 ODUk rate information in compliance with the OIF's OTN over Packet Protocol, ensuring interoperability with 3rd party implementations. PMC-Sierra extends this functionality with the OPSA algorithm which provides enhanced jitter performance over the full range of ODUk rates. The rest of this section discusses OPSA, with references to OPFP where necessary.

Figure 5 Retaining OTN Rate Information Through a Fabric



4.2.1 OPSA Building Blocks

The OPSA process is achieved by three fundamental building blocks.


1. **Packet Maker:** Segments incoming data into artificial packets that vary in size from nominal by ± 1 byte.
2. **Packet Monitor:** Reassembles the data stream from received packets, also measures the size of incoming packets.
3. **Adaptive Clock Block:** Low pass filter function that takes packet size information from the packet monitor to control a voltage controlled oscillator (VCO).

PMC-Sierra's low pass filter features an ultra-low cutoff frequency in the sub 1 Hz range, that, in combination with the Packet Maker and Monitor, generates unmatched jitter and wander performance.

Packet Maker

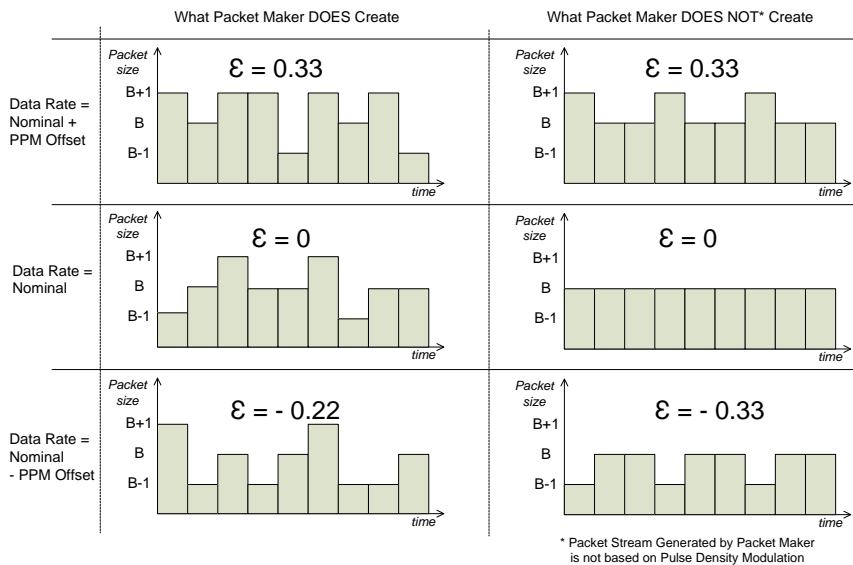
The Packet Maker block takes continuous OTN data in and segments it into packets at a size B, B+1 or B-1 bytes in length.

Figure 6 Packet Maker Block



The decision to make the packet a certain size is done every T reference clock cycles using a Delta-Sigma Modulator (DSM). The DSM noise shapes the packet size decisions such that packets of size B-1, B, and B+1 are always being made, which avoids low frequency phase discontinuities caused by periods where no packet size changes occur. DSM ensures frequent packet size decision changes which are easier to filter out at the egress line card. Figure 7 below shows the difference in packet sizes when a DSM is used vs. not being used. In this figure, EPSILON (ϵ) is used to define the average value of the packet size decisions.


Figure 7 Packet Maker Packets



Packet Monitor

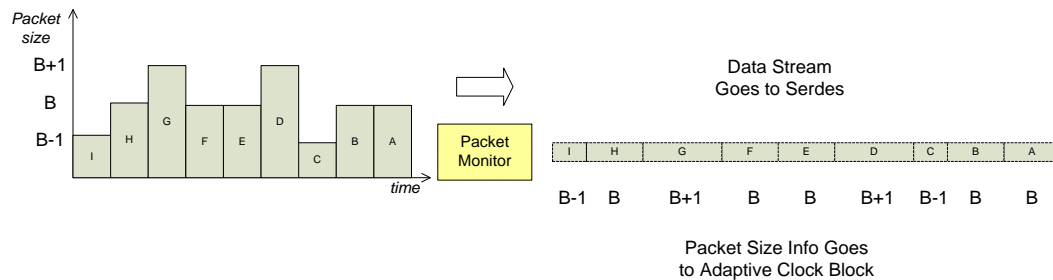
After traversing the packet/cell fabric, the OTN packets arrive at the egress line card and must pass through the Packet Monitor block that reassembles the original OTN data stream and also collects the size of each packet received from the fabric.

Figure 8 Packet Monitor Block



The reassembled OTN packet stream is usually sent to a downstream OTN framer prior to being clocked out of the SERDES by the clock reconstructed by the Adaptive Clock Block (ACB) from the packet size information.

Figure 9 Reassembled ODUk and Packet Sizes



Adaptive Clock Block

All PMC-Sierra OTN processors feature an integrated Adaptive Clock Block (ACB) on each line and client interface.

The ACB takes the packet size information from the Packet Monitor block and filters the high frequency shaping noise introduced at the far-end Packet Maker. The result is then used to tune an integrated voltage controlled oscillator (VCO) which re-creates the original ODUk clock to drive the integrated SERDES.

Integration of PLLs on each line and client interface allows a single timing reference to time the card and eliminates the need for external clocking circuitry.

4.2.2 Jitter Performance of OPSA

Recovering the timing information from a packet fabric can lead to unwanted timing jitter. The OIF requires that the clock re-generated using the OTN over Packet Fabric Protocol meets the jitter requirements of the ITU-T G.8251 specification.

PMC’s OPSA algorithm exceeds these requirements to ensure that there is additional margin in the optical system’s jitter budget. The Hypothetical Reference Model (HRM) requirement is met through noise shaping at the Packet Maker and with the ACB’s low pass filter set to 300 Hz. OPSA’s jitter performance exceeds that of the HRM by utilizing the ACB’s ability to improve its low pass filter to the sub-Hz range.

The following tables show the G.8251 jitter generation and transfer requirements. Actual measured results are available from PMC-Sierra upon request.

Table 3 G.8251 Jitter Generation Requirements









Rate	Jitter Frequency Range	G.8251 Requirement	Exceeded by PMC-Sierra
OTU1	5 kHz to 20 MHz	0.3 UI	
OTU1	1 MHz to 20 MHz	0.1 UI	
OTU2	20 kHz to 80 MHz	0.3 UI	
OTU2	4 MHz to 80 MHz	0.1 UI	

Table 4 G.8251 Jitter Transfer Requirements

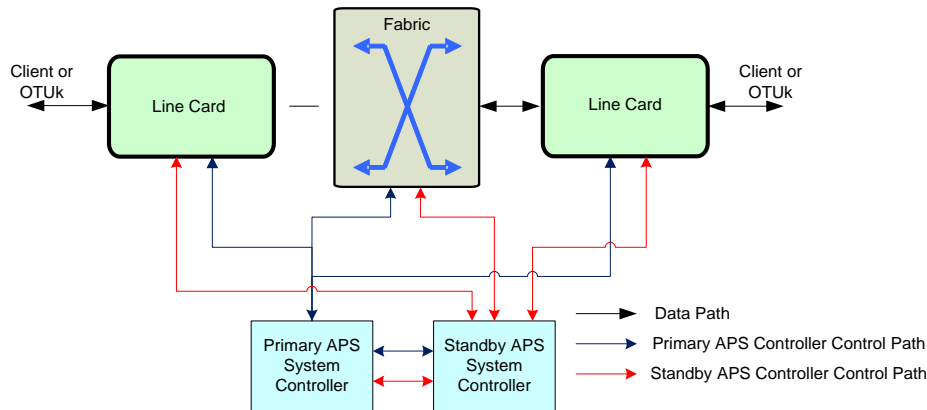
Rate	Jitter Frequency Range	G.8251 Requirement	Exceeded by PMC-Sierra
OTU1	2.5 kHz to 250 kHz	0.1 dB	
OTU1	250 kHz to 20 MHz	-20 dB/decade roll off	
OTU2	10 kHz to 1 MHz	0.1 dB	
OTU2	1 MHz to 80 MHz	-20 dB/decade roll off	

4.3 OTN Protection with CBRI and OPSA

Packet/cell fabrics switching OTN traffic must provide the same level of protection as SONET/SDH fabrics do when switching STS/AU traffic. These requirements are met when utilizing CBRI and OPSA.

In the generic fabric protection scheme, OTUk interface protection is localized to the individual line cards, whereas line and fabric card protection needs to involve the fabric's Automatic Protection Switching (APS) control system due to the need to co-ordinate alarms and switching functions across multiple cards in the system. The APS controller continually monitors the working and protect hardware. When a failure is detected, it makes provisions to the system to route traffic from the failure to the redundant protection hardware. This minimizes the downtime for any enabled services being transported and switched through the system.

Figure 10 Generic System APS Topology



4.3.1 OTN Line Card and OTN Fabric Card Protection

The following two figures show working and protect OTN line and fabric cards utilizing CBRI as the OTN processor to fabric interconnect and OPSA to generate the OTN packets. Under normal conditions, only the working card has OTN traffic enabled and the protect card does not in order to ensure the fabric is not unnecessarily loaded with the protect traffic. When a failure is detected on the working card, its traffic is dropped for a period greater than or equal to the fabric's maximum latency before the traffic from the protect card is enabled. This allows the egress line card to identify the first packet from the protection card easily and re-initialize the PDV FIFOs using the CBRI timestamps. Depending upon the fabric, this recovery process should only require 10s of microseconds of a 50ms total protection switch budget.

Figure 11 OTN Line Card Protection

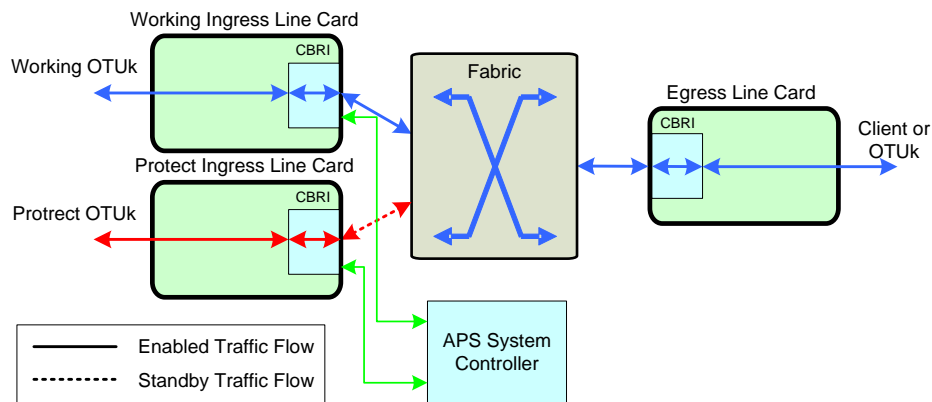



Figure 12 OTN Fabric Card Protection



5 Conclusion

Service providers worldwide are planning for and deploying OTN switching in their next-generation transport networks, built around Packet Optical Transport Platforms (P-OTPs) with a converged packet fabric. These systems support the convergence of both packet and TDM services and have the breadth of client support and scalability to meet aggregation, metro and long haul network requirements.

PMC-Sierra was first to market with OTN processors that enabled cost-effective and power efficient OTN switching via a centralized fabric. As a result of its leadership position, PMC-Sierra helped drive the OIF's standardization of the OTN Packet Fabric Protocol Implementation Agreement. Through customer deployment and interoperability testing, PMC-Sierra has proven that our technology, including OPSA and CBRI, meets all standard requirements.

Furthermore, PMC-Sierra's OPSA and CBRI technology offer OEMs significant advantages when deploying P-OTPs:

1. Fully integrated PLLs on all interfaces plus the need for only a single timing reference, which significantly reduces BOM cost and board space requirements
2. Integration of ultra-low frequency low-pass filters ensure optimal jitter and wander performance.

PMC-Sierra's continuing innovation in optical networking will help propel the transformation of networks globally. PMC-Sierra's HyPHY family of OTN processors enable designers to build P-OTPs based on a single fabric which is capable of simultaneously functioning both as an OTN switch element as well as a packet switch element in the transport network of the future.



Contacting PMC-Sierra

PMC-Sierra
8555 Baxter Place
Burnaby, BC
Canada V5A 4V7

Tel: +1 (604) 415-6000

Fax: +1 (604) 415-6200

Document Information: document@pmc-sierra.com

Corporate Information: info@pmc-sierra.com

Technical Support: apps@pmc-sierra.com

Web Site: <http://www.pmc-sierra.com>