

# **PM9351, PM9352, PM9353, PM9354, PM9355, PM9311, PM9312, PM9313, PM9315**

## **TTX and ETT1 Chip Sets**

### **Designing Linecards that Work With Both ETT1 and TTX Switch Cores**

**PRELIMINARY**

**Issue 1: August 2000**

*PRELIMINARY*

*APPLICATION NOTE*

*PMC2001264*



*PM935X TTX, PM931X ETT1*

*ISSUE 1      DESIGNING LINECARDS THAT WORK WITH BOTH ETT1 AND TTX SWITCH CORES*

---

## CONTENTS

1	INTRODUCTION.....	1
2	OVERVIEW OF THE SWITCH CORE.....	2
2.1	ETT1 CHIP SET .....	2
2.1.1	LCS2-192 MODE (10 GBPS LINECARDS) .....	3
2.1.2	LCS2-48T4 MODE (2.5 GBPS LINECARDS).....	4
2.1.3	SYSTEM LATENCIES.....	4
2.1.4	CRC CALCULATION FOR THE LCS HEADER .....	6
2.2	TTX CHIP SET.....	6
2.2.1	LCS2-192 MODE (10 GBPS LINECARDS) .....	6
2.2.2	LCS2-48T4 MODE (2.5 GBPS LINECARDS).....	7
2.2.3	SYSTEM LATENCIES.....	7
2.2.4	CRC CALCULATION FOR THE LCS HEADER .....	9
3	ISSUES TO ADDRESS IN A COMMON INTERFACE DESIGN.....	9
3.1	MAPPING OF LCS2 QUEUES .....	10
3.2	LCS2 PROTOCOL PROCESSING .....	10
3.3	LATENCIES.....	12
3.4	CRC VALIDATION.....	12
3.5	TDM SERVICE.....	12

PRELIMINARY

APPLICATION NOTE

PMC2001264



PM935X TTX, PM931X ETT1

ISSUE 1    DESIGNING LINECARDS THAT WORK WITH BOTH ETT1 AND TTX SWITCH CORES

---

**LIST OF FIGURES**

FIGURE 1    A LINECARD INTERFACING WITH THE SWITCH CORE ..... 1

FIGURE 2    AN ETT1 PORT OPERATING IN LCS2-48T4 MODE..... 3

FIGURE 3    ALLOCATION OF THE TIMING BUDGET FOR AN ETT1 LINECARD TO  
RESPOND TO A GRANT FROM THE SWITCH CORE..... 5

FIGURE 4    FUNCTIONAL BLOCK DIAGRAM OF A TYPICAL SWITCH INTERFACE ..... 10

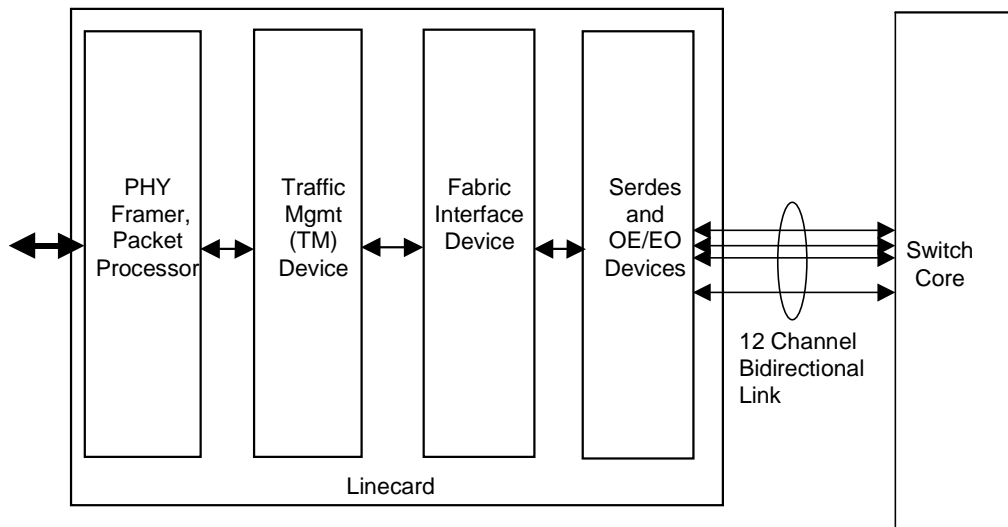
**LIST OF TABLES**

TABLE 1	HOLE REQUEST LATENCIES (IN OC-192C CELL-TIMES) FOR DIFFERENT PRIORITIES FOR VARIOUS CONFIGURATIONS OF TTX SWITCHES. ....	8
TABLE 2	REQUEST TO GRANT MINIMUM LATENCIES (IN OC-192C CELL-TIMES) FOR DIFFERENT PRIORITIES FOR VARIOUS CONFIGURATIONS OF TTX SWITCHES. ....	8
TABLE 3	NUMBER OF QUEUES AND REQUEST COUNTER PAIRS NEEDED ON AN LCS2-192 OR LCS2-48 LINECARD. ....	11

## 1 INTRODUCTION

This document presents the issues that need to be considered when designing an interface to a switch fabric that uses either the ETT1 (PM-9311/2/3/5) Chip Set or the TTX (PM-9351/2/3/4/5) Chip Set. This switch interface is typically located on the linecard between the traffic management device (or devices) and the set of Serdes (serializer/deserializer) and OE/EO converters as shown in Figure 1.

**Figure 1 A Linecard Interfacing with the Switch Core**



The main functions of the interface device will be:

- Allow the linecard to interface with the switch core
- Perform the Linecard-to-Switch (LCS) Version 2 protocol functions
- Provide mapping from the set of queues in the traffic management device to the Virtual Output Queues required by the switch core
- Perform a conversion from the bus format used by the traffic management device to the parallel bus format used by the Serdes devices

Linecards will interface to both the ETT1 and TTX switch cores via a three-way handshake as specified by the LCS protocol. The LCS protocol information is contained within an eight-byte header that is added to every cell. For more information on the LCS protocol, refer to the [LCS Protocol Specification - Protocol Version 2](#), available from PMC-Sierra, Inc.

The physical link between an OC-192c linecard and the switch core is comprised of twelve serial links, operated as an inverse multiplexed bundle. Each 72-byte cell is sliced across the 12 links, so that each serial link carries a 6-byte segment, which is 8b/10b encoded. Before actual data transfer occurs, each serial link must be independently bit-aligned, byte-aligned and frame-aligned using initialization sequences. Details of the physical link are available in the *ETT1 Databook*, available from PMC-Sierra, Inc. Mechanisms are provided to adjust for variations in transmission and reception speeds due to clock frequency differences on the linecard and switch core. The link technology operates at 1.5 Gbaud (150MHz).

This document presents a common switch interface design for ETT1 and TTX switch cores from the perspective of unicast and multicast best-effort service. Details of the TDM service in ETT1/TTX can be found in the *TDM Service in ETT1/TTX*, available from PMC-Sierra, Inc. However, there are certain incompatibilities in the way TDM service works for ETT1 and TTX. Anyone trying to design common linecards for TDM service should get in touch with application engineers within the Carrier Switch Division of PMC-Sierra, Inc.

## **2 OVERVIEW OF THE SWITCH CORE**

Since this document presents a switch interface design that works with both the ETT1 and TTX switches, we now present some of the key features of the ETT1 and TTX switch cores that affect this design.

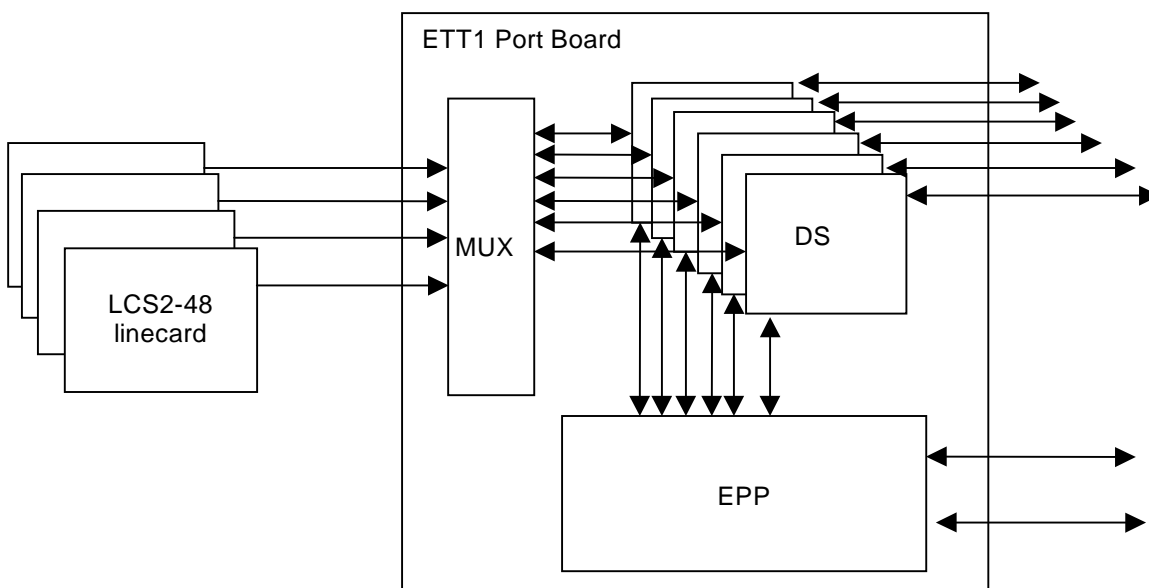
### **2.1 ETT1 Chip Set**

The ETT1 Chip Set provides an input-queued crossbar-based switch that can have up to 32 ports, each port operating at 10 Gbps, yielding an aggregate switching capacity of 320 Gbps. The LCS protocol is used to communicate between the linecard and the switch. The use of LCS makes it possible for the ETT1 switch to have a physical separation between the switch core and the linecards of up to 200ft/70m. The physical interface that has been implemented in the ETT1 Chip Set is designed to use off-the-shelf parts for the physical link. This interface provides a 1.5 Gbps serial link that uses 8b/10b encoded data. Twelve of these links are combined to provide a single LCS link operating at 18 Gbaud, providing an effective data bandwidth that is in excess of an OC-192c link. In the rest of this document, we shall refer to an OC-192c link running LCS as an LCS2-192 link, and the corresponding OC-48c link as an LCS2-48 link.

Each port in the ETT1 switch supports a linecard bandwidth in excess of 10 Gbps, which might be used by a single linecard (say OC-192c), or shared among up to four OC-48c ports, each of 2.5 Gbps. This latter mode, consisting of four 2.5 Gbps linecards, is referred to as the subport mode or the LCS2-48T4 mode. The physical interface into the switch core is always a single channel; in the subport mode, the four LCS2-48 (2.5 Gbps) channels are time multiplexed using strict TDM onto a single LCS2-48T4 (10 Gbps) channel. In other words, in the LCS2-48T4 mode, any four consecutive cells entering or leaving the switch core belong to different LCS2-48 channels. A separate customer-designed four-to-one multiplexer (MUX) device is required to manage the four LCS2-48 channels and merge them so they appear as a single LCS2-48T4 channel to the ETT1 port, as shown in Figure 2. A single ETT1 switch can have some ports in LCS2-192 mode, and some in LCS2-48T4 mode; there is no restriction. Both port configurations support four levels of prioritized best-effort traffic, unicast and multicast, as well as TDM traffic.



**Figure 2 An ETT1 Port Operating in LCS2-48T4 Mode**



In LCS2-48T4 mode, the switch must look at each incoming cell and determine which OC-48c subport has sent the cell. The LCS label field is used to achieve this. Two bits within the label (referred to as MUX bits in the *LCS Protocol Specification – Protocol Version 2*, available from PMC-Sierra, Inc.) are used to denote the source subport, numbered 0 through 3. The MUX bits must be inserted in the LCS label before the cell arrives at the switch port, and could be inserted by the source linecards themselves. Alternatively, the MUX device located between the subport linecards and the switch core could insert them. In this latter case the MUX device might not be able to re-calculate the LCS CRC. An ETT1 switch port can be configured to calculate the LCS CRC either with or without the MUX bits.

## **2.1.1 LCS2-192 Mode (10 Gbps Linecards)**

### **2.1.1.1 Unicast Traffic**

Every ETT1 port has a separate request counter and ingress queue for each (output port, priority) combination to which it can send unicast cells. In a full configuration, an ETT1 port can send to 32 ETT1 ports (including itself), at four priorities, and so will use 128 unicast request counters and ingress queues. The ingress queues are referred to as Virtual Output Queues (VOQs) as each queue holds cells going to just one output port. Every ETT1 port also has the same number of unicast egress queues, one queue for every (input port, priority) combination that can send cells to the port. The 128 unicast egress queues are referred to as Virtual Input Queues (VIQs) since each queue only holds cells that have come from a single input port. The transfer of cells from ingress to egress queues is lossless from a queuing perspective. The best effort unicast ingress and egress queues can each store up to 64 cells.

### **2.1.1.2 Multicast Traffic**

The ETT1 switch also supports multicast traffic. The LCS header will indicate that a cell is multicast, and it will also contain a multicast label. The ETT1 port uses this label to determine the multicast fanout, which is the list of destination ports. In the ETT1 switch, a multicast cell of a given priority will always have priority over a unicast cell of the same priority. Every ETT1 port has a single ingress queue and a single egress queue for multicast traffic at each priority, each of which can hold up to 96 cells. Each port also has a single multicast request counter at each priority.

### **2.1.2 LCS2-48T4 Mode (2.5 Gbps Linecards)**

There are some important differences between the LCS2-192 mode and the LCS2-48T4 mode. One difference is that the LCS header must now identify the OC-48c subport associated with each cell. Two bits in the LCS label field are used for this purpose. A second difference is that separate LCS request counters and separate egress queues must be maintained for each of the subports. So the number of queues increases four-fold in order to preserve the independence of each subport.

#### **2.1.2.1 Unicast Traffic**

The ETT1 switch port can support all four best-effort priorities when operating in LCS2-48T4 mode. Furthermore, an ETT1 switch can have some ports operating in LCS2-48T4 mode, while other ports operate in LCS2-192 mode. Suppose an ETT1 core is configured with all 32 ports in LCS2-48T4 mode, thereby supporting 128 OC-48c linecards. Since the ETT1 Chip Set uses virtual output queues, for each priority, the switch needs to keep track of unicast requests going to each one of the 128 possible output channels. So, 128 counters are needed for each ingress linecard, where each counter reflects the outstanding *requests* that one LCS2-48 ingress linecard has made to the switch. Since each port can support up to four LCS2-48 channels, a total of 512 request counters are needed for each priority.

#### **2.1.2.2 Multicast Traffic**

There is a single cell queue at each switch port, for multicast cells of each priority. Each LCS2-48 channel has its own request counter. On the egress side, there is a special multicast output queue for each subport. The multicast cells forwarded through the crossbar are queued in the order they arrive, enabling cells to be sent out-of-order between egress subports, but ensuring in-order delivery for a given egress subport.

### **2.1.3 System Latencies**

#### **2.1.3.1 Processing Time at the Linecard**

The linecard must respond to the grants from the ETT1 switch core within a certain period of time. In ETT1, the grant-to-cell latency is 39 OC-192c cell-times regardless of whether the port is in LCS2-192 or LCS2-48T4 mode. Figure 3 shows how the total time is divided between the ETT1 chips and other devices in the path. Since the latency from point A to point B in Figure 3 is 39 cell-times (or 1.56  $\mu$ s), the linecard designer needs to determine the latencies of the devices in the path and figure out how to meet this 39

cell-times round trip time. For example, assuming a 70m link between the linecard and the switch core, the fiber delay is roughly 15 cell-times\* (7.5 cell-times in each direction). If each Serdes and OE conversion together takes 1 cell-time, the linecard is left with 20 cell-times (800ns) to respond. For an LCS2-48 linecard, there is also a customer-designed MUX chip on the path, which results in additional latency.

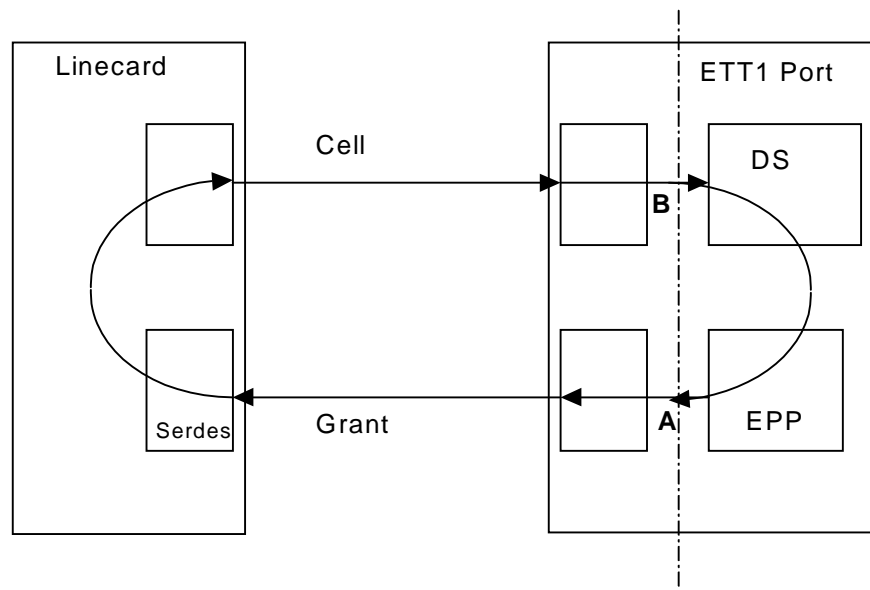
### 2.1.3.2 Hole Request Latency

The linecard can apply flow control to the ETT1 switch core by sending in hole requests. When a hole request is sent at a particular priority, the switch core responds to that request by ensuring that a future egress cell does not contain a valid cell at that priority. For ETT1, the switch core can take up to 64 cell-times to respond to the hole request (refer to the *ETT1 Databook*, available from PMC-Sierra, Inc.) Assuming four cell-times for the four Serdes and OE/EO conversions in the path, and another 15 cell-times for the fiber delay, from the perspective of the interface device, the response time for a hole request is at most 83 cell-times (or 3.32  $\mu$ s).

### 2.1.3.3 Request to Grant Minimum Latency

For ETT1, the minimum latency from when the switch port receives a LCS request from the linecard to when the LCS grant emerges from the port is 16 cell-times (640 ns).

**Figure 3 Allocation of the Timing Budget for an ETT1 Linecard to Respond to a Grant from the Switch Core**



\* The exact figure depends on the refractive index of the fiber, and should be calculated using the numbers obtained from the fiber vendor.

### **2.1.4 CRC Calculation for the LCS Header**

The LCS2 header includes a CRC-16 field (details about the fields covered and the polynomials used are available in the *LCS Protocol Specification - Protocol Version 2*, available from PMC-Sierra, Inc.). Prior to computation, the remainder is set to 0xFFFF, and then the CRC-16 computation is done over all 64 bits of the header with the CRC-16 field set to all zeroes. However, at an ETT1 port, the CRC-16 remainder is truncated to the 8 least significant bits before transmission to the linecard. So, at the linecard, the computation must also be done over the header with the CRC-16 field set to all zeroes, and the least significant 8 bits of the result should be compared with the least significant 8 bits of the header's CRC-16 field. The 8 most significant bits of the CRC-16 field should be ignored.

The ETT1 ports can also be configured to include or ignore (mask to 0) the MUX bits of the Request Label and/or the Grant Label. Checking of the ingress header and generation of the egress header can be configured independently.

## **2.2 TTX Chip Set**

Like the ETT1 Chip Set, the TTX Chip Set is also a single-stage crossbar-based switch with a centralized scheduler. The switch can have up to 256 LCS2-192 ports, or 1024 LCS2-48 ports or 64 LCS2-768 ports at OC-768c rates, or any mix that results in an aggregate switching capacity of up to 2.56 Tbps. The external interface is still LCS, and hence backward compatible with the ETT1 Chip Set. In TTX, the physical separation between the switch core and the linecards can be up to 1000ft /300m. The external interface is comprised of 12 channels that will either operate at 1.5 Gbps or 2.5 Gbps depending on the capabilities of the linecard, the switch core and the size of the cell payload. The 1.5 Gbps interface is the one that is backward compatible with the ETT1 Chip Set, and so is the only one considered in this document.

Each port in the TTX switch can support a linecard bandwidth in excess of 10 Gbps, and hence as in the ETT1 switch, this bandwidth can be used to interface to a single LCS2-192 linecard or four LCS2-48 linecards. There are up to four levels of prioritized best-effort traffic, unicast and multicast, as well as a TDM class of service.

The most significant difference between the ETT1 and TTX switch architectures relates to the presence of cell queues in the interior of the switch. In ETT1, the ingress port processor has an elaborate VOQ-based storage for incoming cells; in the TTX architecture, the cells are all stored on the linecard while the switch core does not store any except for egress OC-48 cells. The LCS protocol, via the three-way handshake, allows the linecard to send in a cell exactly when the crossbar has been configured for it to be transmitted to the egress port.

### **2.2.1 LCS2-192 Mode (10 Gbps Linecards)**

#### **2.2.1.1 Unicast Traffic**

Every TTX port has a separate request counter for each (output port, priority) combination to which unicast cells can be sent. In a full configuration, a TTX port can send to 256 ports (including itself), at four

priorities, and so requires 1024 unicast request counters. As mentioned above, there are no VOQs to store incoming cells. The switch core sends back grants to the linecard for pending requests which causes the linecard to send the corresponding cells into the switch just in time for the cells to pass through the crossbar to the egress side. Unlike ETT1, for LCS2-192 ports, there are no unicast egress queues (VIQs) either. Cells that traverse the crossbar are forwarded directly to the egress linecard.

### **2.2.1.2 Multicast Traffic**

The TTX switch also supports multicast traffic. The LCS header will indicate that an incoming cell is multicast, and it will also contain a multicast label. The TTX switch internally uses this label to determine the list of ports to which the cell should go, i.e., the multicast fanout. Every TTX port has a single FIFO for multicast traffic at each priority that stores the multicast labels of the arriving requests. In TTX, the multicast cells of a given priority will always have priority over unicast cells of the same priority class.

## **2.2.2 LCS2-48T4 Mode (2.5 Gbps Linecards)**

### **2.2.2.1 Unicast Traffic**

In the LCS2-48T4 mode, each TTX port has a separate request counter for each (output subport, priority) combination to which it can send unicast cells. Also, each ingress LCS2-48 subport will have a separate set of counters for each (output subport, priority) combination. In a full configuration, a TTX port can send to 256 ports (including itself) or 1024 subports, at four priorities, and since each port can support up to four ingress LCS2-48 channels, a total of 4096 unicast counters are needed per priority. As mentioned above, there are no VOQs to store incoming cells. However, 16 VIQs are maintained one for each subport at each priority, to temporarily buffer OC-48 cells.

### **2.2.2.2 Multicast Traffic**

The LCS2 header indicates that a cell is multicast, and it will also contain a multicast label. The TTX switch internally uses this label to determine the multicast fanout. Each ingress LCS2-48 channel has a single FIFO for multicast traffic at each priority that stores the multicast labels of the arriving requests.

## **2.2.3 System Latencies**

### **2.2.3.1 Processing Time at the Linecard**

The linecard must respond to the grants from the TTX switch core within a certain period of time. In TTX, the grant-to-cell latency is 256 OC-192c cell-times for an LCS2-192 port, and 128 OC-192c cell-times for an LCS248T4 port. Assuming a 300m link between the linecard and the switch core, the fiber delay is roughly 75 cell-times (37.5 cell-times in each direction). If the four Serdes and OE conversions in the roundtrip path together take up to four cell-times, the linecard is left with 177 cell-times (7.08  $\mu$ s) to respond in LCS2-192 mode. For an LCS2-48 linecard, the linecard has 49 cell-times less the latency due to the additional customer-designed MUX chip on the path.

### **2.2.3.2 Hole Request Latency**

As in ETT1, in TTX, the linecard can apply flow control to the switch core by sending in hole requests. When a hole request is sent at a particular priority, the switch core responds to that request by ensuring that an egress cell does not contain a valid cell at that priority. For TTX, in LCS2-192 mode, the hole request latency varies with the priority and switch configuration as shown in Table 1. For LCS2-48T4 mode, however, it is 306 cell-times for all priorities and all configurations. Using the same numbers as those in Section 2.2.3.1, for a 128-port TTX switch, from the perspective of the linecard, the response time for a hole request is about 385 cell-times for LCS2-48T4 mode and between 457 cell-times (priority 3) and 575 cell-times (priority 0) in LCS2-192 mode.

**Table 1 Hole Request Latencies (in OC-192c cell-times) for Different Priorities for Various Configurations of TTX Switches.**

	64 Port Switch	128 Port Switch	256 Port Switch
Priority 0	390	496	714
Priority 1	377	453	620
Priority 2	364	418	526
Priority 3	351	378	432

### **2.2.3.3 Request to Grant Minimum Latency**

For TTX, the minimum latency from when the switch port receives an LCS request from the linecard to when the LCS grant emerges from the port varies with the priority and switch configuration as shown in Table 2.

**Table 2 Request to Grant Minimum Latencies (in OC-192c cell-times) for Different Priorities for Various Configurations of TTX Switches.**

	64 Port Switch	128 Port Switch	256 Port Switch
Priority 0	102	208	426
Priority 1	89	165	332
Priority 2	76	130	238
Priority 3	63	90	144

### **2.2.4 CRC Calculation for the LCS Header**

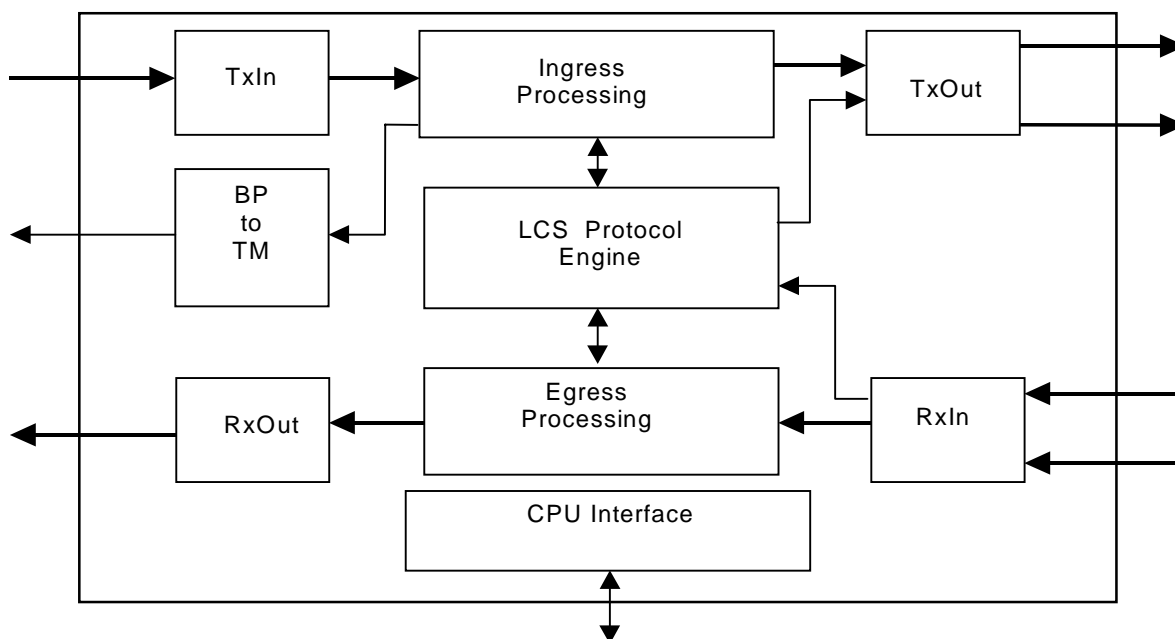
A TTX port uses the standard CRC-16 computation as specified in the *LCS Protocol Specification - Protocol Version 2*, available from PMC-Sierra, Inc. All 16 bits of the CRC are transmitted by the switch port to the linecard, unlike an ETT1 port that truncates the most significant 8 bits. Just like ETT1, the TTX ports can also be configured to include or ignore (mask to 0) the MUX bits of the Request Label and/or the Grant Label. Checking of the ingress header and generation of the egress header can be configured independently.

## **3 ISSUES TO ADDRESS IN A COMMON INTERFACE DESIGN**

Several issues need to be considered when designing a linecard that will work with both ETT1 and TTX Chip Sets. As described in the previous section, the switch port can interface with a single LCS2-192 linecard or with four LCS2-48 linecards. Switch compatibility requirements for each implementation, an LCS2-192 linecard and an LCS2-48 linecard, are covered here. The common switch interface would be typically placed between the traffic management (TM) device (or devices) and the set of Serdes (serializer/deserializer) and OE/EO converters on the linecard.

When the port is configured in the LCS2-48T4 mode, i.e., is interfacing with four LCS2-48 linecards, a separate customer-designed multiplexer (MUX) chip on the ETT1 or TTX port board is used to aggregate the LCS2-48 streams into a single, time-multiplexed 10 Gbps stream. Alternatively, the MUX chip could also reside on a linecard supporting four LCS2-48 streams, and the LCS interface could be located downstream of the MUX device, directly handling an LCS-192 stream which is comprised of four LCS-48 streams that are time-division-multiplexed.

Both LCS2-192 and LCS2-48 switch interfaces would have the same basic architecture as shown in Figure 4, but be slightly different in the details. In order to transmit up to a 10 Gbps stream of data to the switch fabric, an LCS-192 switch interface will receive one cell every 40 ns from the TM device, which is equivalent to 25 Mpps, and will transmit one cell every 40ns to the switch core. For LCS2-48 linecards the cells will be arriving at each linecard (and departing) every 160ns. This will process enough cells so that the four LCS2-48 streams can be time-multiplexed together and maintain the 10 Gb/s (25 Mpps) LCS2-48T4 stream to the ETT1 or TTX switch fabric. Using Figure 4 as a reference, a typical interface would operate as follows. In the ingress direction, the switch interface would receive cells from the TM device via the TxIn block, store and process them in the Ingress Processing block and then send them to the switch core via the TxOut block which would slice the outgoing cells over the 12 links to the Serdes devices. In the egress direction, the interface would receive the cells from the switch core sliced across 12 links via the RxIn block and combine them into cells, do some processing on them in the Egress Processing block, and then send them to the TM device via the RxOut block. The LCS Protocol Engine would control the LCS protocol processing done in both the ingress and egress directions

**Figure 4 Functional Block Diagram of a Typical Switch Interface**


### **3.1 Mapping of LCS2 Queues**

One of the primary functions of the switch interface is to provide a set of queues that map the queuing structure in the TM device to the queuing structure required by the switch core. As seen from Table 3, an LCS2-192 interface that can talk to both ETT1 and TTX switch cores would require 1024 unicast queues (256 for each priority), 4 multicast queues (1 for each priority), at least 1 TDM queue and 1 queue for control packets. When cells arrive from the TM device, the information in the header is used to select one of the queues to store the payload. Similarly, for an LCS2-48 linecard, the switch interface will have to support 4096 unicast queues (1024 for each priority), 4 multicast queues (1 for each priority), at least 1 TDM queue and 1 queue for control packets.

### **3.2 LCS2 Protocol Processing**

Due to the difference in the number of ports supported, the usage of the some fields in the LCS header is different in ETT1 and TTX, and linecard designers should consult the *LCS Protocol Specification – Protocol Version 2*, from PMC-Sierra, Inc., to ensure compatibility. The switch interface has to perform the following functions related to the LCS protocol:

- Maintain request counters.
- Initiate insertion of LCS Control packets like Start/Stop and Request Count packets in the cell flow to the switch core.
- Enable or disable data request generation upon detecting Start/Stop packets in the egress direction.



- Deal with error conditions arriving from request/grant errors.
- Provide a programmable counter to insert idle cells towards the switch core. This is needed for rate adaptation between the linecard and the switch core due to PPM differences between the respective clocks.

**Table 3 Number of Queues and Request Counter Pairs Needed on an LCS2-192 or LCS2-48 Linecard.**

	LCS2-192		LCS2-48T4	
	ETT1	TTX	ETT1	TTX
Max number of ports in the switch, N	32	256	128	1024
Unicast Queues/Request Counter Pairs	$32*4 = 128$	$256*4 = 1024$	$128*4 = 512$	$1024*4 = 4096$
Multicast Queues/RequestCounter Pairs	$1*4 = 4$	$1*4 = 4$	$1*4 = 4$	$1*4 = 4$
TDM Queues/Request Counter Pairs	1 or more	up to 1024	1 or more	up to 1024
Request Counter Width	10 bit	10 bit	8 bit	10 bit

The switch interface needs to maintain state for all the virtual output queues. It does this by having counters that mirror the queue structure. Each queue corresponding to unicast, multicast and control packets requires two counters. The first counter, called *Waiting\_Req*, tracks the requests that are waiting for grants from the switch core while the second counter, *Outstanding\_Req*, tracks the requests that have not yet been submitted to the switch core. For each multicast queue, a *Label\_Queue* is used to store the multicast labels for all the generated requests sent and outstanding, in FIFO order. Similarly, two *Frame Tables* are used to maintain the TDM tags to slots mapping for *Frame 0* and *Frame 1*. The *Waiting\_Req* needs to be 10 bits wide to work with both ETT1 and TTX switch cores.

When a cell from the TM device arrives to be forwarded to the switch core, the corresponding *Outstanding\_Req* counter is incremented. If it is a multicast cell, then the multicast label is also stored in the corresponding *Label\_Queue*. To generate a request, one of the queues, with a non-zero *Outstanding\_Req* value and a *Waiting\_Req* value less than the maximum, is selected based on a given priority/fairness algorithm. The corresponding *Waiting\_Req* counter is then incremented and the *Outstanding\_Req* counter decremented. The cell's *Label* is then obtained – for a unicast request it is generated while for a multicast packet it is obtained from the *Label\_Queue*. If the *Waiting\_Req* count reaches the maximum count, then no more requests are generated for that queue, whereas if the *Outstanding\_Req* count reaches zero, then no more requests can be generated until there are new cell arrivals.

When a grant arrives from the switch core, the LCS Protocol Engine decrements the count in the *Waiting\_Req* counter. If the grant is for a multicast queue, then the label at the head of the corresponding *Label\_Queue* is compared to the *Label* in the grant; if there is a mismatch an exception condition is generated. If it is for a unicast or control queue and the corresponding *Waiting\_Req* value is zero, then an error has occurred, and an exception condition is generated. If there is no error, then the queue number and the sequence number are sent to the Ingress Processor block.

Periodically, LCS Request Count control packets for unicast queues are sent to the switch core. This allows the linecard and the switch to recover from losing synchronization. For multicast queues, a Request Count packet is only generated when a grant error occurs. When a Request Count packet has to be inserted, the LCS Protocol Engine will provide the Ingress Processing block, the Label and the Count to be inserted in the packet.

### **3.3 Latencies**

The grant to cell latency requirement is more stringent in the ETT1 switch, and hence that is the one that should be used for the design of a common switch interface. The maximum latency for ETT1 is 39 cell-times. As discussed in Section 2.1.3, this leaves about 800ns for the linecard to send a cell in response to a grant. On the other hand, the hole request latency is much larger for the TTX switch, and therefore it is one that should be used. The specific numbers are tabulated in Section 2.2.3. Also, the LCS request to grant minimum latencies are larger for TTX (See Section 2.24) and hence those are the ones that should be used.

### **3.4 CRC Validation**

The switch interface needs to check the integrity of the incoming data by validating the CRC on it and also include CRC on the cells going out to the switch core. To be compatible with both ETT1 and TTX, the switch interface would have to calculate the CRC-16 checksum across the 64 bits of the header with CRC field set to all zeros as described in the *LCS Protocol Specification - Protocol Version 2*, available from PMC-Sierra, Inc. The result is then inserted into the CRC field in the LCS header. For validating cells coming from the switch core, in ETT1 only the 8 least significant bits of the CRC-16 field are valid, while in TTX the entire CRC-16 field should be used.

### **3.5 TDM Service**

This document presents a common switch interface design for ETT1 and TTX switch cores from the perspective of unicast and multicast best-effort service. Details of the TDM service in ETT1/TTX can be found in the *TDM Service in ETT1/TTX*, available from PMC-Sierra, Inc. However, there are certain incompatibilities in the way TDM service works for ETT1 and TTX. Anyone trying to design common linecards for TDM service should get in touch with application engineers within the Carrier Switch Division of PMC-Sierra, Inc.

**CONTACTING PMC-SIERRA, INC.**

PMC-Sierra, Inc.  
105-8555 Baxter Place Burnaby, BC  
Canada V5A 4V7

Tel: (604) 415-6000

Fax: (604) 415-6200

Document Information: [document@pmc-sierra.com](mailto:document@pmc-sierra.com)  
Corporate Information: [info@pmc-sierra.com](mailto:info@pmc-sierra.com)  
Application Information: [apps@pmc-sierra.com](mailto:apps@pmc-sierra.com)  
(604) 415-4533  
Web Site: <http://www.pmc-sierra.com>

© 2000  
PMC-Sierra, Inc.  
105 - 8555 Baxter Place  
Burnaby BC Canada V5A 4V7  
Phone 604.415.6000 FAX 604.415.6200

This document is for the internal use of PMC-Sierra, Inc. and PMC-Sierra, Inc. customers only.

In any event, no part of this document may be reproduced in any form without the express written consent of PMC-Sierra, Inc.

PMC-2001264 Issue date: August 2000