



FocalPoint TDM Support

With Lossless Switching and Guaranteed
Bandwidth

Application Note

Nov, 2009

Table of Contents

Table of Contents	2
1.0 Introduction.....	3
2.0 FocalPoint Features for TDM.....	3
2.1 FocalPoint Congestion Management	4
2.2 FocalPoint TDM Frame Format	5
3.0 Converged Packet-TDM System using FocalPoint	6
3.1 Packet Blade	6
3.2 TDM Blade.....	7
3.3 Transport Blade.....	7
4.0 Mapper / Demapper Device	8
4.1 Mapping Function	8
4.2 Demapping Function.....	9
5.0 System Timing	9
6.0 High Availability.....	10

1.0 Introduction

SONET/SDH and OTN frames are prevalent in Carrier and Transport Networks. These Time Domain Multiplexed (TDM) frames are the medium used to carry critical time sensitive data. The design of these networks allow them to provide guarantee bandwidth and loss free delivery.

Today, both the nature of the services carried by these networks, and the networks themselves are changing. The majority of new services and applications are now packet oriented and delivered using Ethernet frames over packet switched networks. Some of these applications in fact are already carrying TDM encapsulated flows such as VoIP. Moving forward, it is only natural to assume that common equipment will be used, and packet network enhancements will evolve to carry TDM services.

Earlier transport systems contained TDM backplane fabrics with support for TSI functionality. Support for packet services was limited, and packets were converted to TDM frames on the line card. Today, as more packet data services are offered to customers, many systems support both packet and TDM fabrics as shown in the figure below. A system like this, however, is costly, as it must support two types of fabrics through the backplane.

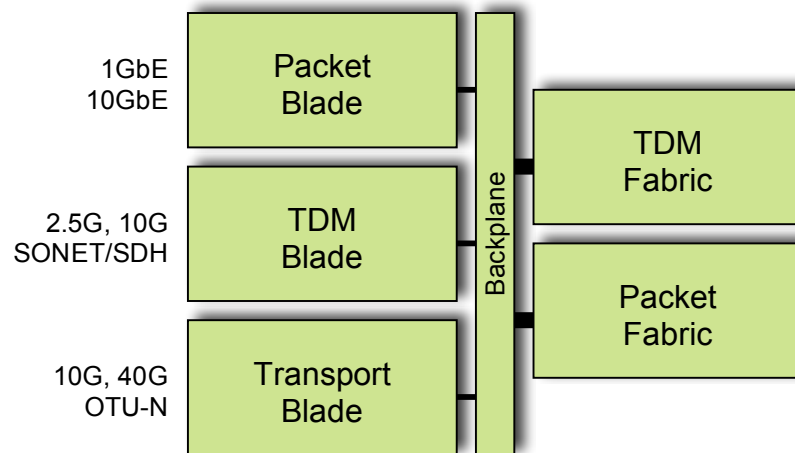


Figure 1. Traditional Mixed Service System Architecture

Over time, the carrier networks will migrate to entirely packet-based technologies. OTN is still the protocol of choice for metro and long-haul optical networks requiring FEC. Because of this, it makes sense to design new transport systems with packet-based fabrics that can support both TDM services, and packet based services

The objective of this application note is to show how Fulcrum’s FocalPoint switch products can be employed to support both TDM and packet based services. The TDM services can be delivered over the switch with no loss, and with guaranteed bandwidth. Using a FocalPoint switch to converge these services can be a very cost effective solution.

2.0 FocalPoint Features for TDM

The FocalPoint family of 10GbE switches provide features that have enabled their use in a variety of telecom backplane applications. This section will discuss several of these

features, which enable the transport of TDM flows along with packet flows within the same backplane fabric.

2.1 FocalPoint Congestion Management

TDM flows require bounded and deterministic latency through the backplane so that network timing can be maintained across the system, allowing buffering mechanisms to smooth the TDM flows on egress. FocalPoint provides a true output queued switch architecture. All packets are forwarded over the switch in a loss free manner, and stored in the output queues such that the output queues are never under-run, and there is sufficient elasticity for clocking compensation

FocalPoint provides several congestion management mechanisms to support this as shown in figure 2. In this figure, only a single ingress and egress port are shown for simplicity.




Figure 2. TDM Congestion Management

As ingress data arrives in the FocalPoint fabric, an advanced TCAM-based classification engine can inspect the first 128 bytes of the frame and use ACL rules to identify a traffic class for each frame. Based on traffic class, frames can be placed into one of several logical shared memory partitions in the switch. For example, TDM traffic can be identified and placed into a special memory partition, while other packet-based traffic can be placed into other memory partitions.

Each memory partition has an independent set of watermarks that can be used to generate link-level flow control back to the ingress link partner using IEEE Priority Flow Control frames. This insures that TDM traffic will not be delayed or dropped if packet traffic becomes congested in the switch.

Each switch egress port contains a separate queue per traffic class along with a multi-level scheduler. The scheduler can be programmed to service a given traffic class as strict priority or give it minimum bandwidth guarantees using Deficit Round Robin (DRR). In addition, traffic classes, or groups of traffic classes can be shaped to limit maximum egress bandwidth. Using these mechanisms, TDM traffic can be serviced as strict priority, or if multiple TDM flows are sent to a given egress port, each one can be given a minimum bandwidth guarantee. By doing this, the maximum latency and latency jitter can be bounded for TDM flows through the fabric. This will be discussed further in section 4.2

For systems that have dedicated TDM and packet interfaces on the switch, the traffic class can be simply associated with the ingress port number of the switch. This eliminates the need to carry priority bits in the TDM packet header. In addition, FocalPoint can be used to create multi-stage fabrics while maintaining lossless operation with minimum bandwidth guarantees for TDM flows.

In all cases of SONET/SDH and OTN packets will need to be segmented and encapsulated into Ethernet frames outside of the switch using an FPGA or other companion device. In this process a tag could be added to uniquely identify the flow and the priority. The added tag should also include a time stamp to ensure that over a multi-stage fabric the packets can be reassembled in order and at the right time.

Since the TDM data is encapsulated with an Ethernet header, overspeed should be provided in the fabric. The fabric should also have overspeed to mitigate delay jitter due to congestion. For example, for a shared egress port, if there is a non-TDM jumbo packet being transmitted, the TDM data must be buffered for a short period of time.

2.2 FocalPoint TDM Frame Format

In order to assure the proper level of bandwidth granularity for TDM flows through the switch, the fixed TDM payload segment size is set to 8-bytes and the Ethernet frame sizes are optimized for the switch memory at 160-bytes as shown in figure 3. The mapping and de-mapping of these segments is discussed in section 4.

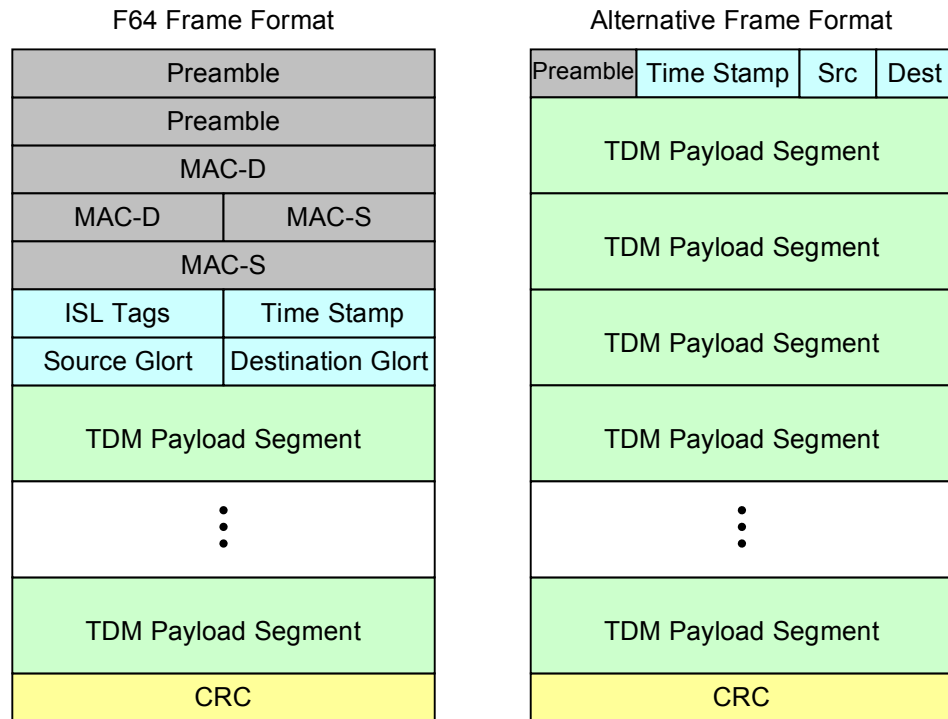


Figure 3. TDM 160-byte Frame Format

The frame on the left is an F64 frame with a standard Fulcrum ISL tag inserted in the header. This frame is set to 160-bytes in order to carry the ISL tag along with sixteen 8-byte TDM payload segments (the 12-bytes of MAC address could also be used for TDM payload if needed). The VLAN tag can be substituted with a frame time stamp, and the source and destination Glort values are used for frame forwarding. The TDM flows are provisioned; therefore the Glort tables are set to a fixed value by software. This frame format provides 80% payload bandwidth utilization, and therefore requires a 1.25x overspeed in the backplane fabric.

Application Note: FocalPoint TDM Support

Some FocalPoint devices can forward the frame using information in the preamble bytes as shown in the frame on the right. Of the 8 preamble bytes, 7 can be parsed and all of them can be forwarded from ingress to egress. Because of this, 4 of the preamble bytes are actually used for payload transport as shown. In this example, 12-bits of time stamp information are mapped into the frame preamble while 6-bits are available for the source address and 6-bits are available for the destination address. This alternative header provides over 95% payload bandwidth utilization, requiring only a 5% speedup factor in the backplane.

3.0 Converged Packet-TDM System using FocalPoint

This section provides an example transport system design, which is similar to the one shown in figure 1, but with a single converged packet/TDM backplane fabric. The fabric cards are not shown in this section, as they would simply contain multiple FocalPoint switches. A number of redundant backplane interfaces are required on the cards shown below. This is for both high availability and for the overspeed discussed above.

3.1 Packet Blade

The packet blade shown in figure 4 aggregates multiple Ethernet streams for transport across the carrier network. These streams can come from systems such as GPON, CMTS or LTE. As shown in the figure, FocalPoint can aggregate multiple 1GbE or 10GbE ports into 10GbE or 40GbE redundant backplane ports. FocalPoint can also pre-classify packets using a TCAM based packet classification engine, and re-direct packets to an attached NPU or CPU for deep packet inspection.

It is assumed that all packets between these blades will be transported across the backplane as full packets and not as TDM frames. As networks evolve into the future, it is also expected that Packet Blades become the dominant line card in the system as TDM services are phased out over time.

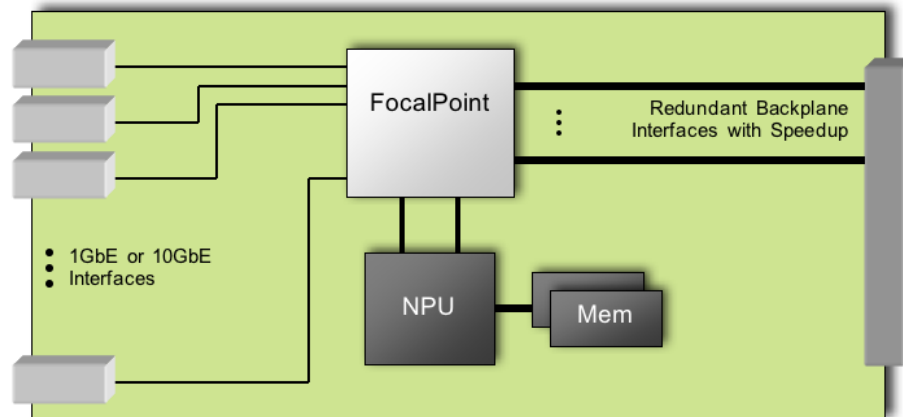


Figure 4. Example Packet Blade

3.2 TDM Blade

The TDM blade is typically the most expensive blade in the system, due to its complexity. This and the fact that it is difficult to scale the components to 40G or 100G, are some of the reasons that SONET and SDH are being replaced by Ethernet in the carrier networks.

In the near future, SONET and SDH infrastructure can be supported using blades like the one shown in the figure below. Here, SONET or SDH streams can be terminated and generated using standard Framer and Pointer Processor silicon. TDM frames are used to send data between these blades across the Ethernet backplane. A Mapper/Demapper device is used to transmit and receive these frames as will be discussed in section 4. This device can also extract network timing from these SONET/SDH streams if needed.

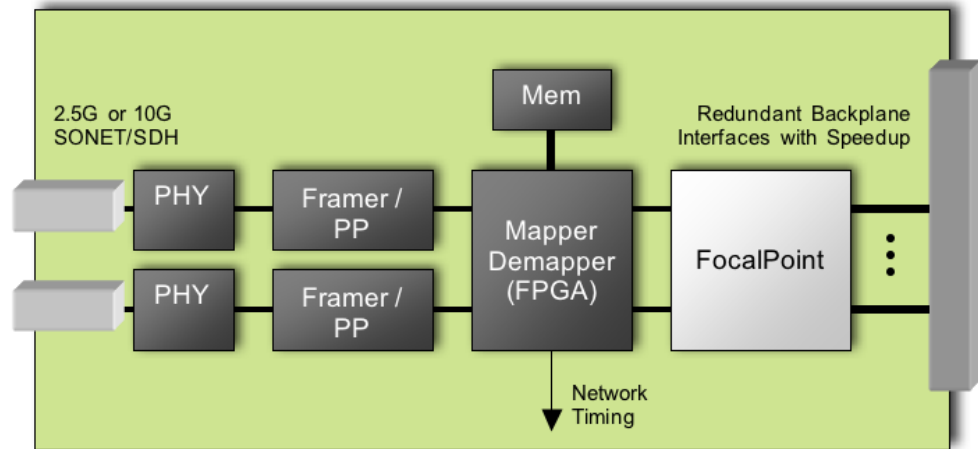


Figure 5. Example TDM Blade

For ingress channelized TDM flows of different bandwidth, the Mapper would have several bins (e.g. 10 for ODU2). Some bins will fill at different rates, so the frame generation out of the Mapper is based on a fixed time interval. The bins represent independent TDM flows, which are shifted out of the Mapper into the fabric.

The TDM blade may also support additional functions that are not shown in the figure above. One function would be Ethernet over SONET framing, where standard Ethernet packets from the Packet Blades are sent to this blade, bypassing the Mapper/Demapper function. Another function is Time Slot Interchange (TSI), which could be performed on this card.

3.3 Transport Blade

The Transport Blade formats data for transportation across the carrier network as shown in the figure below. The OTM Mapper can place a wrapper around SONET/SDH frames or Ethernet Frames for network transport. In this case, the Mapper/Demapper can convert TDM flows to SONET/SDH frames, or operate in bypass mode, presenting Ethernet packets directly to the OTN Mapper for network transport.

Most OTN Mappers today support direct packet interfaces such as XAUI. If the encapsulation is other than PoOTN, then the Mapper/Demapper FPGA is required to translate Demapped SONET/SDH to packets. Eventually, as these systems become

Application Note: FocalPoint TDM Support

completely packet-based, the Mapper/Demapper function will be no longer needed and the OTN Mapper can connect directly to the fabric through the XAUI interface.

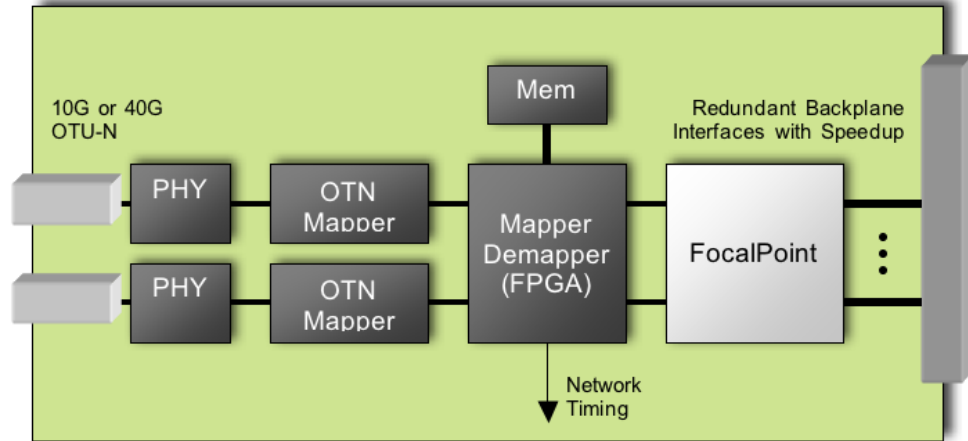


Figure 6. Example Transport Blade

This blade can also provide network timing to the system as will be discussed in section 5. Although the TDM blade can also provide network timing, the Transport Blade can provide network timing for Synchronous Ethernet and IEEE 1588 PTP in a packet-only system once the TDM blades are no longer used.

4.0 Mapper / Demapper Device

The Mapper/Demapper device is used on the TDM blades and the Transport Blades to map TDM payload segments into the TDM frames shown in figure 6. This section will describe these functions in more detail.

4.1 Mapping Function

The mapping function uses time slot positions to identify the destination blades for various TDM segments. Figure 7 shows how this would look for three destination blades in the backplane fabric.

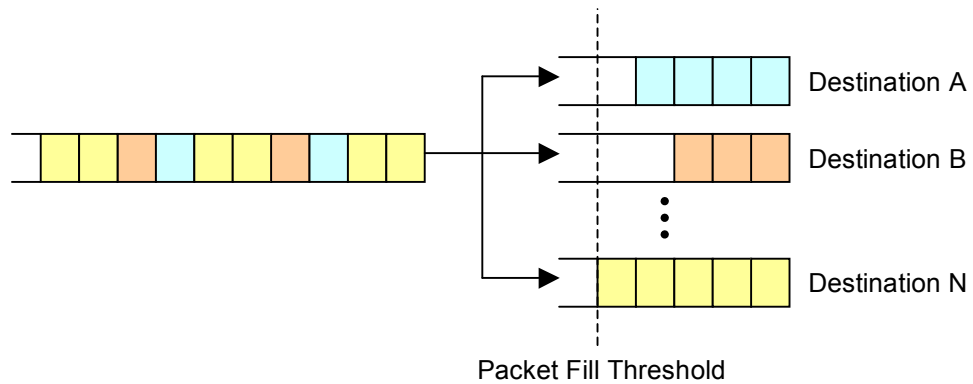


Figure 7. Mapping time slots into TDM packets

Once the packet fill threshold is reached, the segments are encapsulated into a TDM frame and forwarded into the fabric. In order to bound frame jitter on the receive side, a timeout value is used to make sure frames are transmitted even if the packet fill threshold is not reached. It is up to the control plane software to configure the time slot to destination mapping tables. The queues eligible for transmission can be serviced in a round-robin fashion. The TDM frame transmission time is tightly controlled by a global Frame Clock.

4.2 Demapping Function

The Demapping function receives TDM packets from several source blades in the switch, which are identified in the header. Due to the FocalPoint congestion management features, packets from each source will arrive with bounded latency, but also with random delay variations between queues. Because of this, FIFOs are employed which remove jitter and align the frames. This process will be discussed further in section 5. Once the frames are re-aligned, their TDM payload segments can be removed and inserted into an egress stream based on programmed time slot information. It is up to the control plane software to make sure the right amount of bandwidth is allocated throughout the switch to support all the TDM flows.

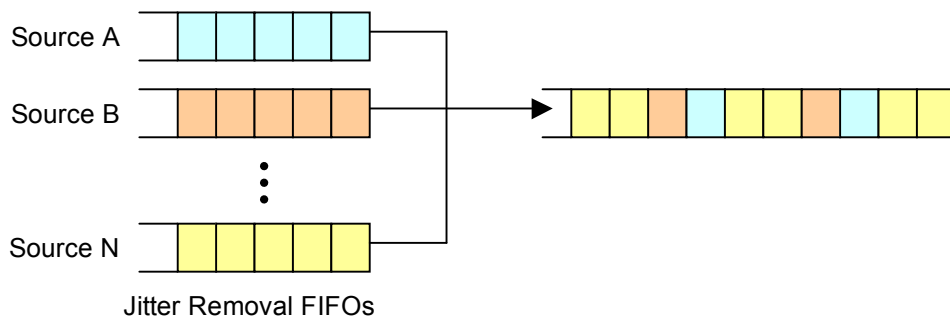


Figure 8. Demapping TDM packets into time slots

The estimated memory size requirement would be on the order of 512Kbytes per 10G flow. This would accommodate delay jitter over 3 fabric stages due to jumbo frame interference and segmentation/reassembly time (160bytes per 1G flow x 2) plus deskew (3 frames).

5.0 System Timing

All TDM flows must be time slot synchronized between ingress and egress blades. The packet fabric does not need to be synchronous as long as the ingress/egress interfaced on the blades are. To do this, the bit clock (REF Clock) and time slot clock (Frame Clock) can be derived from the carrier signals received on the Transport Blades. With two of these blades in the system, both active and backup clocks can be derived as shown in figure 9. The REF Clock can be used by the packet fabric for downstream bit clock timing in Synchronous Ethernet applications if IEEE 1588 support is required. The Frame Clock is used to time the transmission and align the reception of TDM packets across the packet fabric. At the receive side, jitter removal FIFOs are used to align the received packets and remove any latency jitter encountered in the packet fabric.

Application Note: FocalPoint TDM Support

At the ingress, all TDM packets have a frame time stamp inserted in the mapper unit at transmission. This time stamp is based on the frame clock and must be tightly synchronized across all ingress TDM Blades. The jitter removal FIFOs also use this time stamp to determine when to dequeue a received TDM frame in the demapper unit. A maximum jitter latency is assumed, and frames arriving earlier than this must wait in the FIFO until their proper dequeue time. The FIFO must be sized large enough to deal with the worst-case peak-to-peak jitter in the packet fabric. This jitter can be bounded and minimized using the advanced FocalPoint congestion management features discussed earlier.

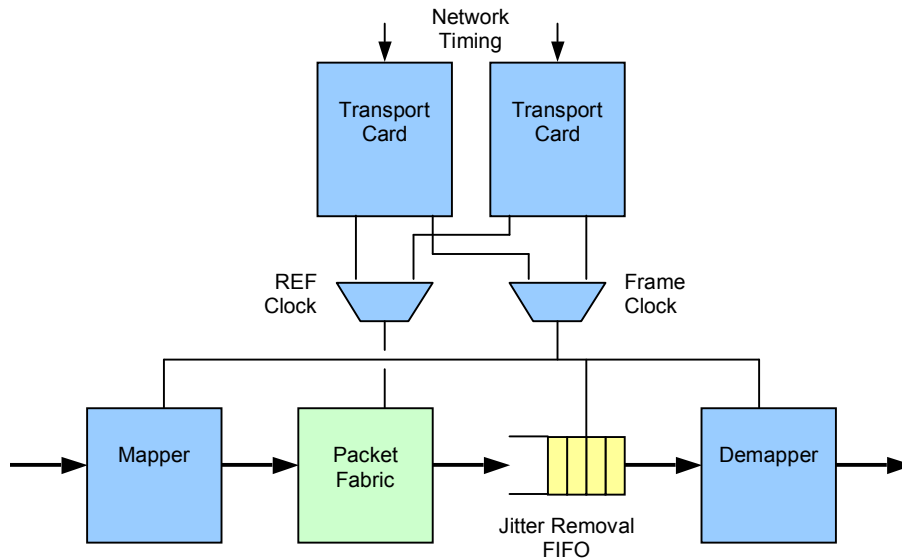


Figure 9. TDM system level timing

6.0 High Availability

All backplane fabrics using in telecom systems must provide high availability features. FocalPoint can provide 1+1, 1:1 or N:1 redundancy by distributing the load across multiple switch cards. Since TDM flows cannot tolerate data loss or latency spikes during failure recovery, an active-active approach must be used with the ingress switch multicasting the same stream across both switch cards. The egress Demapper then must decide which of the two received flows to accept based on link health and packet integrity. This requires a total of 2.5x speedup factor in the backplane.

Packet Blades can use the FocalPoint load balancing features to distribute flows across FocalPoint switches both switch cards. During a switch card failure, data can be automatically re-balanced across the remaining switches.