PM73487

QRT

622 Mbps ATM Traffic Management Device

DATASHEET

Released
Issue 3: JUN 1999

U.S. Patents 5,557,607, 5,570,348, and 5,583,861

Public Revision History

Issue Number	Issue Date	Details of Change
Issue 1	March 1998	Creation of Document
Issue 2	October 1998	This data sheet includes: registers added in B version of device: RX_QUEUE_ENGINE_TEST - bits 26:16 TX_QUEUE_ENGINE_TEST - bits 22:15 QUEUE_ENGINE_CONDITION_PRES_BITS QUEUE_ENGINE_CONDITION_LATCH_BITS QUEUE_ENGINE_INT_MASK RX_LOWER16_SCG_CONFIG RX_LOWER16_SCG_STATE RX_LOWER32_SCG_CONFIG RX_LOWER32_SCG_STATE RX_LOWER48_SCG_CONFIG RX_LOWER48_SCG_STATE TX_LOWER4_SCG_CONFIG TX_LOWER4_SCG_STATE TX_LOWER8_SCG_CONFIG TX_LOWER8_SCG_STATE TX_LOWER12_SCG_CONFIG TX_LOWER12_SCG_STATE Updated RX_SERVICE_TABLE
Issue 3	June 1999	Production Release Version

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

Table of Contents

List of Figures

List of Tables

Product Overview

The PM73487 622 Mbps ATM Traffic Management Device (QRT™) is an advanced communications device capable of supporting very large, high-performance ATM switching systems. The rich feature set of the QRT enables systems to offer many sophisticated network services. The QRT provides 622 Mbps UTOPIA (Level 1 or Level 2) access to switch fabrics composed of PM73488 5 Gbps ATM Switch Fabric Elements (QSEs). Together, these devices can be used to build architectures with capacities from 622 Mbps to 160 Gbps. The QRT can also act as a stand-alone 622 Mbps switch.

The QRT/QSE architecture virtually eliminates head-of-line blocking by means of the QRT's per-Virtual Channel (VC) receive queues and congestion feedback from the QSE™ switch fabric. The distributed architecture acts as an output-buffered switch by incorporating Evil Twin Switching™ (a congestion-reducing routing algorithm in the switch fabric) and a speed-up factor in the switch fabric (running the fabric faster than the line rate).

The QRT uses per-VC receive queues, 64 receive Service Classes (SCs), and 16 transmit SCs per each of the 31 Virtual Outputs (VOs) to enable flexible multi-priority scheduling algorithms. The scheduler can be used to ensure Quality-of-Service (QoS) guarantees for Constant Bit Rate (CBR), Variable Bit Rate (VBR), and Unspecified Bit Rate (UBR) VCs. The QRT also provides five separate congestion thresholds, each with hysteresis, that selectively control AAL5 Early Packet Discard (EPD) and/or Cell Loss Priority (CLP)-based cell dropping for UBR support. Additional highlights of the QRT include full Virtual Path Indicator (VPI)/Virtual Channel Indicator (VCI) header translation, separate input and output cell buffers (up to 64K each), Virtual Path (VP)/VC switching, and support for up to 16K VCs on both the receive and transmit sides.

PMC-Sierra also offers the QRT Device Control Package, which is a software package that harnesses the QRT's rich feature set and shortens system development times.

FEATURES

QUEUING ALGORITHMS

Receive

- Maintains 64 weighted, bandwidth-controlled SCs with per-VC queues.
- Provides round-robin servicing of queues within each SC.
- Provides per-channel (VP or VC), per-SC, and per-direction congested and maximum queue depth limits.
- Provides up to 64K cell buffers.

Transmit

- Provides 31 VOs.
- Maintains 16 SCs for each VO with per-VC accounting.

- Provides per-channel (VP or VC), per-SC Queue (SCQ), per-SC, per-VO, and per-direction congested and maximum queue depth limits.
- Provides up to 64K cell buffers.

CONGESTION MANAGEMENT ALGORITHMS

- Supports EPD and Partial Packet Discard (PPD) for UBR traffic, and as a backup for ABR traffic.
- Supports CLP-based cell discard and Explicit Forward Congestion Indicator (EFCI) cell marking.
- Supports three congestion limits (as well as EPD, CLP, and EFCI, and/or backpressure) for logical multicast on the transmit side.

SWITCHING

- Supports VC and VP switching.
- Supports up to 16K VCs.

ADDRESS MAPPING

- Supports all 12 VP and 16 VC bits through use of a double, indirect lookup table.
- Performs header translation at both the input (receive) and output (transmit) directions. Input header translation is used to pass the output queue channel number through the switch.

MULTICAST

- Supports logical multicast with a superior queue-clearing algorithm.

DIAGNOSTIC/ROBUSTNESS FEATURES

- Checks the header parity.
- Counts tagged cells.
- Runs error checks continually on all fabric lines.
- Checks liveness of control signal lines at both switch fabric and UTOPIA interfaces, working around partial fabric failures.
- Checks Static Random Access Memory (SRAM) and Dynamic Random Access Memory (DRAM) parity.

STATISTICS FEATURES

- In the receive direction, counts cells transmitted and dropped.
- In the transmit direction, counts cells transmitted and dropped on a per-VC basis.

I/O FEATURES

- Provides four switch element interfaces with phase aligners. The phase aligners allow for external serialization of the data stream enabling systems to be built that support device separation of up to 10 meters.
- Provides a UTOPIA Level 2 Multi-PHY (MPHY) 16-bit, 50 MHz interface.
- Provides a 2-level priority servicing algorithm for high and low bandwidth UTOPIA PHY layer devices.
- Provides a multiplexed address/data CPU interface.

- Provides two 100 MHz, 32-bit, synchronous DRAM cell buffer interfaces.
- Provides three 100 MHz, synchronous SRAM control interfaces.
- Provides a JTAG boundary scan interface.

COMPATIBILITY FEATURES

- Compatible with the ATM Forum 3.0, 3.1, and 4.0 specifications.
- Compatible with the ATM Forum UTOPIA Level 1 and Level 2 specifications.
- Compatible with the PM73488 ATM QSE.

PHYSICAL CHARACTERISTICS

- 3.3 V supply voltage.
- 5 V tolerant inputs on the microprocessor and UTOPIA interfaces.
- Available in a 503-pin Enhanced Plastic Ball Grid Array (EPBGA) package.

BLOCK DIAGRAM

Figure 1 shows a QRT system block diagram.

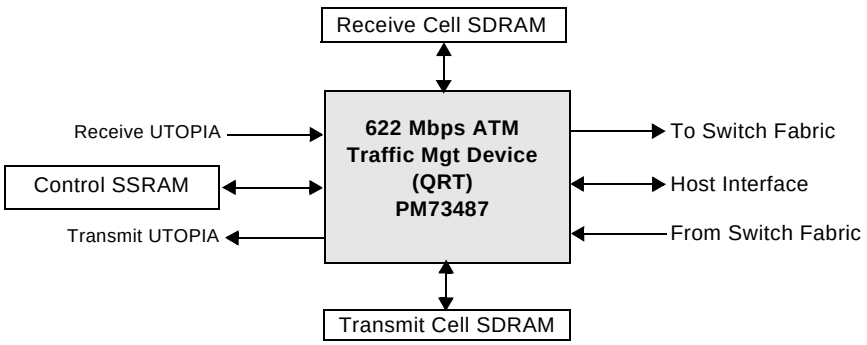

Figure 1. QRT System Block Diagram

1 SYSTEM APPLICATIONS

The QRT, together with the QSE, support a wide range of high-performance ATM switching systems. These systems range in size from 622 Mbps to 160 Gbps. The systems can be developed such that this scalability is provided with linear cost. Another key feature of the QRT/QSE architecture is that it is exceptionally fault-tolerant, both in the switch fabric and the UTOPIA interface.

This section contains a quick overview of the QRT and several example applications:

- a stand-alone 622 Mbps switch using a single QRT,
- a 5 Gbps switch using QRTs and a QSE,
- a 10 Gbps switch using QRTs and QSEs,
- a switch architecture using QRTs and QSEs that scales from 5 Gbps to 20 Gbps,
- a switch architecture using QRTs and QSEs that scales from 5Gbps to 160 Gbps

1.1 QRT System Overview

The QRT provides 622 Mbps of input and output buffered access to switch fabrics composed of QSEs (32 x 32 PM73488s). In addition, the QRT supports a stand-alone, purely output-buffered 622 Mbps switch mode. Head-of-line blocking, commonly associated with input buffers, is virtually eliminated via per-VC receive queues, three types of per-cell switch fabric feedback, and per-VC cell selection algorithms. The QRT also provides eight separate congestion thresholds, each with hysteresis, that selectively control AAL5 Early Packet Discard (EPD)/Packet Tail Discard (PTD), CLP-based cell dropping, and/or EFCI marking. Eight separate maximum thresholds are also supported. Additional highlights of the QRT include full VPI/VCI header translation, separate input and output cell buffers (up to 64K each), Virtual Path Connection (VPC)/Virtual Channel Connection (VCC) connections, and up to 16K VCs. The QRT provides a bidirectional connection between a UTOPIA Level 2 interface and 4-nibble wide, 66 MHz switch fabric interfaces, as shown in Figure 2 on page 5. A significant switch speed-up factor, up to 1.6 times the line rate, is used to support full throughput for many switch fabric configurations.

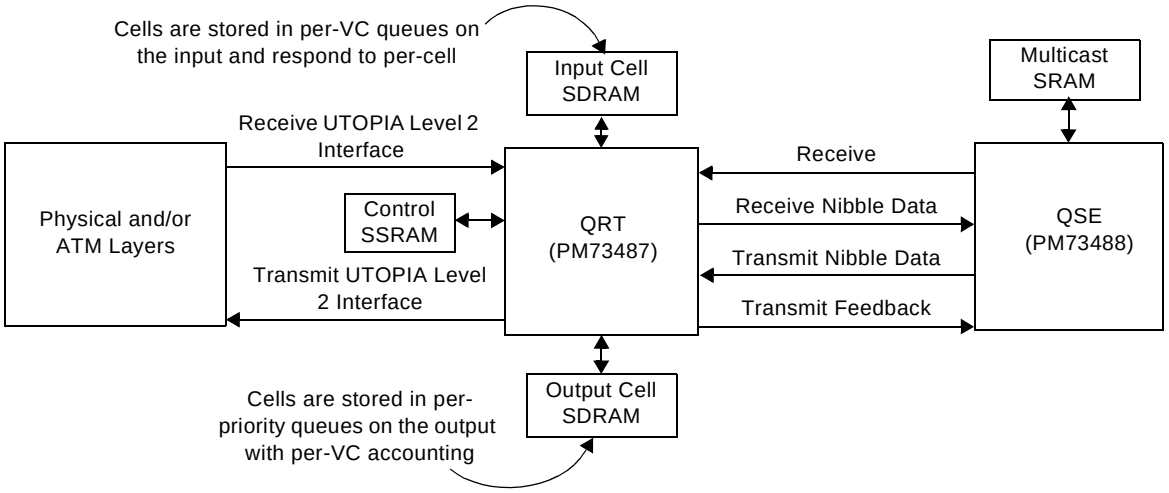

Figure 2. QRT System Overview

1.2 622 Mbps Switch Configuration

The QRT can be used in a stand-alone application that supports ATM switching up to 622 Mbps, as shown in Figure 3. The four switch fabric interfaces are looped back to the QRT, allowing the UTOPIA interface to be fully used. In this application, the QRT operates as an output buffered switch..

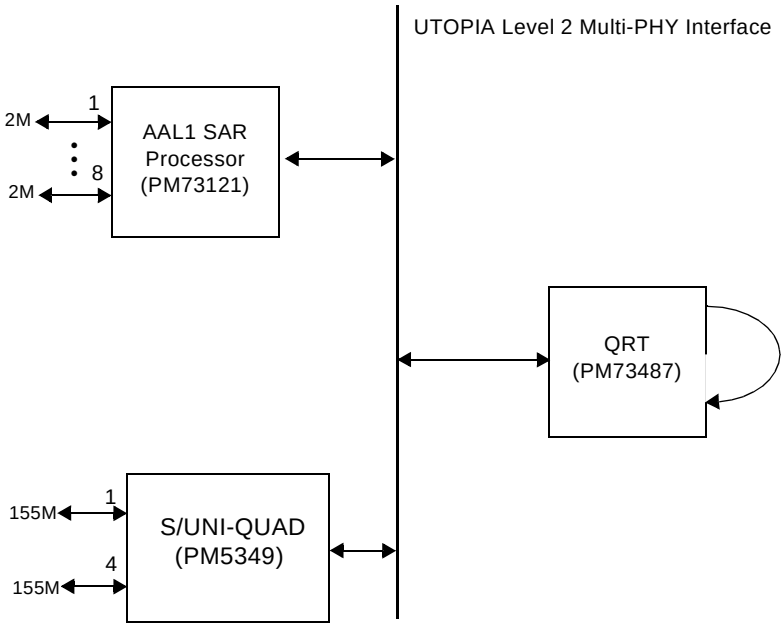

Figure 3. 622 Mbps Switch Configuration

1.3 32 x 32 Switch Application (5 Gbps)

Figure 4 shows a basic 32 × 32 switch application (5 Gbps) using eight QRTs and one QSE.

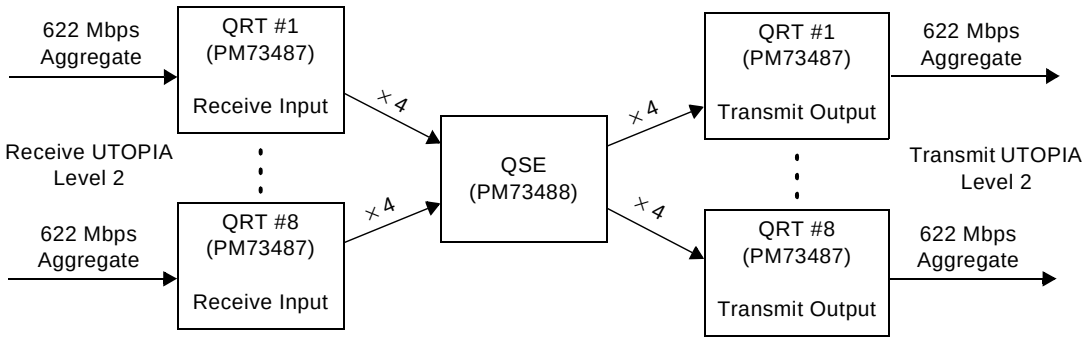

Figure 4. 32 x 32 Switch Application (5 Gbps)

1.4 64 x 64 Switch Application (10 Gbps)

Figure 5 shows a 64 × 64 switch application (10 Gbps) using 16 QRTs and 6 QSEs. This application uses QSEs in a 3-stage fabric. This sized system can be implemented in a single 19-inch rack.

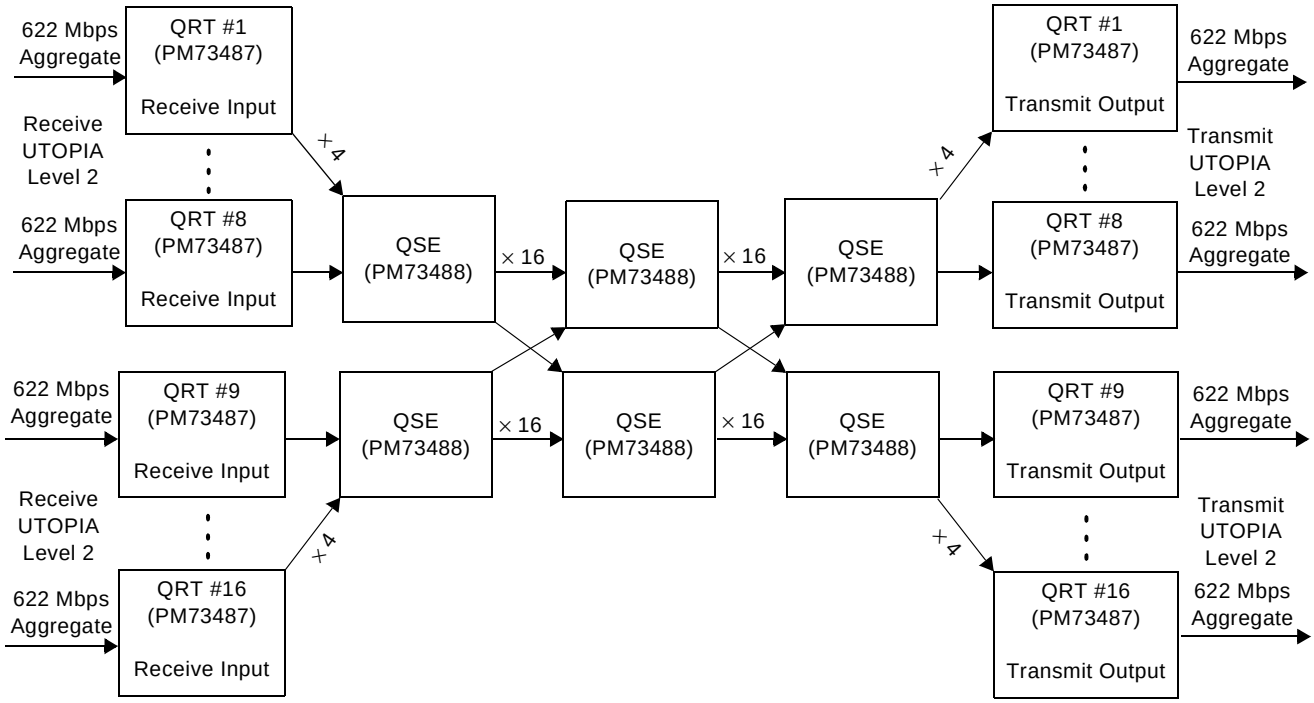

Figure 5. 64 x 64 Switch Application (10 Gbps)

1.5 5 Gbps to 20 Gbps Application Example - Seamless Growth

This section illustrates the modularity of the QRT (PM73487) and QSE (PM73488) architecture. A 5 Gbps system can immediately be created (as shown in Figure 6 on page 7), and then be upgraded to 10 Gbps (as shown in Figure 7 on page 7), or 20 Gbps (as shown in Figure 8 on page 8). Since all these systems are based on a single-stage switch fabric, the per-port cost for each system will remain the same.

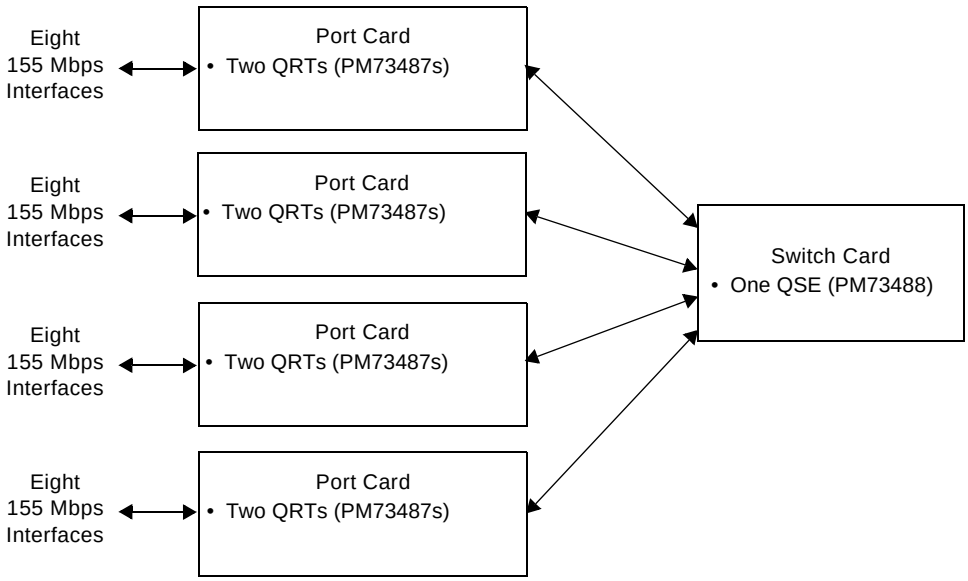

Figure 6. 5 Gbps ATM Switch Using 16 Dual S/UNIs, 8 QRTs, and 1 QSE

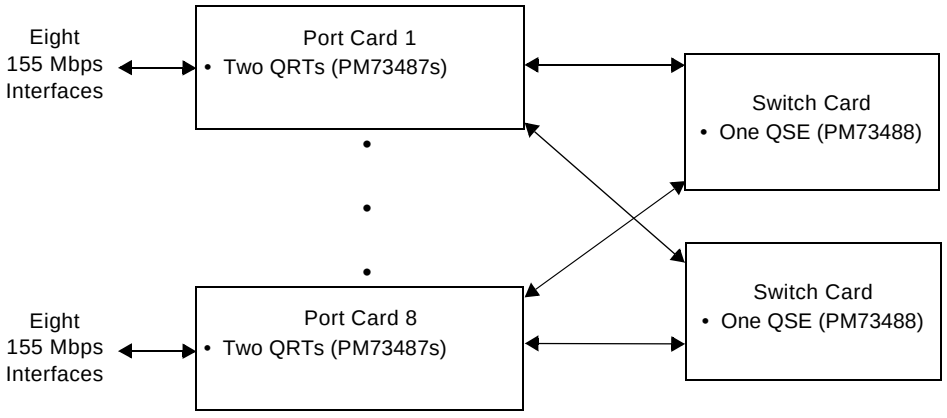

Figure 7. 10 Gbps ATM Switch Using 32 Dual S/UNIs, 16 QRTs, and 2 QSEs

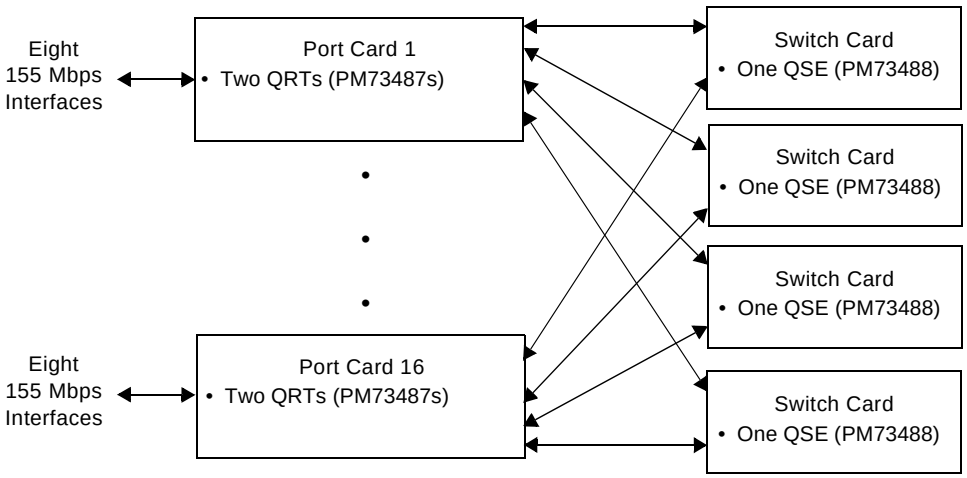

Figure 8. 20 Gbps ATM Switch Using 64 Dual S/UNIs, 32 QRTs, and 4 QSEs

1.6 5 Gbps to 160 Gbps Application Example – LAN-to-WAN

A powerful application of the QRT and the QSE devices is the creation of modules that can be used in a range of switches with only the interconnection changing between different sizes. ATM switches from 5 Gbps to 160 Gbps can be realized with only two unique cards. A port card has one QRT, and a switch card has two QSEs. The switch fabric consists of three stages, each with 32 QSEs (or 16 switch cards). To plan for future scalability, the middle stage must be built-in upfront. This is a one-time cost. Then, in order to scale in 5 Gbps increments, one switch card and its accompanying eight port cards should be added. Finer bandwidth scaling is possible by populating the additional switch card with port cards as needed (in increments of 622 Mbps). With this switch fabric topology, scaling is possible up to 160 Gbps. Once the initial middle stage cost has been incurred, the per-port cost for 5 Gbps through 160 Gbps systems remains almost constant

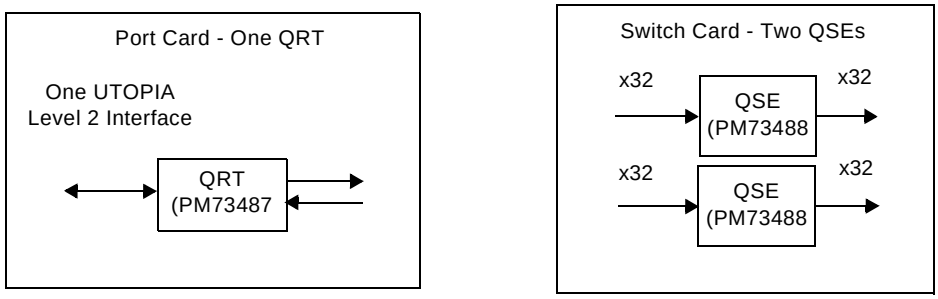

Figure 9. 5 Gbps to 160 Gbps Switches Modeled Using Only Two Cards

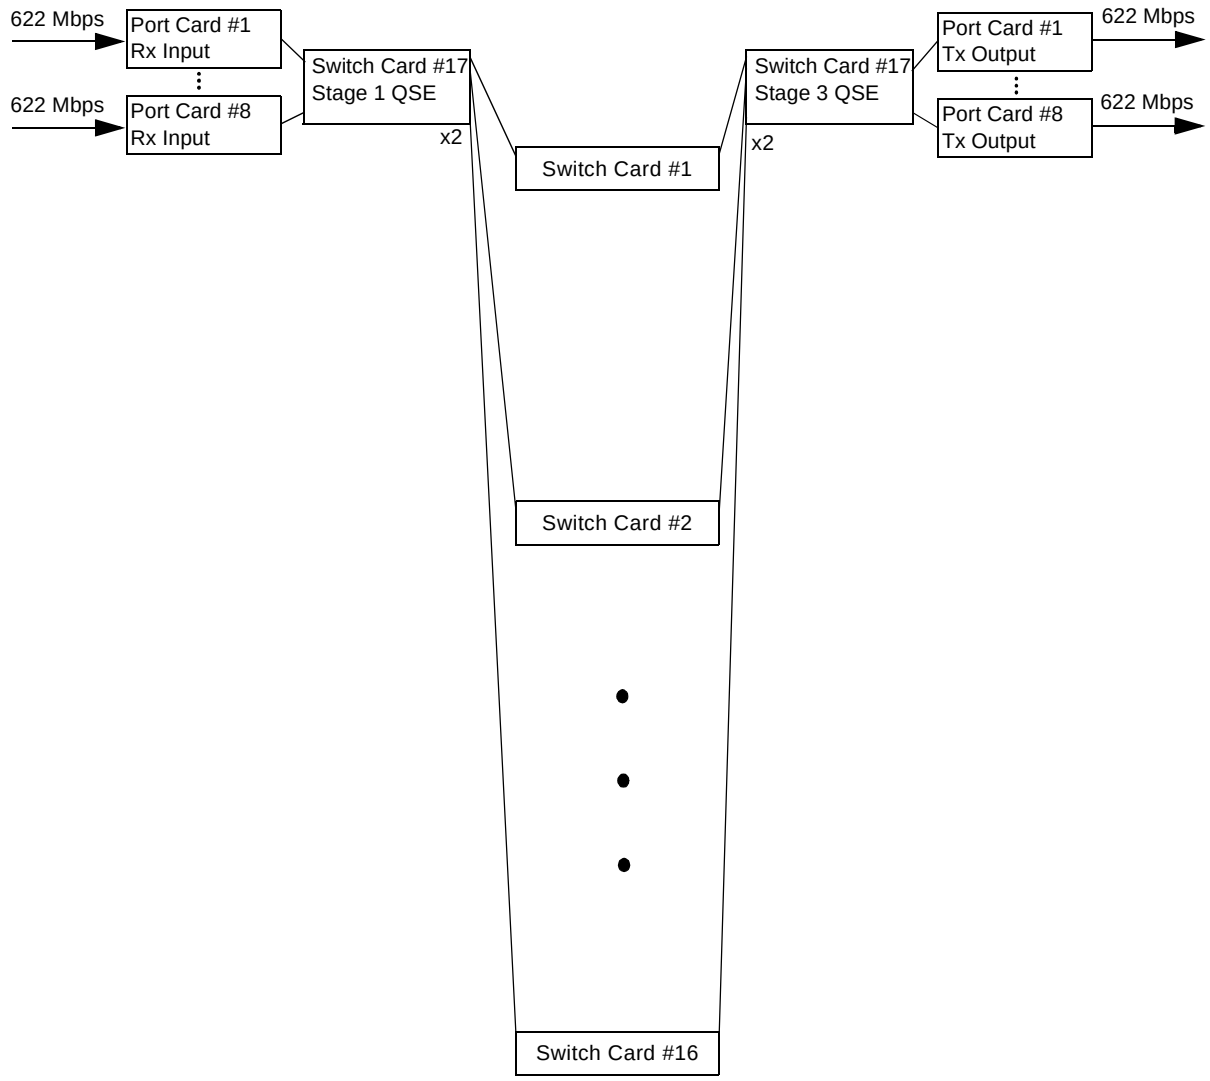

Figure 10. 5 Gbps ATM Switch

Figure 10 shows a 5 Gbps ATM switch using 8 port cards (8 QRTs) and 17 switch cards (34 QSEs). The middle stage is composed of 16 switch cards. The 5 Gbps bandwith is achieved by adding switch card #17 (which is depicted using two boxes: one stage 1 QSE and one stage 3 QSE), and eight port cards (each of which is depicted using two boxes: one for the Rx input side, and one for the Tx output side). Lines between stage 1 and stage 2, and stage 2 and stage 3 switch cards represent two sets of wires, one to each of the QSEs in the middle stage switch cards.

.Figure 11 shows a 10 Gbps ATM switch using 16 port cards (16 QRTs) and 18 switch cards (36 QSEs). Here, another switch card and eight port cards have been added to the 5 Gbps switch depicted in Figure 10.

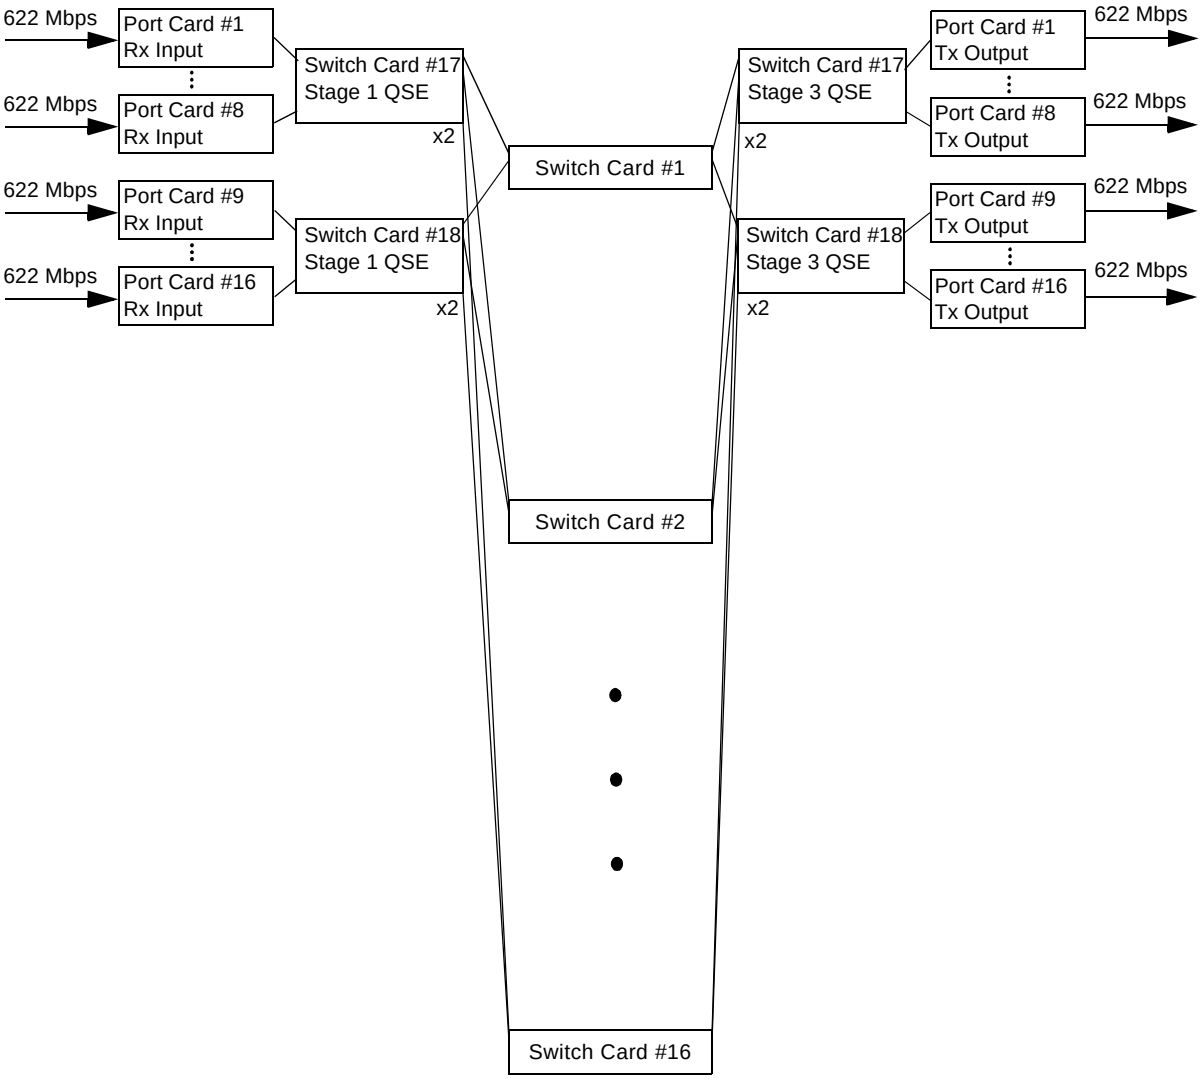

Figure 11. 10 Gbps ATM Switch

Figure 12 shows a 15 Gbps ATM switch using 24 port cards (24 QRTs) and 19 switch cards (38 QSEs).Here, once again, another switch card and eight port cards have been added

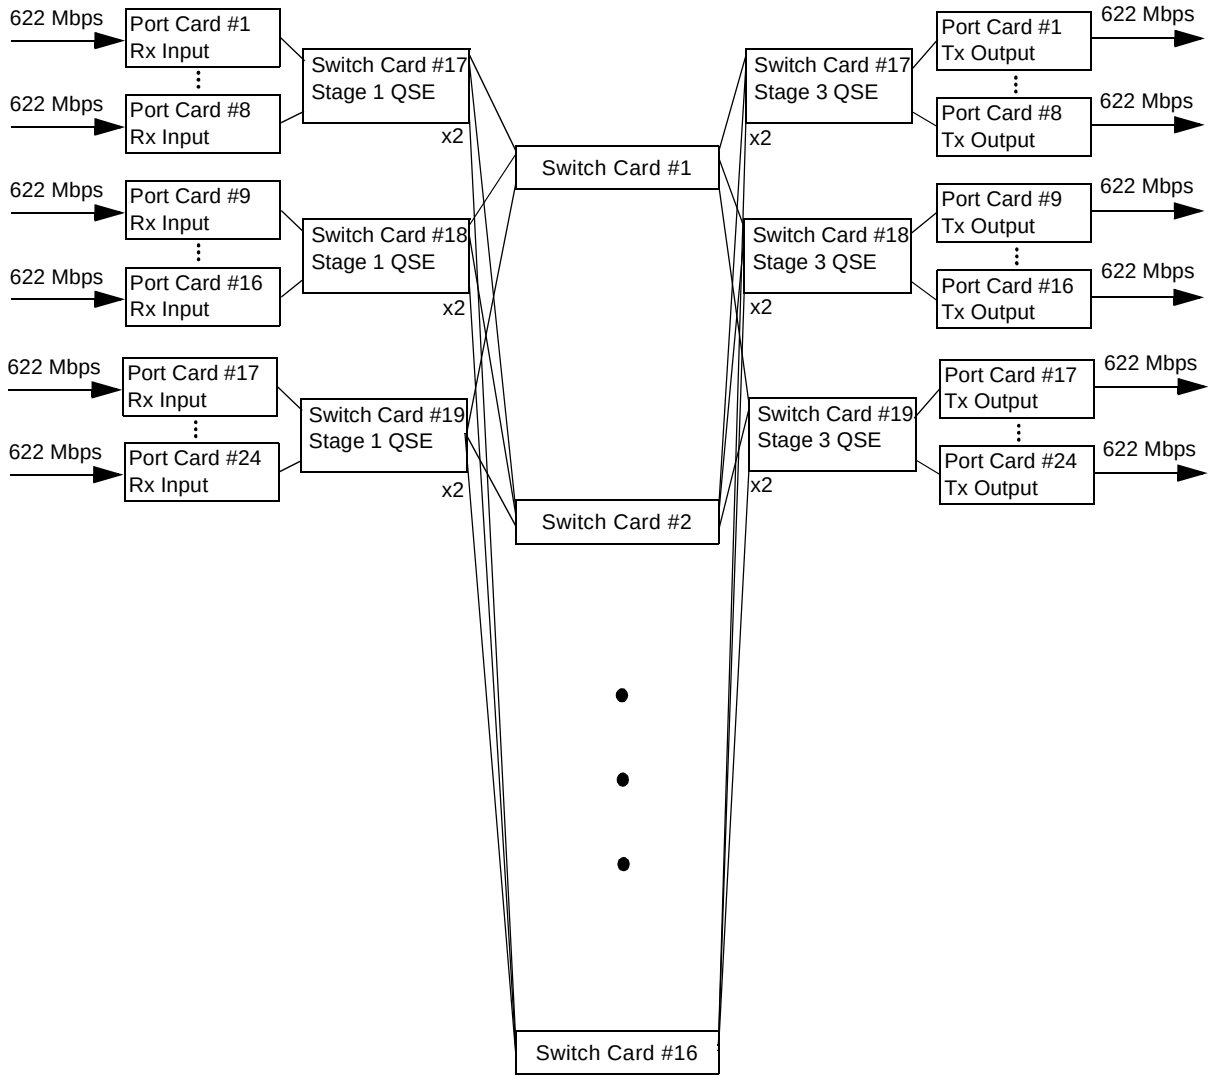

Figure 12. 15 Gbps ATM Switch

Figure 13 shows a 20 Gbps ATM switch composed of 32 port cards (32 QRTs) and 20 switch cards (40 QSEs). By adding additional sets of a switch card and eight port cards in the same manner, this system can scale up to 160 Gbps. .

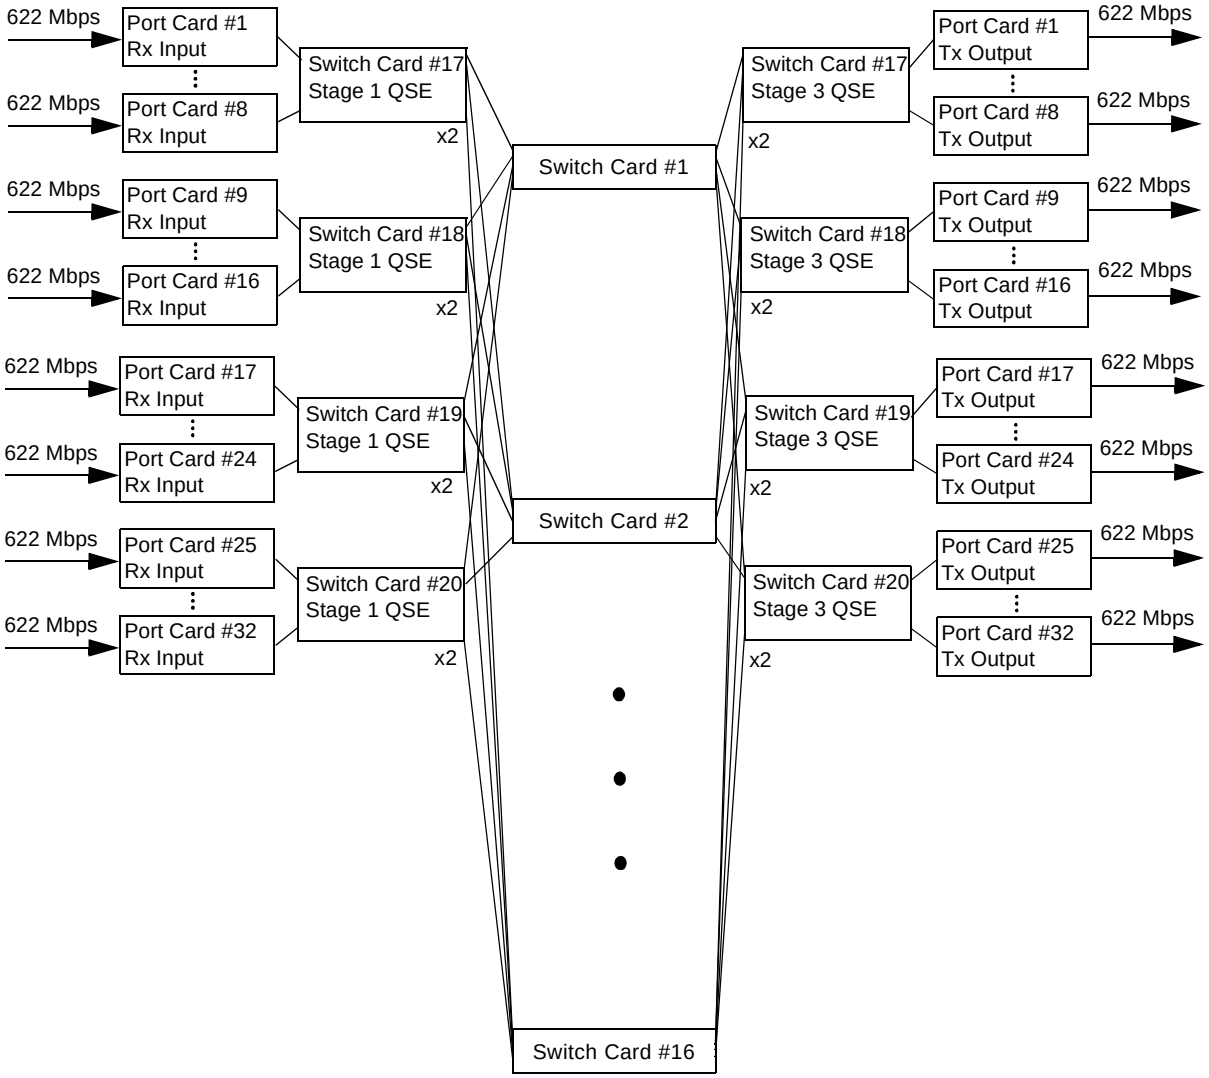

Figure 13. 20 Gbps ATM Switch

.

2 THEORY OF OPERATIONS

2.1 Overview

The QRT is a 622 Mbps, full duplex, intelligent routing table which, when used with a switch fabric composed of either SE or QSE devices, can implement ATM switches from 622 Mbps to 160 Gbps. The QRT supports a 16-bit UTOPIA Level 2 interface for ease of connection to PHY or AAL layer devices. Four nibble-wide data interfaces connect the QRT to the switch interface. External DRAM memory devices provide receive and transmit cell buffering, and external SRAM devices provide control data for the QRT. This section explains the algorithms for the data flow.

Figure 14 shows an overview of the QRT system.

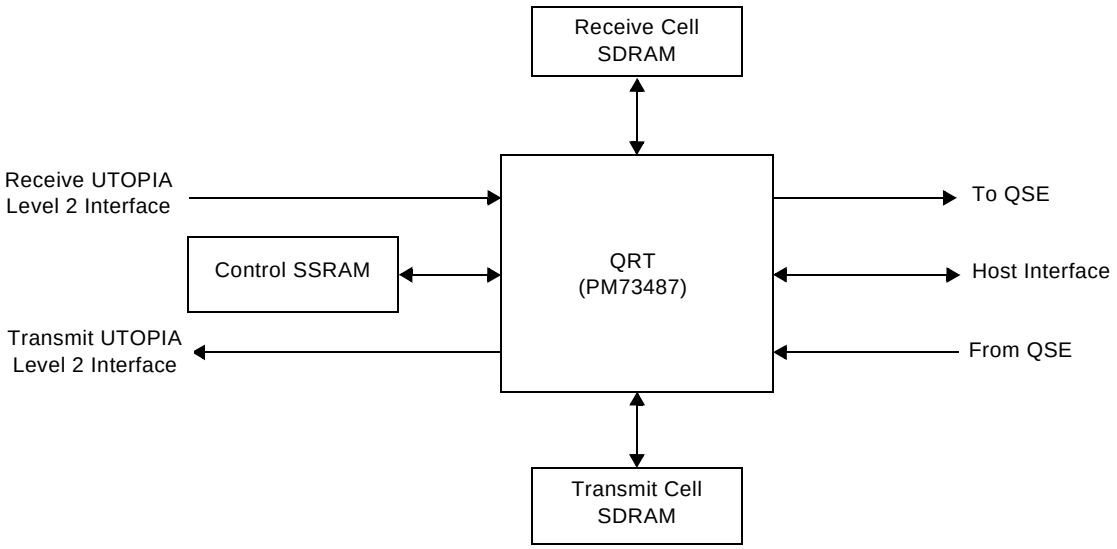

Figure 14. QRT System Overview

2.2 Interface Descriptions

2.2.1 Switch Fabric Interface

The QRT switch fabric interface consists of four groups of signals from both the ingress (receive side) and the egress (transmit side). Each group consists of a Start-Of-Cell (SE_SOC_OUT) signal, a nibble-wide data bus, and a backpressure acknowledgment (BP_ACK_IN) signal. The Start-Of-Cell (SE_SOC_OUT) signal is transmitted at the ingress at the same time as the beginning of a cell. SE_SOC_OUT on the ingress is common to all four groups. The BP_ACK_OUT signal flows from the egress through the switch fabric, in the direction opposite the data, and indicates whether a cell has successfully passed through the switch fabric. Other signals associated with the switch fabric interface are the switch element clock (SE_CLK) and RX_CELL_START. To support the highest possible throughput for various switch fabric configurations, a clock speed-up factor of 1.6 is used. That is, the switch fabric is run at a rate that is effectively 1.6 times faster than the line rate.

2.2.2 Phase Aligners

Phase aligners are used to allow for extended device separation. The technique used is a clock recovery mechanism that requires only the switch fabric to be frequency synchronous. A master clock is distributed to all devices associated with the switch fabric, and the phase of the clock at each interface is dynamically adjusted to account for skew introduced to the signals. The phase aligner circuitry for each interface responds to the cell start and feedback signals, which contain a high number of transitions to ensure accurate phase adjustment of the clock for data and signal sampling.

2.2.3 UTOPIA Interface

The QRT's UTOPIA interface implements the ATM Forum standardized 16-bit, Level 2 configuration, which supports up to 31 Virtual Outputs (VOs) via five address bits. Up to 31 PHY or AAL layer devices with 16-bit UTOPIA Level 2 functionality can be connected to this interface, providing full duplex throughputs of 675 Mbps.

2.2.4 Cell Buffer SDRAM Interface

The QRT supports two Synchronous DRAM (SDRAM or SGRAM) interfaces providing up to 64K of cell buffering in both the receive and transmit directions. Each interface consists of a 32-bit data bus, a 9-bit address bus, two chip select signals, and associated control signals. The frequency of these interfaces is 100 MHz. Both Synchronous Graphic RAM (SGRAM) and SDRAM devices are supported. Clocking for these two interfaces is provided through the device.

2.2.5 Channel RAM (CH_RAM) Interface

The QRT supports up to 16K channels through a Synchronous SRAM (SSRAM) interface. The interface consists of a 32-bit data bus, a 16-bit address bus, and associated control signals. The frequency of this interface is 100 MHz. Clocking for this interface is provided through the device.

2.2.6 Address Lookup RAM (AL_RAM) Interface

The QRT has data structures in the AL_RAM, including VPI/VCI address translation. The interface consists of a 6-bit data bus, a 17-bit address bus, and associated control signals. The frequency of this interface is 100 MHz. Clocking for this interface is provided through the device.

2.2.7 AB_RAM Interface

The QRT stores the per VC head / tail pointers and sent / dropped counters for the receive direction in the AB_RAM. Each interface consists of a 17-bit multiplexed address/data bus and associated control signals. The frequency of this interface is 100 MHz.

2.2.8 Host Processor Interface

The QRT host processor interface allows connection of a microprocessor through a multiplexed 32-bit address/data bus. The suggested microprocessor for this interface is the Intel i960®. The microprocessor has direct access to all of the QRT control registers.

2.2.9 SE_SOC Encodings

The SE_SOC and BP_ACK signals have guaranteed transitions and special encodings, which are defined in this section and in "BP_ACK Encodings" which follows. The SE_SOC_IN and SE_SOC_OUT signals have guaranteed transitions and SOC encodings as shown in Figure 15. The SE_SOC signals carry a repeating pattern of four zeros and four ones to guarantee transitions required by the phase aligner. The "Start-Of-Cell" on the data lines associated with an SE_SOC line is indicated by a change in this pattern. For a valid SE_SOC, the change in pattern is followed by reset of the background pattern such that it is followed by four zeros and four ones. The first nibble (PRES) of the header is coincident with SE_SOC (change in pattern).

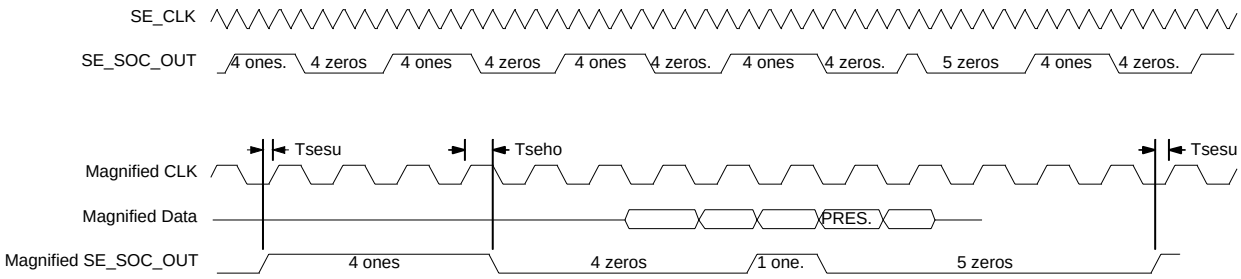

Figure 15. SE_SOC Encodings

2.2.10 BP_ACK Encodings

Figure 16 shows the BP_ACK encodings.

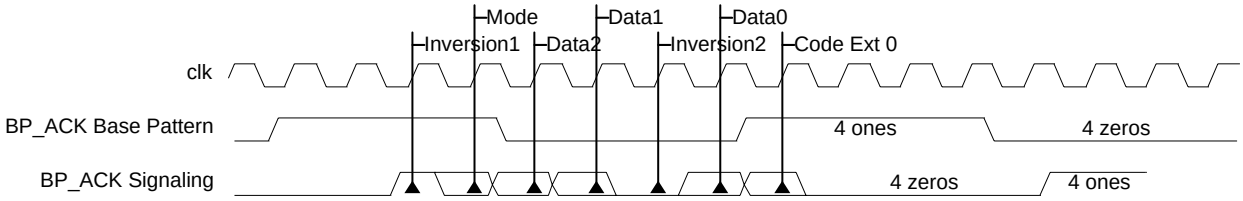

Figure 16. BP_ACK Encodings

The BP_ACK_IN and BP_ACK_OUT signals have guaranteed transitions, and BP and ACK encodings. The BP_ACK signal is used to signal backpressure/cell acknowledgment to the fabric (QSE) at the egress and receive backpressure/cell acknowledgment at the ingress from the fabric (QSE).

To ensure the transitions required by the phase aligner, the BP_ACK signal carries a repeating four zeros, four ones pattern. The actual information is transferred through encoded 7-bit packets that start with a change in this background pattern. The change (an inversion) on the line is followed by a mode bit, followed by two bits of coded message, and a second inversion (inverse of the first inversion). If it is an acknowledgment packet, this is followed by two bits of code exten-

sion (these bits are for future use and currently are required to be “00”). In the case of a backpressure packet, the next bit is the backpressure bit on for low priority multicast cells, followed by one code extension bit.
The background is reset to four zeros and four ones after transmission of each packet.
The QRT and QSE allow back-to-back acknowledgment and backpressure packets. In the case of back-to-back acknowledgment and backpressure packets, the receiving device may see an inverted bit (a “1”) followed by the rest of the packet instead of a reset background pattern.
One backpressure packet and either one or none acknowledgment packet are expected to be received during a cell time. The receipt of multiple acknowledgment or backpressure packets is a failure condition.
Table 1 describes the backpressure and acknowledgment encodings.

Table 1. Backpressure and Acknowledgment Encodings

Mode	Data 2	Data 1	Data 0	Code Ext 0	Description
0	1 = Backpressure on high priority multicast cell.	1 = Backpressure on medium priority multicast cell.	1 = Backpressure on low priority multicast cell.	0	Backpressure information. This signal is present each cell time, regardless of whether a cell was transmitted or not (on that link). This signal is withheld if any problem is detected on the input port.
1	0	0	0	0	Signals no response. Treated as acknowledgment.
1	0	1	0	0	Signals Mid-switch Negative ACKnowledgment (MNACK).
1	1	0	0	0	Signals Output Negative ACKnowledgment (ONACK).
1	1	1	0	0	Signals ACKnowledgment (ACK).

2.2.11 Relation Between External CELL_START and Local CELL_START

Figure 17 shows the relationship between external RX_CELL_START and local CELL_START signals.

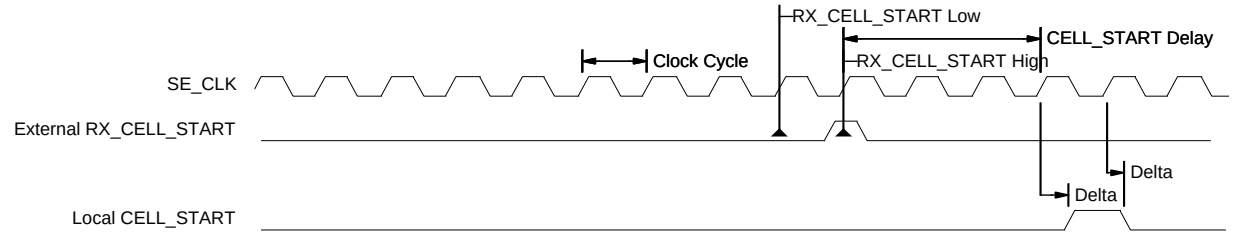

Figure 17. QRT Cell-Level Timing

Delay between the external RX_CELL_START and local CELL_START is programmable through the RX_CELL_START_ALIGN register (refer to "RX_CELL_START_ALIGN (Internal Structure)" on page 122).

The local CELL_START impacts the start of cell transmission to the fabric. It also determines the period within a cell time during which the BP_ACK_IN(3:0) at ingress is valid. As such, the programmable CELL_START delay allows the flexibility to synchronize the QRTs and QSEs in a system.

2.3 Cell Flow Overview

The QRT functions as a 622 Mbps port for an ATM switch fabric composed of either the SE or QSE devices. The QRT transfers cells between a UTOPIA Level 2 interface and a switch fabric interface. The device supports header translation and congestion management. The basic flow of cells through the QRT is as follows (see Figure 18 on page 19):

1. A cell enters the QRT on the receive side from the UTOPIA interface and the channel number is looked up.
2. The cell is then either dropped or transferred to the receive cell buffer DRAM and queued in the receive queue controller depending on six congestion management checks (both maximum and congested thresholds for the device, Service Class Group (SCG), SC, and connection).
3. When an available cell time occurs, four cells are selected by the receive-side scheduler, which reads the cells from the receive cell buffer DRAM and transmits them from the QRT into the switch fabric.
4. Once a cell is received from the switch fabric on the transmit side, it is again either dropped or transferred to the transmit cell buffer DRAM and queued in the transmit queue controller, depending on ten congestion management checks (both maximum and congested thresholds for the device, VO, SC, Service Class Group (SCG),Service Class Queue (SCQ), and connection).
5. When the cell is selected for transmission by the transmit-side scheduler, it is removed from the transmit cell buffer DRAM and processed by the transmit multicast/header mapper for corresponding header translation and distribution.

6. The cell then is sent to the UTOPIA interface and exits the QRT on the transmit side.

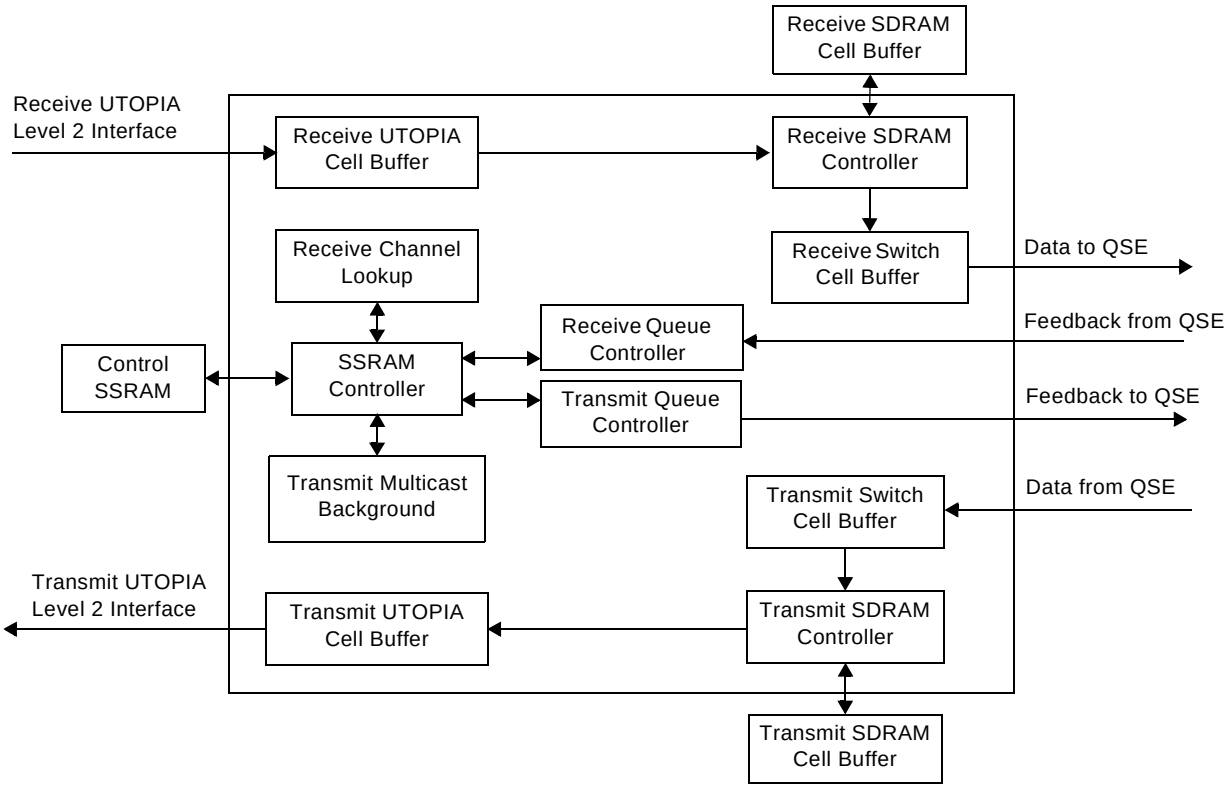

Figure 18. QRT Data Flow Diagram

2.4 UTOPIA Operation

2.4.1 General

Cells received from the UTOPIA interface are first processed by the receive channel lookup block and then queued for transmission within the receive queue controller. The cell waits in the receive cell buffer DRAM for instruction from the receive queue controller to proceed to the switch fabric interface.

2.4.2 UTOPIA Interface

The QRT interfaces directly to a UTOPIA interface device without needing an external FIFO. The receive side UTOPIA has a 4-cell internal FIFO, and the transmit side contains another 4-cell internal FIFO. The QRT UTOPIA interface is 16 bits wide and operates at frequencies up to 50 MHz. It provides the following modes:

- UTOPIA Level 1 single-PHY interface
- UTOPIA Level 2 multi-PHY interface

2.4.2.1 UTOPIA Level 2 Polling

The UTOPIA interface offers three modes of polling, as per the UTOPIA Level 2 specification:

- Standard single cell available polling
- Multiplexed Status Polling (MSP) using four cell available signals
- Direct status indication using four cell available signals

These polling modes allow the QRT to communicate with many different PHY devices. Figure 19 shows the QRT polling PHY devices in a receive UTOPIA operation.

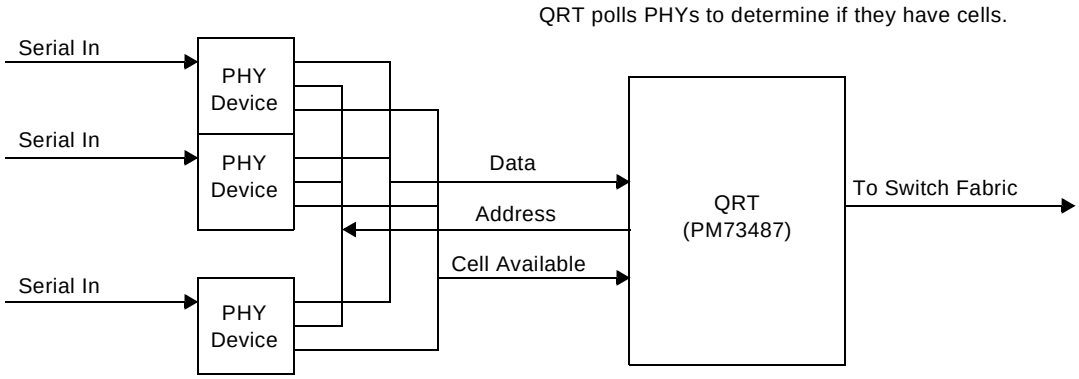

Figure 19. Receive UTOPIA Operation

Figure 20 shows the QRT polling PHY devices in a transmit UTOPIA operation.

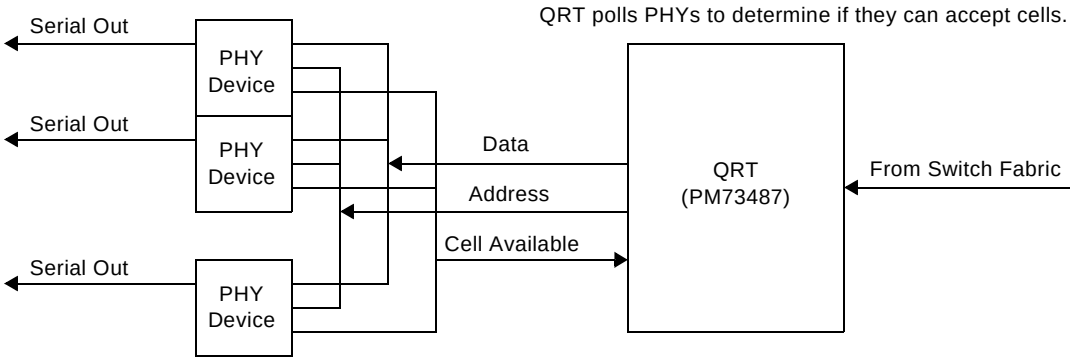

Figure 20. Transmit UTOPIA Operation

2.4.2.1.1 Standard Single Cell Available Polling

In the standard single cell available polling mode, one cell available response occurs every two clocks. Figure 21 shows the receive standard single cell available polling.

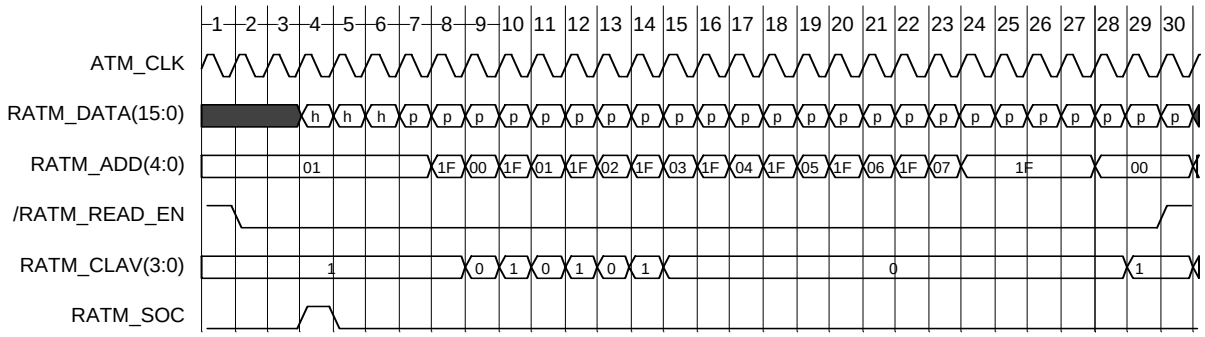

Figure 21. Receive Standard Single Cell Available Polling

Figure 22 shows the transmit standard single cell available polling.

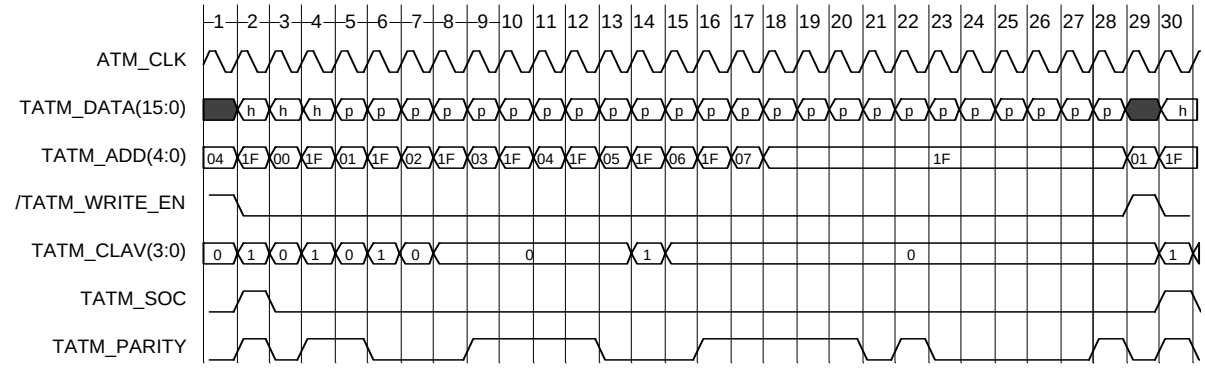

Figure 22. Transmit Standard Single Cell Available Polling

2.4.2.1.2 Multiplexed Status Polling (MSP) Using Four Cell Available Signals

With MSP using four cell available signals, up to four cell available responses occur every two clocks. The advantage offered by the MSP mode is the improved response time for PHY service selection. With this method, it is possible to poll 31 devices in a single cell time. PHY devices, however, must comply with this optional part of the UTOPIA Level 2 specification. A standard PHY device can be configured to use this mode even though it does not support it directly. To effect this, up to eight PHY devices can be configured with the addresses 0, 4, 8, 12, 16, 20, 24, and 28. Figure 23 shows the receive UTOPIA 50 MHz MSP, including cell transfer.

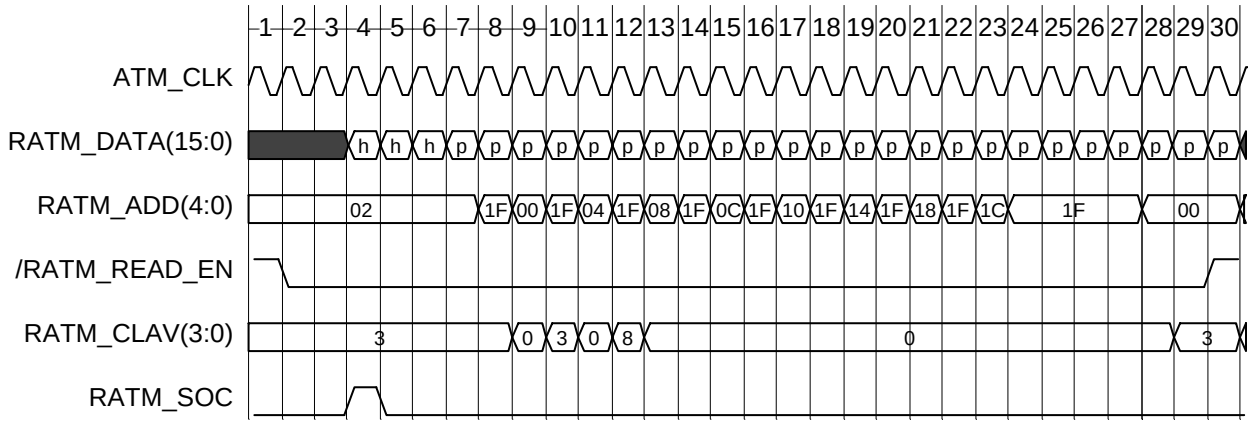

Figure 23. Receive UTOPIA Multiplexed Status Polling (MSP), Including Cell Transfer

Figure 24 shows the transmit UTOPIA 50 MHz MSP including cell transfer.

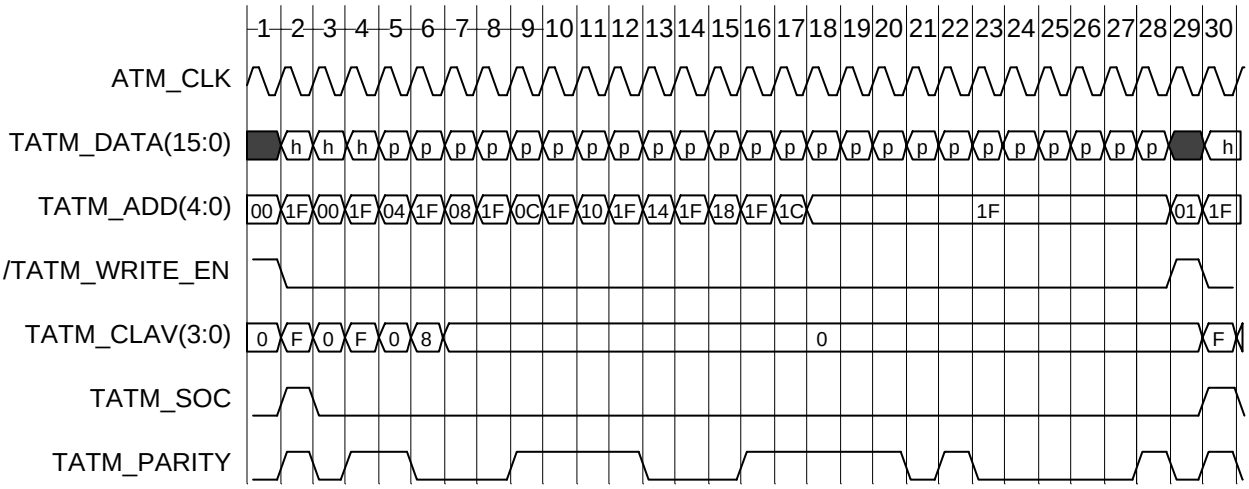

Figure 24. Transmit UTOPIA 50 MHz Multiplexed Status Polling (MSP), Including Cell Transfer

2.4.2.1.3 Direct Status Indication Using Four Cell Available Signals

When configuring the device, setting the MSP mode bit implicitly turns on direct status indication, since it is a subset of the implemented MSP method.

2.4.2.2 Priority Encoding and TDM Table

The Transmit UTOPIA selects PHY devices for service based upon:

- the assigned UT PRIORITY for the PHY(refer to “UT_PRIORITY” on page 119).
- the configuration of the TDM (Time Division Multiplex) table
- per VO presence of the cells in the QEngine
- cell available assertions received from the PHYs.

The use of priority servicing is beneficial when using multi-phy configurations and the UTOPIA bandwidth is nearly fully subscribed.

2.4.2.2.1 Basic 2 Level Priority Algorithm

When TDM is disabled (refer to section 7.2.10 UTOPIA_CONFIG) a PHY device is assigned either a high or low UTOPIA priority based of the bandwidth of the PHY device. Within a priority level (high or low), further control over the service algorithm can be implemented by assigning the lowest numbered PHY addresses to the highest bandwidth PHYs. The general algorithm for deciding which PHY to service is as follows:

1. The High priority encoder has highest service priority. From the high priority PHYs, the lowest address PHY that has indicated it can accept a cell (and for which a cell is present in the QEngine) is selected. If no high priority PHY is selected, then the low priority set is considered next.
2. The Low priority encoder has the next highest service priority. The lowest address PHY that has indicated it can accept a cell (and for which a cell is present in the QEngine) is selected. If no low priority PHY is selected then the cell time is wasted unless the Watchdog is configured for operation, in which case the stale priority set is considered next. The Watchdog is only available on the Transmit side.
3. The Transmit Stale priority encoder has the lowest priority and is created for the PHY devices that the Watchdog deems stale. The lowest address PHY that has been detected dead or "stale" by the Watch-Dog (and for which a cell is present in the Qengine) is selected. The cell is played out on the interface in order to relieve VO queue depth congestion. The Watchdog plays the role of making a best effort delivery, even though the PHY is considered dead.

Caveat: Service selection is performed each cell time with the CLAV information gathered from the previous cell time. This is particularly important, when the standard polling method is used and not all phy's can be polled in a single cell time. In this mode, UTOPIA Priorities have relative meaning within 4 address groups of 8 (0to7, 8to15, 16to24 and 25to31). For example a high priority phy of address=1 will compete for service with a low priority phy of address=7, but will not compete for service against a low priority phy of address=10 since they are in different groups. It is conceivable that a low priority phy can receive as much service as a high priority phy. This could be the case if the phy at address=10 is the only phy in its address group. It will get the entire cell time bandwidth simply because there are no other phys to compete with.

This problem does not exist in MSP mode since all CLAV information is gathered in one cell time.

2.4.2.2.2 TDM and the Basic 2 Level Priority Algorithm

When TDM is enabled, (refer to section 7.2.10 UTOPIA_CONFIG and to section 7.2.12 UT_ENABLE for configuring the TDM Pool) another level is added on top of the Basic 2 Level Priority Algorithm. The TDM table has primary service priority if the UTOPIA interface is configured to use the TDM feature. Each cell time, the TDM pointer is advanced through the TDM Pool in a round-robin fashion. When a PHY is pointed to and a cell is present in the QEngine for the PHY, it will be selected. If a PHY is selected and a cell is NOT present in the QEngine for the PHY, the selection process is deferred to the Basic 2 Level Priority Algorithm so that the cell time is not wasted. The TDM table is most useful when configurations require uniformally distributed bandwidths, such as 4xOC3 configurations. In the event that the TDM bit is not set in the UTOPIA_CONFIG then the servicing algorithm reduces to the Basic 2 level Priority encoding scheme consisting of 1, 2 and 3 above.

The Receive side of the QRT operates in the same fashion as the Transmit side with the exception of the Stale Priority level since there is no Watchdog present in the Receive side.

The UTOPIA Level 2 specification is not designed to support oversubscription due to its lack of multi-priority cell presence indications. The QRT interface assumes this is the case in order to operate correctly.

2.4.2.3 Independently Configurable Interfaces

The receive and transmit sides of the UTOPIA interface are independently configurable for either single-PHY OC-12 or multi-PHY operation. The RX_OC_12C_MODE, TX_OC_12C_MODE, and UTOPIA_2 bits (refer to section 7.2.11 “UTOPIA_CONFIG” starting on page 117) configure the device for such operation. This allows versatility in the types of PHY environments that can be supported (for example, environments that contain high-speed, single-PHY devices can be supported, as well as environments in which the QRT must perform single-chip, multi-PHY to high-speed, single-PHY muxing operations). This versatility is particularly helpful when interfacing to the PMC-Sierra, Inc. PM7322 RCMP-800 Operations, Administration, and Maintenance (OAM) processor, since the output of that device has an interface similar to a single-PHY SATURN interface.

2.4.2.4 Output Channel Number Insertion

The transmit side of the UTOPIA can be configured to insert the QRT output channel identifier in the HEC/UDF field of outgoing cells. The output channel identifier is a value used by the QRT transmit portion to identify cells of a particular cell stream as they come in from the fabric. Insertion is configured by means of setting the UTOPIA_CONFIG(7) register. If the configuration bit is set to 0, the UTOPIA inserts a value of FFFFh in the HEC/UDF field. The transmit UTOPIA does not calculate the HEC for outgoing cells.

2.5 Receiver Operation

2.5.1 Receive Channel Lookup

The receive channel lookup uses two tables: VI_VPI_TABLE (refer to "VI_VPI_TABLE" on page 175) and VCI_TABLE (refer to "VCI_TABLE" on page 176) to generate a channel number for an incoming cell. The channel number in turn is used to access the Channel Control Block (CCB), in the connection table. The CCB contains the configuration and state for the connection.

Figure 25 shows the method used to generate the channel number for VCCs: the Virtual Input (VI) number and the VPI bits are used to index into a VI_VPI_TABLE of up to 4K entries per VI. Each entry contains the base address of a block in the VCI_TABLE for that VP and the size of that block. A VCI_TABLE entry contains a channel number for that VCC. On the other hand, if channel is a VPC, its VI_VPI_TABLE contains the channel number directly (see Figure 26).

The number of active VC bits can be modified during operation of the QRT by creating a new VCI_TABLE and then changing the VC_BASE and VCI_BITS (refer to "VCI_BITS" on page 176) values to point to the new table in one write. This is possible since the BLOCK_OFFSET (refer to "BLOCK_OFFSET" on page 176) is just a pointer to the VCI_TABLE, and the VCI_TABLE holds no state information. Thus, when the first connection arrives, the eventual size of the VCI block can be initially guessed. Later, if the guess proves to be too low and the table grows too big, there is no penalty: a new VCI_TABLE can be created on-the-fly.

This method of determining the CCB allows a flexible and wide range of active VPI and VCI bits without requiring an expensive Content-Addressable Memory (CAM) or causing fragmentation of the CCBs.

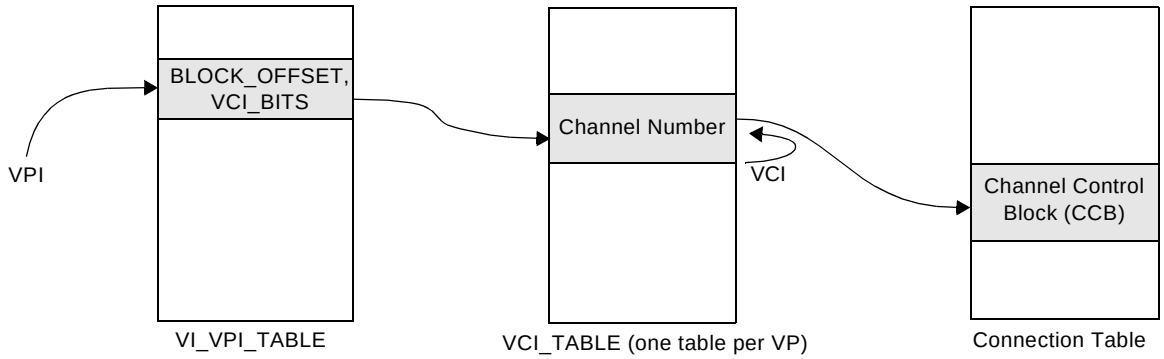

Figure 25. VCC Channel Lookup

Figure 26 shows the mapping for VPCs.

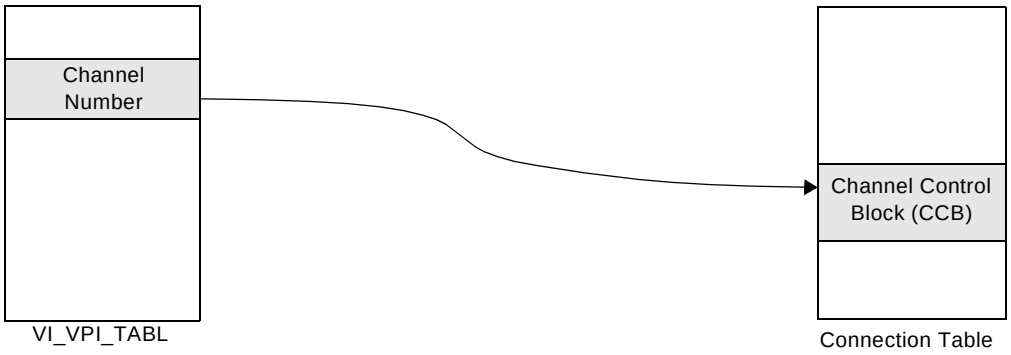

Figure 26. VPC Channel Lookup

2.5.2 Receive VC (Channel) Queuing

Receive cells are enqueued on a per-VC (channel) basis. This means that there are up to 16K queues. Singly-linked lists are used to queue the cells. The head pointers, the tail pointers, and the linked lists are all in external RAM.

Figure 27, Figure 28, and Figure 29 show the operation of the channel linked list structure.

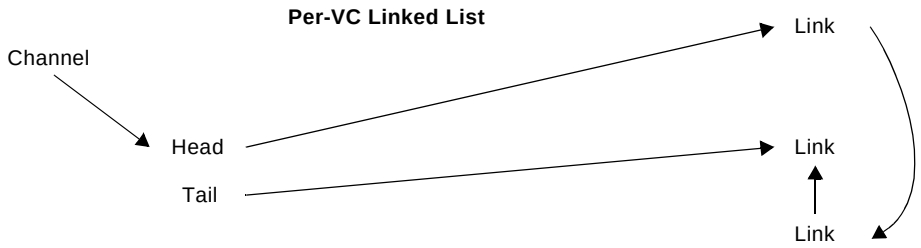

Figure 27. Channel Linked List

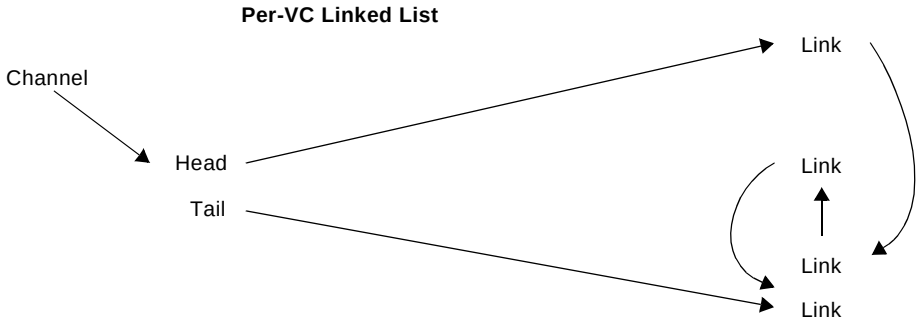

Figure 28. Channel Linked List – a New Cell Arrives

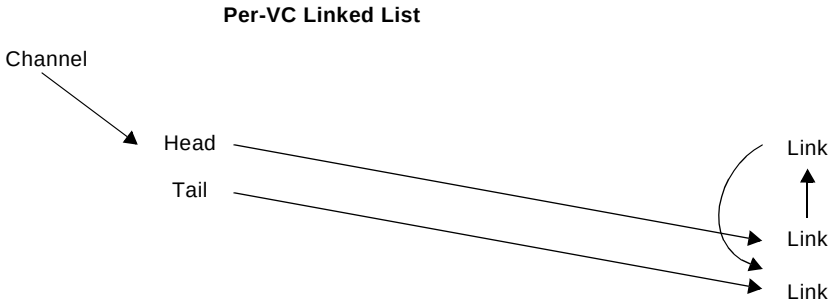

Figure 29. Channel Linked List – a Cell Is Sent to the Fabric

2.5.3 Receive Channel Ring

The list of channels eligible to send a cell to the fabric are kept in per-SC rings. The ring is kept in external memory and pointers to the previous and current channels for each SC are kept in internal memory. A channel number is entered into the ring when the first cell for that channel arrives. While cells for that channel are present in the queuing system, the channel can be removed from the ring by the dequeue process (if the channel is run-limited because of the resequencing algorithm as explained in "Receive Sequencing Algorithm" on page 34) and sometimes re-added to the ring by the process that updates the data structures with the results of the last cell time.

Figure 30, Figure 31, Figure 32, and Figure 33 on page 29 show the operation of the receive channel ring.

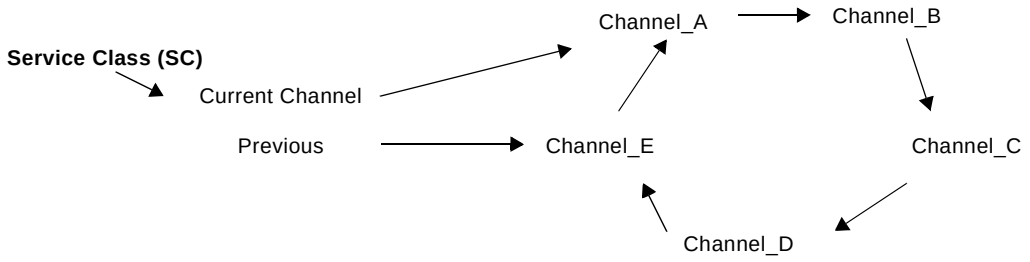

Figure 30. Receive Channel Ring

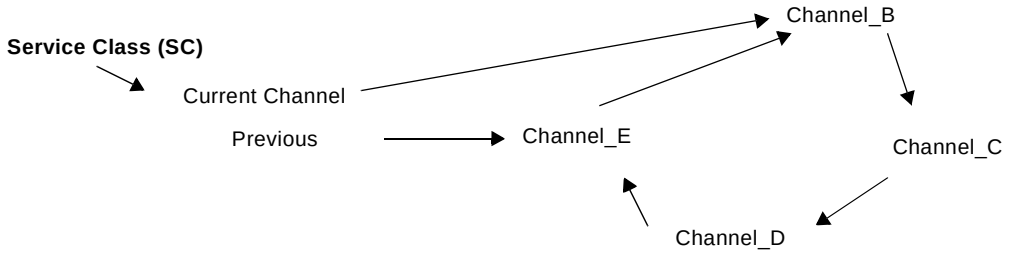

Figure 31. Receive Channel Ring after Channel_A Becomes Run-Limited

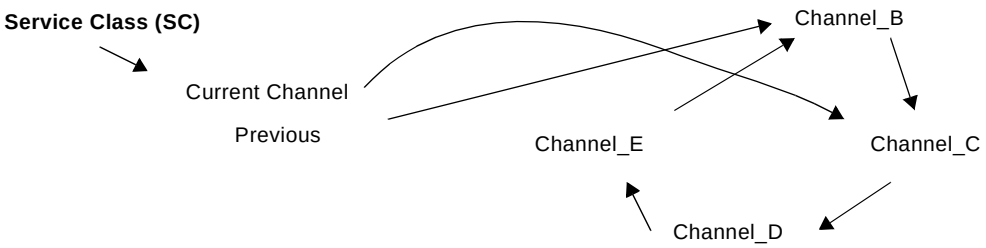

Figure 32. Receive Channel Ring after Channel_B is Served But It is Not Run-Limited

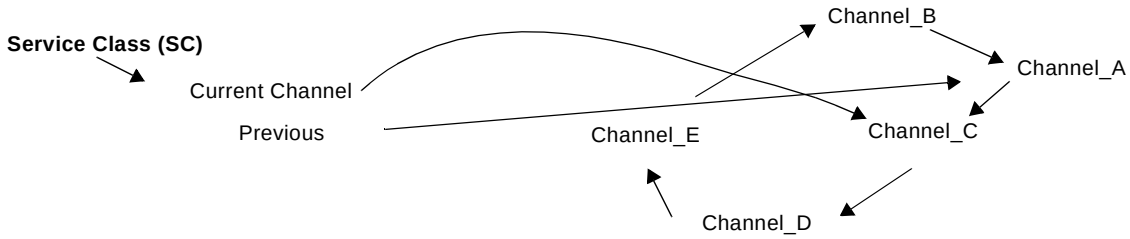

Figure 33. Receive Channel Ring After Channel_A Gets Cell Through Fabric and is Added to Ring

2.5.4 Receive Congestion Management

The receive queue controller maintains current, congested, and maximum queue depth counts of cells on a per-VC, per-SC, and per-device basis. Three congestion management algorithms are available for use on a per-channel basis. In each channel's RX_CH_CONFIG word (refer to section 9.3.1.1 "RX_CH_CONFIG" starting on page 184) there are bits that enable EPD, CLP-based discard, and EFCI. These may be used in combination. In addition, PTD is supported as a mode of the EPD operation.

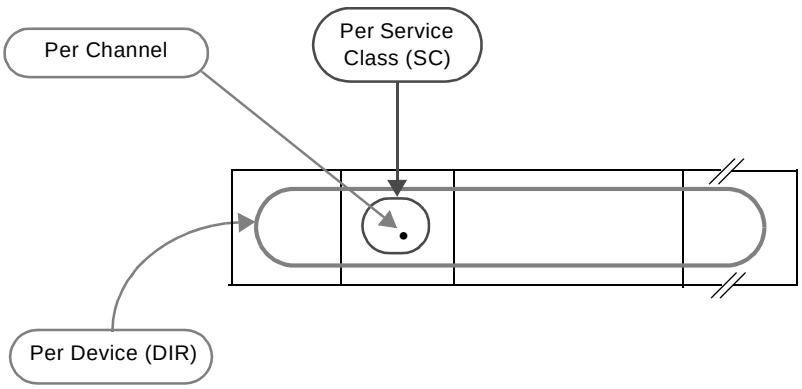

Figure 34. Receive Congestion Limits

A congestion hysteresis bit is kept for each threshold. This bit is set whenever the queue depth exceeds the congestion limit for that threshold. This bit remains asserted until the queue depth falls below one-half of the congestion threshold.

Figure 35 illustrates the operation of EPD/PTD.

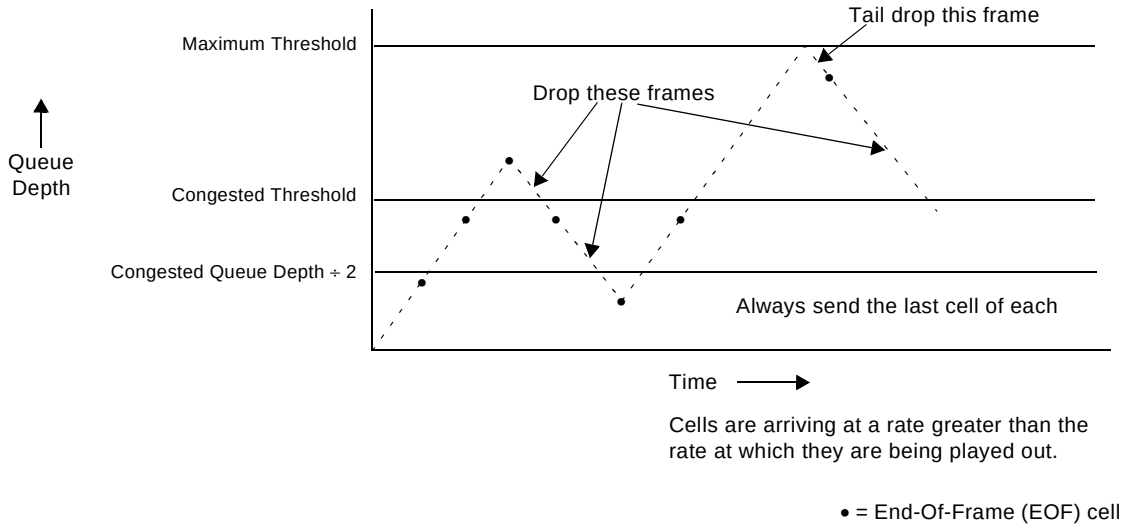

Figure 35. EPD/PTD Operation

Figure 36 shows the operation of EPD in combination with CLP-based dropping.

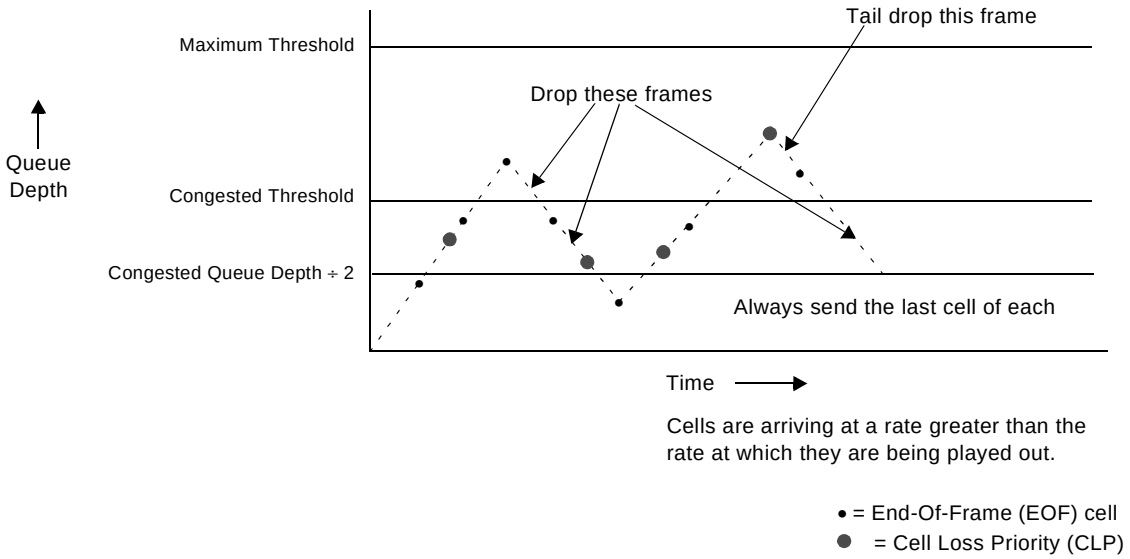

Figure 36. EPD/PTD with CLP Operation

Figure 37 shows the operation of EFCI.

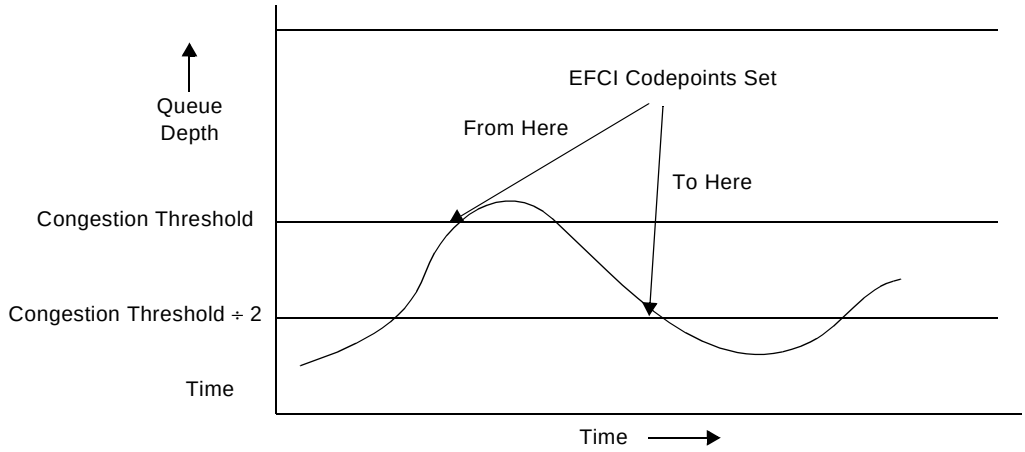

Figure 37. EFCI Operation

The congestion limits are kept in an exponential form. The interpretation of the limits is the same for all measurements, except the device limit. For the other measurements, the value of "0" causes the measurement to always find congestion. The value of "1" may not be used. The value of F_h causes congestion to be found for the limit when the queue depth is 31744. This allows a 15-bit value to be used to store the state of each measurement except the device measurement, which has a 16-bit value.

2.5.5 Receive Queue Service Algorithm

Each switch fabric cell time, the receive queue controller selects up to four cells for transmission to the switch fabric. The controller supports per-channel (per-VC) queues with 64 SCs. The controller addresses the following issues: QoS, Cell Delay Variation (CDV) minimization, Minimum Cell Rate (MCR) guarantees, and fairness maximization. The flexibility of the controller ensures that VCs receive their expected bandwidth in a timely fashion depending upon their traffic requirements.

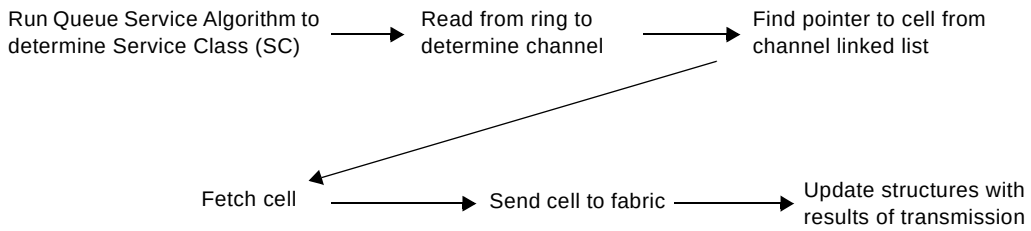

Figure 38. Steps to Send a Cell to the Fabric

The controller has a scheduler that selects cells to be placed in pipelined, “ping-pong” buffers. Once a cell is selected, it is placed in one of these buffers. Each of the four outputs to the switch fabric has two buffers: while a cell in buffer *A* is being transmitted, another cell is selected and placed into buffer *B*. On the subsequent switch fabric cell time, the buffers are “ping-ponged”, and the cell in buffer *B* is sent. Meanwhile, another cell is selected for buffer *A*.

An exception to this process is when the controller receives a negative acknowledgment (NACK) for transmission of a cell. There are two cases: the NACK is an MNACK, indicating cell transmission failed due to collision in the middle of the network, or else the NACK is an ONACK, indicating cell transmission failed due to collision at an output of the network. In the former case, the cell’s switch fabric priority (assigned during VC setup) is compared with that of the cell (if any) in the other ping-pong buffer. Call the first cell *X*, and the second cell *Y*. If the priority of cell *X* is greater than or equal to that of cell *Y*, the buffers are not ping-ponged, and cell *X* will be resent next time. If the priority of cell *X* is less than that of cell *Y*, cell *X* remains in its buffer, and the buffers are ping-ponged as usual, with cell *Y* being sent next. In the latter case, the cell is requeued at the head of its VC’s queue. Thus, the cell will be retransmitted, but at a later time than if the cell was MNACKed.

The switch fabric has been specially designed to minimize the possibility of consecutive collisions at the same place in the middle of the network, and thus a cell’s transmission that failed in that manner stands a good probability of being successful in an immediately subsequent transmission attempt. Collisions at an output of the network are more likely to be recurring for a period of time, and thus the next transmission attempt is delayed.

The scheduler that places cells in the ping-pong buffers operates as follows: The SCs are arranged in a tabular fashion as seen in Figure 39. An SC is designated for either unicast or multicast traffic. Additionally, an SC is designated as either strict priority SC1, strict priority SC2, or General Purpose (GP). Associated with each SC is a weight of either 1, 4, 16, or 64. This information is

used by the controller to decide which SC to service. Following this decision, the selected SC's VCs are serviced in a round-robin manner. The selected VC then transmits the first cell in its queue.

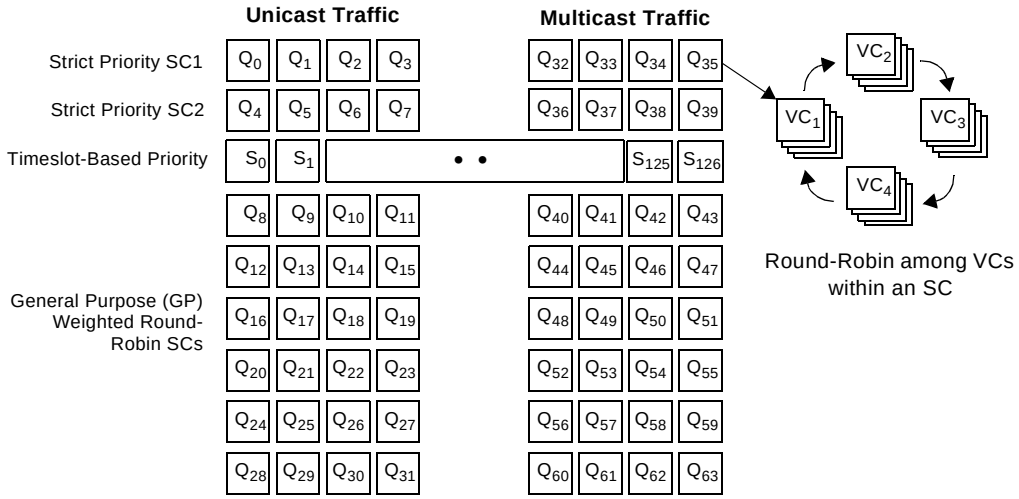

Figure 39. Receive Service Class (SC) Map

The general algorithm for deciding which SC to service is as follows (certain multicast SCs may be ineligible for selection in particular modes or operating conditions; these will be described after the numbered list that follows):

1. Strict priority SC1 has primary service priority. If there is an SC1 with a cell, it will be selected. The SC1 service classes are serviced in a weighted round-robin manner, alternating between unicast and multicast classes (Q_0, Q_{32}, Q_1, Q_{33}, Q_2, Q_{34}, Q_3, Q_{35}, Q_0, ...). The SC1 round-robin pointer will remain pointed at an SC for up to *w* cell selections, where *w* is the SC's weight. If no cells are available in an SC, the round-robin pointer is advanced. Thus, the most time-critical VCs should be placed in an SC1 service class. The pointer for the SC1 service classes is separate from the pointer to the SC2 and GP service classes.
2. Strict priority SC2 has secondary service priority. It is treated in the same fashion as SC1, except it has its own independent round-robin pointer and the weighted round-robin order is: Q_4, Q_{36}, Q_5, Q_{37}, Q_6, Q_{38}, Q_7, Q_{39}, Q_4, ...
3. If no cell exists in the strict priority classes, then the controller accesses the timeslot-based priority table in a round-robin manner. Each entry of this table contains a GP SC number. If the SC pointed to by the active entry has cells, that SC is selected. The active entry is incremented to the next timeslot each time the timeslot table is accessed. The table has 127 entries and wraps around. This servicing mechanism provides the MCR guarantee on a per-SC basis. The number of times an SC is placed in the timeslot table can be used to determine its MCR.
4. If no cell exists in the strict priority classes, and no cell exists in the SC pointed to by the active entry of the timeslot-based priority table, then the GP SCs are serviced in a weighted round-robin manner similar to the SC1 and SC2 classes (Q_8, Q_{40}, Q_9, Q_{41}, Q_{10}, Q_{42}, Q_{11}, Q_{43}, Q_{12}, Q_{44}, ..., Q_{31}, Q_{63}, Q_8, ...). Again this has a separate round-robin pointer than that kept for the SC1 and SC2 service classes.

Certain multicast SCs may be ineligible for selection due to the aggregate mode and the backpressure from the switch fabric. The QRT can be set to a multicast aggregate mode of either 1 or 4. In aggregate mode of 1, each of the switch fabric outputs of the QRT are treated as distinct outputs. Multicast connections must be specifically assigned to an SC in the corresponding column of multicast SCs (there are 32 multicast SCs, with four columns of eight classes each), since all the cells of a multicast VC must use the same output. In this mode, only one column (eight) of the multicast SCs will be eligible for selection (for example, service classes Q_{32}, Q_{36}, Q_{40}, Q_{44}, Q_{48}, Q_{52}, Q_{56}, and Q_{60} correspond to port 0 and service classes Q_{33}, Q_{37}, Q_{41}, Q_{45}, Q_{49}, Q_{53}, Q_{57}, and Q_{61} correspond to port 1). The other three columns of SCs (total of 24 SCs) will be ineligible. In aggregate mode of 4, the four outputs are treated as one logical output, and thus all multicast SCs may be selected for any of the four outputs.

Additional SCs may be ineligible due to backpressure from the switch fabric. There are three types of backpressure: high, medium and low. High backpressure renders the eight SC1 and SC2 multicast SCs ineligible (Q_{32} to Q_{39}). Medium backpressure renders the first eight GP SCs ineligible (Q_{40} to Q_{47}, two rows of four). Low backpressure renders the last 16 GP SCs ineligible (Q_{48} to Q_{63}, four rows of four).

The receive queue controller scheduler provides the following benefits:

- QoS - the strict priority scheme between SC1, SC2, and GP SCs, and the weighted round-robin algorithms allow satisfaction of QoS guarantees.
- CDV minimization - the treatment of the strict priority SCs ensure cells within these SCs get timely service.
- MCR guarantee - the timeslot table ensures all SCs will receive a minimum amount of servicing (clearly, the aggregate bandwidth given to the SC1 and SC2 VCs affects the remaining bandwidth to be divided between the GP SCs).
- Fairness maximization - how SCs (1, 4, 16, or 64) are weighted allows different SCs to support different bandwidth requirements (for example, high bandwidth SCs are assigned 64 and are serviced 64 times as often as low bandwidth SCs, which are assigned 1).

2.5.6 Receive Sequencing Algorithm

One of the service guarantees ATM offers is the FIFO delivery of cells. Since the QRT can send multiple cells from a channel simultaneously across the fabric, and not all of those cells will get through on the first try, the QRT must support an algorithm to make sure the cells can be put back into order. The algorithm it supports is a classic window algorithm where only N cells are allowed to be outstanding without acknowledgment. In the QRT, N is either 1 or 2. This limits the data rate of an individual connection to approximately 155 Mbps. The cells are sequence numbered and reordered at the transmit side.

This algorithm is implemented by removing the channel from the ring of eligible channels whenever two cells are outstanding. The channel is then called run-limited. It also removes the channel from the ring if the last cell present has been sent to the switch fabric. The channel is then called cell-limited. In the former case, it will remain off the ring until the fabric transmission results for

a run-completing cell are known. For N = 1, every cell completes a run. For N = 2, the cell with the modulo lower Sequence Number (SN) is the run-completing cell. At that time it will be added back onto the ring if there are more cells to send or if that cell was ONACKed, in which case that cell can be resent.

The pointers for these cells are stored in two locations in the CCB. When starting from no cells in the fabric, the first cell sent is always in POINTER0 and the second cell is always in POINTER1.

For multicast and unicast cells, use N = 2. The N = 1 setting is available for use, but has lower utility than the N = 2 setting for virtually all situations.

2.6 Transmitter Operation

2.6.1 Transmit Queuing

Transmit cells are enqueued on a per-SC, per-VO basis. As there are 31 VOs, and 16 SCs per VOs, there are a total of 496 queues. Singly linked lists are used to queue the cells. The head and tail pointers are in internal RAM and the linked lists are in external RAM. Figure 40 shows an example transmit per-SCQ linked list.

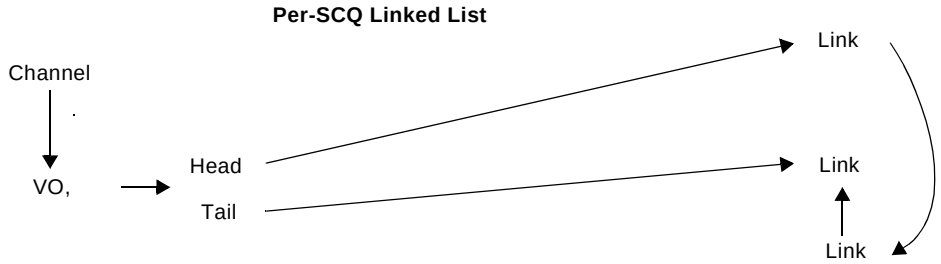

Figure 40. Transmit Per-SCQ Linked List

2.6.2 Transmit Congestion Management

A cell received from the switch fabric interface is queued by the transmit queue controller if it passes ten buffer threshold checks: both maximum and congested thresholds for the device, VO, SC, queue, and channel as shown in Figure 41 on page 36. The cell waits in the transmit cell buffer DRAM until the transmit queue controller selects it for transmit multicast/header mapping. The cell then exits the device through the UTOPIA interface.

A congestion hysteresis bit is kept for each threshold. This bit is set whenever the queue depth exceeds the congestion limit for that threshold. This bit remains asserted until the queue depth falls below one-half of the congestion threshold.

The congestion limits are kept in an exponential form. The interpretation of the limits is the same for all measurements except the device limit. For the other measurements, the value of 0 causes the measurement to always find congestion. The value of 1 may not be used. The value of F_h

causes congestion to be found for the limit when the queue depth is 31744. This allows a 15-bit value to be used to store the state of each measurement except the device measurement, which has a 1-bit value.

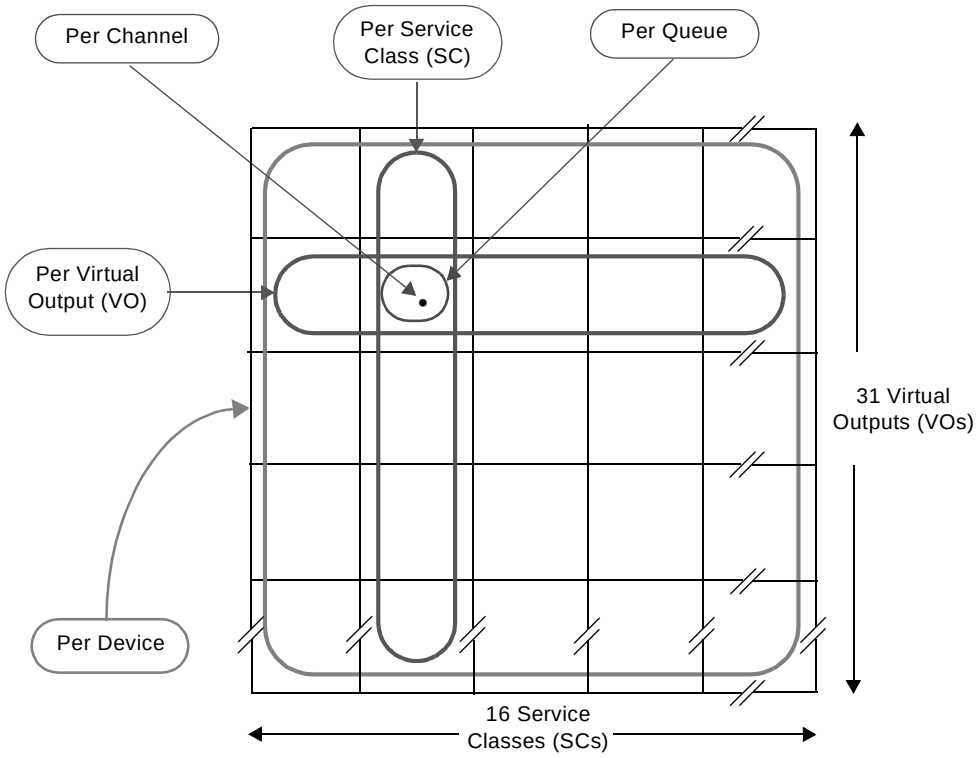

Figure 41. Transmit Maximum and Congested Threshold Checks

Three congestion management algorithms are available for use on a per-channel basis. In each channel's TX_CH_CONFIG word (refer to section 9.3.1.7 "TX_CH_CONFIG" starting on page 189) are bits that enable EPD, CLP-based discard, and EFCI. These may be used in combination. In addition, Packet Tail Discard (PTD) is supported as a mode of the EPD operation. Figure 35 on page 30 illustrates the operation of EPD/PTD. Figure 36 on page 30 illustrates the operation of EPD/PTD with CLP.

As described in "Transmit Resequencing Algorithm" on page 39, there is an interaction between EPD and the resequencing algorithm. Refer to that section for a complete description.

2.6.3 Transmit Queue Service Algorithm

The transmit queue controller supports 16 SCs for each of its 31 VOs (the per-VO structure is shown in Figure 42 on page 38). As with the receive queue controller, the transmit queue controller addresses the following key issues: QoS, CDV minimization, MCR guarantee, fairness maximization, and output isolation.

The VO for which a cell is to be sent is determined first by doing a bit-wise AND of two vectors: one vector indicates the presence of a cell for a VO, and the other vector indicates the willingness of a VO to accept a cell. Of the matching VOs, the lowest numbered VO of high priority is selected if possible; otherwise, the lowest numbered VO is selected.

Once the VO is known, the controller has a scheduler that selects a cell to be transmitted to the UTOPIA interface. The scheduler operates as follows: The SCs are arranged in a tabular fashion as seen in Figure 42 on page 38. An SC is designated for either unicast or multicast traffic. Additionally, an SC is designated as either strict priority SC1, strict priority SC2, or GP. Associated with each SC is a weight of either 1, 4, 16, or 64. This information is used by the controller to decide which SC to service. Following this decision, the selected SC's cells are serviced in a FIFO manner.

The general algorithm for deciding which SC to service is similar to that used by the receive queue controller, and is as follows:

1. Strict priority SC1 has primary service priority. If there is an SC1 service class with a cell, it will be selected. The SC1 service classes are serviced in a weighted round-robin manner, alternating between unicast and multicast classes (Q_0, Q_8, Q_0, ...). The SC1 round-robin pointer will remain pointed at an SC for up to *w* cell selections, where *w* is the SC's weight. If no cells are available in an SC, the round-robin pointer is advanced. Thus, the most time-critical VCs should be placed in an SC1 service class.
2. Strict priority SC2 has secondary service priority. It is treated in the same fashion as SC1, except it has its own independent round-robin pointer, and alternates: Q_1, Q_9, Q_1, ...
3. If no cell exists in the strict priority classes, then the controller accesses the timeslot-based priority table in a round-robin manner. Each entry of this table contains a GP SC number. If the SC pointed to by the active entry has cells, that SC is selected. The active entry is incremented to the next timeslot each time the timeslot table is accessed. The table has 127 entries and wraps around. This servicing mechanism provides the MCR guarantee on a per-SC basis. The number of times an SC is placed in the timeslot table can be used to determine its MCR.
4. If no cell exists in the strict priority classes, and no cell exists in the SC pointed to by the active entry of the timeslot-based priority table, then the GP SCs are serviced in a weighted round-robin manner simi-

lar to the SC1 and SC2 classes (Q_2, Q_{10}, Q_3, Q_{41}, Q_{11}, ..., Q_7, Q_{15}, Q_2, ...).

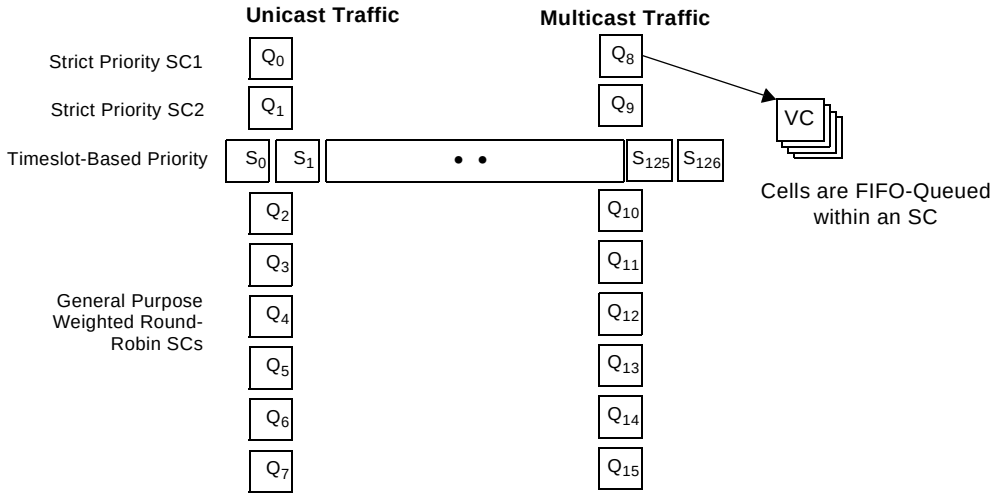

Figure 42. Transmit Service Class (SC) Map (Per VO)

The transmit queue controller scheduler provides the following benefits:

- QoS - the strict priority scheme among SC1, SC2, and GP SCs, and the weighted round-robin algorithms allows satisfaction of QoS guarantees.
- CDV minimization - the treatment of the strict priority SCs ensure that cells within these SCs get timely service.
- MCR guarantee - the timeslot table ensures all SCs will receive a minimum amount of servicing (clearly, the aggregate bandwidth given to the SC1 and SC2 VCs affects the remaining bandwidth to be divided between the GP SCs).
- Fairness maximization - the weights of the SCs (1, 4, 16, or 64) allow different SCs to support different bandwidth requirements (for example, high bandwidth SCs are assigned 64 and are serviced 64 times as often as low bandwidth SCs, which are assigned 1).
- Output isolation - the cells of channels destined for different VOs are kept in separate data structures. This helps isolate the effects of congestion on one VO from causing congestion on another VO.

Figure 43 illustrates the steps that are taken when playing out a cell.

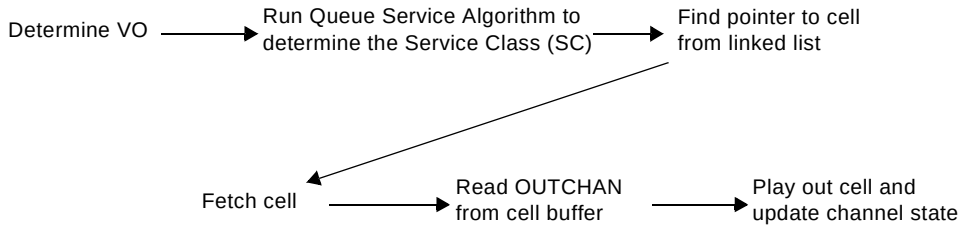

Figure 43. Cell Playout Steps

2.6.4 Transmit Resequencing Algorithm

To guarantee the FIFO delivery of cells, the QRT supports an algorithm to make sure the cells can be put back into order. The algorithm it supports is a classic window algorithm where only N cells are allowed to be outstanding without acknowledgment. In the QRT, N is either 1 or 2. This limits the data rate of an individual connection to approximately 155 Mbps. The transmit end reorders the cells according to their SN.

The resequencing of one algorithm ignores the incoming SN and accepts all cells as if their SN were correct. This can be used for multicast cells as the QSE delivers them in FIFO order.

The resequencing of two algorithms inspects an incoming cell to determine if it has the expected SN, *e*. If it does, the cell is immediately processed. If it has SN *e + 1*, then it is stored to await the run-completing cell (that is, the cell with the original expected SN, *e*). If it has neither SN *e*, nor SN *e + 1*, a recovery algorithm is started which gets the channel back into sequence. This is described in "Transmit Recovery Algorithm" on page 40.

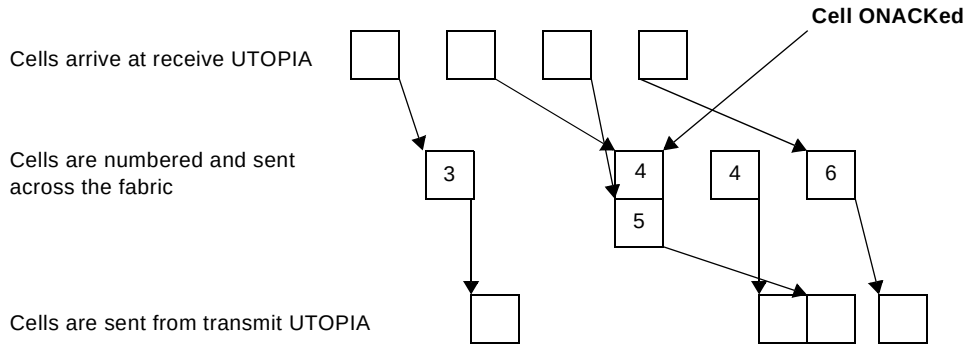

Figure 44. Transmit Resequencing Operation

The resequencing of two algorithms interacts with EPD. When a cell is missing, the algorithm cannot determine if the missing cell is an End-Of-Frame (EOF) cell. It is then necessary to defer the choice of whether or not to send both cells until the run-completing cell is received. The

choice of whether or not to send or drop one or more of the cells is affected by the EOF information, because one frame that is being dropped may end, and another frame that is not to be dropped may start.

2.6.5 Transmit Recovery Algorithm

No recovery algorithm is needed for the resequencing of one algorithm since the SN is ignored.

For resequencing of two algorithms, when a cell with SN *s* is received, and *s* is neither equal to the expected cell number *e*, nor equal to *e + 1*, then the cell is dropped. The new expected SN (for the next cell) is set at *s + 1*. The next time two consecutive cells are received in ascending SN order, the channel will have recovered its sequence. Using this algorithm, some legitimate cells may be dropped while recovering. For example, if the next two cells are legitimate, but are received in descending SN order, they will both be dropped.

2.6.6 Transmit Multicast Cell Background Process

The transmit multicast background process traverses the linked list for that channel and prepares a list of pointers to cells and pointers to headers for multicast cells. This allows the dequeue process to replicate the cell with new headers to each entry in the linked list. This is necessary because multicast cells are bound to different destinations and need different headers.

Figure 45 shows the replication process that occurs, according to how the fields in the MC_LIST word are set.

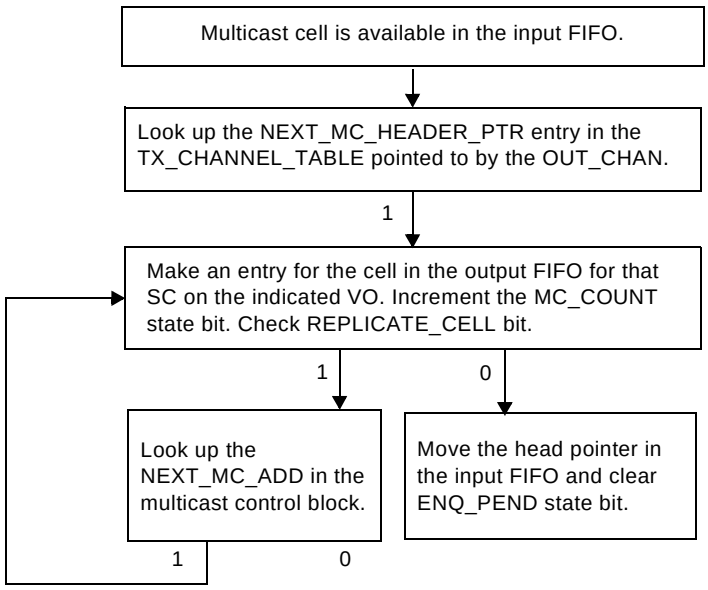

Figure 45. Multicast Background Process

Figure 46 shows the operation of the multicast pointer FIFOs. When a multicast cell arrives, it is immediately stored to RAM. The pointer to that cell buffer and the OUTCHAN for that cell are put onto one of eight input FIFOs. There is one FIFO per input multicast SC. A background pointer replication process which runs at the UTOPIA rate copies pointers from the input FIFOs to the output FIFOs. It does so by traversing the linked list for that OUTCHAN and copying the pointer to the cell buffer to the output FIFO for that SC on the proper VO.

The background process dynamically identifies if any of the output FIFOs are full. If any become full, the process records which VOs are full for that SC and ceases transferring cells for that SC. Transfers still are free to occur for other SCs. Once the dequeue process serves a cell instance from that SC on the bottlenecked VO, the background process is free to continue to do replications for that SC.

The background process runs at exactly the same rate as the UTOPIA interface. This allows it to transmit multicast cells at the full rate out of the interface, even if each multicast cell is only going to one destination on this QRT.

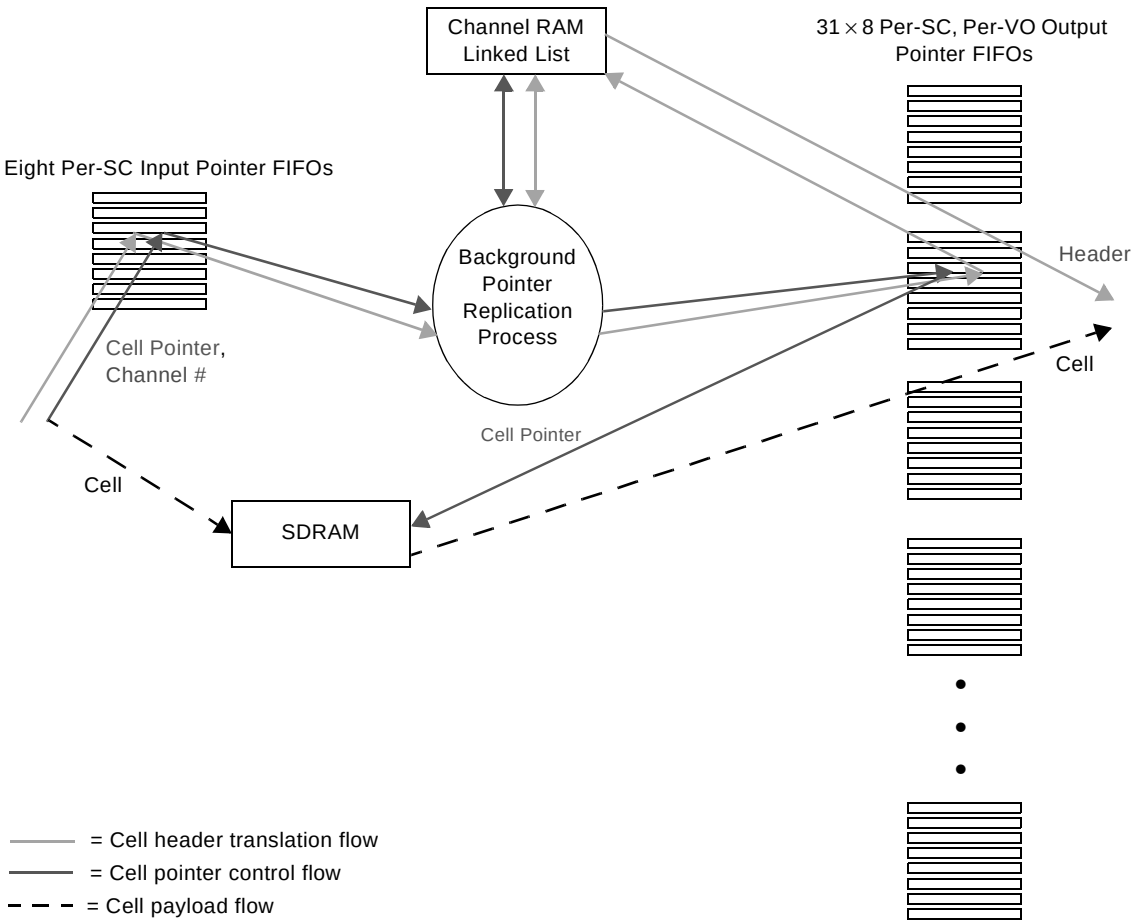

Figure 46. Multicast Pointer FIFO Operation

2.6.7 Transmit Multicast Congestion Management

The transmit multicast can have congestion management applied to it. Three of the five congestion measurements apply: the device, the SC, and the channel. The VO and the SC queue limits do not apply to multicast cells as they do not make sense. This is because only one copy of the cell is kept in the DRAM, regardless of the number of destinations to which the cell is headed. Those counts contain only the number of unicast cells present.

The QRT can be configured to either generate or not generate backpressure on a per-SC basis. If no backpressure is desired, configure TX_EXP_MAX_SC_QD (refer to "TX_EXP_MAX_SC_QD" on page 163) to one-half of the input pointer FIFO depth for that AL_RAM_CONFIG (refer to "AL_RAM_CONFIG" on page 105). This will drop all cells at a depth deeper than this, preventing backpressure from reaching back into the switch fabric. The setting of this is a system-level decision. Preventing backpressure prevents a failure or congestion on one card from affecting the performance of the fabric as a whole. On the other hand, using the backpressure allows more multicast cells to be passed without the fear of dropping in the egress QRT.

The high priority backpressure bit is derived from the near-fullness of queues 8 and 9. The medium priority backpressure bit is derived from the near-fullness of queue 10 and 11. The low priority backpressure bit is derived from the OR of the near-fullness of queues 12 to 15.

EPD, CLP-based dropping, and EFCI are all valid for multicast cells and are configured in the TX_CH_CONFIG word (refer to section 9.3.1.7 "TX_CH_CONFIG" starting on page 189) using the same bits as for unicast connections.

2.7 System Diagram of Internal QRT Blocks and External RAM

Figure 47 shows a system diagram of the internal QRT blocks and the external RAM.

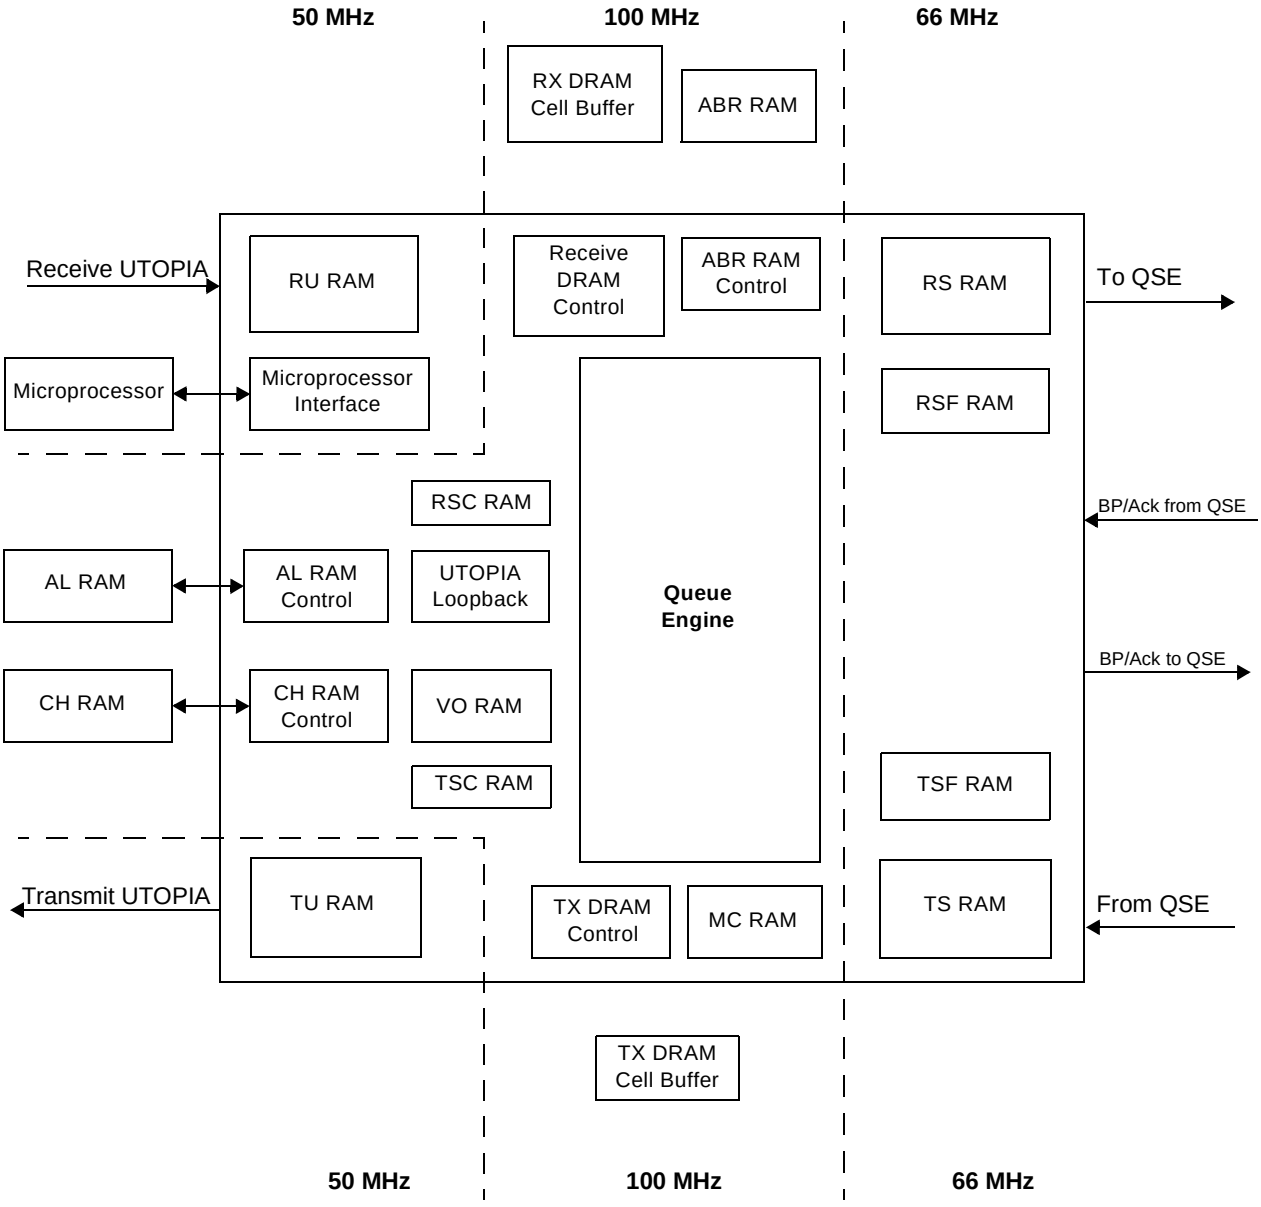

Figure 47. System Diagram of Internal QRT Blocks and External RAM

3 FAULT TOLERANCE

3.1 The Data Path

Figure 48 shows the basic data path through the switch. The SE_D_OUT/IN and SE_SOC_OUT/IN signals are used in the forward path, and the BP_ACK_OUT/IN signals are used in the backward path. Data enters the switch via the ingress or receive side UTOPIA interface and is queued at the Input half of the QRT (the IRT). The receive queue controller selects cells that are then played out to the switch fabric, which consists of one or more stages of QSEs. The cell finally enters the egress QRT where it is queued again at the Output half of the QRT (the ORT). The transmit queue controller selects a cell which is then played out of the switch via the egress or transmit side UTOPIA interface.

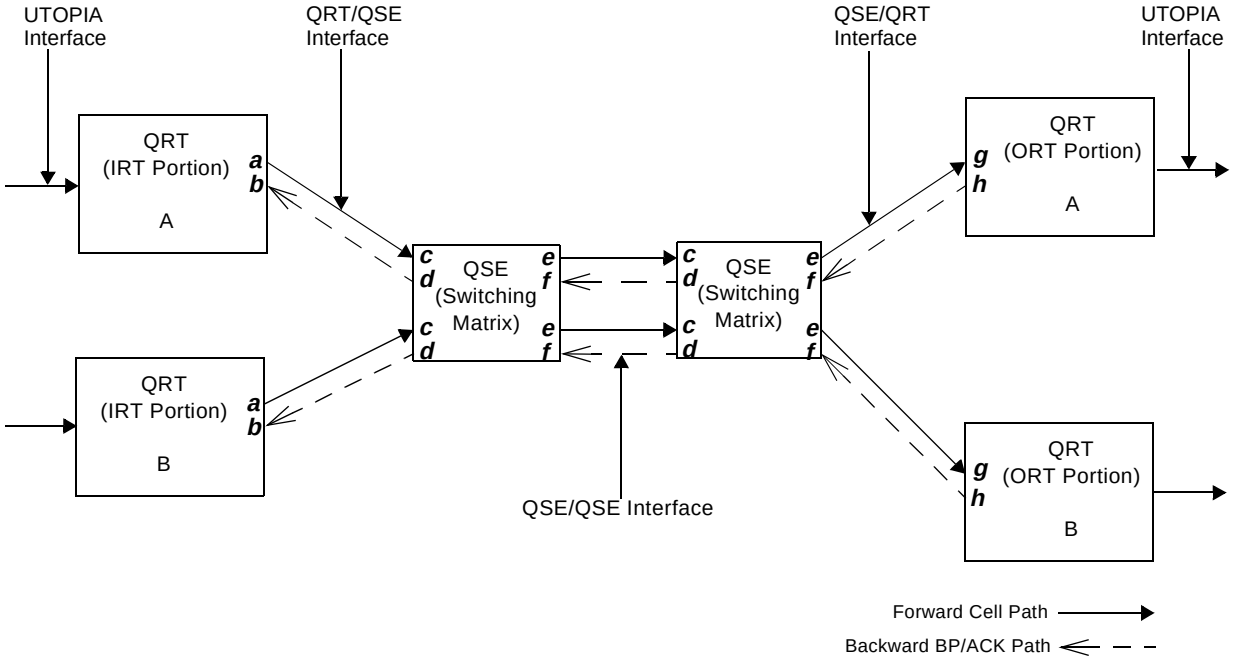

Figure 48. Basic Data Path Through the Switch

3.1.1 UTOPIA Interface

The QRT UTOPIA interface is compatible with the UTOPIA Level 1 specification revision 2.01 and the UTOPIA Level 2 specification in 16-bit mode with cell-level handshaking. An external ATM clock must be provided to this interface with a frequency between 15 MHz and 50 MHz. The lower bound is determined by the ATM_CLK failure detection circuitry. The receive and transmit sides of the interface are independently configurable to operate in either single OC-12 or multi-PHY fashion. The interface also provides several options in polling methods, so bandwidth, servicing fairness, and response time are optimized for any given PHY layer device arrangement.

3.1.2 Switch Fabric Interface

The switch fabric interface of the QRT has four nibble-wide, 50 or 66 MHz interfaces with back-pressure to interface to QSEs (PM73488s). The device can avoid head-of-line blocking by receiving two forms of negative acknowledgment from the switch fabric. One form of negative acknowledgment indicates congestion that is likely to be resolved on the next cell time. This is termed a Mid-switch NACK (or Medium Negative ACKnowledgment - MNACK). When the QRT receives an MNACK, it resends the same cell. The other form of negative acknowledgment indicates congestion that is not likely to be resolved on the next cell time. This is termed Output Negative ACKnowledgment (ONACK). When the QRT receives an ONACK, it skips to another channel and sends a cell from that different channel.

3.2 Fault Detection and Isolation

The data transfers internally between the various RAMs and between the QRT and the QSE are checked by the following mechanisms:

- Memory parity checking
- UTOPIA interface fault detection and recovery mechanisms
- Switch fabric fault detection and recovery mechanisms

3.2.1 Memory Parity Checking

The receive and transmit buffer SDRAMs are checked by multibit parity.

All external SRAMs have parity checking. The parity conditions are checked. There are two kinds of flags (sticky and non-sticky) set for each of these parity error conditions. The sticky error bits are set by the error and are cleared by the processor. The corresponding non-sticky bits are used for debugging purposes.

3.2.2 UTOPIA Interface Fault Detection and Recovery Mechanisms

The QRT uses several mechanisms to ensure cell integrity through the UTOPIA interface and to expediently detect, isolate, and rectify fault conditions.

3.2.2.1 Header Error Check (HEC)

The receive or ingress UTOPIA interface can be configured to perform a HEC calculation using the CHK_HEC bit in the UTOPIA_CONFIG register (refer to section 7.2.11 “UTOPIA_CONFIG” starting on page 117). When a HEC failure is detected and checking is enabled, an interrupt is signaled at the processor interface and the cell is dropped at the UTOPIA interface. Some Segmentation And Reassembly (SAR) and Physical layer (PHY) devices do not produce the correct HEC or use the HEC for other purposes. To connect the QRT to these devices, clear the CHK_HEC bit.

3.2.2.2 Start Of Cell (SOC) Recovery

The receive UTOPIA interface is flexible when dealing with the SOC signal sent from the PHY layer device to the QRT. The QRT can accept a delay of up to four ATM clock cycles in the aligned SOC and data signals after assertion of the Receive UTOPIA ATM

Layer Enable signal (/RATM_READ_EN). The SOC signal can arrive anywhere within this window and the data will be accepted. Customers will find this feature useful if glue logic is used for special PHY layer device adaptations. If, however, the SOC signal arrives after the four-cycle window, the QRT will dump the cell and enter recovery mode. Recovery mode is implemented for both single and multi-PHY configurations and provides robustness to the QRT in the event of a late SOC resulting from a reset PHY or a double SOC resulting from renegade PHY devices. The recovery mode performs precession in the ATM cell cycles that follow. This is necessary to bring a PHY device back into synchronization for slotted cell-level handshaking. SOC recovery performs the same functions for stuck-at faults in the SOC signal. When an SOC failure is detected, an interrupt is signaled at the processor interface.

3.2.2.3 Transmit Watchdog

The QRT transmit or egress UTOPIA interface has a function called the “watchdog”. The watchdog exists to protect the QRT VO queues from overflow if a PHY sink goes offline or stops requesting cells. The watchdog can be configured in the UTOPIA_CONFIG register in the processor interface (refer to “WD_TIME” on page 118). The watchdog can be turned off or set to tolerate either OC-3-, DS1-, or DS0-level outputs. The watchdog operates by observing the liveliness of the Transmit UTOPIA ATM Layer Cell Available signals (TATM_CLAV(3:0)). If the QRT determines a PHY device has stopped accepting cells, the cells intended for that PHY device are played out. Otherwise, if the PHY device can accept these cells and the TATM_CLAV(3:0) signal dormancy is due to a stuck-at fault, normal UTOPIA signaling at the lowest priority occurs whenever spare bandwidth is available.

3.2.2.4 Transmit Parity

The transmit UTOPIA interface performs UTOPIA Level 2 odd parity calculation over the Transmit UTOPIA ATM Layer Data signals (TATM_DATA(15:0)) for the PHY devices to use in error checking.

3.2.2.5 ATM Clock Failure Detection

The UTOPIA interface contains an ATM clock failure detection circuit. The detection circuit samples the ATM clock with the high-frequency system clock and determines if the ATM clock possess signal changes. If the clock failure detection circuit is tripped, an interrupt is signaled at the processor interface.

3.2.2.6 Receive Cell Available Signal Stuck at 1

When a PHY interface device’s cell available signal is stuck at 1, the receive UTOPIA Level 2 interface limits the PHY to approximately one-half the receive side QRT bandwidth. This condition can result from a floating cell available line and should be avoided by designing pull-down resistors for the cell available lines. In the transmit direction, this is not such an issue, because the cell service is dependent on the presence of a cell bound for that PHY device. However, this condition should be minimized in the transmit direction also.

3.2.2.7 Highest Bandwidth Device Support in UTOPIA Level 2 Mode

The QRT UTOPIA Level-2 50 MHz interface was not designed to operate with any device possessing a bandwidth greater than that of an OC-3. For higher bandwidth requirements, the user must use the single-PHY UTOPIA Level-1 mode of operation.

3.2.3 Switch Fabric Fault Detection and Recovery Mechanisms

The QRT uses several mechanisms to ensure cell integrity through the switch fabric and to expediently detect, isolate, and rectify fault conditions.

3.2.3.1 SOC Coding

SOC Coding — A special background pattern "0000111100001..." is generated on the SOC at the ingress QRT and is propagated by the QSE. This background pattern is checked at the egress QRT. If this pattern is inconsistent or missing, the forward cell path is declared bad and the SE_INPUT_PORT_FAIL interrupt (refer to "SE_INPUT_PORT_FAIL" on page 112) is asserted.

3.2.3.2 SOC Inversions

The SOC is indicated by an inversion of the background pattern. Also, the pattern is reset so a valid SOC will always be followed by "000011110000...". This pattern reset is checked by the egress QRT, and if it is inconsistent, the forward path is declared bad and the SE_INPUT_PORT_FAIL interrupt (refer to "SE_INPUT_PORT_FAIL" on page 112) is asserted.

3.2.3.3 Redundant Cell Present Coding

The first nibble of each valid cell has a predetermined format. This format is checked as the cell is received at the egress QRT. If the format is inconsistent, the forward cell path is declared bad and the SE_INPUT_PORT_FAIL interrupt (refer to "SE_INPUT_PORT_FAIL" on page 112) is asserted. This also increases the robustness of the cell presence detection, preventing an all-1s input from creating cells.

3.2.3.4 Idle Cell Pattern Checking

An idle cell at the ingress QRT has a predetermined format. This pattern is checked at the egress QRT when the idle cell is received. If the format is inconsistent, the forward cell path is declared bad and the SE_INPUT_PORT_FAIL interrupt (refer to "SE_INPUT_PORT_FAIL" on page 112) is asserted.

3.2.3.5 Dropping Cells with Bad Header Parity

Odd parity is generated over the first 12 nibbles of every valid cell. The parity bit is embedded in the twelfth nibble. Parity checking can be disabled by asserting the PARITY_FAIL_DIS bit (refer to "PARITY_FAIL_DIS" on page 107). When parity checking is enabled, cells with bad parity are dropped and a failure is reported to the microprocessor via the TX_PARITY_FAIL flag (refer to "TX_PARITY_FAIL" on page 111).

3.2.3.6 Forced Bad Parity

The header parity detection logic can be checked by clearing the P bit in the RX_CH_TAG word (refer to "RX_CH_TAG" on page 185). This forces bad parity (even parity) to be created.

3.2.3.7 Marked Cell Counting

The Mark Bit (MB) in the RX_CH_TAG word (refer to "RX_CH_TAG" on page 185) is set and it is sent between the QRT and QSE. If ACK or NO_RESP feedback is received, the RX_MARKED_CELLS counters count enabled cells (modulo 16) at the ingress QRT. The TX_MARKED_CELLS counters count marked cells (modulo 16) that are received at the egress QRT. The RX_MARKED_CELLS and TX_MARKED_CELLS counters (refer to "MARKED_CELLS_COUNT" on page 110) help identify subtle failures in the fabric and can be used to create strong diagnostic routines. These counters are separate for all four switch fabric interfaces in each of the transmit and receive directions.

3.2.3.8 Remote Data Path Failure Indication

When a cell path is determined to be bad, the egress QRT indicates a remote failure by violating the syntax of the BP_ACK_OUT signal. This is an indication to the ingress QRT that a fabric fault in the forward cell path has been detected. Also, an SE_INPUT_PORT_FAIL interrupt (refer to "SE_INPUT_PORT_FAIL" on page 112) is flagged to the microprocessor. The cell being received at this time is discarded. This is detected as a BP_REMOTE_FAIL (refer to "BP_REMOTE_FAIL" on page 112) by the QRT or QSE on the other end of the link. Withholding backpressure from the ingress QRT prompts it to send only idle cells on the forward cell path until it recognizes a valid backpressure pattern again.

3.2.3.9 Unacknowledged Cell Detection

If no acknowledgment is received for a cell, the ACK_LIVE_FAIL interrupt (refer to "ACK_LIVE_FAIL" on page 111) is asserted. This is an indication of a problem in the end-to-end path through the switch fabric.

3.2.3.10 Switch Fabric Loopback

The internal loopback feature also helps detect and isolate fabric faults. When dribbling errors or other faults are detected, internal loopback can help isolate the fault.

3.2.3.11 Fabric Clock Failure Detection

The switch fabric interface contains a clock failure detection circuit. The detection circuit samples the fabric clock with the high-frequency system clock and determines whether or not it possess signal changes. If the clock failure detection circuit is tripped, an interrupt is signaled at the processor interface.

3.2.3.12 Liveness of Backpressure Signal

The backpressure signal is checked for a "10" pattern at the start of the backpressure signal. For each of the QSEs, the QRT has a Backpressure (BP) liveness indication bit called BP_ACK_FAIL (refer to "BP_ACK_FAIL" on page 112). There are two bits per QSE

port called ACK_LIVE_FAIL and BP_REMOTE_FAIL (refer to "ACK_LIVE_FAIL" and to "BP_REMOTE_FAIL" on page 112) that check the ACK response from the switch fabric and the liveness of the data line. The liveness signal alone will not determine that a QSE is faulty.

3.2.3.13 BP_ACK_IN Pattern Checking

Backpressure and acknowledgment are transmitted from the egress QRT to the ingress QRT in packets on the BP_ACK_OUT line. The format of the packets is as follows:

A background pattern "0000111100001. . ."
: Generated by the QRT at egress and propagated by the QSE. This pattern is checked at the ingress QRT.

First Inversion (of the background pattern)
: Indicates the beginning of the packet.

Mode
: Indicates the nature of the packet (that is, acknowledge or backpressure).

Data1
: Indicates the Most Significant Bit (MSB) of data.

Data0
: Indicates the Least Significant Bit (LSB) of data for acknowledgment - The second bit of data for backpressure.

Second Inversion
: Verify the first inversion was not a glitch and the fabric is not stuck.

Code Ext1
: The LSB of the backpressure data; otherwise, it must be "0" for acknowledgment to be accepted as valid.

Code Ext0
: Must be "0" to be accepted as valid. This is reserved for future use.

If any part of the coding is missing or inconsistent, the BP/ACK path is indicated as bad by asserting the BP_ACK_FAIL interrupt (refer to "BP_ACK_FAIL" on page 112).

3.2.3.14 BP_ACK Inversion Checking

The first inversion and the second inversion in the packet are separated by three bits to ensure a fabric fault, such as stuck at "1" or "0", or a glitch will not result in a false packet. If the second inversion is not consistent (inverse) with the first inversion, a bad path is indicated by asserting the BP_ACK_FAIL interrupt (refer to "BP_ACK_FAIL" on page 112).

When the BP/ACK path is determined to be bad, the ingress QRT withholds issuing valid cells and instead transmits idle cells until the fabric recovers from the fault.

3.2.3.15 BP_ACK Remote Failure Detection

A missing or corrupted (bad second inversion) backpressure packet is reported to the microprocessor by the BP_REMOTE_FAIL flag (refer to "BP_REMOTE_FAIL" on page 112). The BP_REMOTE_FAIL flag is an indication of a broken cell path at the ingress QRT. A missing or corrupted acknowledgment packet is reported to the microprocessor by the ACK_LIVE_FAIL (refer to "ACK_LIVE_FAIL" on page 111).

3.2.3.16 Detection Hysteresis

If the fault detected is the result of a missing or bad background pattern, it takes the QRT a minimum of 8 and a maximum of 12 switch fabric clocks to recover after the pattern has been restored. If a backpressure packet is withheld or detected as bad during a cell time because of built-in hysteresis, it takes two valid backpressure packets in successive cell times for the QRT to recover. If an acknowledgment packet is detected as bad, it takes one cell time for the QRT to recover.

3.2.4 Tables of Switch Fabric Interface Failure Behaviors

3.2.4.1 IRT-to-Switch Fabric Interface

In Figure 48 on page 44, the IRT interface consists of ***a*** and ***b***. In the figure, ***a*** refers to each of the four SE_SOC_OUT and SE_D_OUT#(3:0) data ports, while ***b*** refers to the corresponding BP_ACK_IN signals in the QRT. Table 2 summarizes the failure conditions detected by the IRT on ***b*** and the actions taken.

Table 2. Failure Conditions, IRT-to-Switch Fabric Interface

Fault Detected On *b*	Action Taken	Comment
Cannot lock to special coding and guaranteed transitions on BP_ACK_IN.	Idle cells are sent out on data interface ***a***. Internally to the IRT, cells that would have gone out are ACKed if all ports are failing; else they are ONACKed. No multicast cells are generated for the port. BP_ACK_FAIL (refer to "BP_ACK_FAIL" on page 112) signaled to the microprocessor.	Port treated as dead. Problem is probably with the BP_ACK_IN line.
No backpressure received on BP_ACK_IN.	Idle cells are sent out on data interface ***a***. Internally to the IRT, cells that would have gone out are ACKed if all ports fail; else they are ONACKed. No multicast cells are generated for the port. BP_REMOTE_FAIL (refer to "BP_REMOTE_FAIL" on page 112) signaled to the microprocessor.	Port treated as dead. Problem is with the forward data flow, and the QSE is signaling this back to the IRT.
No ACK, MNACK, or ONACK received, although unicast cell is sent out.	Cell that was transmitted is treated as sent. ACK_LIVE_FAIL (refer to "ACK_LIVE_FAIL" on page 111) signaled to the microprocessor.	

3.2.4.2 QSE Interface, Receive Data Direction

In Figure 48 on page 44, a QSE receive interface consists of ***c*** and ***d***. In the figure, ***c*** refers to each of the four SE_SOC_IN and SE_D_IN#(3:0) data ports, while ***d*** refers to the corresponding BP_ACK_OUT signals in the QSE.

Table 3 summarizes the failure conditions detected by the ORT on ***c*** and the actions taken.

Table 3. Failure Conditions, QSE Receive Interface

Fault Detected On *c*	Action Taken	Comment
Cannot lock to special coding and guaranteed transitions on SE_SOC_IN.	No backpressure sent out on ***d***. All data discarded. SE_INPUT_PORT_FAIL (refer to "SE_INPUT_PORT_FAIL" on page 112) signaled to the microprocessor.	Withholding backpressure on ***d*** signals to the previous stage that the port should not be used.
Invalid cell present coding on SE_D_IN#(3:0).	No backpressure sent out on ***d***. All data discarded. SE_INPUT_PORT_FAIL signaled to the microprocessor.	Probably due to unconnected input lines that are pulled up or down. Withholding backpressure on ***d*** signals to the previous stage that the port should not be used.
Bad idle cell coding on SE_D_IN#(3:0).	No backpressure sent out on ***d***. All data discarded. SE_INPUT_PORT_FAIL signaled to the microprocessor.	Withholding backpressure on ***d*** signals to the previous stage that the port should not be used.
Parity fail.	ONACK sent out on ***d*** for unicast data. Multicast data dropped. PARITY_ERROR (refer to the *QSE Long Form Data Sheet*) signaled to the microprocessor.	The QSE does not necessarily have time to drop the cell by the time it has detected a parity error.

3.2.4.3 QSE Interface, Transmit Data Direction

In Figure 48 on page 44, a QSE transmit interface consists of ***e*** and ***f***. In the figure, ***e*** refers to each of the 32 SE_SOC_OUT and SE_D_OUT#(3:0) data ports, while ***f*** refers to the corresponding BP_ACK_IN signals in the QSE. Table 4 summarizes the failure conditions detected by the QSE on ***f*** and the actions taken.

Table 4. Failure Conditions, QSE Transmit Interface

Fault Detected On *f*	Action Taken	Comment
Cannot lock to special coding and guaranteed transitions on BP_ACK_IN.	Idle cells sent out on data interface ***e***. If possible, data routed around port. Multicast data is dropped if all possible port choices are dead or off. Unicast data is optionally dropped if all possible port choices are dead or off. BP_ACK_FAIL (refer to "BP_ACK_FAIL" on page 112) signaled to the microprocessor.	Port treated as dead. Problem is probably with the BP_ACK_IN line.
No backpressure received on BP_ACK_IN.	Idle cells sent out on data interface ***e***. If possible, data routed around port. Multicast data is dropped if all possible port choices are dead or off. Unicast data is optionally dropped if all possible port choices are dead or off. BP_REMOTE_FAIL (refer to "BP_REMOTE_FAIL" on page 112) signaled to the microprocessor.	Port treated as dead. Problem is with the forward data flow.
No ACK, MNACK, or ONACK received, although the cell sent out is not currently monitored in the QSE.	No action taken.	Lack of ACK, MNACK, or ONACK is not monitored by the QSE.

3.2.4.4 Switch Fabric-to-ORT Interface

In Figure 48 on page 44, an ORT interface consists of ***g*** and ***h***. In the figure, ***g*** refers to each of the four SE_SOC_IN and SE_D_IN#(3:0) data ports, while ***h*** refers to the corresponding BP_ACK_OUT signals in the QRT. Table 5 summarizes the failure conditions detected by the ORT on ***g*** and the actions taken.

Table 5. Failure Conditions, Switch Fabric-to-ORT Interface

Fault Detected On *g*	Action Taken	Comment
Cannot lock to special coding and guaranteed transitions on SE_SOC_IN.	No backpressure sent out on ***h***. All data discarded. SE_INPUT_PORT_FAIL (refer to "SE_INPUT_PORT_FAIL" on page 112) signaled to the microprocessor.	Withholding backpressure on ***h*** signals to the previous stage that the port should not be used.
Invalid cell present coding on SE_D_IN#(3:0).	No backpressure sent out on ***h***. All data discarded. SE_INPUT_PORT_FAIL signaled to the microprocessor.	Probably due to unconnected input lines that are pulled up or down. Withholding backpressure on ***h*** signals to the previous stage that the port should not be used.
Bad idle cell coding on SE_D_IN#(3:0).	No backpressure sent out on ***h***. All data discarded. SE_INPUT_PORT_FAIL signaled to the microprocessor.	Withholding backpressure on ***h*** signals to the previous stage that the port should not be used.

Table 5. Failure Conditions, Switch Fabric-to-ORT Interface (Continued)

Fault Detected On *g*	Action Taken	Comment
Parity fail.	ACK sent out on ***h***. Parity errored cell dropped. TX_PARITY_FAIL (refer to "TX_PARITY_FAIL" on page 111) signaled to the microprocessor.	ACK already sent by the time the QRT has detected a parity error. In this case, a cell that was dropped was ACKed.

3.2.4.5 Types of Failures and Their Manifestations

Table 6 shows possible faults, their effects, and how they affect the network.

Table 6. Faults and Effects on the Network

	Fault	Manifestation	Effect on Network
Wire Connection	Data line from SE_D_IN#(3:0) stuck at 0 or 1.	Invalid idle cell, cells with missing Cell Present and parity error.	Port shut down (interrupt generated and backpressure withheld) on receipt of first 2 consecutive bad idle cells until condition is fixed, as port failure is sent to the source of data by lack of backpressure indication. If only one bad idel cell received, due to the round-trip delay between the QRT and the QSE, the source could send another (user) cell, causing the destination (QRT) to come out of the shut down state. This means the raw interrupt at QRT (SE_INPUT_PORT_FAIL) will be asserted and deasserted. The lateched interrupt (SE_INPUT_PORT_FAIL_LATCH) will still be asserted.
	SE_SOC_IN(3:0) line stuck at 0 or 1.	Loss-of-lock on special coding on SE_SOC_IN(3:0).	Port shut down until condition is fixed, as port failure is sent to the source of data by lack of backpressure indication.
	BP_ACK_IN(3:0) line stuck at 0 or 1.	Loss-of-lock on special coding on BP_ACK_IN(3:0).	Port shut down until condition is fixed.
	Bridging fault within a port.	Invalid idle cell with some 10/01 fail or parity error.	Port shut down on receipt of first bad idle cell until condition is fixed, as port failure is sent to the source of data by lack of backpressure indication.

Table 6. Faults and Effects on the Network (Continued)

	Fault	Manifestation	Effect on Network
QRT and QSE Port Failures	No SE_SOC_OUT generation.	Loss-of-lock on special coding on SE_SOC_IN(3:0).	Port shut down until condition is fixed, as port failure is sent to the source of data by lack of backpressure indication.
	No data or invalid data generated.	Invalid idle cell with some 10/01 fail or parity error.	Port shut down on receipt of first bad idle cell until condition is fixed, as port failure is sent to the source of data by lack of backpressure indication.
	No BP_ACK_OUT(3:0) generation.	Loss-of-lock on special coding on BP_ACK_IN(3:0).	Port shut down until condition is fixed.
QSE Chip Failures	Multicast handling.	Cell loss or generation.	Detection possible using marked cell count.
	Multicast cell pool buffer.	Parity error in header or cell.	Detection only in header; not in payload.
	Partial cell buffers.	Parity error in header and cell.	Parity error.
	Multicast and unicast selection networks.	Cell gets out on wrong port, cell duplicated, or cell lost.	Cell to wrong port may be noticed by receiving QRT, if that VC is not active. Cell duplication and cell loss detection possible using marked cell count.
	Arbiter.	Cell lost.	Detection possible using marked cell count.

4 PIN DESCRIPTIONS

4.1 Package Diagram

Figure 49 (parts 1 and 2) shows the 503-pin Enhanced Plastic Ball Grid Array (EPBGA) package used for the QRT. The package measurements are shown in millimeters.

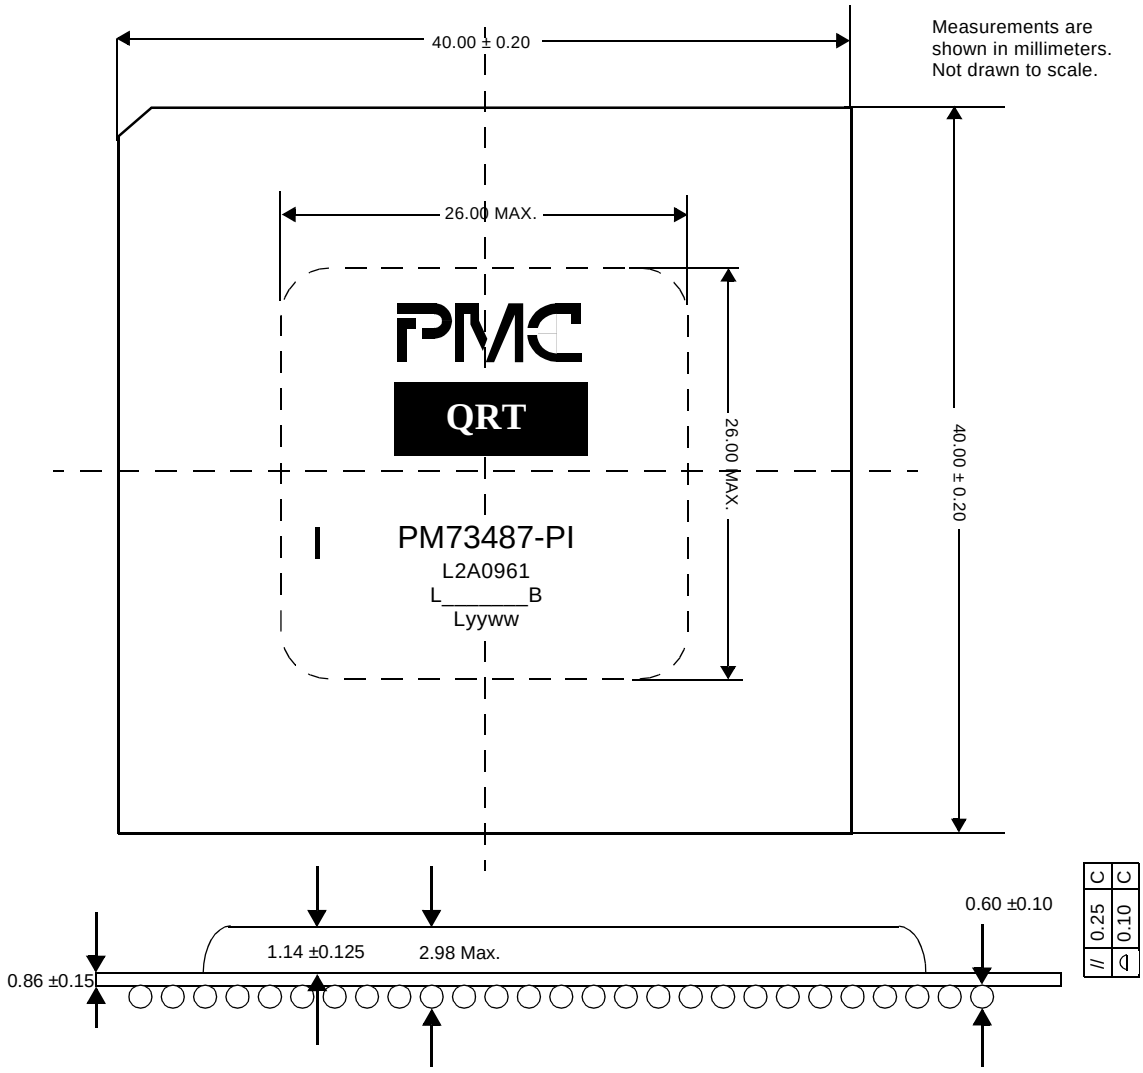

NOTES:

1. "L_____B" is the wafer batch code.
2. "Lyyww" is the assembly date code.
3. Dimensions are for reference.
4. Controlling dimension: millimeter.
5. // = Parallelism tolerance.
6. If you need a measurement not shown in this figure, contact PMC-Sierra.

Figure 49. 503-Pin EPBGA Top and Side Views (Part 1 of 2)

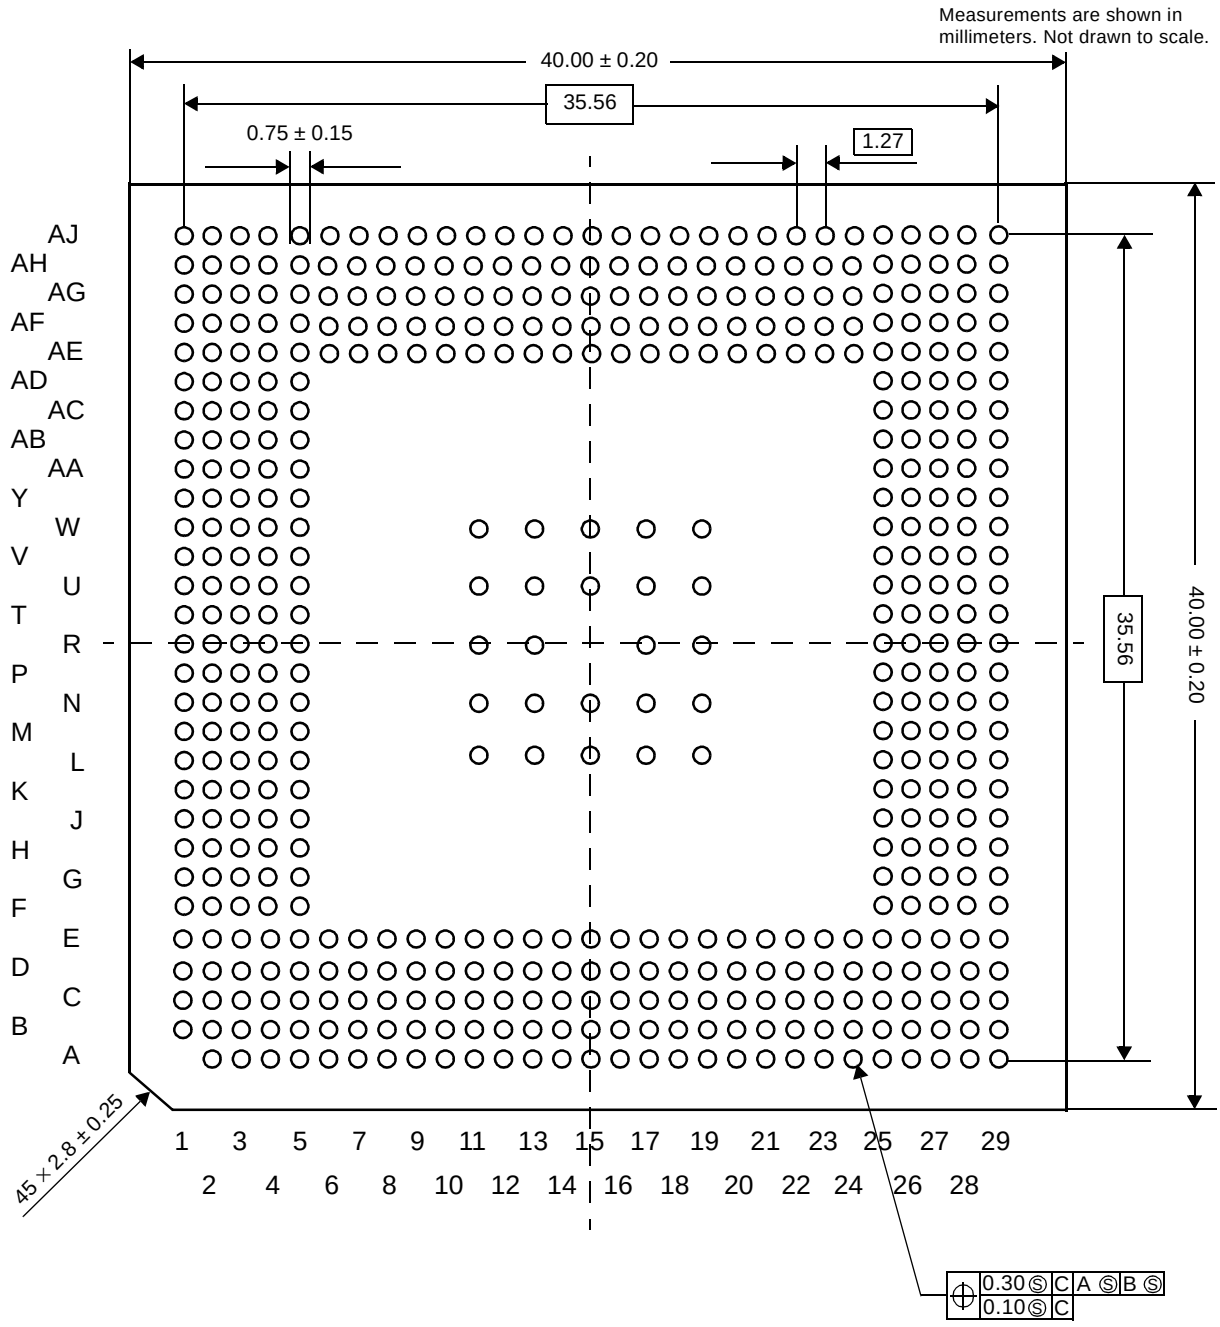

NOTES:

1. Controlling dimension: millimeter.
2. Ⓢ = Regardless of feature size.
3. PCB material: high temperature glass/epoxy resin cloth (that is, driclad, MCL-679, or equivalent). Solder resist: photoimagable (that is, vacrel 8130, DSR3241, PSR4000, or equivalent).
4. If you need a measurement not shown in this figure, contact PMC-Sierra.

Figure 49. 503-Pin EPBGA Bottom View (Part 2 of 2)

4.2 Signal Locations

Table 7. Signal Locations

Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name
A1	(KEY)	B1	VSS	C1	VDD	D1	VSS	E1	VSS
A2	VSS	B2	VDD	C2	ABR_RAM_AD(4)	D2	ABR_RAM_AD(7)	E2	ABR_RAM_AD(12)
A3	VDD	B3	JTAG_TDO	C3	/RESET	D3	ABR_RAM_AD(2)	E3	ABR_RAM_AD(6)
A4	VSS	B4	JTAG_TDI	C4	/JTAG_RESET	D4	/TEST_MODE	E4	ABR_RAM_AD(0)
A5	VDD	B5	JTAG_TCK	C5	JTAG_TMS	D5	/OE	E5	VDD
A6	/CS	B6	/INTR	C6	PCLK	D6	STATS_STRB	E6	VSS
A7	VSS	B7	ADDRDATA(30)	C7	/ADS	D7	/READY	E7	/SCAN_EN
A8	VSS	B8	ADDRDATA(31)	C8	ADDRDATA(28)	D8	ADDRDATA(27)	E8	W_/RD
A9	VSS	B9	ADDRDATA(20)	C9	ADDRDATA(21)	D9	ADDRDATA(29)	E9	ADDRDATA(24)
A10	VDD	B10	ADDRDATA(22)	C10	ADDRDATA(25)	D10	ADDRDATA(23)	E10	ADDRDATA(26)
A11	VDD	B11	ADDRDATA(12)	C11	ADDRDATA(19)	D11	ADDRDATA(15)	E11	ADDRDATA(17)
A12	VSS	B12	ADDRDATA(11)	C12	ADDRDATA(14)	D12	ADDRDATA(10)	E12	ADDRDATA(18)
A13	ADDRDATA(6)	B13	ADDRDATA(8)	C13	ADDRDATA(9)	D13	ADDRDATA(13)	E13	ADDRDATA(16)
A14	VDD	B14	ADDRDATA(5)	C14	ADDRDATA(7)	D14	ADDRDATA(3)	E14	ADDRDATA(2)
A15	VSS	B15	ADDRDATA(1)	C15	ADDRDATA(4)	D15	RATM_DATA(15)	E15	ADDRDATA(0)
A16	VDD	B16	SYS_CLK	C16	RATM_DATA(9)	D16	RATM_DATA(6)	E16	RATM_DATA(2)
A17	RATM_DATA(14)	B17	RATM_DATA(13)	C17	RATM_DATA(12)	D17	RATM_DATA(10)	E17	RATM_DATA(1)
A18	VSS	B18	RATM_DATA(11)	C18	RATM_DATA(4)	D18	RATM_DATA(5)	E18	RATM_DATA(3)
A19	VDD	B19	RATM_DATA(7)	C19	RATM_ADD(4)	D19	RATM_ADDR(0)	E19	RATM_CLAV(2)
A20	VDD	B20	RATM_DATA(8)	C20	/ RATM_READ_EN	D20	RATM_CLAV(0)	E20	RATM_ADD(1)
A21	VSS	B21	RATM_ADD(2)	C21	RATM_ADD(3)	D21	RATM_SOC	E21	RX_DRAM_DATA(30)
A22	VSS	B22	RATM_DATA(0)	C22	RATM_CLAV(3)	D22	RX_DRAM_DATA(29)	E22	RX_DRAM_DATA(25)

Table 7. Signal Locations (Continued)

Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name
A23	VSS	B23	RX_DRAM_DATA(31)	C23	RX_DRAM_DATA(26)	D23	RX_DRAM_DATA(23)	E23	RX_DRAM_DATA(16)
A24	RATM_CLAV(1)	B24	RX_DRAM_DATA(28)	C24	RX_DRAM_DATA(24)	D24	RX_DRAM_DATA(20)	E24	RX_DRAM_DATA(18)
A25	VSS	B25	RX_DRAM_DATA(27)	C25	RX_DRAM_DATA(21)	D25	RX_DRAM_DATA(15)	E25	VDD
A26	VSS	B26	RX_DRAM_DATA(22)	C26	RX_DRAM_DATA(17)	D26	VSS	E26	RX_DRAM_DATA(11)
A27	VDD	B27	RX_DRAM_DATA(19)	C27	RX_DRAM_DATA(14)	D27	RX_DRAM_DATA(8)	E27	RX_DRAM_DATA(9)
A28	VSS	B28	RX_DRAM_DATA(12)	C28	RX_DRAM_DATA(10)	D28	RX_DRAM_DATA(7)	E28	RX_DRAM_DATA(5)
A29	VDD	B29	VSS	C29	VDD	D29	VSS	E29	VSS

Table 7. Signal Locations (Continued)

Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name
F1	/ ABR_RAM_WE	K25	RX_DRAM_AD D(6)	P1	VDD	U25	SE_SOC_IN(0)	AA 1	VSS
F2	ABR_RAM_AD(13)	K26	RX_DRAM_AD D(2)	P2	CH_RAM_DAT A(0)	U26	BP_ACK_OUT(1)	AA 2	CH_RAM_DAT A(18)
F3	ABR_RAM_AD(9)	K27	/ RX_DRAM_RA S	P3	CH_RAM_ADD (9)	U27	BP_ACK_OUT(3)	AA 3	CH_RAM_DAT A(19)
F4	ABR_RAM_AD(5)	K28	RX_DRAM_AD D(0)	P4	CH_RAM_ADD (6)	U28	BP_ACK_IN(1)	AA 4	CH_RAM_DAT A(27)
F5	ABR_RAM_AD(3)	K29	VDD	P5	CH_RAM_ADD (2)	U29	BP_ACK_IN(2)	AA 5	CH_RAM_DAT A(22)
F25	RX_DRAM_DA TA(6)	L1	VDD	P25	SE_D_OUT0(3)	V1	VSS	AA 25	SE_D_IN0(2)
F26	RX_DRAM_DA TA(3)	L2	CH_RAM_ADD (7)	P26	SE_D_OUT1(1)	V2	CH_RAM_DAT A(10)	AA 26	SE_CLK
F27	RX_DRAM_DA TA(1)	L3	/CH_RAM_WE1	P27	SE_D_OUT1(0)	V3	CH_RAM_DAT A(13)	AA 27	SE_D_IN3(1)
F28	RX_DRAM_AD D(7)	L4	/ CH_RAM_ADS C	P28	SE_D_OUT1(2)	V4	CH_RAM_DAT A(9)	AA 28	SE_D_IN3(2)
F29	RX_DRAM_DA TA(2)	L5	/ABR_RAM_OE	P29	VDD	V5	CH_RAM_DAT A(16)	AA 29	VSS
G1	VSS	L11	VSS	R1	VSS	V25	SE_SOC_IN(1)	AB1	VSS
G2	ABR_RAM_AD(16)	L13	VSS	R2	CH_RAM_DAT A(4)	V26	SE_SOC_IN(3)	AB2	CH_RAM_DAT A(29)
G3	ABR_RAM_AD(11)	L15	VSS	R3	CH_RAM_ADD (16)	V27	ATM_CLK	AB3	CH_RAM_DAT A(26)
G4	ABR_RAM_AD(8)	L17	VSS	R4	CH_RAM_ADD (15)	V28	BP_ACK_OUT(0)	AB4	CH_RAM_DAT A(25)
G5	ABR_RAM_AD(1)	L19	VSS	R5	CH_RAM_ADD (17)	V29	VDD	AB5	CH_RAM_DAT A(31)
G25	RX_DRAM_DA TA(13)	L25	DRAM_CKE	R11	VSS	W1	VDD	AB2 5	TX_DRAM_DA TA(31)
G26	RX_DRAM_DA TA(4)	L26	SE_D_OUT3(3)	R13	VSS	W2	CH_RAM_DAT A(11)	AB2 6	SE_D_IN0(3)
G27	RX_DRAM_DA TA(0)	L27	SE_SOC_OUT			W3	CH_RAM_DAT A(17)	AB2 7	SE_D_IN1(3)
G28	RX_DRAM_AD D(5)	L28	SE_D_OUT3(0)	R17	VSS	W4	CH_RAM_DAT A(14)	AB2 8	SE_D_IN1(0)
G29	VSS	L29	VDD	R19	VSS	W5	CH_RAM_PARI TY0	AB2 9	VSS
H1	VSS	M1	VSS	R25	SE_D_OUT0(1)	W1 1	VSS	AC1	VSS

Table 7. Signal Locations (Continued)

Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name
H2	CH_RAM_ADD(0)	M2	CH_RAM_ADD(11)	R26	SE_D_OUT0(0)	W13	VSS	AC2	CH_RAM_DATA(28)
H3	/ABR_RAM_ADV	M3	CH_RAM_ADD(4)	R27	BP_ACK_IN(3)	W15	VSS	AC3	CH_RAM_DATA(30)
H4	ABR_RAM_AD(14)	M4	CH_RAM_ADD(5)	R28	SE_D_OUT0(2)	W17	VSS	AC4	/ALRAM_ADSC
H5	ABR_RAM_AD(10)	M5	CH_RAM_ADD(3)	R29	VSS	W19	VSS	AC5	ALRAM_ADD(4)
H25	/RX_DRAM_CS(1)	M25	/RX_DRAM_CAS	T1	VDD	W25	SE_D_IN2(0)	AC25	TX_DRAM_DATA(22)
H26	RX_DRAM_ADD(4)	M26	SE_D_OUT2(2)	T2	CH_RAM_DATA(6)	W26	SE_D_IN2(3)	AC26	TX_DRAM_DATA(27)
H27	/RX_DRAM_WE	M27	SE_D_OUT2(0)	T3	CH_RAM_DATA(3)	W27	SE_D_IN3(3)	AC27	TX_DRAM_DATA(30)
H28	/RX_DRAM_CS(0)	M28	SE_D_OUT3(1)	T4	CH_RAM_DATA(2)	W28	RX_CELL_START	AC28	SE_D_IN0(1)
H29	VSS	M29	VSS	T5	CH_RAM_DATA(1)	W29	VDD	AC29	VDD
J1	VSS	N1	CH_RAM_ADD(14)	T25	SE_SOC_IN(2)	Y1	VDD	AD1	ALRAM_CLK
J2	/CH_RAM_OE	N2	CH_RAM_ADD(13)	T26	PROC_MON	Y2	CH_RAM_DATA(20)	AD2	ALRAMADD18N
J3	/CH_RAM_WE0	N3	CH_RAM_ADD(12)	T27	BP_ACK_OUT(2)	Y3	CH_RAM_DATA(23)	AD3	CH_RAM_PARITY1
J4	/ABR_RAM_ADSP	N4	CH_RAM_ADD(10)	T28	BP_ACK_IN(0)	Y4	CH_RAM_DATA(21)	AD4	ALRAMADD17N
J5	ABR_RAM_AD(15)	N5	CH_RAM_ADD(1)	T29	VSS	Y5	CH_RAM_DATA(24)	AD5	VSS
J25	RX_DRAM_ADD(1)	N11	VSS	U1	CH_RAM_DATA(5)	Y25	SE_D_IN2(2)	AD25	VSS
J26	RX_DRAM_ADD(8)	N13	VSS	U2	CH_RAM_DATA(7)	Y26	SE_D_IN1(2)	AD26	TX_DRAM_DATA(26)
J27	RX_DRAM_ADD(3)	N15	VSS	U3	CH_RAM_DATA(8)	Y27	SE_D_IN2(1)	AD27	TX_DRAM_DATA(28)
J28	RX_DRAM_CLK	N17	VSS	U4	CH_RAM_DATA(12)	Y28	SE_D_IN3(0)	AD28	SE_D_IN0(0)
J29	VSS	N19	VSS	U5	CH_RAM_DATA(15)	Y29	VDD	AD29	SE_D_IN1(1)
K1	VDD	N25	RX_DRAM_BA	U11	VSS				

Table 7. Signal Locations (Continued)

Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name
K2	CH_RAM_ADD (8)	N26	SE_D_OUT3(2)	U13	VSS				
K3	CH_RAM_CLK	N27	SE_D_OUT2(3)	U15	VSS				
K4	ABR_RAM_CL K	N28	SE_D_OUT1(3)	U17	VSS				
K5	CH_RAM_ADD 17N	N29	SE_D_OUT2(1)	U19	VSS				

Table 7. Signal Locations (Continued)

Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name
AE1	VSS	AF1	VSS	AG1	VDD	AH1	VSS	AJ1	VDD
AE2	/ALRAM_OE	AF2	/ALRAM_WE	AG2	ALRAM_ADD(1)	AH2	VSS	AJ2	VSS
AE3	ALRAM_ADD(0)	AF3	ALRAM_ADD(2)	AG3	ALRAM_ADD(6)	AH3	VDD	AJ3	VDD
AE4	ALRAM_ADD(5)	AF4	ALRAM_ADD(3)	AG4	ALRAM_ADD(10)	AH4	ALRAM_ADD(12)	AJ4	VSS
AE5	VDD	AF5	ALRAM_ADD(7)	AG5	ALRAM_ADD(11)	AH5	ALRAM_ADD(15)	AJ5	VDD
AE6	ALRAM_ADD(8)	AF6	ALRAM_ADD(13)	AG6	ALRAM_ADD(16)	AH6	ALRAM_DATA(0)	AJ6	ALRAM_DATA(5)
AE7	ALRAM_ADD(9)	AF7	ALRAM_ADD(14)	AG7	ALRAM_ADD(17)	AH7	ALRAM_DATA(1)	AJ7	VSS
AE8	ALRAM_ADD(18)	AF8	ALRAM_DATA(3)	AG8	ALRAM_DATA(7)	AH8	ALRAM_DATA(4)	AJ8	VSS
AE9	ALRAM_DATA(2)	AF9	ALRAM_DATA(8)	AG9	ALRAM_DATA(14)	AH9	ALRAM_DATA(15)	AJ9	VSS
AE10	ALRAM_DATA(11)	AF10	ALRAM_DATA(6)	AG10	ALRAM_DATA(10)	AH10	ALRAM_DATA(13)	AJ10	VDD
AE11	ALRAM_DATA(9)	AF11	ALRAM_DATA(12)	AG11	ALRAM_DATA(16)	AH11	TATM_CLAV(3)	AJ11	VDD
AE12	TATM_SOC	AF12	TATM_CLAV(0)	AG12	TATM_CLAV(1)	AH12	TATM_DATA(0)	AJ12	VSS
AE13	TATM_PARITY	AF13	TATM_DATA(1)	AG13	TATM_DATA(3)	AH13	TATM_DATA(5)	AJ13	TATM_DATA(6)
AE14	/ TATM_WRITE_EN	AF14	TATM_CLAV(2)	AG14	TATM_DATA(2)	AH14	TATM_DATA(4)	AJ14	VDD
AE15	TATM_DATA(9)	AF15	TATM_DATA(8)	AG15	TATM_DATA(7)	AH15	TATM_DATA(10)	AJ15	VSS
AE16	TATM_DATA(11)	AF16	TATM_DATA(13)	AG16	TATM_DATA(12)	AH16	TATM_DATA(14)	AJ16	VDD
AE17	TX_DRAM_ADD(0)	AF17	TX_DRAM_BA	AG17	TATM_ADD(3)	AH17	TATM_DATA(15)	AJ17	TATM_ADD(1)
AE18	TX_DRAM_ADD(1)	AF18	TATM_ADD(2)	AG18	TATM_ADD(0)	AH18	TX_DRAM_CLK	AJ18	VSS
AE19	/ TX_DRAM_WE	AF19	/ TX_DRAM_CAS	AG19	/ TX_DRAM_RAS	AH19	TATM_ADD(4)	AJ19	VDD
AE20	TX_DRAM_DATA(0)	AF20	TX_DRAM_ADD(7)	AG20	TX_DRAM_ADD(3)	AH20	TX_DRAM_ADD(5)	AJ20	VDD

Table 7. Signal Locations (Continued)

Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name	Ball	Signal Name
AE21	TX_DRAM_ADD(6)	AF21	TX_DRAM_DATA(2)	AG21	TX_DRAM_ADD(8)	AH21	TX_DRAM_DD(2)	AJ21	VSS
AE22	TX_DRAM_DATA(4)	AF22	/ TX_DRAM_CS(0)	AG22	TX_DRAM_ADD(4)	AH22	TX_DRAM_DATA(3)	AJ22	VSS
AE23	TX_DRAM_DATA(18)	AF23	TX_DRAM_DATA(9)	AG23	TX_DRAM_DATA(5)	AH23	/ TX_DRAM_CS(1)	AJ23	VSS
AE24	TX_DRAM_DATA(11)	AF24	TX_DRAM_DATA(8)	AG24	TX_DRAM_DATA(6)	AH24	TX_DRAM_DATA(1)	AJ24	TX_DRAM_DATA(7)
AE25	VDD	AF25	TX_DRAM_DATA(16)	AG25	TX_DRAM_DATA(14)	AH25	TX_DRAM_DATA(10)	AJ25	VSS
AE26	TX_DRAM_DATA(20)	AF26	TX_DRAM_DATA(17)	AG26	TX_DRAM_DATA(13)	AH26	TX_DRAM_DATA(12)	AJ26	VSS
AE27	TX_DRAM_DATA(25)	AF27	TX_DRAM_DATA(23)	AG27	TX_DRAM_DATA(19)	AH27	TX_DRAM_DATA(15)	AJ27	VDD
AE28	TX_DRAM_DATA(29)	AF28	TX_DRAM_DATA(24)	AG28	TX_DRAM_DATA(21)	AH28	VSS	AJ28	VSS
AE29	VSS	AF29	VSS	AG29	VSS	AH29	VDD	AJ29	VDD

4.3 Signal Descriptions (372 Signal Pins)

All inputs and Bidirectional inputs have internal pull up circuit except for /OE input. /OE has an internal pull down circuit.
All 5V tolerant/ LVTTL inputs have a Schmitt Trigger Hysteresis circuit.
All CMOS inputs are not 5V tolerant.

4.3.1 Processor Interface Signals

Table 8. Processor Interface Signals (38 Pins)

Signal Name	Ball	Type	Drive/ Input Level	Slew Rate	Description
PCLK	C6	In	CMOS		*Processor Clock*
ADDRDATA(31:0)	B8, B7, D9, C8, D8, E10, C10, E9, D10, B10, C9, B9, C11, E12, E11, E13, D11, C12, D13, B11, B12, D12, C13, B13, C14, A13, B14, C15, D14, E14, B15, E15	Bidir	5 ma/ 5 V or LV TTL	Mod-erate (Mod)	*Address/Data Bits 31 to 0* are part of the 32-bit processor address/data bus.
/ADS	C7	In	5 V or LV TTL		*Address/Data Status* is an active low signal that indicates an address state.
W_/RD	E8	In	5 V or LV TTL		*Write_/Read* is an active high signal that selects a write cycle when it is high.
/READY	D7	Out	5 ma	Mod	*Ready* is an active low signal that indicates the processor cycle is finished. When this signal is deasserted, it is driven high, then tristated.
/CS	A6	In	5 V or LV TTL		*Chip Select* is an active low signal that selects the device for processor access.
/INTR	B6	Out	5 ma	Mod	*Interrupt* is an active low signal that indicates an interrupt is present.

4.3.2 Statistics Interface Signal

Table 9. Statistics Interface Signal (1 Pin)

Signal Name	Ball	Type	Drive/ Input Level	Slew Rate	Description
STATS_STRB	D6	Out	8 ma	Mod	*STATS_STRB* is an active high signal that indicates a fixed position in the cell time in the SYS_CLK domain. This can be used to trigger external circuitry.

4.3.3 Switch Element Interface Signals

Table 10. Switch Element Interface Signals (47 Pins)

Signal Name	Ball	Type	Drive/ Input Level	Slew Rate	Description
SE_CLK	AA26	In	CMOS		*Switch Element Clock* is the 50 MHz or 66 MHz clock for nibble transfer.
RX_CELL_START	W28	In	5 V or LV TTL		*Receive Cell Start* indicates the SOC time in the receive direction. It should be driven high every cell time (118 SE_CLKs).
BP_ACK_IN(3:0)	R27, U29, U28, T28	In	5 V or LV TTL		*Backpressure Input 3 down to 0.* It carries the cell acknowledge and backpressure from the switch fabric.
SE_SOC_OUT	L27	Out	8 ma	Mod	*Switch Fabric Start Of Cell Out* indicates the SOC for all four SE_D_OUT signals. This signal precedes the first nibble of cell by one clock. For cells leaving the QRT and entering the switch fabric, this signal indicates the SOC.
SE_D_OUT0(3:0)	P25, R28, R25, R26	Out	5 ma	Mod	*Switch Element Data Out Ports 3 down to 0 Bits 3 down to 0* are four nibble-wide pathways that carry the cell to the QSEs (PM73488).
SE_D_OUT1(3:0)	N28, P28, P26, P27	Out	5 ma	Mod	
SE_D_OUT2(3:0)	N27, M26, N29, M27	Out	5 ma	Mod	
SE_D_OUT3(3:0)	L26, N26, M28, L28	Out	5 ma	Mod	
BP_ACK_OUT(3:0)	U27, T27, U26, V28	Out	8 ma	Mod	*Backpressure Output 3 down to 0* asserts multipriority backpressure and cell acknowledge toward the switch fabric.
SE_SOC_IN(3:0)	V26, T25, V25, U25	In	5 V or LV TTL		*Switch Fabric Start of Cell 3 to 0* indicates the SOC time in the transmit direction for the four incoming SE_D_IN3, SE_D_IN2, SE_D_IN1 and SE_D_IN0, respectively.
SE_D_IN0(3:0)	AB26, AA25, AC28, AD28	In	5 V or LV TTL		*Switch Element Data In Ports 3 to 0 Bits 3 down to 0* are part of the nibble-wide, 50 MHz data pathway that carries the cell from the switch fabric.
SE_D_IN1(3:0)	AB27, Y26, AD29, AB28	In	5 V or LV TTL		
SE_D_IN2(3:0)	W26, Y25, Y27, W25	In	5 V or LV TTL		
SE_D_IN3(3:0)	W27, AA28, AA27, Y28	In	5 V or LV TTL		

4.3.4 CH_RAM Interface Signals

Table 11. CH_ RAM Interface Signals (58 Pins)

Signal Name	Ball	Type	Drive/ Input Level	Slew Rate	Description
CH_RAM_ADD(17:0)	R5, R3, R4, N1, N2, N3, M2, N4, P3, K2, L2, P4, M4, M3, M5, P5, N5, H2	Out	5 ma	Mod	*CH_ RAM Address Bits 17 to 0* are part of the 18-bit SRAM address bus.
CH_RAM_DATA(31:0)	AB5, AC3, AB2, AC2, AA4, AB3, AB4, Y5, Y3, AA5, Y4, Y2, AA3, AA2, W3, V5, U5, W4, V3, U4, W2, V2, V4, U3, U2, T2, U1, R2, T3, T4, T5, P2	Bidir	5 ma/CMOS	Mod	*CH_RAM Data Bits 31 to 0* are part of the 32-bit SRAM data bus.
CH_RAM_PARITY0	W5	Bidir	5 ma/CMOS	Mod	Odd parity bit for CH_RAM_DATA(15:0).
CH_RAM_PARITY1	AD3	Bidir	5 ma/CMOS	Mod	Odd parity bit for CH_RAM_DATA(31:16).
CH_RAM_CLK	K3	Out	8 ma	Fast	*CH_RAM Clock* provides the clock to the CH_RAM.This signal should be terminated with a series resistor before connecting to the RAM modules
CH_RAM_ADD17N	K5	Out	5 ma	Mod	*CH_RAM Not Address Bit 17* reverses bit 17 of CH_RAM_ADD(17:0).
/CH_RAM_OE	J2	Out	8 ma	Fast	*CH_RAM Output Enable* is an active low signal that enables the SRAM to drive the CH_RAM_DATA(31:0), CH_RAM_PARITY0, and CH_RAM_PARITY1.This signal should be terminated with a series resistor before connecting to the RAM modules
/CH_RAM_WE0	J3	Out	5 ma	Mod	*CH_RAM Write Enable 0* is an active low signal that strobes CH_RAM_DATA(15:0) and CH_RAM_PARITY0 into an external SRAM.
/CH_RAM_WE1	L3	Out	5 ma	Mod	*CH_RAM Write Enable 1* is an active low signal that strobes CH_RAM_DATA(31:16) and CH_RAM_PARITY1 into an external SRAM.

Table 11. CH_ RAM Interface Signals (58 Pins) (Continued)

Signal Name	Ball	Type	Drive/ Input Level	Slew Rate	Description
/CH_RAM_ADSC	L4	Out	5 ma	Mod	*CH_RAM Synchronous Address Status Controller* is an active low signal that causes new addresses to be registered within the external SSRAM.

4.3.5 AL_RAM Interface Signals

Table 12. Address Lookup RAM Interface Signals (42 Pins)

Signal Name	Ball	Type	Drive/ Input Level	Slew Rate	Description
ALRAM_ADD(18:0)	AE8, AG7, AG6, AH5, AF7, AF6, AH4, AG5, AG4, AE7, AE6, AF5, AG3, AE4, AC5, AF4, AF3, AG2, AE3	Out	5 ma	Mod	*AL RAM Address Bits 18 to 0* are part of the 19-bit SRAM address bus.
ALRAM_DATA(16:0)	AG11, AH9, AG9, AH10, AF11, AE10, AG10, AE11, AF9, AG8, AF10, AJ6, AH8, AF8, AE9, AH7, AH6	Bidir	5 ma/CMOS	Mod	*AL RAM Data Bits 15 to 0* are part of the 16-bit SRAM data bus. Bit 16 is for parity.
ALRAM_CLK	AD1	Out	8 ma	Fast	*AL RAM Clock* provides the clock to the ALRAM. This signal should be terminated with a series resistor before connecting to the RAM modules
ALRAMADD17N	AD4	Out	5 ma	Mod	*AL RAM Not Address 17* reverses bit 17 of ALRAM_ADD(18:0).
ALRAMADD18N	AD2	Out	5 ma	Mod	*AL RAM Not Address 18* reverses bit 18 of ALRAM_ADD(18:0).
/ALRAM_OE	AE2	Out	8 ma	Fast	*AL RAM Output Enable* is an active low signal that enables the SRAM to drive AL_RAM_DATA(16:0).This signal should be terminated with a series resistor before connecting to the RAM modules
/ALRAM_WE	AF2	Out	5 ma	Mod	*AL RAM Write Enable* is an active low signal that strobes data into an external SRAM.
/ALRAM_ADSC	AC4	Out	5 ma	Mod	*AL RAM Synchronous Address Status Controller* is an active low signal that causes new addresses to be registered within the external SSRAM.

4.3.6 ABR_RAM Interface Signals

Table 13. ABR_RAM Interface Signals (22 Pins)

Signal Name	Ball	Type	Drive/ Input Level	Slew Rate	Description
ABR_RAM_AD(16:0)	G2, J5, H4, F2, E2, G3, H5, F3, G4, D2, E3, F4, C2, F5, D3, G5, E4	Bidir	5 ma/CMOS	Mod	*ABR RAM Address Data Bits 16 to 0* form the time division multiplexed address data bus.
ABR_RAM_CLK	K4	Out	8 ma	Fast	*ABR RAM Clock* provides the clock to the AB RAM.This signal should be terminated with a series resistor before connecting to the RAM modules
/ABR_RAM_ADSP	J4	Out	5 ma	Mod	*ABR RAM Address Data Selection* defines the type of information on the address/data bus (ADDRDATA(31:0)).
/ABR_RAM_OE	L5	Out	8 ma	Fast	*ABR_RAM Output Enable* is an active low signal that enables the RAM to drive ABR_RAM_AD(16:0).This signal should be terminated with a series resistor before connecting to the RAM modules
/ABR_RAM_ADV	H3	Out	5 ma	Mod	*ABR_RAM Advance* is an active low signal that signals the external SSRAM to advance its address.
/ABR_RAM_WE	F1	Out	5 ma	Mod	*ABR RAM Write Enable* is an active low signal that enables a write into the ABR_RAM.

4.3.7 Receive Cell Buffer DRAMs

Table 14. Receive Cell Buffer RAM Interface Signals (49 Pins)

Signal Name	Ball	Type	Drive/ Input Level	Slew Rate	Description
RX_DRAM_ADD(8:0)	J26, F28, K25, G28, H26, J27, K26, J25, K28	Out	5 ma	Mod	*RX DRAM Address Bits 8 to 0* are part of the 9-bit SRAM address bus. Note that TX_DRAM_ADD(8) must be connected to the AutoPre-charge pin.If DRAM_TYPE = 1 (refer to "RX_DRAM_TYPE" on page 104) then connect RX_DRAM_ADD(8) to the DRAM autoprecharge pin, which should be bit 10 of the DRAM address bus
RX_DRAM_DATA(31:0)	B23, E21, D22, B24, B25, C23, E22, C24, D23, B26, C25, D24, B27, E24, C26, E23, D25, C27, G25, B28, E26, C28, E27, D27, D28, F25, E28, G26, F26, F29, F27, G27	Bidir	5 ma/CMOS	Mod	*Receive DRAM Data Bits 31 to 0* are part of the 32-bit SRAM data bus.
RX_DRAM_CLK	J28	Out	8 ma	Fast	*Receive DRAM Clock* provides the clock to the SDRAM.This signal should be terminated with a series resistor before connecting to the RAM modules
DRAM_CKE	L25	Out	5 ma	Mod	*DRAM Clock Enable* provides a clock enable signal for RX_DRAM and TX_DRAM
/RX_DRAM_CS(1:0)	H25, H28	Out	5 ma	Mod	*Receive DRAM Chip Select Bits 1 to 0* enable the SDRAMs. If DRAM_TYPE = 1 (refer to "RX_DRAM_TYPE" on page 104), these are RX_DRAM_ADD(9:8).
RX_DRAM_BA	N25	Out	5 ma	Mod	*Receive DRAM Bank Address* defines the bank to which the opera-tion is addressed.
/RX_DRAM_RAS	K27	Out	5 ma	Mod	*Receive DRAM Row Address Strobe* is an active low signal that writes in the row address.
/RX_DRAM_CAS	M25	Out	5 ma	Mod	*Receive DRAM Column Address Strobe* is an active low signal that writes in the column address.
/RX_DRAM_WE	H27	Out	5 ma	Mod	*Receive DRAM Write Enable* is an active low signal that enables a write into the synchronous DRAM.

NOTE: DQM (I/O mask enables) pins to the SGRAM need to be tied to logic 0.

4.3.8 Transmit Cell Buffer DRAMs

Table 15. Transmit Cell Buffers RAM Interface Signals (48 Pins)

Signal Name	Ball	Type	Drive/ Input Level	Slew Rate	Description
TX_DRAM_ADD(8:0)	AG21, AF20, AE21, AH20, AG22, AG20, AH21, AE18, AE17	Out	5 ma	Mod	*Transmit DRAM Address Bits 8 to 0* are part of the 9-bit DRAM address bus. Note that TX_DRAM_ADD(8) must be connected to the AutoPrecharge pin. If DRAM_TYPE = 1 (refer to "RX_DRAM_TYPE" on page 104) then connect TX_DRAM_ADD(8) to the DRAM autoprecharge pin, which should be bit 10 of the DRAM address bus.
TX_DRAM_DATA(31:0)	AB25, AC27, AE28, AD27, AC26, AD26, AE27, AF28, AF27, AC25, AG28, AE26, AG27, AE23, AF26, AF25, AH27, AG25, AG26, AH26, AE24, AH25, AF23, AF24, AJ24, AG24, AG23, AE22, AH22, AF21, AH24, AE20	Bidir	5 ma/CMOS	Mod	*Transmit DRAM Data Bits 31 to 0* are part of the 32-bit DRAM data bus.
TX_DRAM_CLK	AH18	Out	8 ma	Fast	*Transmit DRAM Clock* provides the clock to the SDRAM.This signal should be terminated with a series resistor before connecting to the RAM modules
/TX_DRAM_CS(1:0)	AH23, AF22	Out	5 ma	Mod	*Transmit DRAM Chip Select Bits 1 and 0* select the SDRAM devices. If DRAM_TYPE = 1 (refer to "RX_DRAM_TYPE" on page 104), then these are TX_DRAM_ADD(9:8).
TX_DRAM_BA	AF17	Out	5 ma	Mod	*Transmit DRAM Bank Address* defines the bank to which the operation is addressed.
/TX_DRAM_RAS	AG19	Out	5 ma	Mod	*Transmit DRAM Row Address Strobe* is an active low signal that writes in the row address.
/TX_DRAM_CAS	AF19	Out	5 ma	Mod	*Transmit DRAM Column Address Strobe* is an active low signal that writes in the column address.

Table 15. Transmit Cell Buffers RAM Interface Signals (48 Pins) (Continued)

Signal Name	Ball	Type	Drive/ Input Level	Slew Rate	Description
/TX_DRAM_WE	AE19	Out	5 ma	Mod	*Transmit DRAM Write Enable* is an active low signal that enables a write into the synchronous DRAM.

NOTE: DQM (I/O mask enables) pins to the SGRAM or SDRAM need to be tied to logic 0.

4.3.9 UTOPIA ATM Layer Interface Signals

4.3.9.1 Transmit UTOPIA ATM Layer Interface Signals

Table 16. Transmit UTOPIA ATM Layer Interface Signals (29 Pins)

Signal Name	Ball	Type	Drive/Input Level	Slew Rate	Description
ATM_CLK	V27	In	CMOS		*UTOPIA ATM Layer Clock* provides timing information for both the transmit and receive UTOPIA interfaces.
TATM_ADD(4:0)	AH19, AG17, AF18, AJ17, AG18	Out	12 ma	Mod	*Transmit UTOPIA ATM Layer Address Bits 4 to 0.*
TATM_DATA(15:0)	AH17, AH16, AF16, AG16, AE16, AH15, AE15, AF15, AG15, AJ13, AH13, AH14, AG13, AG14, AF13, AH12	Out	12 ma	Mod	*Transmit UTOPIA ATM Layer Data Bits 15 to 0* are part of the 16-bit UTOPIA transmit data bus that carries data toward the PHY layer device.
TATM_CLAV(3:0)	AH11, AF14, AG12, AF12	In	5 V or LV TTL		*Transmit UTOPIA ATM Layer Cell Available Bits 3 to 0* are active high signals that indicate the selected PHY layer devices may accept another cell.
TATM_SOC	AE12	Out	12 ma	Mod	*Transmit UTOPIA ATM Layer Start-Of-Cell* marks the start of a cell.
/TATM_WRITE_EN	AE14	Out	12 ma	Mod	*Transmit UTOPIA ATM Layer Write Enable* enables the write of a cell byte.
TATM_PARITY	AE13	Out	12 ma	Mod	*Transmit UTOPIA ATM Layer Odd Parity* bit over TATM_DATA(15:0).

4.3.9.2 Receive UTOPIA ATM Layer Interface Signals

Table 17. Receive UTOPIA ATM Layer Interface Signals (27 Pins)

Signal Name	Ball	Type	Drive/ Input Level	Slew Rate	Description
RATM_ADD(4:0)	C19, C21, B21, E20, D19	Out	12 ma	Mod	*Receive UTOPIA ATM Layer Address Bits 4 to 0*
RATM_DATA(15:0)	D15, A17, B17, C17, B18, D17, C16, B20, B19, D16, D18, C18, E18, E16, E17, B22	In	5V or LV TTL		*Receive UTOPIA ATM Layer Data Bits 15 to 0 are* part of the 16-bit UTOPIA receive data bus that carries data from the PHY layer device.
RATM_SOC	D21	In	5V or LV TTL		*Receive UTOPIA ATM Layer Start-Of-Cell* marks the start of a cell.
RATM_CLAV(3:0)	C22, E19, A24, D20	In	5V or LV TTL		*Receive UTOPIA ATM Layer Cell Available Bits 0 to* 3 indicate the selected PHY layer devices have another cell.
/RATM_READ_EN	C20	Out	12 ma	Mod	*Receive UTOPIA ATM Layer Read Enable* causes a read from the FIFO in the PHY layer device.

4.3.10 Boundary Scan Signals

Table 18. Test Signals (8 Signal Pins)

Signal Name	Ball	Type	Drive/Input Level	Slew Rate	Description
/JTAG_RESET	C4	In	CMOS		*JTAG Reset* is an active low, true asynchronous reset to the JTAG controller.
JTAG_TCK	B5	In	CMOS		*Scan Test Clock* is an independent clock used to drive the internal boundary scan test logic. Connect this signal to VDD through a pull-up resistor.
JTAG_TDI	B4	In	CMOS		*Scan Test Data Input* is the serial input for boundary scan test data and instruction bits. Connect this signal to VDD through a pull-up resistor.
JTAG_TDO	B3	Out	4 ma	Mod	*Scan Test Data Output* is the serial output for boundary scan test data.
JTAG_TMS	C5	In	CMOS		*Scan Test Mode Select* controls the operation of the boundary scan test logic. Connect this signal to VDD through a pull-up resistor.
/TEST_MODE	D4	In	CMOS		This is a manufacturing test mode bit for manufacturing test. It MUST be pulled up for functional mode on the board.
/SCAN_EN	E7	In	CMOS		*Scan Test Enable* is used to enable the internal scan test logic. Connect this signal to VDD through a pull-up resistor.
PROC_MON	T26	Out	N/A	N/A	*Process Monitor* is used for manufacturing test. It is connected to a NAND tree that may be used for VIL/VIH testing.

4.3.11 Miscellaneous Signals

Table 19. Miscellaneous Signals (3 Signal Pins)

Signal Name	Ball	Type	Drive /Input Level	Slew	Description
SYSCLK	B16	In	CMOS		*System Clock* provides a high speed clock input for the state machine and the memory interfaces.
/OE	D5	In	CMOS		*Output Enable* is an active low signal that enables all the outputs of the device. Setting it high will tri-state all outputs except PROCMON and disable all input pull up resistors for in-circuit IDD tests.
/RESET	C3	In	5V or LV TTL		*Reset* is an active low signal used to initialize or re-initialize the device. SE_CLK must be present for the reset to take effect.
VDD	A3, A5, A10, A11, A14, A16, A19, A20, A27, A29, B2, C1, C29, K1, K29, L1, L29, P1, P29, T1, V29, W1, W29, Y1, Y29, AC29, AE5, AG1, AH3, AH29, AJ1, AJ3, AJ5, AJ10, AJ11, AJ14, AJ16, AJ19, AJ20, AJ27, AJ29, E5, E25, AE25	In	N/A		Supply voltage 3.3 ± 5% V.
VSS	A2, A4, A7, A8, A9, A12, A15, A18, A21, A22, A23, A25, A26, A28, B1, B29, D1, D26, D29, E1, E6, E29, G1, G29, H1, H29, J1, J29, L11, L13, L15, L17, L19, M1, M29, N11, N13, N15, N17, N19, R1, R11, R13, R17, R19, R29, T29, U11, U13, U15, U17, U19, V1, W11, W13, W15, W17, W19, AA1, AA29, AB1, AB29, AC1, AD5, AD25, AE1, AE29, AF1, AF29, AG29, AH1, AH2, AH28, AJ2, AJ4, AJ7, AJ8, AJ9, AJ12, AJ15, AJ18, AJ21, AJ22, AJ23, AJ25, AJ26, AJ28	In	N/A		Ground.

5 PHYSICAL CHARACTERISTICS

Symbol	Parameter	Conditions	Min	Max	Unit
V_{DD}	Supply voltage	With respect to GND	-0.3	3.9	V
I_{OUT}	DC output current, per pin		-12	12	mA
T_{STG}	Storage temperature		-65	125	°C
t_R	Input rise time			10	ns
t_F	Input fall time			10	ns

.

Symbol	Parameter	Conditions	Min	Max	Unit
V_{DD}	Supply voltage		3.14	3.46	V
V_O	Output voltage		0	V_{DD}	V
T_A	Operating temperature		-40	85	°C

.

Symbol	Parameter	Conditions	Min	Typ	Max	Unit
V_{IHT}	High-level TTL input voltage	All 5V tolerant /LVTTL inputs	2.15	V_{DD}	5.5V	V
V_{IHC}	High-level CMOS input voltage	All CMOS inputs	2.0	V_{DD}	V_{DD} + 0.3	V
V_{IL}	Low-level input voltage	All 5V tolerant /LVTTL and CMOS inputs	V_{SS} - 0.3	0	0.8	V
V_{OH}	CMOS high-level output voltage	I_{OH} = - Specified DC Drive current (in Pin Description section) assuming the UTOPIA interface drivers are not simultaneously loaded at 12 ma	2.4		V_{DD}	V
V_{OL}	CMOS low-level output voltage	I_{OL} = Specified DC Drive current(in Pin Description section) assuming the UTOPIA interface drivers are not simultaneously loaded at 12 ma			0.4	V
I_{TYP}	Typical operating current, output unloaded	SE_CLK = 50 MHz SYS_CLK= 100 MHZ ATM_CLK = 66 MHZ		600	900	mA

NOTES:
- V_{CC} = 3.3V ± 10%, T_A = -40°C to 85°C for industrial use.
- Typical values are T_A = 25°C and V_{CC} = 3.3 V.

Symbol	Parameter	Conditions	Min	Max	Unit
C_{IN}	Input capacitance			6	pF
C_{OUT}	Output capacitance			6	pF
C_{LOAD}	Load capacitance	To meet timing on any signal.		30	pF
NOTES: • Capacitance measured at 25°C. • Sample tested only.					

Table 20. Estimated Package Thermal Characteristics

Symbol	Parameter	Condition	Typical	Unit
θ_{JC}	Junction to Case thermal resistance		3.3	°C/Watt
θ_{JA}	Junction to Ambient thermal resistance	Still air	14.9	°C/Watt
θ_{JA}	Junction to Ambient thermal resistance	200 lfpm	12.6	°C/Watt
θ_{JA}	Junction to Ambient thermal resistance	400 lfpm	11.8	°C/Watt
θ_{JA}	Junction to Ambient thermal resistance	600 lfpm	11.2	°C/Watt

NOTE:

- The junction temperature must be kept below 125°C while the device is operating.

6 TIMING DIAGRAMS

All pin names are described in section 4 “Pin Descriptions” starting on page 55. Unless otherwise indicated, all output timing delays assume a capacitive loading of 30 pF

6.1 UTOPIA Timing

Figure 50 shows the receive UTOPIA 50 MHz timing (the RATM_ADD and RATM_CLAV values are shown in hexadecimal).

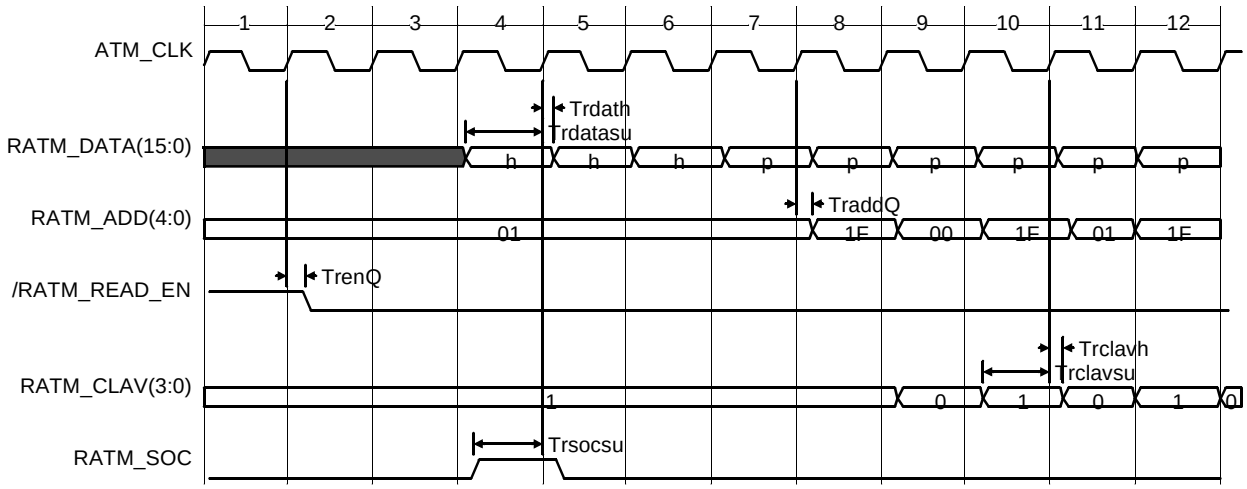

Figure 50. Receive UTOPIA 50 MHz Timing

Symbol	Parameter	Signals	Min	Max	Unit
	ATM_CLK frequency	ATM_CLK		55	MHz
	ATM_CLK duty cycle	ATM_CLK	40	60	%
TraddQ	Clock-to-output valid time	RATM_ADD(4:0)	3.4	10.5	ns
Trclavsu	Input setup time	RATM_CLAV(3:0)	4		ns
Trclavh	Input hold time	RATM_CLAV(3:0)	1		ns
TrenQ	Clock-to-output valid time	/RATM_READ_EN	3.5	9.5	ns
Trsocsu	Input setup time	RATM_SOC	4		ns
Trdatasu	Input setup time	RATM_DATA(15:0)	4		ns
Trdath	Input hold time	RATM_DATA(15:0)	1		ns

Figure 51 shows the transmit UTOPIA 50 MHz timing.

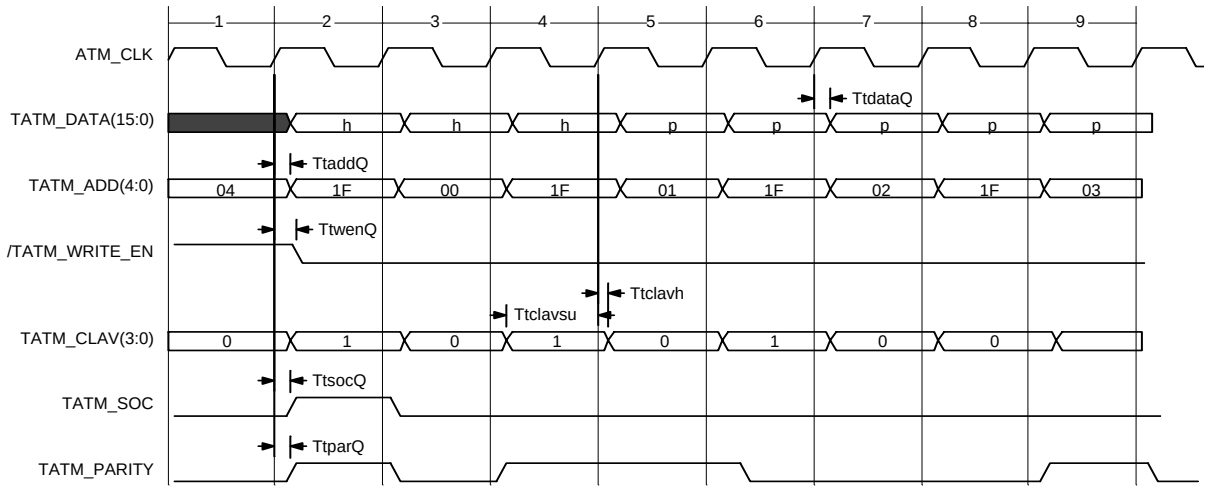

Figure 51. Transmit UTOPIA 50 MHz Timing

Symbol	Parameter	Signals	Min	Max	Unit
	ATM_CLK frequency	ATM_CLK		55	MHz
TtaddQ	Clock-to-output valid time	TATM_ADD(4:0)	3.4	12	ns
Ttclavsu	Input setup time	TATM_CLAV(3:0)	2		ns
TtwenQ	Clock-to-output valid time	/TATM_WRITE_EN	3.3	9.5	ns
TtsocQ	Clock-to-output valid time	TATM_SOC	3.3	9.5	ns
TtdataQ	Clock-to-output valid time	TATM_DATA(15:0)	3.3	12	ns
Ttclavh	Input hold time	TATM_CLAV(3:0)	2		ns
TtparQ	Clock-to-output valid time	TATM_PARITY	3.8	11	ns

6.2 DRAM External Memory Timing

NOTE: All inputs are the minimum required. All outputs are the minimum expected. All inputs and outputs are assume to have a 30 pf capacitive loading. The RX_DRAM_CLK and TX_DRAM_CLK are assume to have a 22 ohm series terminated resister connected to a combined capacitive load of 36 pf.

Figure 52 shows the receive DRAM external memory 100 MHz read timing.

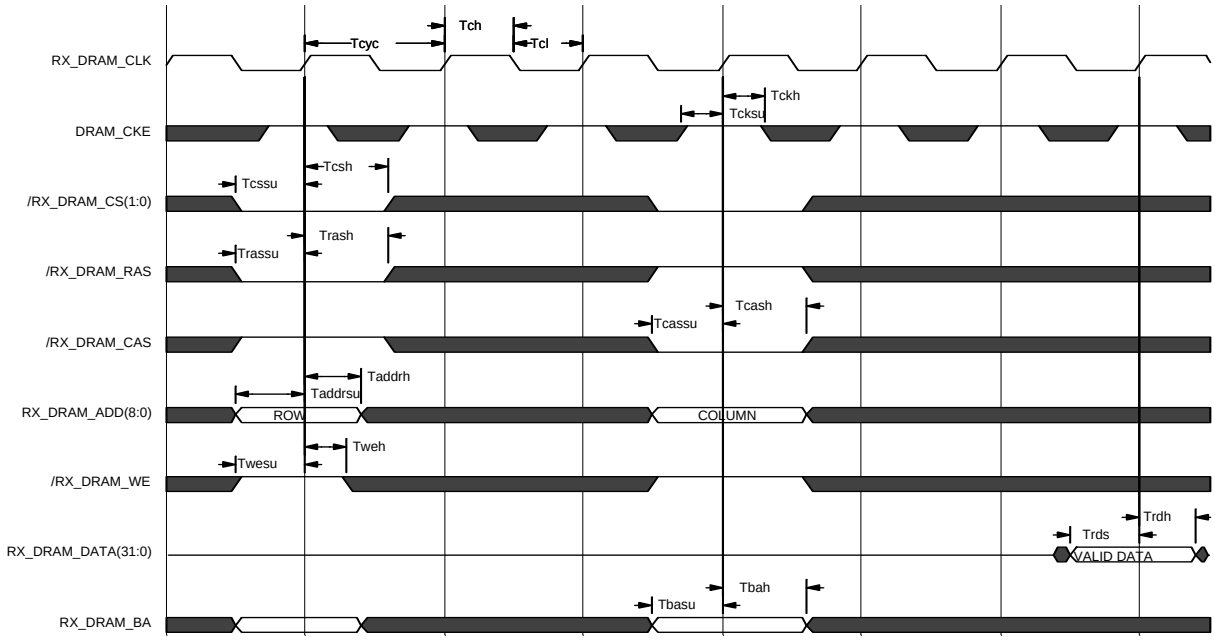

Figure 52. Receive DRAM External Memory 100 MHz Read Timing

Symbol	Parameter	Signals	Min	Max	Unit
Tcyc	Clock period	RX_DRAM_CLK	10		ns
Tch	Clock high period	RX_DRAM_CLK	3		ns
Tcl	Clock low period	RX_DRAM_CLK	3		ns
Taddrsu	Address setup time	RX_DRAM_ADD(8:0)	2.7		ns
Tbasu	Bank address setup time	RX_DRAM_BA	2.9		ns
Taddrh	Address hold time	RX_DRAM_ADD(8:0)	1.3		ns
Tbah	Bank address hold time	RX_DRAM_BA	1.5		ns
Tckh	Enable hold time *	DRAM_CKE			ns
Tcksu	Enable setup time *	DRAM_CKE			ns
Trassu	RAS setup time	/RX_DRAM_RAS	3.1		ns
Trash	RAS hold time	/RX_DRAM_RAS	1.5		ns
Tcassu	CAS setup time	/RX_DRAM_CAS	3.2		ns
Tcash	CAS hold time	/RX_DRAM_CAS	1.5		ns
Tcssu	Chip select setup time	/RX_DRAM_CS(1:0)	2.4		ns
Tcsh	Chip select hold time	/RX_DRAM_CS(1:0)	2.2		ns
Trdh	Required hold time required (read data)	RX_DRAM_DATA(31:0)	2.7		ns

Symbol	Parameter	Signals	Min	Max	Unit
Twesu	Write enable setup time	/RX_DRAM_WE	2.8		ns
Tweh	Write enable hold time	/RX_DRAM_WE	1.5		ns
Trds	Required setup time (read data)	RX_DRAM_DATA(31:0)	0		ns

Figure 53 shows the receive DRAM external memory 100 MHz write timing.

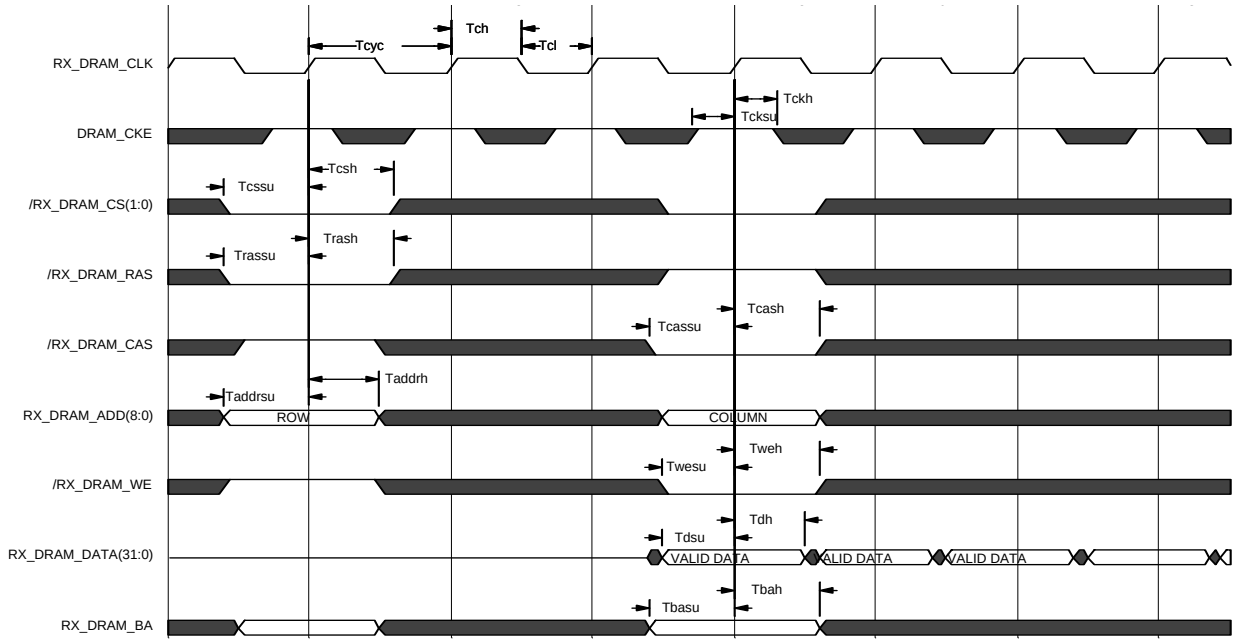

Figure 53. Receive DRAM External Memory 100 MHz Write Timing

Symbol	Parameter	Signals	Min	Max	Unit
Tcyc	Clock period	RX_DRAM_CLK	10		ns
Tch	Clock high period	RX_DRAM_CLK	3		ns
Tcl	Clock low period	RX_DRAM_CLK	3		ns
Taddrsu	Address setup time	RX_DRAM_ADD(8:0)	2.7		ns
Taddrh	Address hold time	RX_DRAM_ADD(8:0)	1.3		ns
Tbasu	Bank address setup time	RX_DRAM_BA	2.9		ns
Tbah	Bank address hold time	RX_DRAM_BA	1.5		ns
Tckh	Clock enable hold time *	DRAM_CKE	*		ns
Tcksu	Clock enable setup time *	DRAM_CKE	*		ns
Trassu	RAS setup time	/RX_DRAM_RAS	3.1		ns
Trash	RAS hold time	/RX_DRAM_RAS	1.5		ns
Tcassu	CAS setup time	/RX_DRAM_CAS	3.2		ns
Tcash	CAS hold time	/RX_DRAM_CAS	1.5		ns
Tcssu	Chip select setup time	/RX_DRAM_CS(1:0)	2.4		ns
Tcsh	Chip select hold time	/RX_DRAM_CS(1:0)	2.2		ns

Symbol	Parameter	Signals	Min	Max	Unit
Twesu	Write enable setup time	/RX_DRAM_WE	2.8		ns
Tweh	Write enable setup time	/RX_DRAM_WE	1.5		ns
Tdsu	(Write) Data valid before clock	RX_DRAM_DATA(31:0)	2.4		ns
Tdh	(Write) Data valid after clock	RX_DRAM_DATA(31:0)	1		ns

* The DRAM_CKE is a functionally static signal. Software should ensure that the DRAM_CKE is set at least 1 microsecond before desserting the SW_RESET bit.

Figure 54 shows the transmit DRAM external memory 100 MHz read timing.

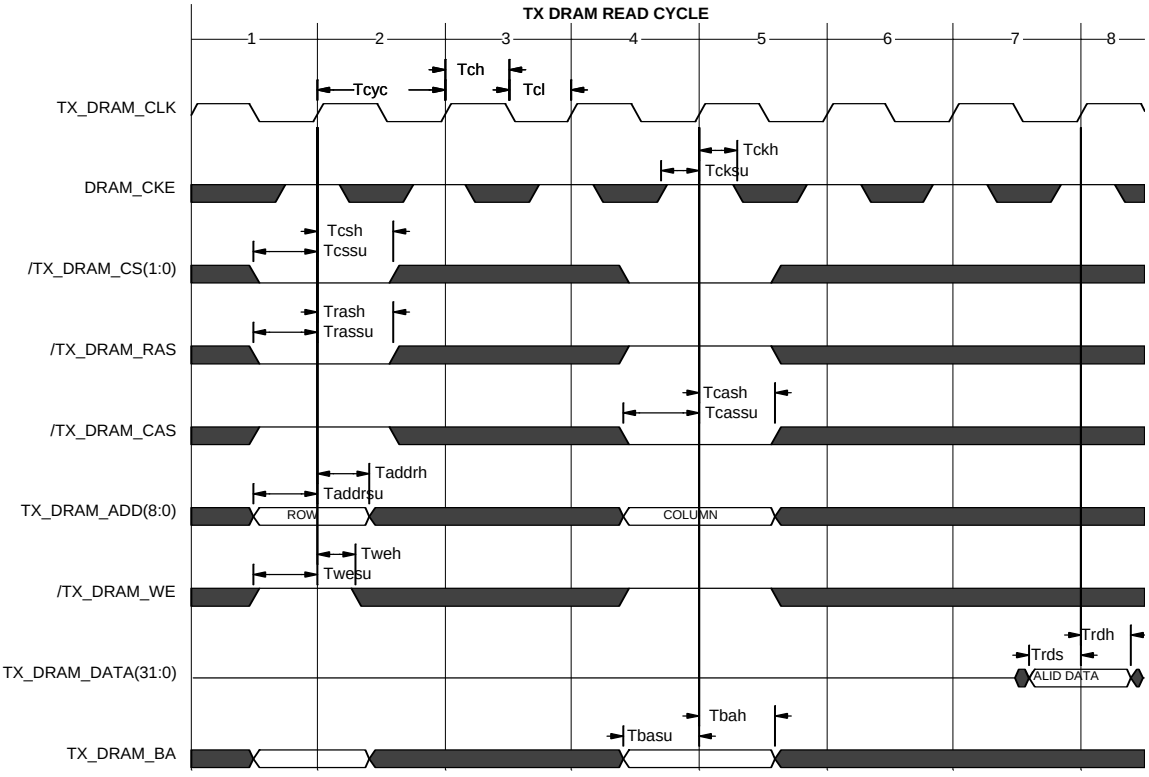

Figure 54. Transmit DRAM External Memory 100 MHz Read Timing

Symbol	Parameter	Signals	Min	Max	Unit
Tcyc	Clock period	TX_DRAM_CLK	10		ns
Tch	Clock high period	TX_DRAM_CLK	3		ns
Tcl	Clock low period	TX_DRAM_CLK	3		ns
Taddrsu	Address setup time	TX_DRAM_ADD(8:0)	3.1		ns
Taddrh	Address hold time	TX_DRAM_ADD(8:0)	1.3		ns
Tbasu	Bank address setup time	TX_DRAM_BA	2.8		ns
Tbah	Bank address hold time	TX_DRAM_BA	1.5		ns
Tckh	Clock enable hold time	DRAM_CKE	*		ns

Symbol	Parameter	Signals	Min	Max	Unit
Tcksu	Clock enable setup time	DRAM_CKE	*		ns
Trassu	RAS setup time	/TX_DRAM_RAS	3.3		ns
Trash	RAS hold time	/TX_DRAM_RAS	1.5		ns
Tcassu	CAS setup time	/TX_DRAM_CAS	3.3		ns
Tcash	CAS hold time	/TX_DRAM_CAS	1.5		ns
Tcssu	Chip select setup time	/TX_DRAM_CS(1:0)	2.7		ns
Tcsh	Chips select hold time	/TX_DRAM_CS(1:0)	1.5		ns
Trdh	(Read) Data Valid required after clock	TX_DRAM_DATA(31:0)	2		ns
Twesu	Write enable setup time	/TX_DRAM_WE	2.9		ns
Tweh	Write enable hold time	/TX_DRAM_WE	1.5		ns
Trds	(Read) Data valid required before clock	TX_DRAM_DATA(31:0)	0		ns

Figure 55 shows the transmit DRAM external memory 100 MHz write timing.

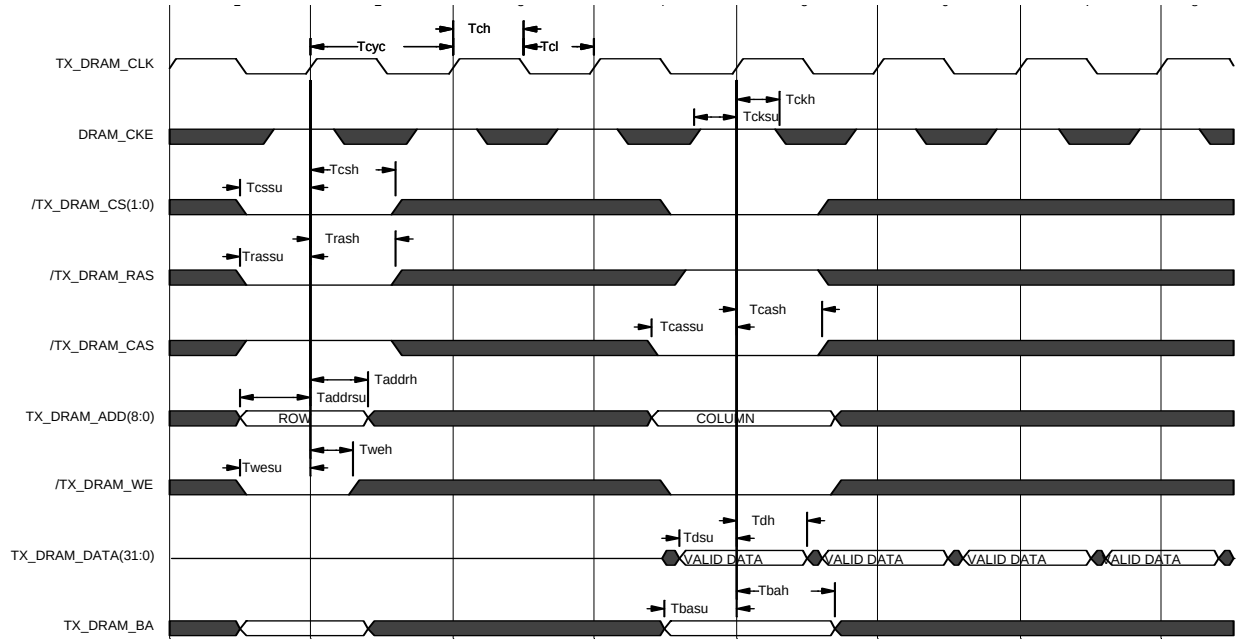

Figure 55. Transmit DRAM External Memory 100 MHz Write Timing

Symbol	Parameter	Signals	Min	Max	Unit
Tcyc	TX_DRAM_CLK period	TX_DRAM_CLK	10		ns
Tch	TX_DRAM_CLK high period	TX_DRAM_CLK	3		ns
Tcl	TX_DRAM_CLK low period	TX_DRAM_CLK	3		ns
Taddrsu	Address setup time	TX_DRAM_ADD(8:0)	3.1		ns
Taddrh	Address hold time	TX_DRAM_ADD(8:0)	1.3		ns
Tbasu	Bank address setup time	TX_DRAM_BA	2.8		ns
Tbah	Bank address hold time	TX_DRAM_BA	1.5		ns

Symbol	Parameter	Signals	Min	Max	Unit
Tckh	CLK enable hold time	DRAM_CKE	*		ns
Tcksu	CLK enable setup time	DRAM_CKE	*		ns
Trassu	RAS setup time	/TX_DRAM_RAS	3.3		ns
Trash	RAS hold time	/TX_DRAM_RAS	1.5		ns
Tcassu	CAS setup time	/TX_DRAM_CAS	3.3		ns
Tcash	CAS hold time	/TX_DRAM_CAS	1.5		ns
Tcssu	Chip select setup time	/TX_DRAM_CS(1:0)	2.7		ns
Tcsh	Chip select hold time	/TX_DRAM_CS(1:0)	1.5		ns
Twesu	Write enable setup time	/TX_DRAM_WE	2.9		ns
Tweh	Write enable hold time	/TX_DRAM_WE	1.5		ns
Tdsu	(Write) Data valid before clock	TX_DRAM_DATA(31:0)	2.2		ns
Tdh	(Write) Data Valid after clock	TX_DRAM_DATA(31:0)	1.5		ns

* The DRAM_CKE is a functionally static signal. Software should ensure that the DRAM_CKE is set at least 1 microsecond before desserting the SW_RESET bit.

6.3 SRAM Timings

All inputs and outputs are assume to have a 30 pf capacitive loading. The ALRAM_CLK, ABRAM_CLK and CHRAM_CLK are assume to have a 22 ohm series terminated resister connected to a combined capacitive load of 36 pf.
Figure 56 shows the AL RAM read timing.

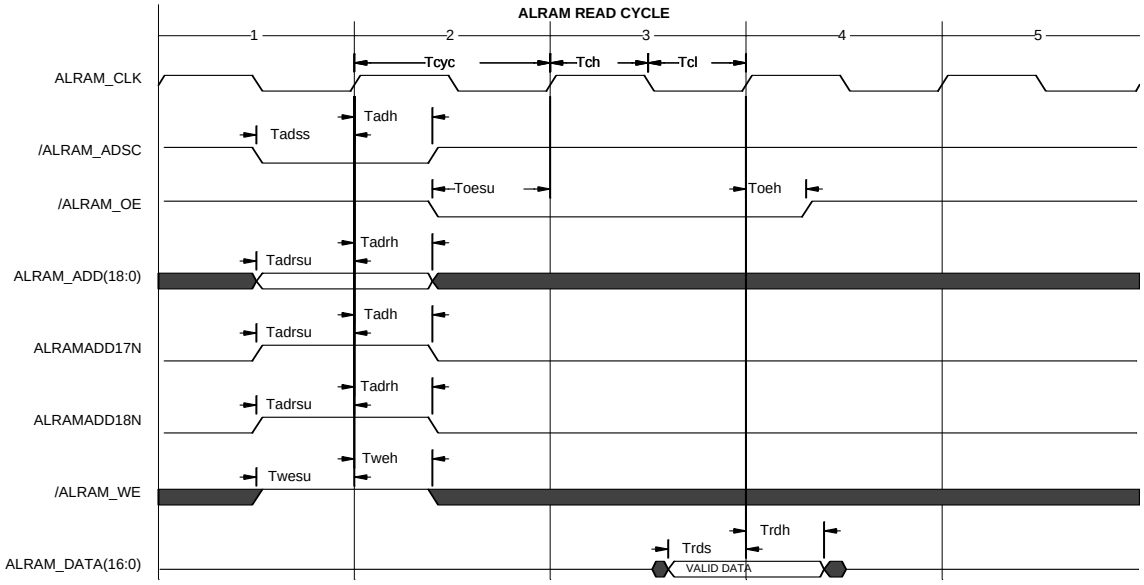

Figure 56. Address Lookup RAM Read Timing

Symbol	Parameter	Signals	Min	Max	Units
Tcyc	Clock period	ALRAM_CLK	10		ns
Tch	Clock high period	ALRAM_CLK	3		ns
Tcl	Clock low period	ALRAM_CLK	3		ns
Tadrsu	Address setup time	ALRAM_ADD(18:0), ALRAMADD17N, ALRAMADD18N	2.7		ns
Tadrh	Address hold time	ALRAM_ADD(18:0), ALRAMADD17N, ALRAMADD18N	1.5		ns
Toesu	Output enable setup time	/ALRAM_OE	3.5		ns
Toeh	Output enable hold time	/ALRAM_OE	-1		ns
Tadss	Address strobe setup time	/ALRAM_ADSC	2.8		ns
Tadh	Address strobe hold time	/ALRAM_ADSC	1.5		ns
Trds	(Read) Data valid required before clock	ALRAM_DATA(16:0)	3.8		ns
Trdh	(Read) Data valid required after clock	ALRAM_DATA(16:0)	1.5		ns
Twesu	Write enable setup time	/ALRAM_WE	2.9		ns
Tweh	Write enable hold time	/ALRAM_WE	1.2		ns

Figure 57 shows the AL RAM write timing.

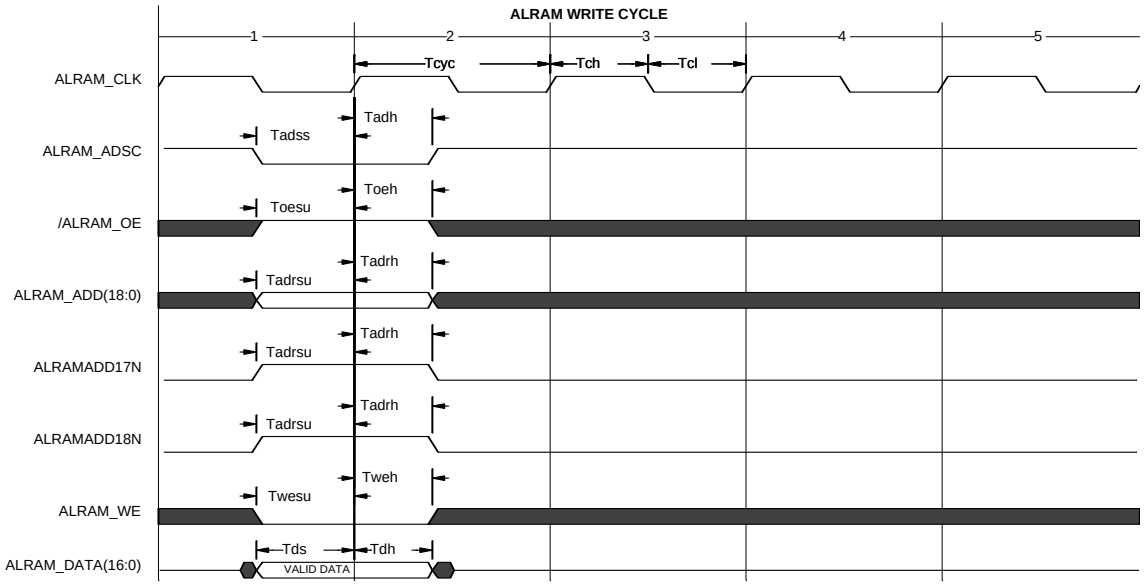

Figure 57. Address Lookup RAM Write Timing

Symbol	Parameter	Signals	Min	Max	Units
Tcyc	Clock period	ALRAM_CLK	10		ns
Tch	Clock high period	ALRAM_CLK	3		ns
Tcl	Clock low period	ALRAM_CLK	3		ns
Tadrsu	Address setup time	ALRAM_ADD(18:0), ALRAMADD17N, ALRAMADD18N	2.7		ns
Tadrh	Address hold time	ALRAM_ADD(18:0), ALRAMADD17N, ALRAMADD18N	1.5		ns
Toesu	Output enable setup time	/ALRAM_OE	3.5		ns
Toeh	Output enable hold time	/ALRAM_OE	-1		ns
Tadss	Address strobe setup time	/ALRAM_ADSC	2.8		ns
Tadh	Address strobe hold time	/ALRAM_ADSC	1.5		ns
Tds	(Write) Data valid before clock	ALRAM_DATA(16:0)	2.6		ns
Tdh	(Write) Data valid after clock	ALRAM_DATA(16:0)	1.5		ns
Twesu	Write enable setup time	/ALRAM_WE	2.9		ns
Tweh	Write enable hold time	/ALRAM_WE	1.2		ns

Figure 58 shows the channel RAM read timing.

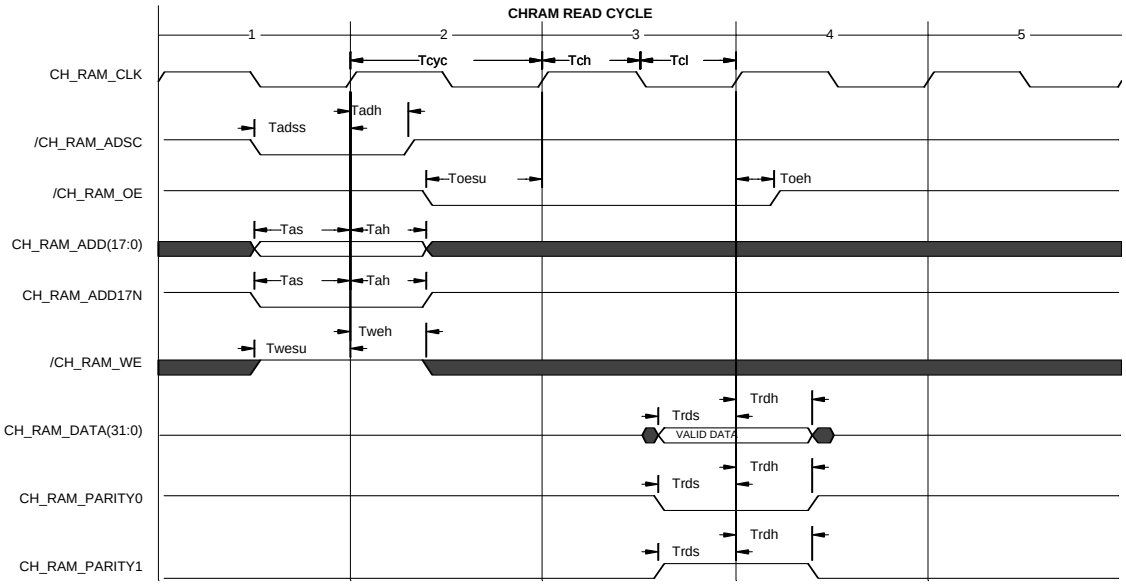

Figure 58. Channel RAM Read Timing

Symbol	Parameter	Signals	Min	Max	Units
Tcyc	Clock period	CH_RAM_CLK	10		ns
Tch	Clock high period	CH_RAM_CLK	3		ns
Tcl	Clock low period	CH_RAM_CLK	3		ns
Tas	Address setup time	CH_RAM_ADD(17:0), CH_RAM_ADD17N	3.1		ns
Tah	Address hold time	CH_RAM_ADD(17:0), CH_RAM_ADD17N	1		ns
Toesu	Output enable setup time	/CH_RAM_OE	3.5		ns
Toeh	Output enable hold time	/CH_RAM_OE	-1		ns
Tadss	Address strobe setup time	/CH_RAM_ADSC	3.2		ns
Tadh	Address strobe hold time	/CH_RAM_ADSC	1		ns
Trds	(Read) Data valid required before clock	CH_RAM_DATA(31:0)	4.7		ns
Trdh	(Read) Data valid required after clock	CH_RAM_DATA(31:0)	1		ns
Twesu	Write enable setup time	/CH_RAM_WE	3.3		ns
Tweh	Write enable hold time	/CH_RAM_WE	1		ns

Figure 59 shows the channel RAM write timing.

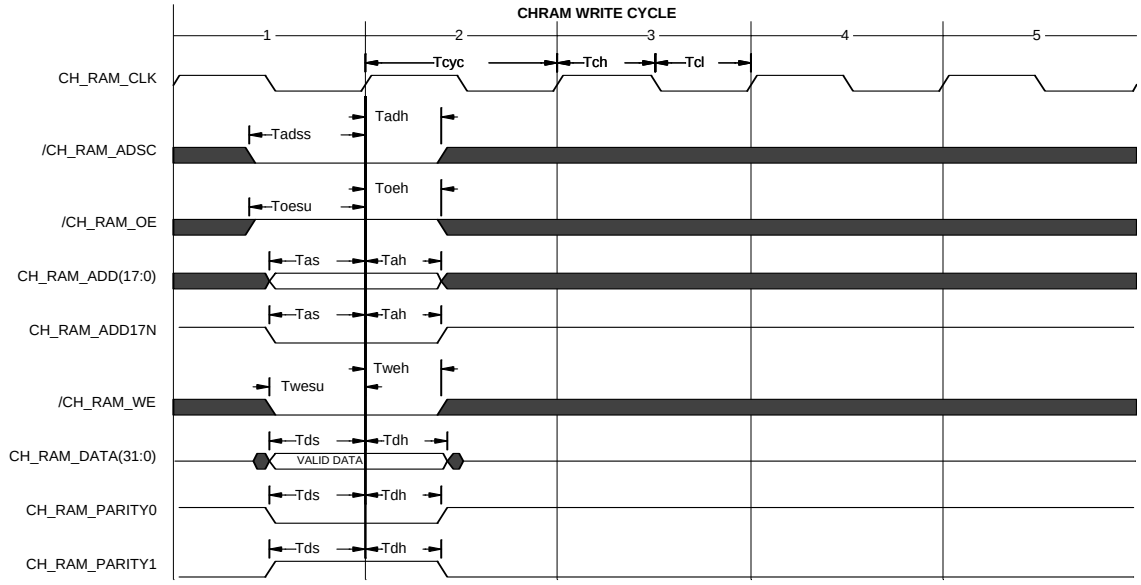

Figure 59. Channel RAM Write Timing

Symbol	Parameter	Signals	Min	Max	Units
Tcyc	Clock period	CH_RAM_CLK	10		ns
Tch	Clock high period	CH_RAM_CLK	3		ns
Tcl	Clock low period	CH_RAM_CLK	3		ns
Tas	Address setup time	CH_RAM_ADD(17:0), CH_RAM_ADD17N	3.1		ns
Tah	Address hold time	CH_RAM_ADD(17:0), CH_RAM_ADD17N	0.8		ns
Toesu	Output enable setup time	/CH_RAM_OE	3.5		ns
Toeh	Output enable hold time	/CH_RAM_OE	-1		ns
Tadss	Address strobe setup time	/CH_RAM_ADSC	3.2		ns
Tadh	Address strobe hold time	/CH_RAM_ADSC	1		ns
Tds	(Write) Data valid before clock	CH_RAM_DATA	2.5		ns
Tdh	(Write) Data valid after clock	CH_RAM_DATA	1		ns
Twesu	Write enable setup time	/CH_RAM_WE	3.3		ns
Tweh	Write enable hold time	/CH_RAM_WE	1		ns

Figure 60 shows the AB RAM read timing.

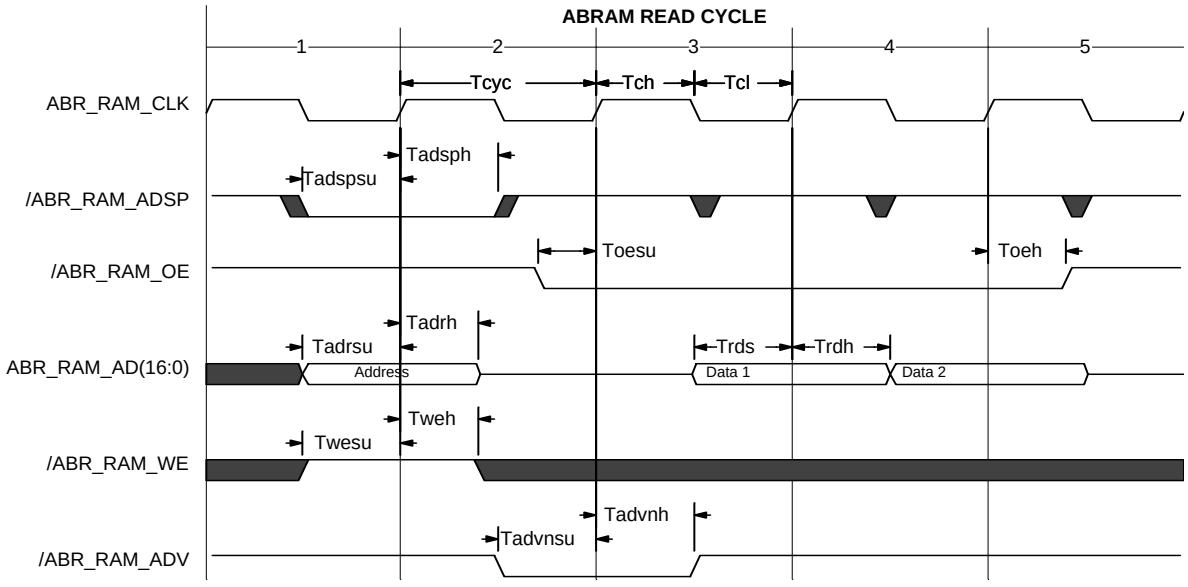

Figure 60. AB RAM Read Timing

Symbol	Parameter	Signals	Min	Max	Units
Tcyc	Clock period	AB_RAM_CLK	10		ns
Tch	Clock high period	AB_RAM_CLK	3		ns
Tcl	Clock low period	AB_RAM_CLK	3		ns
Tadrsu	(Read) Required data valid before clock Adress setup time	AB_RAM_AD(16:0)	2.6		ns
Tadrh	(Read) Required data valid after clock Address hold time	AB_RAM_AD(16:0)	1		ns
Trds	(Read) Required data valid before clock	AB_RAM_AD(16:0)	3.5		
Trdh	(Read) Required data valid after clock	AB_RAM_AD(16:0)	1		
Twesu	Write enable setup time	/AB_RAM_WE	3.5		ns
Tweh	Write enable hold time	/AB_RAM_WE	1		ns
Tadspsu	Address status processor setup time	/AB_RAM_ADSP	3.3		ns
Tadsph	Address status processor hold time	/AB_RAM_ADSP	1		ns
Tadvnsu	Address advance setup time	/AB_RAM_ADV	3.4		ns
Tadvnh	Address advance hold time	/AB_RAM_ADV	1		ns
Toesu	Output enable setup time	/AB_RAM_OE	1.4		ns
Toeh	Output enable hold time	/AB_RAM_OE	-1		ns

Figure 61 shows the AB RAM write timing.

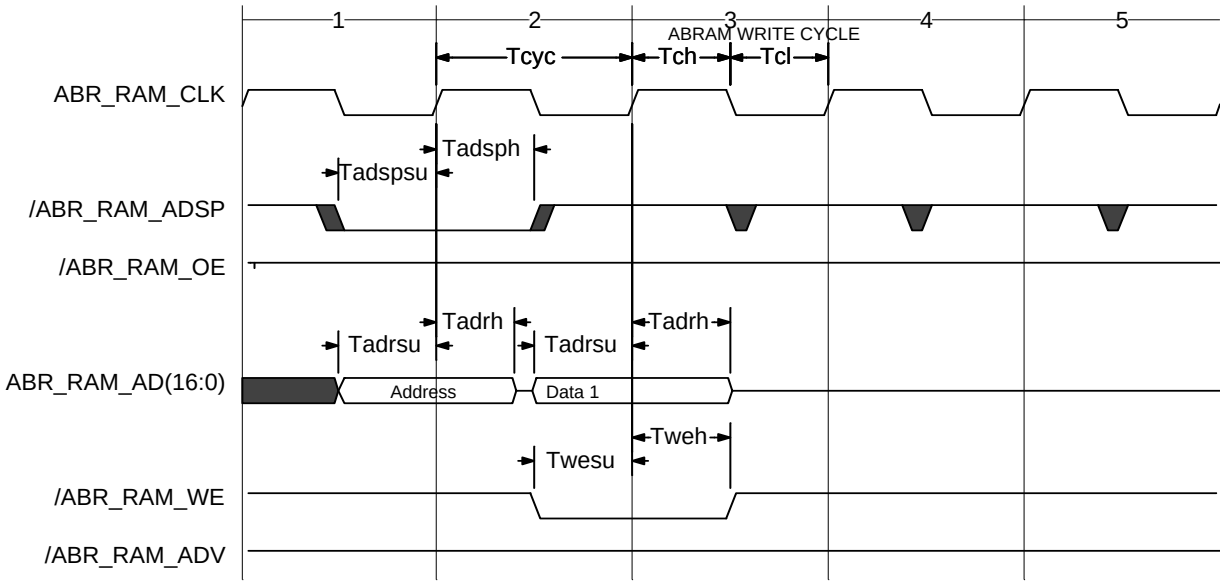

Figure 61. AB RAM Write Timing

Symbol	Parameter	Signals	Min	Max	Units
Tcyc	Clock period	AB_RAM_CLK	10		ns
Tch	Clock high period	AB_RAM_CLK	3		ns
Tcl	Clock low period	AB_RAM_CLK	3		ns
Tadrsu	(Write) Required data valid before clock Adress setup time	AB_RAM_AD(16:0)	2.6		ns
Tadrh	(Write) Required data valid after clock Address hold time	AB_RAM_AD(16:0)	1		ns
Twesu	Write enable setup time	/AB_RAM_WE	3.5		ns
Tweh	Write enable hold time	/AB_RAM_WE	1		ns
Tadspsu	Address status processor setup time	/AB_RAM_ADSP	3.3		ns
Tadsph	Address status processor hold time	/AB_RAM_ADSP	1		ns

6.4 QRT-QSE Interface Timing

Output timing delays assume a capacitive loading of 30 pF on SE_D_OUT(3:0,3:0) and 48 pF on SE_SOC_OUT. Figure 62 shows the bit-level timing for the QRT..

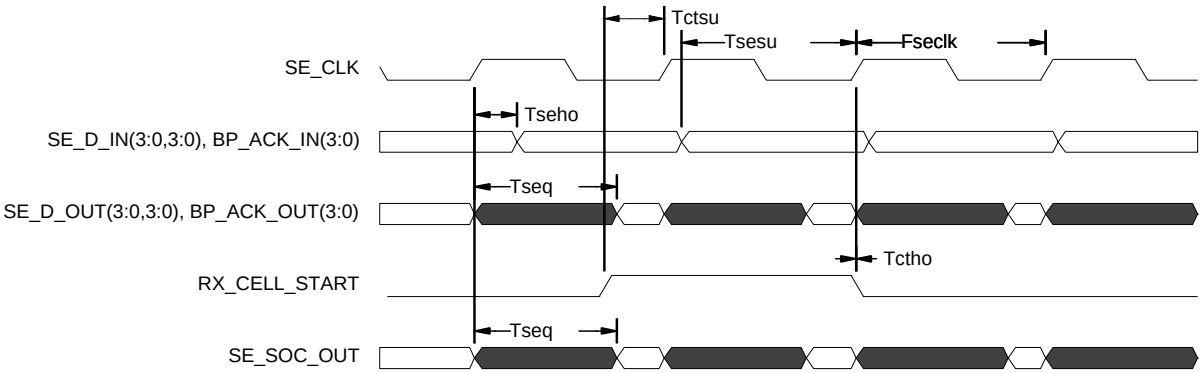

Figure 62. QRT Bit-Level Timing

Symbol	Parameter	Signals	Min	Max	Unit
Fseclk**	Frequency of SE_CLK	SE_CLK	65.4	68	MHz
	clock duty cycle	SE_CLK	40	60	%
	<10 KHz Jitter tolerance *	BP_ACK_IN(3:0), SE_SOC_IN(3:0)		2	clock period
	>10 KHz Jitter tolerance *	BP_ACK_IN(3:0), SE_SOC_IN(3:0)		0.35	clock period
Tctsu	Control signal setup time	RX_CELL_START	1.5		ns
Tctho	Control signal hold time	RX_CELL_START	0		ns
Tsesu*	Setup time before SE_CLK	SE_D_IN(3:0,3:0), SE_SOC_IN(3:0), BP_ACK_IN(3:0)	4.0		ns
Tseho*	Hold time after SE_CLK	SE_D_IN(3:0,3:0), SE_SOC_IN(3:0), BP_ACK_IN(3:0)	1.5		ns
Tseq	Output delay from SE_CLK	SE_D_OUT(3:0,3:0), BP_ACK_OUT(3:0), SE_SOC_OUT	2.9	10.0	ns
	Output delay skew *	SE_D_OUT(0,3:0) and SE_SOC_OUT SE_D_OUT(1,3:0) and SE_SOC_OUT SE_D_OUT(2,3:0) and SE_SOC_OUT SE_D_OUT(3,3:0) and SE_SOC_OUT		1.3	ns
	Input delay skew *	SE_D_IN(0,3:0) and SE_SOC_IN(0) SE_D_IN(1,3:0) and SE_SOC_IN(1) SE_D_IN(2,3:0) and SE_SOC_IN(2) SE_D_IN(3,3:0) and SE_SOC_IN(3)		3	ns

Symbol	Parameter	Signals	Min	Max	Unit
* When the phase aligners are turned on, Tsesu and Tseho are no longer defined. However, the maximum input and output skew and jitter on these signals with respect to the SE_SOC_IN is constrained to specification listed in this table. ** This is assuming the SYSCLK is operating at 100 Mhz. To operate the device fabric interface at other frequencies, refer to "Relationships Among Various Clock Domains" on page 229.					

6.4.1 Switch Fabric Cell Formats

Table 21 on page 91 and Table 22 on page 92 define the switch fabric interface cell formats for data cells and idle cells, respectively.

Table 21. QRT-QSE Interface Cell Format.

Nibble	Symbol	Definition	Comment	
0	Pres(1:0), MC, SP	Pres = 10_b: Valid cell present. Pres = 01_b: Idle cell. Pres = 00_b: Invalid cell. Pres = 11_b: Invalid cell. MC = 1_b: Multicast cell. SP: Spare bit. Programmed via the RX_CH_CONFIG register (refer to section 9.3.1.1 "RX_CH_CONFIG" starting on page 184).		
1	SP(1:0), Priority(1:0)	SP(1:0) Spare bits. Priority = 11_b High-priority cell. 10_b Medium-priority cell. 01_b Low-priority cell. 00_b Undefined. Cell discarded by QSE.	The QRT should be configured never to generate priority 00_b cells, since they are discarded by the QSE.	
2	TAG_0_6	Routing tag 0.	MCG_3	The QSE interprets all four nibbles as the MCG.
3	TAG_1_7	Routing tag 1.	MCG_2	
4	TAG_2	Routing tag 2.	MCG_1	
5	TAG_3	Routing tag 3.	MCG_0	
6	TAG_4	Routing tag 4.		
7	TAG_5	Routing tag 5.		
8	TAG_0_6	Routing tag 6.		
9	TAG_1_7	Routing tag 7.		
10	OUTCHAN_3	Intepreted as OUTCHAN(15:12) by the QRT.	Not used by the QSE.	

Nibble	Symbol	Definition	Comment
11	SP(1:0), MB, P	SP(1:0) Spare bits. MB Mark bit. Cells that are present and have this bit set are counted by the TX_MARKED_CELLS and RX_MARKED_CELLS counters (refer to "MARKED_CELLS_COUNT" on page 110). P Set to odd parity by software over nibbles 0-11.	
12	OUTCHAN_2	Intepreted as OUTCHAN(11:8) by the QRT.	Not used by the QSE.
13	OUTCHAN_1	Intepreted as OUTCHAN(7:4) by the QRT.	Not used by the QSE.
14	OUTCHAN_0	Intepreted as OUTCHAN(3:0) by the QRT.	Not used by the QSE.
15	VCI_3	Intepreted as VCI(15:12) by the QRT.	Not used by the QSE.
16	VCI_2	Intepreted as VCI(11:8) by the QRT.	Not used by the QSE.
17	VCI_1	Intepreted as VCI(7:4) by the QRT.	Not used by the QSE.
18	VCI_0	Intepreted as VCI(3:0) by the QRT.	Not used by the QSE.
19	PTI(2:0)/CLP	Intepreted as the PTI and CLP field from the cell by the QRT.	Not used by the QSE.
20	SEQ_1	Intepreted as SEQ(7:4) by the QRT.	Not used by the QSE.
21	SEQ_0	Intepreted as SEQ(3:0) by the QRT.	Not used by the QSE.
22-117	Payload	Interpreted as 48 bytes of ATM cell payload by the QRT.	Not used by the QSE.

Table 22. QRT-QSE Interface Idle Cell Format

Nibble	Symbol	Definition	Comment
0	Pres(3:0)	Pres = 0100_b: Cell not present.	
1	IDLE_0	IDLE_0 = 0000_b: All 0.	
2	IDLE_1	IDLE_1 = 1000_b: Marching 1.	Marching "1" pattern protects against bridging faults. If the QRT is not in sync with RX_CELL_START, the device outputs 1010_b here.
3	IDLE_2	IDLE_2 = 0100_b: Marching 1.	Marching "1" pattern protects against bridging faults.
4	IDLE_3	IDLE_3 = 0010_b: Marching 1.	
5	IDLE_4	IDLE_4 = 0001_b: Marching 1.	
6	IDLE_5	IDLE_5 = 000_b.	
7	IDLE_6	IDLE_6 = 000_b.	
8 - 15	Reserved	The QRT currently outputs 000_b.	
16-117	Unused	The QRT currently outputs 000_b.	

6.5 Processor Interface Timing

A microprocessor cycle starts when the chip select (/CS) and command (/ADS) are asserted. During read cycles, the device asserts /READY to indicate that data on the data bus is valid and during write cycles the chip asserts /READY to indicate that the write has finished and data from the bus can be removed. The microprocessor can terminate the current cycle at any time. As shown in Figure 63, the device stops driving the bus and deasserts all control lines when the cycle terminates. The current cycle terminates when the device select is deasserted or both read and write are deasserted. A new cycle can start once the /READY has been deasserted. If the cycle was terminated prematurely before the /READY was asserted, then a new microprocessor cycle can start after one clock cycle.

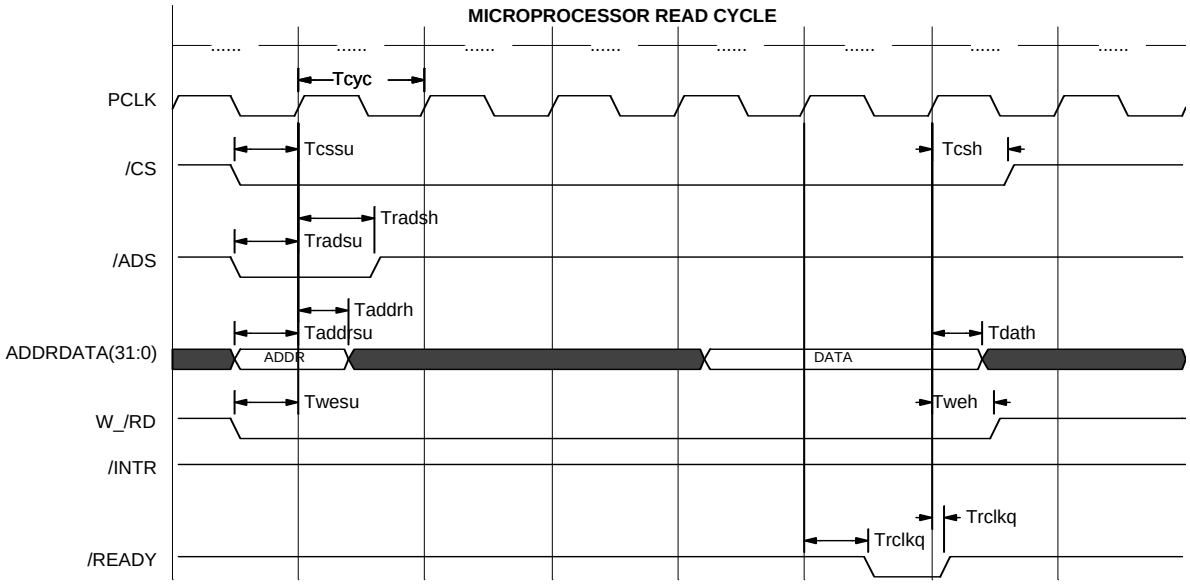

Figure 63. Microprocessor Read Timing

Symbol	Parameter	Signals	Min	Max	Unit
Tcyc	Processor clock frequency	PCLK	12.5	50	MHz
Taddrsu	Input address setup time	ADDRDATA(31:0)	3		ns
Taddrh	Input address hold time	ADDRDATA(31:0)	1		ns
Tradsu	Address strobe setup time	/ADS	3		ns
Tradsh	Address strobe hold time	/ADS	1		ns
Tcssu	Chip select setup time	/CS	3		ns
Tcsh	Chip select hold time	/CS	1		ns

Symbol	Parameter	Signals	Min	Max	Unit
Twesu	Write enable setup	W_/RD	1		ns
Tweh	Write enable hold time	W_/RD	1		ns
Tdath	Data hold time	ADDRDATA(31:0)	2		ns
Trclkq	Acknowledge to output valid	/READY	2	9	ns

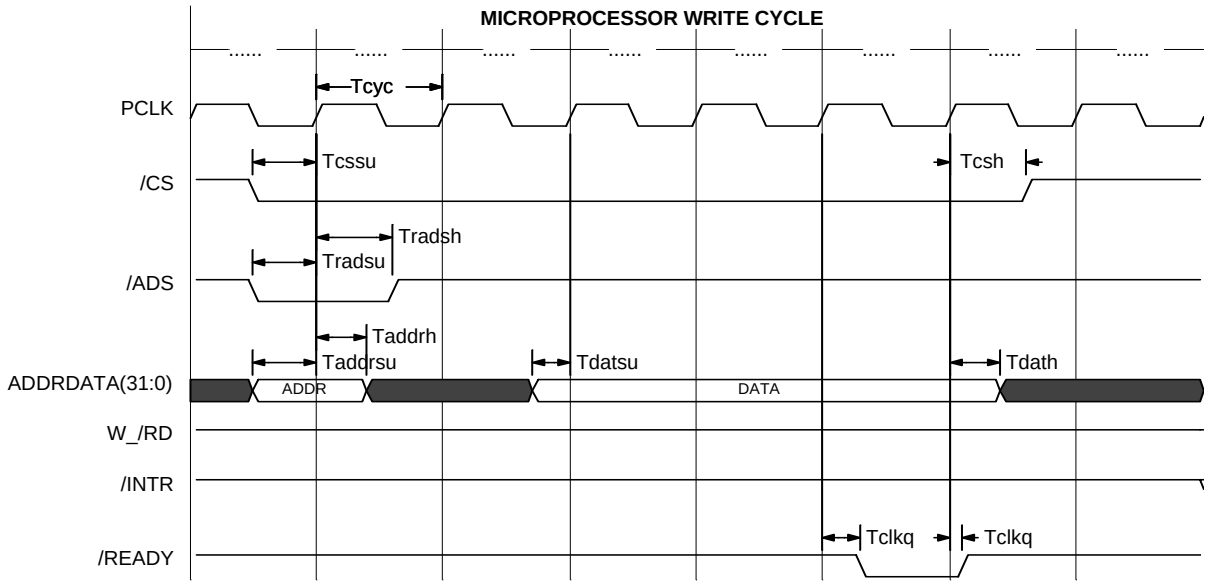

Figure 64. Microprocessor Write Timing

Symbol	Parameter	Signals	Min	Max	Unit
Tcyc	Processor clock frequency	PCLK	12.5	50	MHz
	Clock duty cycle	PCLK	40	60	%
Taddrsu	Input address setup time	ADDRDATA(31:0)	3		ns
Taddrh	Input address hold time	ADDRDATA(31:0)	1		ns
Tradsu	Address strobe setup time	/ADS	3		ns
Tradsh	Address strobe hold time	/ADS	1		ns
Tcssu	Chip select setup time	/CS	3		ns
Tcsh	Chip select hold time	/CS	1		ns
Tclkq	Acknowledge to output valid	/READY	2	9	ns
Tdath	Data hold time	ADDRDATA(31:0)	1		ns

Symbol	Parameter	Signals	Min	Max	Unit
Tdatsu	Data setup time	ADDRDATA(31:0)	3		ns

6.6 Miscellaneous Timing

Figure 65 shows the reset pin (/RESET) timing. The /RESET signal must be asserted for a minimum time (Tres) to be registered properly in the presence of SE_CLK. The QRT remains in reset while /RESET is asserted and starts performing normally after the first RX_CELL_START after Trstproc. From then on RX_CELL_START must occur every 118 clocks. If an invalid RX_CELL_START (that is, one that is not exactly 118 clocks after the preceding RX_CELL_START) is received by the device (as occurs, for example, after a switchover), the first RX_CELL_START is ignored and processing starts with the next RX_CELL_START.

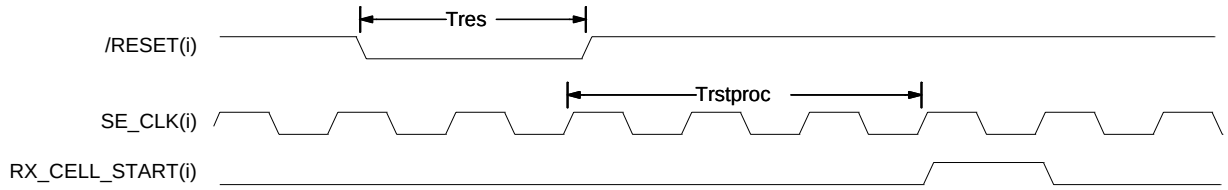

Figure 65. Reset Timing

Symbol	Parameter	Signals	Min	Max	Unit
Tres	Reset assertion time.	/RESET	40		SE_CLK periods
Trstproc	Reset processing time.	/RESET		60	SE_CLK periods

Figure 66 shows the timing for the JTAG port. The /JTAG_RESET signal is asynchronous to JTAG_TCK.

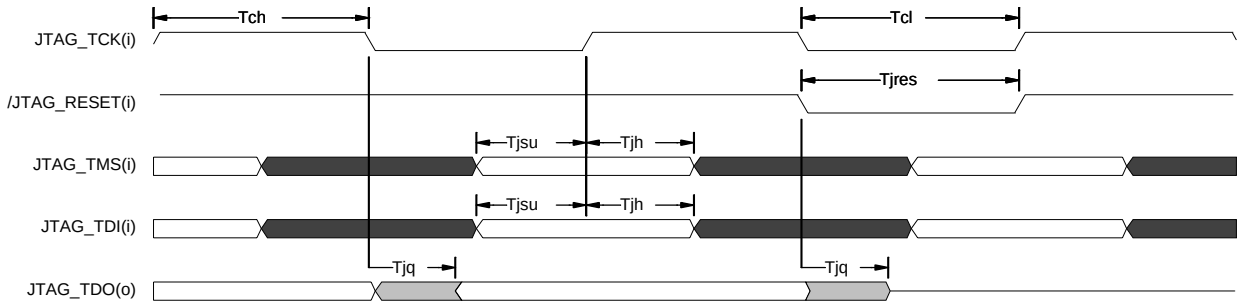

Figure 66. JTAG Timing

Symbol	Parameter	Signals	Min	Max	Unit
	JTAG_TCK frequency			1	MHz
Tch	JTAG_TCK high		80		ns
Tcl	JTAG_TCK low		80		ns
Tjh	JTAG_TCK hold time	JTAG_TMS, JTAG_TDI	40		ns
Tjres	/JTAG_RESET low	/JTAG_RESET	80		ns
Tjsu	JTAG_TCK setup time	JTAG_TMS, JTAG_TDI	40		ns
Tjq	JTAG_TCK-to-output delay /JTAG_RESET-to-output delay	JTAG_TDO	2.3	40	ns

7 MICROPROCESSOR PORTS

7.1 Microprocessor Ports Summary

NOTES:

- The external /RESET signal resets all microprocessor write ports to 0 when asserted.
- All ports marked as “Reserved” must be initialized to 0 (unless otherwise indicated) at initial setup. Software modifications to these locations after setup may cause incorrect operation.
- All read/write port bits marked “Not used” must be written with the value 0 to maintain software compatibility with future versions.
- Most read-only bits marked “Not used” are driven with a 0; however, some may have an unpredictable value on reads. Therefore, all read-only bits marked “Not used” should be masked by the software on reads to maintain compatibility with future versions.

Table 23. Microprocessor Ports Summary

Byte Address	Long Address	Name	Read or Write	Description
0_h	0_h	REVISION	R	Contains the device part number and revision.
4_h	1_h	RESET	R/W	Resets the device, except the microprocessor interface.
8_h	2_h	TEST_CONFIG	R/W	Define various test/diagnostic functions.
C_h	3_h	SRAM_CONFIG	R/W	Defines the SRAM configuration.
10_h	4_h	SWITCH_CONFIG	R/W	Defines the type of switch fabric.
14_h	5_h	RAM BIST RESULT	R/W	RAM BIST completion result
14_h-18_h	5_h-6_h	Not used	R/W	Set to 0 for software compatibility with future devices.
$1C_h$	7_h	MARKED_CELLS_COUNT	R	Indicates the number of cells modulo (mod) 16 that had Tag(9,1) set to 1, sent to or from the switch fabric.
20_h	8_h	CONDITION_PRES_BITS	R	Defines the condition of present bits.
24_h	9_h	CONDITION_LATCH_BITS	R	Defines the condition of latch bits.
28_h	A_h	INTR_MASK	R/W	Defines the interrupt mask bits.
$2C_h$-$3C_h$	B_h-F_h	Reserved	R/W	Do not initialize.
40_h	10_h	UTOPIA_CONFIG	R/W	Defines the mode of the UTOPIA interface.
44_h	11_h	UT_PRIORITY	R/W	Indicates the priority of each virtual input/output.
48_h	12_h	UT_ENABLE	R/W	Enables whether the polling results from the PHY layer device will be considered during the UTOPIA PHY selection mechanism. Also sets the number of devices to be polled when standard polling is used.
$4C_h$	13_h	TX_UT_STAT	R	Indicates the transmit UTOPIA status.
50_h	14_h	TX_UT_WD_ALIVE	R/W	Indicates the transmit UTOPIA watchdog liveness.
54_h-$7F_h$	15_h-$1F_h$	Not used	R	Driven with a 0. Mask on reads to maintain compatibility with future versions.
80_h	20_h	RX_CELL_START_ALIGN	R/W	Contains alignment value for the internal RxCellStart signal from the external RX_CELL_START.

Table 23. Microprocessor Ports Summary (Continued)

Byte Address	Long Address	Name	Read or Write	Description
84_h-$9C_h$	21_h-27_h	Not used	R	Driven with a 0. Mask on reads to maintain compatibility with future versions.
$A0_h$	28_h	RX_QUEUE_ENGINE_TEST	R/W	Receive queue engine test/debug register.
$A4_h$	29_h	TX_QUEUE_ENGINE_TEST	R/W	Transmit queue engine test/debug register.
$A8_h$	$2A_h$	QUEUE_ENGINE_CONDITION_PRES_BITS	R	Current status of queue engine failure detection bits.
AC_h	$2B_h$	QUEUE_ENGINE_CONDITION_LATCH_BITS	R/W	Latched version of queue engine failure detection bits.
$B0_h$	$2C_h$	QUEUE_ENGINE_INT_MASK	R/W	Interrupt mask for the queue engine failure detection bits.
$B4_h$-FC_h	$2D_h$-$3F_h$	Not used	R	Driven with a 0. Mask on reads to maintain compatibility with future versions.
100_h - $7BFFC_h$	40_h - $1EFFFF_h$	Not used		Driven with a 0. Mask on reads to maintain compatibility with future versions.
$7C0000_h$	$1F0000_h$	RX_DIR_CONFIG	R/W	Sets the congestion threshold for the receive direction.
$7C0004_h$	$1F0001_h$	RX_DIR_STATE	R/W	Indicates the current queue depth for this SC.
$7C0008_h$-$7C001F_h$	$1F0002_h$-$1F0007_h$	Reserved	R	Do not initialize.
$7C0020_h$	$1F0008_h$	TX_DIR_CONFIG	R/W	Sets the congestion threshold for the transmit direction.
$7C0024_h$ Or $7C00A4_h$	$1F0009_h$ Or $1F0029_h$	TX_DIR_STATE	R/W	Indicates the current state of the transmit direction.
$7C0028_h$-$7C003C_h$	$1F000A_h$-$1F000F_h$	Reserved	R	Do not initialize.
$7C0040_h$	$1F0010_h$	RX_SENT_CELLS	R	Twenty-four-bit counter of cells sent in the receive direction.
$7C0044_h$	$1F0011_h$	RX_DROPPED_CELLS	R	Twenty-four-bit counter of cells dropped in the receive direction.
$7C0048_h$	$1F0012_h$	TX_SENT_CELLS	R	Twenty-four-bit counter of cells sent in the transmit direction.
$7C004C_h$	$1F0013_h$	TX_DROPPED_CELLS	R	Twenty-four-bit counter of cells dropped in the transmit direction.
$7C0050_h$	$1F0014_h$	RX_LOWER16_SCG_CONFIG	R/W	Configuration of the Rx Lower 16 Service Class Group.
$7C0054_h$	$1F0015_h$	RX_LOWER16_SCG_STATE	R/W	State of the Rx Lower 16 Service Class Group.
$7C0058_h$	$1F0016_h$	RX_LOWER32_SCG_CONFIG	R/W	Configuration of the Rx Lower 32 Service Class Group.
$7C005C_h$	$1F0017_h$	RX_LOWER32_SCG_STATE	R/W	State of the Rx Lower 32 Service Class Group.

Table 23. Microprocessor Ports Summary (Continued)

Byte Address	Long Address	Name	Read or Write	Description
$7C0060_h$	$1F0018_h$	RX_LOWER48_SCG_CONFIG	R/W	Configuration of the Rx Lower 48 Service Class Group.
$7C0064_h$	$1F0019_h$	RX_LOWER48_SCG_STATE	R/W	State of the Rx Lower 48 Service Class Group.
$7C0068_h$	$1F001A_h$	TX_LOWER4_SCG_CONFIG	R/W	Configuration of the Tx Lower 4 Service Class Group.
$7C006C_h$	$1F001B_h$	TX_LOWER4_SCG_STATE	R/W	State of the Tx Lower 4 Service Class Group.
$7C0070_h$	$1F001C_h$	TX_LOWER8_SCG_CONFIG	R/W	Configuration of the Tx Lower 8 Service Class Group.
$7C0074_h$	$1F001D_h$	TX_LOWER8_SCG_STATE	R/W	State of the Tx Lower 8 Service Class Group.
$7C0078_h$	$1F001E_h$	TX_LOWER12_SCG_CONFIG	R/W	Configuration of the Tx Lower 12 Service Class Group.
$7C007C_h$	$1F001F_h$	TX_LOWER12_SCG_STATE	R/W	State of the Tx Lower 12 Service Class Group.
$7C0080_h$-$7C01FC_h$	$1F0020_h$-$1F007F_h$	Reserved	R	Do not initialize.

7.2 Microprocessor Ports Bit Definitions

7.2.1 REVISION

This register contains the device part number and revision.

Address: 0_h (0_h byte)
Type: Read-only
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:16)	Driven with a 0. Mask on reads to maintain compatibility with future versions.
PART_TYPE (15:4)	Indicates the part number. Always returns 487_h.
REVISION (3:0)	Driven with a 1_h to indicate the second QRT version.

7.2.2 RESET

Address: 1_h (4_h byte)
Type: Read/write
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:2)	Write with a 0 to maintain software compatibility with future versions.
HW_RESET (1)	1 Resets the device to a state similar to that resulting from asserting the /RESET pin, except the microprocessor registers are available. The internal or external RAMs are not available. This allows the basic mode of the device to be initialized before the I/O is enabled. 0 This reset is deasserted. Resets to 1_b.
SW_RESET (0)	1 Resets the device, except for the microprocessor interface. The processor may initialize all the registers and internal memories without the cell flow while this is asserted. The switch fabric and UTOPIA interface are running, but there is no cell flow. 0 Normal operation. Resets to 1_b.

7.2.3 TEST_CONFIG

Address: 2_h (8_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:22)	Write with a 0 to maintain software compatibility with future versions.
VO Ram BIST control (21:20)	00 Normal operation. 01 BIST memory test mode. RAM BIST controller will begin testing the VO Ram. This value must be maintained until the VO_RAM_COMPLETE bit in the BIST_RESULT register is '1'. 10 BIST Controller test mode. Perform the controller error detector test. The VO_RAM_BIST_FAIL bit will be set to '1' at the end of this operation 11 Invalid control code. Do not use Resets to 00_b.
RS Ram BIST control (19:18)	00 Normal operation. 01 BIST memory test mode. RAM BIST controller will begin testing the RS Ram. This value must be maintained until the RS_RAM_COMPLETE bit in the BIST_RESULT register is '1'. 10 BIST controller test mode. Perform the controller error detector test. The RS_RAM_BIST_FAIL bit will be set to '1' at the end of this operation 11 Invalid control code. Do not use Resets to 00_b.
TS Ram BIST control (17:16)	00 Normal operation. 01 BIST memory test mode. RAM BIST controller will begin testing the TS Ram. This value must be maintained until the TS_RAM_COMPLETE bit in the BIST_RESULT register is '1'. 10 BIST controller test mode. Perform the controller error detector test. The TS_RAM_BIST_FAIL bit will be set to '1' at the end of this operation 11 Invalid control code. Do not use Resets to 00_b.
RF Ram BIST control (15:14)	00 Normal operation. 01 BIST memory test mode. RAM BIST controller will begin testing the RF Ram. This value must be maintained until the RF_RAM_COMPLETE bit in the BIST_RESULT register is '1'. 10 BIST controller test mode. Perform the controller error detector test. The RF_RAM_BIST_FAIL bit will be set to '1' at the end of this operation 11 Invalid control code. Do not use Resets to 00_b.

Field (Bits)	Description
TF Ram BIST control (13:12)	00 Normal operation. 01 BIST memory test mode. RAM BIST controller will begin testing the TF Ram. This value must be maintained until the RF_RAM_COMPLETE bit in the BIST_RESULT register is '1'. 10 BIST controller test mode. Perform the controller error detector test. The TF_RAM_BIST_FAIL bit will be set to '1' at the end of this operation 11 Invalid control code. Do not use Resets to 00_b.
Not used (11:10)	Write with a 0 to maintain software compatibility with future versions.
Reserved (9)	Write with a 0 to maintain future software compatibility. Resets to 0.
Reserved (8)	Write with a 0 to maintain future software compatibility. Resets to 0.
Not used (7:6)	Write with a 0 to maintain software compatibility with future versions.
UTOP_LOOP (5)	NOTE: For UT loopback, the device must be in OC_12C_MODE (refer to "RX_OC_12C_MODE" on page 118), and the ATM_CLK in the QRT must be synchronously derived from the SYSCLK. 1 Loopback cells from receive UTOPIA to transmit UTOPIA. 0 Normal operation. Resets to 0_b.
SF_LOOP (4)	NOTE: For SF loopback, the device must not use randomization (set DISABLE_RANDOMIZATION = 1; refer to "DISABLE_RANDOMIZATION" on page 107) and turn off the complex phase aligners (set CPA_OFF = 1; refer to "CPA_OFF" on page 107). 1 Loopback at the switch fabric interface. 0 Normal operation. Resets to 0_b.
Not used (3:1)	Write with a 0 to maintain software compatibility with future versions.
Reserved (0)	Resets to 0.

7.2.4 SRAM_CONFIG

Address: 3_h (C_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:24)	Write with a 0 to maintain software compatibility with future versions.
RX_DRAM_TYPE (23)	1 Using 1M × 16 SDRAM or 1M x 32 SGRAM devices. For the function of / RX_DRAM_CS(1:0), refer to "/RX_DRAM_CS(1:0)" on page 69. . 0 Using 256K × 32 SGRAM devices.
TX_DRAM_TYPE (22)	1 Using 1M × 16 SDRAM or 1M x 32 SGRAM devices. For the function of / TX_DRAM_CS(1:0), refer to "/TX_DRAM_CS(1:0)" on page 70. 0 Using 256K × 32 SGRAM devices.
Not used (21:14)	Write with a 0 to maintain software compatibility with future versions.
NUM_VI (13:12)	Defines the number of Virtual Input (VI) bits to use to determine the index into the VI_VPI_TABLE (refer to section 9.2.1 "VI_VPI_TABLE" starting on page 175) as part of the channel lookup. 0_h Use zero VI bits. Use only the VP bits in the calculation for the VI_VPI_TABLE. Only UTOPIA address 0 may be used in this mode. 1_h Use two VI bits (1:0) as the two MSBs in the calculation of the index for the VI_VPI_TABLE. Only UTOPIA addresses 0-3 may be used in this mode. 2_h Use five VI bits (4:0) as the five MSBs in the calculation of the index for the VI_VPI_TABLE. Any of the 31 UTOPIA addresses may be used in this mode. 3_h Not defined. Resets to 0_h.
ALRAM_TYPE (11)	1 Using the Single Cycle Deselect (SCD) type SSRAM for ALRAM with ALRAM_CONFIG = 3. Note: This is only valid when the ALRAM is populated with 2 chips of 256K x 18 type. 0 Using the Double Cycle Deselect (DCD) type SSRAM for ALRAM Resets to 0_h
Not used (10)	Write with a 0 to maintain software compatibility with future versions.

(Continued)

Field (Bits)	Description
AL_RAM_CONFIG (9:8)	Defines the size of the SRAM used to store the address translation tables and a number of other tables. The processor must know or determine the SRAM size. Then, the processor must inform the device of the SRAM size by configuring this field. Initialize to the proper setting. 0_h AL_RAM = 64K × 16 SSRAM. 8K VI_VPI_TABLE (13 bits). Sum of VI and VP bits. Refer to section 9.2.1 "VI_VPI_TABLE" starting on page 175. 4K VCI_TABLE (12 bits). Refer to "VCI_TABLE" on page 176. 4K service table. 8K multicast input pointer FIFOs 2 × 8 × 512. 16K multicast cell output pointer FIFOs 2 × 8 × 32 × 32. 8K transmit cell buffer linked list (one-half of the PHY). 8K receive cell buffer linked list. 1_h AL_RAM = 128K × 16 SSRAM. 16K VI_VPI_TABLE (14 bits). 16K VCI_TABLE (14 bits). 4K service table. 16K multicast input pointer FIFOs 2 × 8 × 1K. 16K multicast cell output pointer FIFOs 2 × 8 × 32 × 32. 16K transmit cell buffer linked list. 16K receive cell buffer linked list. 2_h AL_RAM = 256K × 16 SSRAM. 32K VI_VPI_TABLE (15 bits). 32K VCI_TABLE (15 bits). 4K service table. 16K multicast cell output pointer FIFOs 2 × 8 × 32 × 32. 16K multicast input pointer FIFOs 2 × 8 × 2K. 16K or 32K transmit cell buffer linked list. 16K or 32K or 64K receive cell buffer linked list. 3_h AL_RAM = 512K × 16 SSRAM. 128K VI_VPI_TABLE (17 bits). 64K VCI_TABLE (16 bits). 4K service table. 32K multicast input pointer FIFOs 2 × 8 × 2K. 32K multicast cell output pointer FIFOs 2 × 8 × 32 × 64. 16K or 32K or 64K transmit cell buffer linked list. 16K or 32Kor 64K receive cell buffer linked list. Resets to 0_b.
Not used (7:2)	Write with a 0 to maintain software compatibility with future versions.
CH_RAM_CONFIG (1:0)	Defines the size of the SRAM used by the channel tables. The processor must know or determine the SRAM size. Then, the processor must inform the device of the SRAM size by configuring this field. Initialize to the proper setting. This setting has no effect on the QRT operation, but it may in a future QRT device version. 0_h CH_RAM = 32K × 32 SSRAM. 2K channels. 1_h CH_RAM = 64K × 32 SSRAM. 4K channels. 2_h CH_RAM = 128K × 32 SSRAM. 8K channels. 3_h CH_RAM = 256K × 32 SSRAM. 16K channels. Resets to 0_b.

7.2.5 SWITCH_CONFIG

Address: 4_h (10_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:20)	Write with a 0 to maintain software compatibility with future versions.
ENABLE_PORT(3:0) (19:16)	Bit 19 corresponds to port 3, bit 18 corresponds to port 2, bit 17 corresponds to port 1 and bit 16 corresponds to port 0 as follows: 1 Enable the switch fabric port, and allow cells to be sent to it. Also, if SF_LOOP (refer to "SF_LOOP" on page 103) is asserted, replicate the looped back cells out onto the switch fabric. 0 Disable the switch fabric port. Send no cells to it. Resets to 0.
CLEAR_UPPER_VPI (15)	1 Clears VPI(11:8) bits for all incoming (Rx) cells 0 Not clear VPI(11:8) bits for all incoming (Rx) cells Resets to 0.
ENFORCE_UNUSED_VCI (14)	1 Drop cells with any non-zero value set in the unused VCI bits. 0 Do not check for non-zero value in the unused VCI bits Resets to 0.
TX_COUNT_SENT_CLP (13)	1 Store the CLP = 0 sent cells in TX_CELLS_SENT. Store the CLP = 1 sent cells in TX_CH_CELLS_DROPPED for all channels (refer to "TX_CH_CELLS_SENT" on page 193 and "TX_CH_CELLS_DRPD" on page 193). 0 Use the TX_CELLS_SENT and TX_CH_CELLS_DRPD counters as documented in "TX_CH_CELLS_SENT" on page 193 and "TX_CH_CELLS_DRPD" on page 193 respectively.
Reserved (12:8)	Initialize to 0 at initial setup. Software modifications to these locations after setup may cause incorrect operation.
TIME_SLOT_PRIORITY (7)	1 Automatically increment the RX_CH_SWITCH_GROUP "RX_CH_CONFIG" on page 184 by 1 (1 to 2, and 2 to 3, 3 will not be incremented) when an Unicast cell was selected from the RX_SERVICE_TABLE ("RX_SERVICE_TABLE" on page 182). RX_CH_SWITCH_GROUP is used to set the PRIORITY(1:0) field in the SE_D_OUT "QRT-QSE Interface Cell Format." on page 91 nibble 1. 0 Use the RX_CH_SWITCH_GROUP value for PRIORITY "RX_CH_CONFIG" on page 184 Resets to 0.
GLOBAL_SF_SPARE (6)	1 Set the SF_SPARE ("QRT-QSE Interface Cell Format." on page 91) bit to 1 (nibble 0 of the SE_D_OUT) for all cells sent to the switch fabric 0 Use the RX_CH_SF_SPARE "RX_CH_CONFIG" on page 184 value in the RX_CH_CONFIG(4) word for the SF_SPARE bit Resets to 0.

Field (Bits)	Description
SEPARATE_OUT_MC (5)	1 Treat switch fabric outputs as separate outputs for multicast traffic. 0 Group switch fabric outputs as a single output for multicast traffic. Resets to 0.
Reserved (4)	Initialize to 0 at initial setup. Software modifications to this location after setup may cause incorrect operation. Resets to 0.
PARITY_FAIL_DIS (3)	1 Disable parity checking on the switch fabric. 0 Enable parity checking on the switch fabric. Resets to 0.
CPA_OFF (2)	1 Turn off complex phase aligner at the switch fabric interface. 0 Use complex phase aligner to retime signals at the switch fabric interface. Resets to 0.
DISABLE_RANDOMIZATION (1)	1 Disable the randomization at the switch fabric interface. 0 Enable the randomization at the switch fabric interface.
Reserved (0)	Initialize to 1 at initial setup. Software modifications to this location after setup may cause incorrect operation. Resets to 0.

7.2.6 RAM BIST RESULT

Address: 5_h (14_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:21)	Write with a 0 to maintain software compatibility with future versions.
TF_RAM_BIST_FAIL (20)	If TF_RAM_BIST_CONTROL is set to 01_b 0 RAM BIST memory test has PASSED. 1 RAM BIST memory test has FAILED. If TF_RAM_BIST_CONTROL is set to 10_b 0 RAM BIST controller test has FAILED. 1 RAM BIST controller test has PASSED. Note: This bit must be cleared by writing a '0' before executing the RAM_BIST memory or controller tests. Resets to 0_b.
RF_RAM_BIST_FAIL (19)	If RF_RAM_BIST_CONTROL is set to 01_b 0 RAM BIST memory test has PASSED. 1 RAM BIST memory test has FAILED. If RF_RAM_BIST_CONTROL is set to 10_b 0 RAM BIST controller test has FAILED. 1 RAM BIST controller test has PASSED. Note: This bit must be cleared by writing a '0' before executing the RAM_BIST memory or controller tests. Resets to 0_b.
VO_RAM_BIST_FAIL (18)	If VO_RAM_BIST_CONTROL is set to 01_b 0 RAM BIST memory test has PASSED. 1 RAM BIST memory test has FAILED. If VO_RAM_BIST_CONTROL is set to 10_b 0 RAM BIST controller test has FAILED. 1 RAM BIST controller test has PASSED. Note: This bit must be cleared by writing a '0' before executing the RAM_BIST memory or controller tests. Resets to 0_b.

Field (Bits)	Description
RS_RAM_BIST_FAIL (17)	If RS_RAM_BIST_CONTROL is set to 01_b 0 RAM BIST memory test has PASSED. 1 RAM BIST memory test has FAILED. If RS_RAM_BIST_CONTROL is set to 10_b 0 RAM BIST controller test has FAILED. 1 RAM BIST controller test has PASSED. Note: This bit must be cleared by writing a '0' before executing the RAM_BIST memory or controller tests. Resets to 0_b.
TS_RAM_BIST_FAIL (16)	If TS_RAM_BIST_CONTROL is set to 01_b 0 RAM BIST memory test has PASSED. 1 RAM BIST memory test has FAILED. If TS_RAM_BIST_CONTROL is set to 10_b 0 RAM BIST controller test has FAILED. 1 RAM BIST controller test has PASSED. Note: This bit must be cleared by writing a '0' before executing the RAM_BIST memory or controller tests. Resets to 0_b.
Not used (15:5)	Write with a 0 to maintain software compatibility with future versions.
TF_RAM_BIST_COMPLETE (4)	0 TF_RAM_BIST memory test is in progress 1 TF_RAM_BIST memory test has completed Resets to 0_b.
RF_RAM_BIST_COMPLETE (3)	0 RF_RAM_BIST memory test is in progress 1 RF_RAM_BIST memory test has completed Resets to 0_b.
VO_RAM_BIST_COMPLETE (2)	0 VO_RAM_BIST memory test is in progress 1 VO_RAM_BIST memory test has completed Resets to 0_b.
RS_RAM_BIST_COMPLETE (1)	0 RS_RAM_BIST memory test is in progress 1 RS_RAM_BIST memory test has completed Resets to 0_b.
TS_RAM_BIST_COMPLETE (0)	0 TS_RAM_BIST memory test is in progress 1 TS_RAM_BIST memory test has completed Resets to 0_b.

7.2.7 MARKED_CELLS_COUNT

TX_MARKED_CELLS is a 4-bit counter for each of the four switch fabric input ports. TX_MARKED_CELLS counts cells that have the MB set in their tag.
RX_MARKED_CELLS is a 4-bit counter for each of the four switch fabric output ports. RX_MARKED_CELLS counts cells that have the MB set in their tag.

Address: 7_h ($1C_h$ byte)
Type: Read-only
Format: Refer to the following table.

Field (Bits)	Description
TX_MARKED_CELLS(3) (31:28)	The number of cells mod 16 with the MB set in the tag that were received through SE_D_IN(3,3:0).
TX_MARKED_CELL(2) (27:24)	The number of cells mod 16 with the MB set in the tag that were received through SE_D_IN(2,3:0).
TX_MARKED_CELLS(1) (23:20)	The number of cells mod 16 with the MB set in the tag that were received through SE_D_IN(1,3:0).
TX_MARKED_CELLS(0) (19:16)	The number of cells mod 16 with the MB set in the tag that were received through SE_D_IN(0,3:0).
RX_MARKED_CELLS(3) (15:12)	The number of cells mod 16 that were successfully transmitted out of SE_D_OUT(3,3:0).
RX_MARKED_CELLS(2) (11:8)	The number of cells mod 16 that were successfully transmitted out of SE_D_OUT(2,3:0).
RX_MARKED_CELLS(1) (7:4)	The number of cells mod 16 that were successfully transmitted out of SE_D_OUT(1,3:0).
RX_MARKED_CELLS(0) (3:0)	The number of cells mod 16 that were successfully transmitted out of SE_D_OUT(0,3:0).

7.2.8 CONDITION_PRES_BITS

Each interrupt has a condition present bit that indicates the condition is present now. Some failures are transitory, so this bit may read clear even though the condition is occurring frequently.

Address: 8_h (20_h byte)

Type: Read-only. All bits reset to 0, unless otherwise specified.

Format: Refer to the following table.

Field (Bits)	Description
ACK_LIVE_FAIL (31:28)	Bit 31 corresponds to port 3, bit 30 corresponds to port 2, bit 29 corresponds to port 1 and bit 28 corresponds to port 0 as follows: 1 No ACK, MNACK, or ONACK was received on the same cell time that a unicast cell was sent. This indicates that the end-to-end path may be bad. 0 Normal operation. Refer to "Unacknowledged Cell Detection" on page 48, "Liveness of Backpressure Signal" on page 48, "BP_ACK Remote Failure Detection" on page 50, and to Table 2 on page 50 for more information about the ACK_LIVE_FAIL interrupt.
UT_SOC_FAIL (27)	1 Failure was detected on the RATM_SOC signal. 0 Normal operation.
UT_CLK_FAIL (26)	1 Failure was detected on the UTOPIA clock (ATM_CLK). 0 Normal operation.
SF_CLK_FAIL (25)	1 Failure was detected on the switch fabric clock (SE_CLK). 0 Normal operation.
RX_UTOP_HEC_FAIL (24)	1 The last HEC on the receive UTOPIA interface was bad. 0 The last HEC on the receive UTOPIA interface was good.
RX_DRAM_PARITY_FAIL (23)	1 The last parity check on the RX_DRAM was bad. 0 The last parity check on the RX_DRAM was good.
TX_DRAM_PARITY_FAIL (22)	1 The last parity check on the TX_DRAM was bad. 0 The last parity check on the TX_DRAM was good.
AL_RAM_PARITY_FAIL (21)	1 The last parity check on the AL_RAM was bad. 0 The last parity check on the AL_RAM was good.
CH_RAM_PARITY_FAIL (20)	1 The last parity check on the CH_RAM was bad. 0 The last parity check on the CH_RAM was good.
TX_PARITY_FAIL (19:16)	Bit 19 corresponds to port 3, bit 18 corresponds to port 2, bit 17 corresponds to port 1 and bit 16 corresponds to port 0 as follows: 1 The transmit parity calculated over the first 12 nibbles of the cell header coming into the QRT on SE_D_IN#(3:0) was bad during the last valid cell time. 0 The transmit parity was good during the last valid cell time.
BP_IGNORED (15)	1 The backpressure was ignored during the last cell time, because a cell was received in violation of the backpressure that had been given. 0 The backpressure was not ignored during the last cell time.

(Continued)

Field (Bits)	Description
SE_INPUT_PORT_FAIL (14:11)	Bit 14 corresponds to port 3, bit 13 corresponds to port 2, bit 12 corresponds to port 1 and bit 11 corresponds to port 0 as follows: 1 Indicates one of the following: • No SOC or "10" pattern was present on a switch fabric port input, or • invalid cell code, or • bad idle cell code, or • fabric out of synch. 0 The switch fabric input was normal. Refer to "SOC Coding" on page 47, "SOC Inversions" on page 47, "Redundant Cell Present Coding" on page 47, "Idle Cell Pattern Checking" on page 47, Table 3 on page 51, and to Table 5 on page 52 for more information about the SE_INPUT_PORT_FAIL interrupt.
BP_REMOTE_FAIL (10:7)	Bit 10 corresponds to port 3, bit 9 corresponds to port 2, bit 8 corresponds to port 1 and bit 7 corresponds to port 0 as follows: 1 The forward data path from the QRT to the QSE was found to be bad through signaling on the BP_ACK_IN(3:0) line. 0 The forward data path from the QRT to the QSE was good. Refer to "Remote Data Path Failure Indication" on page 48, "Liveness of Backpressure Signal" on page 48, "BP_ACK Remote Failure Detection" on page 50, Table 2 on page 50, and to Table 4 on page 52 for more information about the BP_REMOTE_FAIL interrupt.
BP_ACK_FAIL (6:3)	Bit 6 corresponds to port 3, bit 5 corresponds to port 2, bit 4 corresponds to port 1 and bit 3 corresponds to port 0 as follows: 1 BP_ACK_IN(3:0) did not have the special code. 0 BP_ACK_IN(3:0) had the special code. Refer to "Liveness of Backpressure Signal" on page 48, "BP_ACK_IN Pattern Checking" on page 49, "BP_ACK Inversion Checking" on page 49, Table 2 on page 50, and to Table 4 on page 52 for more information about the BP_ACK_FAIL interrupt.
OUT_WD_INT (2)	1 At least one of the UTOPIA output watchdogs is currently expired. 0 None of the UTOPIA output watchdogs is currently expired.
EQD_INT (1)	1 The device is in the empty congestion state. 0 The device is not in the empty congestion state.
Reserved (0)	Initialize to 0 at initial setup. Software modifications to these locations after setup may cause incorrect operation.

7.2.9 CONDITION_LATCH_BITS

Each interrupt has a latch bit indicating there has been a failure detected since the last time that the latch word was read.

Address: 9_h (24_h byte)
Type: Read-only – Cleared on Read.
Reset Value: All bits reset to 0, unless specified.
Format: Refer to the following table.

Field (Bits)	Description
ACK_FAIL_LATCH (31:28)	Bit 31 corresponds to port 3, bit 30 corresponds to port 2, bit 29 corresponds to port 1 and bit 28 corresponds to port 0 as follows: 1 No ACK, MNACK, or ONACK was received on the same cell time that a unicast cell was sent. 0 Normal operation. Refer to section 3 “Fault Tolerance” starting on page 44 for more information.
UT_SOC_FAIL_LATCH (27)	1 Failure was detected on the RATM_SOC signal. 0 Normal operation.
UT_CLK_FAIL_LATCH (26)	1 Failure was detected on the UTOPIA clock (ATM_CLK). 0 Normal operation.
SF_CLK_FAIL_LATCH (25)	1 Failure was detected on the switch fabric clock (SE_CLK). 0 Normal operation.
RX_UTOP_HEC_FAIL_LATCH (24)	1 The HEC of the receive UTOPIA interface was bad at least once since this latch was read. 0 The HEC of the receive UTOPIA interface was good every time it was measured since this latch was read.
RX_DRAM_PARITY_FAIL_LATCH (23)	1 The parity of the RX_DRAM was bad at least once since this latch was read. 0 The parity of the RX_DRAM was good every time it was measured since this latch was read.
TX_DRAM_PARITY_FAIL_LATCH (22)	1 The parity of the TX_DRAM was bad at least once since this latch was read. 0 The parity of the TX_DRAM was good every time it was measured since this latch was read.
AL_RAM_PARITY_FAIL_LATCH (21)	1 The parity of the AL_RAM was bad at least once since this latch was read. 0 The parity of the AL_RAM was good every time it was measured since this latch was read.
CH_RAM_PARITY_FAIL_LATCH (20)	1 The parity of the CH_RAM was bad at least once since this latch was read. 0 The parity of the CH_RAM was good every time it was measured since this latch was read.
TX_PARITY_FAIL_LATCH (19:16)	Bit 19 corresponds to port 3, bit 18 corresponds to port 2, bit 17 corresponds to port 1 and bit 16 corresponds to port 0 as follows: 1 The transmit parity calculated over the first 12 nibbles of the cell header coming into the QRT on SE_D_IN#(3:0) was bad at least once since this latch was read. 0 The transmit parity was good in every cell since this latch was read.

(Continued)

Field (Bits)	Description
BP_IGNORED_FAIL_LATCH (15)	1 Backpressure was ignored at least once since this latch was read. 0 Backpressure was not ignored since this latch was read. Refer to section 3 “Fault Tolerance” starting on page 44 for more information.
SE_INPUT_PORT_FAIL_LATCH (14:11)	Bit 14 corresponds to port 3, bit 13 corresponds to port 2, bit 12 corresponds to port 1 and bit 11 corresponds to port 0 as follows: 1 One of the following was detected at least once since this latch was read: • Bad SOC pattern • invalid cell code • bad idle cell code • fabric out of synch 0 None of the above has been detected since this latch was read. Refer to section 3 “Fault Tolerance” starting on page 44 for more information.
BP_REMOTE_FAIL_LATCH (10:7)	Bit 10 corresponds to port 3, bit 9 corresponds to port 2, bit 8 corresponds to port 1 and bit 7 corresponds to port 0 as follows: 1 The forward data path was bad at least once since this latch was read. 0 The forward data path was not bad at least once since this latch was read. Refer to section 3 “Fault Tolerance” starting on page 44 for more information.
BP_ACK_FAIL_LATCH (6:3)	Bit 6 corresponds to port 3, bit 5 corresponds to port 2, bit 4 corresponds to port 1 and bit 3 corresponds to port 0 as follows: 1 The BP_ACK_IN(3:0) line did not have the required pattern at least once since this latch was read. 0 The BP_ACK_IN(3:0) line did have the required pattern at least once since this latch was read.
OUT_WD_INT_LATCH (2)	1 At least one of the UTOPIA output watchdogs expired since the last time this register was read. 0 None of the UTOPIA output watchdogs expired since the last time this register was read.
EQD_INT_LATCH (1)	1 The empty queue depth threshold was crossed in either direction since this latch was read. 0 The empty queue depth threshold was not crossed in either direction since this latch was read. The congestion is declared to have been relieved when the queue depth crossed back above twice the configured congestion threshold.
Reserved (0)	Initialize to 0 at initial setup. Software modifications to these locations after setup may cause incorrect operation.

7.2.10 INTR_MASK

Each interrupt has a mask bit that prevents the interrupt pin (/INTR) from being asserted, even though the corresponding latch bit is set.

Address: A_h (28_h byte)
Type: Read/Write
Reset Value:1
Format: Refer to the following table.

Field (Bits)	Description
MASK_ACK_FAIL (31:28)	Bit 31 corresponds to port 3, bit 30 corresponds to port 2, bit 29 corresponds to port 1 and bit 28 corresponds to port 0 as follows: 1 Mask the ACK_LIVE_FAIL interrupt (refer to "ACK_LIVE_FAIL" on page 111). 0 Enable ACK_LIVE_FAIL.
MASK_UT_SOC_FAIL (27)	1 Mask the receive ATM SOC fail interrupt (refer to "UT_SOC_FAIL" on page 111). 0 Normal operation.
MASK_UT_CLK_FAIL (26)	1 Mask the UTOPIA clock fail interrupt (refer to "UT_CLK_FAIL" on page 111). 0 Enable the UTOPIA clock fail interrupt.
MASK_SF_CLK_FAIL (25)	1 Mask the switch fabric clock fail interrupt (refer to "SF_CLK_FAIL" on page 111). 0 Enable the switch fabric clock fail interrupt.
MASK_RX_UTOP_HEC_FAIL (24)	1 Mask the receive UTOPIA HEC fail interrupt (refer to "RX_UTOP_HEC_FAIL" on page 111). 0 Enable the receive UTOPIA HEC fail interrupt.
MASK_RX_RAM_PARITY_FAIL (23)	1 Mask the receive RAM parity fail interrupt (refer to "RX_DRAM_PARITY_FAIL" on page 111). 0 Enable the RX RAM parity fail interrupt.
MASK_TX_RAM_PARITY_FAIL (22)	1 Mask the transmit RAM parity fail interrupt (refer to "TX_DRAM_PARITY_FAIL" on page 111). 0 Enable the TX RAM parity fail interrupt.
MASK_AL_RAM_PARITY_FAIL (21)	1 Mask the AL RAM parity fail interrupt (refer to "AL_RAM_PARITY_FAIL" on page 111). 0 Enable the AL RAM parity fail interrupt.
MASK_CH_RAM_PARITY_FAIL (20)	1 Mask the CH RAM parity fail interrupt (refer to "CH_RAM_PARITY_FAIL" on page 111). 0 Enable the CH RAM parity fail interrupt.
MASK_TX_PARITY_FAIL (19:16)	Bit 19 corresponds to port 3, bit 18 corresponds to port 2, bit 17 corresponds to port 1 and bit 16 corresponds to port 0 as follows: 1 Mask the transmit parity fail interrupt (refer to "TX_PARITY_FAIL" on page 111). 0 Enable the transmit parity fail interrupt.
MASK_BP_IGNORED (15)	1 Mask the backpressure ignored interrupt (refer to "BP_IGNORED" on page 111). 0 Enable the backpressure ignored interrupt.

(Continued)

Field (Bits)	Description
MASK_SE_INPUT_PORT_FAIL (14:11)	Bit 14 corresponds to port 3, bit 13 corresponds to port 2, bit 12 corresponds to port 1 and bit 11 corresponds to port 0 as follows: 1 Mask the switch element input port interrupt (refer to "SE_INPUT_PORT_FAIL" on page 112). 0 Enable the switch element input port interrupt.
MASK_BP_REMOTE_FAIL (10:7)	Bit 10 corresponds to port 3, bit 9 corresponds to port 2, bit 8 corresponds to port 1 and bit 7 corresponds to port 0 as follows: 1 Mask the backpressure remote fail interrupt (refer to "BP_REMOTE_FAIL" on page 112). 0 Enable the backpressure remote fail interrupt.
MASK_BP_ACK_FAIL (6:3)	Bit 6 corresponds to port 3, bit 5 corresponds to port 2, bit 4 corresponds to port 1 and bit 3 corresponds to port 0 as follows: 1 Mask the backpressure fail interrupt (refer to "BP_ACK_FAIL" on page 112). 0 Enable the backpressure fail interrupt.
MASK_OUT_WD_INT (2)	1 Mask the output watchdog interrupt (refer to "OUT_WD_INT" on page 112). 0 Enable the output watchdog interrupt.
MASK_EQD_INT (1)	1 Mask the empty queue depth interrupt (refer to "EQD_INT" on page 112). 0 Enable the empty queue depth interrupt.
Reserved (0)	Set to 1 for future software compatibility.

7.2.11 UTOPIA_CONFIG

The UTOPIA_CONFIG register configures both the receive and transmit UTOPIA interface during reset. Since this register initializes an internal state machine, alteration during non-reset conditions is not recommended. The transmit UTOPIA watchdog is configured from this register as well as the use of OUTCHAN insertion in the outgoing cells.
The MSP_MODE bit selects between the following polling methods:

- UTOPIA Level 2 standard polling that uses a single cell available signal (TATM_CLAV(0) and RATM_CLAV(0)).
- UTOPIA Level 2 optional MSP that uses four cell available signals (TATM_CLAV(3:0) and RATM_CLAV(3:0)) to poll as many as 31 devices in one cell time.

The CHK_HEC bit enables the checking of the HEC of incoming cells. If the device finds a HEC error, an interrupt is asserted and the cell is dropped in the receive UTOPIA block.

NOTE: The BACKPLANE_MODE bit was excluded from the QRT testing program.

The TX_OC_12C_MODE, RX_OC_12C_MODE, and UTOPIA_2 bits allow the user to select the configuration of receive and transmit side interfaces independently. The *_OC_12C_MODE bits take precedence over the UTOPIA_2 bit. For example, for a QRT configured to be RX_OC_12C_MODE and not TX_OC_12C_MODE and UTOPIA_2, the interface treats the receive side as a single-PHY interface and the transmit side as a multi-PHY interface requiring addressing. This feature is useful when interfacing to the PMC-Sierra PM7322 RCMP-800 device. Table 24 explains all combination of these bits.

Table 24. Various ways to configure UTOPIA interface

UTOPIACONFIG(2:0)	Explanation
111	Tx and Rx side configured for single PHY.
110	Tx and Rx side configured for single PHY.
101	Tx side configured for single PHY, Rx side configured for MPHY.
100	Tx and Rx side configured for single PHY.
011	Tx side configured for MPHY, Rx side configured for single PHY. Useful for RCMP applications.
010	Tx and Rx side configured for single PHY.
001	Tx and Rx side configured for MPHY.
000	Tx and Rx side configured for single PHY.

Address: 10_h (40_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:10)	Driven with a 0. Mask on reads to maintain software compatibility with future versions.
WD_TIME (9:8)	The period during which time the TATM_CLAV(3:0) must be asserted for a given output. Initialize to a setting that allows the slowest PHY device attached to the UTOPIA interface to respond in time. 0_h Set to a setting tolerant to DS0 outputs; tolerance equals 52,000 cell times. 1_h Set to a setting tolerant to DS1 outputs; tolerance equals 5000 cell times. 2_h Set to a setting tolerant to OC-3 outputs; tolerance equals 32 cell times. 3_h Disable the watchdog. Resets to 00_b.
OUT_CHAN_INSERTION (7)	1 Insert Tx Output Channel number in the 16 bit HEC/UDF field to be used downstream 0 Insert FFFFh in the HEC/UDF field. Resets to 0.
TDM (6)	1 Perform UTOPIA TDM Plus servicing. The TDM feature provides an additional level of servicing to compliment 2-level priority servicing. The suite of phy addresses included in the TDM table is based on UTENABLE 0 Perform only 2-level priority servicing. Resets to 0.
MSP_MODE (5)	1 Perform MSP (four cell available lines used). Use this setting when using direct status indication. 0 Perform UTOPIA Level 2 MPHY standard polling (one cell available line used). Resets to 0.
CHK_HEC (4)	1 Check the HEC of incoming cells. 0 Ignore the HEC of incoming cells. Use this setting when connecting to devices that do not drive the proper HEC on the UTOPIA bus or when the HEC field is used for other purposes. Resets to 0.
BACKPLANE_MODE (3)	NOTE: This feature was excluded from the QRT testing program. 1 Allow extra latency across the UTOPIA interface necessary to implement a 25 MHz (maximum) backplane. 0 Use normal UTOPIA Level 2 timing. Resets to 0.
TX_OC_12C_MODE (2)	1 The device is connected to an OC-12C line at transmit UTOPIA address 0. This overrides the setting of the UTOPIA_2 bit. 0 The device is connected to a normal transmit UTOPIA Level 2 interface. Resets to 0.
RX_OC_12C_MODE (1)	1 The device is connected to an OC-12C line at receive UTOPIA address 0. This overrides the setting of the UTOPIA_2 bit. 0 The device is connected to a normal receive UTOPIA Level 2 interface. Resets to 0.

(Continued)

Field (Bits)	Description
UTOPIA_2 (0)	1 The interface is in UTOPIA Level 2 mode. 0 The interface is in UTOPIA Level 1 mode. Resets to 0.

7.2.12 UT_PRIORITY

The UTOPIA interface provides at the very least, 2-level priority assignment scheme. UT_PRIORITY configures each UTOPIA interface to belong to either the high bandwidth or low bandwidth category. PHY devices with UT_PRIORITY = 1 are serviced with strict priority over PHY devices with UT_PRIORITY = 0. PHY devices within each bandwidth category are served lowest numbered PHY device address first. This register is configurable independent of software reset and can be changed during cell flow.

Address: 11_h (44_h byte)

Type: Read/Write

Format: Refer to the following table.

Field (Bits)	Description
Not used (31)	Write with a 0 to maintain software compatibility with future versions. PHY devices16-30 are defined to be at low priority.
UT_PRIORITY (30)	1 Service PHY device 30 at high priority. 0 Service PHY device 30 at low priority.
UT_PRIORITY (29)	1 Service PHY device 29 at high priority. 0 Service PHY device 29 at low priority.
. . .	. . .
UT_PRIORITY (0)	1 Service PHY device 0 at high priority. 0 Service PHY device 0 at low priority.

7.2.13 UT_ENABLE

The UT_ENABLE register enables whether or not the polling results from the PHY device will be considered during the UTOPIA PHY device selection mechanism. The UT_ENABLE register is also used to set the number of devices that will be polled when the standard polling method is used. Enabling only the devices that are used and configuring the system so the device addresses are grouped together in the lower address range provides the best results for the implemented standard polling algorithm. For MSP mode, all devices are polled regardless of whether or not they are enabled, but again, the enable is used to mask out those that are either not on the bus or are turned off. This register is configurable independent of software reset and can be changed during cell flow. When the UTOPIA_CONFIG TDM bit is set, the UT_ENABLE register is also used to define the TDM Priority Pool.

Address: 12_h (48_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description	
Not used (31)	Write with a 0 to maintain software compatibility with future versions.	
UT_ENABLE (30)	1	Enable cell transmission from PHY device 30.
	0	Disable cell transmission from PHY device 30.
UT_ENABLE (29)	1	Enable cell transmission from PHY device 29.
	0	Disable cell transmission from PHY device 29.
. . .	. . .	
UT_ENABLE (0)	1	Enable cell transmission from PHY device 0.
	0	Disable cell transmission from PHY device 0.

A transition on any of these bits clears the corresponding entry in the TX_UT_WD_ALIVE register (refer to “TX_UT_WD_ALIVE” on page 121).

7.2.14 TX_UT_STAT

TX_UT_STAT is a register that reflects the current polling status of the PHY device requests for service. The register is used as an input to the transmit UTOPIA watchdog register (refer to "TX_UT_WD_ALIVE" on page 121).

Address: 13_h ($4C_h$ byte)

Type: Read-only

Format: Refer to the following table.

Field (Bits)	Description
Not used (31)	Will be set to 0. Mask for future software compatibility.
TX_UT_STAT (30)	The PHY device at address 30 is ready to receive a cell.
TX_UT_STAT (29)	The PHY device at address 29 is ready to receive a cell.
. . .	. . .
TX_UT_STAT (0)	The PHY device at address 0 is ready to receive a cell.

7.2.15 TX_UT_WD_ALIVE

The TX_UT_WD_ALIVE bits in this register indicate a given PHY device has indicated it would accept a cell within the WD_TIME (refer to "WD_TIME" on page 118) when polled. Cells bound for PHY devices that fail this test are flushed in the background.

Address: 14_h (50_h byte)

Type: Read/Write

Format: Refer to the following table.

Field (Bits)	Description
Not used (31)	Write with a 0 to maintain software compatibility with future versions.
TX_UT_WD_ALIVE (30)	1 The PHY device at address 30 has responded to polling with an assertion of TATM_CLAV(3:0) for a period less than WD_TIME. 0 The PHY device at address 30 has not responded to a poll with an assertion of TATM_CLAV(3:0) for a period longer than WD_TIME. PHY device 30 has expired, therefore drain remaining cells intended for PHY device 30 at lowest priority.
TX_UT_WD_ALIVE (29)	1 The PHY device at address 29 has responded to polling with an assertion of TATM_CLAV(3:0) for a period less than WD_TIME. 0 The PHY device at address 29 has not responded to a poll with an assertion of TATM_CLAV(3:0) for a period longer than WD_TIME. PHY device 29 has expired, therefore drain remaining cells intended for PHY device 29 at lowest priority.
TX_UT_WD_ALIVE (28:1)	. . .

(Continued)

Field (Bits)	Description
TX_UT_WD_ALIVE (0)	1 The PHY device at address 0 has responded to polling with an assertion of TATM_CLAV(3:0) for a period less than WD_TIME. 0 The PHY device at address 0 has not responded to a poll with an assertion of TATM_CLAV(3:0) for a period longer than WD_TIME. PHY device 0 has expired, therefore drain remaining cells intended for PHY device 0 at lowest priority.

7.2.16 RX_CELL_START_ALIGN (Internal Structure)

RX_CELL_START_ALIGN configures the delay for the internal RxCellStart signal by taking the external RX_CELL_START signal, adding the 7-bit offset, and generating the internal RxCellStart signal. This allows flexibility and simplicity in the distribution of RX_CELL_START to QRTs on different cards or on different shelves.

Address: 20_h (80_h byte)

Type: Read/Write

Format: Refer to the following table.

Field (Bits)	Description
Not used (31:7)	Write with a 0 to maintain software compatibility with future versions.
RX_CELL_START_ALIGN (6:0)	Alignment value for the internal RxCellStart signal from the external RX_CELL_START signal.

7.2.17 RX_QUEUE_ENGINE_TEST

Address: 28_h ($A0_h$ byte)

Type: Read/Write

Function: The bits in this register disable individual functions within the queue engine for system debug purposes.

- Each instance of the "B" process (the UTOPIA receive Buffering process) performs the VPI/VCI mapping of a new cell from the UTOPIA interface to a receive CCB channel number.
- The "U" process (the UTOPIA enqueue process) uses the receive CCB channel number from the "B" process and peforms the cell enqueue function, including any congestion control.
- The "S" process (the receive switch fabric cell Sending process) fetches a cell for switch fabric dequeue.
- The "P" process (the receive switch fabric acknowledge Post-processing process) analyzes the ACK response from the previous cell transmission.

Reset Value: 0

Format: Refer to the following table.

Field (Bits)	Description
RESERVED (31:27)	Write with 0 to maintain future software compatibility.
RX_4_RR_PTRS (26)	1 Use 4 separate sets of round robin pointers, one set for each S process. Use this to achieve more predictable multicast performance when the bit SEPARATE_OUT_MC is set. This also allows more effective timeslot entries to be used when SEPARATE_OUT_MC is cleared. This setting can only be used when AL_RAM_CONFIG is not equal to 0. 0 Use 1 set of round robin pointers.
RX_STATS (25)	1 Enable the storage of RX_CH_STATS in the ABRAM 0 Disable the storage of RX_CH_STATS in the ABRAM. Use this setting when only 64K x 16 of memory is present or if the RX_CH_STATS are not needed.
RX_LONG_SVC_TBL (24)	1 Use 511 entry long RX_SERVICE_TABLE. This allows finer control of the bandwidth. This setting can only be used when AL_RAM_CONFIG is not equal to 0. 0 Use 127 entry long RX_SERVICE_TABLE.
RESERVED (23)	Write with 0 to maintain future software compatibility.
RX_SCG_ENABLE (22)	1 Enforce the Service Class Group congestion limits configured in RX_LWRxx_MAX_QD and RX_LWRxx_EXP_CONG_QD. Note: When this feature is enabled, the maximum queue depth for any of the three service class groups is 31K cells. 0 Ignore the configured SCG congestion limits
RX_RED_ENABLE (21)	1 Enable Random Early packet Discard function. The RX_xxx_EXP_CONG_QD that is used for each of the CH, SC, SCG, and DIR settings is randomly chosen to be the user configured one or the value one codepoint below the configured value. 0 RED disabled

Field (Bits)	Description	
RX_CLP_BEFORE_EPD_ENABLE (20)	1	Enable CLP before EPD dropping. When this feature is enabled, CLP=1 cells are dropped at the queue depth defined by the appropriate the codepoint one below that defined in RX_xxx_EXP_CONG_QD triggering PTD (Packet Tail Dropping). CLP=0 cells are considered for dropping by the EPD algorithm at the queue depth defined by RX_xxx_EXP_CONG_QD.
	0	Disable CLP before EPD dropping.
RX_HYSTERESIS_DISABLE (19)	1	Disable RX DIR, SCG, and SC congestion hyteresis. This setting declares that congestion is relieved when the queue depth drops below the queue depth defined by the appropriate RX_xxx_EXP_CONG_QD. This setting might achieve higher throughput in cases where the incoming traffic characteristics are less bursty.
	0	Enable RX DIR, SCG, and SC congestion hysteresis. This setting declares that congestion is relieved when the queue depth drops below the queue depth defined by the appropriate the codepoint one below that defined in RX_xxx_EXP_CONG_QD. This is a more robust setting which forces the switch to make the most of each dropping event at the expense of access to cell buffers in the steady state.

Field (Bits)	Description
RX_ALTERNATE_ENCODE (18)	1 Enable RX, SCG, and SC alternate congestion encoding. This set of values of the RX_xxx_EXP_MAX_QD and RX_xxx_EXP_CONG_QD defines many of the codepoints with higher values. This increases the resolution of the congestion limits where they are used most of the time. The coding is: 1111- 31K+1 (31745) 1110 - 24K+1 1101 - 16K+1 1100 - 12K+1 1011 - 8K+1 1010 - 6K+1 1001 - 4K+1 1000 - 3K+1 0111 - 2K+1 0110 - 1.5K+1 0101 - 1K+1 0100 - 769 0011 - 513 0010 - 257 0001 - 9 0000 - 0 0 Disable the RX, SCG, and SC alternate encoding. This setting uses the original set of values which are approximately 2**N. The coding is: 1111- 31K+1 (31745) 1110 - 16K+1 1101 - 8K+1 1100 - 4K+1 1011 - 2K+1 1010 - 1K+1 1001 - 513 1000 - 257 0111 - 129 0110 - 65 0101 - 33 0100 - 15 0011 - 9 0010 - 5 0001 - 3 0000 - 0
RX_CH_HYSTERESIS_DISABLE (17)	1 Disable RX_CH congestion hyteresis. This setting declares that congestion is relieved when the queue depth drops below the queue depth defined by the appropriate RX_xxx_EXP_CONG_QD. This setting might achieve higher throughput in cases where the incoming traffic characteristics are less bursty. 0 Enable RX_CH congestion hysteresis. This setting declares that congestion is relieved when the queue depth drops below the queue depth defined by the appropriate the codepoint one below that defined in RX_xxx_EXP_CONG_QD. This is a more robust setting which forces the switch to make the most of each dropping event at the expense of access to cell buffers in the steady state.

Field (Bits)		Description
RX_CH_ALTERNATE_ENCODE (16)	1	Enable RX_CH alternate congestion encoding. This set of values of the RX_CH_EXP_MAX_QD and RX_CH_EXP_CONG_QD defines many of the codepoints with higher values. This increases the resolution of the congestion limits where they are used most of the time. The table is shown in the text for RX_ALTERNATE_ENCODE.
	0	Disable RX_CH alternate encoding. This setting uses the original set of values which are approximately 2**N. The table is shown in the text for RX_ALTERNATE_ENCODE.
DISABLE_P4 (15)	1	Disable the 4th instance of the P process. The P process updates the data structures with the results of crossing the switch fabric. P4 works as a pair with S4.
	0	Enable this process.
DISABLE_P3 (14)	1	Disable the 3rd instance of the P process. P3 works as a pair with S3.
	0	Enable this process.
DISABLE_P2 (13)	1	Disable the 2nd instance of the P process. P2 works as a pair with S2.
	0	Enable this process.
DISABLE_P1 (12)	1	Disable the 1st instance of the P process. P1 works as a pair with S1.
	0	Enable this process.
DISABLE_S4 (11)	1	Disable the 4th instance of the S process. The S process selects the cells to be sent to the switch fabric. When DISABLE_RANDOMIZATION=1, S4 is associated with switch fabric port 3.
	0	Enable this process.
DISABLE_S3 (10)	1	Disable the 3rd instance of the S process.When DISABLE_RANDOMIZATION=1, S3 is associated with switch fabric port 2.
	0	Enable this process.
DISABLE_S2 (9)	1	Disable the 2nd instance of the S process. When DISABLE_RANDOMIZATION=1, S2 is associated with switch fabric port 1.
	0	Enable this process.
DISABLE_S1 (8)	1	Disable the 1st instance of the S process. When DISABLE_RANDOMIZATION=1, S1 is associated with switch fabric port 0.
	0	Enable this process.
DROP_ALL_RX_CELLS (7)	1	Drop all cells received in the RX direction.
	0	Normal operation.
DISABLE_U3 (6)	1	Disable the 3rd instance of the U process. The U process enqueues the cells onto the data structures.
	0	Enable this process.
DISABLE_U2 (5)	1	Disable the 2nd instance of the U process.
	0	Enable this process.
DISABLE_U1 (4)	1	Disable the 1st instance of the U process.
	0	Enable this process.

Field (Bits)	Description
RESERVED (3)	Write with 0 to maintain future software compatibility.
DISABLE_B3 (2)	1 Disable the 3rd instance of the B process. The B process determines the channel number and enqueues the cell into SDRAM. B3 works as a pair with U1. 0 Enable this process.
DISABLE_B2 (1)	1 Disable the 2nd instance of the B process. B2 works as a pair with U3. 0 Enable this process.
DISABLE_B1 (0)	1 Disable the 1st instance of the B process. B1 works as a pair with U2. 0 Enable this process.

7.2.18 TX_QUEUE_ENGINE_TEST

Address: 29_h ($A4_h$ byte)

Type: Read/Write

Function: The bits in this register disable individual functions within the queue engine for system debug purposes.

- The “E” process (the transmit switch fabric Enqueue process) performs the enqueue function for each cell arrived from the switch process. It determines if a cell is eligible for enqueue (that is, SENT or DROP) based on current congestion conditions.
- The “C” process (the update transmit Counters process) updates the per-channel statistics.
- The “M” process (the Multicast replication process) performs the multicast leaf replication task.
- The “D” process (the transmit UTOPIA Dequeue process) fetches a cell for transmit UTOPIA transmission.

Reset Value: 0

Format: Refer to the following table.

Field (Bits)	Description	
RESERVED (31:23)	Write with a 0 to maintain future software compatibility.	
TX_SCG_ENABLE (22)	1	Enforce the Service Class Group congestion limits configured in TX_LWRxx_EXP_MAX_QD and TX_LWRxx_EXP_CONG_QD. Note: When this feature is enabled, the maximum queue depth for any of the three service class groups is 31K cells.
	0	Ignore the configured SCG congestion limits
TX_RED_ENABLE (21)	1	Enable Random Early packet Discard function. The TX_xxx_EXP_CONG_QD that is used for each of the CH, SCQ, SC, SCG, VO and DIR settings is randomly chosen to be the user configured one or the value one codepoint below the configured value.
	0	RED disabled
TX_CLP_BEFORE_EPD_ENABLE (20)	1	Enable CLP before EPD dropping. When this feature is enabled, CLP=1 cells are dropped at the queue depth defined by the appropriate the codepoint one below that defined in TX_xxx_EXP_CONG_QD triggering PTD (Packet Tail Dropping). CLP=0 cells are considered for dropping by the EPD algorithm at the queue depth defined by TX_xxx_EXP_CONG_QD.
	0	Disable CLP before EPD dropping.
TX_HYSTERESIS_DISABLE (19)	1	Disable TX DIR, VO, SCG, SC and SCQ congestion hyteresis. This setting declares that congestion is relieved when the queue depth drops below the queue depth defined by the appropriate TX_xxx_EXP_CONG_QD. This setting might achieve higher throughput in cases where the incoming traffic characteristics are less bursty.
	0	Enable TX DIR, VO, SCG, SC, and SCQ congestion hysteresis. This setting declares that congestion is relieved when the queue depth drops below the queue depth defined by the appropriate the codepoint one below that defined in TX_xxx_EXP_CONG_QD. This is a more robust setting which forces the switch to make the most of each dropping event at the expense of access to cell buffers in the steady state.

Field (Bits)	Description
TX_ALTERNATE_ENCODE (18)	1 Enable TX, SCG, SC and SCQ alternate congestion encoding. This set of values of the TX_xxx_EXP_MAX_QD and TX_xxx_EXP_CONG_QD defines many of the codepoints with higher values. This increases the resolution of the congestion limits where they are used most of the time. The coding is: 1111- 31K+1 (31745) 1110 - 24K+1 1101 - 16K+1 1100 - 12K+1 1011 - 8K+1 1010 - 6K+1 1001 - 4K+1 1000 - 3K+1 0111 - 2K+1 0110 - 1.5K+1 0101 - 1K+1 0100 - 769 0011 - 513 0010 - 257 0001 - 9 0000 - 0 0 Disable the TX, SC, SCG, and SCQ alternate encoding. This setting uses the original set of values which are approximately 2**N. The coding is: 1111- 31K+1 (31745) 1110 - 16K+1 1101 - 8K+1 1100 - 4K+1 1011 - 2K+1 1010 - 1K+1 1001 - 513 1000 - 257 0111 - 129 0110 - 65 0101 - 33 0100 - 15 0011 - 9 0010 - 5 0001 - 3 0000 - 0
TX_CH_HYSTERESIS_DISABLE (17)	1 Disable TX_CH congestion hyteresis. This setting declares that congestion is relieved when the queue depth drops below the queue depth defined by the appropriate TX_xxx_EXP_CONG_QD. This setting might achieve higher throughput in cases where the incoming traffic characteristics are less bursty. 0 Enable TX_CH congestion hysteresis. This setting declares that congestion is relieved when the queue depth drops below the queue depth defined by the appropriate the codepoint one below that defined in TX_xxx_EXP_CONG_QD. This is a more robust setting which forces the switch to make the most of each dropping event at the expense of access to cell buffers in the steady state.

Field (Bits)	Description
TX_CH_ALTERNATE_ENCODE (16)	1 Enable TX_CH alternate congestion encoding. This set of values of the TX_CH_EXP_MAX_QD and TX_CH_EXP_CONG_QD defines many of the codepoints with higher values. This increases the resolution of the congestion limits where they are used most of the time. The table is shown in the text for RX_ALTERNATE_ENCODE. 0 Disable TX_CH alternate encoding. This setting uses the original set of values which are approximately 2**N. The table is shown in the text for RX_ALTERNATE_ENCODE.
TX_DISABLE_BP (15)	1 Do not send BP to the switch fabric when the SC MC infifos are almost full. Drop cells that are delivered to full queues. This prevents one egress port from congesting the fabric through excessive BP. 0 Send BP to the switch fabric when the SC MC infifos are nearly full. This setting allows the egress port to momentarity push back onto the fabric when its INFIFO becomes full.
RESERVED (13:14)	Write with a 0 to maintain future software compatibility.
DROP_ALL_TX_CELLS (12)	1 Drop all cells received in the TX direction. 0 Normal operation.
RESERVED (11)	Write with 0 to maintain future software compatibility.
DISABLE_D3 (10)	1 Disable the 3rd instance of the D process. The D process dequeues cells to the UTOPIA. 0 Enable this process.
DISABLE_D2 (9)	1 Disable the 2nd instance of the D process. 0 Enable this process.
DISABLE_D1 (8)	1 Disable the 1st instance of the D process. 0 Enable this process.
DISABLE_O (7)	1 Disable the the o process. The O process updates the TX CH statistics. 0 Enable this process.
RESERVED (6)	Write with 0 to maintain future software compatibility.
DISABLE_M2 (5)	1 Disable the 2nd instance of the M process. The M process replicates the multicast pointers. 0 Enable this process.
DISABLE_M1_3 (4)	1 Disable the 1st and 3rd instance of the M process. 0 Enable this process.
DISABLE_E4 (3)	1 Disable the 4th instance of the E process. The E process enqueues the cells onto the data structures from the fabric. E4 is associated with switch fabric port 3. 0 Enable this process.
DISABLE_E3 (2)	1 Disable the 3rd instance of the E process. E3 is associated with switch fabric port 2. 0 Enable this process.
DISABLE_E2 (1)	1 Disable the 2nd instance of the E process. E2 is associated with switch fabric port 1. 0 Enable this process.

Field (Bits)	Description	
DISABLE_E1 (0)	1	Disable the 1st instance of the E process. E1 is associated with switch fabric port 1.
	0	Enable this process.

7. 2. 19. QUEUE ENGINE CONDITION_PRES_BITS

Address: $2A_h$ ($A8_h$ byte)
Read-Only. All bits unless specified Resets to 0.
Format:

Field (Bits)	Description	
AB_RAM_PARITY_FAIL (31)	1	The parity on the last AB RAM read was not odd.
	0	The parity on the last AB RAM read was odd.
SF_FABRIC_NOT_SYNC_FAIL (30)	1	The RX_CELL_START pulse did not arrive at exactly 118 SF_CLK clock interval. The fabric interface is out of sync because of this condition.
	0	The RX_CELL_START pulse is occuring at exactly 118 SF_CLK clocks interval. The fabric interface is in sync.
Not used (29:27)	Driven with a 0.	
CH_UNUSED_VCI_SET (26)	1	A cell was received with a VPI with a VPI with VCC_ENTRY=1 and a VCI bit set great than that defined by the value of VCI_BITS for that VPI.
	0	The cell that was just received did not have an unused VCI bit set.
QE_TX_SCG_FAIL (25)	1	An illegal state is present in one of the three Tx Service Class Group counters. This check asserts that whenever there are no cells in the lower 12 services classes, there must also be no cells in the lower 4 service classes. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	The Tx Service Class Group counters are currently in a legal state.
QE_RX_SCG_FAIL (24)	1	An illegal state is present in one of the three Rx Service Class Group counters. This check asserts that whenever there are no cells in the lower 48 services classes, there must also be no cells in the lower 16 service classes. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	The Rx Service Class Group counters are currently in a legal state.
QE_TX_LOWER12_SCG_QD_NEG (23)	1	The Tx Lower12 Service Class Group Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	The Tx Lower12 Service Class Group Queue Depth is non-negative.
QE_TX_LOWER8_SCG_QD_NEG (22)	1	The Tx Lower8 Service Class Group Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	The Tx Lower8 Service Class Group Queue Depth is non-negative.
QE_TX_LOWER4_SCG_QD_NEG (21)	1	The Tx Lower4 Service Class Group Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	The Tx Lower4 Service Class Group Queue Depth is non-negative.
QE_TX_DIR_QD_NEG (20)	1	The Tx Direction Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	The Tx Direction Queue Depth is non-negative.

Field (Bits)	Description	
QE_RX_LOWER48_SCG_QD_NEG (19)	1	The Rx Lower48 Service Class Group Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	The Rx Lower48 Service Class Group Queue Depth is non-negative.
QE_RX_LOWER32_SCG_QD_NEG (18)	1	The Rx Lower32 Service Class Group Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	The Rx Lower32 Service Class Group Queue Depth is non-negative.
QE_RX_LOWER16_SCG_QD_NEG (17)	1	The Rx Lower16 Service Class Group Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	The Rx Lower16 Service Class Group Queue Depth is non-negative.
QE_RX_DIR_QD_NEG (16)	1	The Rx Direction Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	The Rx Direction Queue Depth is non-negative.
Not Used (15)	Driven with a 0	
QE_INVALID_CSTART (14)	1	The internally generated Rx cell start did not arrive within 254 SYSCLKs. This is a sign that the SECLK is absent or is running too slowly.
	0	The internally generated Rx cell start arrived within 254 SYSCLKs.
QE_RX_SEQ_FAIL (13)	1	An illegal state is present in the Rx Sequencing state machine. This is likely a non-recoverable state for the device. The device should be reset and reinitialized.
	0	The Rx Sequencing state machine is currently in a legal state.
QE_TX_CH_QD_NEG (12)	1	A Tx Channel Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	A Tx Channel Queue Depth is non-negative.
QE_TX_SCQ_QD_NEG (11)	1	A Tx Service Class Queue, Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	A Tx Service Class Queue, Queue Depth is non-negative.
QE_TX_SC_QD_NEG (10)	1	A Tx Service Class Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	A Tx Service Class Queue Depth is non-negative.
QE_TX_VO_QD_NEG (9)	1	A Tx Virtual Output Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	A Tx Virtual Output Queue Depth is non-negative.
QE_RX_CH_QD_NEG (8)	1	A Rx Channel Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	A Rx Channel Queue Depth is non-negative.
QE_RX_SC_QD_NEG (7)	1	A Rx Service Class Queue Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	A Rx Service Class Queue Depth is non-negative.

Field (Bits)		Description
QE_RX_CH_COUNT_NEG (6)	1	A Rx Channel Count Depth is negative. This is a non-recoverable state for the device. The device should be reset and reinitialized.
	0	A Rx Channel Count Depth is non-negative.
QE_RX_ETOP_PARITY (5)	1	The pointer to the top of the Rx empty stack has even parity. This is an error if the driver has initialized all of the pointers to be odd. If 64K cell buffers are used, or if the driver did not initialize the pointers to have odd parity by using the MSB as a parity bit, this bit may be set.
	0	The pointer to the top of the Rx empty stack has odd parity.
QE_TX_ETOP_PARITY (4)	1	The pointer to the top of the Tx empty stack has even parity. This is an error if the driver has initialized all of the pointers to be odd. If 64K cell buffers are used, or if the driver did not initialize the pointers to have odd parity by using the MSB as a parity bit, this bit may be set.
	0	The pointer to the top of the Tx empty stack has odd parity.
QE_UNMAPPED_VI (3)	1	The Rx UTOPIA I/F is delivering a cell from a virtual input that it was not supposed to be polling.
	0	.The Rx UTOPIA I/F is not delivering a cell from a VI that it was not supposed to be polling.
QE_VO_MISMATCH (2)	1	The Tx UTOPIA I/F is asking for a cell from a VO which does not have a cell. This is normal if TX_OC_12C_MODE = 1.
	0	The Tx UTOPIA I/F is not asking for a cell from a VO which does not have a cell.
QE_RX_CIRC_LINK (1)	1	A circular link is being written to the RX linked list. This non-recoverable state occurs when a link is added to the RX linked list that points to itself. The device should be reset and reinitialized.
	0	A circular link is not being written to the Rx linked list.
QE_TX_CIRC_LINK (0)	1	A circular link is being written to the Tx linked list. This non-recoverable state occurs when a link is added to the TX linked list that points to itself. The device should be reset and reinitialized.
	0	A circular link is not being written to the Tx linked list.

7. 2. 20. QUEUE_ENGINE_CONDITION_LATCH_BITS

Address: $2B_h$ (AC_h byte)
Read-Only - Cleared on Read. All bits Reset to 0 unless specified.
Format:

Field (Bits)	Description
AB_RAM_PARITY_FAIL_LATCH (31)	1 The AB_RAM_PARITY_FAIL condition occurred since this latch was read 0 The AB_RAM_PARITY_FAIL condition has not occurred since this latch was read
SF_FABRIC_NOT_SYNC_LATCH (30)	1 The RX_CELL_START pulse did not arrive at exactly 118 SF_CLK clock interval. The fabric interface is out of sync because of this condition. 0 The RX_CELL_START pulse is occuring at exactly 118 SF_CLK clocks interval. The fabric interface is in sync.
Not used (29:27)	Driven with a 0.
CH_UNUSED_VCI_SET_LATCH (26)	1 The CH_UNUSED_VCI_SET condition occurred since this latch was read 0 The CH_UNUSED_VCI_SET condition has not occurred since this latch was read
QE_TX_SCG_FAIL_LATCH (25)	1 The QE_TX_SCG_FAIL condition occurred since this latch was read 0 The QE_TX_SCG_FAIL condition has not occurred since this latch was read
QE_RX_SCG_FAIL_LATCH (24)	1 The QE_RX_SCG_FAIL condition occurred since this latch was read 0 The QE_RX_SCG_FAIL condition has not occurred since this latch was read
QE_TX_LOWER12_SCG_QD_NEG_LATCH (23)	1 The QE_TX_LOWER12_SCG_QD_NEG condition occurred since this latch was read 0 The QE_TX_LOWER12_SCG_QD_NEG condition has not occurred since this latch was read
QE_TX_LOWER8_SCG_QD_NEG_LATCH (22)	1 The QE_TX_LOWER8_SCG_QD_NEG condition occurred since this latch was read 0 The QE_TX_LOWER8_SCG_QD_NEG condition has not occurred since this latch was read
QE_TX_LOWER4_SCG_QD_NEG_LATCH (21)	1 The QE_TX_LOWER4_SCG_QD_NEG condition occurred since this latch was read 0 The QE_TX_LOWER4_SCG_QD_NEG condition has not occurred since this latch was read
QE_TX_DIR_QD_NEG_LATCH (20)	1 The QE_TX_DIR_QD_NEG condition occurred since this latch was read 0 The QE_TX_DIR_QD_NEG condition has not occurred since this latch was read

Field (Bits)	Description
QE_RX_LOWER48_SCG_QD_NEG_LATCH (19)	1 The QE_RX_LOWER48_SCG_QD_NEG condition occurred since this latch was read 0 The QE_RX_LOWER48_SCG_QD_NEG condition has not occurred since this latch was read
QE_RX_LOWER32_SCG_QD_NEG_LATCH (18)	1 The QE_RX_LOWER32_SCG_QD_NEG condition occurred since this latch was read 0 The QE_RX_LOWER32_SCG_QD_NEG condition has not occurred since this latch was read
QE_RX_LOWER16_SCG_QD_NEG_LATCH (17)	1 The QE_RX_LOWER16_SCG_QD_NEG condition occurred since this latch was read 0 The QE_RX_LOWER16_SCG_QD_NEG condition has not occurred since this latch was read
QE_RX_DIR_QD_NEG_LATCH (16)	1 The QE_RX_DIR_QD_NEG condition occurred since this latch was read 0 The QE_RX_DIR_QD_NEG condition has not occurred since this latch was read
Not used (15)	Driven with a 0.
QE_INVALID_CSTART_LATCH (14)	1 The QE_INVALID_CSTART condition occurred since this latch was read 0 The QE_INVALID_CSTART condition has not occurred since this latch was read
QE_RX_SEQ_FAIL_LATCH (13)	1 The QE_RX_SEQ_FAIL condition has occurred since this latch was read. 0 The QE_RX_SEQ_FAIL condition has not occurred since this latch was read.
QE_TX_CH_QD_NEG_LATCH (12)	1 The QE_TX_CH_QD_NEG condition has occurred since this latch was read. 0 The QE_TX_CH_QD_NEG condition has not occurred since this latch was read.
QE_TX_SCQ_QD_NEG_LATCH (11)	1 The QE_TX_SCQ_QD_NEG condition has occurred since this latch was read. 0 The QE_TX_SCQ_QD_NEG condition has not occurred since this latch was read.
QE_TX_SC_QD_NEG_LATCH (10)	1 The QE_TX_SC_QD_NEG condition has occurred since this latch was read. 0 The QE_TX_SC_QD_NEG condition has not occurred since this latch was read.
QE_TX_VO_QD_NEG_LATCH (9)	1 The QE_TX_VO_QD_NEG condition has occurred since this latch was read. 0 The QE_TX_VO_QD_NEG condition has not occurred since this latch was read.
QE_RX_CH_QD_NEG_LATCH (8)	1 The QE_RX_CH_QD_NEG condition has occurred since this latch was read. 0 The QE_RX_CH_QD_NEG condition has not occurred since this latch was read.

Field (Bits)	Description	
QE_RX_SC_QD_NEG_LATCH (7)	1	The QE_RX_SC_QD_NEG condition has occurred since this latch was read.
	0	The QE_RX_SC_QD_NEG condition has not occurred since this latch was read.
QE_RX_CH_COUNT_NEG_LATCH (6)	1	The QE_RX_CH_COUNT_NEG condition has occurred since this latch was read.
	0	The QE_RX_CH_COUNT_NEG condition has not occurred since this latch was read.
QE_RX_ETOP_PARITY_LATCH (5)	1	The QE_RX_ETOP_PARITY condition has occurred since this latch was read. If 64K cell buffers are used, or if the driver did not initialize the pointers to have odd parity by using the MSB as a parity bit, this bit may be set.
	0	The QE_RX_ETOP_PARITY condition has not occurred since this latch was read.
QE_TX_ETOP_PARITY_LATCH (4)	1	The QE_TX_ETOP_PARITY condition has occurred since this latch was read. If 64K cell buffers are used, or if the driver did not initialize the pointers to have odd parity by using the MSB as a parity bit, this bit may be set.
	0	The QE_TX_ETOP_PARITY condition has not occurred since this latch was read.
QE_UNMAPPED_VI_LATCH (3)	1	The QE_UNMAPPED_VI condition has occurred since this latch was read.
	0	The QE_UNMAPPED_VI condition has not occurred since this latch was read.
QE_VO_MISMATCH_LATCH (2)	1	The QE_VO_MISMATCH condition has occurred since this latch was read. This is normal if TX_OC_12C_MODE = 1.
	0	The QE_VO_MISMATCH condition has not occurred since this latch was read.
QE_RX_CIRC_LINK_LATCH (1)	1	The QE_RX_CIRC_LINK condition has occurred since this latch was read.
	0	The QE_RX_CIRC_LINK condition has not occurred since this latch was read.
QE_TX_CIRC_LINK_LATCH (0)	1	The QE_TX_CIRC_LINK condition has occurred since this latch was read.
	0	The QE_TX_CIRC_LINK condition has not occurred since this latch was read.

7. 2. 21. QUEUE_ENGINE_INT_MASK

Address: $2C_h$ ($B0_h$ byte)
Read/Write
Function: The following MASK_ bits disable the generation of an interrupt through the INTR pin. They do not, however, affect the operation of the associated interrupt or status bits.
Reset Value:1
Format:

Field (Bits)	Description
MASK_AB_RAM_PARITY_FAIL (31)	1 Mask the AB_RAM_PARITY_FAIL interrupt. 0 Enable this interrupt.
MASK_SF_FABRIC_NOT_SYNC _FAIL (30)	1 Mask the The RX_CELL_START pulse not occuring at exactly 118 SF_CLK clock interval interrupt 0 Enable this interrupt.
Not used (29:27)	Write with a 1 to maintain future software compatibility.
MASK_CH_UNUSED_VCI_SET (26)	1 Mask the CH_UNUSED_VCI_SET interrupt. 0 Enable this interrupt.
MASK_QE_TX_SCG_FAIL (25)	1 Mask the QE_TX_SCG_FAIL interrupt. 0 Enable this interrupt.
MASK_QE_RX_SCG_FAIL (24)	1 Mask the QE_RX_SCG_FAIL interrupt. 0 Enable this interrupt.
MASK_QE_TX_LOWER12_SCG_ QD_NEG (23)	1 Mask the QE_TX_LOWER12_SCG_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_TX_LOWER8_SCG_ QD_NEG (22)	1 Mask the QE_TX_LOWER8_SCG_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_TX_LOWER4_SCG_ QD_NEG (21)	1 Mask the QE_TX_LOWER4_SCG_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_TX_DIR_QD_NEG (20)	1 Mask the QE_TX_DIR_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_RX_LOWER48_SCG_ QD_NEG (19)	1 Mask the QE_RX_LOWER48_SCG_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_RX_LOWER32_SCG_ QD_NEG (18)	1 Mask the QE_RX_LOWER32_SCG_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_RX_LOWER16_SCG_ QD_NEG (17)	1 Mask the QE_RX_LOWER16_SCG_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_RX_DIR_QD_NEG (16)	1 Mask the QE_RX_DIR_QD_NEG interrupt. 0 Enable this interrupt.

Field (Bits)	Description
Not used (15)	Write with a 1 to maintain future software compatibility.
MASK_QE_INVALID_CSTART (14)	1 Mask the QE_INVALID_CSTART interrupt. 0 Enable this interrupt.
MASK_QE_RX_SEQ_FAIL (13)	1 Mask the QE_RX_SEQ_FAIL interrupt. 0 Enable this interrupt.
MASK_QE_TX_CH_QD_NEG (12)	1 Mask the QE_TX_CH_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_TX_SCQ_QD_NEG (11)	1 Mask the QE_TX_SCQ_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_TX_SC_QD_NEG (10)	1 Mask the QE_TX_SC_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_TX_VO_QD_NEG (9)	1 Mask the QE_TX_VO_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_RX_CH_QD_NEG (8)	1 Mask the QE_RX_CH_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_RX_SC_QD_NEG (7)	1 Mask the QE_RX_SC_QD_NEG interrupt. 0 Enable this interrupt.
MASK_QE_RX_CH_COUNT_NEG (6)	1 Mask the QE_RX_CH_COUNT_NEG interrupt 0 Enable this interrupt.
MASK_QE_RX_ETOP_PARITY (5)	1 Mask the QE_RX_ETOP_PARITY interrupt 0 Enable this interrupt.
MASK_QE_TX_ETOP_PARITY (4)	1 Mask the QE_RX_ETOP_PARITY interrupt 0 Enable this interrupt.
MASK_QE_UNMAPPED_VI (3)	1 Mask the QE_UNMAPPED_VI interrupt. 0 Enable this interrupt.
MASK_QE_VO_MISMATCH (2)	1 Mask the QE_VO_MISMATCH interrupt. 0 Enable this interrupt.
MASK_QE_RX_CIRC_LINK (1)	1 Mask the QE_RX_CIRC_LINK interrupt. 0 Enable this interrupt.
MASK_QE_TX_CIRC_LINK (0)	1 Mask the QE_TX_CIRC_LINK interrupt. 0 Enable this interrupt.

7.2.22 RX_DIR_CONFIG

The RX_DIR_EXP_ECONG_QD bit in this register sets the congestion threshold for the receive direction. When the number of free cell buffers falls below this threshold, the configured congestion management actions are taken. A low number sets this threshold to a high number of used cell buffers.

Address: $1F0000_h$ ($7C0000_h$ byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:16)	Write with a 0 to maintain software compatibility with future versions.
Reserved (15:12)	Write with a 3_h.
RX_DIR_EXP_ECONG_QD (11:8)	Exponent of the congested queue depth (this is, the number of empty cell buffers below which the device will enter the empty congestion state). Initialize to the proper setting. The congestion state is entered when the number of free buffers drops below this threshold, and the congestion state is exited when the number of free buffers exceeds twice this threshold.
Not used (7:0)	Write with a 0 to maintain software compatibility with future versions.

7.2.23 RX_DIR_STATE (Internal Structure)

Address: $1F0001_h$ ($7C0004_h$ byte)
Type: Read/Write – Write only during SW_RESET (refer to "SW_RESET" on page 101). Do not update this register.
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:18)	Write with a 0 to maintain software compatibility with future versions.
Reserved (17)	Initialize to 0 at initial setup. Software modifications to this location after setup may cause incorrect operation.
RX_DIR_CONG_STATE (16)	Congestion state. Initialize to 0.
RX_DIR_E_CUR_QD (15:0)	Current count of empty cells in the receive direction. Initialize to the number of cells setting. NOTE: The number of buffers is limited to the physical number of buffers minus three.

7.2.24 TX_DIR_CONFIG

Address: $1F0008_h$ ($7C0020_h$ byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:16)	Write with a 0 to maintain software compatibility with future versions.
Reserved (15:12)	Write with a 3.
TX_DIR_EXP_CONG_QD (11:8)	Exponent of the congested queue depth (this is the number of empty cell buffers below which the device will enter the empty congestion state). Initialize to the proper setting. The congestion state is entered when the number of free buffers drops below this threshold, and the congestion state is exited when the number of free buffers exceeds twice this threshold. NOTE: This field determines the number of empty buffers rather than the number of full buffers.
Not used (7:0)	Write with a 0 to maintain software compatibility with future versions.

7.2.25 TX_DIR_STATE (Internal Structure)

Address: $1F0009_h$ ($7C0024_h$ byte)
Type: Read/Write – Write only during SW_RESET (refer to "SW_RESET" on page 101). Do not update this register.
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:18)	Write with a 0 to maintain software compatibility with future versions.
Reserved (17)	Initialize to 0 at initial setup. Software modifications to this location after setup may cause incorrect operation.
TX_DIR_E_CONG_STATE (16)	Current empty congestion state. Initialize to 0. This bit is read-only.
TX_DIR_E_CUR_QD (15:0)	Current count of empty cells in the transmit direction. Initialize to the number of cells setting. NOTE: The number of buffers is limited to the physical number of buffers minus two.

7.2.26 RX_SENT_CELLS

Address: $1F0010_h$ ($7C0040_h$ byte)
Type: Read-only
Function: Counts the number of cells accepted by the receive direction.
Reset Value:0
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:24)	Write with a 1 to maintain future software compatibility.
RX_SENT_CELLS (23:0)	24-bit rollover counter of the number of cells accepted by the receive direction. This counter is reset to 0 when SW_RESET = 1 or HW_RESET = 1 (refer to "SW_RESET" and "HW_RESET" on page 101).

7.2.27 RX_DROPPED_CELLS

Address: $1F0011_h$ ($7C0044_h$ byte)
Type: Read-only
Function: Counts the number of cells dropped by the receive direction.
Reset Value: 0
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:24)	Write with a 1 to maintain future software compatibility.
RX_DROPPED_CELLS (23:0)	24-bit rollover counter of the number of cells dropped by the receive direction. This counter is reset to 0 when SW_RESET = 1 or HW_RESET = 1 (refer to "SW_RESET" and "HW_RESET" on page 101).

7.2.28 TX_SENT_CELLS

Address: $1F0012_h$ ($7C0048_h$ byte)
Type: Read-only
Function: Counts the number of cells accepted (queued) by the transmit direction.
Reset Value: 0
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:24)	Write with a 1 to maintain future software compatibility.
TX_SENT_CELLS (23:0)	24-bit rollover counter of the number of cells accepted by the transmit direction. This counter is reset to 0 when SW_RESET = 1 or HW_RESET = 1 (refer to "SW_RESET" and "HW_RESET" on page 101).

7.2.29 TX_DROPPED_CELLS

Address: $1F0013_h$ ($7C004C_h$ byte)
Type: Read-only
Function: Counts the number of cells dropped by the transmit direction.
Reset Value: 0
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:24)	Write with a 1 to maintain future software compatibility.
TX_DROPPED_CELLS (23:0)	24-bit rollover counter of the number of cells dropped by the transmit direction. This counter is reset to 0 when SW_RESET = 1 or HW_RESET = 1 (refer to "SW_RESET" and "HW_RESET" on page 101).

7. 2. 30. RX_LOWER16_SCG_CONFIG

Address: $1F0014_h$ ($7C0050_h$)
Read/Write
Function: Congestion configuration register for lower 16 (SC 24-31 and SC 56-63) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:16)	Write with a 0 to maintain future software compatibility.
RX_LWR16_EXP_MAX_QD (15:12)	Maximum QD (exponent or alternate) for service class group 24-31 and 56-63. This limit is enforced if RX_SCG_ENABLE=1.
RX_LWR16_EXP_CONG_QD (11:8)	Congestion QD (exponent or alternate) for service class group 24-31 and 56-63. This limit is enforced if RX_SCG_ENABLE=1.
Not used (7:0)	Write with a 0 to maintain future software compatibility.

7. 2. 31. RX_LOWER16_SCG_STATE (Internal State)

Address: $1F0015_h$ ($7C0054_h$)
Read/Write
Function: Queue state for lower 16 (SC 24-31 and SC 56-63) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:20)	Write with a 0 to maintain future software compatibility.
RX_LWR16_MAXED (19)	This SCG is currently in max congestion. Resets to 0. READ-ONLY.
RX_LWR16_DECR_CONG (18)	This queue depth for this service class group is currently larger than the queue depth defined by one codepoint less than RX_LWR16_CONG_QD. Resets to 0. READ-ONLY.
RX_LWR16_DEF_CONG (17)	This queue depth for this service class group is currently larger than the queue depth defined by RX_LWR16_CONG_QD. Resets to 0. READ-ONLY.
RX_LWR16_CONGESTED (16)	This service class group is currently in congestion as defined by RX_LWR16_CONG_QD, RX_HYSTERESIS_DISABLE, and RX_ALTERNATE_ENCODE. Resets to 0. READ-ONLY.
RESERVED (15)	Write with 0 each time SW_RESET is asserted to maintain future SW compatibility.
RX_LWR16_DATA (14:0)	Current queue depth for this SCG. Initialize to 0000 each time SW_RESET is asserted to maintain software compaitibility. Resets to 0000.

7. 2. 32. RX_LOWER32_SCG_CONFIG

Address: $1F0016_h$ ($7C0058_h$)
Read/Write
Function: Congestion configuration register for lower 32 (SC 16-31 and SC 48-63) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:16)	Write with a 0 to maintain future software compatibility.
RX_LWR32_EXP_MAX_QD (15:12)	Maximum QD (exponent or alternate) for the cells in service class groups 16-31 and 48-63. This limit is enforced if RX_SCG_ENABLE=1.
RX_LWR32_EXP_CONG_QD (11:8)	Congestion QD (exponent or alternate) for the cells in service class groups 16-31 and 48-63. This limit is enforced if RX_SCG_ENABLE=1.
Not used (7:0)	Write with a 0 to maintain future software compatibility.

7. 2. 33. RX_LOWER32_SCG_STATE (Internal State)

Address: $1F0017_h$ ($7C005C_h$)
Read/Write:
Function: Queue state for lower 32 (SC 16-31 and SC 48-63) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:20)	Write with a 0 to maintain future software compatibility.
RX_LWR32_MAXED (19)	This SCG is currently in max congestion. Resets to 0. READ-ONLY.
RX_LWR32_DECR_CONG (18)	This queue depth for this service class group is currently larger than the queue depth defined by one codepoint less than RX_LWR32_CONG_QD. Resets to 0. READ-ONLY.
RX_LWR32_DEF_CONG (17)	This queue depth for this service class group is currently larger than the queue depth defined by RX_LWR32_CONG_QD. Resets to 0. READ-ONLY.
RX_LWR32_CONGESTED (16)	This service class group is currently in congestion as defined by RX_LWR32_CONG_QD, RX_HYSTERESIS_DISABLE, and RX_ALTERNATE_ENCODE. Resets to 0. READ-ONLY.
RESERVED (15)	Write with 0 each time SW_RESET is asserted to maintain future SW compatibility.
RX_LWR32_DATA (14:0)	Current queue depth for this SCG. Initialize to 0000 each time SW_RESET is asserted to maintain software compaitibility. Resets to 0000.

7. 2. 34. RX_LOWER48_SCG_CONFIG

Address: $1F0018_h$ ($7C0060_h$)
Read/Write
Function: Congestion configuration register for lower 48 (SC 8-31 and SC 40-63) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:16)	Write with a 0 to maintain future software compatibility.
RX_LWR48_EXP_MAX_QD (15:12)	Maximum QD (exponent or alternate) for service class group 8-31 and 40-63. This limit is enforced if RX_SCG_ENABLE=1.
RX_LWR48_EXP_CONG_QD (11:8)	Congestion QD (exponent or alternate) for service class group 8-31 and 40-63. This limit is enforced if RX_SCG_ENABLE=1.
Not used (7:0)	Write with a 0 to maintain future software compatibility.

7. 2. 35. RX_LOWER48_SCG_STATE (Internal State)

Address: $1F0019_h$ ($7C0064_h$)
Read/Write
Function: Queue state for lower 48 (SC 8-31 and SC 40-63) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:20)	Write with a 0 to maintain future software compatibility.
RX_LWR48_MAXED (19)	This SCG is currently in max congestion. Resets to 0. READ-ONLY.
RX_LWR48_DECR_CONG (18)	This queue depth for this service class group is currently larger than the queue depth defined by one codepoint less than RX_LWR48_CONG_QD. Resets to 0. READ-ONLY.
RX_LWR48_DEF_CONG (17)	This queue depth for this service class group is currently larger than the queue depth defined by RX_LWR48_CONG_QD. Resets to 0. READ-ONLY.
RX_LWR48_CONGESTED (16)	This service class group is currently in congestion as defined by RX_LWR48_CONG_QD, RX_HYSTERESIS_DISABLE, and RX_ALTERNATE_ENCODE. Resets to 0. READ-ONLY.
RESERVED (15)	Write with 0 each time SW_RESET is asserted to maintain future SW compatibility.
RX_LWR48_DATA (14:0)	Current queue depth for this SCG. Initialize to 0000 each time SW_RESET is asserted to maintain software compaitibility. Resets to 0000.

7. 2. 36. TX_LOWER4_SCG_CONFIG

Address: $1F001A_h$ ($7C0068_h$)
Read/Write
Function: Congestion configuration register for lower 4 (SC 6-7 and 14-15) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:16)	Write with a 0 to maintain future software compatibility.
TX_LWR4_EXP_MAX_QD (15:12)	Maximum QD (exponent or alternate) for service class group 6-7 and 14-15. This limit is enforced if TX_SCG_ENABLE=1.
TX_LWR4_EXP_CONG_QD (11:8)	Congestion QD (exponent or alternate) for service class group 6-7 and 14-15. This limit is enforced if TX_SCG_ENABLE=1.
Not used (7:0)	Write with a 0 to maintain future software compatibility.

7. 2. 37. TX_LOWER4_SCG_STATE (Internal State)

Address: $1F001B_h$ ($7C006C_h$)
Read/Write
Function: Queue state for lower 4 (SC 6-7 and 14-15) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:20)	Write with a 0 to maintain future software compatibility.
TX_LWR4_MAXED (19)	This SCG is currently in max congestion. Resets to 0. READ-ONLY.
TX_LWR4_DECR_CONG (18)	This queue depth for this service class group is currently larger than the queue depth defined by one codepoint less than TX_LWR4_CONG_QD. Resets to 0. READ-ONLY.
TX_LWR4_DEF_CONG (17)	This queue depth for this service class group is currently larger than the queue depth defined by TX_LWR4_CONG_QD. Resets to 0. READ-ONLY.
TX_LWR4_CONGESTED (16)	This service class group is currently in congestion as defined by TX_LWR4_CONG_QD, TX_HYSTERESIS_DISABLE, and TX_ALTERNATE_ENCODE. Resets to 0. READ-ONLY.
RESERVED (15)	Write with 0 each time SW_RESET is asserted to maintain future SW compatibility.
TX_LWR4_DATA (14:0)	Current queue depth for this SCG. Initialize to 0000 each time SW_RESET is asserted to maintain software compaitibility. Resets to 0000.

7. 2. 38. TX_LOWER8_SCG_CONFIG

Address: $1F001C_h$ ($7C0070_h$)
Read/Write
Function: Congestion configuration register for lower 8 (SC 4-7 and 12-15) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:16)	Write with a 0 to maintain future software compatibility.
TX_LWR8_EXP_MAX_QD (15:12)	Maximum QD (exponent or alternate) for service class group 4-7 and 12-15. This limit is enforced if TX_SCG_ENABLE=1.
TX_LWR8_EXP_CONG_QD (11:8)	Congestion QD (exponent or alternate) for service class group 4-7 and 12-15. This limit is enforced if TX_SCG_ENABLE=1.
Not used (7:0)	Write with a 0 to maintain future software compatibility.

7. 2. 39. TX_LOWER8_SCG_STATE (Internal State)

Address: $1F001D_h$ ($7C0074_h$)
Read/Write
Function: Queue state for lower 8 (SC 6-7 and 14-15) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:20)	Write with a 0 to maintain future software compatibility.
TX_LWR8_MAXED (19)	This SCG is currently in max congestion. Resets to 0. READ-ONLY.
TX_LWR8_DECR_CONG (18)	This queue depth for this service class group is currently larger than the queue depth defined by one codepoint less than TX_LWR8_CONG_QD. Resets to 0. READ-ONLY.
TX_LWR8_DEF_CONG (17)	This queue depth for this service class group is currently larger than the queue depth defined by TX_LWR8_CONG_QD. Resets to 0. READ-ONLY.
TX_LWR8_CONGESTED (16)	This service class group is currently in congestion as defined by TX_LWR8_CONG_QD, TX_HYSTERESIS_DISABLE, and TX_ALTERNATE_ENCODE. Resets to 0. READ-ONLY.
RESERVED (15)	Write with 0 each time SW_RESET is asserted to maintain future SW compatibility.
TX_LWR8_DATA (14:0)	Current queue depth for this SCG. Initialize to 0000 each time SW_RESET is asserted to maintain software compaitibility. Resets to 0000.

7. 2. 40. TX_LOWER12_SCG_CONFIG

Address: $1F001E_h$ ($7C0078_h$)
Read/Write
Function: Congestion configuration register for lower 12 (SC 2-7 and 10-15) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:16)	Write with a 0 to maintain future software compatibility.
TX_LWR8_EXP_MAX_QD (15:12)	Maximum QD (exponent or alternate) for service class group 2-7 and 10-15. This limit is enforced if TX_SCG_ENABLE=1.
TX_LWR8_EXP_CONG_QD (11:8)	Congestion QD (exponent or alternate) for service class group 2-7 and 10-15. This limit is enforced if TX_SCG_ENABLE=1.
Not used (7:0)	Write with a 0 to maintain future software compatibility.

7. 2. 41. TX_LOWER12_SCG_STATE (Internal State)

Address: $1F001F_h$ ($7C007C_h$)
Read/Write
Function: Queue state for lower 12 (SC 2-7 and 10-15) Service Class Group
Reset Value:0
Format:

Field (Bits)	Description
Not used (31:20)	Write with a 0 to maintain future software compatibility.
TX_LWR12_MAXED (19)	This SCG is currently in max congestion. Resets to 0. READ-ONLY.
TX_LWR12_DECR_CONG (18)	This queue depth for this service class group is currently larger than the queue depth defined by one codepoint less than TX_LWR12_CONG_QD. Resets to 0. READ-ONLY.
TX_LWR12_DEF_CONG (17)	This queue depth for this service class group is currently larger than the queue depth defined by TX_LWR12_CONG_QD. Resets to 0. READ-ONLY.
TX_LWR12_CONGESTED (16)	This service class group is currently in congestion as defined by TX_LWR12_CONG_QD, TX_HYSTERESIS_DISABLE, and TX_ALTERNATE_ENCODE. Resets to 0. READ-ONLY.
RESERVED (15)	Write with 0 each time SW_RESET is asserted to maintain future SW compatibility.
TX_LWR12_DATA (14:0)	Current queue depth for this SCG. Initialize to 0000 each time SW_RESET is asserted to maintain software compaitibility. Resets to 0000.

8 INTERNAL RAM MEMORY MAP

8.1 Internal RAM Summary

The internal RAM contains the following RAM regions:

- Transmit Service Class RAM (TSC_RAM) 2560 × 16
- Receive Service Class RAM (RSC_RAM) 388 × 16
- Virtual Output Control RAM (VO_RAM) 256 × 16
- Receive Switch Fabric Control RAM (RSF_CONTROL_RAM) 40 × 32
- Transmit Switch Fabric Control RAM (TSF_CONTROL_RAM) 12 × 28
- Receive UTOPIA RAM (RU_RAM) 64 × 32 (4 cell buffers)
- Transmit UTOPIA RAM (TU_RAM) 48 × 32 (3 cell buffers)
- Receive Switch Element RAM (RS_RAM) 128 × 32 (8 cell buffers)
- Transmit Switch Element RAM (TS_RAM) 192 × 32 (12 cell buffers)
- Queue Engine Registers.

Table 25. Internal RAM Summary

Byte Address	Long Address	Name	Description
400000-4027FC$_h$	100000-1009FF$_h$	Transmit Service Class RAM (TSC_RAM)	Contains the transmit SC queue control block.
440000-44060C$_h$	110000-110183$_h$	Receive Service Class RAM (RSC_RAM)	Contains the receive SC RAM.
480000-4803FC$_h$	120000-1200FF$_h$	Virtual Output Control RAM (VO_RAM)	Contains the control tables for the VOs and the direction control.
4C0000-4C00A0$_h$	130000-130028$_h$	Receive Switch Fabric Control RAM (RSF_CONTROL_RAM)	Contains the control table for cell buffers in transmission at the receive switch fabric interfaces.
500000-50002C$_h$	140000-14000B$_h$	Transmit Switch Fabric Control RAM (TSF_CONTROL_RAM)	Contains the control table for cell buffers in transmission at the transmit switch fabric interfaces.
540000-5400FC$_h$	150000-15003F$_h$	Test Access to the Receive UTOPIA RAM (RU_RAM)	Contains the cell buffers for cells headed from the UTOPIA interface.
580000-5800C0$_h$	160000-160030$_h$	Test Access to the Transmit UTOPIA RAM (TU_RAM)	Contains the cell buffers for cells headed toward the UTOPIA interface.
5C0000-5C01FC$_h$	170000-17007F$_h$	Test Access to the Receive Switch Element RAM (RS_RAM)	Contains the cell buffers for cells headed toward the switch fabric.
600000-6002FC$_h$	180000-1800BF$_h$	Test Access to the Transmit Switch Element RAM (TS_RAM)	Contains the cell buffers for cells coming from the switch fabric.
7C0000-7C004C$_h$	1F0000-1F0013$_h$	Queue Engine Registers	For information, refer to sections “QUEUE ENGINE CONDITION_PRES_BITS” on page 132, “RX_DIR_STATE (Internal Structure)” on page 140, “TX_DIR_CONFIG” on page 141, and “TX_DIR_STATE (Internal Structure)” on page 141.

NOTES:

- All bits marked “Reserved” must be initialized to 0 (unless otherwise indicated) at initial setup. Software modifications to these locations after setup may cause incorrect operation.
- All read/write bits marked “Not used” must be written with the value 0 to maintain software compatibility with future versions.
- Most read-only bits marked “Not used” are driven with a 0; however, some may have an unpredictable value on reads. Therefore, all read-only bits marked “Not used” should be masked by the software on reads to maintain compatibility with future versions.
- RAM is not present in bit locations marked “Not present”.

8.2 Transmit Service Class RAM (TSC_RAM) Summary

The TSC_RAM contains the Transmit Service Class Queue (TX SCQ) control block. Figure 67 shows a memory map of the TSC_RAM.

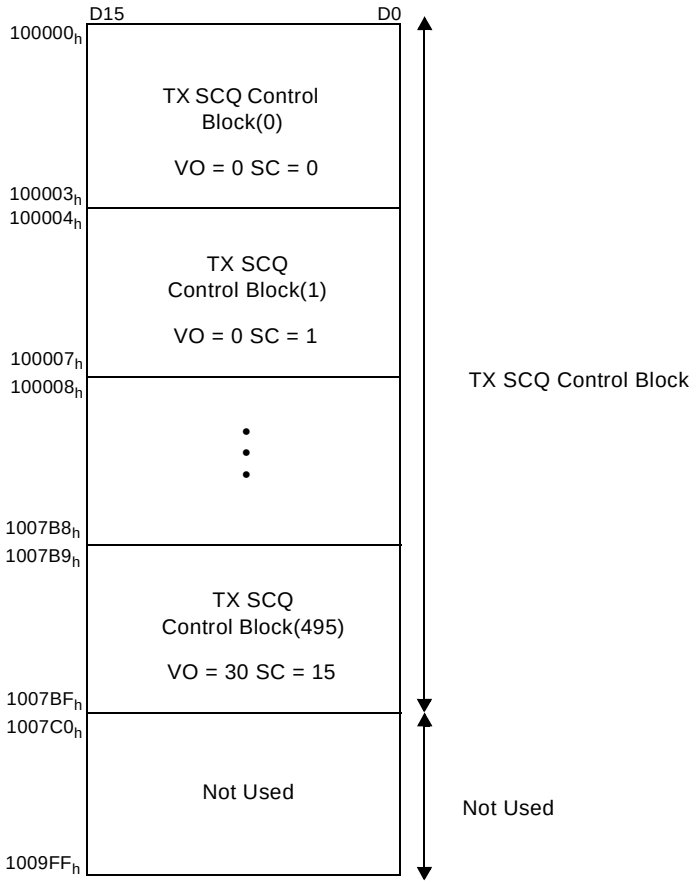

Figure 67. Transmit Service Class RAM (TSC_RAM) Memory Map

8.2.1 Transmit Service Class Queue (TX SCQ) Control Block

Base address: 100000_h (400000_h byte)
Index: 4_h
Number of entries: 496 (1984 words)
Type: Read/Write

Long Address = $100000_h + 40_h \times VO + 4_h \times service\_class + offset$
Byte Address = $400000_h + 100_h \times VO + 10_h \times service\_class + offset$

Table 26. Transmit Service Class Queue (TX SCQ) Control Block Summary

Byte Offset	Long Offset	Name	Read or Write	Description
0_h	0_h	TX_SCQ_CONFIG	R/W	The exponents of the maximum and congested queue depths for this SC.
4_h	1_h	TX_SCQ_STATE	R/W (init only)	The current queue depth and congestion state for this SC.
8_h	2_h	TX_SCQ_HEAD/ TX_SCQ_OUT_FIFO_HEAD	R/W (init only)	For non-multicast queues, this is the head of the linked list. For multicast queues, this is the FIFO read pointer.
C_h	3_h	TX_SCQ_TAIL/ TX_SCQ_OUT_FIFO_TAIL	R/W (init only)	For non-multicast queues, this is the tail of the linked list. For multicast queues, this is the FIFO write pointer.

8.2.1.1 TX_SCQ_CONFIG

Offset: 0_h (0_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_SCQ_EXP_MAX_QD (15:12)	Exponent of the maximum per-SC queue depth. Initialize to the proper setting. Limits the TX_SCQ_CUR_QD to: $TX\_SCQ\_CUR\_QD \leq 1 + 2^{TX\_SCQ\_EXP\_MAX\_QD}$ A value of 0_h causes all cells for this SCQ to be dropped. A value of F_h limits this SC to 31744 cells.
TX_SCQ_EXP_CONG_QD (11:8)	Exponent of the congested per-SCQ queue depth. Initialize to the proper setting. The congestion state is entered when the threshold is exceeded, and the congestion state is exited when the queue length drops below 50 percent of this threshold. Enables congestion management when: $TX\_SCQ\_CUR\_QD > 1 + 2^{TX\_SCQ\_EXP\_CONG\_QD}$ A value of 0_h causes this SCQ to be in congestion at all times. A value of 1_h may not be used. A value of F_h causes this SCQ to enter congestion at a depth of 31744 cells.

(Continued)

Field (Bits)	Description
Not used (7:3)	Write with a 0 to maintain software compatibility with future versions.
TX_SCQ_EXP_WEIGHT (2:1)	Defines the weight of the this SCQ in the queue service decision algorithm. The weight is equal to: $2^{(TX\_SCQEXP\_WEIGHT \times 2)}$ This allows weights of 1, 4, 16, and 64.
Not used (0)	Write with a 0 to maintain software compatibility with future versions.

8.2.1.2 TX_SCQ_STATE (Internal Structure)

Offset: 1_h (4_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_SCQ_STATE (15)	Current SCQ congestion state. Initialize to 0_h.
TX_SCQ_CUR_QD (14:0)	Current SCQ service class unicast queue depth. Initialize to 0000_h. This count does not include cells that are stored in the per-channel resequencing pointer buffers.

8.2.1.3 TX_SCQ_HEAD/TX_SCQ_OUT_FIFO_HEAD (Internal Structure)

Offset: 2_h (8_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following tables.

For Unicast SCs

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_SCQ_HEAD (15:0)	The head of the queue for this SCQ.

For Multicast SCs

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
Not used (15:5)	Write with a 0 to maintain software compatibility with future versions.
TX_SCQ_OUT_FIFO_HEAD (4:0)	The head of the output multicast FIFO (that is, the READ pointer).

8.2.1.4 TX_SCQ_TAIL/TX_SCQ_OUT_FIFO_TAIL (Internal Structure)

Offset: 3_h (C_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following tables.

For Unicast SCs

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_SCQ_TAIL (15:0)	The tail of the queue for this SC.

For Multicast SCs

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
Not used (15:5)	Write with a 0 to maintain software compatibility with future versions.
TX_SCQ_OUT_FIFO_TAIL (4:0)	The tail of the output multicast FIFO (that is, the WRITE pointer).

8.3 Receive Service Class RAM (RSC_RAM) Summary

The RSC_RAM contains the Receive Service Class (RX SC) Control Block. Figure 68 shows a memory map of the RSC_RAM.

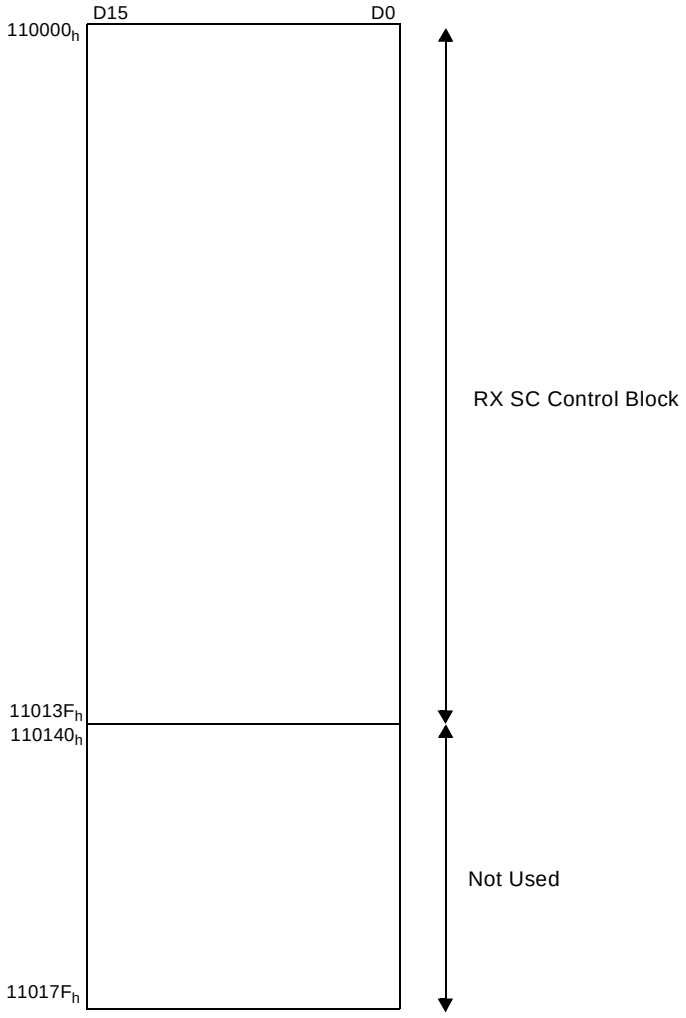

Figure 68. Receive Service Class RAM (RSC_RAM) Memory Map

8.3.1 Receive Service Class (RX SC) Control Block

Base address: 110000_h (440000_h byte)
Index: 1_h
Number of entries: 64 (320 words)
Type: Read/Write
Address =110000_h + *offset* + *service_class*

Table 27. Receive Service Class (RX SC) Control Block Summary

Byte Offset	Long Offset	Name	Read or Write	Description
$0\text{-}FC_h$	$0\text{-}3F_h$	RX_SC_CONFIG	R/W	The exponents of the maximum and congested queue depths for this SC.
$100\text{-}1FC_h$	$40\text{-}7F_h$	RX_SC_STATE	R/W (init only)	The current queue depth for this SC.
$200\text{-}2FC_h$	$80\text{-}BF_h$	RX_SC_CUR_CHAN	R/W (init only)	The current channel being serviced for this SC.
$300\text{-}3FC_h$	$C0\text{-}FF_h$	RX_SC_PREV_CHAN	R/W (init only)	The previous channel serviced for this SC.
$400\text{-}4FC_h$	$100\text{-}13F_h$	RX_SC_CH_COUNT	R/W (init only)	The number of channels with cells currently queued.

8.3.1.1 RX_SC_CONFIG

Offset: 0_h (0_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
RX_SC_EXP_MAX_QD (15:12)	Exponent of the maximum per-SCQ depth. Initialize to the proper setting. Limits the RX_SC_CUR_QD to: $RX\_SC\_CUR\_QD \leq 1 + 2^{RX\_SC\_EXP\_MAX\_QD}$ A value of 0_h causes all cells for this SC to be dropped. A value of F_h limits this SC to 31744 cells.
RX_SC_EXP_CONG_QD (11:8)	Exponent of the congested per-SCQ depth. Initialize to the proper setting. The congestion state is entered when the threshold is exceeded, and the congestion state is exited when the queue length drops below 50 percent of this threshold. Enables congestion management when: $RX\_SC\_CUR\_QD > 1 + 2^{RX\_SC\_EXP\_CONG\_QD}$ A value of 0_h causes this SC to be in congestion at all times. A value of 1_h may not be used. A value of F_h causes this SC to enter congestion at a depth of 31744 cells.
Not used (7:3)	Write with a 0 to maintain software compatibility with future versions.
RX_SC_EXP_WEIGHT (2:1)	Defines the weight of the SC in the queue service decision algorithm. $weight = 2^{(RX\_SC\_EXP\_WEIGHT \times 2)}$ This allows weights of 1, 4, 16, or 64.
Not used (0)	Write with a 0 to maintain software compatibility with future versions.

8.3.1.2 RX_SC_STATE (Internal Structure)

Offset: 40_h (100_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
RX_SC_STATE (15)	Current per-congestion state. Initialize to 0_h.
RX_SC_CUR_QD (14:0)	Current per-SCQ depth. Initialize to 0. This count includes both unicast and multicast cells, as well as those cells in the resequence buffers.

8.3.1.3 RX_SC_CUR_CHAN (Internal Structure)

Offset: 80_h (200_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
Not used (15:14)	Write with a 0 to maintain software compatibility with future versions.
RX_SC_CUR_CHAN (13:0)	Points to the channel that is currently being serviced or that was just serviced by the round-robin service. Initialize to 0000_h.

8.3.1.4 RX_SC_PREV_CHAN (Internal Structure)

Offset: $C0_h$ (300_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
Not used (15:14)	Write with a 0 to maintain software compatibility with future versions.
RX_SC_PREV_CHAN (13:0)	Points to the channel that was just served before the current channel in the round-robin service. Initialize to 0000_h.

8.3.1.5 RX_SC_CH_COUNT (Internal Structure)

Offset:100_h (400_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
Not used (15:14)	Write with a 0 to maintain software compatibility with future versions.
RX_SC_CH_COUNT (13:0)	The number of channels with cells currently queued. Initialize to 0000_h. • Increment when a new channel is added (that is, when a cell received from the receive UTOPIA interface is the first cell of the channel). • Decrement when the last cell of a channel has been successfully transmitted to the switch fabric.

8.4 Virtual Output Control RAM (VO_RAM) Summary

The Virtual Output Control RAM (VO_RAM) contains the following regions:

- Transmit VO Control Block
- Transmit SC Control Block
- Transmit Multicast SC Control Block

Figure 69 shows the VO_RAM memory map.

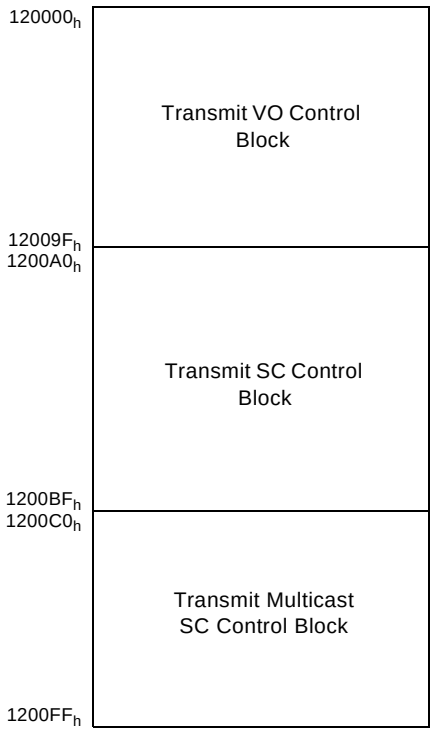

Figure 69. Virtual Output Control RAM (VO_RAM) Memory Map

8.4.1 Transmit VO Control Block

Address base: 120000_h (480000_h byte)
Index: 1_h
Number of entries: $A0_h$
Type: Read/Write

Table 28. Transmit VO Control Block Summary

Byte Offset	Long Offset	Name	Read or Write	Description
$0\text{-}78_h$	$0\text{-}1E_h$	TX_VO_CONFIG	R/W	VO configuration.
$7C_h$	$1F_h$	Not used	R/W	Write with a 0 to maintain software compatibility with future versions.
$80\text{-}F8_h$	$20\text{-}3E_h$	TX_VO_STATE	R/W	VO state.
FC_h	$3F_h$	Not used	R/W	Write with a 0 to maintain software compatibility with future versions.
$100\text{-}178_h$	$40\text{-}5E_h$	TX_VO_CELLS_PRES	R/W	Bit field of the SCs with cells present for each VO.
$17C_h$	$5F_h$	Not used	R/W	Write with a 0 to maintain software compatibility with future versions.
$180\text{-}1F8_h$	$60\text{-}7E_h$	TX_VO_SERVICE_STATE_0	R/W	State word 0 of the service algorithm for this VO.
$1FC_h$	$7F_h$	Not used	R/W	Write with a 0 to maintain software compatibility with future versions.
$200\text{-}278_h$	$80\text{-}9E_h$	TX_VO_SERVICE_STATE_1	R/W	State word 1 of the service algorithm for this VO.
$27C_h$	$9F_h$	Not used	R/W	Write with a 0 to maintain software compatibility with future versions.

8.4.1.1 TX_VO_CONFIG

Offset: 0_h (0_h byte)
Type: Read/Write
Index: 1_h
Number of entries: 31
Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_VO_EXP_MAX_QD (15:12)	Exponent of the maximum per-VO queue depth. Initialize to the proper setting. Limits the TX_VO_CUR_QD to: $TX\_VO\_CUR\_QD \leq 1 + 2^{TX\_VO\_EXP\_MAX\_QD}$ A value of 0_h causes all cells for this VO to be dropped. A value of F_h limits this SC to 31744 cells.

(Continued)

Field (Bits)	Description
TX_VO_EXP_CONG_QD (11:8)	Exponent of the congested VO queue depth. Initialize to the proper setting. The congestion state is entered when the threshold is exceeded, and the congestion state is exited when the queue length drops below 50 percent of this threshold. Enables congestion management when: $TX\_VO\_CUR\_QD > 1 + 2^{TX\_VO\_EXP\_CONG\_QD}$ A value of 0_h causes this VO to be in congestion at all times. A value of 1_h may not be used. A value of F_h causes this VO to enter congestion at a depth of 31744 cells.
Not used (7:0)	Write with a 0 to maintain software compatibility with future versions.

8.4.1.2 TX_VO_STATE (Internal Structure)

Offset: 20_h (80_h byte)

Type: Read/Write – Do not write after SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Index: 1_h

Number of entries: 31

Format: Refer to the following table.

Field (Bits)	Description
TX_VO_CONG_ST (15)	Congestion state bit for the same SC congestion determination. Initialize to 0. This bit is read-only while cells are flowing.
TX_VO_CUR_QD (14:0)	Current VO queue depth. This is the count of the unicast cells queued for this VO. Initialize to 0000_h. This field is maintained by a state machine and is read-only after initialization.

8.4.1.3 TX_VO_CELLS_PRES (Internal Structure)

Offset: 40_h (100_h byte)

Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Index: 1_h

Number of entries: 31

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_VO_CELLS_PRES(15) (15)	Cells are present in the transmit SCQ 15 for VO = index. Initialize to 0.
TX_VO_CELLS_PRES(14) (14)	Cells are present in transmit SCQ 14 for VO = index. Initialize to 0.

(Continued)

Field (Bits)	Description
• • •	• • •
TX_VO_CELLS_PRES(0) (0)	Cells are present in the tranmsit SCQ 0 for VO = index. Initialize to 0.

8.4.1.4 TX_VO_SERVICE_STATE_0 (Internal Structure)

Offset: 60_h (180_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_VO_GP_PTR (15:12)	Pointer to 1 of the 12 GP SC queues. Initialize to 0000.
TX_VO_GP_WEIGHT_CNT (11:6)	Current weight count for the GP SC being serviced. Initialize to 000000.
TX_VO_SC2_WEIGHT_CNT (5:0)	Current weight count for the strict priority 2 class being serviced. Initialize to 000000.

8.4.1.5 TX_VO_SERVICE_STATE_1 (Internal Structure)

Offset:80_h (200_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_VO_SC1_WEIGHT_CNT (15:10)	Current weight count for the strict priority 1 class being serviced. Initialize to 000000.
TX_VO_SC2_PTR (9)	Points to one of the two strict priority 2 SCQs. Initialize to 0.
TX_VO_SC1_PTR (8)	Points to one of the two strict priority 1 SCQs. Initialize to 0.
Not used (7)	Write with a 0 to maintain software compatibility with future versions.
TX_VO_TIME_SLOT_PTR (6:0)	Points to one of the 128 timeslot SCs in the TX_SERVICE_TABLE (refer to "TX_SERVICE_TABLE" on page 181). Initialize to 0000000.

8.4.2 Transmit SC Control Block

Base address: $1200A0_h$ (480280_h byte)
Index: 2_h
Number of entries: 16 (32 words)
Type: Read/Write

Long address = $1200A0_h + 2_h \times service\_class + offset$
Byte address = $480280_h + 8_h \times service\_class + offset$

Table 29. Transmit SC Control Block Summary

Byte Offset	Long Offset	Name	Read or Write	Description
0_h	0_h	TX_SC_CONFIG	R/W	SC configuration word. This word defines the configuration of a set of SCs that stretch across all the transmit VOs. This field is not present in the receive direction. Initialize to the proper value. This word may be modified during cell flow.
4_h	1_h	TX_SC_STATE	R/W	SC state. This word controls the congestion state of the set of same SCs. Initialize to 0.

8.4.2.1 TX_SC_CONFIG

Offset: 0_h (0_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_EXP_MAX_SC_QD (15:12)	Exponent of the maximum per-SC queue depth. Initialize to the proper setting. Limits the TX_SC_CUR_QD to: $TX\_SC\_CUR\_QD \leq 1 + 2^{TX\_SC\_EXP\_MAX\_QD}$ A value of 0_h causes all cells for this SC to be dropped. A value of F_h limits this SC to 31744 cells.
TX_EXP_CONG_SC_QD (11:8)	Exponent of the congested per-SC queue depth. Initialize to the proper setting. The congestion state is entered when the threshold is exceeded, and the congestion state is exited when the queue length drops below 50 percent of this threshold. Enables congestion management when: $TX\_SC\_CUR\_QD > 1 + 2^{TX\_SC\_EXP\_CONG\_QD}$ A value of 0_h causes this SC to be in congestion at all times. A value of 1_h may not be used. A value of F_h causes this SC to enter congestion at a depth of 31744 cells.
Not used (7:0)	Write with a 0 to maintain software compatibility with future versions.

8.4.2.2 TX_SC_STATE (Internal Structure)

Offset: 1_h (4_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_SC_CONG_ST (15)	Congestion state bit for the SC congestion determination. Initialize to 0. This bit is read-only while cells are flowing.
TX_SC_CUR_QD (14:0)	Current SCQ depth. This is the count of the cells queued in all SC N across the VOs. Initialize to 0000_h. This field is maintained by a state machine and is read-only after initialization.

8.4.3 Transmit Multicast SC Control Block Summary

Base address: $1200C0_h$ (480300_h byte)

Index: 4_h

Number of entries: 8 (64 words)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Long address = $200C0_h + 8_h \times (service\_class\ mod\ 8) + offset$

Byte address = $480300_h + 20_h \times (service\_class\ mod\ 8) + offset$

Table 30. Transmit Multicast SC Control Block Summary

Byte Offset	Long Offset	Name	Read or Write	Description
0_h	0_h	TX_SC_MC_IN_FIFO_HEAD	R/W	Head of the multicast input FIFO for this SC.
4_h	1_h	TX_SC_MC_IN_FIFO_TAIL	R/W	Tail of the multicast input FIFO for this SC.
8_h	2_h	TX_SC_MC_BOTTLE	R/W	Bottlenecked VO for multicast in this SC.
C_h	3_h	TX_SC_MC_NEXT_HEADER_PTR	R/W	The next multicast linked list entry (refer to section "MC_LIST" on page 195) to be transferred into an SCQ output FIFO.
10-$1C_h$	4-7_h	Reserved	R/W	Initialize to 0 at initial setup. Software modifications to this location after setup may cause incorrect operation.

8.4.3.1 TX_SC_MC_IN_FIFO_HEAD (Internal Structure)

Offset: 0_h (0_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
Not used (15:12)	Write with a 0 to maintain software compatibility with future versions.
TX_SC_MC_IN_FIFO_HEAD (11:0)	Head of the multicast input FIFO for this SC. Initialize to 0.

8.4.3.2 TX_SC_MC_IN_FIFO_TAIL (Internal Structure)

Offset: 1_h (4_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
Not used (15:12)	Write with a 0 to maintain software compatibility with future versions.
TX_SC_MC_IN_FIFO_TAIL (11:0)	Tail of the multicast IN_FIFO for this SC. Initialize to 0.

8.4.3.3 TX_SC_MC_BOTTLE (Internal Structure)

Offset: 2_h (8_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
Not used (15:5)	Write with a 0 to maintain software compatibility with future versions.
TX_SC_MC_BOTTLE (4:0)	Bottlenecked SCQ for multicast in this SC. Initialize to 0.

8.4.3.4 TX_SC_MC_NEXT_HEADER_PTR (Internal Structure)

Offset: 3_h (C_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_SC_MC_NEXT_HEADER_PTR (15:0)	The next header of the multicast linked list entry (refer to section "MC_LIST" on page 195) to be transferred out to an SCQ output FIFO. Initialize to 0.

8.5 Receive Switch Fabric Control RAM (RSF_CONTROL) Summary

NOTE: All registers in this RAM contain state information. No configuration or initialization is necessary. The only foreseen uses for this information are in system- and board-level diagnostics and/or debug.

Base address:130000_h ($4C0000_h$ byte)
Index: 5_h
Number of entries: 8 (40 words)
Type: Read/Write
Long address = 130000_h + *long_offset* + *RSF_buffer_number*
Byte address = $4C0000_h$ + *byte_offset* + *RSF_buffer_number*

Table 31. Receive Switch Fabric Control RAM (RSF_CONTROL) Summary

Byte Offset	Long Offset	Name	Read or Write	Description
0_h	0_h	RX_RSF_CONFIG	R/W	The channel configuration of the cell in the Receive Switch Fabric (RSF) transmission.
20_h	8_h	RX_RSF_TAG	R/W	The tag word for the cell intended for RSF transmission. NOTE: This word also contains OUTCHAN(12:0).
40_h	10_h	RX_RSF_NEW_VPI_VCI	R/W	The new VPI and VCI for the cell intended for RSF transmission. NOTE: This word also contains OUTCHAN(15:13).
60_h	18_h	RX_RSF_SN_CHAN	R/W	The sequence number and channel number for the cell intended for RSF transmission
80_h	20_h	RX_RSF_ER_CELL_PTR	R/W	The Explicit Rate (ER) and cell buffer pointer for the cell intended for RSF transmission.

8.5.1 RX_RSF_CONFIG (Internal Structure)

Offset: 0_h (0_h byte)
Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.
Format: Refer to the following table.

Field (Bits)	Description
RX_RSF_CONFIG (31:0)	The snapshot of the RX_CH_CONFIG of the channel where this cell originated. For information on RX_CH_CONFIG, refer to section "RX_CH_TAG" on page 185.

8.5.2 RX_RSF_TAG (Internal Structure)

Offset: 8_h (20_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
RX_RSF_TAG (31:0)	The contents of the RX_CH_TAG of the channel from which this cell originated. Bits (7:4) contain OUT-CHAN(15:12). For information on RX_CH_TAG, refer to section "RX_CH_TAG" on page 185.

8.5.3 RX_RSF_NEW_VPI_VCI (Internal Structure)

Offset: 10_h (40_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
VPMODE_CH (31)	Indicates the entry is valid VPC.
VC_TRANS (30)	Indicates the VC is to be translated. Initialize to the proper setting.
Not used (29:28)	Write with a 0 to maintain software compatibility with future versions.
RSF_OUTCHAN_NEW_VPI (27:16)	The OUTCHAN(11:0) or VPI value to be used.
NEW_VCI (15:0)	The VCI value to be used when translation is enabled by VC_TRANS. Initialize to the proper setting.

8.5.4 RX_RSF_SN_CHAN (Internal Structure)

Offset: 18_h (60_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
EFCI (31)	Explicit Forward Congestion Indication bit to OR into the outgoing data cells.
Not used (30:22)	Write with a 0 to maintain software compatibility with future versions.
RX_RSF_SN (21:16)	The SN for the cell pending, or that has completed, RSF transmission.

(Continued)

Field (Bits)	Description
RX_RSF_CHAN (13:0)	The receive channel number from which this cell originated.

8.5.5 RX_RSF_ER_CELL_PTR (Internal Structure)

Offset: 20_h (80_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
RX_RSF_ER (31:16)	The ER value to be sent in the payload of this cell pending RSF transmission if it is an RM cell.
RX_RSF_CELL_PTR (15:0)	The cell pointer address of the cell pending, or that has completed, RSF transmission.

8.6 Transmit Switch Fabric Control RAM (TSF_CONTROL_RAM) Summary

NOTE: All registers in this RAM contain state information. No configuration or initialization is necessary. The only foreseen uses for this information are in system- and board-level diagnostics and/or debug.

8.6.1 TX_TSF_SN_CHAN (Internal Structure)

Base address: 140000_h (500000_h byte)

Index: 1_h

Number of entries: 12

Type: Read-only

Long address = 140000_h + *TSF_buffer_number*

Byte address = $500000_h + 4 \times$ *TSF_buffer_number*

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:28)	RAM is not present in these bit locations.
TX_TSF_CLP_PTI (27:24)	The CLP and PTI for the cell received from the Transmit Switch Fabric (TSF).
Not used (23:22)	Driven with a 0. Mask on reads to maintain software compatibility with future versions.
TX_TSF_SN (21:16)	The sequence number of the cell received from the TSF.
TX_TSF_OUTCHAN (15:0)	The transmit OUTCHAN channel number of the cell received from the TSF.

8.7 Test Access to the Receive UTOPIA RAM (RU_RAM)

Base address: 150000_h (540000_h byte)
Index: 1_h
Number of entries: 64 (four cells)
Type: Read/Write during SW_RESET (refer to "SW_RESET" on page 101).
Format: Refer to the following table

Field (Bits)	Description
RX_UTOPIA_RAM (31:0)	Test access to the RU_RAM.

Receive UTOPIA Cell Buffer Summary (Internal Structure)
Base address: 1500000_h (5400000_h byte)
Index: 10_h
Entry number: 4_h
Type: Read/Write: during SW_RESET (refer to "SW_RESET" on page 101).
Long address = 1500000_h + entry_number $\times$ 10_h

Table 32. Receive UTOPIA Cell Buffers Summary

Offset	Cell			
	31:24	23:16	15:8	7:0
0	VPI(11:4)	VPI(3:0), VCI(15:12)	VCI(11:4)	VCI(3:0), PTI, CLP
1	payload 0	payload 1	payload 2	payload 3
2	payload 4	payload 5	payload 6	payload 7
3	payload 8	payload 9	payload 10	payload 11
4	payload 12	payload 13	payload 14	payload 15
5	payload 16	payload 17	payload 18	payload 19
6	payload 20	payload 21	payload 22	payload 23
7	payload 24	payload 25	payload 26	payload 27
8	payload 28	payload 29	payload 30	payload 31
9	payload 32	payload 33	payload 34	payload 35
10	payload 36	payload 37	payload 38	payload 39
11	payload 40	payload 41	payload 42	payload 43
12	payload 44	payload 45	payload 46	payload 47
13	00_h	000, VI	Not used	Not used
14	Not used	Not used	Not used	Not used
15	Not used	Not used	Not used	Not used

8.8 Test Access to the Transmit UTOPIA RAM (TU_RAM)

Address: 160000_h (580000_h byte)
Index: 1_h
Number of entries: 64(four cells)
Type: Read/Write during SW_RESET (refer to "SW_RESET" on page 101).
Format: Refer to the following table.

Field (Bits)	Description
TX_UTOPIA_RAM (31:0)	Test access to the TU_RAM.

Transmit UTOPIA Cell Buffer Summary (Internal Structure)
Base address: 160000_h (580000_h byte)
Index:4_h
Type: Read/Write during SW_RESET (refer to "SW_RESET" on page 101).
Long address = 160000_h + entry_number $\times$ 10_h

Table 33. Transmit UTOPIA Cell Buffers Summary

Offset	Cell			
	31:24	23:16	15:8	7:0
0	VPI(11:4)	VPI(3:0), VCI(15:12)	VCI(11:8)	VCI(3:0), PTI, CLP
1	payload 0	payload 1	payload 2	payload 3
2	payload 4	payload 5	payload 6	payload 7
3	payload 8	payload 9	payload 10	payload 11
4	payload 12	payload 13	payload 14	payload 15
5	payload 16	payload 17	payload 18	payload 19
6	payload 20	payload 21	payload 22	payload 23
7	payload 24	payload 25	payload 26	payload 27
8	payload 28	payload 29	payload 30	payload 31
9	payload 32	payload 33	payload 34	payload 35
10	payload 36	payload 37	payload 38	payload 39
11	payload 40	payload 41	payload 42	payload 43
12	payload 44	payload 45	payload 46	payload 47
13	Not used	Not used	Not used	Not used
14	Not used	Not used	Not used	Not used
15	Not used	Not used	Not used	Not used

8.9 Test Access to Receive Switch Element RAM (RS_RAM)

Address:170000_h ($5C0000_h$ byte)
Index: 1_h
Number of entries: 128 (eight cells)
Type: Read/Write during SW_RESET (refer to "SW_RESET" on page 101).
Format: Refer to the following table.

Field (Bits)	Description
RX_SWITCH_ELEMENT_RAM (31:0)	Test access to the RS_RAM.

8.10 Test Access to Transmit Swith Element RAM (TS_RAM)

Address: 180000_h (600000_h byte)
Index: 1_h
Number of entries: 192 (12 cells)
Type: Read/Write during SW_RESET (refer to "SW_RESET" on page 101).
Format: Refer to the following table.

Field (Bits)	Description
TX_SWITCH_ELEMENT_RAM (31:0)	Test access to the TS_RAM.

9 EXTERNAL RAM MEMORY MAP

9.1 External RAM Summary

The external RAM contains:

- Address Lookup RAM (AL_RAM)
- Channel RAM (CH_RAM)
- Channel Head/Tail and Statistics RAM (AB_RAM)
- Receive Cell Buffer SDRAM/SGRAM (RX_DRAM)
- Transmit Cell Buffers SDRAM/SGRAM (TX_DRAM)

The following is a memory map of the external RAM, including addresses for the possible SRAM_CONFIG values.

Byte Address	Long Address	Name	Description
800000_h	200000_h	Address Lookup RAM (AL_RAM)	Contains the address lookup tables, linked lists, and multicast pointer FIFOs.
1000000_h	400000_h	Channel RAM (CH_RAM)	Contains the channel tables.
1800000_h	600000_h	AB_RAM	Contains the head and tail pointers for the receive channel queues.
2000000_h	800000_h	Receive Cell Buffer SDRAM/SGRAM (RX_DRAM_REGISTER)	Receive buffer SDRAM/SGRAM.
3000000_h	$C00000_h$	Transmit Cell Buffers SDRAM/SGRAM (TX_DRAM_REGISTER)	Transmit buffer SDRAM/SGRAM.

NOTES:

- All ports marked as "Reserved" must be initialized to 0 at initial setup. Software modifications to these locations after setup may cause incorrect operation.
- All read/write port bits marked "Not used" must be written with the value 0 to maintain software compatibility with future versions.
- All read-only port bits marked "Not used" are driven with a 0 and should be masked off by the software to maintain compatibility with future versions.

9.2 Address Lookup RAM (AL_RAM)

Table 34 summarizes the contents of the AL_RAM.

Table 34. Address Lookup RAM (AL_RAM) Summary

Byte Offset	Long Offset	Name	Read or Write	Description
800000_h	200000_h	VI_VPI_TABLE	R/W	The first lookup is to this table. For VPCs, this lookup determines the channel number directly. For VCCs, this determines the number of VCI bits to use and the start of a block in the channel number table.
If AL_RAM_CONFIG = 0: 808000_h If AL_RAM_CONFIG = 1: 810000_h If AL_RAM_CONFIG = 2: 820000_h If AL_RAM_CONFIG = 3: 880000_h	If AL_RAM_CONFIG = 0: 202000_h If AL_RAM_CONFIG = 1: 204000_h If AL_RAM_CONFIG = 2: 208000_h If AL_RAM_CONFIG = 3: 220000_h	VCI_TABLE	R/W	Lookup table to determine the receive channel number (RX_CHAN_NUM) for VCCs.
Refer to "Multicast Cell Instance Control Block (Internal Structure)" on page 177.	Refer to "Multicast Cell Instance Control Block (Internal Structure)" on page 177.	Multicast Cell Instance Control Block (internal structure)	R/W	Header and cell pointer for multicast input and output FIFO entries.
If AL_RAM_CONFIG = 0: 838000_h If AL_RAM_CONFIG = 1: 870000_h If AL_RAM_CONFIG = 2: $8C0000_h$ If AL_RAM_CONFIG = 3: $9C0000_h$	If AL_RAM_CONFIG = 0: $20E000_h$ If AL_RAM_CONFIG = 1: $21C000_h$ If AL_RAM_CONFIG = 2: 230000_h If AL_RAM_CONFIG = 3: 270000_h	RX_NEXT_CELL (internal structure)	R/W	Index to the next cell in the receive per-channel linked list.
If AL_RAM_CONFIG = 0: 830000_h If AL_RAM_CONFIG = 1: 860000_h If AL_RAM_CONFIG = 2: $8A0000_h$ If AL_RAM_CONFIG = 3: 980000_h	If AL_RAM_CONFIG = 0: $20C000_h$ If AL_RAM_CONFIG = 1: 218000_h If AL_RAM_CONFIG = 2: 228000_h If AL_RAM_CONFIG = 3: 260000_h	Transmit Cell Buffer Control Block (internal structure)	R/W	For unicast cells, this is the index to the next cell in the transmit per-SCQ linked list. For multicast cells, this is the background processor control block.
If AL_RAM_CONFIG = 0: $80C000_h$ If AL_RAM_CONFIG = 1: 820000_h If AL_RAM_CONFIG = 2: 840000_h If AL_RAM_CONFIG = 3: 900000_h	If AL_RAM_CONFIG = 0: 203000_h If AL_RAM_CONFIG = 1: 208000_h If AL_RAM_CONFIG = 2: 210000_h If AL_RAM_CONFIG = 3: 240000_h	Service Order Control Block (internal structure)	R/W	Timeslot-based priority table for receive and transmit SCs.

9.2.1 VI_VPI_TABLE

The first lookup is to this table. For VPCs, this lookup determines the channel number directly. For VCCs, this determines the number of VCI bits to use and the start of a block in the channel number table.

Base address: 200000_h (800000_h byte)
Index: 1_h
Number of entries: Up to 128K
Long address: 200000_h + Shift_left ((VI mod VI_CONF), NUM_VPI) + VP mod NUM_VPI

For NUM_VI, refer to "NUM_VI" on page 104.
VI_CONF = 1 for NUM_VI=0
VI_CONF = 4 for NUM_VI=1
VI_CONF = 32 for NUM_VI=2
For NUM_VPI values, refer to Table 35:

Table 35. Determining the NUM_VPI Value

If AL_RAM_CONFIG =	If NUM_VI =		
	0	1	2
0	NUM_VPI = 12	NUM_VPI = 11	NUM_VPI = 8
1	NUM_VPI = 12	NUM_VPI = 12	NUM_VPI = 9
2	NUM_VPI = 12	NUM_VPI = 12	NUM_VPI = 10
3	NUM_VPI = 12	NUM_VPI = 12	NUM_VPI = 12

Type: Read/Write
Format: Refer to the following tables.

Table 36. VI_VPI_TABLE Entry if VPC_ENTRY = 1

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
VPC_ENTRY (15)	Indicates the entry is a valid VPC. Initialize to the proper setting.
Not used (14)	Write with a 0 to maintain software compatibility with future versions.
RX_CHAN_NUM (13:0)	The receive channel number.

Table 37. VI_VPI_TABLE Entry if VPC_ENTRY = 0

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
VPC_ENTRY (15)	Indicates the entry is a valid VPC. Initialize to the proper setting.

Table 37. VI_VPI_TABLE Entry if VPC_ENTRY = 0 (Continued)

Field (Bits)	Description
VCC_ENTRY (14)	Indicates that at least one valid VCC is referenced in the table pointed to by this entry. Initialize to the proper setting.
VCI_BITS (13:10)	The number of VCI bits to use minus 1. The valid entries are 15 down to 4.
BLOCK_OFFSET (9:0)	The offset ÷ 32 within the channel number block for this VI/VP within the channel number table. For AL_RAM_CONFIG = 3 (refer to "AL_RAM_CONFIG" on page 105), a BLOCK_OFFSET(10) bit is needed. If VCI_BITS = F_h, then BLOCK_OFFSET(10) = 0. Otherwise, BLOCK_OFFSET(10) = incoming cell VPI (0).

9.2.2 VCI_TABLE

The VCI_TABLE determines the channel number for VCCs.

Base address (refer to "AL_RAM_CONFIG" on page 105):

If AL_RAM_CONFIG = 0, then the base address is 202000_h (808000_h byte).
If AL_RAM_CONFIG = 1, then the base address is 204000_h (810000_h byte).
If AL_RAM_CONFIG = 2, then the base address is 208000_h (820000_h byte).
If AL_RAM_CONFIG = 3, then the base address is 220000_h (880000_h byte).

Index: 1_h

Number of entries:

If AL_RAM_CONFIG = 0, then the entry number is 4K.
If AL_RAM_CONFIG = 1, then the entry number is 16K.
If AL_RAM_CONFIG = 2, then the entry number is 32K.
If AL_RAM_CONFIG = 3, then the entry number is 64K.

Type: Read/Write

Address = *base_address + entry_number*

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
Not used (15:14)	Write with a 0 to maintain software compatibility with future versions.
RX_CHAN_NUM (13:0)	The receive channel number.

9.2.3 Multicast Cell Instance Control Block (Internal Structure)

Base address: x_h

For input FIFOS (refer to "AL_RAM_CONFIG" on page 105):
If AL_RAM_CONFIG = 0_h, the base address is 206000_h (818000_h byte).
If AL_RAM_CONFIG = 1_h, the base address is 210000_h (840000_h byte).
If AL_RAM_CONFIG = 2_h, the base address is 220000_h (880000_h byte).
If AL_RAM_CONFIG = 3_h, the base address is 250000_h (940000_h byte).
For output FIFOS:
If AL_RAM_CONFIG = 0_h, the base address is 208000_h (820000_h byte).
If AL_RAM_CONFIG = 1_h, the base address is 214000_h (850000_h byte).
If AL_RAM_CONFIG = 2_h, the base address is 214000_h (850000_h byte).
If AL_RAM_CONFIG = 3_h, the base address is 258000_h (960000_h byte).

Index: 2_h

Entry number:

For input FIFOS (refer to "AL_RAM_CONFIG" on page 105):
If AL_RAM_CONFIG = 0_h, the entry number is 512.
If AL_RAM_CONFIG = 1_h, the entry number is 1K.
If AL_RAM_CONFIG = 2_h, the entry number is 1K.
If AL_RAM_CONFIG = 3_h, the entry number is 2K.
For output FIFOS:
If AL_RAM_CONFIG = 0_h, the entry number is 32.
If AL_RAM_CONFIG = 1_h, the entry number is 32.
If AL_RAM_CONFIG = 2_h, the entry number is 32.
If AL_RAM_CONFIG = 3_h, the entry number is 32.

Type: Read/Write

Byte Offset	Long Offset	Name	Read or Write	Description
0_h	0_h	MC_HEADER_PTR	R/W (init only)	For the input FIFO, this is the OUTCHAN from the cell. For the output FIFO, this is the pointer to the entry in the MC header translation table for this entry.
4_h	1_h	MC_CELL_PTR	R/W (init only)	Pointer to the cell buffer containing the multicast cell to be replicated.

9.2.3.1 MC_HEADER_PTR (Internal Structure)

Offset: 0_h (0_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
MC_HEADER_PTR (15:0)	For the input FIFO, this is the OUTCHAN from the cell. For the output FIFO, this is the pointer to the entry in the MC header translation table for this entry.

9.2.3.2 MC_CELL_PTR (Internal Structure)

Offset: 1_h (4_h byte)

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
MC_CELL_PTR (15:0)	The pointer to the cell buffer containing the multicast cell to be replicated.

9.2.4 RX_NEXT_CELL (Internal Structure)

Base address (refer to "AL_RAM_CONFIG" on page 105):

If AL_RAM_CONFIG = 0_h, the base address is $20E000_h$ (838000_h byte).
If AL_RAM_CONFIG = 1_h, the base address is $21C000_h$ (870000_h byte).
If AL_RAM_CONFIG = 2_h, the base address is 230000_h ($8C0000_h$ byte).
If AL_RAM_CONFIG = 3_h, the base address is 270000_h ($9C0000_h$ byte).

Entry number:

If AL_RAM_CONFIG = 0_h, the entry number is 8K.
If AL_RAM_CONFIG = 1_h, the entry number is 16K.
If AL_RAM_CONFIG = 2_h, the entry number is 64K.
If AL_RAM_CONFIG = 3_h, the entry number is 64K.

Type: Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
RX_NEXT_CELL (15:0)	Index to the next cell in the receive per-channel linked list. Initialize in an incrementing pattern starting with 1_h in the first location.

9.2.5 Transmit Cell Buffer Control Block (Internal Structure)

Base address (refer to "AL_RAM_CONFIG" on page 105):

If AL_RAM_CONFIG = 0_h, the base address is $20C000_h$ (830000_h byte).
If AL_RAM_CONFIG = 1_h, the base address is 218000_h (860000_h byte).
If AL_RAM_CONFIG = 2_h, the base address is 228000_h ($8A0000_h$ byte).
If AL_RAM_CONFIG = 3_h, the base address is 260000_h (980000_h byte).

Index: 1_h

Entry number:

If AL_RAM_CONFIG = 0_h, the entry number is 8K.
If AL_RAM_CONFIG = 1_h, the entry number is 16K.
If AL_RAM_CONFIG = 2_h, the entry number is 32K.
If AL_RAM_CONFIG = 3_h, the entry number is 64K.

Read/Write – Do not write while SW_RESET (refer to "SW_RESET" on page 101) is deasserted.

If the Arriving Cell is Unicast

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
TX_NEXT_CELL (15:0)	Index to the next cell in the transmit per-SCQ linked list. Initialize in an incrementing pattern starting with 1_h in the first location.

If Arriving Cell is Multicast

Field (Bits)	Description
Not Present (31:16)	RAM is not present in these bit locations.
TX_MC_ENQ_PEND (15)	1 The background processor is still queuing cells. 0 The background processor has finished queuing this cell.
Not used (14)	Write with a 0 to maintain future software compatibility.
TX_MC_COUNT (13:0)	Count of times this cell has been enqueued.

9.2.6 Service Order Control Block (Internal Structure)

If AL_RAM_CONFIG = 0_h, then the base address is 203000_h ($80C000_h$ byte).
If AL_RAM_CONFIG = 1_h, then the base address is 208000_h (820000_h byte).
If AL_RAM_CONFIG = 2_h, then the base address is 210000_h (840000_h byte).
If AL_RAM_CONFIG = 3_h, then the base address is 240000_h (900000_h byte).

Index: 1_h
Number of entries: 32
Type: Read/Write

Table 38. Service Order Control Block Summary

Index	Name	Read or Write	Description
$0\text{-}1E_h$	TX_SERVICE_TABLE	R/W	Table of the SC to serve if there are no cells in any of the strict SCs for the transmit direction for VO = Index.
$1F_h$	RX_SERVICE_TABLE	R/W	Table of the SC to serve if there are no cells in any of the strict SCs for the receive direction.

9.2.6.1 TX_SERVICE_TABLE

The TX_SERVICE_TABLE is configured to provide an SCQ with a minimum cell rate. When no strict cell is present, the queue service algorithm examines the SC named and serves a cell from that class if there is one, advancing the pointer to this table. One of these tables exists for each VO. Each table is 127 entries long (the 128^{th} entry is not used).

Type: Read/Write
Long base offset: *base_address* + 128 × VO + entry
Index: 1_h
Number of entries: 128
Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
Not used (15:4)	Write with a 0 to maintain software compatibility with future versions.
FIRST_SC (3:0)	SC to service if there are no strict service cells. To create a null entry, enter one of the strict SCs.

9.2.6.2 RX_SERVICE_TABLE

The RX_SERVICE_TABLE is configured to provide an SCQ with a minimum cell rate. When no strict cell is present, the queue service algorithm examines the SC named and serves a cell from that class if there is one, advancing the pointer to this table.

Type: Read/Write
If RX_LONG_SVC_TBL=0
Long base offset: *base_address* + $F80_h$
Number of entries: 127
If RX_LONG_SVC_TBL=1 (not valid if ALRAM_CONFIG=0)
Long base offset: *base_address* + 1000_h
if RX_4_RR_PTRS = 0
Number of entries: 511
if RX_4_RR_PTRS = 1
Number of entries: 4x511
Index: 1_h
Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	RAM is not present in these bit locations.
Not used (15:6)	Write with a 0 to maintain software compatibility with future versions.
FIRST_SC (5:0)	SC to service if there are no strict service cells. To create a null entry, enter one of the strict SCs.

9.3 Channel RAM (CH_RAM)

Table 39 summarizes the contents of the CH_RAM.

Table 39. Channel RAM (CH_RAM) Summary

Byte Offset	Long Offset	Name	Read or Write	Description
01000000_h	400000_h	Channel Control Block (CCB)	R/W	Receive and transmit per-channel control block.
01000000_h	400000_h	Multicast Control Block	R/W	Control block for transmit multicast linked lists.

9.3.1 Channel Control Block

Base address: 400000_h (01000000_h byte)

Index:10_h

Entry number (refer to "CH_RAM_CONFIG" on page 105)

For CH_RAM_CONFIG = 0_h, the entry number is from 0 to 2047.

For CH_RAM_CONFIG = 1_h, the entry number is from 0 to 4095.

For CH_RAM_CONFIG = 2_h, the entry number is from 0 to 8191.

For CH_RAM_CONFIG = 3_h, the entry number is from 0 to 16385.

Type: Read/Write

Long address:

In the receive direction, the long address = 400000_h + $RX\_CHAN\_NUM \times 10_h$+ *long_offset*

In the transmit direction, the long address = 400000_h + $OUTCHAN \times 10_h$+ *long_offset*.

Table 40. Channel Control Block Summary

Byte Offset	Long Offset	Name	Read or Write	Description
0_h	0_h	RX_CH_CONFIG	R/W	Overall controls for channel.
4_h	1_h	RX_CH_TAG	R/W	Switch fabric tag.
8_h	2_h	RX_CH_NEW_VPI_VCI	R/W	VPI/VCI to which you want to translate.
C_h	3_h	RX_CH_STATE	R/W (init only)	Current channel state.
10_h	4_h	Not used	R/W	Write with a 0 to maintain software compatibility with future versions.
14_h	5_h	RX_CH_NEXT_CH	R/W (init only)	The next channel in the round-robin linked list.
18_h	6_h	Reserved	R/W (init only)	Initialize to 0 at initial setup. Software modifications to this location after setup may cause incorrect operation.
$1C_h$	7_h	Reserved	R/W (init only)	Initialize to 0 at initial setup. Software modifications to this location after setup may cause incorrect operation.
20_h	8_h	RX_CH_SEQ_ST	R/W (init only)	State of the sequencing algorithm.
24_h	9_h	Not used	R/W	Write with a 0 to maintain software compatibility with future versions.
28_h	A_h	TX_CH_CONFIG	R/W	Overall controls for the channel.

Table 40. Channel Control Block Summary (Continued)

Byte Offset	Long Offset	Name	Read or Write	Description
$2C_h$	B_h	TX_NEW_VPI_RESEQ_PTR	R/W	VPI to which you want to translate and the pointer to the resequencing buffer.
30_h	C_h	TX_CH_STATE	R/W (init only)	Current channel state.
34_h	D_h	TX_CH_RESEQ_ST	R/W (init only)	State of the resequencing algorithm.
38_h	E_h	TX_CH_CELLS_SENT	R/W (init only)	Count of transmit cells sent.
$3C_h$	F_h	TX_CH_CELLS_DRPD	R/W (init only)	Count of transmit cells dropped.

9.3.1.1 RX_CH_CONFIG

Offset: 0_h (0_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
RX_CH_EXP_MAX_QD (31:28)	Exponent of the maximum per-SC queue depth. Initialize to the proper setting. Limits RX_CH_CUR_QD (refer to "RX_CH_CUR_QD" on page 187) to: $RX\_CH\_CUR\_QD \leq 1 + 2^{RX\_CH\_EXP\_MAX\_QD}$ A value of 0_h causes all cells for this channel to be dropped. A value of F_h limits this channel to 31744 cells.
RX_CH_EXP_CONG_QD (27:24)	Exponent of the congested channel queue depth. Initialize to the proper setting. The congestion state is entered when the threshold is exceeded, and the congestion state is exited when the queue length drops below 50 percent of this threshold. Enables congestion management when: $RX\_CH\_CUR\_QD > 1 + 2^{RX\_CH\_EXP\_CONG\_QD}$ A value of 0_h causes this channel to be in congestion at all times. A value of 1_h may not be used. A value of F_h causes this channel to enter congestion at a depth of 31744 cells.
RX_CH_ENTRY_PRESENT (23)	Initialize to 0, set other variables, then initialize to the proper setting. 1 Indicates the entry is present and valid. 0 Indicates the entry is disabled. Drop any cells received at this index.
Not Used (22)	Initialize to 0 to maintain future compatibility
RX_CH_SERVICE_CLASS (21:16)	Indicates the SC.
RX_CH_EN_CLP_DROP (15)	Initialize to the proper setting. 1 Enables CLP dropping when either the channel, the SC, or the device is in the congested state. 0 Disables CLP dropping above the congestion threshold.
Not used (14)	Write with a 0 to maintain software compatibility with future versions.

(Continued)

Field (Bits)	Description
RX_CH_EN_AAL5_EPD (13)	Initialize to the proper setting. 1 Enables AAL5 Early Packet Discard (EPD) and Packet Tail Discard (PTD) when either the channel, the SC, or the device is in the congested state. If the maximum threshold is reached, then the tail of the packet is discarded. If CLP causes a cell to be dropped, then the tail is discarded. 0 Disables AAL5 EPD and PTD.
RX_CH_EN_EFCI (12)	Enable EFCI insertion. 0 Disables EFCI insertion. 1 Enable EFCI insertion.
Reserved (11:8)	Initialize to 0 at initial setup. Software modifications to these locations after setup may cause incorrect operation.
RX_CH_SEQ_TYPE (7:6)	Initialize to the proper setting. 11 Reserved. 10 Reserved. 01 Supports sequencing of up to two cells allowing a 155 Mbps VC. 00 No sequencing; allows a 78 Mbps VC. The routing table at the other end of the connection must support the same resequencing mode for unicast connections. For multicast connections, this field may be set to either 00 or 01.
RX_CH_SF_MULTICAST (5)	Insert in the multicast (MC) cell bit in nibble 0 (refer to Table on page 91). If 1, this channel is treated as multicast by the switch fabric. This must be the same as RX_CH_SERVICE_CLASS(5) (refer to "TX_CH_SERVICE_CLASS" on page 189).
RX_CH_SF_SPARE (4)	Insert into the spare (SP) bit in nibble 0 (refer to Table on page 91).
RX_CH_SWITCH_GROUP (3:0)	Switch group to be used in nibble 1 as the cell is sent to the switch fabric. 3_h For RX_CH_SERVICE_CLASS from 00_h to 07_h or 20_h to 27_h 2_h For RX_CH_SERVICE_CLASS from 08_h to $0F_h$ or 28_h to $2F_h$. 1_h For RX_CH_SERVICE_CLASS from 10_h to $1F_h$ or 30_h to $3F_h$. 0_h Reserved. The above settings are highly recommended for unicast service classes (00 to 1F) to insure that high priority cells from the Rx get high priority through the switch fabric (QSE), and similarly for medium and low priority cells. The above settings are REQUIRED for the multicast service classes (20 - 3F) to insure that the multicast backpressure mechanism works properly between the switch fabric and the QRT. The switch fabric gives high, medium or low BP, and the QRT must refrain from sending cells on the corresponding service classes.

9.3.1.2 RX_CH_TAG

Offset: 1_h (4_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
RX_CH_TAG_0_6(3:0) (31:28)	Tag information inserted into SE_D_OUT(3:0) immediately after the SWITCH_GROUP/QUEUE field (refer to Table starting on page 91). It is also inserted after TAG_5(3:0). Initialize to the proper setting.

Field (Bits)	Description
RX_CH_TAG_1_7(3:0) (27:24)	Tag information inserted into SE_D_OUT(3:0) immediately after TAG_0(3:0). It is also inserted after the TAG_0_6 is inserted the second time. Initialize to the proper setting.
RX_CH_TAG_2(3:0) (23:20)	Tag information inserted into SE_D_OUT(3:0) immediately after TAG_1(3:0). Initialize to the proper setting.
RX_CH_TAG_3(3:0) (19:16)	Tag information inserted into SE_D_OUT(3:0) immediately after TAG_2(3:0). Initialize to the proper setting.
RX_CH_TAG_4(3:0) (15:12)	Tag information inserted into SE_D_OUT(3:0) immediately after TAG_3(3:0). Initialize to the proper setting.
RX_CH_TAG_5(3:0) (11:8)	Tag information inserted into SE_D_OUT(3:0) immediately after TAG_4(3:0). Initialize to the proper setting.
RX_CH_TAG_8(3:0) (7:4)	Tag information inserted into SE_D_OUT(3:0) immediately after TAG_7(3:0). Initialize to the proper setting. When connected across the fabric to a QSE, this field is OUTCHAN(15:12).
RX_CH_TAG_9(3:0) SP(1:0),MB,P (3:0)	Tag information inserted into SE_D_OUT(3:0) immediately after TAG_8(3:0). Set MB to 1 to enable the marked cell counters. Set P to 1 to create even parity. Set P to 0 to create odd parity to test the parity checker on the transmit side. Initialize to the proper setting.

NOTES:

- The QSE interprets TAG_2, TAG_3, TAG_0_6, and TAG_1_7 as the MCG.

9.3.1.3 RX_CH_NEW_VPI_VCI

Offset: 2_h (8_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not used (31)	Write with a 0 to maintain software compatibility with future versions.
RX_CH_VCI_TRANS (30)	Indicates that the VCI is to be translated. Initialize to the proper setting.
Not used (29:28)	Write with a 0 to maintain software compatibility with future versions.
RX_CH_OUTCHAN_NEW_VPI (27:16)	OUTCHAN(11:0). This field is always used. Initialize to the proper setting.
RX_CH_NEW_VCI (15:0)	The VCI value to be used when translation is enabled by RX_CH_VCI_TRANS. Initialize to the proper setting.

9.3.1.4 RX_CH_STATE (Internal Structure)

Offset: 3_h (C_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
RX_CH_AAL5_STATE (31:30)	Current AAL5 state. 00_b The last cell was not an end-of-frame and was not dropped. 01_b The last cell received was an end-of-frame. 10_b The last cell was dropped. 11_b Undefined. Initialize to 00_b. This field is maintained by a state machine and is read-only after initialization.
RX_CH_CELL_RCVD (29)	A cell has been received for this channel. Initialize to 0. This bit may not be cleared.
Not used (28:22)	Write with a 0 to maintain software compatibility with future versions.
RX_CH_LAST_CLEARED_SN (21:16)	Last cleared sequence numbered. Initialize to 0.
RX_CH_CONGESTED (15)	1 This channel is currently congested. 0 This channel is currently not congested. Initialize to 0. This field is maintained by a state machine and is read-only after initialization.
RX_CH_CUR_QD (14:0)	Current per-channel queue depth. Initialize to 0000_h. This field is maintained by a state machine and is read-only after initialization.

9.3.1.5 RX_CH_NEXT_CH/RX_CH_SEQ_PTR1 (Internal Structure)

Offset: 5_h (14_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
RX_CH_SEQ_CELL_PTR1 (31:16)	Pointer to the second outstanding cell. Initialize to 0. This field is maintained by a state machine and is read-only after initialization.
Not used (15:14)	Write with a 0 to maintain software compatibility with future versions.
RX_CH_NEXT_CH (13:0)	Next channel in the round-robin service. Initialize to 0000_h.

9.3.1.6 RX_CH_SEQ_ST (Internal Structure)

Address: 8_h (20_h byte)
Type: Read/Write
Format: Refer to the following tables.

If RX_CH_SEQ_TYPE = 00_b (One Cell Allowed to be Outstanding)

Field (Bits)	Description
RX_CH_RUN_LIMITED (31)	Indicates that this channel has been removed from the ring of channels for this SC because it has reached its limit.
Not used (30:18)	Write with a 0 to maintain software compatibility with future versions.
RX_CH_SEQ_STATE0 (17:16)	State of the cell sequencing state machine for cell pointer 0. Initialize to 0. This field is maintained by a state machine and is read-only after initialization. 00_b RX_CH_SEQ_CELL_PTR0 is empty. 01_b RX_CH_SEQ_CELL_PTR0 is currently pending transmission. 10_b RX_CH_SEQ_CELL_PTR0 received an ONACK on last transmission. 11_b RX_CH_SEQ_CELL_PTR0 received an ACK on last transmission.
RX_CH_SEQ_CELL_PTR0 (15:0)	Pointer to the first outstanding cell. Initialize to 0. This field is maintained by a state machine and is read-only after initialization.

If RX_CH_SEQ_TYPE = 01_b (Two Cells Allowed to be Outstanding)

Field (Bits)	Description
RX_CH_RUN_LIMITED (31)	Indicates this channel has been removed from the ring of channels for this SC because it has reached its limit.
RX_CH_PTR_NEXT_TO_CLR (30)	Resequencing state bit. Initialize to 0.
Not used (29:20)	Write with a 0 to maintain software compatibility with future versions.
RX_CH_SEQ_STATE (19:16)	State of the cell reassembly state machines. Initialize to 0. This field is maintained by a state machine and is read-only after initialization. (19:18) State bits for RX_CH_SEQ_CELL_PTR1. (17:16) State bits for RX_CH_SEQ_CELL_PTR0.
RX_CH_SEQ_CELL_PTR0 (15:0)	Pointer to the outstanding cell. Initialize to 0. This field is maintained by a state machine and is read-only after initialization.

9.3.1.7 TX_CH_CONFIG

Address: A_h (28_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
TX_CH_EXP_MAX_QD (31:28)	Exponent of the maximum per-SC queue depth. Initialize to the proper setting. Limits the TX_CH_CUR_QD to: $TX\_CH\_CUR\_QD \leq 1 + 2^{TX\_CH\_EXP\_MAX\_QD}$ A value of 0_h causes all cells for this channel to be dropped. A value of F_h limits this channel to 31744 cells.
TX_CH_EXP_CONG_QD (27:24)	Exponent of the congested per-SC queue depth. Initialize to the proper setting. The congestion state is entered when the threshold is exceeded, and the congestion state is exited when the queue length drops below 50 percent of this threshold. Enables congestion management when: $TX\_CH\_CUR\_QD > 1+2^{TX\_CH\_EXP\_CONG\_QD}$ A value of 0_h causes this CH to be in congestion at all times. A value of 1_h may not be used. A value of F_h causes this CH to enter congestion at a depth of 31744 cells.
TX_CH_ENTRY_PRESENT (23)	Initialize to 0, set other variables, then initialize to the proper setting. 1 Indicates this channel is present and valid. 0 Indicates this channel is disabled. Drop any cells received at this index, but still update TX_CH_CELLS_SENT and TX_CH_CELLS_DRPD (refer to sections "TX_CH_CELLS_SENT" on page 193 and "TX_CH_CELLS_DRPD" on page 193).
Not used (22:20)	Write with a 0 to maintain software compatibility with future versions.
TX_CH_SERVICE_CLASS (19:16)	Indicates the transmit SC.
TX_CH_EN_CLP_DROP (15)	Initialize to the proper setting. 1 Enables CLP dropping when the channel, the SC, the SSC, the VO, or the device is in the congested state. 0 Disables CLP dropping above the congestion threshold.
Not used (14)	Write with a 0 to maintain software compatibility with future versions.
TX_CH_EN_AAL5_EPD (13)	Initialize to the proper setting. 1 Enables AAL5 Early Packet Discard (EPD) when the channel, the SCQ, the SC, the VO, or the device is in the congested state. If the queue depth reaches the maximum threshold, or if CLP = 1 and the queue depth reaches the congestion threshold, the tail of the frame is dropped. 0 Disables AAL5 EPD.
TX_CH_EN_EFCI (12)	Written with a 0 to maintain compatibility with future software versions. 1 Enable EFCI. 0 Disable EFCI.
Reserved (11:8)	Initialize to 0 at initial setup. Software modifications to these locations after setup may cause incorrect operation.

(Continued)

Field (Bits)	Description
TX_CH_RESEQ_TYPE (7:6)	Initialize to the proper setting. 11 Reserved. 10 Reserved. 01 Supports resequencing of up to two cells, allowing a 155 Mbps VC. 00 No resequencing, allowing a 78 Mbps VC. The routing table at the other end of the connection must support the same resequencing mode for unicast connections. **This field must be set to 00 for multicast connections.**
Not used (5)	Write with a 0 to maintain software compatibility with future versions.
TX_CH_UC_VO (4:0)	The transmit UTOPIA level 2 address of the PHY interface for unicast cells. For multicast connections, this is the VO to use for VO congestion management. VO congestion management indicates only the unicast cells bound for that VO.

9.3.1.8 TX_NEW_VPI_RESEQ_PTR

Address: B_h ($2C_h$ byte)
Type: Read/Write
Format: Refer to the following tables.

If TX_CH_SERVICE_CLASS(3) = 0 (Unicast)

Field (Bits)	Description
TX_CH_VPI_TRANS (31)	Indicates the VPI is to be translated. Initialize to the proper setting.
Not used (30:28)	Write with a 0 to maintain software compatibility with future versions.
TX_CH_NEW_VPI (27:16)	The VPI value to be used. Initialize to the proper setting.
Not used (15:0)	Write with a 0 to maintain software compatibility with future versions.

If TX__CH_SERVICE_CLASS(3) = 1 (Multicast)

Field (Bits)	Description
TX_CH_MC_VO (31:27)	The transmit UTOPIA Level 2 address of the PHY interface for the first multicast cell in this linked list. Initialize to the proper setting.
Not used (26:17)	Write with a 0 to maintain software compatibility with future versions.
TX_CH_REPLICATE_CELL (16)	1 Make another copy of this cell. 0 Do not make another copy of this cell.
TX_CH_NEXT_MC_HEADER_PTR (15:0)	Index to the next translation field. Initialize to the proper setting.

9.3.1.9 TX_CH_STATE (Internal Structure)

Address: C_h (30_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
TX_CH_AAL5_STATE (31:30)	If TX_CH_RESEQ_TYPE = 00_b (refer to "TX_CH_RESEQ_TYPE" on page 190), the current AAL5 congestion state. 00_b The last cell received was not an end-of-frame and was not dropped. 01_b The last cell received was an end-of-frame. 10_b The last cell received was dropped. 11_b Undefined. Initialize to 00_b. This field is maintained by a state machine and is read-only after initialization.
TX_CH_CELL_RECEIVED (29)	A cell was received for this channel. Initialize to 0. This bit may not be cleared.
Reserved (28)	Initialize to 0 at initial setup. Software modifications to this location after setup may cause incorrect operation.
TX_CH_IN_DROPPING (27)	History bit that contains the result of past congestion management action when in 2-cell resequencing. Initialize to 0.
Reserved (26:22)	Initialize to 0 at initial setup. Software modifications to this location after setup may cause incorrect operation.
TX_CH_RUN_START_SN (21:16)	Sequence number where the next run will start. Initialize to "01_h".
TX_CH_CONGESTED (15)	1 This channel is currently congested. 0 This channel is currently not congested. Initialize to 0. This field is maintained by a state machine and is read-only after initialization.
TX_CH_CUR_QD (14:0)	Current per-channel queue depth. Initialize to 0000_h. This field is maintained by a state machine and is read-only after device initialization.

9.3.1.10 TX_CH_RESEQ_ST

Address: D_h (34_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not used (31:29)	Write with a 0 to maintain software compatibility with future versions.
TX_CH_RECOVERY_STATE (28)	State of the resequencer error recovery state machine. Initialize to 0_b. This field is maintained by a state machine and is read-only after initialization. 0_b The machine indicates normal operation. 1_b The machine has seen a sequence number error and is attempting to recover.
Not used (27:25)	Write with a 0 to maintain software compatibility with future versions.
TX_CH_STAT_STORE_STATE (24:23)	Initialize to 0_b. This field is maintained by a state machine and is read-only after initialization. When TX_COUNT_SENT_CLP = 0 (refer to "TX_COUNT_SENT_CLP" on page 106), then bit 23 means: 0_b If TX_CH_RESEQ_STATE = EMPTY (bit 22 in this register is set to 0_b), then the counts are accurate. Otherwise, TX_CH_CELLS_SENT (refer to "TX_CH_CELLS_SENT" on page 193) is one too high. 1_b If TX_CH_RESEQ_STATE ≠ EMPTY, then the send count is one too high. Otherwise, the TX_CH_CELLS_SENT is one too high and TX_CH_CELLS_DRPD (refer to "TX_CH_CELLS_DRPD" on page 193) is one too low. When TX_COUNT_SENT_CLP = 1, bit 23 is for the CLP = 0 counter and bit 24 is for the CLP = 1 counter. 0_b The machine has observed a cell count too many. 1_b The machine has observed correct accounting.
TX_CH_RESEQ_STATE (22)	State of the cell reassembly state machine. Initialize to 0_b. This field is maintained by a state machine and is read-only after initialization. 0_b The TX_CH_RESEQ_CELL_PTR is empty. 1_b The TX_CH_RESEQ_CELL_PTR contains a RECEIVED cell.
TX_CH_DROPPED_STORE (21)	0_b Indicates a cell was not dropped when the resequencer machine intended to store the cell. 1_b Indicates a cell was dropped after the resequencer machine intended to store the cell.
Not used (20:17)	Write with a 0 to maintain software compatibility with future versions.
TX_CH_RESEQ_OAM_CELL (16)	0_b Indicates the cell pending resequencing is not an OAM cell. 1_b Indicates the cell pending resequencing is an OAM cell.
TX_CH_RESEQ_CELL_PTR (15:0)	The pointer to the outstanding cell. Initialize to 0. This field is maintained by a state machine and is read-only after initialization.

9.3.1.11 TX_CH_CELLS_SENT

Address: E_h (38_h byte)
Type: Read/Write (Read-only after initialization)
Format: Refer to the following table.

Field (Bits)	Description
TX_CH_CELLS_SENT (31:0)	If TX_COUNT_SENT_CLP = 0 (refer to "TX_COUNT_SENT_CLP" on page 106), count of transmit cells sent. If TX_COUNT_SENT_CLP = 1, count of transmit CLP = 0 cells sent. Count rolls over at $FFFFFFFF_h$. Initialize to 00000000_h. Read-only after initialization.

9.3.1.12 TX_CH_CELLS_DRPD

Address: F_h ($3C_h$ byte)
Type: Read/Write (Read-only after initialization)
Format: Refer to the following table.

Field (Bits)	Description
TX_CH_CELLS_DRPD (31:0)	If TX_COUNT_SENT_CLP = 0 (refer to "TX_COUNT_SENT_CLP" on page 106), count of transmit cells dropped. If TX_COUNT_SENT_CLP = 1, count of transmit CLP = 1 cells sent. Count rolls over at $FFFFFFFF_h$. Initialize to 00000000_h. Read-only after initialization.

9.3.2 Multicast Control Block

Base address:400000_h (1000000_h byte)
Index:10_h
Entry number (refer to "CH_RAM_CONFIG" on page 105):
For CH_RAM_CONFIG = 0_h, the entry number is from 0 to 8191.
For CH_RAM_CONFIG = 1_h, the entry number is from 0 to 16385.
For CH_RAM_CONFIG = 2_h, the entry number is from 0 to 32767.
For CH_RAM_CONFIG = 3_h, the entry number is from 0 to 65534.

NOTE: The entry numbers listed above refer to a block of four words. Table 41 describes how four entries fit into the space of one Channel Control Block (CCB). Each block of 16 words can be used for either one CCB or four multicast control blocks.

Type: Read/Write

When multicasting in either direction, the long address = 400000_h + *TX_CH_NEXT_MC_HEADER_PTR* $\times 4_h$
(Refer to "MC_HEADER_PTR" on page 177).

Table 41. Multicast Control Block Summary

Byte Address	Long Offset	Name	Read or Write	Description
0_h	0_h	MC_LIST	R/W	Multicast linked list.
4_h	1_h	MC_NEW_VPI_VCI	R/W	Multicast VPI/VCI to which you want to translate.
8_h-C_h	2-3_h	Not used	R/W	Write with a 0 to maintain software compatibility with future versions.
10_h	4_h	MC_LIST	R/W	Multicast linked list.
14_h	5_h	MC_NEW_VPI_VCI	R/W	Multicast VPI/VCI to which you want to translate.
18_h-$1C_h$	6-7_h	Not used	R/W	Write with a 0 to maintain software compatibility with future versions.
20_h	8_h	MC_LIST	R/W	Multicast linked list.
24_h	9_h	MC_NEW_VPI_VCI	R/W	Multicast VPI/VCI to which you want to translate.
28_h-$2C_h$	A-B_h	Not used	R/W	Write with a 0 to maintain software compatibility with future versions.
30_h	C_h	MC_LIST	R/W	Multicast linked list.
34_h	D_h	MC_NEW_VPI_VCI	R/W	Multicast VPI/VCI to which you want to translate.
38_h-$3C_h$	E-F_h	Not used	R/W	Write with a 0 to maintain software compatibility with future versions.

9.3.2.1 MC_LIST

Address: 0_h or 4_h or 8_h or C_h (0_h or 10_h or 20_h or 30_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
TX_MC_VO (31:27)	The transmit UTOPIA VO for the next entry in the linked list.
Not used (26:17)	Write with a 0 to maintain software compatibility with future versions.
TX_MC_REPLICATE_CELL (16)	1 Make another copy of this cell. 0 Do not make another copy of this cell.
TX_MC_NEXT_MC_HEADER_PTR (15:0)	Index to the next translation field. Initialize to the proper setting.

9.3.2.2 MC_NEW_VPI_VCI

Offset: 1_h or 5_h or 9_h or D_h (4_h or 14_h or 24_h or 34_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
TX_MC_VPI_TRANS (31)	Indicates the VPI is to be translated. Initialize to the proper setting.
TX_MC_VCI_TRANS (30)	Indicates the VCI is to be translated. Initialize to the proper setting.
Not used (29:28)	Write with a 0 to maintain software compatibility with future versions.
TX_MC_NEW_VPI (27:16)	The VPI value to be used when translation is enabled by TX_MC_VPI_TRANS. Initialize to the proper setting.
TX_MC_NEW_VCI (15:0)	The VCI value to be used when translation is enabled by TX_MC_VCI_TRANS. Initialize to the proper setting.

9.4 ABR_RAM

Table 42 summarizes the contents of the AB_RAM.

Table 42. AB_RAM Summary

Byte Offset	Long Offset	Name	Read or Write	Description
0180000_h	600000_h	Receive Channel Queue Block	R/W	Receive channel head and tail pointers.
0184000_h	610000_h	Receive Channel Sent/Drop counters	R/W	Receive channel sent/drop cell counters

9.4.1 Receive Channel Queue Block

Base address: 600000_h (0180000_h byte)
Index: 2_h
The long address = 600000_h + *RX_CHAN_NUM* × 2_h+ *long_offset*

Table 43. Receive Channel Queue Block Summary

Address	Name	Read or Write	Description
0_h	RX_CH_HEAD	R/W	Head pointer of the receive channel queue.
1_h	RX_CH_TAIL	R/W	Tail pointer of the receive channel queue.

9.4.1.1 RX_CH_HEAD (Internal Structure)

Offset: 0_h (0_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	No RAM is present in these bit positions.
RX_QUEUE_HEAD (15:0)	The head of the channel queue.

9.4.1.2 RX_CH_TAIL (Internal Structure)

Offset: 1_h (4_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
Not present (31:16)	No RAM is present in these bit positions.
RX_QUEUE_TAIL (15:0)	The tail of the channel queue.

9.4.2 Receive Channel Statistics Block

Base address: 610000_h (0184000_h byte)
Index: 4_h
The long address = 610000_h + *RX_CHAN_NUM* × 4_h+ *long_offset*

The receive channel statistics counters are only valid if the RX_STATS bit in the RX_QUEUE_ENGINE_TEST register ("RX_QUEUE_ENGINE_TEST" on page 123) is set to 1

Table 44. Receive Channel Queue Block Summary

Address	Name	Read or Write	Description
0_h	RX_CH_SENT	R/W	Cells sent from this channel.
2_h	RX_CH_DROPPED	R/W	Cells dropped from this channel.

9.4.2.1 RX_CH_SENT (Internal Structure)

Offset: 0_h (0_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
RX_CH_SENT (31:0)	Total cells sent from this channel. Initialize to 0

9.4.2.2 RX_CH_DROPPED (Internal Structure)

Offset: 2_h (8_h byte)
Type: Read/Write
Format: Refer to the following table.

Field (Bits)	Description
RX_CH_DROPPED (31:0)	Total cells dropped from this channel. initialize to 0

9.5 SDRAM/SGRAM Interface Description

The SDRAMs/SGRAMs are used in the QRT system to store the cells in the receive and transmit directions. Before the SDRAMs/SGRAMs can be used, they must be prepared for operation. This section discusses setting up the SDRAM/SGRAM for normal operations, as well as test mode access by the microprocessor to check the data integrity of the SDRAM/SGRAM.

The receive and transmit SDRAMs/SGRAMs are each mapped to a memory location of the microprocessor address space, and the addresses within the RAMs are accessed by an indirect addressing method. The RX_DRAM_REGISTER is mapped to address 2000000 and TX_DRAM_REGISTER is mapped to address 3000000. At this address space exists a 16-bit, read/write register. For all set-up operations and write operations, the microprocessor writes to this address port. For reading a SDRAM/SGRAM location, writes are initially performed to set up the read cycles, and then this SDRAM/SGRAM register is read to obtain the value at the desired SDRAM/SGRAM location.

To perform a SDRAM/SGRAM operation, the corresponding access code (along with the address and data) is written to the microprocessor address port corresponding to the SDRAM/SGRAM. The SDRAM/SGRAM controller inside the QRT converts this access code to a SDRAM/ SGRAM access cycle obeying the signaling requirements of the SDRAM/SGRAM.

Access Code	Instruction
0000	NO OP
0001	READ BANK 0
0010	WRITE BANK 0
0011	ACTIVE BANK 0
0100	Reserved
0101	READ BANK 1
0110	WRITE BANK 1
0111	ACTIVE BANK 1
1000	Reserved
1001	PRECHARGE
1010	AUTOREF
1011	Reserved
1100	LOAD REG
1101	Reserved
1110	SELF REFRESH

After powerup or after a sustained software reset, the transmit and receive SDRAMs/SGRAMs need to be initialized. The RX_DRAM_REGISTER runs in a burst-of-four mode, whereas the TX_DRAM_REGISTER runs in a burst-of-eight mode, so the MODE register needs to be initialized differently.

To initialize the SDRAM/SGRAM:	Chip 0	Chip 1(SGRAM only)
W-No op	00000000_h	00004000_h
W-Precharge both banks	93000000_h	93004000_h
W-Autorefresh	$A2000000_h$	$A2004000_h$
W-Autorefresh	$A2000000_h$	$A2004000_h$
W-Load register C2YY (YY = 32 for RX_DRAM_REGISTER, 33 for TX_DRAM_REGISTER)	$C2320000_h$	$C2324000_h$
W-No op	00000000_h	00004000_h

To write a location to the SDRAM/SGRAM (for example, location 0_h in bank 0 with 55555555_h):

W-Active bank 0	32000000_h	32004000_h
W-Write bank 0	26000000_h	26004000_h

To read the location of the SDRAM/SGRAM (for example, location 0_h in bank 0):

W-Active bank 0	32000000_h	32004000_h
W-Read bank 0	12000000_h	12004000_h
R-	value of the DRAM location	

To write a location to the SDRAM/SGRAM (for example, location 1234_h in bank 1with $AAAAAAAA_h$):

W-Active bank 1	72120000_h	72124000_h
W-Write bank 1	$6A340000_h$	$6A344000_h$

9.5.1 RX_DRAM_REGISTER

Address: 2000000_h (8000000_h byte)
Type: Read/Write

Field (Bits)	Description
RX_DRAM_ACCESS_CODE (31:28)	The access code to control the type of SDRAM/SGRAM cycle.
RX_DRAM_DATA (27:26)	The data to be written to the DRAM. This 2-bit field is replicated 16 times and written to the 32-bit wide SDRAM/SGRAM. Data of 00, 01, 10, and 11 correspond to DRAM data 00000000_h, 55555555_h, $AAAAAAAA_h$, and $FFFFFFFF_h$. No other values of data can be written to the SDRAM/SGRAM.
RX_DRAM_MI_SELECT (25)	0 Normal operation. 1 Enables microprocessor access to the SDRAM/SGRAM.

(Continued)

Field (Bits)	Description
RX_DRAM_ADDRESS (24:16)	Corresponds to the row address during active cycles and column address during READ or WRITE cycles. The MSB is autoprecharge during Channel Associated Signaling (CAS) cycles (READ or WRITE).
RX_DRAM_UPPER_ADDR (15:14)	For SGRAM: Buffer RAM is organized as 2 chips with 512 rows each 00_b Select SGRAM chip 0. 01_b Select SGRAM chip 1. 10_b Invalid. 11_b Invalid. For SDRAM: Buffer RAM is organized as 2048 rows 00_b Select Rows 0 to 511 01_b Select Rows 512 to 1023. 10_b Select Rows 1024 to 1535 11_b Select Rows 1536 to 2047
Reserved (13:0)	Initialize to 0 at initial setup. Software modifications to this location after setup may cause incorrect operation.

9.5.2 TX_DRAM_REGISTER

Address: 3000000_h ($C000000_h$ byte)
Type: Read/Write

Field (Bits)	Description
TX_DRAM_ACCESS_CODE (31:28)	The access code to control the type of SDRAM/SGRAM cycle.
TX_DRAM_DATA (27:26)	The data to be written to the SDRAM/SGRAM. This 2-bit field is replicated 16 times and written to the 32-bit wide SDRAM/SGRAM. Data of 00, 01, 10, and 11 correspond to DRAM data 00000000_h, 55555555_h, $AAAAAAAA_h$, and $FFFFFFFF_h$. No other values of data can be written to the SDRAM/SGRAM.
TX_DRAM_MI_SELECT (25)	Select for microprocessor access to the SDRAM/SGRAM. Write with a 0 for normal operation.
TX_DRAM_ADDRESS (24:16)	Corresponds to the row address during ACTIVE cycles and column address during READ or WRITE cycles. The MSB is autoprecharge during CAS cycles (READ or WRITE).
TX_DRAM_UPPER_ADDR (15:14)	For SGRAM: Buffer RAM is organized as 2 chips with 512 rows each 00_b Select SGRAM chip 0. 01_b Select SGRAM chip 1. 10_b Invalid. 11_b Invalid. For SDRAM: Buffer RAM is organized as 2048 rows 00_b Select Rows 0 to 511 01_b Select Rows 512 to 1023. 10_b Select Rows 1024 to 1535 11_b Select Rows 1536 to 2047
Reserved (13:0)	Initialize to 0 at initial setup. Software modifications to this location after setup may cause incorrect operation.

9.5.3 Receive Cell Buffer SDRAM/SGRAM Summary (Internal Structure)

Base address: 2000000_h (8000000_h byte)
Index: 10_h
Entry number (refer to "AL_RAM_CONFIG" on page 105):
If AL_RAM_CONFIG = 0_h, the entry number is 8K.
If AL_RAM_CONFIG = 1_h, the entry number is 16K.
If AL_RAM_CONFIG = 2_h, the entry number is 64K.
If AL_RAM_CONFIG = 3_h, the entry number is 64K.
Read/Write: use RX_DRAM_REGISTER
Long address = 2000000_h + entry_number $\times$ 10_h

Table 45. Receive Cell Buffers SDRAM/SGRAM Summary

Offset	Cell			
	31:24	23:16	15:8	7:0
0	VPI(11:4)	VPI(3:0), VCI(15:12)	VCI(11:4)	VCI(3:0) PTI, CLP
1	payload 0	payload 1	payload 2	payload 3
2	payload 4	payload 5	payload 6	payload 7
3	payload 8	payload 9	payload 10	payload 11
4	payload 12	payload 13	payload 14	payload 15
5	payload 16	payload 17	payload 18	payload 19
6	payload 20	payload 21	payload 22	payload 23
7	payload 24	payload 25	payload 26	payload 27
8	payload 28	payload 29	payload 30	payload 31
9	payload 32	payload 33	payload 34	payload 35
10	payload 36	payload 37	payload 38	payload 39
11	payload 40	payload 41	payload 42	payload 43
12	payload 44	payload 45	payload 46	payload 47
13	00_h	00_h	00_h	0_h, Parity(3:0)
14	Not used	Not used	Not used	Not used
15	Not used	Not used	Not used	Not used

9.5.4 Transmit Cell Buffer SDRAM/SGRAM Summary (Internal Structure)

Base address: 3000000_h ($C000000_h$ byte)
Index: 10_h
Entry number (refer to "AL_RAM_CONFIG" on page 105):

If AL_RAM_CONFIG = 0_h, the entry number is 8K.
If AL_RAM_CONFIG = 1_h, the entry number is 16K.
If AL_RAM_CONFIG = 2_h, the entry number is 32K.
If AL_RAM_CONFIG = 3_h, the entry number is 64K.

Read/Write: use TX_DRAM_REGISTER
Long address = 3000000_h + entry_number $\times$ 10_h

Table 46. Transmit Cell Buffers SDRAM/SGRAM Summary

Offset	Cell			
	31:24	23:16	15:8	7:0
0	Parity(7:0)	00_h	OUTCHAN(15:8)	OUTCHAN(7:0)
1	VPI(11:4)	VPI(3:0), VCI(15:12)	VCI(11:4)	VCI(3:0), PTI, CLP
2	payload 0	payload 1	payload 2	payload 3
3	payload 4	payload 5	payload 6	payload 7
4	payload 8	payload 9	payload 10	payload 11
5	payload 12	payload 13	payload 14	payload 15
6	payload 16	payload 17	payload 18	payload 19
7	payload 20	payload 21	payload 22	payload 23
8	payload 24	payload 25	payload 26	payload 27
9	payload 28	payload 29	payload 30	payload 31
10	payload 32	payload 33	payload 34	payload 35
11	payload 36	payload 37	payload 38	payload 39
12	payload 40	payload 41	payload 42	payload 43
13	payload 44	payload 45	payload 46	payload 47
14	Not used	Not used	Not used	Not used
15	Not used	Not used	Not used	Not used

10 JTAG

JTAG Support

The QRT supports the IEEE Boundary Scan Specification as described in the IEEE 1149.1 standards. The Test Access Port (TAP) consists of the five standard pins, TRSTB, TCK, TMS, TDI and TDO, used to control the TAP controller and the boundary scan registers. The TRSTB input is the active low reset signal used to reset the TAP controller. TCK is the test clock used to sample data on input, TDI and to output data on output, TDO. The TMS input is used to direct the TAP controller through its states. The basic boundary scan architecture is shown below.

Figure 70. Boundary Scan Architecture

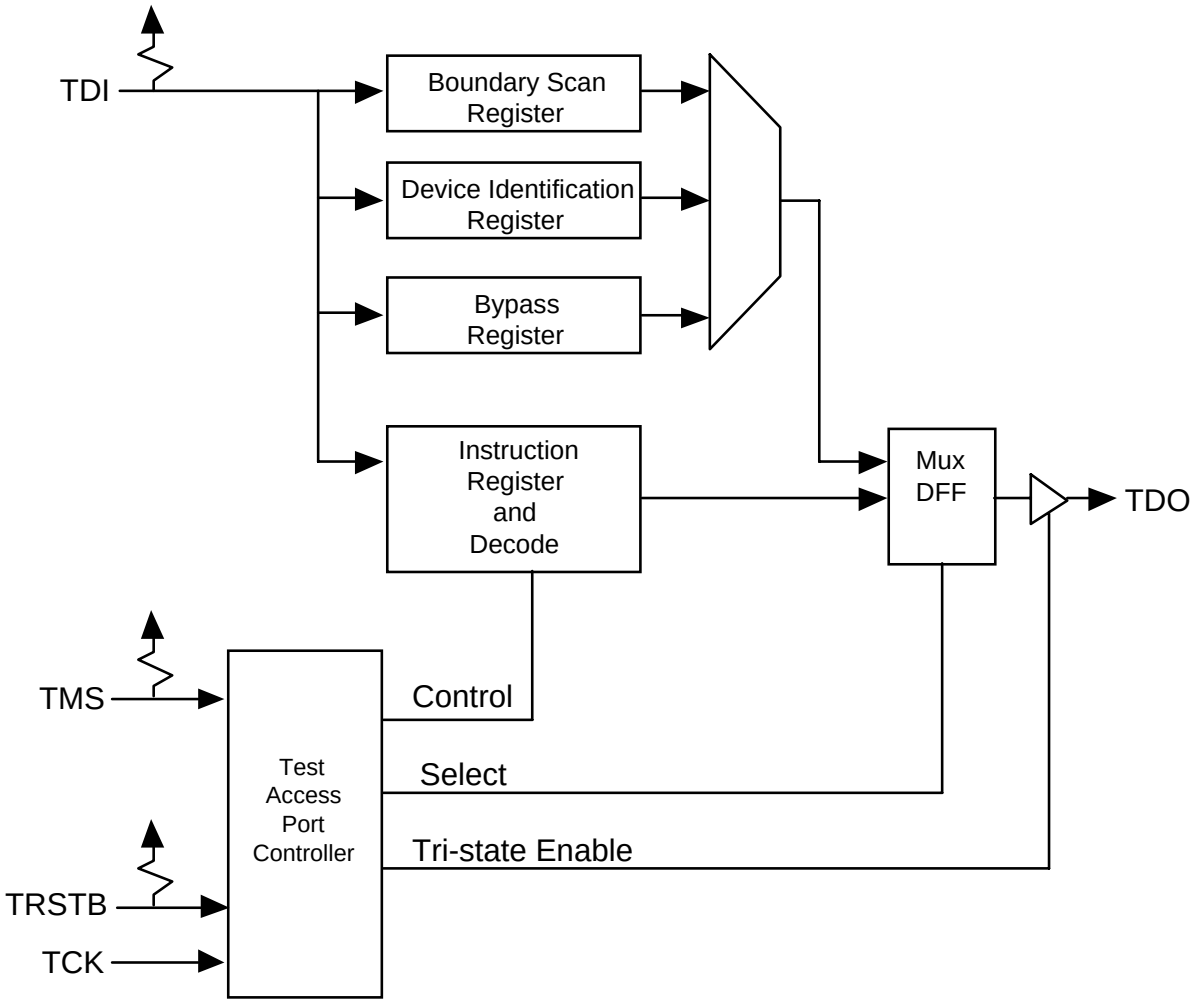

The boundary scan architecture consists of a TAP controller, an instruction register with instruction decode, a bypass register, a device identification register and a boundary scan register. The TAP controller interprets the TMS input and generates control signals to load the instruction and data registers. The instruction register with instruction decode block is used to select the test to be executed and/or the register to be accessed. The bypass register offers a single bit delay from pri-

mary input, TDI to primary output, TDO. The device identification register contains the device identification code.

The boundary scan register allows testing of board inter-connectivity. The boundary scan register consists of a shift register place in series with device inputs and outputs. Using the boundary scan register, all digital inputs can be sampled and shifted out on primary output, TDO. In addition, patterns can be shifted in on primary input, TDI and forced onto all digital outputs.

TAP Controller

The TAP controller is a synchronous finite state machine clocked by the rising edge of primary input, TCK. All state transitions are controlled using primary input, TMS. The finite state machine is described below.

Figure 71. TAP Controller Finite State Machine

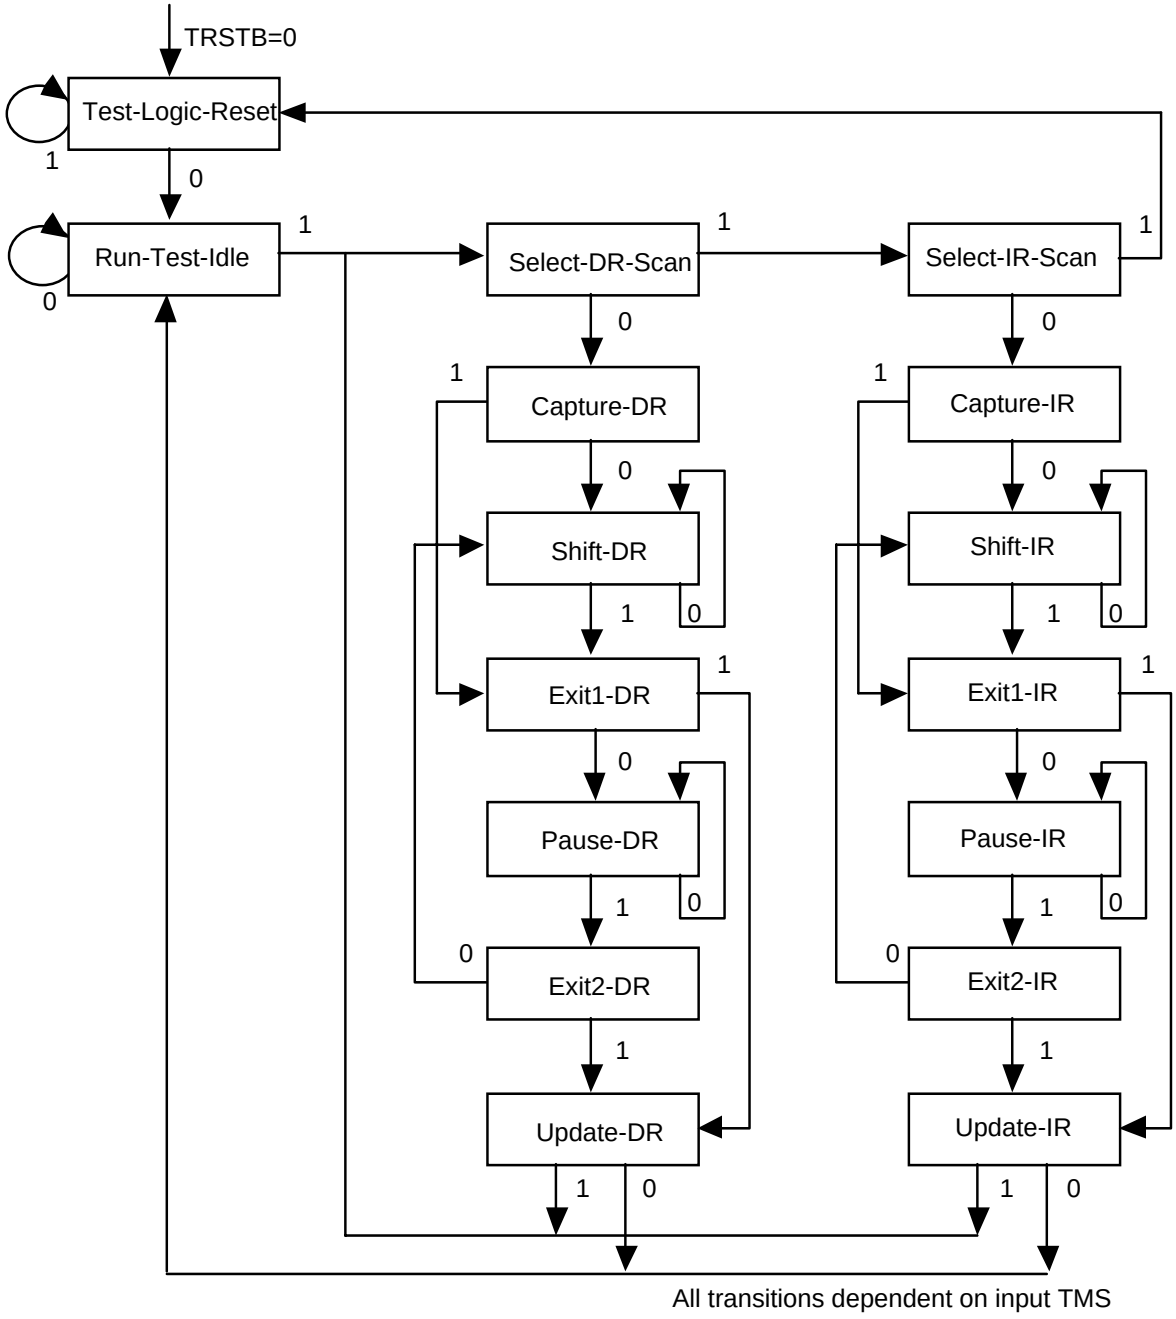

Test-Logic-Reset:

The test logic reset state is used to disable the TAP logic when the device is in normal mode operation. The state is entered asynchronously by asserting input, TRSTB. The state is entered synchronously regardless of the current TAP controller state by forcing input, TMS high for 5 TCK clock cycles. While in this state, the instruction register is set to the IDCODE instruction.

Run-Test-Idle:

The run test/idle state is used to execute tests.

Capture-DR:

The capture data register state is used to load parallel data into the test data registers selected by the current instruction. If the selected register does not allow parallel loads or no loading is required by the current instruction, the test register maintains its value. Loading occurs on the rising edge of TCK.

Shift-DR:

The shift data register state is used to shift the selected test data registers by one stage. Shifting is from MSB to LSB and occurs on the rising edge of TCK.

Update-DR:

The update data register state is used to load a test register's parallel output latch. In general, the output latches are used to control the device. For example, for the EXTEST instruction, the boundary scan test register's parallel output latches are used to control the device's outputs. The parallel output latches are updated on the falling edge of TCK.

Capture-IR:

The capture instruction register state is used to load the instruction register with a fixed instruction. The load occurs on the rising edge of TCK.

Shift-IR:

The shift instruction register state is used to shift both the instruction register and the selected test data registers by one stage. Shifting is from MSB to LSB and occurs on the rising edge of TCK.

Update-IR:

The update instruction register state is used to load a new instruction into the instruction register. The new instruction must be scanned in using the Shift-IR state. The load occurs on the falling edge of TCK.

The Pause-DR and Pause-IR states are provided to allow shifting through the test data and/or instruction registers to be momentarily paused.

The TDO output is enabled during states Shift-DR and Shift-IR. Otherwise, it is tri-stated.

Boundary Scan Instructions

The following is a description of the standard instructions. Each instruction selects an serial test data register path between input, TDI, and output, TDO.

BYPASS

The bypass instruction shifts data from input TDI to output TDO with one TCK clock period delay. The instruction is used to bypass the device.

EXTEST

The external test instruction allows testing of the interconnection to other devices. When the current instruction is the EXTEST instruction, the boundary scan register is place between input TDI and output TDO. Primary device inputs can be sampled by loading the boundary scan register using the Capture-DR state. The sampled values can then be viewed by shifting the boundary scan register using

the Shift-DR state. Primary device outputs can be controlled by loading patterns shifted in through input TDI into the boundary scan register using the Update-DR state.

SAMPLE

The sample instruction samples all the device inputs and outputs. For this instruction, the boundary scan register is placed between TDI and TDO. Primary device inputs and outputs can be sampled by loading the boundary scan register using the Capture-DR state. The sampled values can then be viewed by shifting the boundary scan register using the Shift-DR state.

IDCODE

The identification instruction is used to connect the identification register between TDI and TDO. The device's identification code can then be shifted out using the Shift-DR state.

STCTEST

The single transport chain instruction is used to test out the TAP controller and the boundary scan register during production test. When this instruction is the current instruction, the boundary scan register is connected between TDI and TDO. During the Capture-DR state, the device identification code is loaded into the boundary scan register. The code can then be shifted out on output TDO using the Shift-DR state.

Table 47. Boundary Scan Pin Order

- Pin Type Output Enable -- controls the enable pin of 3 state drivers
- Pin Type Clock -- Observable only cell
- Pin Type Input -- Observable and controllable input cell
- Pin Type Output3 -- Bidirectional 3 state output
- Pin Type Output2 -- Three state output

Order #	Pin #	Pin name	Pin Type
1		HIZ	Output enable
2		HIZ	Output enable
3		HIZ	Output enable
4		HIZ	Output enable
5		HIZ	Output enable
6		HIZ	Output enable
7		HIZ	Output enable
8		HIZ	Output enable

9	E7	SCANENBN	clock
10	C3	RESETN	clock
11	E4	ABRAMADD.0	output3
12	E4	ABRAMADD.0	input
13	G5	ABRAMADD.1	output3
14	G5	ABRAMADD.1	input
15	D3	ABRAMADD.2	output3
16	D3	ABRAMADD.2	input
17	F5	ABRAMADD.3	output3
18	F5	ABRAMADD.3	input
19	C2	ABRAMADD.4	output3
20	C2	ABRAMADD.4	input
21	F4	ABRAMADD.5	output3
22	F4	ABRAMADD.5	input
23	E3	ABRAMADD.6	output3
24	E3	ABRAMADD.6	input
25	D2	ABRAMADD.7	output3
26	D2	ABRAMADD.7	input
27	G4	ABRAMADD.8	output3
28	G4	ABRAMADD.8	input
29	F3	ABRAMADD.9	output3
30	F3	ABRAMADD.9	input
31	H5	ABRAMADD.10	output3
32	H5	ABRAMADD.10	input
33	G3	ABRAMADD.11	output3
34	G3	ABRAMADD.11	input
35	E2	ABRAMADD.12	output3
36	E2	ABRAMADD.12	input
37	F2	ABRAMADD.13	output3
38	F2	ABRAMADD.13	input
39	H4	ABRAMADD.14	output3

40	H4	ABRAMADD.14	input
41	J5	ABRAMADD.15	output3
42	J5	ABRAMADD.15	input
43	G2	ABRAMADD.16	output3
44	G2	ABRAMADD.16	input
45	K4	ABRAMCLK	output2
46	F1	ABRAMWEN	output2
47	L5	ABRAMOEN	output2
48	H3	ABRAMADVN	output2
49	J4	ABRAMADSPN	output2
50	K3	CHRAMCLK	output2
51	L4	CHRAMADSCN	output2
52	K5	CHRAMADD17N	output2
53	J2	CHRAMOEN	output2
54	J3	CHRAMWE0N	output2
55	L3	CHRAMWE1N	output2
56	H2	CHRAMADD.0	output2
57	N5	CHRAMADD.1	output2
58	P5	CHRAMADD.2	output2
59	M5	CHRAMADD.3	output2
60	M3	CHRAMADD.4	output2
61	M4	CHRAMADD.5	output2
62	P4	CHRAMADD.6	output2
63	L2	CHRAMADD.7	output2
64	K2	CHRAMADD.8	output2
65	P3	CHRAMADD.9	output2
66	N4	CHRAMADD.10	output2
67	M2	CHRAMADD.11	output2
68	N3	CHRAMADD.12	output2

69	N2	CHRAMADD.13	output2
70	N1	CHRAMADD.14	output2
71	R4	CHRAMADD.15	output2
72	R3	CHRAMADD.16	output2
73	R5	CHRAMADD.17	output2
74	P2	CHRAMDATA.0	output3
75	P2	CHRAMDATA.0	input
76	T5	CHRAMDATA.1	output3
77	T5	CHRAMDATA.1	input
78	T4	CHRAMDATA.2	output3
79	T4	CHRAMDATA.2	input
80	T3	CHRAMDATA.3	output3
81	T3	CHRAMDATA.3	input
82	R2	CHRAMDATA.4	output3
83	R2	CHRAMDATA.4	input
84	U1	CHRAMDATA.5	output3
85	U1	CHRAMDATA.5	input
86	T2	CHRAMDATA.6	output3
87	T2	CHRAMDATA.6	input
88	U2	CHRAMDATA.7	output3
89	U2	CHRAMDATA.7	input
90	U3	CHRAMDATA.8	output3
91	U3	CHRAMDATA.8	input
92	V4	CHRAMDATA.9	output3
93	V4	CHRAMDATA.9	input
94	V2	CHRAMDATA.10	output3
95	V2	CHRAMDATA.10	input
96	W2	CHRAMDATA.11	output3
97	W2	CHRAMDATA.11	input
98	U4	CHRAMDATA.12	output3
99	U4	CHRAMDATA.12	input

100	V3	CHRAMDATA.13	output3
101	V3	CHRAMDATA.13	input
102	W4	CHRAMDATA.14	output3
103	W4	CHRAMDATA.14	input
104	U5	CHRAMDATA.15	output3
105	U5	CHRAMDATA.15	input
106	W5	CHRAMPARITY.0	output3
107	W5	CHRAMPARITY.0	input
108	V5	CHRAMDATA.16	output3
109	V5	CHRAMDATA.16	input
110	W3	CHRAMDATA.17	output3
111	W3	CHRAMDATA.17	input
112	AA2	CHRAMDATA.18	output3
113	AA2	CHRAMDATA.18	input
114	AA3	CHRAMDATA.19	output3
115	AA3	CHRAMDATA.19	input
116	Y2	CHRAMDATA.20	output3
117	Y2	CHRAMDATA.20	input
118	Y4	CHRAMDATA.21	output3
119	Y4	CHRAMDATA.21	input
120	AA5	CHRAMDATA.22	output3
121	AA5	CHRAMDATA.22	input
122	Y3	CHRAMDATA.23	output3
123	Y3	CHRAMDATA.23	input
124	Y5	CHRAMDATA.24	output3
125	Y5	CHRAMDATA.24	input
126	AB4	CHRAMDATA.25	output3
127	AB4	CHRAMDATA.25	input
128	AB3	CHRAMDATA.26	output3
129	AB3	CHRAMDATA.26	input
130	AA4	CHRAMDATA.27	output3

131	AA4	CHRAMDATA.27	input
132	AC2	CHRAMDATA.28	output3
133	AC2	CHRAMDATA.28	input
134	AB2	CHRAMDATA.29	output3
135	AB2	CHRAMDATA.29	input
136	AC3	CHRAMDATA.30	output3
137	AC3	CHRAMDATA.30	input
138	AB5	CHRAMDATA.31	output3
139	AB5	CHRAMDATA.31	input
140	AD3	CHRAMPARITY.1	output3
141	AD3	CHRAMPARITY.1	input
142	AD1	ALRAMCLK	output2
143	AC4	ALRAMADSCN	output2
144	AD4	ALRAMADD17N	output2
145	AD2	ALRAMADD18N	output2
146	AE2	ALRAMOEN	output2
147	AF2	ALRAMWE0N	output2
148	AE3	ALRAMADD.0	output2
149	AG2	ALRAMADD.1	output2
150	AF3	ALRAMADD.2	output2
151	AF4	ALRAMADD.3	output2
152	AC5	ALRAMADD.4	output2
153	AE4	ALRAMADD.5	output2
154	AG3	ALRAMADD.6	output2
155	AF5	ALRAMADD.7	output2
156	AE6	ALRAMADD.8	output2
157	AE7	ALRAMADD.9	output2
158	AG4	ALRAMADD.10	output2
159	AG5	ALRAMADD.11	output2
160	AH4	ALRAMADD.12	output2
161	AF6	ALRAMADD.13	output2

162	AF7	ALRAMADD.14	output2
163	AH5	ALRAMADD.15	output2
164	AG6	ALRAMADD.16	output2
165	AG7	ALRAMADD.17	output2
166	AE8	ALRAMADD.18	output2
167	AH6	ALRAMDATA.0	output3
168	AH6	ALRAMDATA.0	input
169	AH7	ALRAMDATA.1	output3
170	AH7	ALRAMDATA.1	input
171	AE9	ALRAMDATA.2	output3
172	AE9	ALRAMDATA.2	input
173	AF8	ALRAMDATA.3	output3
174	AF8	ALRAMDATA.3	input
175	AH8	ALRAMDATA.4	output3
176	AH8	ALRAMDATA.4	input
177	AJ6	ALRAMDATA.5	output3
178	AJ6	ALRAMDATA.5	input
179	AF10	ALRAMDATA.6	output3
180	AF10	ALRAMDATA.6	input
181	AG8	ALRAMDATA.7	output3
182	AG8	ALRAMDATA.7	input
183	AF9	ALRAMDATA.8	output3
184	AF9	ALRAMDATA.8	input
185	AE11	ALRAMDATA.9	output3
186	AE11	ALRAMDATA.9	input
187	AG10	ALRAMDATA.10	output3
188	AG10	ALRAMDATA.10	input
189	AE10	ALRAMDATA.11	output3
190	AE10	ALRAMDATA.11	input
191	AF11	ALRAMDATA.12	output3
192	AF11	ALRAMDATA.12	input

193	AH10	ALRAMDATA.13	output3
194	AH10	ALRAMDATA.13	input
195	AG9	ALRAMDATA.14	output3
196	AG9	ALRAMDATA.14	input
197	AH9	ALRAMDATA.15	output3
198	AH9	ALRAMDATA.15	input
199	AG11	ALRAMDATA.16	output3
200	AG11	ALRAMDATA.16	input
201	AE13	TATMPARITY	output2
202	AE12	TATMSOC	output2
203	AE14	TATMWEN	output2
204	AF12	TATMCLAV.0	clock
205	AG12	TATMCLAV.1	clock
206	AF14	TATMCLAV.2	clock
207	AH11	TATMCLAV.3	clock
208	AH12	TATMDATA.0	output2
209	AF13	TATMDATA.1	output2
210	AG14	TATMDATA.2	output2
211	AG13	TATMDATA.3	output2
212	AH14	TATMDATA.4	output2
213	AH13	TATMDATA.5	output2
214	AJ13	TATMDATA.6	output2
215	AG15	TATMDATA.7	output2
216	AF15	TATMDATA.8	output2
217	AE15	TATMDATA.9	output2
218	AH15	TATMDATA.10	output2
219	AE16	TATMDATA.11	output2
220	AG16	TATMDATA.12	output2
221	AF16	TATMDATA.13	output2
222	AH16	TATMDATA.14	output2
223	AH17	TATMDATA.15	output2

224	AG18	TATMADD.0	output2
225	AJ17	TATMADD.1	output2
226	AF18	TATMADD.2	output2
227	AG17	TATMADD.3	output2
228	AH19	TATMADD.4	output2
229	AH18	TXDRAMCLK	output2
230	AF17	TXDRAMBA	output2
231	AF19	TXDRAMCASN	output2
232	AG19	TXDRAMRASN	output2
233	AE19	TXDRAMWEN	output2
234	AE17	TXDRAMADD.0	output2
235	AE18	TXDRAMADD.1	output2
236	AH21	TXDRAMADD.2	output2
237	AG20	TXDRAMADD.3	output2
238	AG22	TXDRAMADD.4	output2
239	AH20	TXDRAMADD.5	output2
240	AE21	TXDRAMADD.6	output2
241	AF20	TXDRAMADD.7	output2
242	AG21	TXDRAMADD.8	output2
243	AF22	TXDRAMCSN.0	output2
244	AH23	TXDRAMCSN.1	output2
245	AE20	TXDRAMDATA.0	output3
246	AE20	TXDRAMDATA.0	input
247	AH24	TXDRAMDATA.1	output3
248	AH24	TXDRAMDATA.1	input
249	AF21	TXDRAMDATA.2	output3
250	AF21	TXDRAMDATA.2	input
251	AH22	TXDRAMDATA.3	output3
252	AH22	TXDRAMDATA.3	input
253	AE22	TXDRAMDATA.4	output3
254	AE22	TXDRAMDATA.4	input

255	AG23	TXDRAMDATA.5	output3
256	AG23	TXDRAMDATA.5	input
257	AG24	TXDRAMDATA.6	output3
258	AG24	TXDRAMDATA.6	input
259	AJ24	TXDRAMDATA.7	output3
260	AJ24	TXDRAMDATA.7	input
261	AF24	TXDRAMDATA.8	output3
262	AF24	TXDRAMDATA.8	input
263	AF23	TXDRAMDATA.9	output3
264	AF23	TXDRAMDATA.9	input
265	AH25	TXDRAMDATA.10	output3
266	AH25	TXDRAMDATA.10	input
267	AE24	TXDRAMDATA.11	output3
268	AE24	TXDRAMDATA.11	input
269	AH26	TXDRAMDATA.12	output3
270	AH26	TXDRAMDATA.12	input
271	AG26	TXDRAMDATA.13	output3
272	AG26	TXDRAMDATA.13	input
273	AG25	TXDRAMDATA.14	output3
274	AG25	TXDRAMDATA.14	input
275	AH27	TXDRAMDATA.15	output3
276	AH27	TXDRAMDATA.15	input
277	AF25	TXDRAMDATA.16	output3
278	AF25	TXDRAMDATA.16	input
279	AF26	TXDRAMDATA.17	output3
280	AF26	TXDRAMDATA.17	input
281	AE23	TXDRAMDATA.18	output3
282	AE23	TXDRAMDATA.18	input
283	AG27	TXDRAMDATA.19	output3
284	AG27	TXDRAMDATA.19	input
285	AE26	TXDRAMDATA.20	output3

286	AE26	TXDRAMDATA.20	input
287	AG28	TXDRAMDATA.21	output3
288	AG28	TXDRAMDATA.21	input
289	AC25	TXDRAMDATA.22	output3
290	AC25	TXDRAMDATA.22	input
291	AF27	TXDRAMDATA.23	output3
292	AF27	TXDRAMDATA.23	input
293	AF28	TXDRAMDATA.24	output3
294	AF28	TXDRAMDATA.24	input
295	AE27	TXDRAMDATA.25	output3
296	AE27	TXDRAMDATA.25	input
297	AD26	TXDRAMDATA.26	output3
298	AD26	TXDRAMDATA.26	input
299	AC26	TXDRAMDATA.27	output3
300	AC26	TXDRAMDATA.27	input
301	AD27	TXDRAMDATA.28	output3
302	AD27	TXDRAMDATA.28	input
303	AE28	TXDRAMDATA.29	output3
304	AE28	TXDRAMDATA.29	input
305	AC27	TXDRAMDATA.30	output3
306	AC27	TXDRAMDATA.30	input
307	AB25	TXDRAMDATA.31	output3
308	AB25	TXDRAMDATA.31	input
309	AD28	SEDIN0.0	clock
310	AC28	SEDIN0.1	clock
311	AA25	SEDIN0.2	clock
312	AB26	SEDIN0.3	clock
313	AB28	SEDIN1.0	clock
314	AD29	SEDIN1.1	clock
315	Y26	SEDIN1.2	clock
316	AB27	SEDIN1.3	clock

317	W25	SEDIN2.0	clock
318	Y27	SEDIN2.1	clock
319	Y25	SEDIN2.2	clock
320	W26	SEDIN2.3	clock
321	Y28	SEDIN3.0	clock
322	AA27	SEDIN3.1	clock
323	AA28	SEDIN3.2	clock
324	W27	SEDIN3.3	clock
325	U25	SESOCI.0	clock
326	V25	SESOCI.1	clock
327	T25	SESOCI.2	clock
328	V26	SESOCI.3	clock
329	W28	RXCELLSTART	clock
330	V28	BPACKOUT.0	output2
331	U26	BPACKOUT.1	output2
332	T27	BPACKOUT.2	output2
333	U27	BPACKOUT.3	output2
334	T28	BPACKIN.0	clock
335	U28	BPACKIN.1	clock
336	U29	BPACKIN.2	clock
337	R27	BPACKIN.3	clock
338	R26	SDOUT0.0	output2
339	R25	SDOUT0.1	output2
340	R28	SDOUT0.2	output2
341	P25	SDOUT0.3	output2
342	P27	SDOUT1.0	output2
343	P26	SDOUT1.1	output2
344	P28	SDOUT1.2	output2
345	N28	SDOUT1.3	output2
346	M27	SDOUT2.0	output2
347	N29	SDOUT2.1	output2

348	M26	SDOUT2.2	output2
349	N27	SDOUT2.3	output2
350	L28	SDOUT3.0	output2
351	M28	SDOUT3.1	output2
352	N26	SDOUT3.2	output2
353	L26	SDOUT3.3	output2
354	L27	SSOCOUT	output2
355	L25	DRAMCKE	output2
356	N25	RXDRAMBA	output2
357	M25	RXDRAMCASN	output2
358	J28	RXDRAMCLK	output2
359	K27	RXDRAMRASN	output2
360	H27	RXDRAMWEN	output2
361	K28	RXDRAMADD.0	output2
362	J25	RXDRAMADD.1	output2
363	K26	RXDRAMADD.2	output2
364	J27	RXDRAMADD.3	output2
365	H26	RXDRAMADD.4	output2
366	G28	RXDRAMADD.5	output2
367	K25	RXDRAMADD.6	output2
368	F28	RXDRAMADD.7	output2
369	J26	RXDRAMADD.8	output2
370	H28	RXDRAMCSN.0	output2
371	H25	RXDRAMCSN.1	output2
372	G27	RXDRAMDATA.0	output3
373	G27	RXDRAMDATA.0	input
374	F27	RXDRAMDATA.1	output3
375	F27	RXDRAMDATA.1	input
376	F29	RXDRAMDATA.2	output3
377	F29	RXDRAMDATA.2	input
378	F26	RXDRAMDATA.3	output3

379	F26	RXDRAMDATA.3	input
380	G26	RXDRAMDATA.4	output3
381	G26	RXDRAMDATA.4	input
382	E28	RXDRAMDATA.5	output3
383	E28	RXDRAMDATA.5	input
384	F25	RXDRAMDATA.6	output3
385	F25	RXDRAMDATA.6	input
386	D28	RXDRAMDATA.7	output3
387	D28	RXDRAMDATA.7	input
388	D27	RXDRAMDATA.8	output3
389	D27	RXDRAMDATA.8	input
390	E27	RXDRAMDATA.9	output3
391	E27	RXDRAMDATA.9	input
392	C28	RXDRAMDATA.10	output3
393	C28	RXDRAMDATA.10	input
394	E26	RXDRAMDATA.11	output3
395	E26	RXDRAMDATA.11	input
396	B28	RXDRAMDATA.12	output3
397	B28	RXDRAMDATA.12	input
398	G25	RXDRAMDATA.13	output3
399	G25	RXDRAMDATA.13	input
400	C27	RXDRAMDATA.14	output3
401	C27	RXDRAMDATA.14	input
402	D25	RXDRAMDATA.15	output3
403	D25	RXDRAMDATA.15	input
404	E23	RXDRAMDATA.16	output3
405	E23	RXDRAMDATA.16	input
406	C26	RXDRAMDATA.17	output3
407	C26	RXDRAMDATA.17	input
408	E24	RXDRAMDATA.18	output3
409	E24	RXDRAMDATA.18	input

410	B27	RXDRAMDATA.19	output3
411	B27	RXDRAMDATA.19	input
412	D24	RXDRAMDATA.20	output3
413	D24	RXDRAMDATA.20	input
414	C25	RXDRAMDATA.21	output3
415	C25	RXDRAMDATA.21	input
416	B26	RXDRAMDATA.22	output3
417	B26	RXDRAMDATA.22	input
418	D23	RXDRAMDATA.23	output3
419	D23	RXDRAMDATA.23	input
420	C24	RXDRAMDATA.24	output3
421	C24	RXDRAMDATA.24	input
422	E22	RXDRAMDATA.25	output3
423	E22	RXDRAMDATA.25	input
424	C23	RXDRAMDATA.26	output3
425	C23	RXDRAMDATA.26	input
426	B25	RXDRAMDATA.27	output3
427	B25	RXDRAMDATA.27	input
428	B24	RXDRAMDATA.28	output3
429	B24	RXDRAMDATA.28	input
430	D22	RXDRAMDATA.29	output3
431	D22	RXDRAMDATA.29	input
432	E21	RXDRAMDATA.30	output3
433	E21	RXDRAMDATA.30	input
434	B23	RXDRAMDATA.31	output3
435	B23	RXDRAMDATA.31	input
436	D20	RATMCLAV.0	clock
437	A24	RATMCLAV.1	clock
438	E19	RATMCLAV.2	clock
439	C22	RATMCLAV.3	clock
440	D21	RATMSOC	clock

441	C20	RATMREN	output2
442	D19	RATMADD.0	output2
443	E20	RATMADD.1	output2
444	B21	RATMADD.2	output2
445	C21	RATMADD.3	output2
446	C19	RATMADD.4	output2
447	B22	RATMDATA.0	clock
448	E17	RATMDATA.1	clock
449	E16	RATMDATA.2	clock
450	E18	RATMDATA.3	clock
451	C18	RATMDATA.4	clock
452	D18	RATMDATA.5	clock
453	D16	RATMDATA.6	input
454	B19	RATMDATA.7	clock
455	B20	RATMDATA.8	clock
456	C16	RATMDATA.9	clock
457	D17	RATMDATA.10	clock
458	B18	RATMDATA.11	clock
459	C17	RATMDATA.12	clock
460	B17	RATMDATA.13	clock
461	A17	RATMDATA.14	clock
462	D15	RATMDATA.15	clock
463	E15	ADDRDATA.0	output3
464	E15	ADDRDATA.0	input
465	B15	ADDRDATA.1	output3
466	B15	ADDRDATA.1	input
467	E14	ADDRDATA.2	output3
468	E14	ADDRDATA.2	input
469	D14	ADDRDATA.3	output3
470	D14	ADDRDATA.3	input
471	C15	ADDRDATA.4	output3

472	C15	ADDRDATA.4	input
473	B14	ADDRDATA.5	output3
474	B14	ADDRDATA.5	input
475	A13	ADDRDATA.6	output3
476	A13	ADDRDATA.6	input
477	C14	ADDRDATA.7	output3
478	C14	ADDRDATA.7	input
479	B13	ADDRDATA.8	output3
480	B13	ADDRDATA.8	input
481	C13	ADDRDATA.9	output3
482	C13	ADDRDATA.9	input
483	D12	ADDRDATA.10	output3
484	D12	ADDRDATA.10	input
485	B12	ADDRDATA.11	output3
486	B12	ADDRDATA.11	input
487	B11	ADDRDATA.12	output3
488	B11	ADDRDATA.12	input
489	D13	ADDRDATA.13	output3
490	D13	ADDRDATA.13	input
491	C12	ADDRDATA.14	output3
492	C12	ADDRDATA.14	input
493	D11	ADDRDATA.15	output3
494	D11	ADDRDATA.15	input
495	E13	ADDRDATA.16	output3
496	E13	ADDRDATA.16	input
497	E11	ADDRDATA.17	output3
498	E11	ADDRDATA.17	input
499	E12	ADDRDATA.18	output3
500	E12	ADDRDATA.18	input
501	C11	ADDRDATA.19	output3
502	C11	ADDRDATA.19	input

503	B9	ADDRDATA.20	output3
504	B9	ADDRDATA.20	input
505	C9	ADDRDATA.21	output3
506	C9	ADDRDATA.21	input
507	B10	ADDRDATA.22	output3
508	B10	ADDRDATA.22	input
509	D10	ADDRDATA.23	output3
510	D10	ADDRDATA.23	input
511	E9	ADDRDATA.24	output3
512	E9	ADDRDATA.24	input
513	C10	ADDRDATA.25	output3
514	C10	ADDRDATA.25	input
515	E10	ADDRDATA.26	output3
516	E10	ADDRDATA.26	input
517	D8	ADDRDATA.27	output3
518	D8	ADDRDATA.27	input
519	C8	ADDRDATA.28	output3
520	C8	ADDRDATA.28	input
521	D9	ADDRDATA.29	output3
522	D9	ADDRDATA.29	input
523	B7	ADDRDATA.30	output3
524	B7	ADDRDATA.30	input
525	B8	ADDRDATA.31	output3
526	B8	ADDRDATA.31	input
527	C7	ADSN	clock
528	E8	WRITE	clock
529	A6	CSN	clock
530	D7	READYN	output2
531	B6	INTRN	output2
532	D6	STATSTRB	output2

11 APPLICATION NOTES

11.1 Connecting QRTs to QSEs Using Gigabit Ethernet Transceivers

QRTs can be connected to QSEs using Gigabit Ethernet transceivers. Several companies, including AMCC® and Vitesse™ Semiconductor Corporation, make Gigabit Ethernet transceivers. Vitesse has three devices: two old (the VSC7135 and the VSC7136) and one new (the VSC7214). The PHY layer for Gigabit Ethernet is taken from Fiber Channel, so these devices are derived from Fiber Channel devices.

The new VSC7214 is a quad device with integrated 8B/10B encoders. The information about the new VSC7214 is at:

http://www.vitesse.com/news/052297.html

The existing VSC7135 and VSC7136 devices require an external 8B/10B encoder/decoder. The web site for the existing VSC7135 and VSC7136 (and the VCS7214) is at:

http://www.vitesse.com/products/stdprod.html#1

The general interconnection method is to replace the first nibble of four cells with an 8B/10B comma character, copying the Present (Pres) and Multicast (MC) cell bits of all four to the second byte, and not sending the inverted Pres bit or the Spare (SP) bit. By doing this replacement, a full byte is freed, leaving room for the comma character and allowing a simple interface. The SOC signal then needs to be reproduced at the far end.

These serialization devices introduce significant latency, which eats into the one-celltime budget for round-trip delay. Sending the BP_ACK_IN(3:0) signal without serialization helps reduce the round-trip latency.

Figure 72 shows an example of connecting the QRT to the QSE using the Vitesse VSC7135 and the VSC7214.

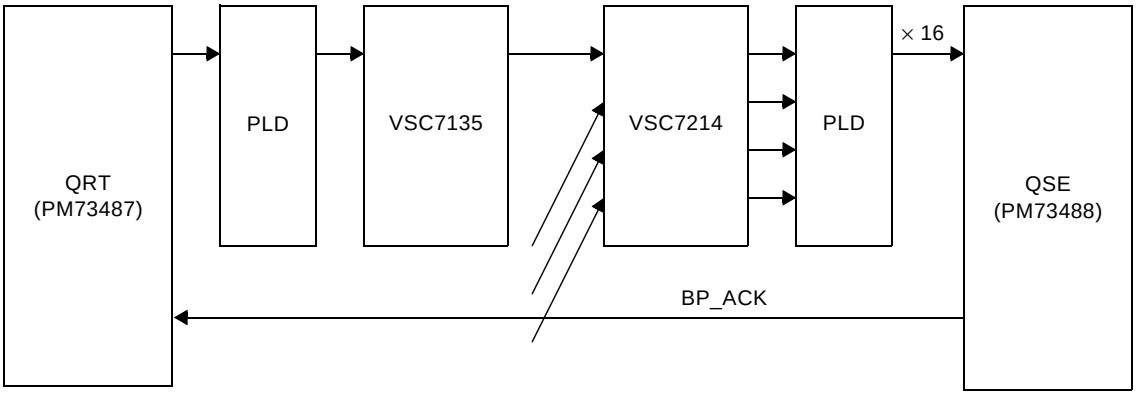

Figure 72. Connecting the QRT to Gigabit Ethernet Transceivers

Another interconnection method is to put all the signals from two or three interfaces over a link, including the SOC line.

NOTE: The BP_ACK_IN(3:0) signal does not always go to the same card as the data and SOC signals. For example, the data signals might go to the first stage of the fabric and the BP_ACK_IN(3:0) signal would go to the third stage of the fabric. For this reason, the BP_ACK_IN(3:0) signal often cannot be combined with the data and SOC signals in a serialized interface. These signals can be combined in a switch fabric that had the first and last stages of the fabric on the same card.

11.2 Connecting to Standard Serializer/Deserializer Chipsets

The QRT/QSE can be connected using standard serializer/deserializer chipsets such as those provided by Hewlett Packard® and National Semiconductor®.

The Hewlett Packard serializer/deserializer chipsets, HDMP-1012/1014, HDMP-1022/1024, and HDMP-1032/1034 (preliminary), are ECL, 5V and 3.3V devices, respectively, that can run at 66 Mhz.

These devices are ideally suited for gang four applications, where the 16 data lines and a single copy of the SOC can be turned into a single differential pair.

In certain 3-stage fabrics, the BPACK and data can share the serializer: 16 data lines, a single copy of the SOC and the four BPACK signals can be turned into a single differential pair when the devices run at 62.5 Mhz.

More infomation about these devices can be found at:

http://www.hp.com/HP-COMP/fiber/sg/gen_chip.html

The National Semiconductor® serializer/deserializer chipsets, the DS90CR283/284 and the DS90CR285/286, are 5V and 3.3 V devices, respectively, that run at 66 MHz. They use the EIA-644 Low Voltage Differential Standard (LVDS) to send the signals differentially at a rate of up to 462 Mbit/s per LVDS data chennel. The setup and hold times match leaving the QRT/QSE. Coming into the QRT/QSE, the phase aligners are turned on.

These devices are best suited for gang four applications, where the 16 data lines and a single copy of the SOC can be accommodated with a single chipset.

In certain 3-stage fabrics, the BPACK and data can share the serializer: 16 data lines, a single copy of the SOC and the four BPACK signals. This occurs in situations where the first and last stages are on the same card and the second stage is on another card.

The chipset spec requires that skew must be tightly controlled between all the LVDS pairs, and all of the outputs must go to the same place.

The web site for the devices is:

http://www.national.com/pf/DS/DS90CR283.html
http://www.national.com/pf/DS/DS90CR284.html
http://www.national.com/pf/DS/DS90CR285.html
http://www.national.com/pf/DS/DS90CR286.html

11.3 Connecting the QRT to the S/UNI-ATLAS PM7324

The S/UNI-ATLAS PM7324 is a full duplex Fault Management (FM) OAM, Performance Management (PM) OAM, and policing device. It has a UTOPIA interface on the PHY side and a SATURN interface on the other side. It can interface directly to the QRT by utilizing the QRT's OC-12 single phy mode. In the UTOPIA_CONFIG register (refer to "7.2.11 UTOPIA_CONFIG" starting on page 117), RX_OC_12C_MODE should be set to 1, TX_OC_12C_MODE can be set to either 0 or 1, and UTOPIA_2 should be set to 1. Please refer to the PMC document PM980583 for a detailed description of this application

Figure 73 shows an example of connecting the QRT to the S/UNI-ATLAS.

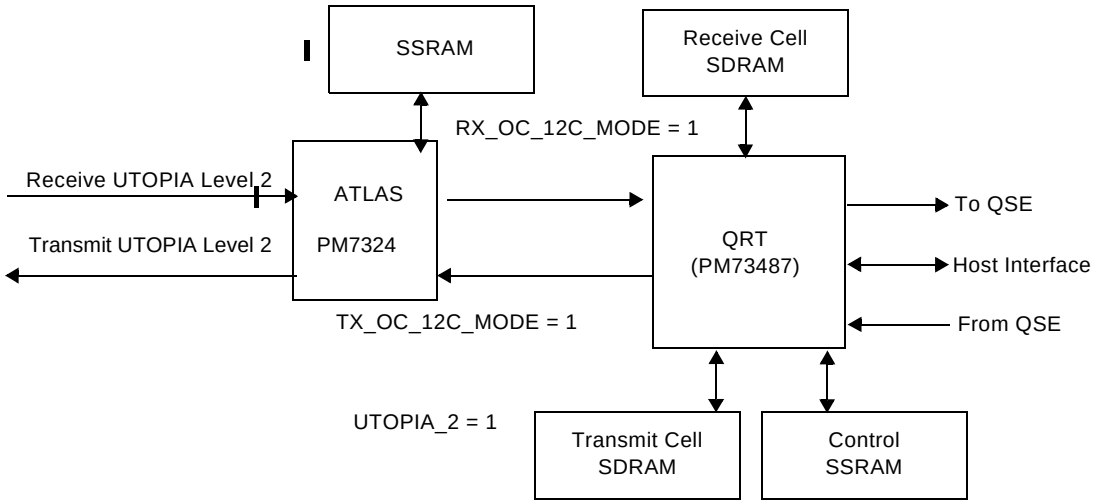

Figure 73. Connecting the QRT to the RCMP-800

11.4 Relationships Among Various Clock Domains

Relationships among various clock domains on the QRT are complex and require careful consideration for effective use of the device and optimal performance. Duty cycle tolerance for SYSCLK, ATMCLK, SE_CLK and PCLK is 60/40.

.

The maximum operating frequency of the QRT is 100 MHz; however, it is possible to operate the QRT at frequencies below 100 MHz if various relationship constraints are met. Frequency relationships also determine the switch speed-up factor and device throughput. Phase aligners only converge in a certain frequency range. DRAMs may potentially lose data below a certain operating frequency, and during processor transactions data integrity is guaranteed if the ATM_CLK-to-SYSCLK relationship is maintained.

11.4.1 Relationship Among the SYSCLK, SE_CLK, and the Switch Speed-Up Factor

The SE_CLK relationship with the SYSCLK is governed by the following equation.

$$SE\_CLK = (66 \div 100) \times SYSCLK$$

Thus, if SYSCLK is lowered in frequency, SE_CLK frequency must be lowered in the same proportion to maintain switch fabric speed-up factor at 1.6.

Switch speed-up factor requirements are switch-topology dependent. For a detailed speed-up factor discussion, refer to the *QSE (PM73488) Long Form Data Sheet*.

11.4.2 The Phase Aligner SE_CLK Frequency Constraint

Phase aligners at the QRT switch fabric interfaces are designed to operate in the SE_CLK frequency range of 50 MHz to 66 MHz. Phase aligners are not expected to converge beyond this frequency range. If an application requires the SE_CLK to be less than 50 MHz, it should operate the device with the phase aligners turned off.

11.4.3 The SYSCLK DRAM Refresh Constraint

In normal operating mode, the QRT refreshes one row of DRAM per cell time. A cell time is defined as the sum of 118 SE_CLK periods.

$$Cell\ Time = 118 \times (1 \div SE\_CLK)$$

A typical SGRAM requires 1024 rows to be refreshed in 17 ms. A typical SDRAM requires 2048 rows to be refreshed in 32 ms. This implies a row refresh every 15.6 µs. The SE_CLK therefore, should never be operated below 7.552 MHz. This also implies that the SYSCLK never be operated below 11.44 MHz.

11.4.4 Relationship Between ATM_CLK and SYSCLK

The maximum frequency allowed at the UTOPIA interface per the UTOPIA Level 2 specification is 50 MHz., However, the QRT is designed to operate at up to 55 MHz if the ATM layer devices or PHY devices permit it.

The maximum bandwidth of the QRT is 720 Mbit/s when SYSCLK is at 100 MHz, ATM_CLK is at 50 MHz, and SECLK is at 66 MHz. Theoretically, at these clock rates, throughput is a function of SYSCLK, SE_CLK, and the switch topology. Speeding up the UTOPIA interface can improve the Cell Transfer Delay (CTD), but it cannot improve the throughput.

The relationship between the ATM_CLK and SYSCLK is defined by the following equation for maximum throughput:

$$\text{ATM\_CLK} \leq (52 \div 100) \times \text{SYSCLK}$$

If ATM_CLK operates below this ratio, the throughput of the device drops linearly in the same proportion. When ATM_CLK is operated above this ratio, the QRT UTOPIA interface throttles back, disallowing bandwidth greater than what the device could handle.

11.4.5 Relationship Between the PCLK and the SYSCLK

The upper bound on the PCLK is defined by the following equation:

$$\text{PCLK} \leq (50 \div 100) \times \text{SYSCLK}$$

If PCLK exceeds this bound, the handshake protocol between the processor interface and the rest of the device could be violated.

The lower bound on the PCLK is defined by the following equation:

$$\text{PCLK} > (1 \div 8) \times \text{SYSCLK}$$

If this constraint is not met, incorrect data could be registered during write operations.

APPENDIX A NOMENCLATURE

A.1 Definitions

Transmit (egress) signals: all signals related to processing the data heading towards the optical/ electrical layer.
Receive (ingress) signals: all signals related to processing the data heading towards the ATM layer.

A.2 Signal Name Prefixes

Many pins have prefixes indicating the function with which they are associated (refer to Table 48).

Table 48. Prefixes and Associated Functions

Pin Name Prefixes	Associated Functions
RATM	Receive direction of UTOPIA ATM
RAM	RAM Interface
RX	Receive
JTAG	JTAG Interface
TATM	Transmit direction of UTOPIA ATM
TX	Transmit

A.3 Numbers

- Hexadecimal numbers are followed by the suffix "$_h$", for example: 1_h, $2C_h$.
- Binary numbers are followed by the suffix "$_b$", for example: 00_b.
- Decimal numbers appear without suffixes.

A.4 Glossary of Abbreviations

Abbreviation	Description
AAL	ATM Adaptation Layer
AAL1	ATM Adaptation Layer 1
AAL5	ATM Adaptation Layer 5
ABR	Available Bit Rate
ACK	Acknowledgment
AL	Address Lookup (as in AL RAM)
ATM	Asynchronous Transfer Mode or ATM Layer
BECN	Backwards Explicit Congestion Notification
BP	Backpressure
BPACK	Backpressure Acknowledgment

Abbreviation	Description
CAM	Content-Addressable Memory
CAS	Channel Associated Signaling
CBR	Constant Bit Rate
CCB	Channel Control Block
CDV	Cell Delay Variation
CH	Channel (as in CH RAM)
CLAV	Cell Available
CLP	Cell Loss Priority
CMOS	Complementary Metal Oxide Semiconductor
CPU	Central Processing Unit
CTD	Cell Transfer Delay
DQM	I/O mask enables
DRAM	Dynamic Random Access Memory
DS0	Digital Signal Level 1
DS1	Digital Signal Level 1
DSI	Direct Status Indication
EFCI	Explicit Forward Congestion Indication
EIA	Electronic Industries Association
EOF	End-Of-Frame
EPBGA	Enhanced Plastic Ball Grid Array
EPD	Early Packet Discard
ER	Explicit Rate
FECN	Forward Explicit Congestion Notification
FIFO	First In, First Out
FM	Fault Management
GFC	Generic Flow Control
GP	General Purpose
HEC	Header Error Check
IRT	Input half of QRT
LAN	Local Area Network
LSB	Least Significant Bit
LSW	Least Significant Word
LV	Low Voltage Differential
LVDS	Low Voltage Differential Standard
MB	Mark Bit

Abbreviation	Description
MBS	Maximum Burst Size
MC	Multicast
MCR	Minimum Cell Rate
MCTG	Multicast Translation Group
MNACK	Medium Negative ACKnowledgment or Mid-switch Negative ACKnowledgment
Mod	Modulo
MPHY	Multi-PHY
MSB	Most Significant Bit
MSP	Multiplexed Status Polling
MSW	Most Significant Word
NACK	Negative ACKnowledgment
nrtVBR	Non Real-Time Variable Bit Rate
OAM	Operations, Administration, and Maintenance
OC-3	Optical Carrier level 3 (155.52 Mbit/s)
OC-12	Optical Carrier level 12 (622.08 Mbit/s)
ONACK	Output Negative ACKnowledgment
ORT	Output half of QRT
PB	Proportional Bandwidth
PHY	Physical or Physical Layer or Physical Layer Device
PLD	Programmable Logic Device
PM	Performance Management
PPD	Partial Packet Discard
PTD	Packet Tail Discard
PTI	Payload Type Indicator
QoS	Quality-of-Service
QRT	PM73487
QSE	PM73488
RAM	Random Access Memory
RCMP-800	Routing Control, Monitoring and Policing for 800 Mbit/s (PMC-Sierra device)
RSF	Receive Switch Fabric
rtVBR	Real-Time Variable Bit Rate
RX	Receive
SAR	Segmentation And Reassembly
SATURN	SONET/ATM User Network (refers to the interface defined by the SATURN Development Group)

Abbreviation	Description
SC	Service Class
SCR	Sustainable Cell Rate
SCQ	Service Class Queue
SDRAM	Synchronous Dynamic Random Access Memory
SESAR	Switched Ethernet Segmentation And Reassembly
SGRAM	Synchronous Graphic Random Access Memory
SN	Sequence Number
SOJ	Small Outline J-lead
SOC	Start-Of-Cell
SP	Spare
SRAM	Static Random Access Memory
SSRAM	Synchronous Static Random Access Memory
Tbit/s	Terabits per second
TSF	Transmit Switch Fabric
TTL	Transistor-toTransitor Logic
TX	Transmit
UBR	Unspecified Bit Rate
UNI	User-to-Network Interface
UPC	Usage Parameter Control
UTOPIA	Universal Test and Operations Interface for ATM
VBR	Variable Bit Rate
VC	Virtual Channel
VCC	Virtual Channel Connection
VCI	Virtual Channel Identifier
VI	Virtual Input
VIH	Input High Voltage
VIL	Input Low Voltage
VO	Virtual Output
VP	Virtual Path
VPC	Virtual Path Connection
VPI	Virtual Path Identifier
UNI	User-to-Network Interface
VBR	Variable Bit Rate
WAN	Wide Area Network

APPENDIX B REFERENCES

- ATM Forum, *ATM User-Network Interface Specification*, V3.0, September 10, 1993.
- ATM Forum, *UTOPIA – An ATM PHY Data Path Interface, Level 1*, V2.01, February, 1994.
- ATM Forum, *UTOPIA – An ATM - PHY Interface Specification, Level 2*, V1.0, June 1995.
- IEEE 1149.1, *Standard Test Access Port and Boundary Scan Architecture*, May 21, 1990.
- ITU (CCITT) Recommendation I.432, *B-ISDN User-Network Interface - Physical Interface Specification*, June 1990.
- PMC-Sierra, Inc., *ATM Switch Fabric Element (QSE – PM73488) Long Form Data Sheet*.

APPENDIX C RAM SELECTION

Table 49. QRT RAM Selection

NAME	Description	Type	Size	Suggested RAMs	Interface	NOTES
AB_RAM	Contains head & tail pointers and sent/drop counters for RX queues	SSRAM	128kx18	MT58LC**128k18**D9-10 MT58LC**128k18**C6-10	17bit multiplex bus.	See note "B"
AL_RAM	VPI/VCI Address Lookup tables. Cell buffer pointers, service order tables and Multicast pointer FIFOs. (Different ALRAM_CONFIG modes correspond to different VPI/VCI ranges and cell buffer sizes)	SSRAM	128kx18 256kx18* 512kx18* 1. ALRAM_CONFIG mode 0 = 64Kx18 2. ALRAM_CONFIG mode 1 = 128Kx 18 3. ALRAM_CONFIG mode 2 = 256Kx18 4. ALRAM_CONFIG mode 3 = 512Kx18	MT58LC**128k18**D9-10 MT58LC**128k18**C6-10 MT58LC**256k18**D9-10 MT58LC**256k18**C6-10	17bit data, 18bit address	See note "B" Can use 2 256kx18 SCD chips by setting ALRAM_TYPE = 1 (see datasheet section 7.2.4)
CH_RAM	Per Tx/Rx channel table info.	SSRAM	64kx36 = 4k VCs 128kx36 = 8k VCs* 256kx36 = 16k VCs*	MT58LC**128k18**C6-10 MT58LC**256k18**C6-10	34bit data, 17bit address	See note "B"
Rx DRAM	Receive cell buffers. Per VC queuing. 64k cell buffers.	SDRAM or SGRAM	1Mx16 SDRAM * 256kx32 SGRAM	MT48LC**1M16**A1-10 (SDRAM) MT41LC**256k32**D5-10 (SGRAM)	32 bit data, 9bit address, 2bit CS.	SDRAM, must use 2 chips (64k cells). SGRAM, can use 1 or 2 chips (16k or 32k cells).
Tx DRAM	Transmit cell buffers. Per Virtual Output & Service Class queuing. 64k cell buffers.	SDRAM or SGRAM	1Mx16 SDRAM * 256kx32 SGRAM	MT48LC**1M16**A1-10 (SDRAM) MT41LC**256k32**D5-10 (SGRAM)	32 bit data, 9bit address, 2bit CS.	SDRAM, must use 2 chips (64k cells). SGRAM, can use 1 or 2 chips (16k or 32k cells).

* Most common memory configurations.

Note "B" : SCD SSRAM can be used if the entire address space of the Ram bank is within 1 RAM chip.
SCD = Single Cycle Deselect memory. The notation 'D9' can be seen in the name of these SSRAM.
DCD = Double Cycle Deselect memory. The notation 'C6' can be seen in the name of these SSRAM.

ORDERING INFORMATION

Table 50 lists the ordering information.

Table 50. Ordering Information

Part Number	Description
PM73487-PI	503-pin Enhanced Plastic Ball Grid Array (EPBGA) package

CONTACTING PMC-SIERRA, INC.

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

Tel: (604) 415-6000
Fax: (604) 415-6200

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

PMC-980618 (R3) ref PMC-981001(R2) Issue date: June 1999