IBM®

IBM PowerPRS Q-64G Packet Routing Switch

Datasheet

Preliminary

December 20, 2001

The following are trademarks of International Business Machines Corporation in the United States, or other countries, or both.

HyperBGA	IBM Logo
IBM	PowerPRS

InfiniBand is a trademark of the InfiniBand Trade Association.

IEEE is a registered trademark of the Institute of Electrical and Electronics Engineers (IEEE).

Other company, product and service names may be trademarks or service marks of others.

All information contained in this document is subject to change without notice. The products described in this document are NOT intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change IBM product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of IBM or third parties. All information contained in this document was obtained in specific environments, and is presented as an illustration. The results obtained in other operating environments may vary.

While the information contained herein is believed to be accurate, such information is preliminary, and should not be relied upon for accuracy or completeness, and no representations or warranties of accuracy or completeness are made.

Note: This document contains information on products in the sampling and/or initial production phases of development. This information is subject to change without notice. Verify with your IBM field applications engineer that you have the latest version of this document before finalizing a design.

THE INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED ON AN “AS IS” BASIS. In no event will IBM be liable for damages arising directly or indirectly from any use of the information contained in this document.

IBM Microelectronics Division
1580 Route 52, Bldg. 504
Hopewell Junction, NY 12533-6351

The IBM home page can be found at
http://www.ibm.com

The IBM Microelectronics Division home page
can be found at http://www.chips.ibm.com

prsq-64g.01title.fm.
December 20, 2001

Contents

List of Figures

List of Tables

1. General Information

1.1 Features

- Nonblocking, self-routing, single-stage switch
- Either 16 or 32 input and output ports
- High performance:
  - Throughput of 2 Gbps per port without speed expansion
  - Aggregate throughput of up to 64 Gbps for one device
- Speed expansion:
  - 16-Gbps logical interface using multiple devices configured for speed expansion
  - 512-Gbps aggregate throughput for eight devices configured for external speed expansion (32 ports at 16 Gbps)
  - 256-Gbps aggregate throughput for four devices configured for internal and external speed expansion (16 ports at 16 Gbps)
  - Other configurations (employing up to seven devices) can be used in specific applications
- Serial data communication of up to 2.5 Gbps, compatible with InfiniBand™ physical layer standards
- Multicast support without packet duplication in the shared memory
- Configurable number of traffic priorities (from one to four)
- Flow control based on a grant mechanism
- Programmable flow control thresholds
- Subport flow control support
- Support for redundant switch-plane operation, including a scheduled switchover facility that operates without packet loss
- Serial processor interface (serial host interface)
- Packet header of two or three bytes, containing destination bitmap, packet priority, and switch redundancy support information, all protected by a parity bit
- Shared memory comprised of a dynamically shared buffer with a total capacity of:
  - Up to 4096 packets of 64, 72, or 80 bytes for eight devices
  - Up to 2048 packets of 64, 72, or 80 bytes for four devices
- 8b/10b encoding for link synchronization and supervision
- Reception of control packets destined for the local processor on any input port
- Transmission of control packets from the local processor to any output port
- Detection of link liveness by reception of specific packets
- Programmable byte shuffling in egress packets
- CMOS 7SF (SA-27E) technology (L_{drawn} = 0.18 µm, L_{eff} = 0.11 µm): 1.8-V LVCMOS-compatible I/O for low-speed signals
- IEEE® Standard 1149.1 boundary scan to facilitate circuit-board testing
- 624-ball IBM HyperBGA™ package

1.2 Description

The IBM PowerPRS™ Q-64G Packet Routing Switch is one of a family of third-generation switching devices designed for high-performance, nonblocking, fixed-length packet switching. It enables the development of scalable switch fabrics with an aggregate bandwidth of 256 to 512 Gbps.

The PowerPRS Q-64G receives packets on up to 32 input ports and routes them to up to 32 output ports based on bitmap information contained in the packet header. To accomplish this, each PowerPRS Q-64G contains four 16 × 16 subswitch elements connected internally for port expansion. The physical links between the PowerPRS Q-64G and the attached devices are high-speed serial links called Unilinks.

The PowerPRS Q-64G is designed to provide OC-192 attachment. To meet these transmission requirements and provide 16-Gbps throughput per port, multiple PowerPRS Q-64Gs are configured for speed expansion. There are two standard multiple-device configurations:

- 512-Gbps configuration (with eight devices providing 32 input and output ports)
- 256-Gbps configuration (with four devices providing 16 input and output ports)

Synchronization is not required between input ports. However, packets on a given port are always received or transmitted at a fixed rate according to the packet length. Four levels of packet priority provide quality-of-service support. A serial grant mechanism controls ingress and egress data flow.

The PowerPRS Q-64G supports redundant switch-plane operation. It includes color-coded scheduled switchover that operates without packet loss. Scheduled switchover is a system-level function that requires hardware and software interaction. The PowerPRS Q-64G performs the hardware assist for this function.

1.3 Ordering Information

Part Number	Description	Single-Device Throughput	Aggregate Throughput
IBM3229P2815	IBM Packet Routing Switch	64 Gbps	For OC-192 attachment: 256 Gbps (four devices) or 512 Gbps (eight devices)

1.4 Conventions and Notation

Throughout this document, standard IBM notation is used, meaning that bits and bytes are numbered in ascending order from left to right. For a four-byte word, bit 0 is the most significant bit (MSB) and bit 31 is the least significant bit (LSB).

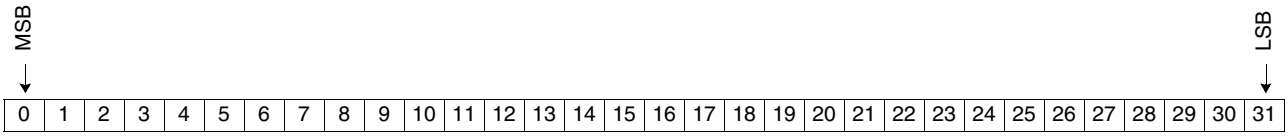

Notation for bit encoding is as follows:

- Hexadecimal values are preceded by an *x* and enclosed in single quotation marks. For example: x‘0A00’.
- Binary values in sentences appear in single quotation marks. For example: ‘1010’.

Differential pairs are designated by an _P for the positive signal and an _N for the negative signal at the end of the signal name. For example: PortDataIn[0]_P and PortDataIn[0]_N.

Nondifferential signals that are active low are designated by a # symbol at the end of the signal name. For example: InterruptOut#.

2. Architecture

2.1 System Application

The IBM PowerPRS Q-64G Packet Routing Switch enables the construction of nonblocking scalable switch fabrics through repeated instances of the same switch element. It is designed for a wide variety of applications, including campus, wide-area network (WAN) edge, access, and backbone switches. When connected to the IBM PowerPRS C192 Common Switch Interface, the PowerPRS Q-64G provides a complete redundant switch fabric for the attachment of OC-48 and OC-192 protocol engines. An example of this architecture is shown in *Figure 2-1.*

Figure 2-1. System View of the PowerPRS Q-64G with the PowerPRS C192 (configured with redundant 512-Gbps switch planes)

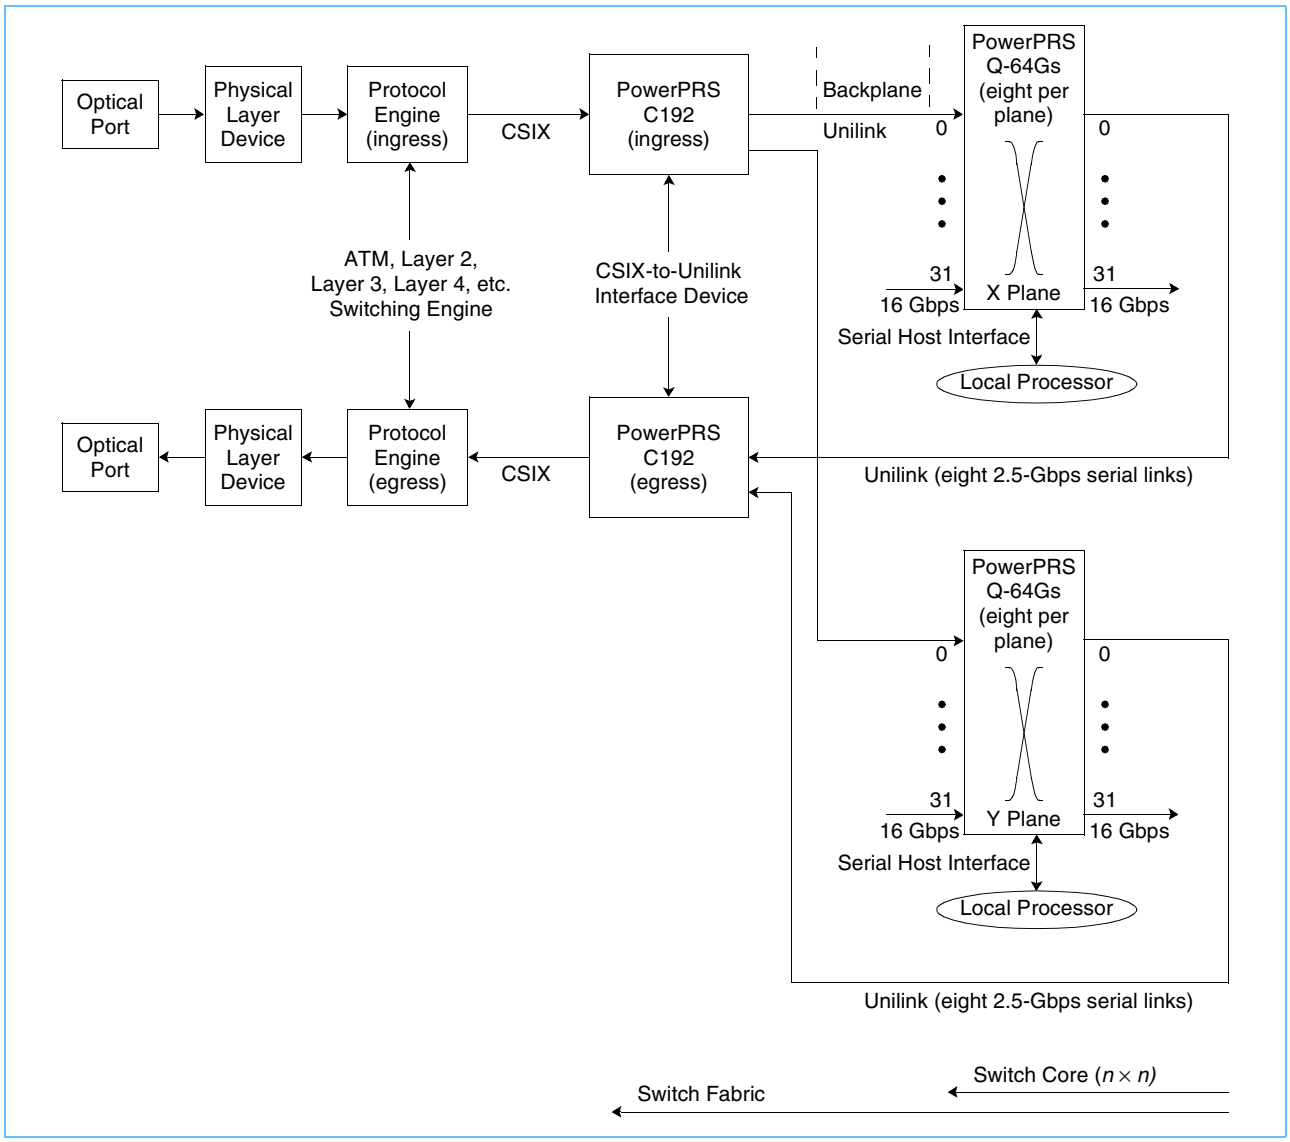

2.2 Internal Structure

The internal structure of the PowerPRS Q-64G is shown in *Figure 2-2,* which depicts a single PowerPRS Q-64G operating without speed expansion. *Figure 2-3* shows the structure of a subswitch element.

Figure 2-2. PowerPRS Q-64G Block Diagram

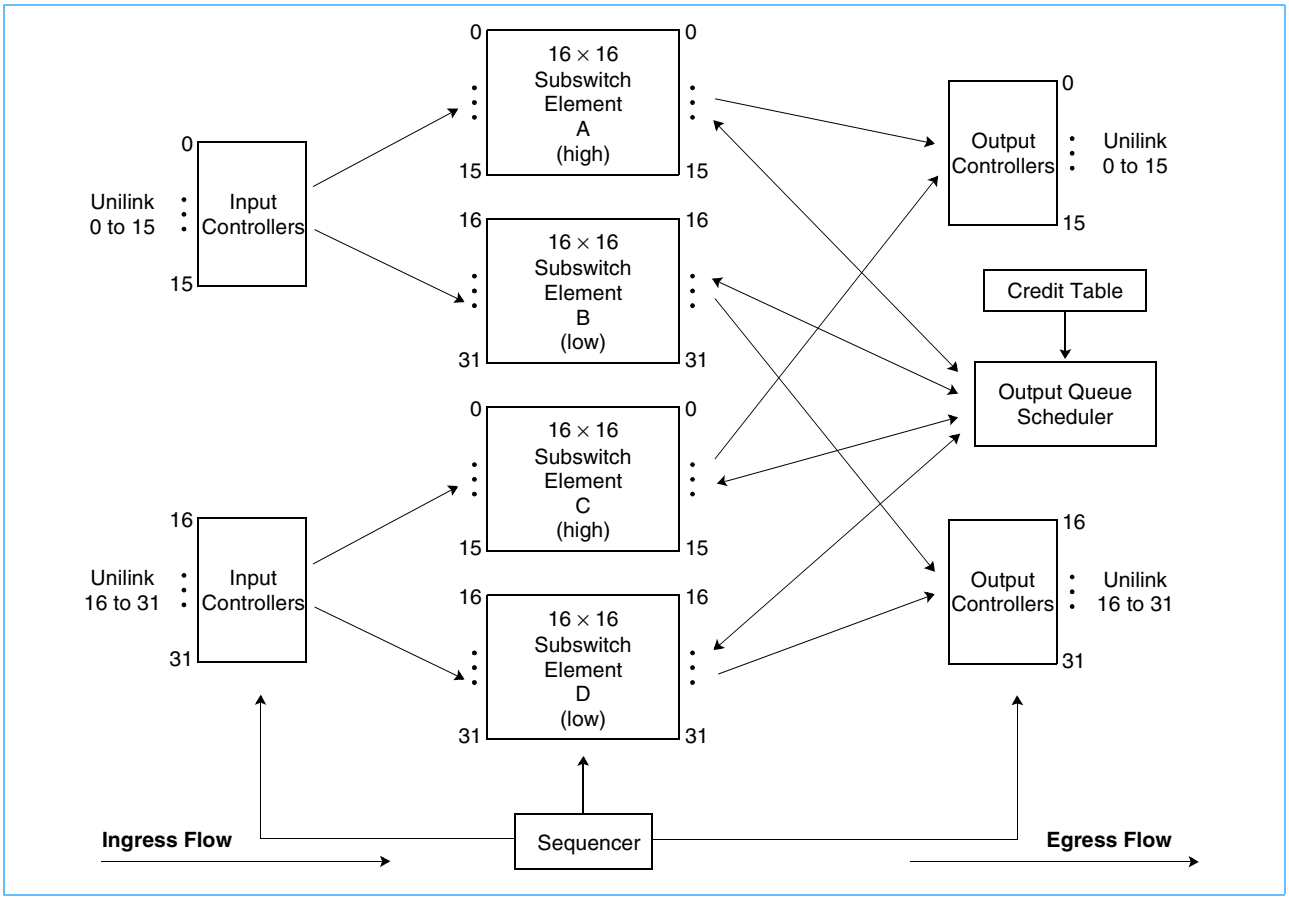

Figure 2-3. 16 × 16 Subswitch Element Block Diagram

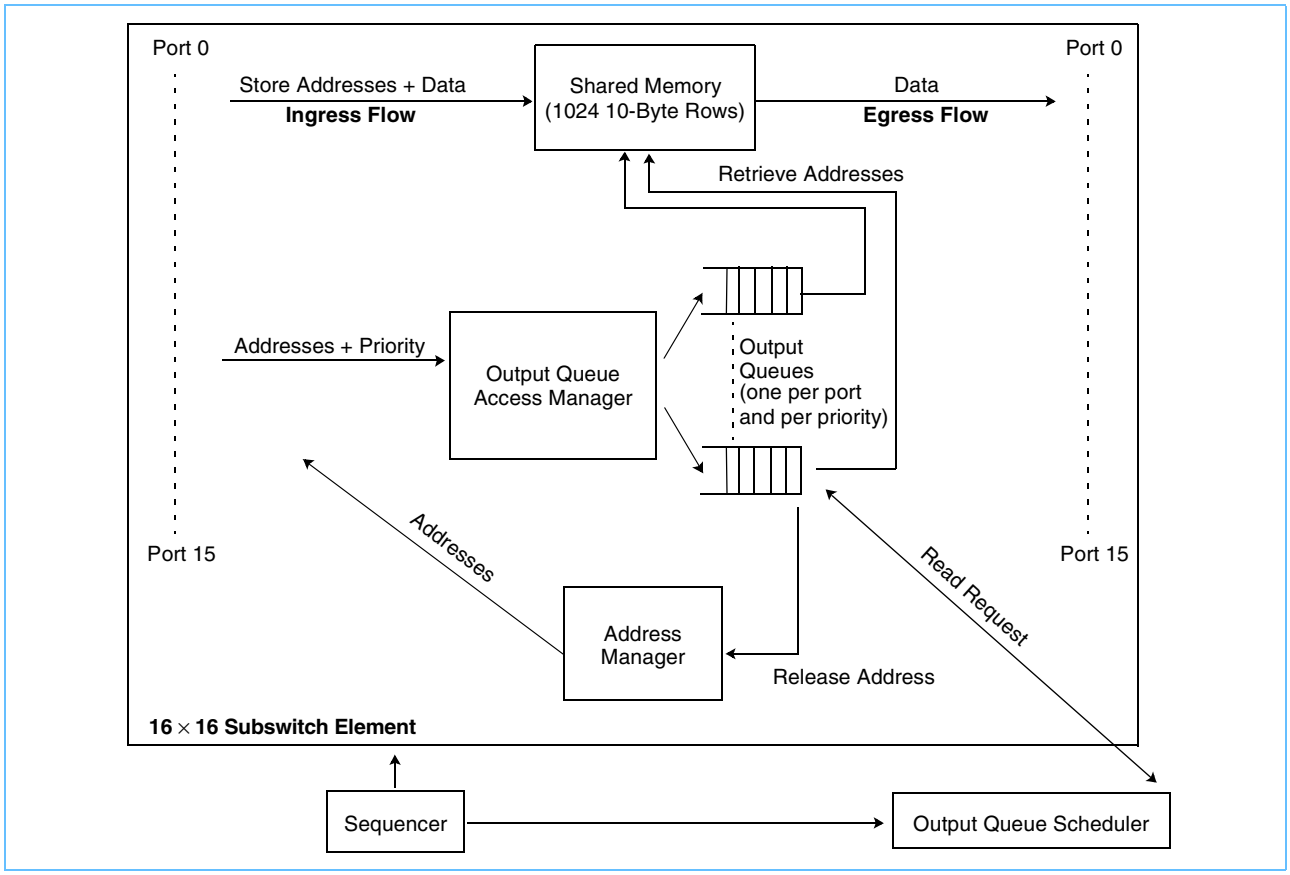

The main components of the PowerPRS Q-64G are:

- Thirty-two input controllers
- Thirty-two output controllers
- Four self-routing subswitch elements, each housing a shared memory bank and a control section comprised of:
  - One address manager
  - One output queue access manager
  - Sixteen output queues (one per output port)
- Device control section, which includes a:
  - Sequencer
  - Output queue scheduler
  - Credit table
- Unilink interface between the PowerPRS Q-64G and attached devices

2.2.1 Unilink Interface

The physical links between the PowerPRS Q-64G and attached devices are high-speed serial links called Unilinks. The Unilink physical interface minimizes the number of pins. There is one Unilink per PowerPRS Q-64G device port, or 32 Unilinks per device. Each Unilink is comprised of two pairs of differential

lines; one differential pair carries ingress flow and the other differential pair carries egress flow. Each pair of differential lines (that is, each device port) carries two packets and has a total throughput of 2 Gbps. Therefore, each Unilink carries two ingress packets and two egress packets at a time. The PowerPRS Q-64G physical interface carries the two ingress or two egress packets in separate streams; one packet is carried on the high channel and one packet is carried on the low channel. See *Section 3.2 Physical Interface and Packet Processing* on page 26 for more information.

2.2.2 Shared Memory

Each PowerPRS Q-64G device includes four 16 × 16 self-routing subswitch elements, denoted A, B, C, and D (see *Figure 2-2).* The four subswitch elements house the shared memory, which stores the packets that the PowerPRS Q-64G has received but has not yet transmitted. The shared memory on each subswitch element consists of 1024 10-byte rows, and has two read ports and two write ports.

Subswitch elements A and B store packets from input ports 0 to 15, and subswitch elements C and D store packets from input ports 16 to 31. For port expansion, each pair of subswitch elements is connected internally; that is, they are connected in parallel, in a single stage, to increase the number of ports without changing the port speed. This provides the 32 × 32-port device configuration. The *high* subswitch elements (A and C) store packets destined for output ports 0 to 15, and the *low* subswitch elements (B and D) store packets destined for output ports 16 to 31 (see *Figure 2-2* on page 16*).*

2.2.3 Sequencer

The sequencer controls the PowerPRS Q-64G internal data flow by granting shared memory access to the input and output ports. Sequencer operation is based on time-division multiplexing (TDM). The sequencer cycles concurrently among the input and output ports, granting shared memory access to two input ports and two output ports at a time (one from ports 0 to 15 and one from ports 16 to 31) and visiting each port once per cycle. During each shared memory access, one packet is transmitted to or from each of the port's two subswitch elements.

Packets are transmitted and stored in equal lengths called logical units (LUs). The standard PowerPRS Q-64G configurations include either four or eight devices, in which one device is the master and the rest of the devices are slaves. In these configurations, packets are divided into eight LUs (one master LU and seven slave LUs) and distributed over all the devices.

During each shared memory access, 8 to 10 bytes of data are processed (read or written) per subswitch element, depending on packet length. In the standard multiple-device configurations, processing an entire LU requires one shared memory access. The sequencer cycle equals the time required to process the data associated with one shared memory access. All sequencer cycles are equal in length.

The sequencer ensures that packets on a given port are always processed at a fixed interval according to their LU length; therefore, no synchronization is required between input ports. The slave device sequencers are synchronized to the master device sequencer so that all the LUs for a particular port (or packet) are processed at the same time. See *Section 2.3.3 Master/Slave Synchronization with Multiple Devices* on page 21 for more information.

2.2.4 Address Managers

Each subswitch element has an address manager that tracks the available shared memory addresses on the subswitch element and provides new store addresses to the input controllers. When an address manager provides a store address to an input controller, it removes that address from the available shared memory

address pool. After the packet is transmitted, the output queue returns the address to the address manager, which returns it to the available address pool. For multicast packets, one store address is sent to multiple output queues. The address manager tracks the number of output queues holding each store address and, when the count reaches zero, returns the address to the available shared memory address pool.

2.2.5 Input Controllers

The PowerPRS Q-64G has 32 input controllers, one input controller per port. Each input controller processes two packets at a time, one packet on the high channel and one packet on the low channel. When a packet arrives, the input controller of the master device extracts the header information (including packet priority and destination) from the master LU. It checks the master LU header integrity using a parity bit on the header bytes. If the packet is valid, the input controller stores it in the shared memory when access is granted by the sequencer.

An input controller stores a packet in the shared memory of one of its two subswitch elements, depending on the packet's destination. Packets stored in subswitch elements A and C are destined for output ports 0 to 15, and packets stored in subswitch elements B and D are destined for output ports 16 to 31 (see *Figure 2-2).* The input controller uses the store address provided by the address manager of the subswitch element. The input controller also forwards the shared memory address, packet priority, and packet destination to the output queue access manager. Packets arrive with a priority of 0 to 3, with 0 being the highest priority. Note that multicast packets have only one priority for all destinations.

In multiple-device configurations, the input controller on the master device forwards information such as the shared memory store address and subswitch element ID to the input controllers on the slave devices.

2.2.6 Output Queue Access Managers

Each subswitch element has an output queue access manager that receives the packet store address, priority, and destination from the input controllers and forwards the information to the output queues for the subswitch element. Each output queue access manager also maintains the counters that the PowerPRS Q-64G uses to control ingress traffic flow. For each output queue, a counter tracks the total number of packets enqueued for that output, regardless of priority. Another counter tracks the total number of packets stored in the shared memory, regardless of output or priority.

2.2.7 Output Queues

The output queues contain the shared memory addresses of packets awaiting transmission from the PowerPRS Q-64G. In each of the four subswitch elements, each of the 16 output ports has one output queue per priority. Each output queue is organized into two address banks: one bank holds addresses written by even ports and the other bank holds addresses written by odd ports. Packet addresses are organized in a first-in-first-out (FIFO) queuing structure in each address bank. Each output queue can store up to 1024 addresses.

A unicast packet address is stored in one output queue, and a multicast packet address is stored in two or more output queues.

2.2.8 Output Queue Scheduler and Credit Table

The output queue scheduler determines which output queue will provide the next egress packet retrieve address and notifies each selected output queue in turn. Each selected output queue then forwards its next retrieve address to the shared memory.

The output queue scheduler selects an output queue using several pieces of information. For each subswitch element output port, the output queues provide an output queue status (that is, output queue empty), one per priority, to the output queue scheduler. The output queue scheduler also receives the send grants that control egress traffic flow to the attached devices. In general, the output queue scheduler selects the highest-priority of the occupied output queues (so that high-priority packets overtake low-priority packets). However, a fixed amount of bandwidth can be assigned to low-priority packets by altering priority scheduling in the credit table (see *Section 3.5.3 Credit Table* on page 53).

Note: As discussed in *Section 2.2.7,* each output queue contains two address banks (those written by even ports and those written by odd ports). If both address banks are occupied for a single priority, the output queue scheduler toggles between the two banks to select and notify the entire output queue. Packet addresses are processed on a FIFO basis within an address bank.

2.2.9 Output Controllers

The PowerPRS Q-64G has 32 output controllers, one output controller per port. Each port transmits two packets at a time, one packet on the high channel and one packet on the low channel. When access is granted by the sequencer, the output controller retrieves the next two packets to be transmitted on a port from the retrieve addresses that the output queue scheduler forwarded to the shared memory. The output controller inserts ingress and subport flow control information (that is, it inserts the grants) into the packet header before forwarding the packets to the physical interface for serialization.

Note that each output controller retrieves packets from two 16×16 subswitch elements and merges the traffic to a single output port. Output controllers 0 to 15 transmit from subswitch elements A and C; output controllers 16 to 31 transmit from subswitch elements B and D (see *Figure 2-2* on page 16*).*

2.3 Multiple-Device Configurations

The PowerPRS Q-64G is designed to provide OC-192 attachment. To meet these transmission requirements and provide 16-Gbps throughput per port, multiple PowerPRS Q-64Gs are configured for speed expansion. There are two standard multiple-device configurations:

- 512-Gbps configuration (eight devices that provide 32 input and output ports)
- 256-Gbps configuration (four devices that provide 16 input and output ports)

In both of these multiple-device configurations, one device is the master and the rest of the devices are slaves. The master device performs packet routing and queueing and forwards packet synchronization and shared memory address information to the slave devices. Because the slave devices only store slave LUs and do not perform packet routing or queuing, their control sections are inactive to minimize power consumption.

Table 2-1 on page 21 presents some of the features, including the shared memory capacity, of the two standard multiple-device configurations. Descriptions of these two configurations follow.

Table 2-1. Multiple-Device Configuration Summary

Device Configuration	Number of Ports	Port Speed (Gbps)	Packet Length (bytes)	LU Size (bytes)	Shared Memory Capacity (packets)
512-Gbps configuration (eight devices with external speed expansion)	32 × 32	16	64, 72, or 80	8, 9, or 10	4096
256-Gbps configuration (four devices with internal and external speed expansion)	16 × 16	16	64, 72, or 80	8, 9, or 10	2048

2.3.1 512-Gbps Configuration

In the 512-Gbps configuration, eight PowerPRS Q-64Gs are configured for external speed expansion (see *Figure 2-4).* For external speed expansion, multiple devices are connected in parallel and the like-numbered ports on all the devices are grouped. The total number of ports remains the same as on a single device, but the throughput per port equals the throughput per port for a single device times the number of devices. A single PowerPRS Q-64G features 32 ports at 2 Gbps per port; therefore, this configuration provides 32 ports at 16 Gbps per port (for an aggregate throughput of 512 Gbps). In this configuration, the eight devices are assembled on two switch cards.

2.3.2 256-Gbps Configuration

In the 256-Gbps configuration, four PowerPRS Q-64Gs are configured for both internal and external speed expansion (see *Figure 2-5* on page 23*).* For internal speed expansion, two ports within a device are paired. This doubles the port speed but halves the number of ports. Because a single PowerPRS Q-64G features 32 ports at 2 Gbps per port, this configuration provides 16 ports at 16 Gbps per port (for an aggregate throughput of 256 Gbps). In this configuration, the four devices are assembled on one switch card.

2.3.3 Master/Slave Synchronization with Multiple Devices

2.3.3.1 Sequencers

Each PowerPRS Q-64G contains a sequencer. When multiple devices are configured for external speed expansion, the slave device sequencers must be synchronized to the master device sequencer to ensure that the LUs for a particular port (or packet) are processed at the same time on all the devices. This synchronization is done with the SyncIn (slave device input) and SyncOut (master device output) pins. The SyncIn/Out pin mode bit in the *Configuration 1 Register* (page 112) sets the operating mode for these pins.

Note: In the 256-Gbps and 512-Gbps configurations, LU (and packet) transmission requires one sequencer cycle.

Figure 2-4. 512-Gbps Configuration (external speed expansion)

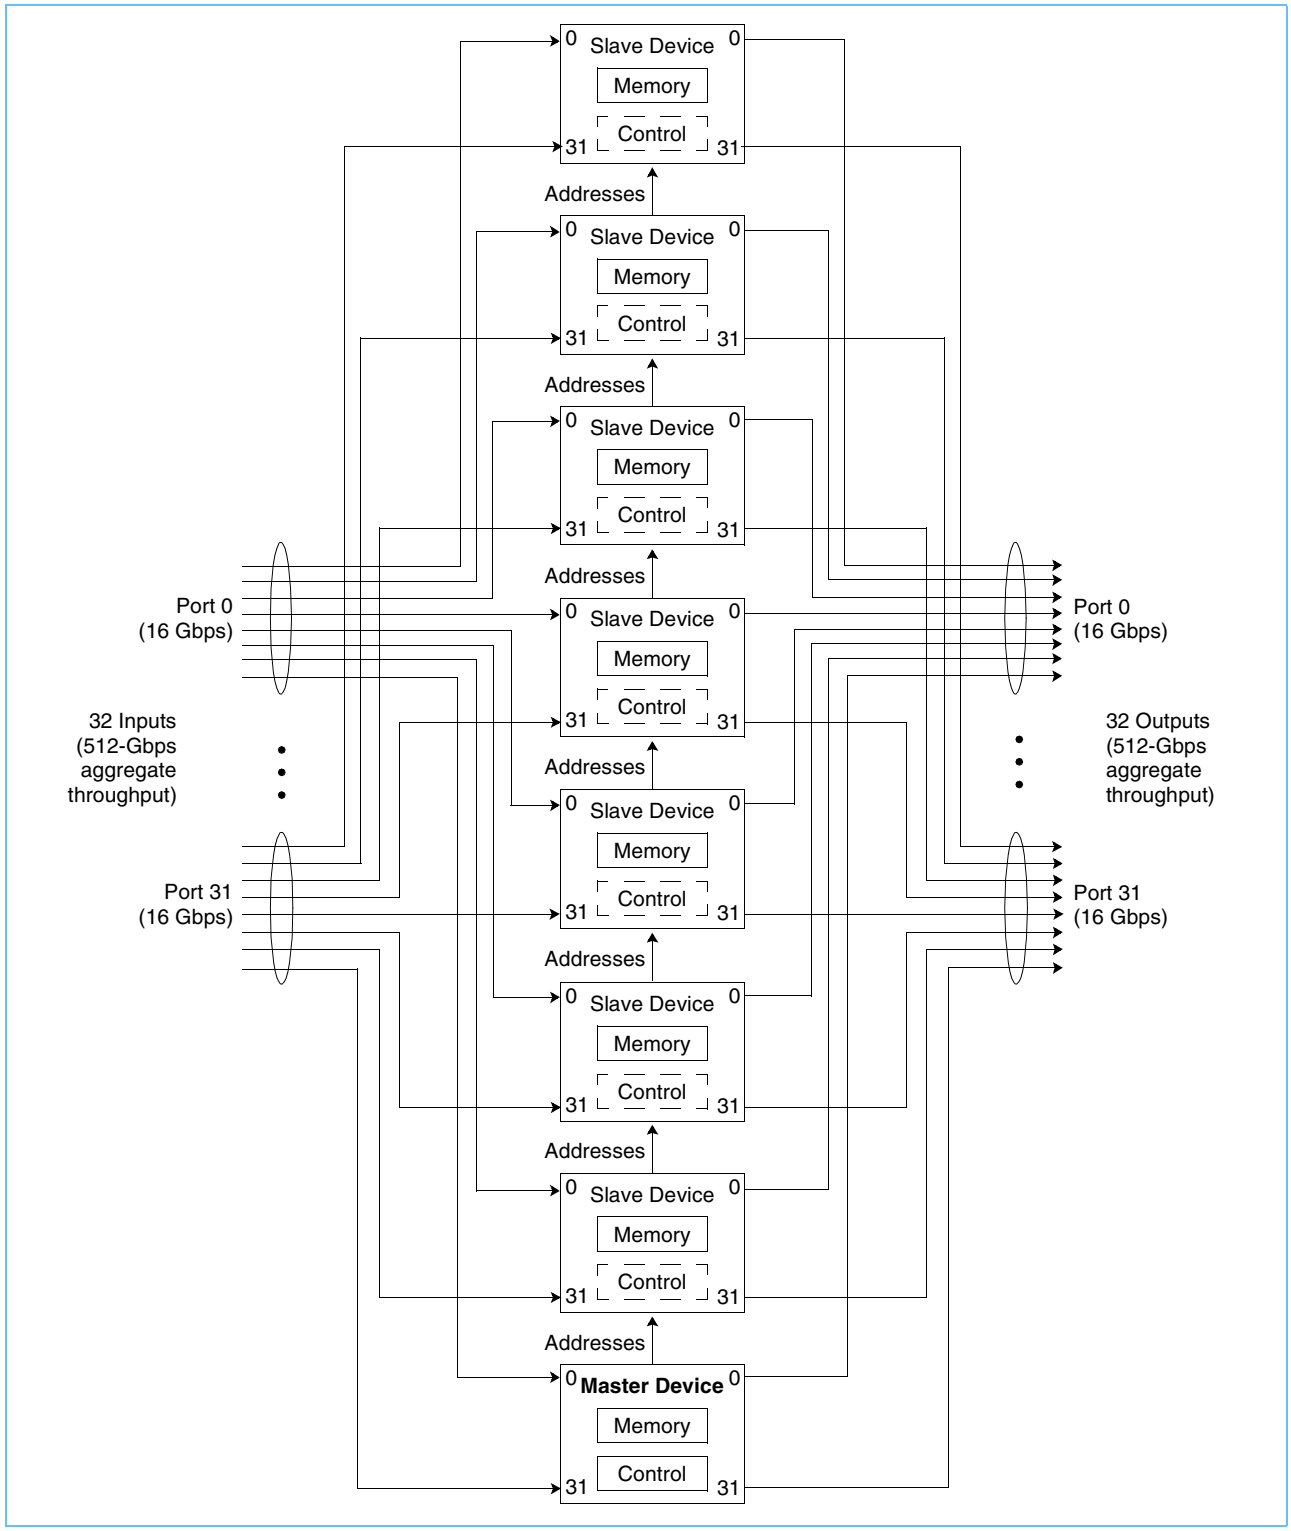

Figure 2-5. 256-Gbps Configuration (internal and external speed expansion)

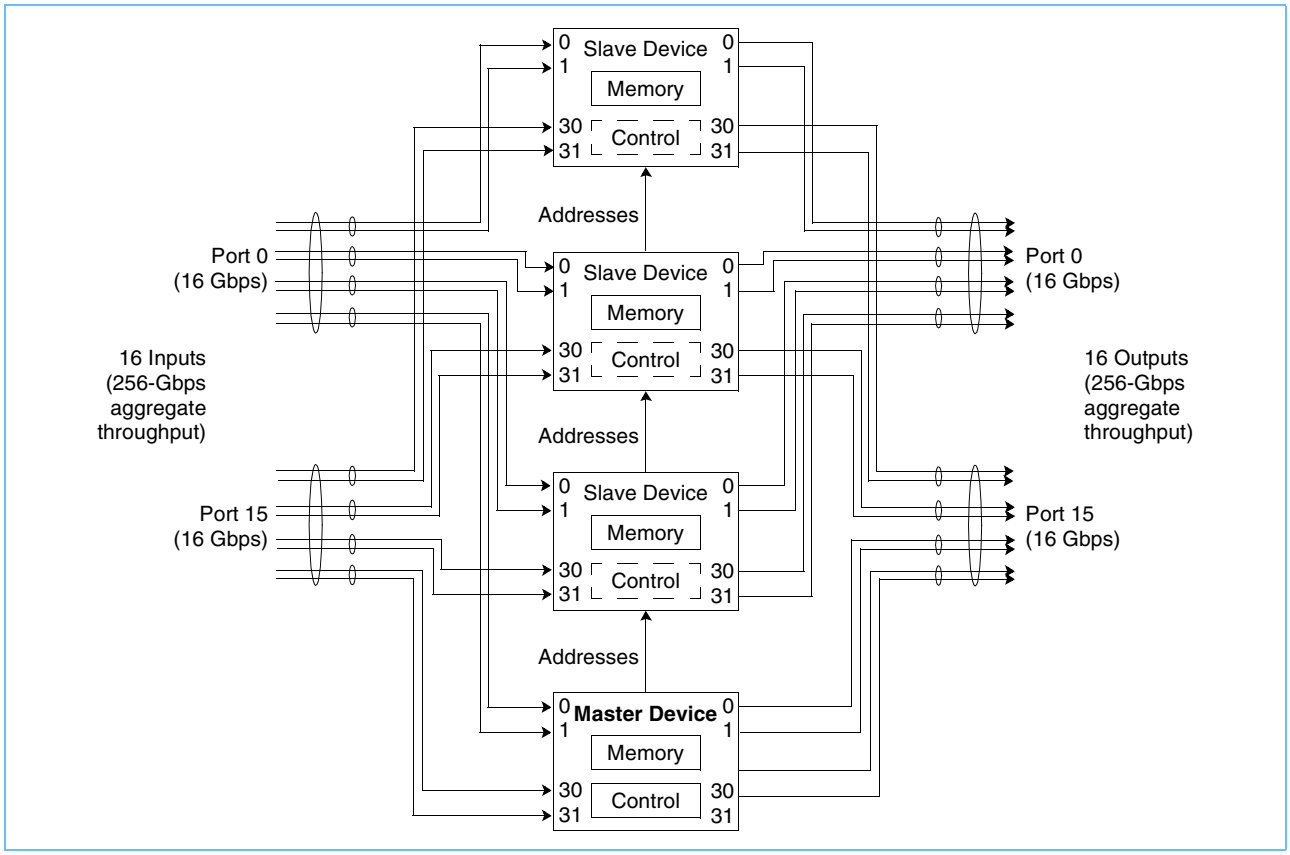

2.3.3.2 Shared Memory Addresses

When multiple devices are configured for external speed expansion, the input and output controllers of the master device forward shared memory addresses to the input and output controllers of the slave devices. The master device conveys this information via two speed-expansion buses:

- *SpexDataIn* is the speed-expansion bus that enters each slave device. It is comprised of eight Unilinks for ingress addresses and eight Unilinks for egress addresses.
- *SpexDataOut* is the speed-expansion bus that exits the master device and all but the last slave device. It is also comprised of eight Unilinks for ingress addresses and eight Unilinks for egress addresses.

The devices are connected serially to the speed-expansion buses. The master device generates addresses and provides them to the first slave device. The first slave device then conveys the addresses to the next slave device, and so forth, until the last slave device receives the addresses.

3. Functional Description

This section describes basic PowerPRS Q-64G functionality, including information about:

- Packet type
- Physical interface and packet processing
- Packet format according to packet type
- Ingress, egress, and subport flow control
- Packet reception and transmission
- Side communication channel
- Switchover support

3.1 Packet Types

There are four types of packets:

- Data packets
- Control packets
- Service packets
- Idle packets

All the packet types can be carried on either the high channel or the low channel. These two channels simultaneously carry one packet apiece, and the packets can be different types.

Note: The grants used for ingress, egress, and subport flow control are carried in the header and payload of ingress and egress packets. See *Section 3.3 Packet Format According to Packet Type* on page 28 for more information.

3.1.1 Data Packets

Data packets contain user data to be switched from an input to one or more outputs. Data packets have a priority of 0, 1, 2, or 3, with 0 being the highest priority. In the basic configuration, the PowerPRS Q-64G output queue scheduler prioritizes packet transmission for each output port in the following order:

1. Service packets
2. Control packets
3. Priority 0 data packets
4. Priority 1 data packets
5. Priority 2 data packets
6. Priority 3 data packets

Besides the packet priority, data packets also carry routing information (destination bitmap), filtering information used for switchover support (color coding), and a “best-effort discard” flag.

3.1.2 Control Packets

Control packets carry the communications between the local processor and the protocol engine. They do not have a specific priority. Control packet transmission is relatively infrequent and does not affect the performance of high-priority traffic because the local processor access is slow compared to the data packet traffic rate.

Ingress control packets originate at the protocol engine. The PowerPRS Q-64G can receive control packets on either the high channel or the low channel of any input port. An input controller identifies an ingress packet as a control packet when the destination bitmap value is all zeros. Ingress control packets are stored in the shared memory by the input controller and then transferred to the local processor using registers (see *Section 3.8.2 Control and Service Packet Reception* on page 55).

Egress control packets originate at the local processor. The local processor can transmit control packets on any output port. Control packets are always transmitted on an output before any other packets stored in the shared memory destined for that output. Egress control packets are transferred from the local processor to output controllers using registers (see *Section 3.9.2 Control and Service Packet Transmission* on page 56).

3.1.3 Service Packets

Service packets carry the communications between the local processor and the attached devices. They are used to test the continuity of the links between the PowerPRS Q-64G and the attached devices and to gain access to attached device internal resources.

There are three types of service packets: event-1, event-2, and command. Only command service packets contain payload. Service packets are received and transmitted in the same manner as control packets (that is, ingress packets are stored by an input controller and exchanged with the local processor using registers).

Note: When the PowerPRS Q-64G is attached to the PowerPRS C192, event-1 service packets are used for link liveness, event-2 service packets are used for write acknowledgement and switchover, and command service packets are used for gaining read access to PowerPRS C192 resources.

3.1.4 Idle Packets

Idle packets do not carry user data. They are transmitted on a port only when there are no data, control, or service packets available for transmission or when these packets cannot be transmitted (for example, during switchover). Idle packets are also used to perform link synchronization. The attached devices generate *ingress* idle packets; the PowerPRS Q-64G generates *egress* idle packets. Idle packet color coding is used for switchover support.

3.2 Physical Interface and Packet Processing

As discussed in *Section 2.2.1,* the physical links between the PowerPRS Q-64G and attached devices are high-speed serial links called Unilinks. There is one Unilink per PowerPRS Q-64G device port, or 32 Unilinks per device. Each Unilink is comprised of two pairs of differential lines; one differential pair carries ingress flow and the other differential pair carries egress flow. Each pair of differential lines (that is, each device port) carries two packets at a time.

Packets are transmitted in equal lengths called logical units (LUs). The number of LUs depends on the device configuration. The LU bytes of two packets carried over a Unilink are serialized in the following order:

- LU for packet 0, byte 0
- LU for packet 1, byte 0
- LU for packet 0, byte 1
- LU for packet 1, byte 1

This pattern continues through the final two LU bytes, which are serialized as follows:

- LU for packet 0, byte (LU length - 1)
- LU for packet 1, byte (LU length - 1)

Each Unilink differential pair operates at 2.5 Gbps. The Unilink data stream is a single 400-ps bit stream that carries 8b/10b encoding for link synchronization and supervision.

On the ingress path, the PowerPRS Q-64G physical interface deserializes the Unilink data stream into a 10-bit, 4-ns data stream. The physical interface performs byte alignment and removes the 8b/10b code, translating the 10-bit stream into an 8-bit stream. It also demultiplexes the 4-ns data stream into two 8-ns (125 MBps) byte streams, one byte stream for each LU. The two byte streams are carried over separate buses to the input controller. The high-channel bus carries the LU for packet 0 and the low-channel bus carries the LU for packet 1. Although the Unilink physical throughput is 2.5 Gbps, the Unilink logical throughput is only 2 Gbps because the Unilink data stream carries 8b/10b encoding.

On the egress path, this process is reversed. The PowerPRS Q-64G physical interface multiplexes the two 8-ns byte streams into a single 4-ns byte stream, adds 8b/10b encoding, and serializes the data into a single 400-ps bit stream.

In the standard 256-Gbps or 512-Gbps PowerPRS Q-64G configuration, each 16-Gbps port is comprised of eight Unilinks. Each packet is divided into eight LUs, and each Unilink carries one LU from each packet. The first LU of a packet is the master, and the other LUs are the slaves. The master LU carries packet header bytes, which contain packet control information and precede packet payload bytes. Slave LUs carry only payload bytes. The LU length is either 8, 9, or 10 bytes depending on the device configuration.

Figure 3-1 illustrates how the master and slave LUs of two data packets are distributed over the eight Unilinks that comprise a port. The master Unilink carries the master LUs of both packets. Each slave Unilink carries one slave LU from each of the two packets.

Because PowerPRS Q-64G data is transmitted with a known clock, only bit phase alignment and packet delineation must be performed. All eight packet LUs are received or transmitted on all eight Unilink ports simultaneously. All eight LUs must arrive at the device port input pins within a 4-ns time frame; that is, the difference between the time the first bit of the first LU arrives at its input pin and the time the first bit of the last LU arrives at its input pin must be within 4 ns. External software running on the local processor compensates for any skew between the multiple Unilink ports.

When an input port simultaneously receives two data packets destined for the same output port, the output queue stores the address of the packet received on the high channel before it stores the address of the packet received on the low channel because, internally, the packet received on the high channel is considered to have arrived first. Similarly, when an output port transmits two data packets simultaneously, the attached device processes the packet carried on the high channel first to ensure correct packet serialization through the switch. The LUs of successive packets are transported sequentially, with no gap between them.

Figure 3-1. Packet Format for a 16-Gbps Port

		Byte 0	Byte 1	Byte 2	Byte 3	Byte 4	Byte 5	Byte 6	Byte 7	Byte 8	Byte 9
Master Unilink	Packet 0, LU 0	H0	H1	H2/D	D	D	D	D	D	D	D
	Packet 1, LU 0	H0	H1	H2/D	D	D	D	D	D	D	D
Slave 1 Unilink	Packet 0, LU 1	D	D	D	D	D	D	D	D	D	D
	Packet 1, LU 1	D	D	D	D	D	D	D	D	D	D
•											
•											
•											
Slave 7 Unilink	Packet 0, LU 7	D	D	D	D	D	D	D	D	D	D
	Packet 1, LU 7	D	D	D	D	D	D	D	D	D	D

Notes:
1. Packet 0 is carried on the high channel and packet 1 is carried on the low channel.
2. There are eight logical units (LUs) per packet, and LU 0 is the master LU.
3. H0, H1, and H2 are packet header bytes; H2 exists only in the 512-Gbps configuration.
4. D is user data.

3.3 Packet Format According to Packet Type

3.3.1 General Packet Format Information

3.3.1.1 Packet Header

The master LU carries the packet header, which is comprised of the packet qualifier byte (the first byte, H0) and either one byte (H1, in the 256-Gbps configuration) or two bytes (H1 and H2, in the 512-Gbps configuration) of additional information. Depending on the packet type, the packet qualifier byte may contain information about:

- Packet type
- Packet priority
- Header parity
- Packet color (for switchover support)
- Packet filtering information (for switchover support)
- Best-effort discard
- Extended bitmap (for packet routing)
- Flow control flywheels

The packet qualifier byte can also include reserved bits. Reserved bits pass through the device unmodified. Unless otherwise specified, reserved bits must be set to ‘0’.

For all packets, odd parity is calculated on the entire packet header (H0 and H1 or H0 through H2, depending on the configuration), including reserved bits. Parity calculation always includes the number of bytes defined by the header length, even when a header byte does not contain any information. Because this is an odd-parity device, the parity bit is set to '1' when the packet header byte calculation results in an odd number of '1' bits. Parity checking ensures that the packet header is valid, and ignores the additional information carried in idle packet headers.

Idle packet header bytes (H1 and H2) contain flow control flywheel information, and ingress data packet header bytes (H1 and H2) contain the destination bitmap. Flow control grants are also carried in the headers of data, control, service, and idle packets, as well as in byte 6 of idle packets. Byte 5 of idle packets carries side communication channel (SCC) information.

Note: A flow control grant is active when the assigned bit is set to '1' in the packet header. See *Section 3.4 Ingress Flow Control* on page 47, *Section 3.5 Egress Flow Control* on page 52, and *Section 3.6 Subport Flow Control* on page 53 for more information about flow control grants.

3.3.1.2 Flow Control Flywheels

The grants used for ingress, egress, and multicast flow control are carried in data, control, service, and idle packet headers. Flow control flywheels determine which grants are carried during each ingress or egress packet cycle because multiple packet cycles are required to receive or transmit all the grants.

For example, when PowerPRS Q-64G output controllers insert flow control grants into the egress packet headers, four internal flywheels determine which grants are carried during each egress packet cycle. To extract the correct information from the egress packet headers, the attached devices contain corresponding flywheels synchronized to those in the PowerPRS Q-64G. Egress idle packet headers carry the PowerPRS Q-64G flywheel status, which is used to synchronize the attached device flywheels to those in the PowerPRS Q-64G. See *Section 3.3.6 Flow Control Flywheels for Grants Carried in Egress Packets* on page 36 for more information.

Similarly, two attached device flywheels determine which grants are carried during each ingress packet cycle. To extract the correct information from the ingress packet headers, the PowerPRS Q-64G contains corresponding flywheels synchronized to those in the attached devices. The attached device flywheel status used for flywheel synchronization is carried in the ingress idle packet headers. See *Section 3.3.2 Flow Control Flywheels for Grants Carried in Ingress Packets* (below) for more information.

Note: Each PowerPRS Q-64G port has a set of flywheels synchronized to a corresponding set of flywheels on its attached device. No flywheel synchronization exists between ports. The flow control flywheels increment according to the number of data packet priorities enabled in the number of priorities field in the *Configuration 1 Register* (page 112).

3.3.2 Flow Control Flywheels for Grants Carried in Ingress Packets

There are two flow control flywheels associated with the grants carried in ingress packet headers:

- Subport grant type/subport flywheel
- Grant priority flywheel

The subport grant type/subport flywheel and the grant priority flywheel are used to extract grants from ingress data packet and control packet headers, which carry one subport grant and one send grant for the port per packet cycle.

These two flywheels do not apply to ingress idle packets, each of which has the capacity to carry all 20 subport grants and up to 4 send grants for the port. When an idle packet is transmitted, these flywheels continue to increment. If a port receives only data packets and control packets and all four priorities are enabled, 20 packet cycles are required to update the complete set of subport grants and 4 packet cycles are required to update the complete set of send grants.

3.3.2.1 Subport Grant Type/Subport Flywheel

The subport grant type/subport flywheel determines the type of subport grant carried by the current ingress packet cycle. This flywheel cycles as follows:

000	Subport output queue grants, subport 0
001	Subport output queue grants, subport 1
010	Subport output queue grants, subport 2
011	Subport output queue grants, subport 3
100	Subport multicast grants

For each of the above subport grant types, there is one grant per priority (20 subport grants total per port, when all four priorities are enabled). The subport grant carried by the current packet cycle is determined by the subport grant type/subport flywheel value in combination with the grant priority value. The subport grant type/subport flywheel is incremented each time the grant priority flywheel returns to ‘000’. The grant priority flywheel returns to ‘000’ after reaching ‘100’.

To synchronize the PowerPRS Q-64G flywheels with those in the attached devices, the subport grant type/subport flywheel status and grant priority flywheel status are inserted into the packet qualifier byte (H0) of ingress idle packets. The subport grants are inserted into header bytes H1 and H2 and byte 6 of ingress idle packets and into byte H0 of ingress data and control packets. See the following packet format descriptions for more information.

3.3.2.2 Grant Priority Flywheel

The grant priority flywheel determines the priority of the grant carried by the current ingress packet cycle:

00	Priority 0
01	Priority 1
10	Priority 2
11	Priority 3

This priority is used for both the send grant and the subport grant. There is only one send grant per priority. The grant priority flywheel is incremented every packet cycle, and returns to ‘00’ after reaching the number of priorities defined in the number of priorities field in the *Configuration 1 Register* (page 112).

3.3.3 Ingress Idle Packet Format

The ingress idle packet format for the PowerPRS Q-64G 16-Gbps configuration is presented in *Figure 3-1*. Bytes 0:2 and bytes 5:6 of the master LU carry control information, as described in *Tables 3-1* through *3-3*. Ingress and egress idle packets both carry SCC information. See *Section 3.10 Side Communication Channel* on page 57 for more information.

Figure 3-2. Ingress Idle Packet Format

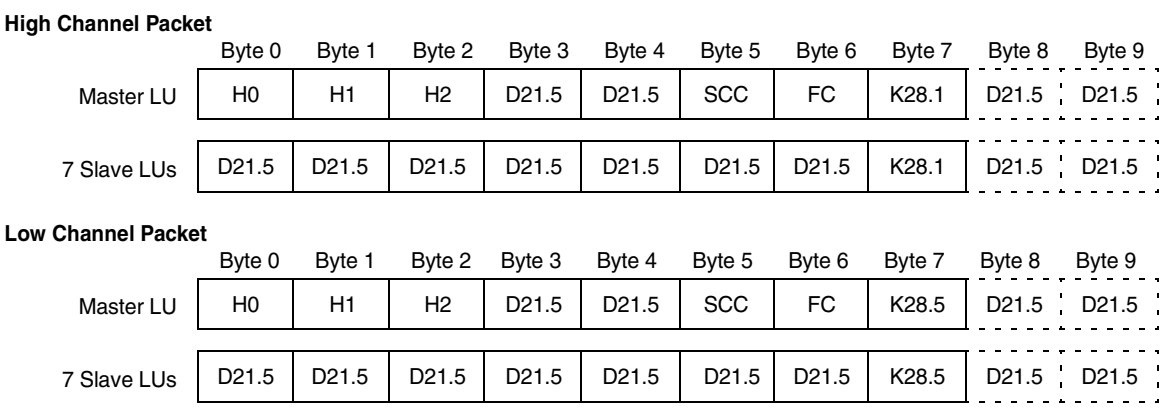

Notes:

1. H0, H1, and H2 are packet header bytes.
2. D21.5, K28.1, and K28.5 characters are 8b/10b control characters. K28.1 and K28.5 characters are used for link synchronization (byte alignment and packet clock recovery) and supervision (byte alignment and synchronization lost). D21.5 characters guarantee continuous transition on the line.
3. SCC is side communication channel data. Bits 4:7 are copies of bits 0:3.
4. FC is flow control information in addition to that transmitted in the packet header.

Table 3-1. Ingress Idle Packet, Byte H0

Byte	Information Carried				
	Bit 0	Bit 1	Bits 2:3	Bits 4:5	Bits 6:7
H0 (high channel)	Subport grant type/ subport flywheel status (bit 0 of three bits)	Parity	Protection	Color	Subport grant type/ subport flywheel status (bits 1:2 of three bits)
H0 (low channel)	Reserved	Parity	Protection	Color	Grant priority flywheel status

Table 3-2. Ingress Idle Packet, Byte H0 Field Descriptions (Page 1 of 2)

Bit(s)	Field Name	Description
0	Subport Grant Type/Subport Flywheel Status (high channel)	See the description for bits 6:7.
	Reserved (low channel)	Reserved.
1	Parity	Header parity.
2:3	Protection	Must be set to ‘00’ for idle packets.

Note: On the high channel, bits 0, 6, and 7 of byte H0 carry the three-bit subport grant type/subport flywheel status. On the low channel, bit 0 is reserved, and bits 6:7 carry the subport grant type/subport flywheel status.

Table 3-2. Ingress Idle Packet, Byte H0 Field Descriptions (Page 2 of 2)

Bit(s)	Field Name	Description
4:5	Color	Identifies the idle packet color for switchover support: 00 Blue 01 Red Others Reserved
6:7	Subport Grant Type/Subport Flywheel Status (high channel)	Reports the status of the subport grant type/subport flywheel: 000 Subport output queue grants, subport 0 001 Subport output queue grants, subport 1 010 Subport output queue grants, subport 2 011 Subport output queue grants, subport 3 100 Subport multicast grants See *Section 3.3.2.1 Subport Grant Type/Subport Flywheel* on page 30 for more information.
	Grant Priority Flywheel Status (low channel)	Reports the status of the grant priority flywheel: 00 Priority 0 01 Priority 1 10 Priority 2 11 Priority 3 See *Section 3.3.2.2 Grant Priority Flywheel* on page 30 for more information.
Note: On the high channel, bits 0, 6, and 7 of byte H0 carry the three-bit subport grant type/subport flywheel status. On the low channel, bit 0 is reserved, and bits 6:7 carry the subport grant type/subport flywheel status.		

Table 3-3. Ingress Idle Packet, Bytes H1, H2, SCC, and FC

Byte	Information Carried							
	Bit 0	Bit 1	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7
H1 (high and low channel)	Subport output queue grant: subport 0, priority 0	Subport output queue grant: subport 0, priority 1	Subport output queue grant: subport 0, priority 2	Subport output queue grant: subport 0, priority 3	Subport output queue grant: subport 2, priority 0	Subport output queue grant: subport 2, priority 1	Subport output queue grant: subport 2, priority 2	Subport output queue grant: subport 2, priority 3
H2 (high and low channel)	Subport output queue grant: subport 1, priority 0	Subport output queue grant: subport 1, priority 1	Subport output queue grant: subport 1, priority 2	Subport output queue grant: subport 1, priority 3	Subport output queue grant: subport 3, priority 0	Subport output queue grant: subport 3, priority 1	Subport output queue grant: subport 3, priority 2	Subport output queue grant: subport 3, priority 3
SCC (high and low channel)	SSC data: bit 0	SCC data: bit 1	SCC data: bit 2	SCC data: bit 3	SCC data: bit 0	SCC data: bit 1	SCC data: bit 2	SCC data: bit 3
FC (high and low channel)	Subport multicast grant: priority 0	Subport multicast grant: priority 1	Subport multicast grant: priority 2	Subport multicast grant: priority 3	Send grant: priority 0	Send grant: priority 1	Send grant: priority 2	Send grant: priority 3

3.3.4 Ingress Data Packet and Control Packet Format

The ingress data packet and control packet format for the PowerPRS Q-64G 16-Gbps configuration is presented in *Figure 3-3*. The packet qualifier byte (H0) is described in *Tables 3-4* and *3-5*.

Figure 3-3. Ingress Data Packet and Control Packet Format

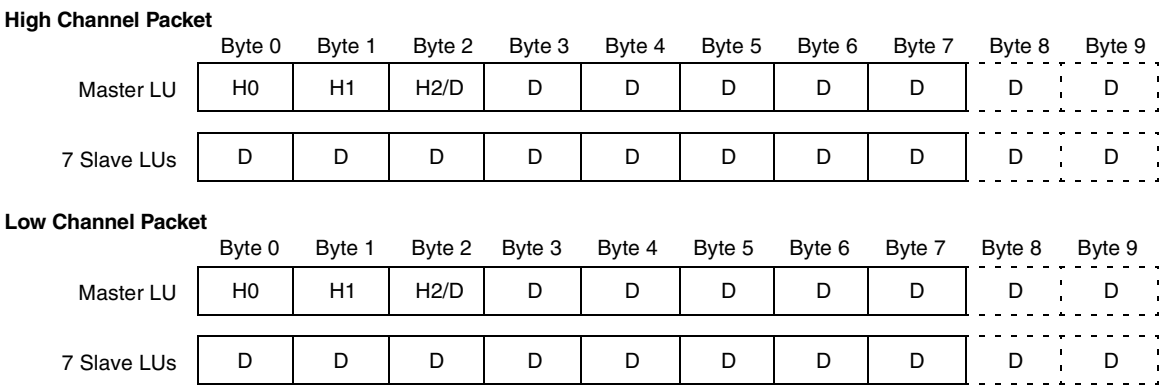

Table 3-4. Ingress Data Packet and Control Packet, Byte H0

Byte	Information Carried					
	Bit 0	Bit 1	Bits 2:3	Bit 4	Bit 5	Bits 6:7
H0 (high channel)	Extended bitmap	Parity	Protection	Best effort	Send grant	Packet priority
H0 (low channel)	Extended bitmap	Parity	Protection	Best effort	Subport grant	Packet priority

Table 3-5. Ingress Data Packet and Control Packet, Byte H0 Field Descriptions

Bit(s)	Field Name	Description
0	Extended Bitmap	Designates the range of output ports addressed by the destination bitmap in the data packet header: 256-Gbps Configuration: 0 Ports 0 to 7; 1 Ports 7 and 15 512-Gbps Configuration: 0 Ports 0to 15; 1 Ports 15 to 31 Output port mapping for the 256-Gbps configuration is detailed in *Table 3-6,* and output port mapping for the 512-Gbps configuration is detailed in *Table 3-7*. Note that each ingress packet header addresses only half the output ports. This reduces the destination bitmap field length and increases the packet payload size. Because only half the ports are addressed in a single packet cycle, broadcast operation requires two packet cycles. A multicast packet requires either one or two packet cycles, depending on whether the active bits are distributed over half the ports or all the ports.
1	Parity	Header parity.
2:3	Protection	Specifies the traffic type (red data packet, blue data packet, or idle packet) and the application method of the bitmap filter to the ingress packet destination bitmap. The bitmap filter is specified with the *Bitmap Filter Register* (page 125). The PowerPRS Q-64G uses the resulting masked destination bitmap to route the packet. If the resulting bitmap is all zeros, the PowerPRS Q-64G ignores the packet. Note that control packets are detected before the bitmap filter is applied. Protection Field / Color / Bitmap Filter 00 / Not applicable / Not applicable (packet is an idle packet or a service packet). 01 / Red (backup) / Packet destination bitmap is bitwise ANDed with bitwise complement of bitmap filter. 10 / Red (active) / Packet destination bitmap is bitwise ANDed with bitmap filter. 11 / Blue (unfiltered) / Packet destination bitmap is used as is (unfiltered). This filtering function supports switchover and load balancing (see *Section 3.11 Switchover Support* on page 57).
4	Best Effort	Flags traffic as best-effort bandwidth traffic when the best-effort discard function is enabled. 1 The packet is flagged as best-effort bandwidth traffic, and the PowerPRS Q-64G can discard it, if necessary. 0 The packet is flagged as guaranteed bandwidth traffic, and the PowerPRS Q-64G cannot discard it. When the best-effort discard function is enabled, excess output port congestion triggers a mechanism that discards best-effort bandwidth traffic to provide output port access to guaranteed bandwidth traffic. The best-effort discard function is enabled with the best-effort discard enable bit in the *Configuration 0 Register* (page 110). See *Section 3.4.6 Best-Effort Discard* on page 49 for more information.
5	Send Grant	Reports the status of the send grant for the priority specified by the grant priority flywheel.
	Subport Grant	Reports the status of the subport grant of the type specified by the subport grant type/subport flywheel and the priority specified by the grant priority flywheel.
6:7	Packet Priority	Specifies the packet priority for data packets: 00 Priority 0 (highest priority) 01 Priority 1 10 Priority 2 11 Priority 3 (lowest priority) This field is ignored for control packets because control packets do not have a specific priority (see *Section 3.1.2* on page 26).

For data packets, the protection field, best-effort bit, and packet priority field of the packet qualifier byte (H0) are processed without modification. The best-effort bit is ignored unless the best-effort discard function is enabled (see *Section 3.4.6 Best-Effort Discard* on page 49). The protection field and packet priority field are always processed by the PowerPRS Q-64G, and the user must set them to the appropriate values.

For ingress data packets, header bytes H1 and H2 contain the packet destination bitmap, which designates the output ports to which the packet is destined and, thereby, the output queues into which the packet will be enqueued (see *Tables 3-6* and *3-7).* For the PowerPRS Q-64G, the destination bitmap is a logical bitmap. Each logical port can be mapped to any physical port of the same subswitch using the *Bitmap Mapping Register* (page 128). For example, logical bitmap 0 can address physical port 1 (rather than physical port 0). When a destination bitmap bit is set to '1', the packet is routed to the corresponding logical port. The bitmap field can point to multiple output ports. If the destination bitmap value is all zeros, the packet is recognized as a control packet.

Table 3-6. Ingress Data Packet and Control Packet, Byte H1 in the 256-Gbps Configuration

Byte	Packet Destination							
	Bit 0	Bit 1	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7
H1 (high or low channel) for extended bitmap (byte H0) = 0	Output port 0	Output port 1	Output port 2	Output port 3	Output port 4	Output port 5	Output port 6	Output port 7
H1 (high or low channel) for extended bitmap (byte H0) = 1	Output port 8	Output port 9	Output port 10	Output port 11	Output port 12	Output port 13	Output port 14	Output port 15

In the 256-Gbps configuration, four PowerPRS Q-64Gs are configured for internal speed expansion, in which the device ports are paired as follows: physical port 0 is paired with physical port 1 to form logical port 0, physical port 2 is paired with physical port 3 to form logical port 1, and so forth. For packet routing, if bit *n* of the packet destination bitmap is set to '1', the packet is routed to physical output ports ($n \times 2$) and [($n \times 2$) + 1]. For example, if bit 3 of the destination bitmap is set to '1', the packet is routed to physical output ports 6 and 7.

Table 3-7. Ingress Data Packet and Control Packet, Bytes H1 and H2 in the 512-Gbps Configuration

Byte	Packet Destination							
	Bit 0	Bit 1	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7
H1 (high or low channel) for extended bitmap (byte H0) = 0	Output port 0	Output port 1	Output port 2	Output port 3	Output port 4	Output port 5	Output port 6	Output port 7
H2 (high or low channel) for extended bitmap (byte H0) = 0	Output port 8	Output port 9	Output port 10	Output port 11	Output port 12	Output port 13	Output port 14	Output port 15
H1 (high or low channel) for extended bitmap (byte H0) = 1	Output port 16	Output port 17	Output port 18	Output port 19	Output port 20	Output port 21	Output port 22	Output port 23
H2 (high or low channel) for extended bitmap (byte H0) = 1	Output port 24	Output port 25	Output port 26	Output port 27	Output port 28	Output port 29	Output port 30	Output port 31

3.3.5 Ingress Service Packet Format

The ingress service packet format for the PowerPRS Q-64G 16-Gbps configuration is presented in *Figure 3-7*. The packet qualifier byte (H0) is described in *Tables 3-8* and *3-9.*

Figure 3-4. Ingress Service Packet Format

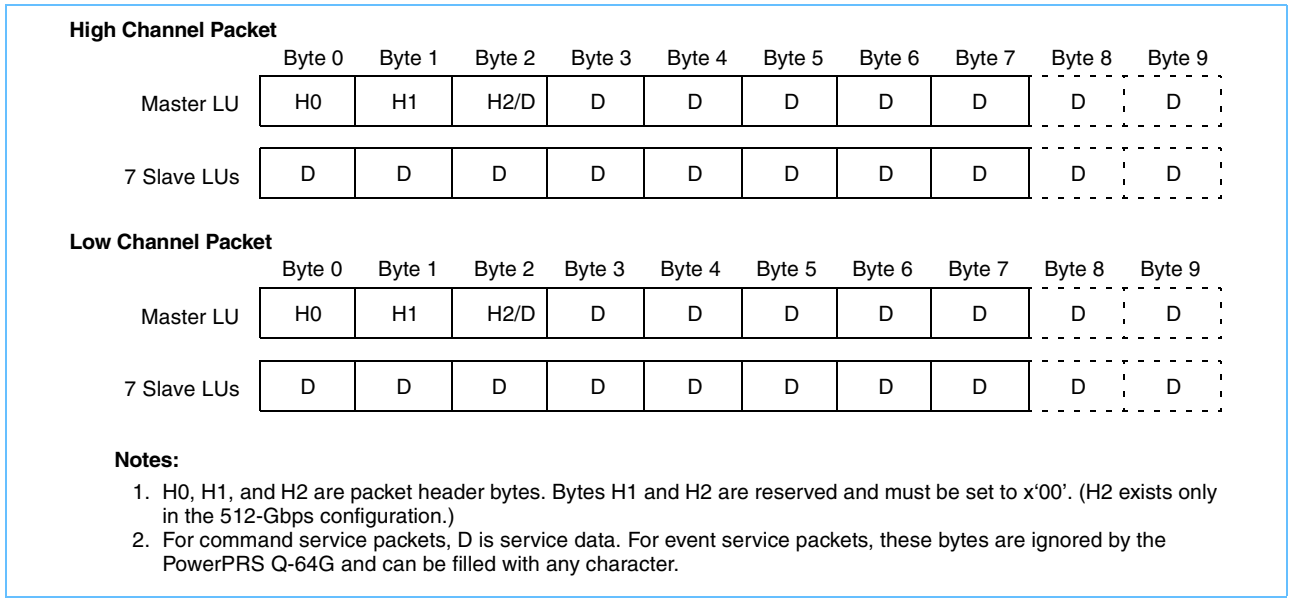

Notes:

1. H0, H1, and H2 are packet header bytes. Bytes H1 and H2 are reserved and must be set to x‘00’. (H2 exists only in the 512-Gbps configuration.)
2. For command service packets, D is service data. For event service packets, these bytes are ignored by the PowerPRS Q-64G and can be filled with any character.

Table 3-8. Ingress Service Packet, Byte H0

Byte	Information Carried				
	Bit 0	Bit 1	Bits 2:3	Bits 4:6	Bit 7
H0 (high and low channel)	Reserved	Parity	Protection	Service packet type	Reserved

Table 3-9. Ingress Service Packet, Byte H0 Field Descriptions

Bit(s)	Field Name	Description
0	Reserved	Reserved.
1	Parity	Header parity.
2:3	Protection	Must be set to ‘00’ for service packets.
4:6	Service Packet Type	Specifies the type of service packet: 100 Event-1 service packet 101 Command service packet 110 Event-2 service packet 111 Reserved
7	Reserved	Reserved.

3.3.6 Flow Control Flywheels for Grants Carried in Egress Packets

There are four flow control flywheels associated with the grants carried in egress packets:

- Extended output queue grant flywheel
- Output queue grant priority flywheel
- Subport grant type/priority flywheel
- Subport grant port flywheel

The extended output queue grant flywheel and the output queue grant priority flywheel carry related information as do the subport grant type/priority flywheel and the subport grant port flywheel.

3.3.6.1 Extended Output Queue Grant Flywheel

The extended output queue grant flywheel determines the set of ports for the output queue grants carried by the current egress packet cycle. For the 256-Gbps configuration, this flywheel cycles as follows:

0 Ports 0 to 7
1 Ports 8 to 15

For the 512-Gbps configuration, it cycles as follows:

0 Ports 0 to 15
1 Ports 16 to 31

There is one output queue grant per port and per priority. The particular set of output queue grants carried by the current packet cycle is determined by this flywheel in combination with the output queue grant priority flywheel. This flywheel is incremented every packet cycle. One packet cycle carries the grants for half the ports and one priority. When all four priorities are enabled, eight packet cycles are required to update the complete set of output queue grants.

To synchronize the PowerPRS Q-64G flywheels with those in the attached devices, the extended output queue grant flywheel status and output queue grant priority flywheel status are inserted into the packet qualifier byte (H0) of egress idle packets carried on the low channel. The output queue grants are inserted into header bytes H1 and, for the 512-Gbps configuration, H2 of egress data, control, service, and idle packets carried on the low channel. See the following packet format descriptions for more information.

3.3.6.2 Output Queue Grant Priority Flywheel

The output queue grant priority flywheel determines the priority of the output queue grants carried by the current egress packet cycle. This flywheel cycles as follows:

00 Priority 0
01 Priority 1
10 Priority 2
11 Priority 3

This flywheel is incremented each time the extended output queue grant flywheel returns to ‘0’. This flywheel returns to ‘00’ after reaching the number of priorities defined in the number of priorities field in the *Configuration 1 Register* (page 112).

3.3.6.3 Subport Grant Type/Priority Flywheel

The subport grant type/priority flywheel determines the type of subport grants carried by the current egress packet cycle. This flywheel cycles as follows:

000 Subport output queue grants, priority 0
001 Subport output queue grants, priority 1
010 Subport output queue grants, priority 2
011 Subport output queue grants, priority 3
100 Subport multicast grants

The particular set of subport grants carried by the current packet cycle is determined by this flywheel in combination with the subport grant port flywheel. This flywheel is incremented each time the subport grant port flywheel returns to ‘000’, and reset to ‘100’ after it reaches the number of priorities defined in the number of priorities field in the *Configuration 1 Register*. After reaching ‘100’, this flywheel returns to ‘000’ after only one packet cycle. In effect, the combined value of the subport grant type/priority flywheel and the subport grant port flywheel is a six-bit field that counts to ‘100000’ and then returns to ‘000000’.

Transmission of a complete set of subport grants requires 8, 16, 24, or 32 packet cycles, depending on the number of enabled priorities. Every 61 packet cycles, PowerPRS Q-64G multicast grants (not *subport* multicast grants) are inserted in place of the subport grants. (Packet cycle 61 was chosen so that the multicast grants would not preempt subport grants of the same priority every time the multicast grants were sent.) The insertion of multicast grants is indicated by the grant type bit in the packet qualifier byte (H0) of egress packets on the high channel. Multicast grants are also sent immediately when the multicast grant status changes. The subport grant type/priority flywheel and subport grant port flywheel are incremented even when multicast grants, rather than subport grants, are transmitted.

To synchronize the PowerPRS Q-64G flywheels with those in the attached devices, the subport grant/type priority flywheel status and subport grant port flywheel status are inserted into byte 6 (that is, the FC byte) of egress idle packets. The subport grants are inserted into header bytes H1 and, for the 512-Gbps configuration, H2 of egress data, control, service, and idle packets carried on the high channel. See the following packet format descriptions for more information.

3.3.6.4 Subport Grant Port Flywheel

The subport grant port flywheel determines the ports for the subport output queue grants carried by the current egress packet cycle. Each packet cycle carries grants for two ports in the 256-Gbps configuration and for four ports in the 512-Gbps configuration. For the 256-Gbps configuration, this flywheel cycles as follows:

000	Ports 0 and 1
001	Ports 2 and 3
010	Ports 4 and 5
011	Ports 6 and 7
100	Ports 8 and 9
101	Ports 10 and 11
110	Ports 12 and 13
111	Ports 14 and 15

For the 512-Gbps configuration, it cycles as follows:

000	Ports 0, 1, 2, and 3
001	Ports 4, 5, 6, and 7
010	Ports 8, 9, 10, and 11
011	Ports 12, 13, 14, and 15
100	Ports 16, 17, 18, and 19
101	Ports 20, 21, 22, and 23
110	Ports 24, 25, 26, and 27
111	Ports 28, 29, 30, and 31

This flywheel is incremented every packet cycle.

3.3.7 Egress Idle Packet Format

The egress idle packet format for the PowerPRS Q-64G 16-Gbps configuration is presented in *Figure 3-5*. Bytes 0:2 and bytes 5:6 of the master LU carry control information, as described in *Tables 3-10* through *3-19*.

Figure 3-5. Egress Idle Packet Format

High Channel Packet

	Byte 0	Byte 1	Byte 2	Byte 3	Byte 4	Byte 5	Byte 6	Byte 7	Byte 8	Byte 9
Master LU	H0	H1	H2/D21.5	D21.5	D21.5	SCC	FC	K28.1	D21.5	D21.5
7 Slave LUs	D21.5	D21.5	D21.5	D21.5	D21.5	D21.5	D21.5	K28.1	D21.5	D21.5

Low Channel Packet

	Byte 0	Byte 1	Byte 2	Byte 3	Byte 4	Byte 5	Byte 6	Byte 7	Byte 8	Byte 9
Master LU	H0	H1	H2/D21.5	D21.5	D21.5	SCC	FC	K28.5	D21.5	D21.5
7 Slave LUs	D21.5	D21.5	D21.5	D21.5	D21.5	D21.5	D21.5	K28.5	D21.5	D21.5

Notes:

1. H0, H1, and H2 are packet header bytes. H2 exists only in the 512-Gbps configuration.
2. D21.5, K28.1, and K28.5 characters are 8b/10b control characters. K28.1 and K28.5 characters are used for link synchronization (byte alignment and packet clock recovery) and supervision (byte alignment and synchronization lost). D21.5 characters guarantee continuous transition on the line.
3. SCC is side communication channel data. Bits 4:7 are copies of bits 0:3.
4. FC is flow control information in addition to that transmitted in the packet header.

Table 3-10. Egress Idle Packet, Byte H0

Byte	Information Carried				
	Bit 0	Bit 1	Bits 2:3	Bits 4:5	Bits 6:7
H0 (high channel)	Grant type	Parity	Protection	Color	Reserved
H0 (low channel)	Extended output queue grant flywheel status	Parity	Protection	Color	Output queue grant priority flywheel status

Table 3-11. Egress Idle Packet, Byte H0 Field Descriptions

Bit(s)	Field Name	Description
0	Grant Type (high channel)	Identifies the type of grants carried in byte H1 of the packet header: 0 Subport grants (default value) 1 Multicast grants (set to '1' every 61 packet cycles or when the multicast grant status changes) The default value is '0'. In this case, the set of subport grants carried by the egress idle packet is determined by the subport grant type/priority flywheel and the subport grant port flywheel. The FC byte carries the status of these flywheels (see *Tables 3-18* and *3-19*). Every 61 packet cycles, multicast grants are sent in place of the subport grants. Multicast grants are also sent immediately when the multicast grant status changes. The grant type/priority and subport grant port flywheels are incremented even when the multicast grants are inserted in place of the subport grants. **Note:** A subport multicast grant is the logical AND of the individual subport multicast grants that were received in ingress packets for all the ports in the range (see *Section 3.6 Subport Flow Control* on page 53).
	Extended Output Queue Grant Flywheel Status (low channel)	Reports the status of the extended output queue grant flywheel: 256-Gbps Configuration: 0 Ports 0 to 7; 1 Ports 8 to 15 512-Gbps Configuration: 0 Ports 0 to 15; 1 Ports 16 to 31 See *Section 3.3.6.1 Extended Output Queue Grant Flywheel* on page 37 for more information.
1	Parity	Header parity.
2:3	Protection	Must be set to '00' for idle packets.
4:5	Color	Identifies the packet color for switchover support. See the ingress idle packet color field description in *Table 3-2*.
6:7	Output Queue Grant Priority Flywheel Status	Reports the status of the output queue grant priority flywheel: 00 Priority 0 01 Priority 1 10 Priority 2 11 Priority 3 See *Section 3.3.6.2 Output Queue Grant Priority Flywheel* on page 37 for more information.

Table 3-12. Egress Idle Packet, Byte H1 on the High Channel in the 256-Gbps Configuration
(This table also applies to egress data packets, control packets, and service packets.)

Byte	Information Carried							
	Bit 0	Bit 1	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7
H1 (high channel) for grant type (byte H0) = 0 Subport grant type/ priority flywheel (FC byte) = 0 to 3 (grant priority)	Subport output queue grant: port *n*, subport 0	Subport output queue grant: port *n*, subport 1	Subport output queue grant: port *n*, subport 2	Subport output queue grant: port *n*, subport 3	Subport output queue grant: port *n* + 1, subport 0	Subport output queue grant: port *n* + 1, subport 1	Subport output queue grant: port *n* + 1, subport 2	Subport output queue grant: port *n* + 1, subport 3
H1 (high channel) for grant type (byte H0) = 0 Subport grant type/ priority flywheel (FC byte) = 4 (subport multicast)	Subport multicast grant: ports 0 to 7, priority 0	Subport multicast grant: ports 0 to 7, priority 1	Subport multicast grant: ports 0 to 7, priority 2	Subport multicast grant: ports 0 to 7, priority 3	Subport multicast grant: ports 8 to 15, priority 0	Subport multicast grant: ports 8 to 15, priority 1	Subport multicast grant: ports 8 to 15, priority 2	Subport multicast grant: ports 8 to 15, priority 3
H1 (high channel) for grant type (byte H0) = 1	Multicast grant: high subswitch element, priority 0	Multicast grant: high subswitch element, priority 1	Multicast grant: high subswitch element, priority 2	Multicast grant: high subswitch element, priority 3	Multicast grant: low subswitch element, priority 0	Multicast grant: low subswitch element, priority 1	Multicast grant: low subswitch element, priority 2	Multicast grant: low subswitch element, priority 3

Note: *n* = subport grant port flywheel (FC byte) = 0 to 7

Table 3-13. Egress Idle Packet, Byte H1 on the High Channel in the 512-Gbps Configuration
(This table also applies to egress data packets, control packets, and service packets.)

Byte	Information Carried							
	Bit 0	Bit 1	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7
H1 (high channel) for grant type (byte H0) = 0 Subport grant type/ priority flywheel (FC byte) = 0 to 3 (grant priority)	Subport output queue grant: port *n*, subport 0	Subport output queue grant: port *n*, subport 1	Subport output queue grant: port *n*, subport 2	Subport output queue grant: port *n*, subport 3	Subport output queue grant: port *n* + 1, subport 0	Subport output queue grant: port *n* + 1, subport 1	Subport output queue grant: port *n* + 1, subport 2	Subport output queue grant: port *n* + 1, subport 3
H1 (high channel) for grant type (byte H0) = 0 Subport grant type/ priority flywheel (FC byte) = 4 (subport multicast)	Subport multicast grant: ports 0 to 15, priority 0	Subport multicast grant: ports 0 to 15, priority 1	Subport multicast grant: ports 0 to 15, priority 2	Subport multicast grant: ports 0 to 15, priority 3	Subport multicast grant: ports 16 to 31, priority 0	Subport multicast grant: ports 16 to 31, priority 1	Subport multicast grant: ports 16 to 31, priority 2	Subport multicast grant: ports 16 to 31, priority 3
H1 (high channel) for grant type (byte H0) = 1	Multicast grant: high subswitch element, priority 0	Multicast grant: high subswitch element, priority 1	Multicast grant: high subswitch element, priority 2	Multicast grant: high subswitch element, priority 3	Multicast grant: low subswitch element, priority 0	Multicast grant: low subswitch element, priority 1	Multicast grant: low subswitch element, priority 2	Multicast grant: low subswitch element, priority 3

Note: *n* = subport grant port flywheel (FC byte) = 0 to 7

Table 3-14. Egress Idle Packet, Byte H2 on the High Channel in the 512-Gbps Configuration
(This table also applies to egress data packets, control packets, and service packets.)

Byte	Information Carried							
	Bit 0	Bit 1	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7
H2 (high channel) for subport grant type/ priority flywheel (FC byte) = 0 to 3 (grant priority)	Subport output queue grant: port *n* + 2, subport 0	Subport output queue grant: port *n* + 2, subport 1	Subport output queue grant: port *n* + 2, subport 2	Subport output queue grant: port *n* + 2, subport 3	Subport output queue grant: port *n* + 3, subport 0	Subport output queue grant: port *n* + 3, subport 1	Subport output queue grant: port *n* + 3, subport 2	Subport output queue grant: port *n* + 3, subport 3
H2 (high channel) for subport grant type/ priority flywheel (FC byte) = 4 (subport multicast)	Reserved	Reserved	Reserved	Reserved	Reserved	Reserved	Reserved	Reserved

Note: *n* = subport grant port flywheel (FC byte) = 0 to 7

Table 3-15. Egress Idle Packet, Byte H1 on the Low Channel in the 256-Gbps Configuration
(This table also applies to egress data packets, control packets, and service packets.)

Byte	Information Carried							
	Bit 0	Bit 1	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7
H1 (low channel) for extended output queue grant flywheel = 0	Output queue grant: port 0	Output queue grant: port 1	Output queue grant: port 2	Output queue grant: port 3	Output queue grant: port 4	Output queue grant: port 5	Output queue grant: port 6	Output queue grant: port 7
H1 (low channel) for extended output queue grant flywheel = 1	Output queue grant: port 8	Output queue grant: port 9	Output queue grant: port 10	Output queue grant: port 11	Output queue grant: port 12	Output queue grant: port 13	Output queue grant: port 14	Output queue grant: port 15

Table 3-16. Egress Idle Packet, Bytes H1 and H2 on the Low Channel in the 512-Gbps Configuration
(This table also applies to egress data packets, control packets, and service packets.)

Byte	Information Carried							
	Bit 0	Bit 1	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7
H1 (low channel) for extended output queue grant flywheel = 0	Output queue grant: port 0	Output queue grant: port 1	Output queue grant: port 2	Output queue grant: port 3	Output queue grant: port 4	Output queue grant: port 5	Output queue grant: port 6	Output queue grant: port 7
H2 (low channel) for extended output queue grant flywheel = 0	Output queue grant: port 8	Output queue grant: port 9	Output queue grant: port 10	Output queue grant: port 11	Output queue grant: port 12	Output queue grant: port 13	Output queue grant: port 14	Output queue grant: port 15
H1 (low channel) for extended output queue grant flywheel = 1	Output queue grant: port 16	Output queue grant: port 17	Output queue grant: port 18	Output queue grant: port 19	Output queue grant: port 20	Output queue grant: port 21	Output queue grant: port 22	Output queue grant: port 23
H2 (low channel) for extended output queue grant flywheel = 1	Output queue grant: port 24	Output queue grant: port 25	Output queue grant: port 26	Output queue grant: port 27	Output queue grant: port 28	Output queue grant: port 29	Output queue grant: port 30	Output queue grant: port 31

Table 3-17. Egress Idle Packet, SCC Byte

Byte	Information Carried							
	Bit 0	Bit 1	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7
SCC (high or low channel)	SCC data: bit 0	SCC data: bit 1	SCC data: bit 2	SCC data: bit 3	SCC data: bit 0	SCC data: bit 1	SCC data: bit 2	SCC data: bit 3

Table 3-18. Egress Idle Packet, FC Byte

Byte	Information Carried			
	Bit 0	Bit 1	Bits 2:4	Bits 5:7
FC (high or low channel)	Memory grant	Reserved	Subport grant type/ priority flywheel status	Subport grant port flywheel status

Table 3-19. Egress Idle Packet, FC Byte Field Descriptions

Bit(s)	Field Name	Description
0	Memory Grant	Reports the status of the memory grant.
1	Reserved	Reserved.
2:4	Subport Grant Type/ Priority Flywheel Status	Reports the status of the subport grant type/priority flywheel: 000 Subport output queue grants, priority 0 001 Subport output queue grants, priority 1 010 Subport output queue grants, priority 2 011 Subport output queue grants, priority 3 100 Subport multicast grants See *Section 3.3.6.3 Subport Grant Type/Priority Flywheel* on page 37 for more information.
5:7	Subport Grant Port Flywheel Status	Reports the status of the subport grant port flywheel: 256-Gbps Configuration: 000 Ports 0 and 1; 001 Ports 2 and 3; 010 Ports 4 and 5; 011 Ports 6 and 7; 100 Ports 8 and 9; 101 Ports 10 and 11; 110 Ports 12 and 13; 111 Ports 14 and 15 512-Gbps Configuration: 000 Ports 0, 1, 2, and 3; 001 Ports 4, 5, 6, and 7; 010 Ports 8, 9, 10, and 11; 011 Ports 12, 13, 14, and 15; 100 Ports 16, 17, 18, and 19; 101 Ports 20, 21, 22, and 23; 110 Ports 24, 25, 26, and 27; 111 Ports 28, 29, 30, and 31 See *Section 3.3.6.4 Subport Grant Port Flywheel* on page 38 for more information.

3.3.8 Egress Data Packet and Control Packet Format

The egress data packet and control packet format for the PowerPRS Q-64G 16-Gbps configuration is presented in *Figure 3-6*. The packet qualifier byte (H0) is described in *Tables 3-20* and *3-21.* Header bytes H1 and H2 of egress data packets and control packets are identical to those of egress idle packets (see *Tables 3-12* through *3-16).*

Figure 3-6. Egress Data Packet and Control Packet Format

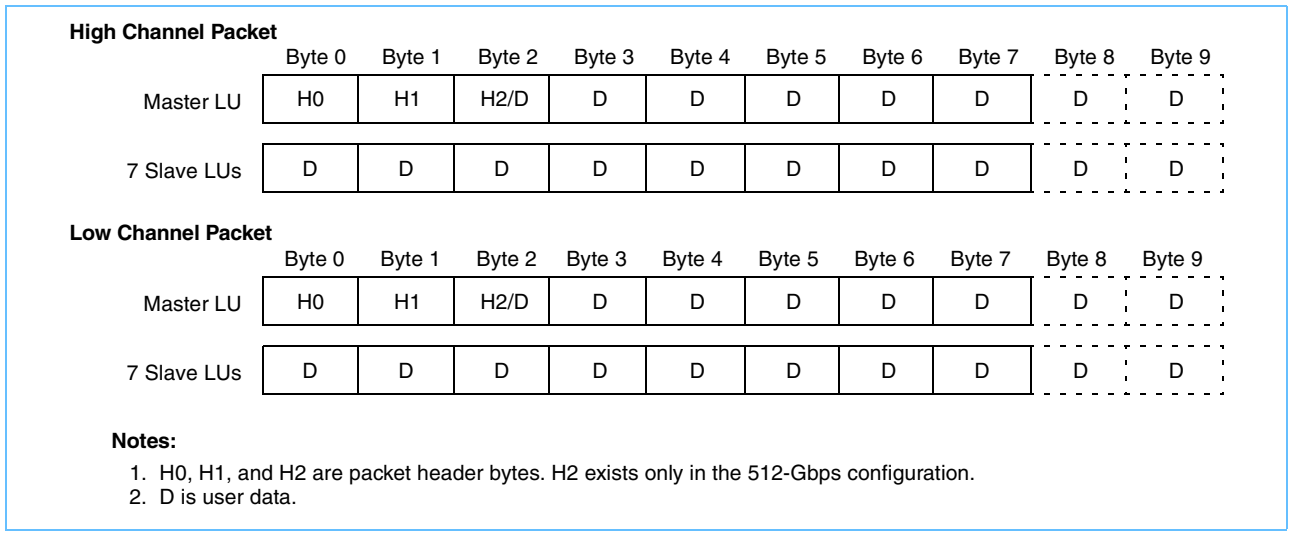

Notes:

1. H0, H1, and H2 are packet header bytes. H2 exists only in the 512-Gbps configuration.
2. D is user data.

Table 3-20. Egress Data Packet and Control Packet, H0 Byte

Byte	Information Carried					
	Bit 0	Bit 1	Bits 2:3	Bit 4	Bit 5	Bits 6:7
H0 (high channel)	Grant type	Parity	Protection	Best effort	Reserved	Packet priority
H0 (low channel)	Memory grant	Parity	Protection	Best effort	Reserved	Packet priority

Table 3-21. Egress Data Packet and Control Packet, Byte H0 Field Descriptions

Bit(s)	Field Name	Description
0	Grant Type (high channel)	Identifies the type of grants carried in byte H1 of the packet header: 0 Subport grants (default value) 1 Multicast grants (set to '1' every 61 packet cycles or when the multicast grant status changes) These bit settings operate the same as those for egress idle packets (see *Table 3-11*). Byte H1 is also formatted the same as it is for egress idle packets (see *Tables 3-12* through *3-16*).
	Memory Grant (low channel)	Reports the status of the memory grant.
1	Parity	Header parity.
2:3	Protection	This field is processed without modification. See the ingress data packet and control packet byte H0 field descriptions in *Table 3-5* on page 34.
4	Best Effort	This bit is processed without modification. See the ingress data packet and control packet byte H0 field descriptions in *Table 3-5*.
	Reserved	Reserved.
6:7	Packet Priority	This field is processed without modification. See the ingress data packet and control packet byte H0 field descriptions in *Table 3-5*.

3.3.9 Egress Service Packet Format

The egress service packet format for the PowerPRS Q-64G 16-Gbps configuration is presented in *Figure 3-7*. The packet qualifier byte (H0) is described in *Tables 3-22* and *3-23*. Header bytes H1 and H2 of egress data packets and control packets are identical to those of egress idle packets (see *Tables 3-12* through *3-16).*

Figure 3-7. Egress Service Packet Format

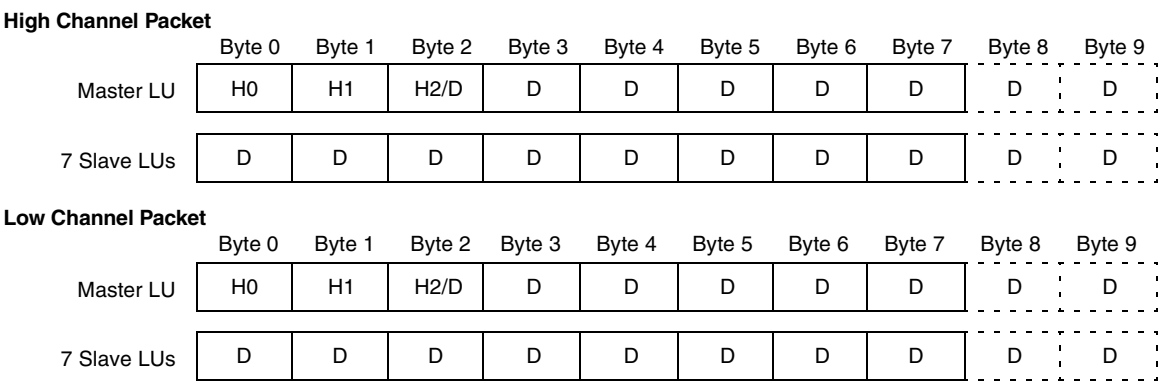

Notes:

1. H0, H1, and H2 are packet header bytes. H2 exists only in the 512-Gbps configuration.
2. For command service packets, D is service data. For event service packets, these bytes are ignored by the attached device and can be filled with any character.

Table 3-22. Egress Service Packet, Byte H0

Byte	Information Carried				
	Bit 0	Bit 1	Bits 2:3	Bits 4:6	Bit 7
H0 (high channel)	Grant type	Parity	Protection	Service packet type	Reserved
H0 (low channel)	Memory grant	Parity	Protection	Service packet type	Reserved

Table 3-23. Egress Service Packet, Byte H0 Field Descriptions

Bit(s)	Field Name	Description
0	Grant Type (high channel)	Identifies the type of grants carried in bytes H1 and H2 of the packet header: 0 Subport grants (default value) 1 Multicast grants (set to '1' every 61 packet cycles or when the multicast grant status changes) These bit settings operate the same as those for egress idle packets (see *Table 3-11*). Bytes H1 and H2 are also formatted the same as they are for egress idle packets (see *Tables 3-12* through *3-16).*
	Memory Grant (low channel)	Reports the status of the memory grant.
1	Parity	Header parity.
2:3	Protection	Must be set to '00' for service packets.
4:6	Service Packet Type	Specifies the type of service packet: 100 Event-1 service packet 101 Command service packet 110 Event-2 service packet 111 Reserved
7	Reserved	Reserved.

3.4 Ingress Flow Control

Note: In the flow control sections (that is, *Sections 3.4* through *3.7),* references to 32 ports and port ranges of 0 to 15 and 16 to 31 are specific to the PowerPRS Q-64G 512-Gbps configuration. Substitute 16 ports and port ranges of 0 to 7 and 8 to 15 for the 256-Gbps configuration.

Ingress traffic flow to the PowerPRS Q-64G is controlled by a variety of mechanisms, primarily:

- Output queue grants that reflect the output queue occupancy. There is one output queue grant issued per subswitch element pair, per output, and per priority.
- Memory grants that reflect the shared memory occupancy. There is one memory grant issued for input ports 0 to 15 and one issued for input ports 16 to 31.
- Multicast grants that reflect the multicast packet use of shared memory. There is one multicast grant issued per subswitch element and per priority.

As discussed in *Section 3.3,* the PowerPRS Q-64G inserts the grant status into egress data, control, service, and idle packet headers to convey these grants to the ingress side of the attached devices. These grants allow the devices to transmit packets to the PowerPRS Q-64G. Ingress traffic flow is also controlled by the flow control latency and best-effort discard functions. Each of these flow control mechanisms is discussed below.

Note: An attached device transmits a *unicast* packet to the PowerPRS Q-64G only when it has received both the memory grant and the output queue grant for the destination output. An attached device transmits a *multicast* packet to the PowerPRS Q-64G when it has received the memory grant, regardless of the output queue grant status and the multicast grant status. The PowerPRS Q-64G discards ingress packets destined for an output that lacks the correct output queue grant status (see *Section 3.4.5 Flow Control Latency* on page 49).

3.4.1 Output Queue Grants

Output queue grants:

- Prevent packets of a single priority destined for the same output from filling the shared memory.
- Prevent low-priority packets from occupying too much shared memory.

The PowerPRS Q-64G issues one output queue grant per subswitch element pair (A and B, and C and D), per output port (0 to 31), and per priority (0 to 3). For a subswitch element, the PowerPRS Q-64G issues an output queue grant for an output priority when the total number of packets in the output queue is below the output queue threshold for that priority. The PowerPRS Q-64G inserts the output queue grant status into egress data, control, service, and idle packet headers. Ports 0 to 15 receive the status of the output queues in subswitch elements A and B; ports 16 to 31 receive the status of the output queues in subswitch elements C and D. The PowerPRS Q-64G removes the grant when the number of packets is equal to or greater than the threshold (there is no hysteresis). The four programmable output queue thresholds, one for each output priority, are accessed through the *Threshold Access Register* (page 114).

To generate output queue grants, the PowerPRS Q-64G continuously compares the total packet count in each output queue to the four output queue thresholds. However, the PowerPRS Q-64G output queue grant table is refreshed only once per packet cycle. This guarantees that all the attached devices receive the same flow control information.

3.4.2 Memory Grants

Memory grants provide ingress flow control by preventing packet transmission from the attached devices when the PowerPRS Q-64G shared memory is full. The PowerPRS Q-64G issues one memory grant for input ports 0 to 15 and one for input ports 16 to 31. Because the packets from an attached device are routed to the shared memory of two subswitch elements (A and B, or C and D), the PowerPRS Q-64G generates these grants as follows:

- The PowerPRS Q-64G issues the memory grant for input ports 0 to 15 when the total packet count in the subswitch element A shared memory and in the subswitch element B shared memory are *both* below the shared memory threshold. The memory grant status is inserted into the headers of egress data, control, service, and idle packets on ports 0 to 15.
- The PowerPRS Q-64G issues the memory grant for input ports 16 to 31 when the total packet count in the subswitch element C shared memory and in the subswitch element D shared memory are *both* below the shared memory threshold. The memory grant status is inserted into the headers of egress data, control, service, and idle packets on ports 16 to 31.

The PowerPRS Q-64G removes a memory grant when the number of packets in either subswitch element is equal to or greater than the threshold (there is no hysteresis). When the packet count in a subswitch element shared memory reaches the shared memory threshold, the event is reported via the corresponding shared memory threshold exceeded bit in the *Status Register* (page 120). This bit generates an interrupt unless it is masked by the corresponding bit in the *Interrupt Mask Register* (page 122). The programmable shared memory threshold is accessed through the *Threshold Access Register* (page 114).

When programming the shared memory threshold, the number of packet store addresses reserved by the input controllers must be subtracted from the total number of packet store addresses available in the shared memory. (These reserved addresses store new packets during header extraction and packet retrieval.) In the 512-Gbps configuration, 32 addresses are reserved; in the 256-Gbps configuration, 16 addresses are reserved. Threshold programming must also consider flow control latency on the Unilinks.

To generate memory grants, the PowerPRS Q-64G continuously compares the packet count in each subswitch element to the shared memory threshold. However, the PowerPRS Q-64G memory grant table is refreshed only once per packet cycle. This guarantees that all the attached devices receive the same flow control information.

3.4.3 Multicast Grants

Multicast grants provide PowerPRS Q-64G ingress flow control by preventing multicast packets from filling the entire shared memory. The PowerPRS Q-64G issues one multicast grant per subswitch element (A, B, C, and D) and per priority (0 to 3). The multicast grant status for subswitch elements A and B is inserted into the headers of egress data, control, service, and idle packets on ports 0 to 15; the multicast grant status for subswitch elements C and D is inserted into the headers of egress data, control, service, and idle packets on ports 16 to 31. For each of the four subswitch elements, an internal register, MCCount, contains the difference between the total number of packet addresses in the output queues (regardless of priority or input port) and the total number of packets stored in the shared memory.

Each priority has a multicast high threshold and a multicast low threshold that apply to all four subswitch elements. These thresholds are accessed through the *Threshold Access Register*. A multicast grant for a priority is issued when the MCCount register value falls below the multicast low threshold for that priority; a multicast grant for a priority is removed when the MCCount register value is equal to or greater than the

multicast high threshold for that priority. The difference between the multicast high and low thresholds for a priority causes hysteresis. Hysteresis reduces the number of times a change in multicast grant status preempts the insertion of subport grants in egress packet headers.

To generate multicast grants, the PowerPRS Q-64G continuously compares the MCCount register value to the four multicast high thresholds and the four multicast low thresholds. However, the PowerPRS Q-64G multicast grant table is refreshed only once per packet cycle. This guarantees that all the attached devices receive the same flow control information.

3.4.4 Shared Memory Overrun

The PowerPRS Q-64G receives any incoming packet that has a store address, regardless of output queue and shared memory occupancy (unless the flow control latency function requires packet discard; see *Section 3.4.5*). If a packet is received when the input controller does not have an available store address, the input controller discards the packet. This condition is reported in the no address interrupt bit in the *Status Register* (page 120). An interrupt is generated if the condition is not masked with the no address interrupt bit in the *Interrupt Mask Register* (page 122). This error occurs only if the shared memory threshold is programmed incorrectly, or if the attached device is not responding to the memory grant information.

3.4.5 Flow Control Latency

The input controller flow control latency function detects operational errors with the attached devices. This function is enabled with the flow control latency field in the *Configuration 0 Register* (page 110). When the flow control latency function is enabled, the input controller checks whether an incoming unicast packet is destined to an output for which neither an output queue grant nor a memory grant has been issued in the past *n* packet cycles. For a multicast packet, the input controller checks only the memory grant information. If the packet is destined for an output for which no grants have been issued, the packet is discarded. This error is reported via the flow control violation bit in the *Status Register* and, unless masked in the *Interrupt Mask Register*, the error generates a flow control violation interrupt. The violating ports are identified by the corresponding bits in the *Flow Control Violation Register* (page 132).

Note: The PowerPRS Q-64G cannot verify whether the attached devices are processing the multicast grants. The PowerPRS Q-64G accepts multicast packets even when the multicast grant is removed.

3.4.6 Best-Effort Discard

In some applications, certain low-priority traffic is more important than the high-priority traffic that monopolizes the output port and prevents low-priority traffic from accessing the port. The best-effort discard function attempts to correct this situation by categorizing incoming traffic as either “guaranteed bandwidth” or “best-effort bandwidth.” Best-effort bandwidth traffic is discarded at the input controllers, when necessary, to provide output port access to guaranteed traffic. The best-effort discard function helps to ensure guaranteed bandwidth traffic quality of service. This function is enabled with the best-effort discard enable bit in the *Configuration 0 Register* (page 110).

When the PowerPRS Q-64G is operating as a lossy switch, the best-effort discard flow control function is activated as soon as the aggregate traffic load for an output port exceeds its capacity for a given time period. The PowerPRS Q-64G discards only traffic flagged as best-effort; guaranteed bandwidth traffic is never discarded. Guaranteed bandwidth traffic congestion is managed through the normal flow control mechanism. Best-effort discard from the shared memory is completed in a single burst to minimize the number of affected packets.

3.4.6.1 Best-Effort Discard Counters

Operation of the lossy switch is based on a set of five counters (see *Table 3-24)* per output port that are incremented when the port's output queues receive specific packet types.

Table 3-24. Best-Effort Discard Counters

Counter	Description
Counter 1	The main counter. Counts each incoming packet.
Counter 2	Counts each incoming packet that is either guaranteed bandwidth of any priority or best-effort bandwidth of priority 0, 1, or 2.
Counter 3	Counts each incoming packet that is either guaranteed bandwidth of any priority or best-effort bandwidth of priority 0 or 1.
Counter 4	Counts each incoming packet that is either guaranteed bandwidth of any priority or best-effort bandwidth of priority 0.
Counter 5	Counts each incoming packet that is guaranteed bandwidth of any priority.

Best-effort discard counters are decremented at a packet speed equivalent to the throughput of the port (the line rate is provided by a register defined at port initialization). The counters are accessible via the *Best-Effort Resources Access Register* (page 118).

3.4.6.2 Best-Effort Discard Thresholds

Three thresholds govern the best-effort discard process:

- *Enable discard threshold*. When counter 1 reaches this threshold, the discard process starts.
- *Halt discard threshold*. When counter 1 reaches this threshold, the discard process stops.
- *Priority discard threshold*. This threshold, in comparison to counters 1, 2, 3, and 4, determines which packets are discarded:
  - When counter 1 is above the priority discard threshold, best-effort packets of priority 3 are discarded.
  - When counter 2 is above the priority discard threshold, best-effort packets of priority 2 are discarded.
  - When counter 3 is above the priority discard threshold, best-effort packets of priority 1 are discarded.
  - When counter 4 is above the priority discard threshold, best-effort packets of priority 0 are discarded.

The priority discard threshold only determines which packets are discarded; the halt discard threshold determines when the discard process stops. The best-effort discard thresholds are accessed via the *Best-Effort Resources Access Register*.

Figure 3-8 illustrates how the best-effort discard thresholds operate. In the figure, best-effort discard starts because counter 1 for the output port exceeds the enable discard threshold. Counters 1, 2, and 3 are all above the priority discard threshold and, consequently, the best-effort packets of priority 1, 2, and 3 destined for that output port are discarded at the input controller. Discard continues until counter 1 falls below the halt discard threshold.

Figure 3-8. Best-Effort Discard Counters and Thresholds

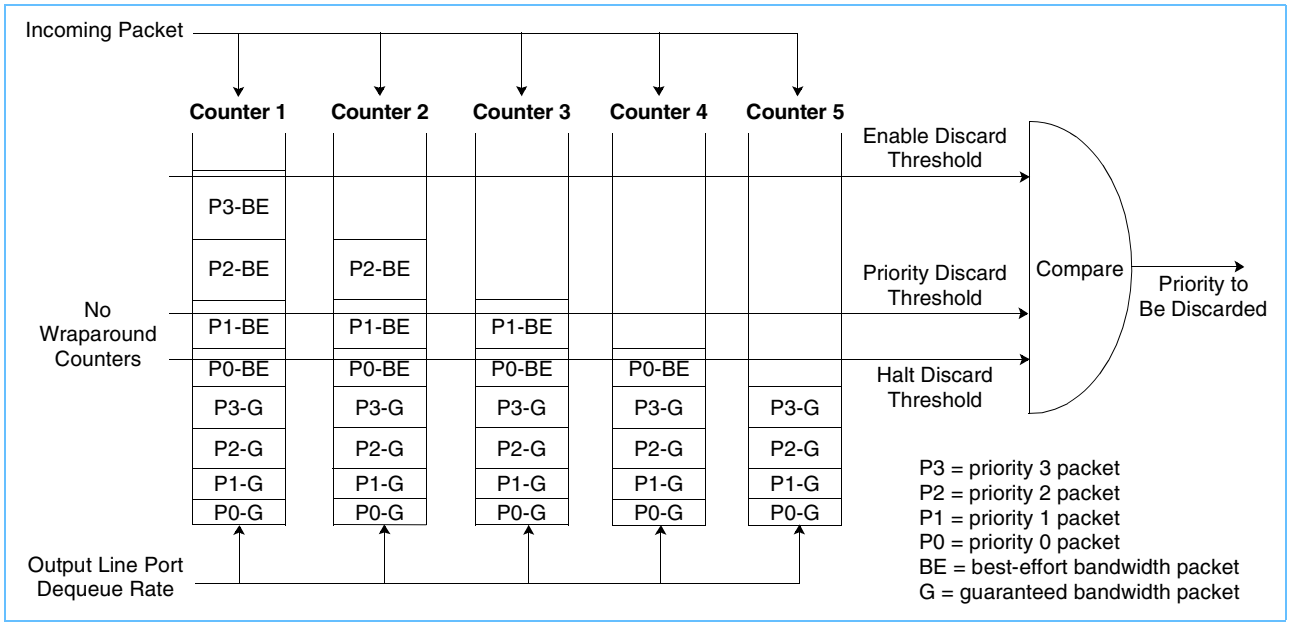

The assumption is that guaranteed bandwidth should be engineered so that it never exceeds the priority discard threshold—it dequeues traffic without the need for flow control. However, situations exist in which the traffic pattern changes before the halt discard threshold is reached. In this case, if another counter exceeds the priority discard threshold, the input controllers discard additional packets.

3.4.6.3 Best-Effort Discard Filters

Best-effort discard occurs within the input controllers, which filter the destination bitmaps of incoming packets (see *Figure 3-9).* There is one best-effort discard filter per priority. The best-effort discard filter is used as a destination bitmap mask, with the discard set to ‘0’. For each combination of destination output port and priority, there is a 20-bit counter that provides discard quantity and rate information. These drop counters are enabled with the best-effort drop counters enable bit in the *Configuration 0 Register* (page 110) and are accessible via the *Best-Effort Resources Access Register* (page 118).

When an incoming packet arrives at the PowerPRS Q-64G, its packet priority field, best-effort bit, and destination bitmap are provided to the best-effort filter logic. If the packet is read as best-effort, the best-effort discard filter for the packet priority is applied to the destination bitmap. If the combination is ‘0’, the packet is discarded rather than enqueued, and the corresponding drop counter (or counters for a multicast packet) is updated.

Figure 3-9. Best-Effort Discard Drop Filters

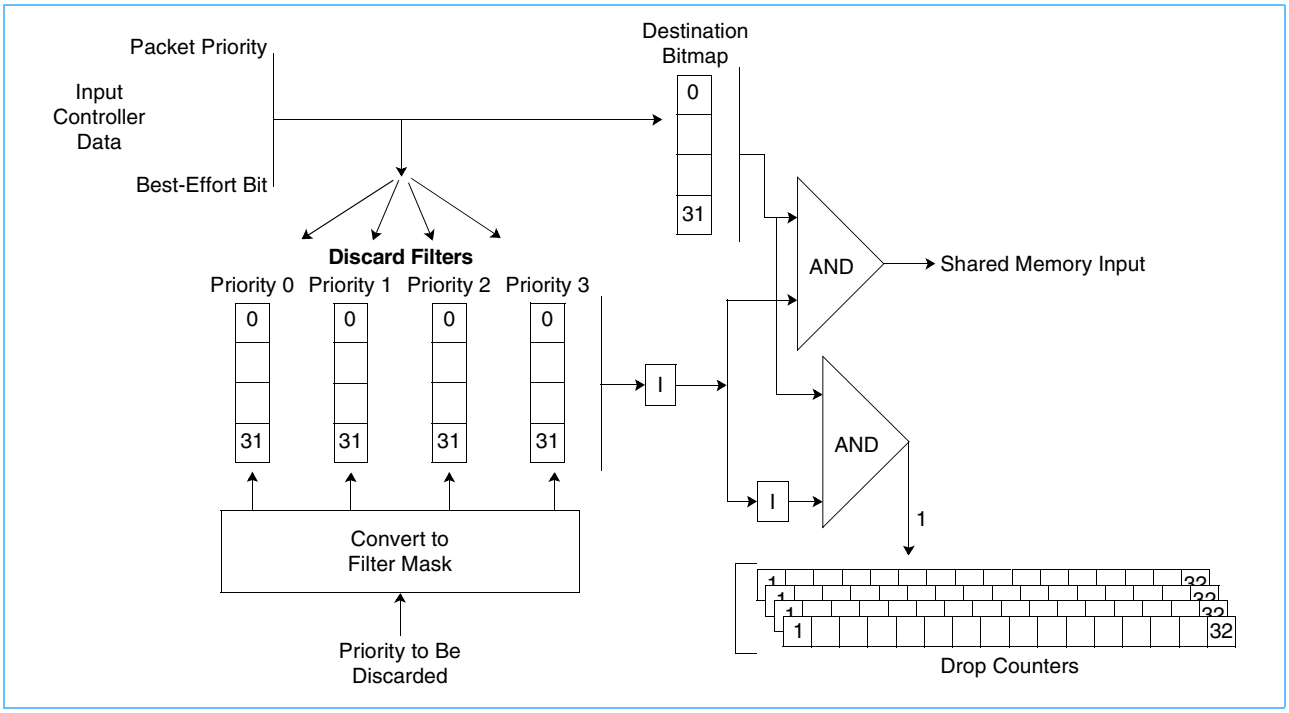

3.5 Egress Flow Control

Egress traffic flow from the PowerPRS Q-64G is controlled by a variety of mechanisms, including send grants, the send grant antistreaming function, and the credit table.

3.5.1 Send Grants

Egress traffic flow from the PowerPRS Q-64G to an attached device is controlled primarily by a grant mechanism. The PowerPRS Q-64G transmits a packet on an output port when the attached device issues a send grant for that port and priority. The attached device inserts the send grant status for a port into the ingress idle packet payload and ingress data packet headers (see *Section 3.3 Packet Format According to Packet Type* on page 28). The PowerPRS Q-64G extracts the send grant status from the ingress packets and stores it in internal tables. When the send grant is removed, the PowerPRS Q-64G generates idle packets on the port.

Note: The PowerPRS Q-64G sends control packets to an output port if the send grant status for that port is active for at least one priority. It sends service packets to an output port regardless of the send grant status.

3.5.2 Send Grant Antistreaming

The optional send grant antistreaming function is intended to prevent a defective attached device from indefinitely removing the send grant. This function is enabled with the send grant antistreaming enable bit in the *Configuration 0 Register* (page 110). If the send grant is not issued for any priority for a given number of contiguous packet cycles, the antistreaming function (if enabled) internally forces the send grant to active for all priorities until the attached device issues another send grant. The number of contiguous packet cycles (from

16 to 2048) is programmable using the send grant antistreaming threshold field in the *Configuration 0 Register*. When the number of packet cycles is exceeded, a violation is reported via the send grant violation bit in the *Status Register* (page 120), and an interrupt is raised if the violation is not masked with the send grant violation bit in the *Interrupt Mask Register* (page 122). The port with the violation is reported in the *Send Grant Violation Register* (page 131).

When the send grant antistreaming function is enabled, the output queue scheduler also considers the send grant active and sets an internal send grant to avoid congestion inside the PowerPRS Q-64G. The internal send grant is removed when the attached device resumes the normal send grant.

3.5.3 Credit Table

In the basic PowerPRS Q-64G configuration, the output queue scheduler schedules packet transmission for each output port in the following order:

1. Service packets
2. Control packets
3. Priority 0 data packets
4. Priority 1 data packets
5. Priority 2 data packets
6. Priority 3 data packets

The credit table provides a weighted cycling mechanism that can be programmed to guarantee minimum bandwidth for low-priority packets. For each output port, the credit table indicates the packet priority to be transmitted during each packet cycle. The credit table includes 256 entries, or credits, per port. The credit table is read once per packet cycle; that is, the priority entry is read for the current credit number, which generates a credit for that priority. The credit number is incremented by one for every packet cycle. After credit number 255, the credit number returns to 0.

When a credit is generated for a priority, a packet of that priority is sent on the output port if there is a packet in the output queue and the send grant is active for that priority. If either the output queue is empty or the send grant is inactive, the basic algorithm applies.

Use of the credit table is specified by the credit table enable field in the *Configuration 1 Register* (page 112). This field also specifies which packet priorities can be preempted by lower-priority packets. Indirect access to the credit table is provided via the *Credit Table Access Register* (page 116).

3.6 Subport Flow Control

A single PowerPRS C192 can be attached to four OC-48 protocol engines, or *subports*. The PowerPRS C192 uses two types of grants to control egress traffic flow to the subports:

- Subport output queue grants
- Subport multicast grants

Although these grants function much like the PowerPRS Q-64G output queue grants and multicast grants discussed in *Sections 3.4.1* and *3.4.3,* they reflect egress packet occupancy of the attached PowerPRS C192s. A PowerPRS C192 inserts the subport grant status into ingress data, control, and idle packets. The PowerPRS Q-64G extracts the subport grant status and stores it in an internal table, and then it inserts the subport grant status into egress data, control, service, and idle packet headers to broadcast it to the ingress side of all the attached PowerPRS C192s. The PowerPRS C192s use the subport grant status to control ingress traffic to the PowerPRS Q-64G.

The ingress side of the PowerPRS Q-64G receives one subport output queue grant per port (0 to 31), per subport (0 to 3), and per priority (0 to 3) and one subport multicast grant per port (0 to 31) and per priority (0 to 3). Although the PowerPRS Q-64G broadcasts all the subport output queue grants via egress packets, it broadcasts only two subport multicast grants per priority (one for ports 0 to 15 and one for ports 16 to 31). The subport multicast grants broadcast to ports 0 to 15 and to ports 16 to 31 are the logical AND of the ingress flow control information for their respective port ranges. Consequently, a subport multicast grant for a port range and priority is issued only if the output queue grant status for all 16 ports in the range is active.

The internal subport grant table is refreshed only once per packet cycle. This guarantees that all the attached devices receive the same flow control information.

3.7 Flow Control Information Summary

The PowerPRS Q-64G stores the current status of the various flow control grants in internal tables. Each grant requires one bit of table capacity. *Table 3-25* summarizes the types of flow control grants, the components for which each type of grant is issued, and the resulting number of bits required to store the grant status in the table.

Table 3-25. Flow Control Information Summary

Flow Control Grant Type	Components for Which Grants are Issued				Total Number of Table Bits Required
	Subswitch Elements	Ports	Subports	Priorities	
Output queue grants	A/B and C/D	0 to 31		0 to 3	256
Memory grants	A/B and C/D				2
Multicast grants	A, B, C, and D			0 to 3	16
Send grants		0 to 31		0 to 3	128
Subport output queue grants		0 to 31	0 to 3	0 to 3	512
Subport multicast grants		0 to 31		0 to 3	128

3.8 Packet Reception

Packet reception on a particular input port is not synchronized with packet reception on the other input ports. When a packet arrives at an input port, the input controller analyzes the packet header and extracts various fields. If the header parity is incorrect, the entire packet is discarded. This error is reported via the header parity error bit in the *Status Register* (page 120), and the affected port is identified in the *Header Parity Error Register* (page 132). An interrupt is generated unless the error is masked with the header parity error bit in the *Interrupt Mask Register* (page 122). If the header parity is correct, the packet type is analyzed and the flow control information is extracted. Further processing depends on the packet type and is discussed in *Sections 3.8.1* and *3.8.2*.

In speed expansion configurations, the input controller on the master device conducts the packet header analysis and extraction, and then forwards the packet control information to the input controllers on the slave devices.

Note: A multicast packet is routed to only one subswitch element and stored only once in the shared memory; however, its shared memory address is enqueued in each output queue designated by the destination bitmap. In the 256-Gbps configuration, all destination output ports must be within the 0 to 7 port range or the

8 to 15 port range. In the 512-Gbps configuration, all destination output ports must be within the 0 to 15 port range or the 16 to 31 port range. If an ingress packet destination bitmap contains output ports in both ranges, the attached device duplicates the packet. Each output port transmits a multicast packet according to a first-first-out queuing structure and, consequently, the multicast packet is not necessarily transmitted at the same time on every port. A multicast packet has only one priority, which applies to all its destinations.

3.8.1 Data Packet Reception

PowerPRS Q-64G input controllers discard an ingress data packet when:

- All packet destination output ports are disabled.
- The input controller does not have a store address available for packet storage in the shared memory. This flow control error is reported via the no address interrupt bit in the *Status Register* (page 120). This bit generates an interrupt unless the error is masked with the corresponding no address interrupt bit in the *Interrupt Mask Register* (page 122). This error occurs only if the shared memory threshold is programmed incorrectly, or if the attached device does not respond to the memory grant information.
- Output queue grants or memory grants have not been issued, as required. This flow control error is reported via the flow control violation bit in the *Status Register* (page 120). This bit generates an interrupt unless the error is masked with the corresponding flow control violation bit in the *Interrupt Mask Register* (page 122). This error occurs only if the attached device does not respond to the ingress flow control information.
- The destination bitmap value is all zeros after packet filtering.
- The input controller best-effort discard filter is enabled and the discard condition is met (see *Section 3.4.6.3 Best-Effort Discard Filters* on page 51).

3.8.2 Control and Service Packet Reception

When the PowerPRS Q-64G receives a control packet or command service packet, the packet is stored in input controller internal registers (each input controller can store one control packet or command service packet per channel). The bit corresponding to the input port is set in the appropriate *Ingress Control Packet or Command Service Packet Received Register* (page 136), and the control packet received or command service packet received bit is set in the *Status Register*. After the packet payload is transferred to the local processor using the *Ingress Control Packet and Service Packet Source Register* (page 137) and the *Ingress Control Packet and Service Packet Payload Registers* (page 138), a new control packet or command service packet can be stored. If an input controller receives a control packet or command service packet before the previous one has been processed, the input controller discards the incoming packet and sets the control packet discard bit or command service packet discard bit in the *Status Register*.

When the PowerPRS Q-64G receives an event service packet, the bit corresponding to the input port is set in the appropriate *Ingress Event Service Packet Received Register* (page 139). When the value of the *Ingress Event Service Packet Received Register* is equal to the value of the *Ingress Event Service Packet Mask Register* (page 139), the appropriate “all event service packets received” bit is set in the *Status Register* (page 120).

3.9 Packet Transmission

3.9.1 Output Port Servicing

PowerPRS Q-64G output ports are designed for continuous packet transmission. Packet transmission starts at a fixed point in time, which is synchronized by the internal sequencer (see *Table 3-26).* If neither a data packet nor a control packet is available for transmission on an output port and the send grant is removed, the port transmits an idle packet. Control packets are always transmitted on an output port before any other packets in the shared memory destined for that port, but they do not affect the performance of high-priority traffic. Control packet transmission is relatively infrequent because the local processor access is slow compared to the data packet traffic rate. In the external speed expansion configuration, packet transmission starts at the same time on the master and slave output controllers.

Table 3-26. Packet Transmission Time

Internal Sequencer Cycle (8-ns)	Output Port Number			
0	0, 16	1, 17	8, 24	9, 25
1	2, 18	3, 19	10, 26	11, 27
2	4, 20	5, 21	12, 28	13, 29
3	6, 22	7, 23	14, 30	15, 31
4 to 7 for an 8-byte LU 4 to 8 for a 9-byte LU 4 to 9 for a 10-byte LU	–			

3.9.2 Control and Service Packet Transmission

The local processor prepares control packet and service packet payload using the *Egress Control Packet and Service Packet Payload Registers* (page 134) and specifies the destination bitmap with the *Egress Control Packet and Service Packet Destination Register* (page 135). Control packet or service packet transmission is requested when the *Egress Control Packet and Service Packet Destination Register* is written.

3.9.3 Idle Packet Transmission

An output port transmits idle packets only when there are no data, control, or service packets available, or when the send grant is removed. Idle packets are also transmitted to support switchover (see *Section 3.11 Switchover Support* on page 57).

3.9.4 Look-Up Table

The look-up table allows the byte transmission sequence of egress data, control, and service packets to be rearranged. It identifies if and when, and how many times, a data byte will be sent. *Table 3-27* presents an example of how the bytes can be rearranged. Note that any byte can be exchanged with or overwritten by another byte.

All output ports use the same look-up table, and the byte transmission sequence in the look-up table applies to all data streams. The *Look-Up Table Registers* (page 146) provide indirect access to the look-up table.

Table 3-27. Example of Byte Reordering Using the Look-Up Table

Byte Sequence before Ordering	Look-Up Table Entry Sequence	Byte Sequence after Ordering
Byte 0	3	Byte 3
Byte 1	4	Byte 4
Byte 2	5	Byte 5
Byte 3	3	Byte 3
Byte 4	4	Byte 4
Byte 5	5	Byte 5
Byte 6	9	Byte 9
Byte 7	8	Byte 8
Byte 8	8	Byte 8
Byte 9	0	Byte 0

3.10 Side Communication Channel

A four-bit side communication channel (SCC) allows communication between the attached devices and the local processor. SCC information requires minimal PowerPRS Q-64G processing, and is transferred in-band in the idle packet master LU (in byte 6, bits 0:3 and 4:7).

On the path from the attached devices to the PowerPRS Q-64G, an attached device inserts SCC information into all ingress idle packets. An attached device can generate an idle packet to guarantee that an information change is propagated in a minimum amount of time. When the PowerPRS Q-64G receives an idle packet, it extracts and compares bits 0:3 and 4:7. If the values are identical, an internal register that contains this information is refreshed, and the information is made available through the read-only *Side Communication Channel Input Reporting Registers* (page 133).

On the path from the PowerPRS Q-64G to the attached devices, the PowerPRS Q-64G inserts SCC information from four input pins (SCCIn[0:3]) into all egress idle packets. All output ports send the same SCC information. The PowerPRS Q-64G automatically generates an idle packet to all the ports as soon as an edge is detected on the SCCIn pins to guarantee that the information change is propagated in a minimum period of time.

3.11 Switchover Support

In redundant switch-plane operation, PowerPRS Q-64G switchover support is provided through a color mechanism. This mechanism conducts scheduled switchovers without packet loss. During normal operation, data packets and idle packets are coded red. Red traffic includes data packets with direct filtering, and link-liveness packets with either direct filtering or reverse filtering according to the mask set in the *Bitmap Filter Register* (page 125). By setting reverse filtering in the packet protection field, attached devices can send link-liveness packets to the ports on the backup switch plane (and thereby supervise the backup path). The PowerPRS Q-64G registers and bits involved in the switchover process are described in *Table 3-28* and in the associated register descriptions in *Section 5.*

Note: The PowerPRS Q-64G processes the packets used to test link liveness between the ingress and egress protocol engines as data packets. Event-1 service packets are used to test link liveness between the PowerPRS Q-64G and the PowerPRS C192.

Table 3-28. Registers and Bits Used for Switchover Support

Register or Register Field	Description
Idle Color Force Bit in the *Color Command Register* (page 127)	1 Transmits all egress idle packets with the color specified by the idle color bit, regardless of the expected color bit setting. When the color mechanism is not used, this bit must be set to '1'. 0 Enables the switchover mechanism to determine the color of egress idle packets by setting the expected color bit. The PowerPRS Q-64G sends idle packets of the color specified by the expected color bit on output port *n,* if an expected color packet has been received on all active inputs since the color clear command was last issued and output queue *n* is empty. Otherwise, the PowerPRS Q-64G sends idle packets of the opposite color on output port *n*. The color clear command is sent through the *Color Command Register*.
Idle Color Bit in the *Color Command Register*	Specifies the color assigned to all idle packets when the idle color force bit is set to '1': 0 Blue idle packets 1 Red idle packets In this case, the PowerPRS Q-64G-generated idle packets will not change color during normal operation.
Expected Color Bit in the *Color Command Register*	Specifies the expected color of incoming packets after a color clear command is initiated: 0 Blue packets 1 Red packets
Color Clear Bit in the *Color Command Register*	Processed as a command (action is taken on the rising edge) to clear the idle packet color state machine in preparation for packet color-change detection.
Color Detection Disable Register (page 125)	When a bit is active, disables the input port color detection mechanism and sets the corresponding bit in the *Expected Color Received Register*. This mask reports if an input port is enabled and active during the color-based switchover process.
Expected Color Received Register (page 126)	1 Either the expected color has been received on the input port since the last color clear command or the corresponding bit is set in the *Color Detection Disable Register*. 0 The opposite color is still being received on the input port.
Bitmap Filter Register (page 125)	Specifies the mask applied to the ingress packet destination bitmap for switchover support and load balancing in redundant switch-plane operation. Application of the bitmap filter depends on the packet protection field (bits 2:3) of the packet qualifier byte, H0. The incoming packet bitmap is logically ANDed with either a specified mask (red active packets) or its complement value (red backup packets), or it is left unfiltered (blue packets; see *Table 3-29).* For more information about operating this mask, see *Table 3-5* in *Section 3.3.4 Ingress Data Packet and Control Packet Format* on page 32.

Table 3-29. Ingress Data Packet Protection Field

Protection Field Value	Packet Color	Filtering
01	Red (backup)	Packet destination bitmap is bitwise ANDed with the bitwise complement of the bitmap filter.
10	Red (active)	Packet destination bitmap is bitwise ANDed with the bitmap filter.
11	Blue (unfiltered)	Packet destination bitmap is used unfiltered.

3.11.1 Scheduled Switchover Process

Redundant switch planes operate under one of two conditions:

- When one switch plane is active and the other is a backup. Data traffic flows only through the active switch path.
- When both switch planes are operating under load balancing. Data traffic is split between the two switch planes, which have complementary values in their *Bitmap Filter Registers* (page 125).

When two switch planes are initially operating under load balancing, the switchover process includes three phases:

1. Rerouting traffic to one switch plane
2. Modifying the load-balancing configuration
3. Resuming traffic on both switch planes

The scheduled switchover process for the active/backup initial operating condition is very similar. Minor differences are identified below.

3.11.1.1 Phase 1: Rerouting Traffic to One Switch Plane

Before the scheduled switchover begins, data traffic is routed through both switch planes under the load-balancing configuration specified in the *Bitmap Filter Registers.* During phase 1 of the switchover process, traffic is rerouted so that it flows through only one switch plane. For this discussion, the Y switch plane is dropped and the X switch plane remains active. Phase 1 is initiated on red traffic and is complete when all traffic is blue.

To reroute traffic to one switch plane:

1. Because all current traffic (idle and data) is red, the local processor configures each PowerPRS Q-64G to detect the color blue by changing the expected color bit and issuing a color clear command.
2. All ingress devices change the packet qualifier byte of their incoming packets to change the color from red (normal traffic) to blue (no filtering). In addition:
   - On the Y path, the ingress devices send all their buffered red data packets to the switch core, and then start generating blue idle packets to the switch core.
   - On the X path, all ingress devices send all their buffered red packets (regardless of priority) to the switch core, and then begin sending their blue data packets to the switch core.

   Simultaneously, the egress devices block data packet reception from the X switch plane by locking their peer buffer (which connects the X and Y paths). This prevents blue traffic reception before red traffic is fully exhausted. (If the switch planes were operating in an active/backup condition, this switch plane would be the backup, and the only traffic would be link-liveness packets.)
3. When at least one blue packet (idle or data) has been received on each active input port of the X or Y switch plane, then all the red data packets have been delivered to that switch plane and all the active input ports will be receiving only blue packets (either idle or data). Each of the two local processors attached to the serial host interface (SHI) is informed, through polling, that its switch core is detecting only blue packets.

   Note: The *Expected Color Received Register* (page 126) reports the receipt of a blue packet since the last color clear command on each port that has not been tagged as inactive by the *Color Detection Disable Register* (page 125). The color blue was set by the expected color bit. When all bits are set in the *Expected Color Received Register,* the switch core is detecting only blue packets.
4. On the Y switch plane, when at least one blue packet has been detected on each active input port and output queue *n* is empty, the PowerPRS Q-64G begins to continuously generate blue (rather than red) idle packets to port *n*. On the X switch plane, when at least one blue packet has been detected on each active input port, the PowerPRS Q-64G begins generating blue (rather than red) idle packets, as necessary.

   At this point, all egress packets from both switch planes are blue (idle packets or data packets on the X path, and idle packets on the Y path). All data traffic is blue (unfiltered) and carried on the X path.
5. When all active egress devices have detected the arrival of a blue idle packet from the Y switch core and these devices have no more packets to send from their packet buffer queue for that switch plane to the

attached traffic manager, then they unlock their peer buffer. This step unblocks traffic transmission from the X switch plane, which has blue packets waiting for transmission to the attached devices.

Phase 1 of the switchover is complete for the entire switch fabric. The egress devices convey this status to their attached processor.

3.11.1.2 Phase 2: Modifying the Load-Balancing Configuration

When phase 1 of the scheduled switchover is complete, all traffic through both switch planes is blue and data traffic flows through only one plane. Both local processors can now safely modify the content of the *Bitmap Filter Register* (page 125) in accordance with the new configuration parameters, which may specify new port assignments for a different load-balancing configuration.

3.11.1.3 Phase 3: Resuming Traffic on Both Switch Planes

Phase 3 of the switchover starts the new load-balancing configuration. This phase is similar to phase 1, except it is initiated on blue traffic and is complete when all traffic is red. For this discussion, the Y switch plane is dropped and the X switch plane remains active. During phase 3, split traffic is resumed on both switch planes.

To resume traffic on both switch planes:

1. Because all current traffic (idle and data) is blue, the local processor configures each PowerPRS Q-64G to detect the color red by changing the expected color bit and issuing a color clear command.
2. All ingress devices stop changing the packet qualifier byte of their incoming packets so that the packet color remains red. In addition:
   - On the Y path, all ingress devices start generating red (rather than blue) idle packets to the switch core.
   - On the X path, all ingress devices send all their buffered blue packets (regardless of priority) to the switch core, and then begin sending their red data packets to the switch core.

   Simultaneously, the egress devices block data packet reception from the Y switch plane by locking their peer buffer. This step prevents red traffic reception before blue traffic is fully exhausted.
3. When at least one red packet (idle or data) has been received on each active input port of the X or Y switch plane, then all the blue data packets have been delivered to that switch plane and all the active input ports will be receiving only red packets (either idle or data). Each of the two local processors attached to the SHI is informed, through polling, that its switch core is detecting only red packets.
4. On each switch plane, when at least one red packet has been detected on each active input port, the PowerPRS Q-64G on that switch plane begins to generate red (rather than blue) idle packets, as necessary. At this point, all egress packets from both switch planes are red.
5. When all active egress devices have detected the arrival of a red idle packet from the X switch core and these devices have no more packets to send from their packet buffer queue for that switch plane to the attached traffic manager, then they unlock their peer buffer. This step unblocks traffic transmission from the Y switch plane, which has red packets waiting for transmission to the attached devices.

   Switchover is complete for the entire switch fabric. The egress devices convey this status to their attached processor.

4. Programming Interface

The serial host interface (SHI) is the programming interface between the local processor and the PowerPRS Q-64G. It provides access to all PowerPRS Q-64G internal resources through four signals:

- SHIClockIn:
- SHISelectIn#
- SHISerialDataIn
- SHISerialDataOut

The SHI and the SHI internal logic are synchronized to the SHI clock, which is asynchronous to the system clock (see *Section 7.2.2 SHI Signals* on page 166). The *SHI Instruction Register* and the power-on registers (addresses x‘01’ to x‘08’) are reset when the PowerOnResetIn# signal is activated.

4.1 SHI Instruction Register

An instruction scanned into the SHI is decoded into four parts:

- Data field (32 bits)
- Address field (7 bits)
- OpCode bit
- Parity bit

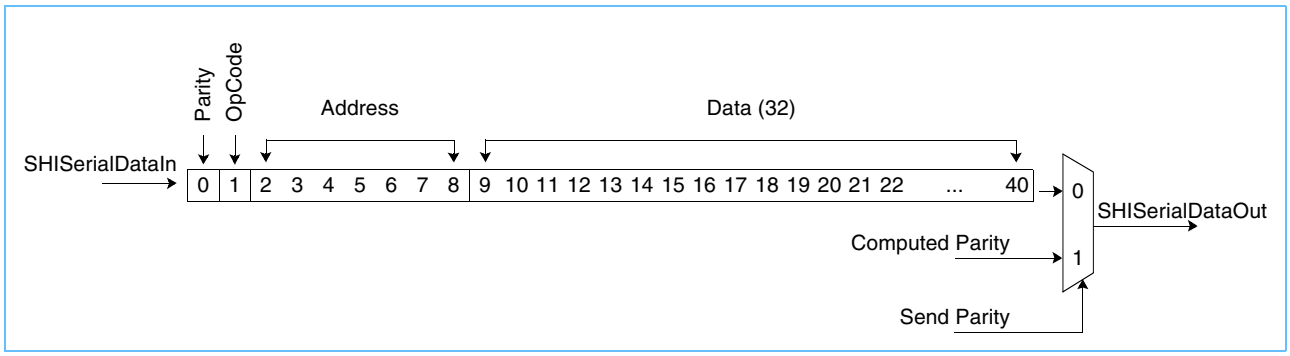

Bit(s)	Field Name	Description
0	Parity	If the parity is correct, executes the instruction. If the parity is incorrect, inhibits the instruction.
1	OpCode	Specifies the SHI command to be executed. See *Table 4-1* for descriptions of these commands.
2:8	Address	Specifies the internal register to be read or written.
9:40	Data (32)	Contains the value that will be written or the value that has been returned.

Table 4-1. SHI OpCode Commands

OpCode	Operation	Type	Address	Description
0	No-Op	33-bit	x‘00’	Required after a read register command. Clears the data out of the *SHI Instruction Register* (page 61).
0	Read register	9-bit	x‘01’ to x‘7F’	Loads the content of the register specified by the address field in the *SHI Instruction Register* into the data field of the *SHI Instruction Register*. This command requires a no-op command (OpCode ‘0’) to send the read result and parity over the SHISerialDataOut signal and clear the data field in the *SHI Instruction Register*.
1	Write register	41-bit	x‘01’ to x‘7F’	Writes the value of the data field in the *SHI Instruction Register* into the register specified by the address field in the *SHI Instruction Register*.
Note: Each read register command must be followed by the no-op command.				

4.2 SHI Instruction Execution

An SHI instruction is invoked when the SHISelectIn# input is set to ‘0’. During execution, data transmitted over SHISerialDataIn is shifted into the *SHI Instruction Register.* Shifted serial data must begin with the least significant bit and end with the most significant bit of the instruction to be executed. This scan operation is synchronized with SHIClockIn. Instructions always execute one SHI clock cycle after the SHISelectIn# signal changes from an active to an inactive state.

4.3 SHI Parity Checking

Each instruction scanned into the *SHI Instruction Register* has one bit of parity protection. The parity bit is the most significant bit (bit 0) of the *SHI Instruction Register,* and is the last bit scanned.

The *SHI Instruction Register* checks whether an incoming instruction has the required odd parity. If a parity error is detected on a received instruction, the execution of that instruction is inhibited and the SHI parity error bit is set in the *Status Register* (page 120). All SHI command bits are protected by the parity bit (that is, if SHISelectIn# is active during *n* SHI clock cycles, the parity is checked on *n* bits).

4.4 SHI Parity Generation

Both incoming and outgoing data carry odd parity. This parity is computed for each SHI clock cycle when the SHISelectIn# signal is active. The computed parity is sent on SHISerialDataOut when the SHISelectIn# signal is deactivated.

5. Register Descriptions

This section describes the registers, including field definitions, that provide the mechanism for PowerPRS Q-64G configuration specification and status reporting.

Table 5-1 identifies each register and provides the page number where the corresponding description is located. In the register descriptions, *reserved* bits and addresses are not implemented. Reserved bits return '0' when read and ignore all write values.

Registers x'01' to x'08' are implemented in the serial host interface (SHI) logic and are reset by activating the PowerOnResetIn# input signal. These registers can be accessed before the phase-locked loop (PLL) is started or the flush is complete. All the bits in the remaining registers are set to '0' during a flush, unless otherwise specified.

Note: In the 256-Gbps internal speed expansion configuration, the ports in each of the four PowerPRS Q-64Gs are paired as follows: physical ports 0 and 1 are paired to form logical port 0, physical ports 2 and 3 are paired to form logical port 1, and so forth. Because all the registers report physical information, information to or from logical queue *L* is reported using bit $L \times 2$ for the registers in which a bit corresponds to a port. For example, to enable the output queue for logical port 3, bit 6 in the *Output Queue Enable Register* (page 124) must be set.

Table 5-1. Register Map (Page 1 of 4)

Register Name	Address	Access	Page
Reserved for no-op operation	x'00'		
SHI Internal Registers: x'01' to x'08'			
Internal PLL Programming Register	x'01'	Read/Write	67
Internal PLL Status Register	x'02'	Read Only	67
Unilink PLL Programming Register	x'03'	Read/Write	68
Unilink PLL Status Register	x'04'	Read Only	68
Reset Register	x'05'	Read/Write	69
BIST Counter Register	x'06'	Read/Write	70
BIST Data Register	x'07'	Read/Write	70
BIST Select Register	x'08'	Read/Write	71
Unilink Programming Registers: x'09' to x'32'			
UL Global Register	x'09'	Read/Write	72
UL Errors Register	x'0A'	Read Only	74
UL PLL Unlock Register	x'0B'	Read Only	76
UL TxPort Driver Enable Register	x'0C'	Read/Write	78
UL TxPort Attachment Enable Register	x'0D'	Read/Write	78
UL TxPort Parameters Register	x'0E'	Read/Write	79
UL TxPort BIST Request Register	x'0F'	Read/Write	80
UL TxPort BIST Error Register	x'10'	Read Only	80

Table 5-1. Register Map (Page 2 of 4)

Register Name	Address	Access	Page
UL TxPort Reset BIST Error Register	x‘11’	Read/Write	81
UL RxPort Attachment Enable Register	x‘12’	Read/Write	81
UL RxPort Byte Alignment Done Register	x‘13’	Read Only	82
UL RxPort K28.5 Spacing OK Register	x‘14’	Read Only	82
UL RxPort LU Deskew Register	x‘15’	Read/Write	83
UL RxPort Data Mode Register	x‘16’	Read/Write	84
UL RxPort Data Valid Register	x‘17’	Read Only	84
UL RxPort Signal Lost Register	x‘18’	Read Only	85
UL RxPort Invalid K Character Register	x‘19’	Read/Clear	85
UL RxPort Synchronization Lost Register	x‘1A’	Read/Clear	86
UL RxPort Code Violation Register	x‘1B’	Read/Clear	86
UL RxPort BIST Request Register	x‘1C’	Read/Write	87
UL RxPort BIST Error Register	x‘1D’	Read Only	87
UL RxPort Reset BIST Error Register	x‘1E’	Read/Write	88
UL RxPort BIST Wrap Register	x‘1F’	Read/Write	88
UL TxSpex Bus Driver Enable Register	x‘20’	Read/Write	89
UL TxSpex Bus Attachment Enable Register	x‘21’	Read/Write	90
UL TxSpex Bus Parameters Register	x‘22’	Read/Write	91
UL TxSpex Bus BIST Request Register	x‘23’	Read/Write	93
UL TxSpex Bus BIST Error Register	x‘24’	Read Only	94
UL TxSpex Bus Reset BIST Error Register	x‘25’	Read/Write	95
UL RxSpex Bus Attachment Enable Register	x‘26’	Read/Write	96
UL RxSpex Bus Byte Alignment Done Register	x‘27’	Read Only	97
UL RxSpex Bus K28.5 Spacing OK Register	x‘28’	Read Only	98
UL RxSpex Bus Latency Programming Register	x‘29’	Read/Write	99
UL RxSpex Bus Data Mode Register	x‘2A’	Read/Write	101
UL RxSpex Bus Signal Lost Register	x‘2B’	Read Only	102
UL RxSpex Bus Invalid K Character Register	x‘2C’	Read/Clear	103
UL RxSpex Bus Synchronization Lost Register	x‘2D’	Read/Clear	104
UL RxSpex Bus Code Violation Register	x‘2E’	Read/Clear	105
UL RxSpex BIST Request Register	x‘2F’	Read/Write	106
UL RxSpex Bus BIST Error Register	x‘30’	Read Only	107

Table 5-1. Register Map (Page 3 of 4)

Register Name	Address	Access	Page
UL RxSpex Bus Reset BIST Error Register	x‘31’	Read/Write	108
UL RxSpex BIST Wrap Register	x‘32’	Read/Write	109
Functional Registers: x‘33’ to x‘50’			
Configuration 0 Register	x‘33’	Read/Write	110
Configuration 1 Register	x‘34’	Read/Write	112
Threshold Access Register	x‘35’	Read/Write	114
Credit Table Access Register	x‘36’	Read/Write	116
Best-Effort Resources Access Register	x‘37’	Read/Write	118
Status Register	x‘38’	Read/Clear	120
Interrupt Mask Register	x‘39’	Read/Write	122
Output Queue Enable Register	x‘3A’	Read/Write	124
Input Controller Enable Register	x‘3B’	Read/Write	124
Bitmap Filter Register	x‘3C’	Read/Write	125
Color Detection Disable Register	x‘3D’	Read/Write	125
Expected Color Received Register	x‘3E’	Read Only	126
Color Command Register	x‘3F’	Read/Write	127
Bitmap Mapping Register	x‘40’	Read/Write	128
Output Queue Status Registers	x‘41’ to x‘48’	Read/Clear	129
Best-Effort Discard Alarm Register	x‘49’	Read/Clear	131
Send Grant Violation Register	x‘4A’	Read/Clear	131
Header Parity Error Register	x‘4B’	Read/Clear	132
Flow Control Violation Register	x‘4C’	Read/Clear	132
Side Communication Channel Input Reporting Registers	x‘4D’ to x‘50’	Read Only	133
Control Packet and Service Packet Transmission Registers: x‘51’ to x‘57’			
Egress Control Packet and Service Packet Payload Registers	x‘51’ to x‘56’	Read/Write	134
Egress Control Packet and Service Packet Destination Register	x‘57’	Read/Write	135
Control Packet and Service Packet Reception Registers: x‘58’ to x‘63’			
Ingress Control Packet or Command Service Packet Received Registers	x‘58’ to x‘5B’	Read Only	136
Ingress Control Packet and Service Packet Source Register	x‘5C’	Read/Write	137
Ingress Control Packet and Service Packet Payload Registers	x‘5D’ to x‘5F’	Read Only	138
Ingress Event Service Packet Received Registers	x‘60’ and x‘61’	Read/Clear	139
Ingress Event Service Packet Mask Registers	x‘62’ and x‘63’	Read/Write	139

Table 5-1. Register Map (Page 4 of 4)

Register Name	Address	Access	Page
Debug Facilities Registers: x‘64’ to x‘7E’			
Debug Bus Select Register	x‘64’	Read/Write	140
Send Grant Disable Register	x‘65’	Read/Write	142
Force Send Grant Register	x‘66’	Read/Write	142
Send Grant Status Registers	x‘67’ to x‘6A’	Read Only	143
Subswitch Element Occupancy (1) Registers	x‘6B’ to x‘6E’	Read Only	144
Subswitch Element Occupancy (2) Registers	x‘6F’ x‘72’	Read Only	145
Look-Up Table Registers	x‘73’ and x‘74’	Read/Write	146
Blue Idle Packet or Data Packet Received Register	x‘75’	Read/Clear	147
Red Idle Packet or Data Packet Received Register	x‘76’	Read/Clear	147
Miscellaneous Debug Register	x‘77’	Read/Write	148
Force Packet Capture Ports Register	x‘78’	Read/Write	150
Force Packet Capture Header Register	x‘79’	Read/Write	151
Force Packet Capture Mask Register	x‘7A’	Read/Write	152
Packet Captured Registers	x‘7B’ and x‘7C’	Read Only	152
Unilink Debug Control Register	x‘7D’	Read/Write	153
Unilink Force Error Register	x‘7E’	Read/Write	154
Reserved	x‘7F’		

5.1 SHI Internal Registers

5.1.1 Internal PLL Programming Register

Address x‘01’

Access Type Read/Write

Reset Value ‘1000 0000 0000 0000 0000 0000 0000 0000’

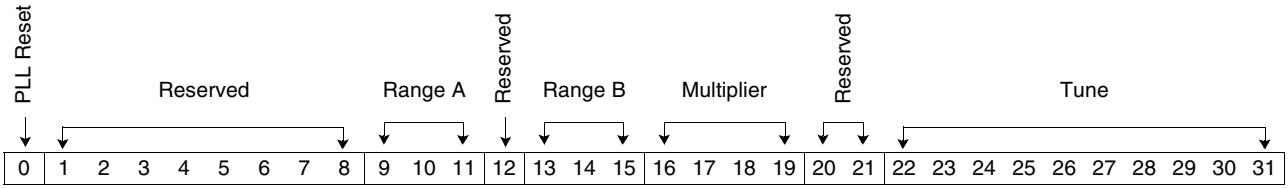

Bit(s)	Field Name	Description
0	PLL Reset	When set to ‘1’, holds the PLL in a reset state. Should not be released until the reference clock is stable and the PLL is programmed correctly.
1:8	Reserved	Reserved.
9:11	Range A (2:0)	Used to select the PLL output frequency. Must be set to ‘111’.
12	Reserved	Reserved.
13:15	Range B (2:0)	Not used. Must be set to ‘111’.
16:19	Multiplier (3:0)	Defines the PLL feedback divider. Must be set to ‘0001’.
20:21	Reserved	Reserved.
22:31	Tune (9:0)	Used to optimize PLL stability and jitter. Must be set to ‘01 0011 1000’.

5.1.2 Internal PLL Status Register

Address x‘02’

Access Type Read Only

Reset Value ‘u000 0000 0uuu uuuu uuuu uuuu uuuu uuuu’, where ‘u’ = undefined

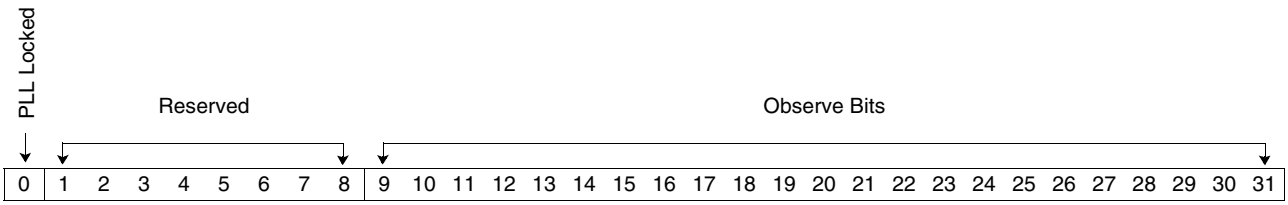

Bit(s)	Field Name	Description
0	PLL Locked	When set to ‘1’, the feedback clock is in phase with the reference clock.
1:8	Reserved	Reserved.
9:31	Observe Bits	Used for testing (23 bits [22:0]).

5.1.3 Unilink PLL Programming Register

Address x‘03’

Access Type Read/Write

Reset Value ‘1000 0000 0000 0000 0000 0000 0000 0000’

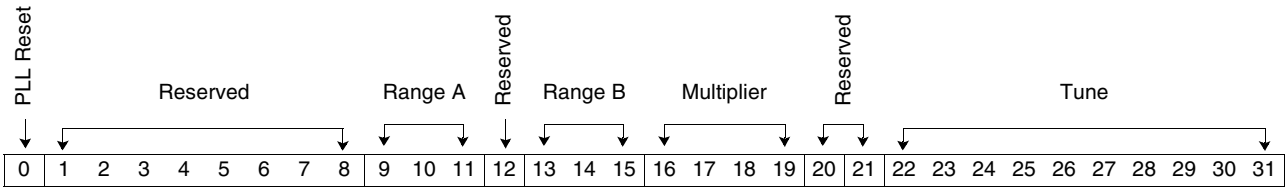

Bit(s)	Field Name	Description
0	PLL Reset	When set to ‘1’, holds the PLL in a reset state. Should not be released until the reference clock is stable and the PLL is programmed correctly.
1:8	Reserved	Reserved.
9:11	Range A (2:0)	Used to select the PLL output frequency. Must be set to ‘111’.
12	Reserved	Reserved.
13:15	Range B (2:0)	Not used. Must be set to ‘111’.
16:19	Multiplier (3:0)	Defines the PLL feedback divider. Must be set to ‘0101’.
20:21	Reserved	Reserved.
22:31	Tune (9:0)	Used to optimize PLL stability and jitter. Must be set to ‘01 1011 1110’.

5.1.4 Unilink PLL Status Register

Address x‘04’

Access Type Read Only

Reset Value ‘u000 0000 0uuu uuuu uuuu uuuu uuuu uuuu’, where ‘u’ = undefined

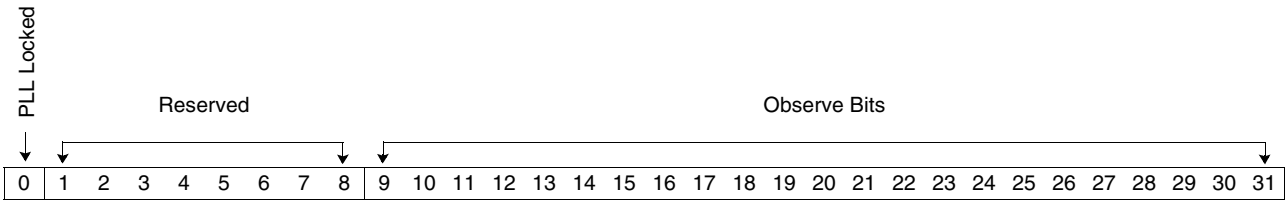

Bit(s)	Field Name	Description
0	PLL Locked	When set to ‘1’, the feedback clock is in phase with the reference clock.
1:8	Reserved	Reserved.
9:31	Observe Bits	Used for testing (23 bits [22:0]).

5.1.5 Reset Register

Address x‘05’

Access Type Read/Write

Reset Value ‘1110 0010 0000 0000 0000 0000 0000 0uuu’, where ‘u’ = undefined

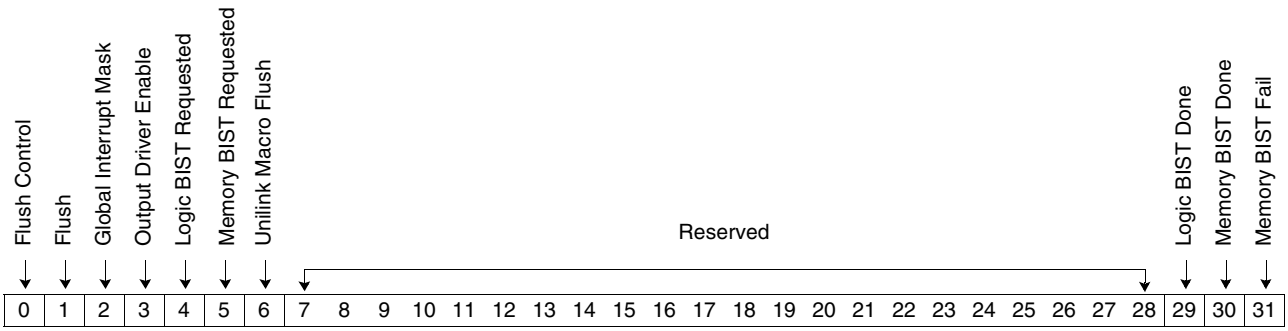

Bit(s)	Field Name	Description
0	Flush Control	When set to ‘1’, keeps the device control section in a reset state. For slave devices, this bit must be set to ‘1’.
1	Flush	When set to ‘1’, keeps the functions common to the master and slave devices in a reset state.
2	Global Interrupt Mask	1 Disables event- and error-generated interrupts to the local processor. The device interrupt signal (active low) is tristated and pulled up with an external resistor. The *Status Register* (page 120) bits are asserted when the corresponding events or errors occur. 0 Enables event- and error-generated interrupts to the local processor.
3	Output Driver Enable	1 Enables all device drivers until another configuration disables them. 0 Disables (tristates) all device drivers except for the SHISerialDataOut driver.
4	Logic BIST Requested	When set to ‘1’, enables the built-in self-test (BIST) controller to start executing the internal logic BIST as soon as the flush bit is deactivated. This bit can be asserted only while the flush bit is active. Logic BIST completion is reported via the logic BIST done bit (bit 29). See *Section 6.4 Logic BIST Execution Sequence* on page 160 for more information.
5	Memory BIST Requested	When set to ‘1’, enables the BIST controller to start executing the memory BIST as soon as the flush bit is deactivated. This bit can be asserted only while the flush bit is active. Memory BIST completion and results are reported via the memory BIST done and memory BIST fail bits (bits 30 and 31). See *Section 6.5 Memory BIST Execution Sequence* on page 161 for more information.
6	Unilink Macro Flush	When set to ‘1’, keeps all Unilink macros in a reset state.
7:28	Reserved	Reserved.
29	Logic BIST Done (read only)	Set to ‘1’ when the BIST controller completes internal processing after a logic BIST request command.
30	Memory BIST Done (read only)	Set to ‘1’ when the BIST controller completes internal processing after a memory BIST request command.
31	Memory BIST Fail (read only)	Set to ‘1’ when, after completion of the memory BIST process, at least one memory BIST check failed on at least one RAM. This bit is valid only when the memory BIST done bit is asserted.

5.1.6 BIST Counter Register

Address x‘06’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

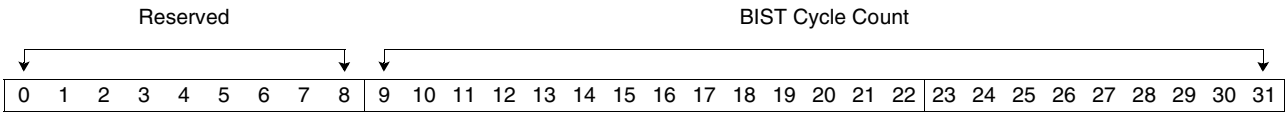

Bit(s)	Field Name	Description
0:8	Reserved	Reserved.
9:31	BIST Cycle Count	Specifies the number of BIST cycles to be performed.

5.1.7 BIST Data Register

This register, along with the *BIST Select Register* (page 71), provides indirect read/write access to the internal pseudorandom pattern generator (PRPG) and multiple-input signature (MISR) registers.

Address x‘07’

Access Type Read/Write

Reset Value Undefined

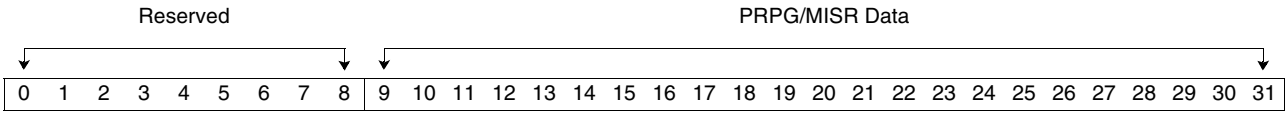

Bit(s)	Field Name	Description
0:8	Reserved	Reserved.
9:31	PRPG/MISR Data	Contains the data that has been or will be exchanged using the settings provided in the *BIST Select Register*.

5.1.8 BIST Select Register

This register, along with the *BIST Data Register* (page 70), provides indirect read/write access to the internal PRPG and MISR registers.

Write access to an internal PRPG or MISR register requires two SHI commands:

1. Write the BIST register select field of the *BIST Select Register* with the value specifying which internal PRPG or MISR register is to be accessed.
2. Write the *BIST Data Register* with the value desired for the internal PRPG or MISR register specified in step 1. The internal PRPG or MISR register is loaded.

Read access to an internal PRPG or MISR register requires two SHI commands:

1. Write the BIST register select field in the *BIST Select Register* with the value specifying which internal PRPG or MISR register is to be accessed.
2. Read the *BIST Data Register*. The value for the internal PRPG or MISR register specified in step 1 is returned.

Address x‘08’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

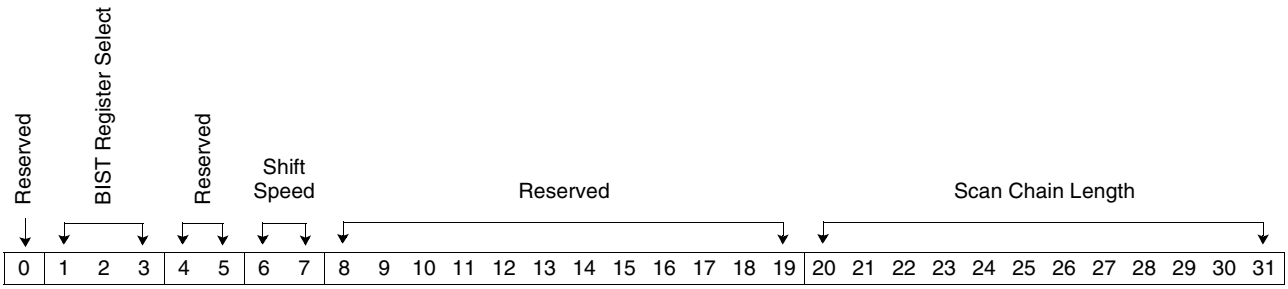

Bit(s)	Field Name	Description
0	Reserved	Reserved.
1:3	BIST Register Select	Specifies the BIST register: 000 PRPG0 100 MISR0 001 PRPG1 101 MISR1 010 PRPG2 110 MISR2 011 PRPG3 111 MISR3
4:5	Reserved	Reserved.
6:7	Shift Speed	Defines the delay between the A and B clock pulses while shifting occurs during the BIST: 00 8 ns 01 16 ns 10 24 ns 11 32 ns
8:19	Reserved	Reserved.
20:31	Scan Chain Length	Specifies the scan chain length.

5.2 Unilink Programming Registers

The *Unilink Programming Registers* are used to synchronize and supervise all Unilink port and speed-expansion bus logic. For information about using these registers, see *Section 6.2 Speed-Expansion Bus Initialization* on page 155 and *Section 6.3 Port Initialization and Operation* on page 159.

Unilink port logic contains 32 Unilink receivers and 32 Unilink transmitters. Each of the eight Unilink *receive* macros includes four Unilink receivers and one internal phase-locked loop (PLL), and each of the eight Unilink *transmit* macros includes four Unilink transmitters and one internal PLL.

Unilink speed-expansion bus logic contains 16 Unilink receivers and 16 Unilink transmitters (that is, 8 of each for the *ingress* speed-expansion bus and 8 of each for the *egress* speed-expansion bus).

Note: In the register notation that follows:

- UL TxPort = Unilink port transmitter
- UL RxPort = Unilink port receiver
- UL TxSpex Bus = Unilink speed-expansion bus transmitter
- UL RxSpex Bus = Unilink speed-expansion bus receiver

5.2.1 UL Global Register

This register is used to specify the Unilink global port and speed-expansion bus settings.

Address x‘09’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

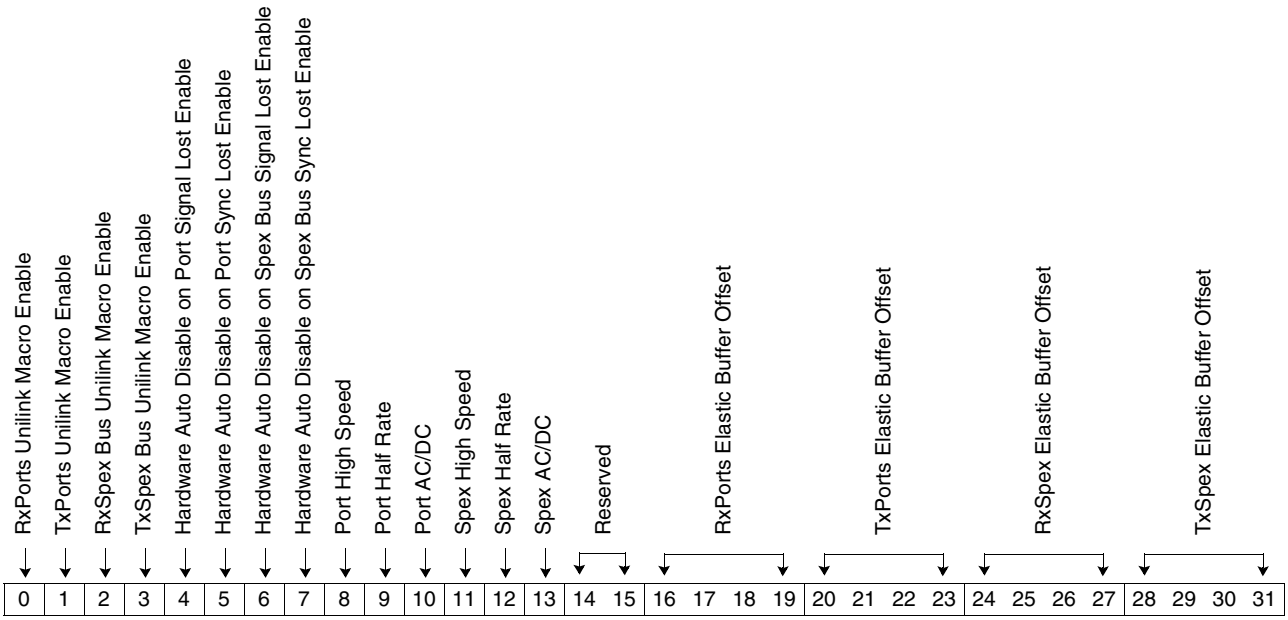

Bit(s)	Field Name	Description
0	RxPorts Unilink Macro Enable	When set to ‘0’, keeps the Unilink port receive macros in a reset state.
1	TxPorts Unilink Macro Enable	When set to ‘0’, keeps the Unilink port transmit macros in a reset state.
2	RxSpex Bus Unilink Macro Enable	When set to ‘0’, keeps the Unilink speed-expansion bus receive macros in a reset state.
3	TxSpex Bus Unilink Macro Enable	When set to ‘0’, keeps the Unilink speed-expansion bus transmit macros in a reset state.
4	Hardware Auto Disable on Port Signal Lost Enable	When set to ‘1’, enables the Unilink port attachment logic to reset the line when a signal is lost (used for testing). Must be set to ‘0’.
5	Hardware Auto Disable on Port Sync Lost Enable	When set to ‘1’, enables the Unilink port attachment logic to reset the line when synchronization is lost (used for testing). Must be set to ‘0’.
6	Hardware Auto Disable on Spex Bus Signal Lost Enable	When set to ‘1’, enables the Unilink speed-expansion bus attachment logic to reset the line when a signal is lost (used for testing). Must be set to ‘0’.
7	Hardware Auto Disable on Spex Bus Sync Lost Enable	When set to ‘1’, enables the Unilink speed-expansion bus attachment logic to reset the line when synchronization is lost (used for testing). Must be set to ‘0’.
8	Port High Speed	Reserved for testing. Must be set to ‘0’.
9	Port Half Rate	When set to ‘1’, the Unilink macro runs at half the nominal speed (used for testing). Must be set to ‘0’.
10	Port AC/DC	Selects the Unilink port receive macro coupling type: 0 AC coupling 1 DC coupling Must be set to ‘1’.
11	Spex High Speed	Reserved for testing. Must be set to ‘0’.
12	Spex Half Rate	When set to ‘1’, the Unilink macro runs at half the nominal speed (used for testing). Must be set to ‘0’.
13	Spex AC/DC	Selects the Unilink speed-expansion bus receive macro coupling type: 0 AC coupling 1 DC coupling Must be set to ‘1’.
14:15	Reserved	Reserved.
16:19	RxPorts Elastic Buffer Offset	Defines the offset between the Unilink clock deskew memory device read and write pointers. Must be set to 8.
20:23	TxPorts Elastic Buffer Offset	Defines the offset between the Unilink clock deskew memory device read and write pointers. Must be set to 8.
24:27	RxSpex Elastic Buffer Offset	Defines the offset between the Unilink clock deskew memory device read and write pointers. Must be set to x‘C’.
28:31	TxSpex Elastic Buffer Offset	Defines the offset between the Unilink clock deskew memory device read and write pointers. Must be set to 4.

5.2.2 UL Errors Register

This register is used to report the occurrence of certain Unilink errors. When a bit is set in this register, at least one bit is set in the corresponding error register.

Address x‘0A’

Access Type Read Only

Reset Value ‘1111 0000 0000 0000 0000 0000 0000 0000’

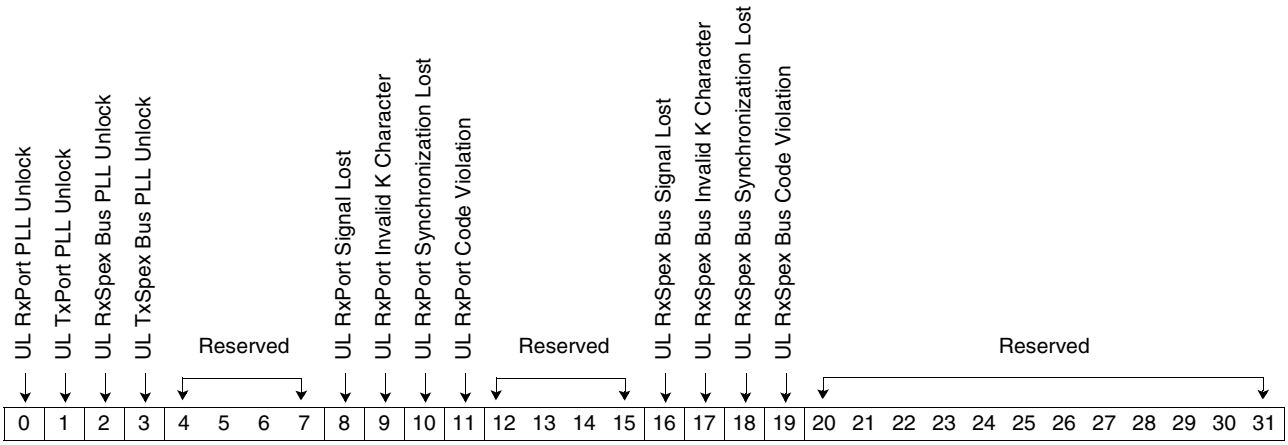

Bit(s)	Field Name	Description
0	UL RxPort PLL Unlock	When set to ‘1’, a PLL unlock error has occurred in a Unilink port receiver. See the *UL PLL Unlock Register* (page 76).
1	UL TxPort PLL Unlock	When set to ‘1’, a PLL unlock error has occurred in a Unilink port transmitter. See the *UL PLL Unlock Register* description.
2	UL RxSpex Bus PLL Unlock	When set to ‘1’, a PLL unlock error has occurred in a Unilink speed-expansion bus receiver. See the *UL PLL Unlock Register* description.
3	UL TxSpex Bus PLL Unlock	When set to ‘1’, a PLL unlock error has occurred in a Unilink speed-expansion bus transmitter. See the *UL PLL Unlock Register* description.
4:7	Reserved	Reserved.
8	UL RxPort Signal Lost	When set to ‘1’, a signal loss error has occurred in a Unilink port receiver. See the *UL RxPort Signal Lost Register* (page 85).
9	UL RxPort Invalid K Character	When set to ‘1’, an invalid K character error has occurred in a Unilink port receiver. See the *UL RxPort Invalid K Character Register* (page 85).
10	UL RxPort Synchronization Lost	When set to ‘1’, a synchronization loss error has occurred in a Unilink port receiver. See the *UL RxPort Synchronization Lost Register* (page 86).
11	UL RxPort Code Violation	When set to ‘1’, a code violation error has occurred in a Unilink port receiver. See the *UL RxPort Code Violation Register* (page 86).
12:15	Reserved	Reserved.
16	UL RxSpex Bus Signal Lost	When set to ‘1’, a signal loss error has occurred in a Unilink speed-expansion bus receiver. See the *UL RxSpex Bus Signal Lost Register* (page 102).
17	UL RxSpex Bus Invalid K Character	When set to ‘1’, an invalid K character error has occurred in a Unilink speed-expansion bus receiver. See the *UL RxSpex Bus Invalid K Character Register* (page 103).

Bit(s)	Field Name	Description
18	UL RxSpex Bus Synchronization Lost	When set to '1', a synchronization loss error has occurred in a Unilink speed-expansion bus receiver. See the *UL RxSpex Bus Synchronization Lost Register* (page 104).
19	UL RxSpex Bus Code Violation	When set to '1', a code violation error has occurred in a Unilink speed-expansion bus receiver. See the *UL RxSpex Bus Code Violation Register* (page 105).
20:31	Reserved	Reserved.

5.2.3 UL PLL Unlock Register

Address x‘0B’

Access Type Read Only

Reset Value ‘1111 1111 1111 1111 1111 1111 0000 0000’

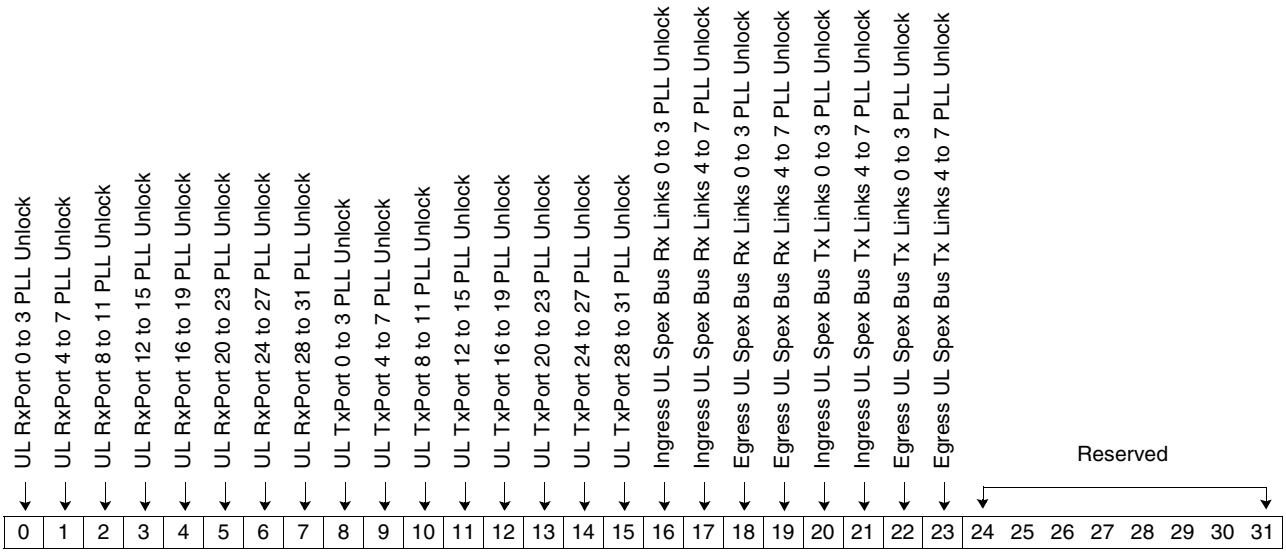

Bit(s)	Field Name	Description
0	UL RxPort 0 to 3 PLL Unlock	When set to ‘1’, the Unilink hard macro internal PLL is not locked.
1	UL RxPort 4 to 7 PLL Unlock	See the description for bit 0.
2	UL RxPort 8 to 11 PLL Unlock	See the description for bit 0.
3	UL RxPort 12 to 15 PLL Unlock	See the description for bit 0.
4	UL RxPort 16 to 19 PLL Unlock	See the description for bit 0.
5	UL RxPort 20 to 23 PLL Unlock	See the description for bit 0.
6	UL RxPort 24 to 27 PLL Unlock	See the description for bit 0.
7	UL RxPort 28 to 31 PLL Unlock	See the description for bit 0.
8	UL TxPort 0 to 3 PLL Unlock	See the description for bit 0.
9	UL TxPorts 4 to 7 PLL Unlock	See the description for bit 0.
10	UL TxPorts 8 to 11 PLL Unlock	See the description for bit 0.

Bit(s)	Field Name	Description
11	UL TxPorts 12 to 15 PLL Unlock	See the description for bit 0.
12	UL TxPorts 16 to 19 PLL Unlock	See the description for bit 0.
13	UL TxPort 20 to 23 PLL Unlock	See the description for bit 0.
14	UL TxPort 24 to 27 PLL Unlock	See the description for bit 0.
15	UL TxPort 28 to 31 PLL Unlock	See the description for bit 0.
16	Ingress UL Spex Bus Rx Links 0 to 3 PLL Unlock	See the description for bit 0.
17	Ingress UL Spex Bus Rx Links 4 to 7 PLL Unlock	See the description for bit 0.
18	Egress UL Spex Bus Rx Links 0 to 3 PLL Unlock	See the description for bit 0.
19	Egress UL Spex Bus Rx Links 4 to 7 PLL Unlock	See the description for bit 0.
20	Ingress UL Spex Bus Tx Links 0 to 3 PLL Unlock	See the description for bit 0.
21	Ingress UL Spex Bus Tx Links 4 to 7 PLL Unlock	See the description for bit 0.
22	Egress UL Spex Bus Tx Links 0 to 3 PLL Unlock	See the description for bit 0.
23	Egress UL Spex Bus Tx Links 4 to 7 PLL Unlock	See the description for bit 0.
24:31	Reserved	Reserved.

5.2.4 UL TxPort Driver Enable Register

To enable a Unilink port driver, the following three conditions must be met:

- The output driver enable bit must be set in the *Reset Register* (page 69).
- The FullyInsertedIn# signal must be active (low level).
- The corresponding bit must be set in the *UL TxPort Driver Enable Register.*

If any of these conditions is not met, the Unilink output driver is tristated.

Address x‘0C’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

UL TxPort Driver Enable (for port *n* = bit *n)*

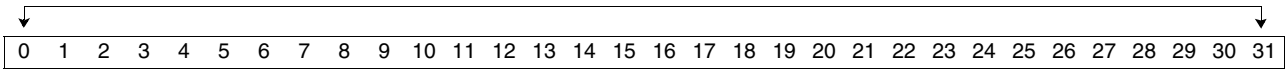

Bit(s)	Field Name	Description
0:31	UL TxPort Driver Enable (for port *n* = bit *n)*	When set to ‘1’, the Unilink port driver is enabled, as described above.

5.2.5 UL TxPort Attachment Enable Register

Address x‘0D’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

UL TxPort Attachment Enable (for port *n* = bit *n)*

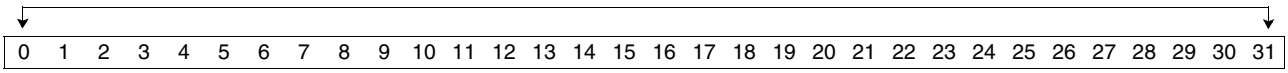

Bit(s)	Field Name	Description
0:31	UL TxPort Attachment Enable (for port *n* = bit *n)*	1 Allows normal data transmission on the port. 0 Forces data sent by the port to ‘0’.

5.2.6 UL TxPort Parameters Register

This register provides access to the Unilink port transmitter parameters, that is, the finite impulse response (FIR) coefficients and driver power level. The FIR coefficients are used to adjust the driver pre-emphasis FIR filter.

Write access to the *UL TxPort Parameters Register* requires one SHI command:

1. Write the *UL TxPort Parameters Register* with either the write bit *or* the global write bit set to '1', the port specified in the port number field, and the required values specified in the FIR coefficient and driver power level fields.

Read access to the *UL TxPort Parameters Register* requires two SHI commands:

1. Write the *UL TxPort Parameters Register* with the global write bit *and* the write bit set to '0', and the port specified in the port number field.
2. Read the *UL TxPort Parameters Register* to return the FIR coefficients and driver power level.

Address x'0E'

Access Type Read/Write

Reset Value '0000 0000 0000 0000 0000 0000 0000 0000'

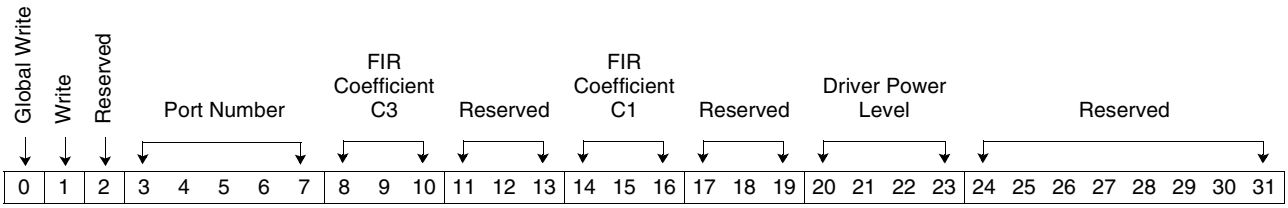

Bit(s)	Field Name	Description
0	Global Write	1 Specifies a write to the parameters of all the ports. This setting enables all the ports to be set with the same parameters in a single operation. 0 When the write bit is also set to '0', specifies a read of the parameters of the port specified in the port number field.
1	Write	1 Specifies a write to the parameters of the port specified in the port number field. 0 When the global write bit is also set to '0', specifies a read of the parameters of the port specified in the port number field.
2	Reserved	Reserved.
3:7	Port Number	Specifies the port.
8:10	FIR Coefficient C3 (2:0)	Specifies the value of FIR coefficient C3.
11:13	Reserved	Must be set to '000'.
14:16	FIR Coefficient C1 (2:0)	Specifies the value of FIR coefficient C1.
17:19	Reserved	Must be set to '000'.
20:23	Driver Power Level (3:0)	Adjusts the transmit core driver output power.
24:31	Reserved	Reserved.

5.2.7 UL TxPort BIST Request Register

Address x‘0F’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

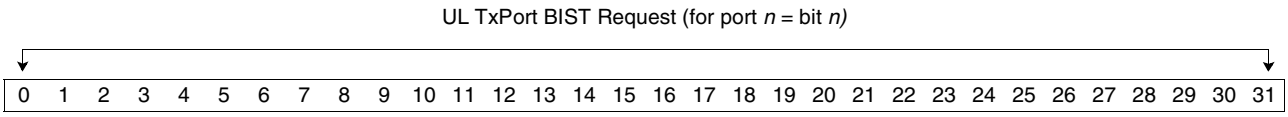

Bit(s)	Field Name	Description
0:31	UL TxPort BIST Request (for port *n* = bit *n)*	When set to ‘1’, executes the Unilink internal BIST on the port.

5.2.8 UL TxPort BIST Error Register

Address x‘10’

Access Type Read Only

Reset Value Undefined

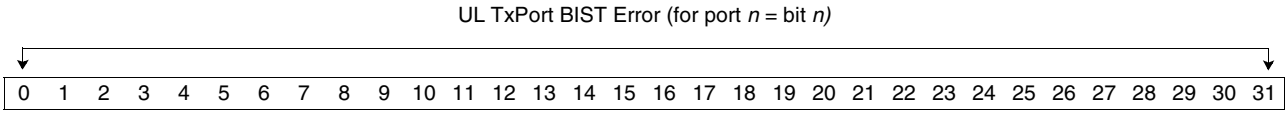

Bit(s)	Field Name	Description
0:31	UL TxPort BIST Error (for port *n* = bit *n)*	When set to ‘1’, reports a Unilink internal BIST failure on the port.

5.2.9 UL TxPort Reset BIST Error Register

Address x‘11’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

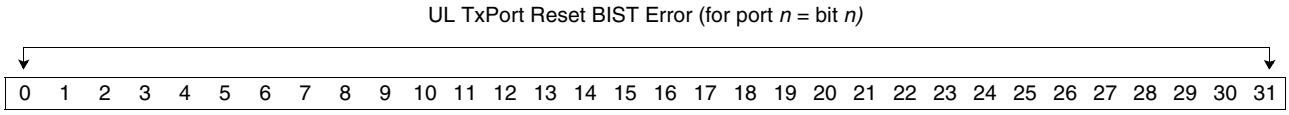

Bit(s)	Field Name	Description
0:31	UL TxPort Reset BIST Error (for port *n* = bit *n)*	When set to ‘1’, resets the Unilink internal BIST logic on the port.

5.2.10 UL RxPort Attachment Enable Register

Address x‘12’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

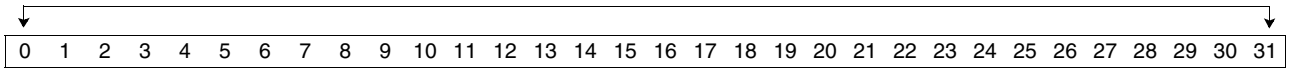

Bit(s)	Field Name		Description
0:31	UL RxPort Attachment Enable (for port *n* = bit *n)*	0	Keeps the Unilink receive synchronization and supervision logic in a reset state. In this state, the port cannot receive packets.
		1	Starts link synchronization.

5.2.11 UL RxPort Byte Alignment Done Register

Address x‘13’

Access Type Read Only

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

UL RxPort Byte Alignment Done (for port *n* = bit *n)*

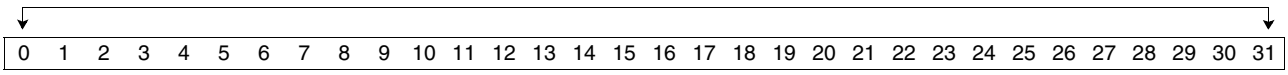

Bit(s)	Field Name	Description
0:31	UL RxPort Byte Alignment Done (for port *n* = bit *n)*	When set to ‘1’, the link synchronization logic has completed the byte-recovery phase of link synchronization on the port.

5.2.12 UL RxPort K28.5 Spacing OK Register

Address x‘14’

Access Type Read Only

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

UL RxPort K28.5 Spacing OK (for port *n* = bit *n)*

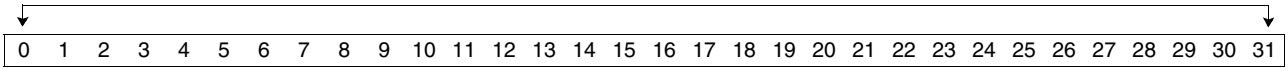

Bit(s)	Field Name	Description
0:31	UL RxPort K28.5 Spacing OK (for port *n* = bit *n)*	When set to ‘1’, the link synchronization logic has completed the K28.5-position verification phase of link synchronization on the port. The link is ready and LU deskew can begin.

5.2.13 UL RxPort LU Deskew Register

This register provides access to the port transmitter logical unit (LU) deskew settings.

Write access to the *UL RxPort LU Deskew Register* requires one SHI command:

1. Write the *UL RxPort LU Deskew Register* with the write bit set to ‘1’, the port specified in the port number field, and the port LU deskew setting specified in the LU deskew command field.

Read access to the *UL RxPort LU Deskew Register* requires two SHI commands:

1. Write the *UL RxPort LU Deskew Register* with the write bit set to ‘0’ and the port specified in the port number field.
2. Read the *UL RxPort LU Deskew Register* to return the LU deskew command field.

For more information about LU deskew, see *Section 6.3.1 Initializing Unilink Ports* on page 159.

Address x‘15’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

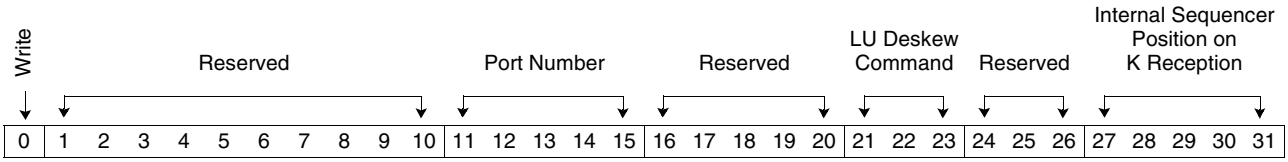

Bit(s)	Field Name	Description
0	Write	1 Specifies a write to the LU deskew setting of the port specified in the port number field. 0 Specifies a read of the LU deskew setting of the port specified in the port number field.
1:10	Reserved	Reserved.
11:15	Port Number	Specifies the port.
16:20	Reserved	Reserved.
21:23	LU Deskew Command	Specifies the LU deskew setting. When this field is set to *n,* the pipeline delay is $n \times 4$ ns.
24:26	Reserved	Reserved.
27:31	Internal Sequencer Position on K Reception (read only)	Reports the internal sequencer position (in 4-ns increments) when the K character of an idle packet is received. This position indicates the skew between LUs carried on different Unilinks of the same port.

5.2.14 UL RxPort Data Mode Register

Address x‘16’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

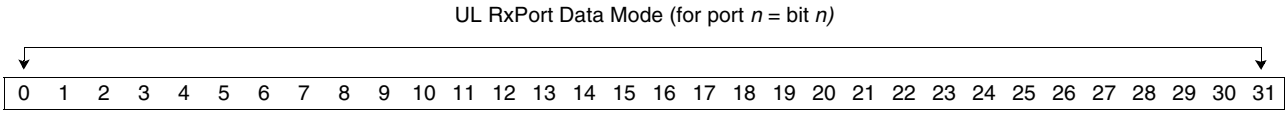

Bit(s)	Field Name	Description
0:31	UL RxPort Data Mode (for port *n* = bit *n)*	Set to ‘1’ when LU deskew is complete. This bit starts link supervision and packet reception on the port.

5.2.15 UL RxPort Data Valid Register

Address x‘17’

Access Type Read Only

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

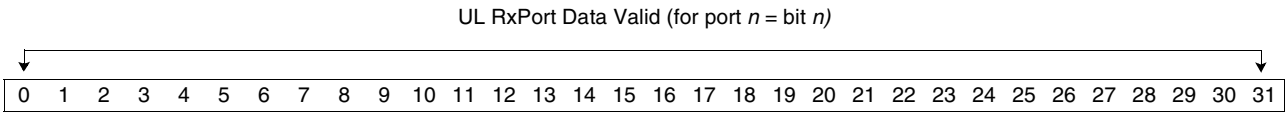

Bit(s)	Field Name	Description
0:31	UL RxPort Data Valid (for port *n* = bit *n)*	When set to ‘1’, the input controller can process packets. In normal operation, this register is a copy of the *UL RxPort Data Mode Register* (page 84). When either a port signal loss or synchronization loss error exists *and* when the hardware auto disable bit in the *UL Global Register* (page 72) is set, the hardware automatically clears this bit to keep the input controller from processing invalid packets.

5.2.16 UL RxPort Signal Lost Register

Address x‘18’

Access Type Read Only

Reset Value Undefined

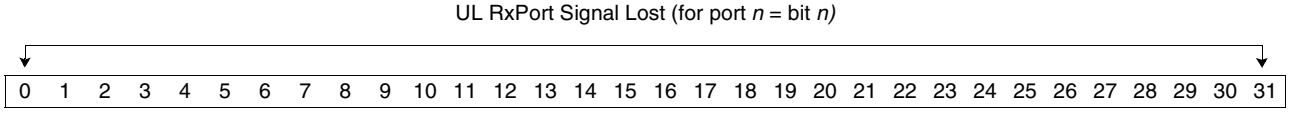

Bit(s)	Field Name	Description
0:31	UL RxPort Signal Lost (for port *n* = bit *n)*	When set to ‘1’, the Unilink port receiver does not detect a signal. If the “hardware auto disable on port signal lost enable” bit in the *UL Global Register* (page 72) is set, the corresponding bit in the *UL RxPort Data Valid Register* (page 84) is automatically cleared and the port does not receive packets.

5.2.17 UL RxPort Invalid K Character Register

Address x‘19’

Access Type Read/Clear

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

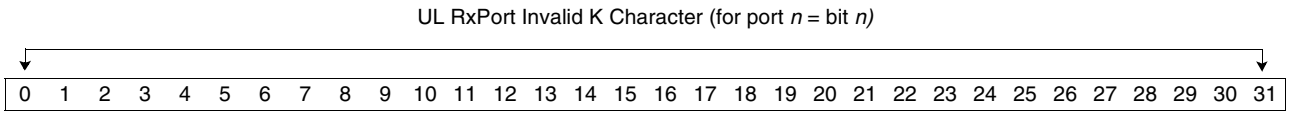

Bit(s)	Field Name	Description
0:31	UL RxPort Invalid K Character (for port *n* = bit *n)*	When set to ‘1’, the Unilink supervision logic has detected at least one invalid K character (either an unaligned K28.5 character or a K28.1/K28.5 character in the wrong place) during the port data mode.

5.2.18 UL RxPort Synchronization Lost Register

Address x‘1A’

Access Type Read/Clear

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

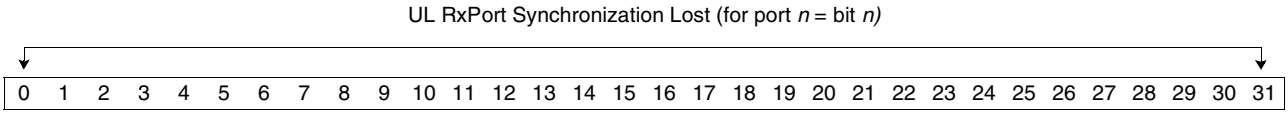

Bit(s)	Field Name	Description
0:31	UL RxPort Synchronization Lost (for port *n* = bit *n)*	When set to ‘1’, the Unilink supervision logic has detected a synchronization loss condition (either three consecutive unaligned K28.5 characters or three consecutive K28.1/K28.5 characters in the wrong place) during the port data mode. If the “hardware auto disable on port sync lost enable” bit in the *UL Global Register* (page 72) is set, the corresponding bit in the *UL RxPort Data Valid Register* (page 84) is automatically cleared and the port does not receive packets.

5.2.19 UL RxPort Code Violation Register

Address x‘1B’

Access Type Read/Clear

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

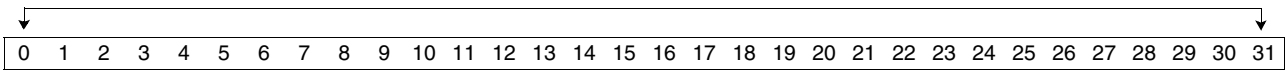

Bit(s)	Field Name	Description
0:31	UL RxPort Code Violation (for port *n* = bit *n)*	When set to ‘1’, the Unilink supervision logic has detected an 8b/10b code violation during the port data mode.

5.2.20 UL RxPort BIST Request Register

Address x‘1C’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

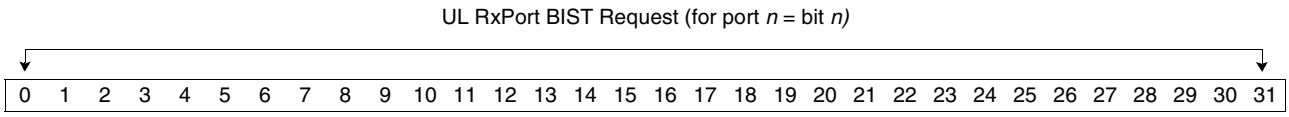

Bit(s)	Field Name	Description
0:31	UL RxPort BIST Request (for port *n* = bit *n)*	When set to ‘1’, executes the Unilink internal BIST on the port.

5.2.21 UL RxPort BIST Error Register

Address x‘1D’

Access Type Read Only

Reset Value Undefined

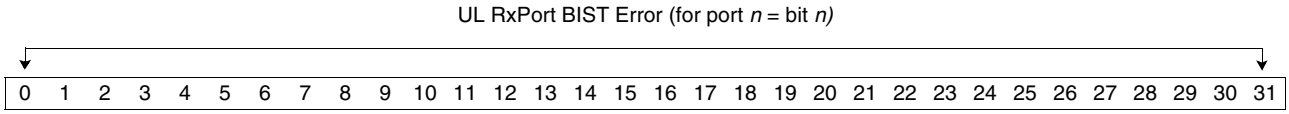

Bit(s)	Field Name	Description
0:31	UL RxPort BIST Error (for port *n* = bit *n)*	When set to ‘1’, reports a Unilink internal BIST failure on the port.

5.2.22 UL RxPort Reset BIST Error Register

Address x‘1E’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

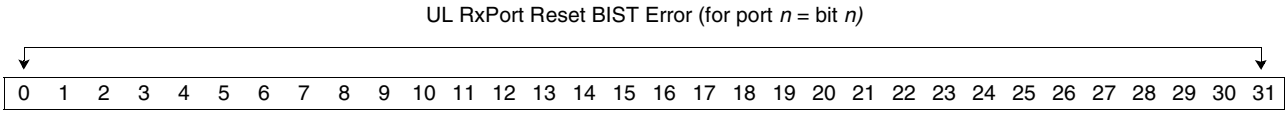

Bit(s)	Field Name	Description
0:31	UL RxPort Reset BIST Error (for port *n* = bit *n)*	When set to ‘1’, resets the Unilink internal BIST logic on the port.

5.2.23 UL RxPort BIST Wrap Register

Address x‘1F’

Access Type Read/Write

Reset Value Undefined

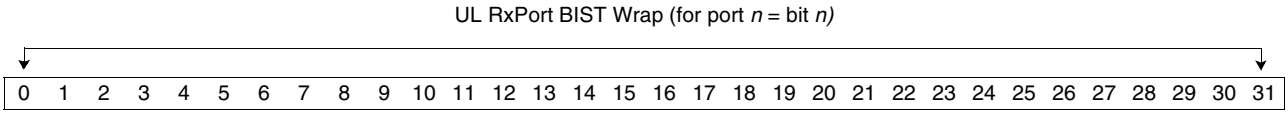

Bit(s)	Field Name	Description
0:31	UL RxPort BIST Wrap (for port *n* = bit *n)*	When set to ‘1’, enables the Unilink port receive macro to run the BIST with patterns generated by the Unilink transmitter (rather than with Unilink receiver internal patterns). The PowerPRS Q-64G compares the pattern received against the pattern generated to test a complete transmission from the Unilink transmitter to the Unilink receiver.

5.2.24 UL TxSpex Bus Driver Enable Register

To enable a Unilink speed-expansion bus driver, the following two conditions must be met:

1. The output driver enable bit must be set in the *Reset Register* (page 69).
2. The corresponding bit must be set in the *UL TxSpex Bus Driver Enable Register*.

If either of these conditions is not met, the Unilink output driver is tristated.

Address x‘20’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

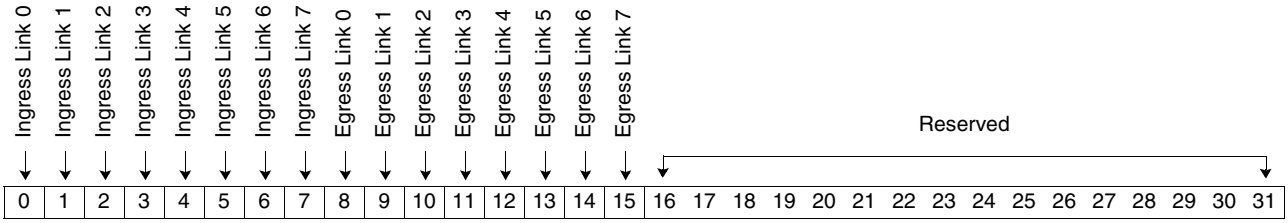

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, the Unilink speed-expansion bus driver is enabled, as described above.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.25 UL TxSpex Bus Attachment Enable Register

Address x‘21’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

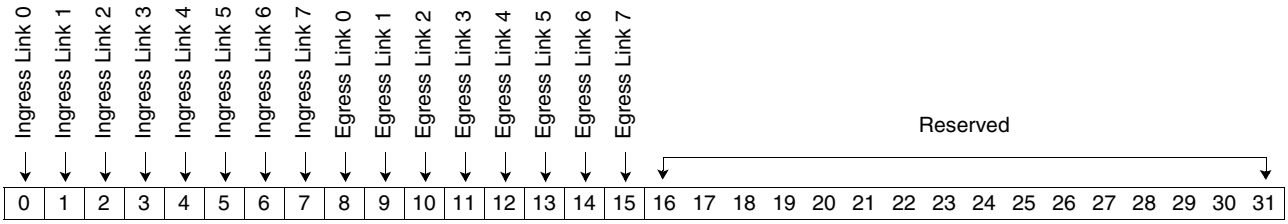

Bit(s)	Field Name	Description
0	Ingress Link 0	1 Allows normal transmission on the link. 0 Forces data sent over the link to ‘0’.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.26 UL TxSpex Bus Parameters Register

This register provides access to the Unilink speed-expansion bus transmitter parameters, that is, the FIR coefficients and driver power level. The FIR coefficients are used to adjust the driver pre-emphasis FIR filter.

Write access to the *UL TxSpex Bus Parameters Register* requires one SHI command:

1. Write the *UL TxSpex Bus Parameters Register* with either the write bit *or* the global write bit set to '1', the speed-expansion bus specified by the ingress/egress bus bit, the link specified in the link number field, and the required values specified in the FIR coefficient and driver power level fields.

Read access to the *UL TxSpex Bus Parameters Register* requires two SHI commands:

1. Write the *UL TxSpex Bus Parameters Register* with the global write bit *and* the write bit set to '0', the speed-expansion bus specified by the ingress/egress bus bit, and the link specified in the link number field.
2. Read the *UL TxSpex Bus Parameters Register* to return the FIR coefficients and power driver level.

Address x'22'

Access Type Read/Write

Reset Value '0000 0000 0000 0000 0000 0000 0000 0000'

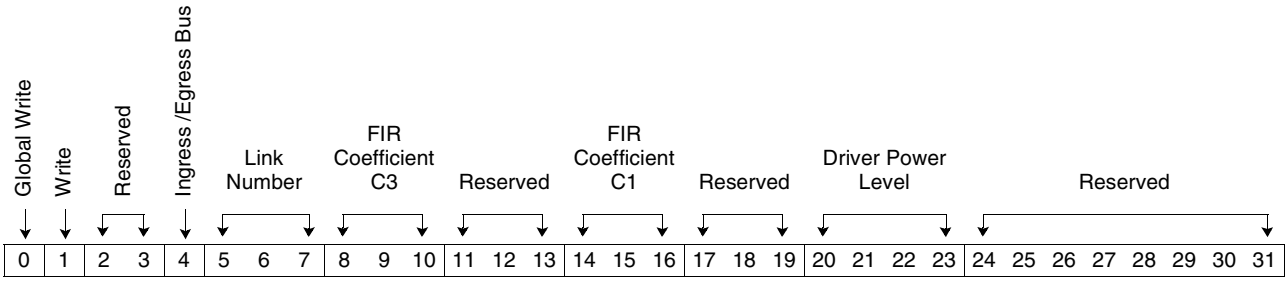

Bit(s)	Field Name	Description
0	Global Write	1 Specifies a write to the parameters of all the links on the speed-expansion bus specified by the ingress/egress bus bit. This setting enables all the links to be set with the same parameters using a single operation. 0 When the write bit is also set to '0', specifies a read of the parameters of the link specified in the link number field.
1	Write	1 Specifies a write to the parameters of the link specified in the link number field. 0 When the global write bit is also set to '0', specifies a read of the parameters of the link specified in the link number field.
2:3	Reserved	Reserved.
4	Ingress/Egress Bus	Specifies the speed-expansion bus: 1 Egress speed-expansion bus 0 Ingress speed-expansion bus
5:7	Link Number	Specifies the link.
8:10	FIR Coefficient C3 (2:0)	Specifies the value of FIR coefficient C3.
11:13	Reserved	Must be set to '000'.
14:16	FIR Coefficient C1 (2:0)	Specifies the value of FIR coefficient C1.

Bit(s)	Field Name	Description
17:19	Reserved	Must be set to ‘000’.
20:23	Driver Power Level (3:0)	Adjusts the transmit core driver output power.
24:31	Reserved	Reserved.

5.2.27 UL TxSpex Bus BIST Request Register

Address x‘23’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

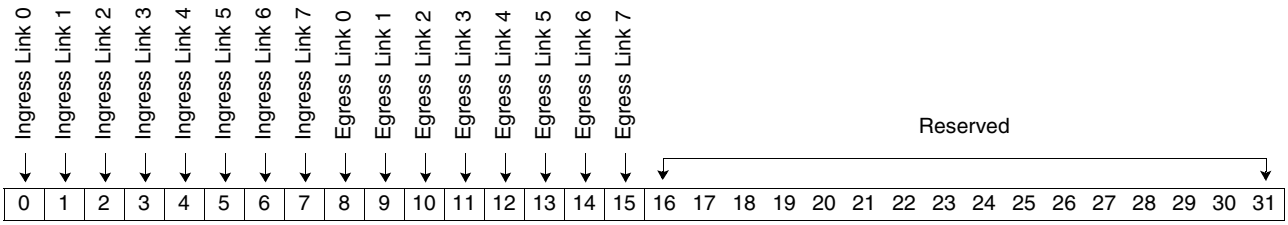

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, executes the Unilink internal BIST on the link.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.28 UL TxSpex Bus BIST Error Register

Address x‘24’

Access Type Read Only

Reset Value ‘uuuu uuuu uuuu uuuu 0000 0000 0000 0000’, where ‘u’ = undefined

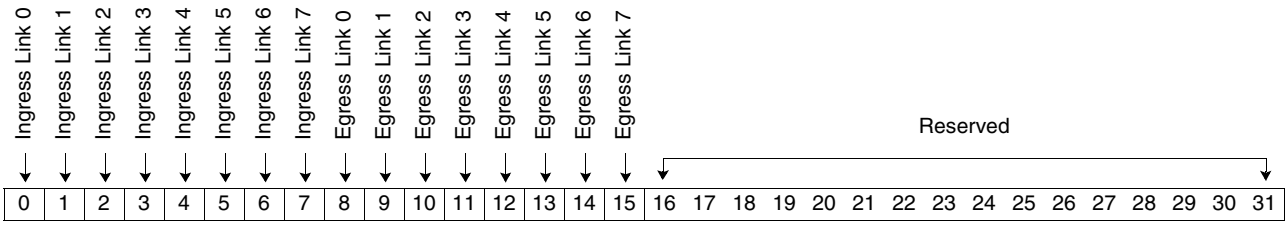

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, reports a Unilink internal BIST failure on the port.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.29 UL TxSpex Bus Reset BIST Error Register

Address x‘25’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

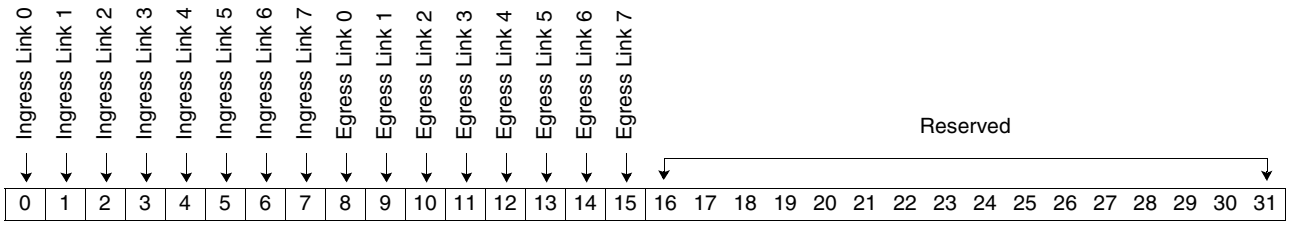

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, resets the Unilink internal BIST logic on the link.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.30 UL RxSpex Bus Attachment Enable Register

Address x‘26’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

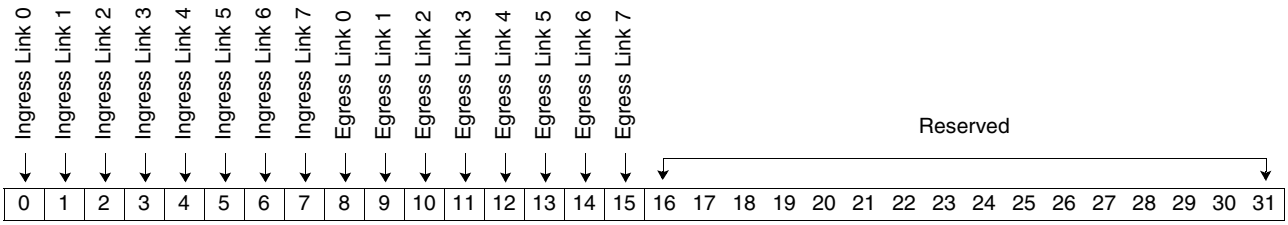

Bit(s)	Field Name	Description
0	Ingress Link 0	0 Keeps the Unilink receive synchronization and supervision logic in a reset state. 1 Starts link synchronization.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.31 UL RxSpex Bus Byte Alignment Done Register

Address x‘27’

Access Type Read Only

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

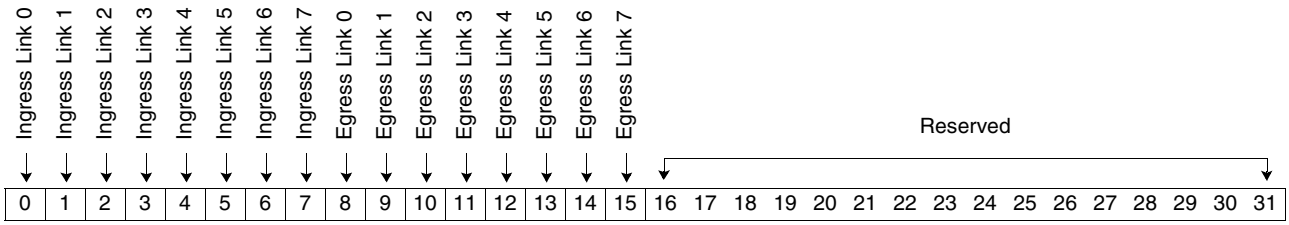

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, the link synchronization logic has completed the byte-recovery phase of link synchronization on the link.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.32 UL RxSpex Bus K28.5 Spacing OK Register

Address x‘28’

Access Type Read Only

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

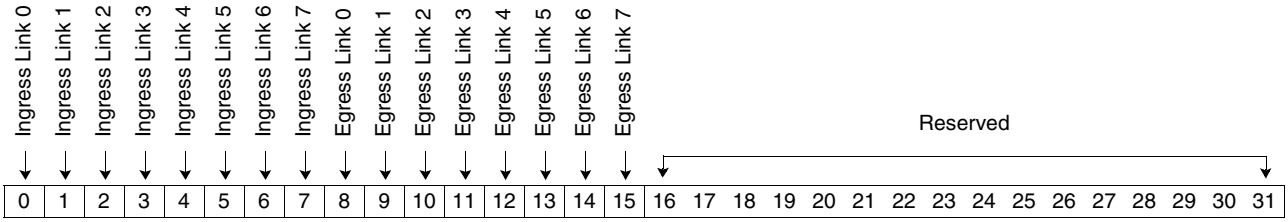

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, the link synchronization logic has completed the K28.5-position verification phase of link synchronization on the link. The link is ready and link latency programming can begin.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.33 UL RxSpex Bus Latency Programming Register

This register provides access to the Unilink speed-expansion bus latency settings. For more information about speed-expansion bus latency, see *Section 6.2.2 Setting the Speed-Expansion Bus Data Latency* on page 156.

Write access to the *UL RxSpex Bus Latency Programming Register* requires one SHI command:

1. Write the *UL RxSpex Bus Latency Programming Register* with the write bit set to '1', the speed-expansion bus specified by the ingress/egress bit, the link specified in the link number field, and the data latency setting specified in the data latency field.

Read access to the *UL RxSpex Bus Latency Programming Register* requires two SHI commands:

1. Write the *UL RxSpex Bus Latency Programming Register* with the write bit set to '0', the speed-expansion bus specified by the ingress/egress bit, and the link specified in the link number field.
2. Read the *UL RxSpex Bus Latency Programming Register* to return the data latency value and speed-expansion bus delay.

Address x'29'

Access Type Read/Write

Reset Value '0000 0000 0000 0000 0000 0000 0000 0000'

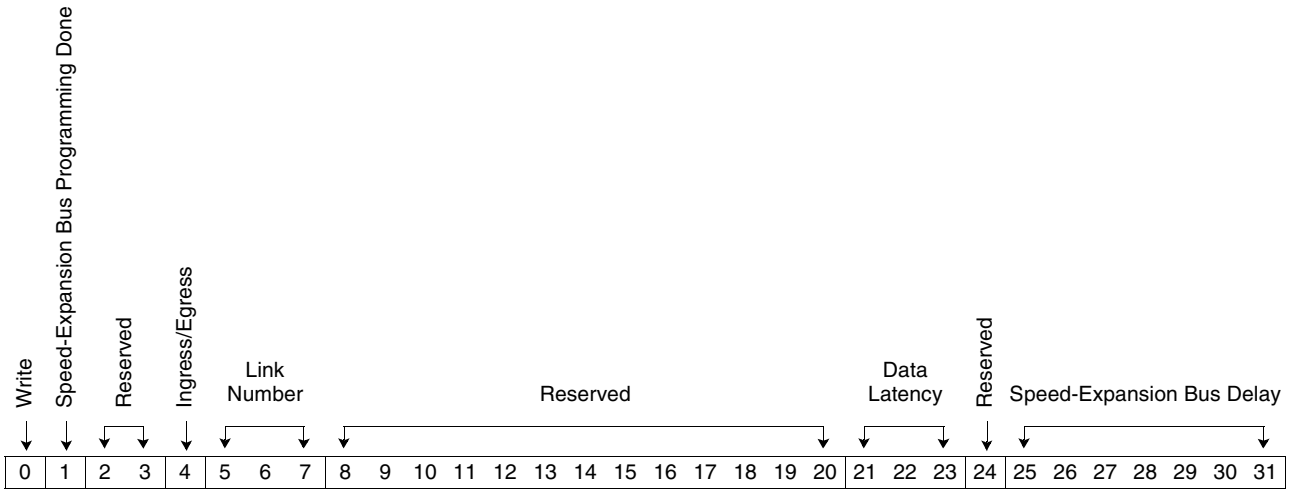

Bit(s)	Field Name	Description
0	Write	1 Specifies a write to the speed-expansion bus latency settings. 0 Specifies a read from the speed-expansion bus latency settings.
1	Speed-Expansion Bus Programming Done	This bit must be set after speed-expansion bus latency programming is complete. It must be set first in the master device and then in the slave devices. 1 Speed-expansion bus programming is complete. Data packet addresses can be propagated through the speed-expansion bus links. There are no more K characters on the link, and address protection is provided by the address parity bit. 0 Speed-expansion bus programming is in progress. Idle packets are propagated for link synchronization.
2:3	Reserved	Reserved.

Bit(s)	Field Name	Description
4	Ingress/Egress	Specifies the speed expansion bus: 0 Ingress speed-expansion bus 1 Egress speed-expansion bus
5:7	Link Number	Specifies the link number (0 to 7).
8:20	Reserved	Reserved.
21:23	Data Latency	Specifies the data latency. When this field is set to *n,* the packet delay is *n* packets.
24	Reserved	Reserved.
25:31	Speed-Expansion Bus Delay (read only)	Reports the speed-expansion bus delay. When this field is set to *n,* the speed-expansion bus delay is $n \times 8$ ns.

5.2.34 UL RxSpex Bus Data Mode Register

Address x‘2A’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

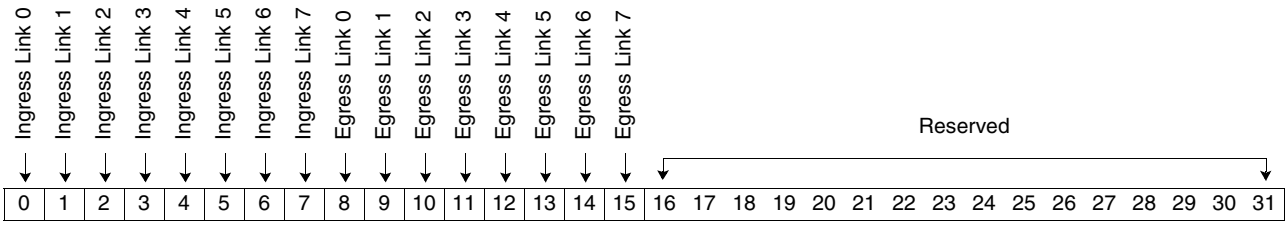

Bit(s)	Field Name	Description
0	Ingress Link 0	Set to ‘1’ when link synchronization is complete. This bit starts master-to-slave communication via the speed-expansion bus.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.35 UL RxSpex Bus Signal Lost Register

Address x‘2B’

Access Type Read Only

Reset Value Undefined

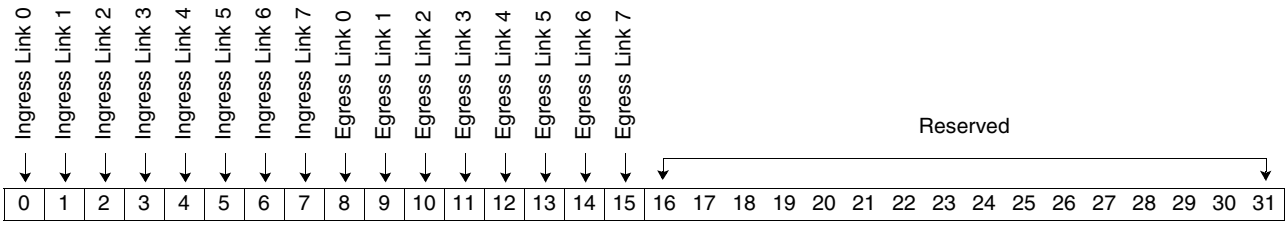

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, the Unilink speed-expansion bus receiver does not detect a signal. If the “hardware auto disable on spex bus signal lost enable” bit in the *UL Global Register* (page 72) is set, valid data is automatically cleared and the link no longer receives addresses.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.36 UL RxSpex Bus Invalid K Character Register

Address x‘2C’

Access Type Read/Clear

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

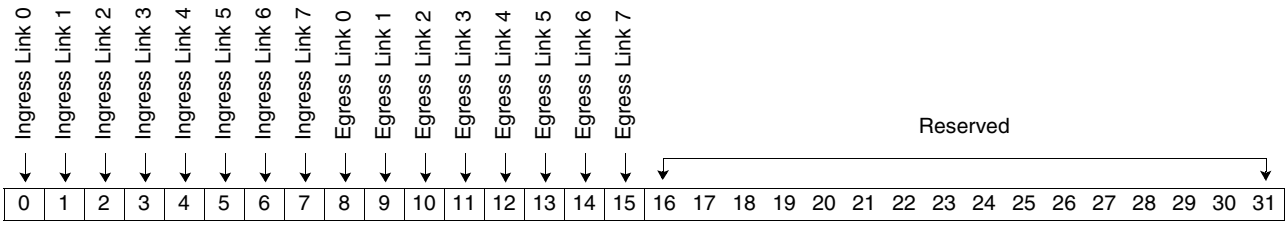

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, the Unilink supervision logic has detected at least one invalid K character (either an unaligned 28.5 character or a K28.1/K28.5 character in the wrong place) during the link data mode.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.37 UL RxSpex Bus Synchronization Lost Register

Address x‘2D’

Access Type Read/Clear

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

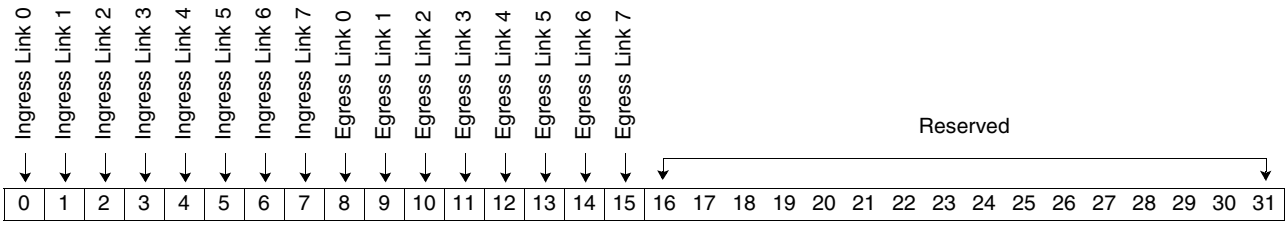

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, the Unilink supervision logic has detected a loss of synchronization condition (either three consecutive unaligned K28.5 characters or three consecutive K28.1/K28.5 characters in the wrong place) during the link data mode. If the “hardware auto disable on speed-expansion bus sync lost enable” bit in the *UL Global Register* (page 72) is set, valid data is automatically reset and the link does not receive addresses.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.38 UL RxSpex Bus Code Violation Register

Address x‘2E’

Access Type Read/Clear

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

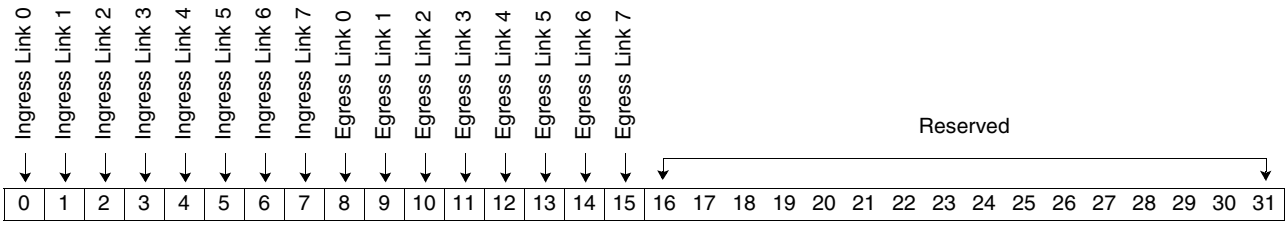

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, the Unilink supervision logic has detected an 8b/10b code violation during the link data mode.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.39 UL RxSpex BIST Request Register

Address x‘2F’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

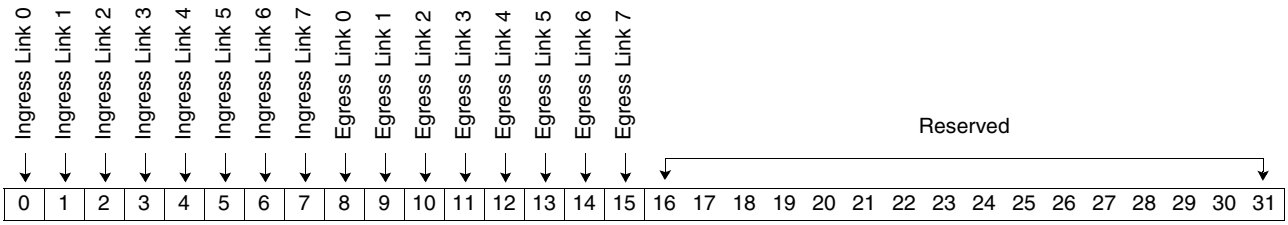

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, executes the Unilink internal BIST on the link.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.40 UL RxSpex Bus BIST Error Register

Address x‘30’

Access Type Read Only

Reset Value ’uuuu uuuu uuuu uuuu 0000 0000 0000 0000’, where ‘u’ = undefined

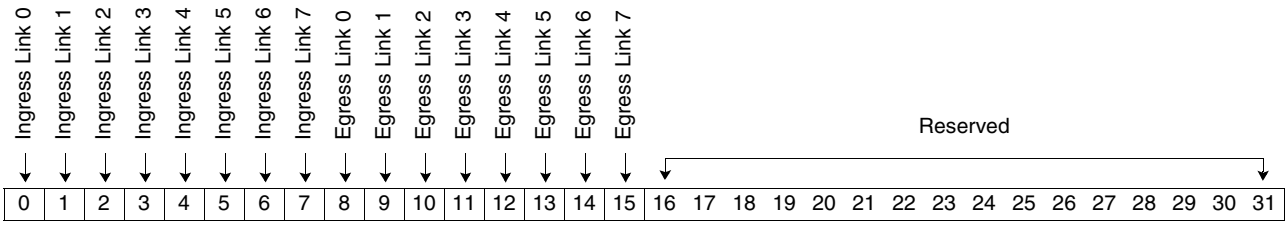

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, reports a Unilink internal BIST failure on the link.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.41 UL RxSpex Bus Reset BIST Error Register

Address x‘31’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

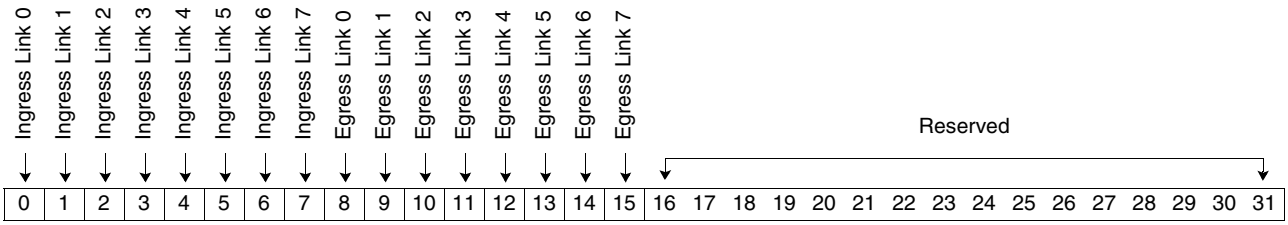

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, resets the Unilink internal BIST logic on the link.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.2.42 UL RxSpex BIST Wrap Register

Address x‘32’

Access Type Read/Write

Reset Value Undefined

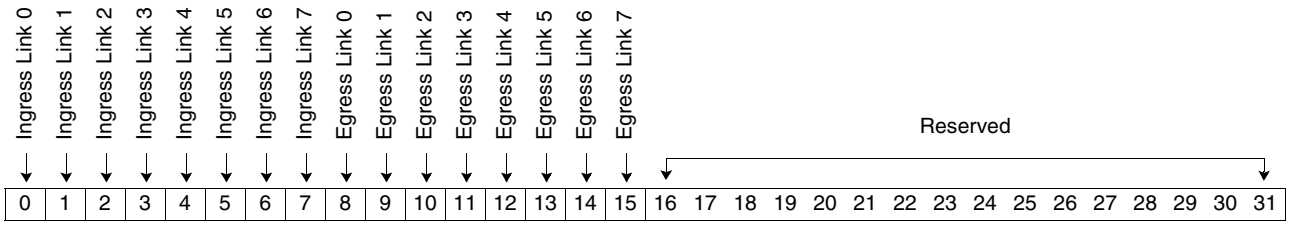

Bit(s)	Field Name	Description
0	Ingress Link 0	When set to ‘1’, enables the Unilink port receive macro to run the BIST with patterns generated by the Unilink transmitter on the master device and checked by the Unilink receiver on the slave device.
1	Ingress Link 1	See the description for bit 0.
2	Ingress Link 2	See the description for bit 0.
3	Ingress Link 3	See the description for bit 0.
4	Ingress Link 4	See the description for bit 0.
5	Ingress Link 5	See the description for bit 0.
6	Ingress Link 6	See the description for bit 0.
7	Ingress Link 7	See the description for bit 0.
8	Egress Link 0	See the description for bit 0.
9	Egress Link 1	See the description for bit 0.
10	Egress Link 2	See the description for bit 0.
11	Egress Link 3	See the description for bit 0.
12	Egress Link 4	See the description for bit 0.
13	Egress Link 5	See the description for bit 0.
14	Egress Link 6	See the description for bit 0.
15	Egress Link 7	See the description for bit 0.
16:31	Reserved	Reserved.

5.3 Functional Registers

5.3.1 Configuration 0 Register

Address x‘33’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0001’

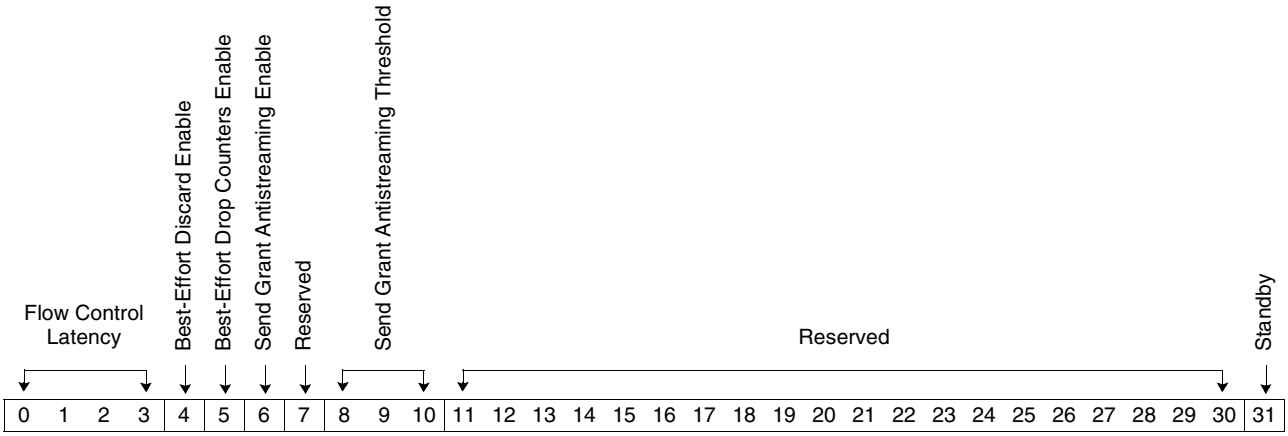

Bit(s)	Field Name	Description
0:3	Flow Control Latency	When enabled, the input controllers check flow control. For a unicast packet, the input controller checks whether an incoming packet is destined to an output for which neither an output queue grant nor a memory grant has been issued in the past *n* packet cycles. For a multicast packet, the input controller checks only the memory grant information. If the packet is destined for an output lacking grants, the packet is discarded. The error is reported via the flow control violation bit in the *Status Register* (page 120) and, unless masked in the *Interrupt Mask Register* (page 122), the error generates a flow control violation interrupt. The violating ports are identified by the corresponding bits in the *Flow Control Violation Register* (page 132). 0000 Function is disabled 0001 Function is enabled with *n* = 8 0010 Function is enabled with *n* = 10 0011 Function is enabled with *n* = 12 0100 Function is enabled with *n* = 14 0101 Function is enabled with *n* = 16 0110 Function is enabled with *n* = 18 0111 Function is enabled with *n* = 20 1000 Function is enabled with *n* = 22 1001 Function is enabled with *n* = 24 1010 Function is enabled with *n* = 26 1011 Function is enabled with *n* = 28 1100 Function is enabled with *n* = 30 1101 Function is enabled with *n* = 32 1110 Function is enabled with *n* = 34 1111 Function is enabled with *n* = 36
4	Best-Effort Discard Enable	1 Enables the input controllers to discard best-effort traffic, depending on the best-effort discard thresholds and the best-effort counter values. See *Section 3.4.6 Best-Effort Discard* on page 49. 0 Clears the best-effort discard counters.
5	Best-Effort Drop Counters Enable	1 Enables the best-effort drop counters to count the discarded best-effort packets. 0 Clears the best-effort drop counters.
6	Send Grant Antistreaming Enable	When set to ‘1’, enables the send grant antistreaming function under the parameters defined in the send grant antistreaming threshold field.
7	Reserved	Reserved.

Bit(s)	Field Name	Description
8:10	Send Grant Antistreaming Threshold	Provides protection when the attached device removes the send grant for an extended period. If the send grant is not issued for any priority for *n* contiguous packet cycles, this bit is internally forced to active for all priorities until the attached device issues another send grant: 000 *n* = 16 100 *n* = 256 001 *n* = 32 101 *n* = 512 010 *n* = 64 110 *n* = 1024 011 *n* = 128 111 *n* = 2048 The *Send Grant Violation Register* (page 131) identifies the ports for which send grant antistreaming is in progress.
11:30	Reserved	Reserved.
31	Standby	Freezes the device while the configuration is reset. This bit is asserted after the power-on reset is completed and is released as described in *Section 6.1 Reset Sequence* on page 155.

5.3.2 Configuration 1 Register

Address x‘34’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

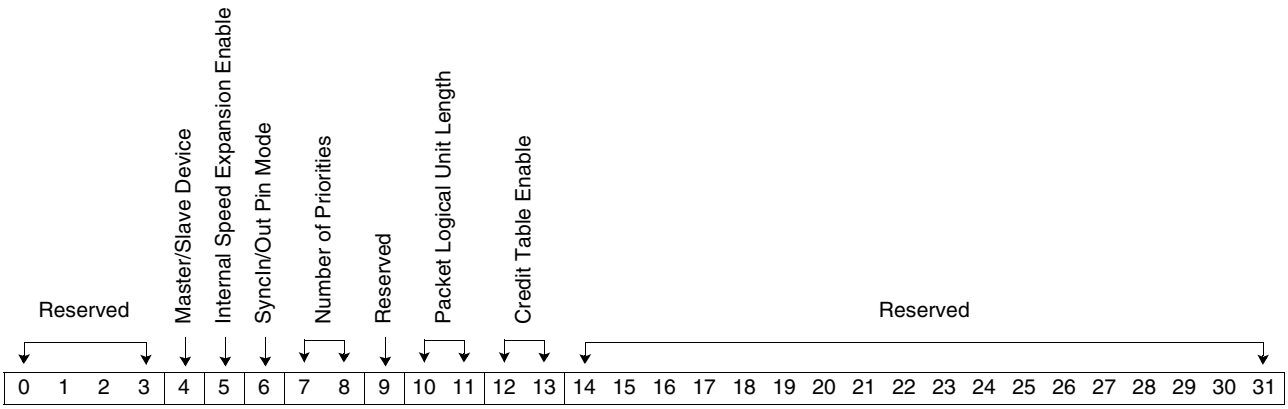

Bit(s)	Field Name	Description
0:3	Reserved	Reserved.
4	Master/Slave Device	When multiple PowerPRS Q-64Gs are configured for external speed expansion, designates whether a device is the master or a slave: 1 Master device 0 Slave device
5	Internal Speed Expansion Enable	When set to ‘1’, enables internal speed expansion. This bit is used with four PowerPRS Q-64Gs in the 256-Gbps configuration.
6	SyncIn/Out Pin Mode	Specifies SyncIn and SyncOut pin operation: 1 The SyncOut pin is enabled and the SyncIn pin is not used. SyncOut is an output signal generated by the internal sequencer, which is running independently. This is the master device setting. 0 The SyncIn pin is enabled and the SyncOut pin is tristated. SyncIn is an input signal used to synchronize the internal sequencer. This is the slave device setting.
7:8	Number of Priorities	Defines the number of priorities for which grants are issued: 00 Priority 0 only is enabled 01 Priorities 0 and 1 are enabled 10 Priorities 0, 1, and 2 are enabled 11 Priorities 0, 1, 2, and 3 are enabled This bit controls the cycling of flow control flywheels used to transmit grants to the attached devices (see *Section 3.3 Packet Format According to Packet Type* on page 28).
9	Reserved	Reserved.
10:11	Packet Logical Unit Length	Specifies the LU length: 00 8 bytes 01 9 bytes 10 10 bytes 11 Reserved

Bit(s)	Field Name	Description
12:13	Credit Table Enable	Defines the range of priorities that can be preempted by the priority provided by the credit table: 00 Credit table is disabled (and the application does not initialize the credit table) 01 Priorities 0, 1, and 2 can be preempted by a lower-priority credit 10 Priorities 1 and 2 can be preempted by a lower-priority credit 11 Only priority 2 can be preempted by a lower-priority credit The credit table is programmed using the *Credit Table Access Register* (page 116).
14:31	Reserved	Reserved.

5.3.3 Threshold Access Register

This register provides indirect access to the internal threshold registers.

Write access to an internal threshold register requires one SHI command:

1. Write the *Threshold Access Register* with the write bit set to ‘1’, the threshold select field specifying the internal threshold register to be written, and the threshold value field specifying the value.

Read access to an internal threshold register requires two SHI commands:

1. Write the *Threshold Access Register* with the write bit set to '0' and the threshold select field specifying the internal threshold register to be read.
2. Read the *Threshold Access Register* to return the corresponding threshold value.

There are three types of internal thresholds:

- *Shared memory threshold.* The shared memory threshold is comprised of one 11-bit field and has a value ranging from 0 to 1024. Shared memory threshold use is described in *Section 3.4.2 Memory Grants* on page 48.
- *Multicast thresholds.* There are two multicast thresholds (one “high” and one “low”) per priority. Each multicast threshold is comprised of four 15-bit fields and has a value ranging from 0 to 15360 (that is, from 0 to 1024 × 15). Multicast threshold use is described in *Section 3.4.3 Multicast Grants* on page 48.
- *Output queue thresholds.* There are four output queue thresholds, one per priority. Each output queue threshold is comprised of four 11-bit fields and has a value ranging from 0 to 1024. Output queue threshold use is described in *Section 3.4.1 Output Queue Grants* on page 47.

Address x‘35’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

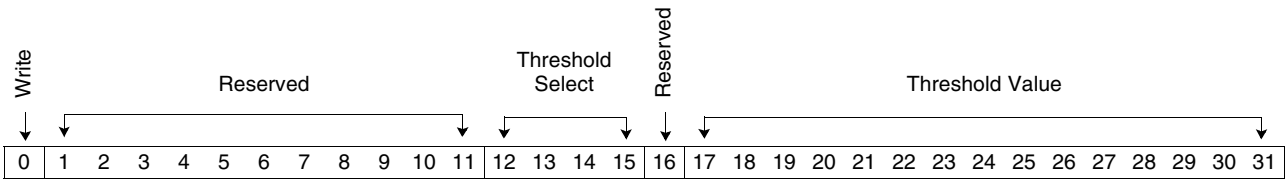

Bit(s)	Field Name	Description
0	Write	1 Specifies a write to the internal threshold register. 0 Specifies a read from the internal threshold register.
1:11	Reserved	Reserved.

Bit(s)	Field Name	Description
12:15	Threshold Select	Specifies the internal threshold register to be read or written. Encoding to specify the shared memory threshold: 0000 Shared memory threshold Encoding to specify a multicast threshold: 0001 Multicast high threshold for priority 0 0010 Multicast high threshold for priority 1 0011 Multicast high threshold for priority 2 0100 Multicast high threshold for priority 3 0101 Multicast low threshold for priority 0 0110 Multicast low threshold for priority 1 0111 Multicast low threshold for priority 2 1000 Multicast low threshold for priority 3 Encoding to specify an output queue threshold: 1001 Output queue threshold for priority 0 1010 Output queue threshold for priority 1 1011 Output queue threshold for priority 2 1100 Output queue threshold for priority 3 Other values are reserved.
16	Reserved	Reserved.
17:31	Threshold Value	Contains the value written to, or the value read from, the specified internal threshold register. Bits 17:31 are used for multicast thresholds, and bits 21:31 are used for shared memory and output queue thresholds (bits 17:20 are reserved).

5.3.4 Credit Table Access Register

This register provides indirect access to the credit table (see *Section 3.5.3 Credit Table* on page 53). There are 256 credits per port. Each port has 32 16-bit addresses, and each address contains eight credits. The one-credit fields designate the priority for which a credit is generated. Credits are generated during the packet cycle that corresponds to the credit number.

The credit table is not accessible while the PowerPRS Q-64G is in standby (that is, when the standby bit of the *Configuration 0 Register* [page 110] is set). Credit table access is allowed only while the corresponding port is active (that is, when the port is receiving or transmitting data).

Write access to the credit table requires one SHI command:

1. Write the *Credit Table Access Register* with the write bit set to '1', the port number field specifying the output port number, the credit table address field specifying the credit table address to be written, and the eight one-credit fields set with eight credits.

Read access to the credit table requires two SHI commands:

1. Write the *Credit Table Access Register* with the write bit set to '0', the port number field specifying the port number, and the credit table address field specifying the credit table address to be read.
2. Read the *Credit Table Access Register* to return the eight corresponding one-credit field values.

Address x'36'

Access Type Read/Write

Reset Value '0000 0000 0000 0000 uuuu uuuu uuuu uuuu', where 'u' = undefined

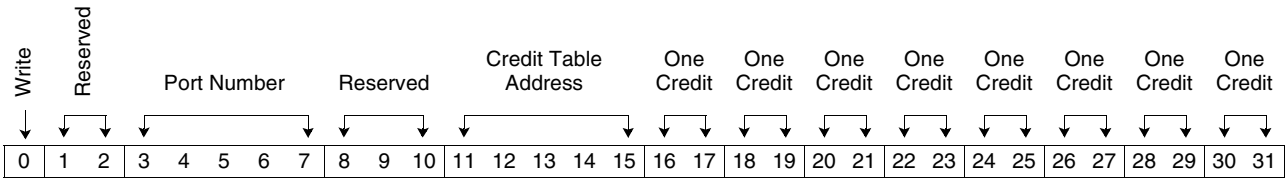

Bit(s)	Field Name	Description
0	Write	1 Specifies a write to the credit table. 0 Specifies a read from the credit table.
1:2	Reserved	Reserved.
3:7	Port Number	Specifies the port number.
8:10	Reserved	Reserved.
11:15	Credit Table Address	Specifies the credit table address.
16:17	One Credit	Specifies the priority.
18:19	One Credit	Specifies the priority.
20:21	One Credit	Specifies the priority.
22:23	One Credit	Specifies the priority.
24:25	One Credit	Specifies the priority.
26:27	One Credit	Specifies the priority.

Bit(s)	Field Name	Description
28:29	One Credit	Specifies the priority.
30:31	One Credit	Specifies the priority.

5.3.5 Best-Effort Resources Access Register

This register provides indirect access to the resources associated with the best-effort discard function (see *Section 3.4.6 Best-Effort Discard* on page 49).

Write access to the *Best-Effort Resources Access Register* requires one SHI command:

1. Write the register with the write bit set to '1', the output port number field specified, the resource select field specified, and the resource value field specified.

Read access to the *Best-Effort Resources Access Register* requires two SHI commands:

1. Write the register with the write bit cleared to '0', the output port number field specified, and the resource select field specified.
2. Read the register to return the corresponding resource value field value.

Note: To write to this register, the best-effort discard function must be enabled (using the best-effort discard enable bit of the *Configuration 0 Register* [page 110]).

Address x'37'

Access Type Read/Write

Reset Value '0000 0000 0000 0000 0000 0000 0000 0000'

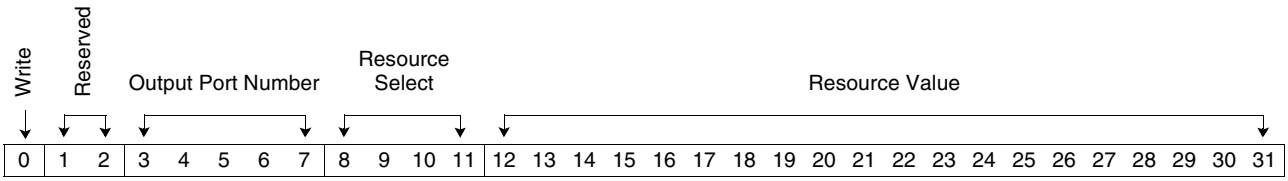

Bit(s)	Field Name	Description
0	Write	1 Specifies a write to the *Best-Effort Resources Access Register*. 0 Specifies a read from the *Best-Effort Resources Access Register*.
1:2	Reserved	Reserved.
3:7	Output Port Number	Specifies the output port number.
8:11	Resource Select	Specifies the resource to be accessed: 0000 Best-effort discard counter 1 0001 Best-effort discard counter 2 0010 Best-effort discard counter 3 0011 Best-effort discard counter 4 0100 Best-effort discard counter 5 0101 Best-effort enable discard threshold 0110 Best-effort priority discard threshold 0111 Best-effort halt discard threshold 1000 Best-effort drop counter for priority 0 1001 Best-effort drop counter for priority 1 1010 Best-effort drop counter for priority 2 1011 Best-effort drop counter for priority 3 1100 Protocol engine virtual packet clock (in this case, only bits 22:31 are valid) Others Reserved

Bit(s)	Field Name	Description
12:31	Resource Value	When the resource select field specifies a counter, reports the counter value. When the resource select field specifies the protocol engine virtual packet clock, the encoded value of this field is: x'00000' No protocol engine packet clock (protocol engine packet clock frequency is equal to 0). x‘00001' Single-shot protocol engine packet clock. This generates a pulse on the internal protocol engine packet clock (used for testing). x‘00002’ Reserved. x‘00003’ Reserved. Any other value x*‘nnnnn’* The protocol engine packet clock has a period of *“nnnnn”* × 8 ns (x*‘nnnnn’* must be greater than x‘00011’ to define a protocol engine packet clock).

5.3.6 Status Register

This register reports error-related events. Each bit in the *Status Register* can generate an interrupt to the local processor; however, an interrupt can be masked by setting the corresponding bit in the *Interrupt Mask Register* (page 122). The occurrence of an event for which the interrupt mask bit is set to '1' sets the corresponding bit in the *Status Register* but does not activate the InterruptOut# signal. The InterruptOut# signal, which interrupts the local processor, is activated only when the global interrupt mask bit *is not* set and the output driver enable bit *is* set in the *Reset Register* (page 69).

Address x‘38’

Access Type Read/Clear. Write is allowed for testing.

Reset Value ‘0000 0000 0000 0000 0010 0000 0000 0000’

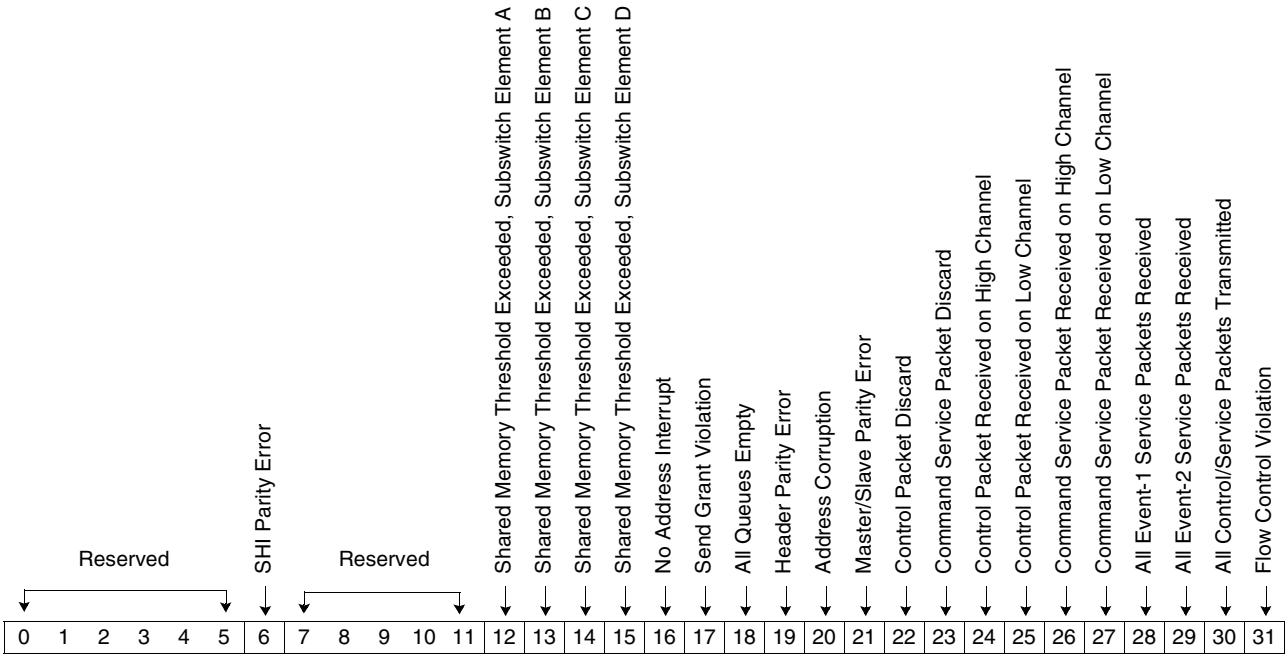

Bit(s)	Field Name	Description
0:5	Reserved	Reserved.
6	SHI Parity Error	Set to ‘1’ when the serial host interface (SHI) detects a parity error in the SHI instruction.
7:11	Reserved	Reserved.
12	Shared Memory Threshold Exceeded, Subswitch Element A	Indicates that the shared memory occupancy is equal to or greater than the shared memory threshold. This is an event, not a status; therefore, it occurs when the shared memory occupancy reaches the threshold and does not change when the occupancy falls below the threshold.
13	Shared Memory Threshold Exceeded, Subswitch Element B	See the description for bit 12.
14	Shared Memory Threshold Exceeded, Subswitch Element C	See the description for bit 12.

Bit(s)	Field Name	Description
15	Shared Memory Threshold Exceeded, Subswitch Element D	See the description for bit 12.
16	No Address Interrupt	Set to ‘1’ when an interrupt is generated because a packet is received on an input port when no store address is available.
17	Send Grant Violation	Set to ‘1’ when an attached device removes the send grant for all priorities for the number of packet cycles defined in the send grant antistreaming threshold field in the *Configuration 0 Register* (page 110). The port is identified in the *Send Grant Violation Register* (page 131). This function is enabled only when the link is synchronized and the output queue is enabled.
18	All Queues Empty	Set to ‘1’ by the edge detection that occurs as soon as all the output queues are empty for all four subswitch elements.
19	Header Parity Error	Set to ‘1’ when a parity error is detected in an incoming packet header. The port is identified via the *Header Parity Error Register* (page 132).
20	Address Corruption	Detected by the address manager. Set to ‘1’ when a corrupted address is detected. When an address is corrupted, one or more addresses are lost from the pool of addresses available to store packets. A reset is required to recover lost addresses.
21	Master/Slave Parity Error	Set to ‘1’ when a parity error is detected on the speed-expansion bus.
22	Control Packet Discard	Set to ‘1’ when an incoming control packet is discarded because the control packet or command service packet previously stored in the input controller has not yet been read.
23	Command Service Packet Discard	Set to ‘1’ when an incoming command service packet is discarded because the control packet or command service packet previously stored in the input controller has not yet been read.
24	Control Packet Received on High Channel	Set to ‘1’ when a new control packet is received on a high channel. The port number is identified in the *Control Packet Received on High Channel Register* (page 136).
25	Control Packet Received on Low Channel	Set to ‘1’ when a new control packet is received on a low channel. The port number is identified in the *Control Packet Received on Low Channel Register* (page 136).
26	Command Service Packet Received on High Channel	Set to ‘1’ when a new command service packet is received on a high channel. The port number is identified in the *Command Service Packet Received on High Channel Register* (page 136).
27	Command Service Packet Received on Low Channel	Set to ‘1’ when a new command service packet is received on a low channel. The port number is identified in the *Command Service Packet Received on Low Channel Register* (page 136).
28	All Event-1 Service Packets Received	Set to ‘1’ when the value of the ingress *Event-1 Service Packet Received Register* (page 139) is equal to the value of the ingress *Event-1 Service Packet Mask Register* (page 139).
29	All Event-2 Service Packets Received	Set to ‘1’ when the value of the ingress *Event-2 Service Packet Received Register* (page 139) is equal to the value of the ingress *Event-2 Service Packet Mask Register* (page 139).
30	All Control/Service Packets Transmitted	Set to ‘1’ when a control packet or a service packet has been successfully transmitted to all destinations.
31	Flow Control Violation	Set to ‘1’ when a flow control violation interrupt is generated. For unicast packets, this interrupt is generated when a packet is destined to an output for which neither an output queue grant nor a memory grant has been issued in the past *n* packet cycles. For multicast packets, only the memory grant information is used to detect violations. The flow control violation function is enabled, and *n* is set, in the flow control latency field in the *Configuration 0 Register*. The violating ports are identified by the corresponding bits in the *Flow Control Violation Register* (page 132).

5.3.7 Interrupt Mask Register

This register sets masks for the *Status Register* (page 120) application bits. Note that the occurrence of an event for which the mask bit is set to '1' sets the corresponding bit in the *Status Register* but does not generate an interrupt. For information about an event or error masked here, see the *Status Register* bit descriptions.

Address x‘39’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

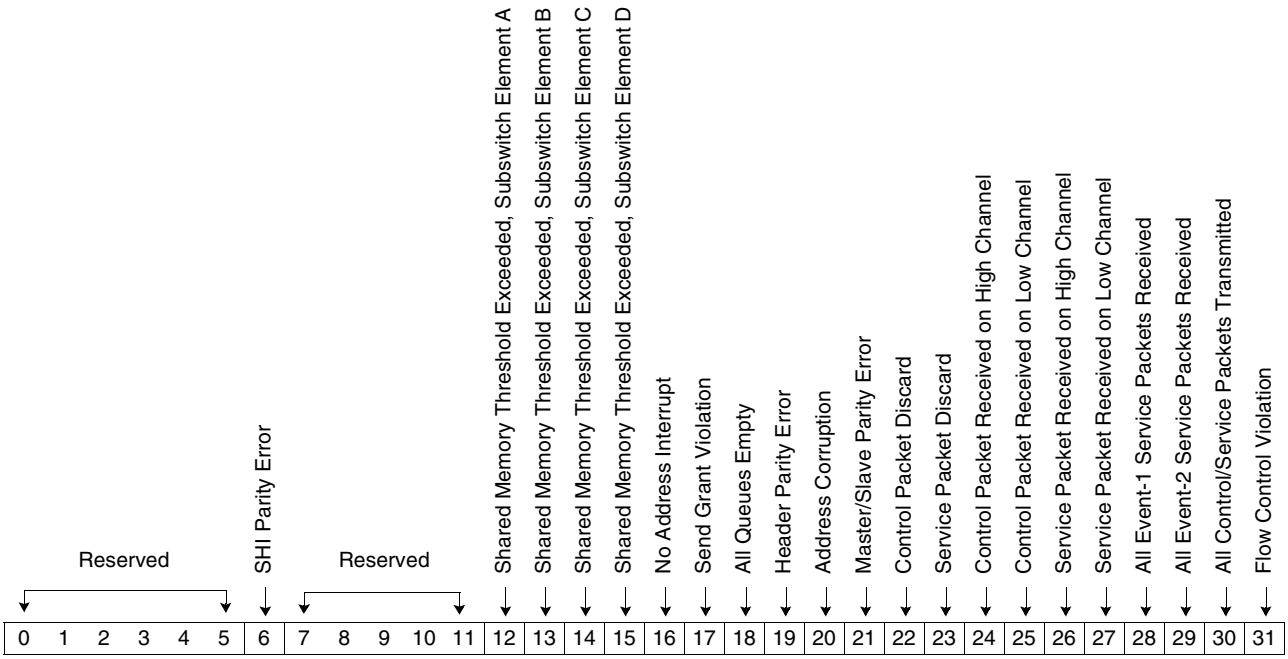

Bit(s)	Field Name	Description
0:5	Reserved	Reserved.
6	SHI Parity Error	When set to ‘1’, masks this interrupt.
7:11	Reserved	Reserved.
12	Shared Memory Threshold Exceeded, Subswitch Element A	When set to ‘1’, masks this interrupt.
13	Shared Memory Threshold Exceeded, Subswitch Element B	When set to ‘1’, masks this interrupt.
14	Shared Memory Threshold Exceeded, Subswitch Element C	When set to ‘1’, masks this interrupt.
15	Shared Memory Threshold Exceeded, Subswitch Element D	When set to ‘1’, masks this interrupt.

Bit(s)	Field Name	Description
16	No Address Interrupt	When set to ‘1’, masks this interrupt.
17	Send Grant Violation	When set to ‘1’, masks this interrupt.
18	All Queues Empty	When set to ‘1’, masks this interrupt.
19	Header Parity Error	When set to ‘1’, masks this interrupt.
20	Address Corruption	When set to ‘1’, masks this interrupt.
21	Master/Slave Parity Error	When set to ‘1’, masks this interrupt.
22	Control Packet Discard	When set to ‘1’, masks this interrupt.
23	Service Packet Discard	When set to ‘1’, masks this interrupt.
24	Control Packet Received on High Channel	When set to ‘1’, masks this interrupt.
25	Control Packet Received on Low Channel	When set to ‘1’, masks this interrupt.
26	Service Packet Received on High Channel	When set to ‘1’, masks this interrupt.
27	Service Packet Received on Low Channel	When set to ‘1’, masks this interrupt.
28	All Event-1 Service Packets Received	When set to ‘1’, masks this interrupt.
29	All Event-2 Service Packets Received	When set to ‘1’, masks this interrupt.
30	All Control/Service Packets Transmitted	When set to ‘1’, masks this interrupt.
31	Flow Control Violation	When set to ‘1’, masks this interrupt.

5.3.8 Output Queue Enable Register

Address x‘3A’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

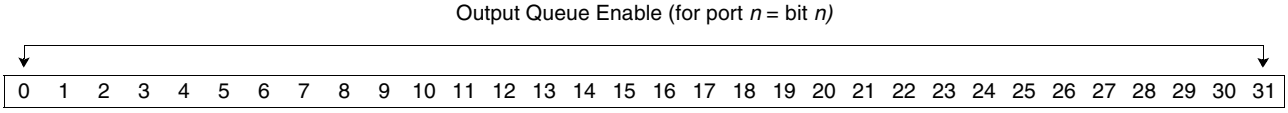

Bit(s)	Field Name	Description
0:31	Output Queue Enable (for port *n* = bit *n)*	1 Enables the output queue for each priority, and the packets destined for that output are enqueued. 0 Disables the output queue for each priority, and the following actions occur: 1. Unicast packets destined to a disabled output queue are discarded. Multicast packets are enqueued only in enabled output queues. 2. A disabled output queue is “slow flushed.” Addresses are dequeued and recycled as in normal operation. The slow flush takes place regardless of the send grant status, as long as the queue is disabled. 3. Output queue grants corresponding to a disabled queue are forced to ‘1’. 4. Idle packets are transmitted if the corresponding bits are set in the *UL TxPort Driver Enable Register* (page 78) and the *UL TxPort Attachment Enable Register* (page 78). **Note:** Control packets and service packets can be transmitted on a port when the output queues are disabled.

5.3.9 Input Controller Enable Register

Address x‘3B’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

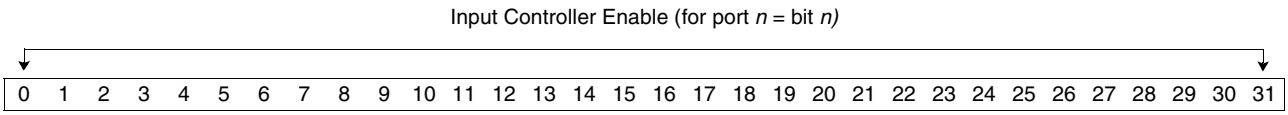

Bit(s)	Field Name	Description
0:31	Input Controller Enable (for port *n* = bit *n)*	1 Enables the port to receive data, control, and service packets and to extract control information (normal setting). 0 Disables the port from receiving data, control, and service packets. The corresponding send grant, subport output queue grants, and subport multicast grants are forced to active. **Note:** This bit is required only for the master device.

5.3.10 Bitmap Filter Register

Address x‘3C’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

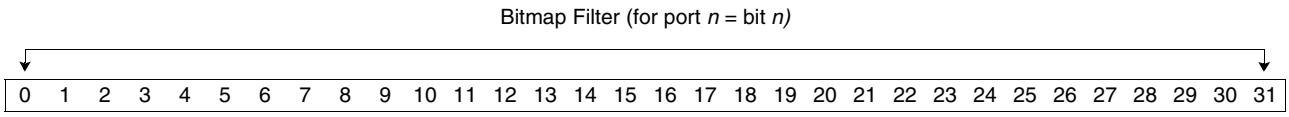

Bit(s)	Field Name	Description
0:31	Bitmap Filter (for port *n* = bit *n)*	Specifies the mask to apply to the received packet destination bitmap for switchover support and load balancing in redundant switch-plane operation. Application of the bitmap filter depends on the packet protection field (see *Table 3-5* on page 34 for more information).

5.3.11 Color Detection Disable Register

Address x‘3D’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

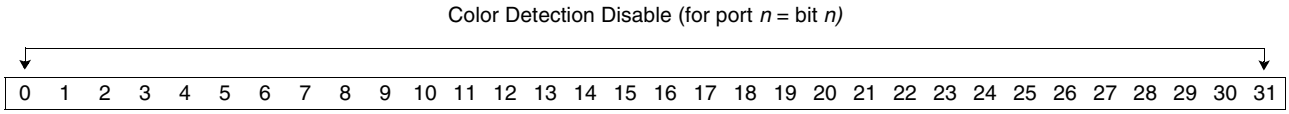

Bit(s)	Field Name	Description
0:31	Color Detection Disable (for port *n* = bit *n)*	When set to ‘1’, disables the input port color detection mechanism and sets the corresponding bit in the *Expected Color Received Register* (page 126). See *Section 3.11 Switchover Support* on page 57.

5.3.12 Expected Color Received Register

Address x‘3E’

Access Type Read Only

Reset Value Undefined

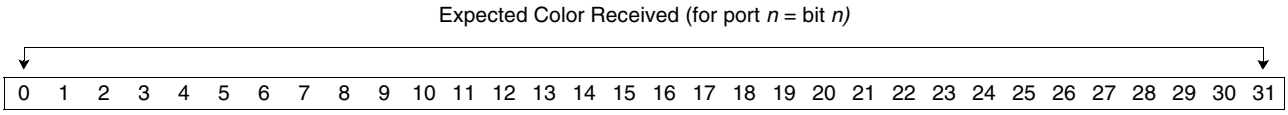

Bit(s)	Field Name	Description
0:31	Expected Color Received (for port *n* = bit *n)*	1 Either the expected color has been received on the input port since the last color clear command or the corresponding bit is set in the *Color Detection Disable Register* (page 125). 0 The color opposite of the expected color is being received on the input port. See *Section 3.11 Switchover Support* on page 57 for more information.

5.3.13 Color Command Register

Address x‘3F’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

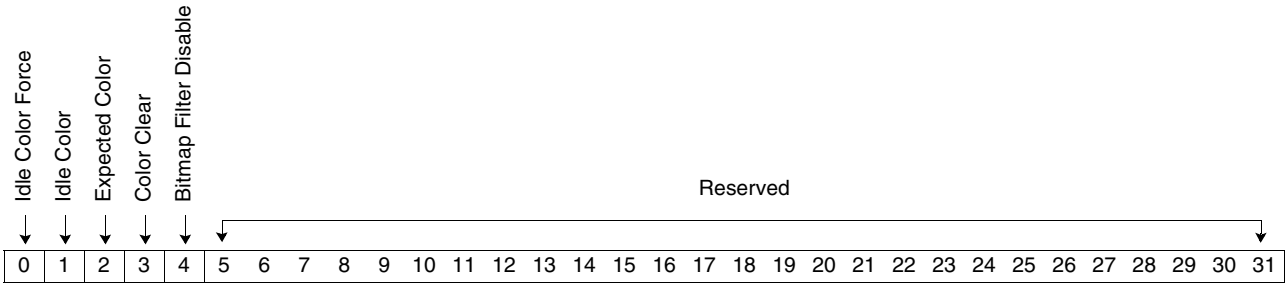

Bit(s)	Field Name	Description
0	Idle Color Force	1 Transmits all output idle packets with the color specified by the idle color bit, regardless of the expected color bit setting. When the color mechanism is not used, this bit must be set to ‘1’. 0 Allows the switchover mechanism to determine the color of output idle packets.
1	Idle Color	Specifies the color given to all idle packets when the idle color force bit is set: 0 Blue idle packets 1 Red idle packets
2	Expected Color	Specifies the expected color of incoming packets after a color clear command is initiated (see bit 3): 0 Blue packets 1 Red packets
3	Color Clear	When set to ‘1’, clears the idle packet color state machine. After this command is initiated, idle packets are transmitted with the color opposite that specified in the expected color bit. Transmission on a given output port takes place only when both of the following conditions are satisfied: • At least one packet of the expected color has been received on all inputs. • The corresponding output queue is empty. If these conditions are not met, idle packets of the opposite color are transmitted on the output port.
4	Bitmap Filter Disable	When set to ‘1’, ingress data packets are processed without bitmap filtering; that is, ingress data packet filtering based on the protection field in the packet qualifier byte (see *Table 3-5* on page 34) and the bitmap filter is *not* performed.
5:31	Reserved	Reserved.

Note: For more information about the function of these bits, see *Section 3.11 Switchover Support* on page 57.

5.3.14 Bitmap Mapping Register

This register is used to map a logical port to any physical port on the same subswitch. It defines the mapping between a bit position in the packet header (bytes H1 and H2) and a physical output queue. That is, all packets received with the bitmap bit specified by the logical bitmap bit position field set to '1' are routed to the physical output queue specified by the physical queue field. For example, if the *Bitmap Mapping Register* is written with the logical bitmap bit position field set to 3 and the physical queue field set to 7, then all the packets received with bitmap bit 3 set to '1' are routed to physical output queue 7.

The *Bitmap Mapping Register* reset value is not rearranged (that is, bitmap bit *n* points to physical output queue *n).* This register cannot be written while the device is in standby.

Write access to the *Bitmap Mapping Register* requires one SHI command:

1. Write the register with the write bit set to '1', the logical bitmap bit position field specified, and the physical queue field specified.

Read access to the *Bitmap Mapping Register* requires two SHI commands:

1. Write the register with the write bit cleared to '0' and the logical bitmap bit position field specified.
2. Read the register to return the corresponding physical queue field value.

Note: Although bitmap mapping is used to map one logical port to one physical port, it must not be used to map one logical port to multiple physical ports or to map multiple logical ports to one physical port. *Failure to follow these requirements will have unknown results.*

In the 256-Gbps configuration, logical ports 0 to 7 must be mapped only to physical ports 0 to 7, and logical ports 8 to 15 must be mapped only to physical ports 8 to 15. In the 512-Gbps configuration, logical ports 0 to 15 must be mapped only to physical ports 0 to 15, and logical ports 16 to 31 must be mapped only to physical ports 16 to 31. *Failure to follow these requirements will have unknown results.*

Address x'40'

Access Type Read/Write

Reset Value '0000 0000 0000 0000 0000 0000 0000 0000'

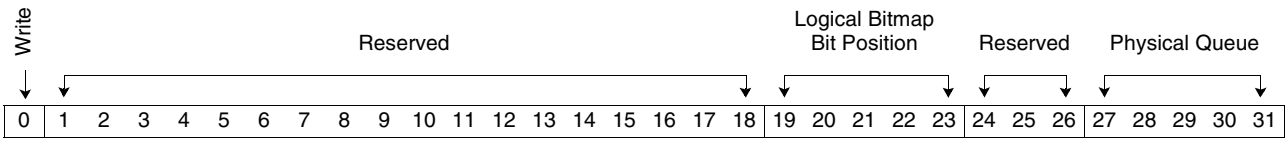

Bit(s)	Field Name	Description
0	Write	1 Specifies a write to the *Bitmap Mapping Register*. 0 Specifies a read from the *Bitmap Mapping Register*.
1:18	Reserved	Reserved.
19:23	Logical Bitmap Bit Position	Specifies the logical bitmap bit position.
24:26	Reserved	Reserved.
27:31	Physical Queue	Specifies the physical output queue.

5.3.15 Output Queue Status Registers

Eight registers provide the status of the 64 output queues (16 output queues per subswitch element). Each of the eight registers is comprised of eight four-bit ranges; each four-bit range describes a single output queue status. Within the four-bit range for an output queue, the first bit is the output queue empty bit, which reports whether the output queue has been emptied since the last read operation, and the other three bits comprise the status field, which reports the highest occupancy level reached by the output queue since the previous read operation.

Address	x‘41’	Subswitch element A output queues 0 to 7	Offset = 0
	x‘42’	Subswitch element A output queues 8 to 15	Offset = 8
	x‘43’	Subswitch element B output queues 0 to 7	Offset = 0
	x‘44’	Subswitch element B output queues 8 to 15	Offset = 8
	x‘45’	Subswitch element C output queues 0 to 7	Offset = 0
	x‘46’	Subswitch element C output queues 8 to 15	Offset = 8
	x‘47’	Subswitch element D output queues 0 to 7	Offset = 0
	x‘48’	Subswitch element D output queues 8 to 15	Offset = 8

Access Type Read/Clear

Reset Value ‘1000 1000 1000 1000 1000 1000 1000 1000’

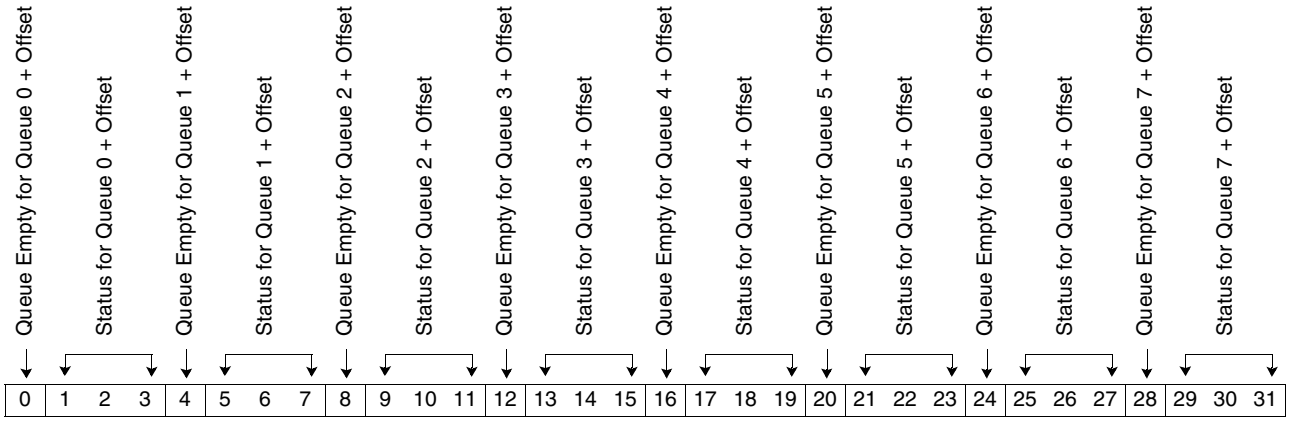

Bit(s)	Field Name	Description
0	Queue Empty for Queue 0 + Offset	When set to ‘1’, the output queue has been emptied since the last read.
1:3	Status for Queue 0 + Offset	Reports that the total number of packets in the output queue has exceeded the threshold for the specified priority(ies). Each time a priority threshold of the output queue is exceeded, the status field for that output queue is updated. Note that the status field does not change when the number of packets of a particular priority falls below the threshold; this field continues to show the output queue status for the highest priority that has exceeded the threshold since the last read. 000 Output queue is not full for any priority 001 Output queue is full for priority 3 010 Output queue is full for priorities 2 and 3 011 Output queue is full for priorities 1, 2, and 3 100 Output queue is full for priorities 0, 1, 2, and 3 Others Reserved
4	Queue Empty for Queue 1 + Offset	See the description for bit 0.

Bit(s)	Field Name	Description
5:7	Status for Queue 1 + Offset	See the description for bits 1:3.
8	Queue Empty for Queue 2 + Offset	See the description for bit 0.
9:11	Status for Queue 2 + Offset	See the description for bits 1:3.
12	Queue Empty for Queue 3 + Offset	See the description for bit 0.
13:15	Status for Queue 3 + Offset	See the description for bits 1:3.
16	Queue Empty for Queue 4 + Offset	See the description for bit 0.
17:19	Status for Queue 4 + Offset	See the description for bits 1:3.
20	Queue Empty for Queue 5 + Offset	See the description for bit 0.
21:23	Status for Queue 5 + Offset	See the description for bits 1:3.
24	Queue Empty for Queue 6 + Offset	See the description for bit 0.
25:27	Status for Queue 6 + Offset	See the description for bits 1:3.
28	Queue Empty for Queue 7 + Offset	See the description for bit 0.
29:31	Status for Queue 7 + Offset	See the description for bits 1:3.

5.3.16 Best-Effort Discard Alarm Register

Address x‘49’

Access Type Read/Clear

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

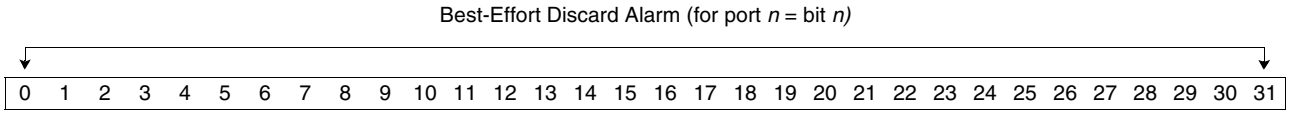

Bit(s)	Field Name	Description
0:31	Best-Effort Discard Alarm (for port *n* = bit *n)*	When set to ‘1’, the best-effort discard logic is in a phase that allows best-effort packets to be discarded for the port.

5.3.17 Send Grant Violation Register

Address x‘4A’

Access Type Read/Clear

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

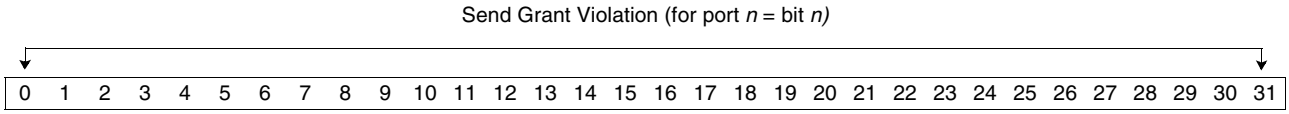

Bit(s)	Field Name	Description
0:31	Send Grant Violation (for port *n* = bit *n)*	When set to ‘1’, reports that the port has removed the send grant for too long and send grant antistreaming is in progress (see the send grant antistreaming enable and send grant antistreaming threshold fields in the *Configuration 0 Register* [page 110]). This register is cleared when read, but the violations detected during the read/clear are not lost.

5.3.18 Header Parity Error Register

Address x‘4B’

Access Type Read/Clear

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

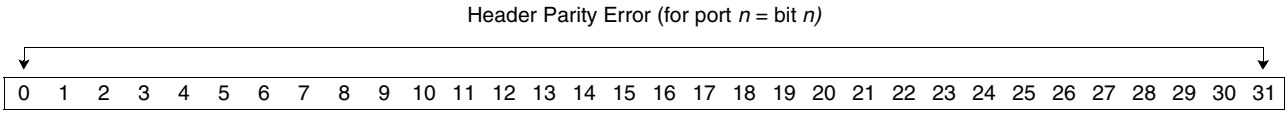

Bit(s)	Field Name	Description
0:31	Header Parity Error (for port *n* = bit *n)*	When set to ‘1’, the input port has detected a packet header parity error. This register is cleared when read, but the errors detected during the read/clear are not lost.

5.3.19 Flow Control Violation Register

Address x‘4C’

Access Type Read/Clear

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

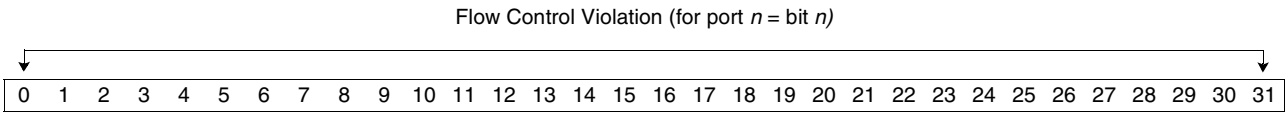

Bit(s)	Field Name	Description
0:31	Flow Control Violation (for port *n* = bit *n)*	When set to ‘1’, the input port has generated a flow control violation interrupt. This register is cleared when read, but the violations detected during the read/clear are not lost.

5.3.20 Side Communication Channel Input Reporting Registers

Address

x‘4D’ Reports SCC bit 0
x‘4E’ Reports SCC bit 1
x‘4F’ Reports SCC bit 2
x‘50’ Reports SCC bit 3

Access Type Read Only

Reset Value Undefined

Side Communication Channel Input Reporting (for port *n* = bit *n)*

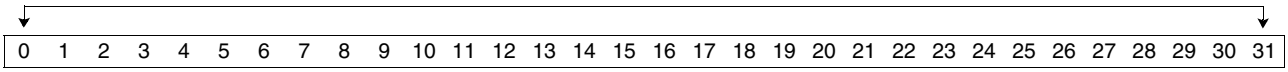

Bit(s)	Field Name	Description
0:31	Side Communication Channel Input Reporting (for port *n* = bit *n)*	Reports the port information extracted from idle packet SCC input bit *y,* where *y* equals 0, 1, 2, or 3 depending on the register.

5.4 Control Packet and Service Packet Transmission Registers

5.4.1 Egress Control Packet and Service Packet Payload Registers

The local processor uses six registers to prepare the payload for egress control packets and service packets. The packet type is defined in the packet qualifier byte (H0). For control packets and command service packets, all eight LUs must be prepared before packet transmission is requested. Transmission is requested when the *Egress Control Packet and Service Packet Destination Register* (page 135) is written. For event-1 and event-2 service packets, only the LU that contains the packet qualifier byte (that is, the first LU) is required; the remaining LU bytes are ignored.

Address

x‘51’ Master LU bytes 0 to 3
x‘52’ Master LU bytes 4 to 7
x‘53’ Master LU bytes 8 and 9, plus 16 reserved bits

Three additional registers are used for internal speed expansion:

x‘54’ Slave LU bytes 0 to 3
x‘55’ Slave LU bytes 4 to 7
x‘56’ Slave LU bytes 8 and 9, plus 16 reserved bits

Access Type

Read/Write

Reset Value

‘0000 0000 0000 0000 0000 0000 0000 0000’

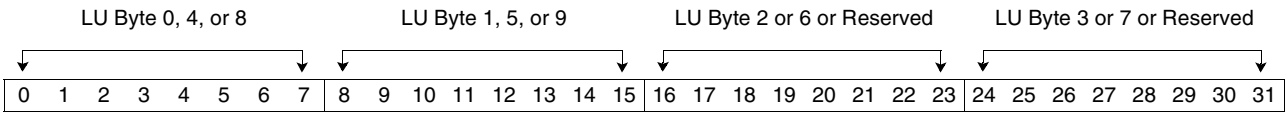

Bit(s)	Field Name	Description
0:7	LU Byte 0, 4, or 8	Depending on the register address, specifies LU byte 0, 4, or 8 of the egress control packet or service packet payload.
8:15	LU Byte 1, 5, or 9	Depending on the register address, specifies LU byte 1, 5, or 9 of the egress control packet or service packet payload.
16:23	LU Byte 2 or 6 or Reserved	Depending on the register address, specifies LU byte 2 or 6 of the egress control packet or service packet payload, or is reserved.
24:31	LU Byte 3 or 7 or Reserved	Depending on the register address, specifies LU byte 3 or 7 of the egress control packet or service packet payload, or is reserved.

5.4.2 Egress Control Packet and Service Packet Destination Register

The local processor uses this register to specify the destination bitmap for egress control packets and service packets. When a port transmits either a control packet or a service packet, the corresponding bit is cleared in this register. When the last port transmits the packet (that is, when the register value returns to x‘0000 0000’), an “all control/service packets transmitted” interrupt (see the *Status Register* [page 120]) is generated. If, for any reason, a port does not transmit the control packet or service packet, the application can reset the corresponding bit in this register without generating an interrupt.

Note: This register is mapped to physical ports. With internal speed expansion, logical port *n* corresponds to physical port $n \times 2$. For example, to send a packet to logical port 3, bit 6 must be set in this register.

Address x‘57’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

Egress Control/Service Packet Destination (for port *n* = bit *n)*

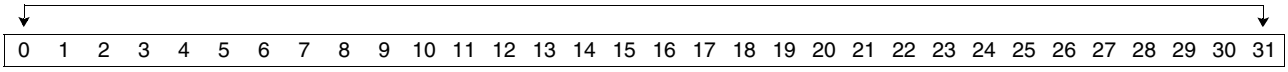

Bit(s)	Field Name	Description
0:31	Egress Control/Service Packet Destination (for port *n* = bit *n)*	When set to ‘1’, specifies the egress control packet or service packet destination.

5.5 Control Packet and Service Packet Reception Registers

Ingress control packets and command service packets are stored in input controller internal registers. Each input controller can process one of these packets from the high channel and one from the low channel. The packet payload is transferred to the local processor via the *Ingress Control Packet and Service Packet Payload Registers* (page 138). The input controller and channel of the packet being processed are identified in the *Ingress Control Packet and Service Packet Source Register* (page 137).

5.5.1 Ingress Control Packet or Command Service Packet Received Registers

These registers report that a control packet or a command service packet has been received on a port. Each time a bit is set in one of these registers, the corresponding bit (that is, control packet received on high channel, control packet received on low channel, service packet received on high channel, or service packet received on low channel) is set in the *Status Register* (page 120).

After the local processor reads the packet payload, the corresponding bit in the *Ingress Control Packet or Command Service Packet Received Registers* is automatically cleared, and the input controller can store a new packet for that channel.

Address

x‘58’	Control packet received on high channel
x‘59’	Control packet received on low channel
x‘5A’	Command service packet received on high channel
x‘5B’	Command service packet received on low channel

Access Type Read Only

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

Ingress Control/Command Service Packet Received (for port *n* = bit *n)*

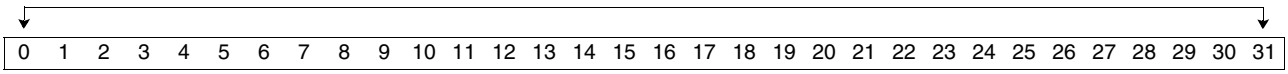

Bit(s)	Field Name	Description
0:31	Ingress Control/ Command Service Packet Received (for port *n* = bit *n)*	When set to ‘1’, reports that a control packet or command service packet has been received on the corresponding port and channel.

5.5.2 Ingress Control Packet and Service Packet Source Register

This register identifies the source (that is, the input controller and channel) of the received packet being processed using the *Ingress Control Packet and Service Packet Payload Registers* (page 138).

The *Ingress Control Packet or Command Service Packet Received Registers* (page 136) identify all input controllers with a stored control packet or command service packet. The local processor determines from which of these input controllers to retrieve a packet and sets this register accordingly.

Address x‘5C’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

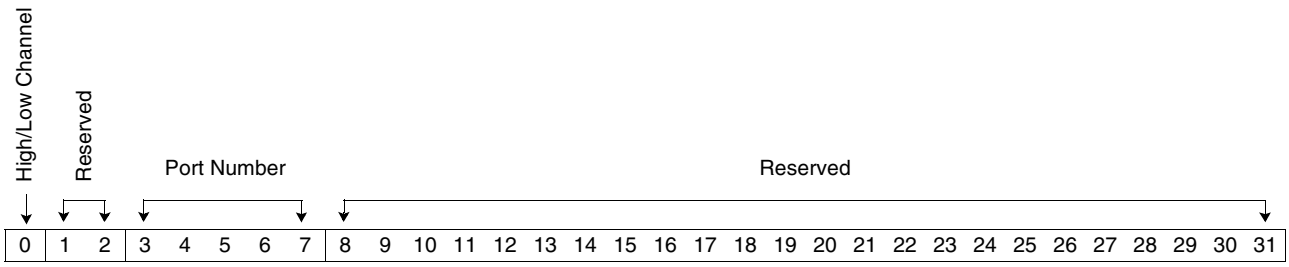

Bit(s)	Field Name	Description
0	High/Low Channel	Specifies the channel: 1 High channel 0 Low channel
1:2	Reserved	Reserved.
3:7	Port Number	Specifies the port.
8:31	Reserved	Reserved.

5.5.3 Ingress Control Packet and Service Packet Payload Registers

These three registers transfer the payload of ingress control packets and command service packets from the input controller to the local processor.

When the x‘5D’ register is read on the master device, another control packet or command service packet can be stored by the input controller. Consequently, the local processor must read the slave device LUs before it reads the master device LUs. In the master device, the local processor must read the *Ingress Control Packet and Service Packet Payload Registers* in the following order:

1. x‘5F’
2. x‘5E’
3. x‘5D’

Reading the x‘5D’ register also resets the corresponding bit in the *Ingress Control Packet or Command Service Packet Received Register* (page 136). If this register is still not empty after the bit is reset, then another packet must be processed, and the corresponding interrupt bit is asserted in the *Status Register* (page 120).

Address

x‘5D’	LU bytes 0 to 3
x‘5E’	LU bytes 4 to 7
x‘5F’	LU bytes 8 and 9, plus 16 reserved bits (required only if the packet size is greater than 8 LUs)

Access Type Read Only

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

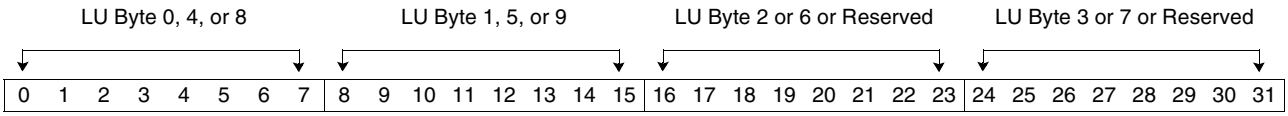

Bit(s)	Field Name	Description
0:7	LU Byte 0, 4, or 8	Depending on the register address, reports LU byte 0, 4, or 8 of the ingress control packet or service packet payload.
8:15	LU Byte 1, 5, or 9	Depending on the register address, reports LU byte 1, 5, or 9 of the ingress control packet or service packet payload.
16:23	LU Byte 2 or 6 or Reserved	Depending on the register address, reports LU byte 2 or 6 of the ingress control packet or service packet payload, or is reserved.
24:31	LU Byte 3 or 7 or Reserved	Depending on the register address, reports LU byte 3 or 7 of the ingress control packet or service packet payload, or is reserved.

5.5.4 Ingress Event Service Packet Received Registers

These registers report that an event-1 or event-2 service packet has been received on a port.

Address x‘60’ Event-1 service packet received
x‘61’ Event-2 service packet received

Access Type Read/Clear

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

Ingress Event Service Packet Received (for port *n* = bit *n)*

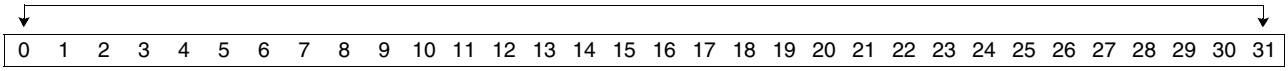

Bit(s)	Field Name	Description
0:31	Ingress Event Service Packet Received (for port *n* = bit *n)*	When set to ‘1’, the port has received an event-1or event-2 service packet.

5.5.5 Ingress Event Service Packet Mask Registers

These registers set the masks applied to the *Ingress Event Service Packet Received Registers* for *Status Register* (page 120) reporting.

Address x‘62’ Event-1 service packet mask
x‘63’ Event-2 service packet mask

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

Ingress Event Service Packet Mask (for port *n* = bit *n)*

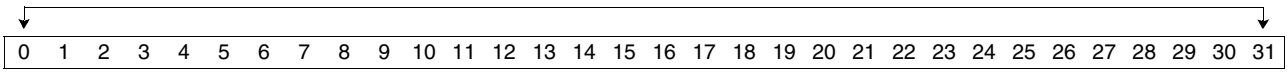

Bit(s)	Field Name	Description
0:31	Ingress Event Service Packet Mask (for port *n* = bit *n)*	Sets the mask applied to the *Ingress Event Service Packet Received Register* for *Status Register* (page 120) reporting. When the value of the *Ingress Event Service Packet Received Register* is equal to the value of the *Ingress Event Service Packet Mask Register,* the “all event service packets received” bit is set in the *Status Register*.

5.6 Debug Facilities Registers

5.6.1 Debug Bus Select Register

Address x‘64’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

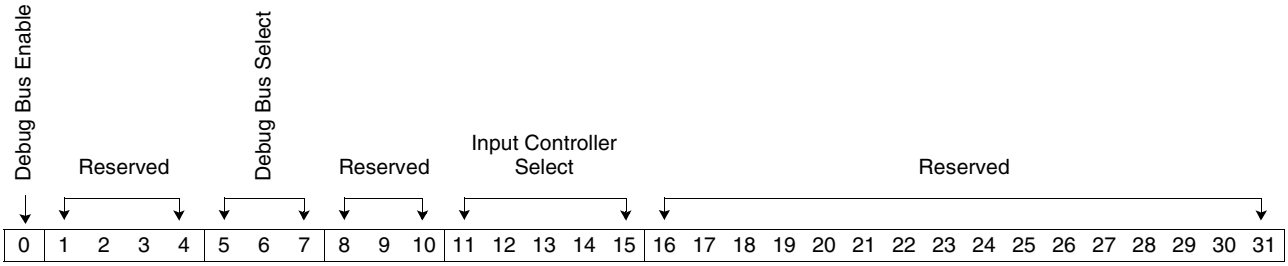

Bit(s)	Field Name	Description
0	Debug Bus Enable	When set to ‘1’, enables the debug bus drivers and the Osc625TestPointOut driver.
1:4	Reserved	Reserved.
5:7	Debug Bus Select	Specifies the device element for which the DebugBusOut[0:15] pins provide information: 000 Sequencer information 001 Input controller information (selected by the input controller select field) 010 Packet routing switch general information Others Reserved *Table 5-2* presents the DebugBusOut[0:15] pin information for each of the device elements.
8:10	Reserved	Reserved.
11:15	Input Controller Select	Specifies the input controller.
16:31	Reserved	Reserved.

Table 5-2. DebugBusOut[0:15] Pin Information by Debug Bus Select Field Value (Page 1 of 2)

Debug Bus Select Field Value	DebugBusOut[0:15] Pin(s)	Description
‘000’	**Sequencer Information**	
	DebugBusOut[0:3]	Sequencer internal count.
	DebugBusOut[4]	Not used.
	DebugBusOut[5]	NotReset (NotStandby).
	DebugBusOut[6]	Sequencer reset pulse from 16-ns clock domain.
	DebugBusOut[7]	Sequencer packet clock.
	DebugBusOut[8:15]	Sequencer counter for speed-expansion bus delay measurement.

Table 5-2. DebugBusOut[0:15] Pin Information by Debug Bus Select Field Value (Page 2 of 2)

Debug Bus Select Field Value	DebugBusOut[0:15] Pin(s)	Description
‘001’ (input controller select field selecting 1 of 32 input controllers)	**Input Controller Information**	
	DebugBusOut[0:3]	High-channel byte counter.
	DebugBusOut[4]	High-channel control packet received.
	DebugBusOut[5]	High-channel command service packet received.
	DebugBusOut[6]	High-channel data packet routed to high subswitch element.
	DebugBusOut[7]	High-channel data packet routed to low subswitch element.
	DebugBusOut[8]	Low-channel control packet received.
	DebugBusOut[9]	Low-channel command service packet received.
	DebugBusOut[10]	Low-channel data packet routed to high subswitch element.
	DebugBusOut[11]	Low-channel data packet routed to low subswitch element.
	DebugBusOut[12]	High channel has no address.
	DebugBusOut[13]	Low channel has no address.
	DebugBusOut[14]	High-channel flow control violation.
	DebugBusOut[15]	Low-channel flow control violation.
‘010’	**Packet Routing Switch General Information**	
	DebugBusOut[0]	Synchronous flush.
	DebugBusOut[1]	Core PLL feedback
	DebugBusOut[2]	B clock at 16 ns.
	DebugBusOut[3]	C clock at 16 ns.
	DebugBusOut[4]	B clock control packet at 8 ns.
	DebugBusOut[5]	C clock control packet at 8 ns.
	DebugBusOut[6]	B clock data packet at 8 ns.
	DebugBusOut[7]	C clock data packet at 8 ns.
	DebugBusOut[8]	Memory BIST controller ACLK pin.
	DebugBusOut[9]	Memory BIST controller BCLK pin.
	DebugBusOut[10]	Memory BIST controller CCLK pin.
	DebugBusOut[11]	Memory BIST controller STCLK pin.
	DebugBusOut[12]	Memory BIST controller ENABLE pin.
	DebugBusOut[13]	Memory BIST controller LBIST pin.
	DebugBusOut[14]	Memory BIST controller TESTM1 pin.
	DebugBusOut[15]	Memory BIST controller TESTM3 pin.

5.6.2 Send Grant Disable Register

Address x‘65’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

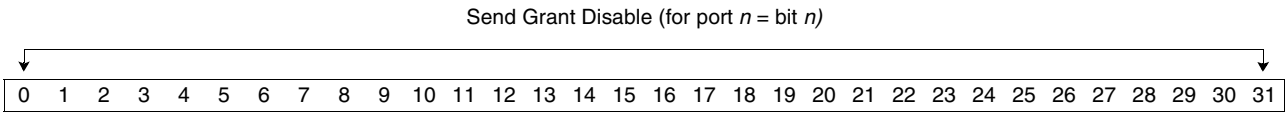

Bit(s)	Field Name	Description
0:31	Send Grant Disable (for port *n* = bit *n)*	1 Disables packet transmission on the port regardless of the value of the send grant signal, unless the corresponding bit is set in the *Force Send Grant Register.* 0 Enables normal send grant signal decoding and packet transmission on the port if the retrieved send grant information is active.

5.6.3 Force Send Grant Register

Address x‘66’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

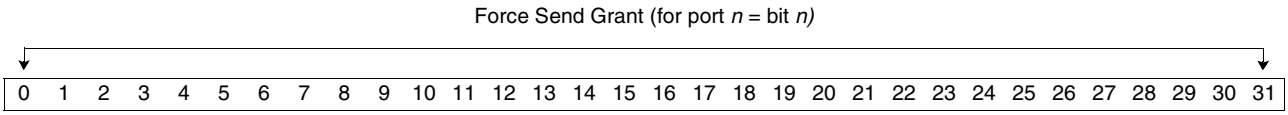

Bit(s)	Field Name	Description
0:31	Force Send Grant (for port *n* = bit *n)*	When set to ‘1’, forces the send grant to active for all priorities on the port regardless of the send grant status extracted from the ingress packet headers. This register is used only when the corresponding bit is set to ‘1’ in the *Send Grant Disable Register.*

5.6.4 Send Grant Status Registers

These four registers report the send grant status, which is extracted from ingress packet headers.

Addresses

x‘67’ Priority 0
x‘68’ Priority 1
x‘69’ Priority 2
x‘6A’ Priority 3

Access Type

Read Only

Reset Value

‘1111 1111 1111 1111 1111 1111 1111 1111’ for x‘67’

‘0000 0000 0000 0000 0000 0000 0000 0000’ for x‘68’, x‘69’, and x‘6A’

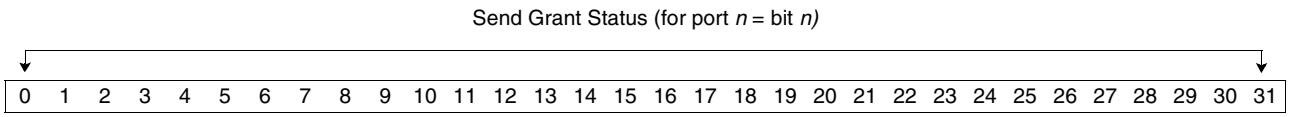

Bit(s)	Field Name	Description
0:31	Send Grant Status (for port *n* = bit *n)*	1 The send grant is active for the port and priority combination. 0 The send grant is inactive for the port and priority combination.

5.6.5 Subswitch Element Occupancy (1) Registers

Address

x‘6B’ Subswitch element A
x‘6C’ Subswitch element B
x‘6D’ Subswitch element C
x‘6E’ Subswitch element D

Access Type Read Only

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

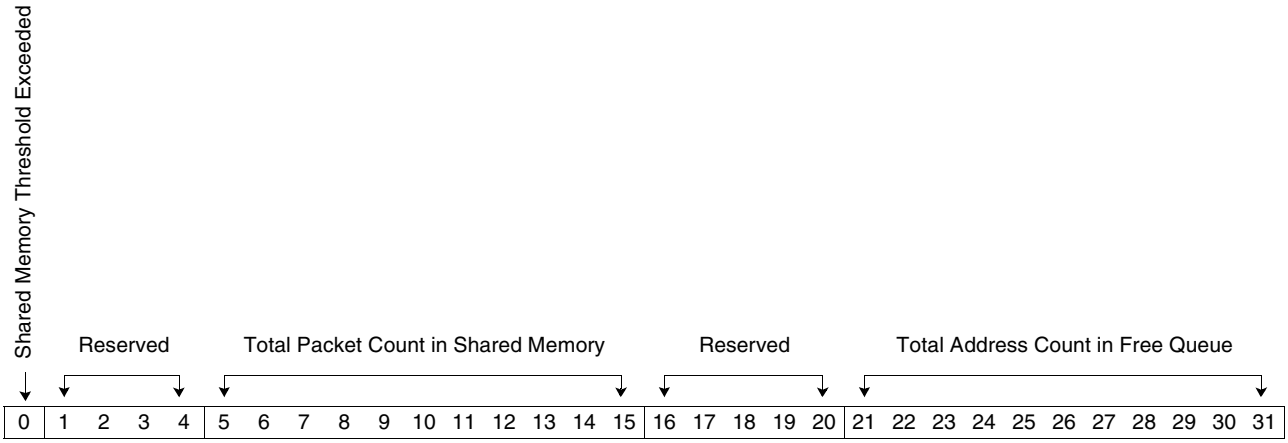

Bit(s)	Field Name	Description
0	Shared Memory Threshold Exceeded	When set to ‘1’, indicates that the number of packets currently in the subswitch element shared memory is greater than or equal to the shared memory threshold. This field is continuously refreshed.
1:4	Reserved	Reserved.
5:15	Total Packet Count in Shared Memory	Reports the number of packets that currently occupy the subswitch element shared memory. This field is continuously refreshed.
16:20	Reserved	Reserved.
21:31	Total Address Count in Free Queue	Reports the number of shared memory addresses currently available in the subswitch element. This field is continuously refreshed.

5.6.6 Subswitch Element Occupancy (2) Registers

Address

x‘6F’ Subswitch element A
x‘70’ Subswitch element B
x‘71’ Subswitch element C
x‘72’ Subswitch element D

Access Type Read Only

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

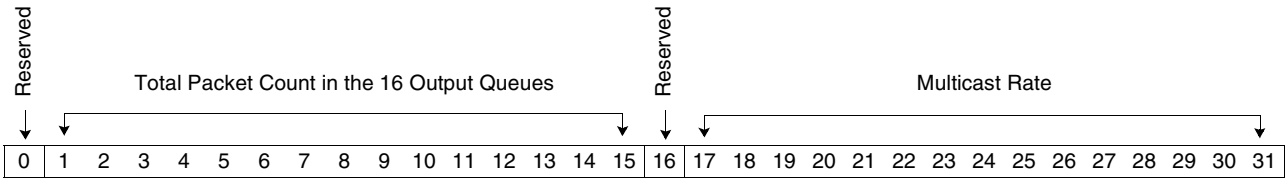

Bit(s)	Field Name	Description
0	Reserved	Reserved.
1:15	Total Packet Count in the 16 Output Queues	Reports the number of packets that currently occupy the 16 subswitch element output queues. This field is continuously refreshed.
16	Reserved	Reserved.
17:31	Multicast Rate	Reports the difference between the number of packets in the 16 subswitch element output queues and the number of packets in the subswitch element shared memory. This field is continuously refreshed.

5.6.7 Look-Up Table Registers

Two *Look-Up Table Registers* provide indirect access to one look-up table. The look-up table has ten entries that specify how the first ten data bytes of each data row of a byte stream are rearranged before transmission. The entry at location *n* of the look-up table specifies the data byte sent as the n^{th} byte in the data stream. The x‘73’ register provides access to look-up table entries 0 through 4, and the x‘74’ register provides access to look-up table entries 5 through 9. The look-up table reset value specifies the normal byte order, from 0 to 9 (no rearrangement). The look-up table cannot be written while the device is in standby.

Address x‘73’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0001 0010 0011 0100’

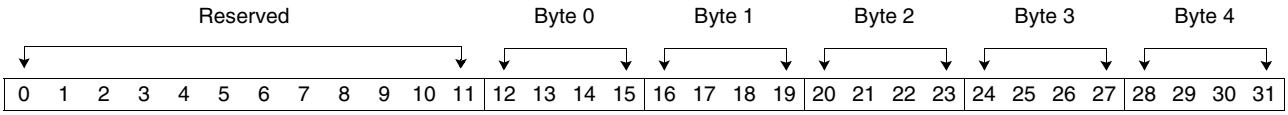

Bit(s)	Field Name	Description
0:11	Reserved	Reserved.
12:15	Byte 0	Look-up table byte 0, which specifies the byte sent in place of LU byte 0.
16:19	Byte 1	Look-up table byte 1, which specifies the byte sent in place of LU byte 1.
20:23	Byte 2	Look-up table byte 2, which specifies the byte sent in place of LU byte 2.
24:27	Byte 3	Look-up table byte 3, which specifies the byte sent in place of LU byte 3.
28:31	Byte 4	Look-up table byte 4, which specifies the byte sent in place of LU byte 4.

Address x‘74’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0101 0110 0111 1000 1001’

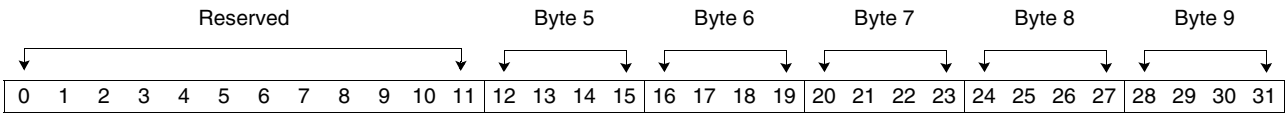

Bit(s)	Field Name	Description
0:11	Reserved	Reserved.
12:15	Byte 5	Look-up table byte 5, which specifies the byte sent in place of LU byte 5.
16:19	Byte 6	Look-up table byte 6, which specifies the byte sent in place of LU byte 6.
20:23	Byte 7	Look-up table byte 7, which specifies the byte sent in place of LU byte 7.
24:27	Byte 8	Look-up table byte 8, which specifies the byte sent in place of LU byte 8.
28:31	Byte 9	Look-up table byte 9, which specifies the byte sent in place of LU byte 9.

5.6.8 Blue Idle Packet or Data Packet Received Register

Address x‘75’

Access Type Read/Clear

Reset Value Undefined

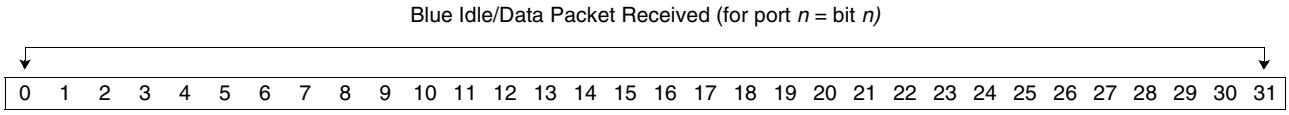

Bit(s)	Field Name	Description
0:31	Blue Idle/Data Packet Received (for port *n* = bit *n)*	When set to ‘1’, the port has received at least one blue idle packet or data packet. The register is cleared when read, but the packet color received during the read/clear is not lost. This register is used for testing.

5.6.9 Red Idle Packet or Data Packet Received Register

Address x‘76’

Access Type Read/Clear

Reset Value Undefined

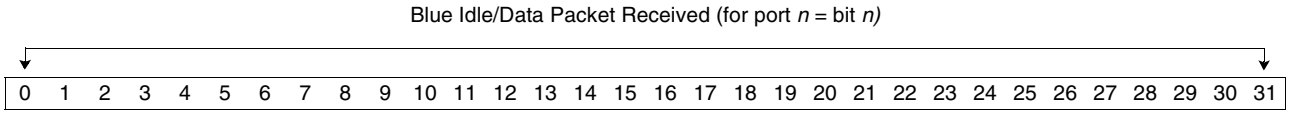

Bit(s)	Field Name	Description
0:31	Blue Idle/Data Packet Received (for port *n* = bit *n)*	When set to ‘1’, the port has received at least one red idle packet or data packet. The register is cleared when read, but the packet color received during the read/clear is not lost. This register is used for testing.

5.6.10 Miscellaneous Debug Register

Address x‘77’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

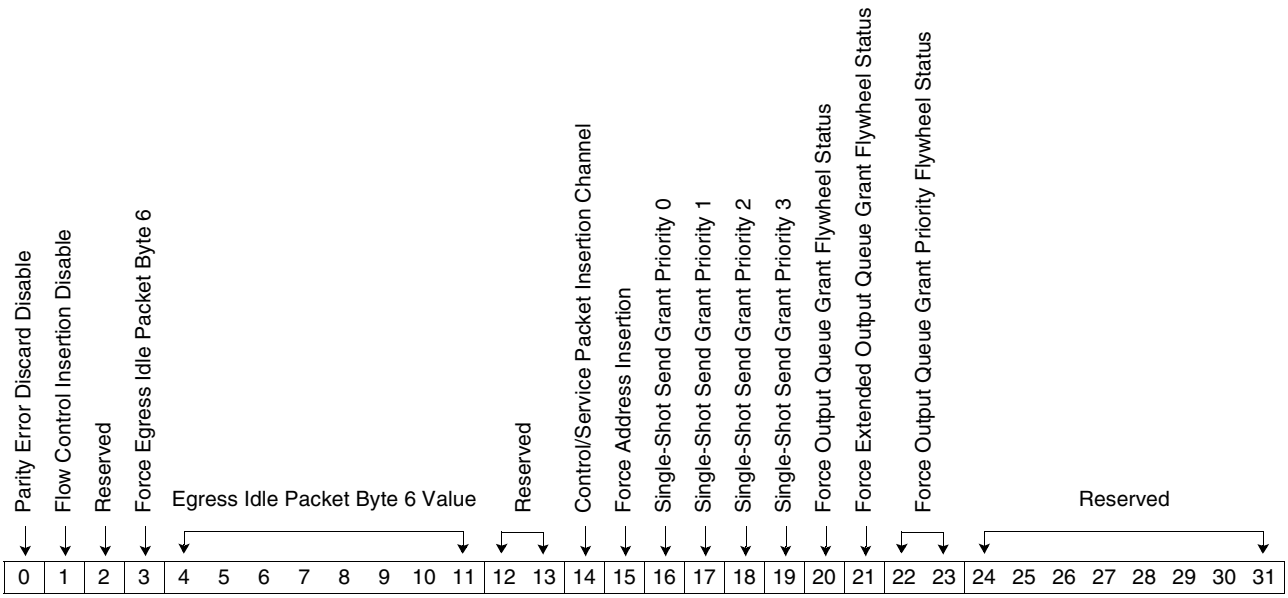

Bit(s)	Field Name	Description
0	Parity Error Discard Disable	1 Disables the input controllers from discarding ingress packets with invalid header parity. These ingress packets are received as normal packets instead. 0 Enables the input controllers to discard ingress packets with invalid header parity. In either case, the header parity error bit is asserted in the *Status Register* (page 120) and the bit corresponding to the input port is set in the *Header Parity Error Register* (page 132).
1	Flow Control Insertion Disable	1 Disables the insertion of flow control information in egress packet headers. Egress data packet headers still contain the destination bitmap. 0 Enables the insertion of flow control information in egress packet headers.
2	Reserved	Reserved.
3	Force Egress Idle Packet Byte 6	When set to ‘1’, replaces the value of byte 6 in egress idle packets with the value defined in the egress idle packet byte 6 value field. This option is used only in the switch loopback test configuration.
4:11	Egress Idle Packet Byte 6 Value	Specifies the value that replaces the value of byte 6 in egress idle packets when the “force egress idle packet byte 6 enable” bit is set to ‘1’.
12:13	Reserved	Reserved.
14	Control/Service Packet Insertion Channel	Specifies the channel used to transmit egress control packets and service packets from the output controller to the attached device: 1 Forces the insertion of control packets and service packets on the high channel. 0 Forces the insertion of control packets and service packets on the low channel.

Bit(s)	Field Name	Description
15	Force Address Insertion	When set to ‘1’, enables the input controller to insert the input channel and port number into the H1 header byte of ingress data packets before storing the packets in the shared memory. The H1 header byte is encoded as follows: Bits 0:1 = ‘00’ Bit 2 = ‘1’ for the high channel or ‘0’ for the low channel Bits 3:7 = input port number (0 to 31) This bit is used for testing.
16	Single-Shot Send Grant Priority 0	When set to ‘1’, enables the *Force Send Grant Register* (page 142; when set to ‘1’) to generate an internal send grant for a single packet of priority 0. (This bit cannot be used to generate a send grant for control packet transmission.)
17	Single-Shot Send Grant Priority 1	When set to ‘1’, enables the *Force Send Grant Register* (when set to ‘1’) to generate an internal send grant for a single packet of priority 1. (This bit cannot be used to generate a send grant for control packet transmission.)
18	Single-Shot Send Grant Priority 2	When set to ‘1’, enables the *Force Send Grant Register* (when set to ‘1’) to generate an internal send grant for a single packet of priority 2. (This bit cannot be used to generate a send grant for control packet transmission.)
19	Single-Shot Send Grant Priority 3	When set to ‘1’, enables the *Force Send Grant Register* (when set to ‘1’) to generate an internal send grant for a single packet of priority 3. (This bit cannot be used to generate a send grant for control packet transmission.)
20	Force Output Queue Grant Flywheel Status	When set to ‘1’: • Forces the value of the extended output queue grant flywheel status bit of the packet qualifier byte (H0) of low-channel egress idle packets to the value specified in the force extended output queue grant flywheel status bit (see bit 21). • Forces the value of the output queue grant priority flywheel status field of the packet qualifier byte of low-channel egress idle packets to the value specified in the force output queue grant priority flywheel status field (see bits 22:23). This bit is used for testing. For idle packet qualifier byte field descriptions, see *Table 3-11* on page 40.
21	Force Extended Output Queue Grant Flywheel Status	Specifies the forced value of the extended output queue grant flywheel status bit of the packet qualifier byte of low-channel egress idle packets when the force output queue grant flywheel status bit is set to ‘1’.
22:23	Force Output Queue Grant Priority Flywheel Status	Specifies the forced value of the output queue grant priority flywheel status field of the packet qualifier byte of low-channel egress idle packets when the force output queue grant flywheel status bit is set to ‘1’.
24:31	Reserved	Reserved.

5.6.11 Force Packet Capture Ports Register

This register is used to request that an input controller on the master device capture an ingress packet in the *Ingress Control Packet and Service Packet Payload Registers* (page 138). The following registers are also involved:

- *Force Packet Capture Header Register* (page 151)
- *Force Packet Capture Mask Register* (page 152)
- *Packet Captured Registers* (page 152)

The *Force Packet Capture Ports Register* designates the input controllers that are to capture the ingress packet. An input controller starts the process on the rising edge of the corresponding register bit. Each time the master input controller receives an LU after it detects the force packet capture request, the master input controller checks whether the following is true:

(ingress packet header AND force packet capture mask) = (force packet capture header AND force packet capture mask)

If the above is true for an ingress packet (regardless of packet type), then the packet LUs are saved in the *Ingress Control Packet and Service Packet Payload Registers* and the corresponding bit is set in the requested *Packet Captured Register* (either the high- or low-channel register, as appropriate). Force packet capture is performed only if the *Ingress Control Packet and Service Packet Payload Registers* are available.

For example, to capture a packet with the packet qualifier byte (H0) bit 0 set to ‘1’ and the packet header byte (H1) bit 3 set to ‘0’ on any port, the required operations are as follows:

1. Write the *Force Packet Capture Header Register* with the value x‘8000 0000’.
2. Write the *Force Packet Capture Mask Register* with the value x‘8010 0000’.
3. Write the *Force Packet Capture Ports Register* with the value x‘FFFF FFFF’.
4. Poll the *Packet Captured Registers* until a ‘1’ appears. This step identifies the port on which the requested packet was captured.
5. Read the packet payload using the *Ingress Control Packet and Service Packet Payload Registers*.
6. Reset the *Force Packet Capture Ports Register*.

Address x‘78’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

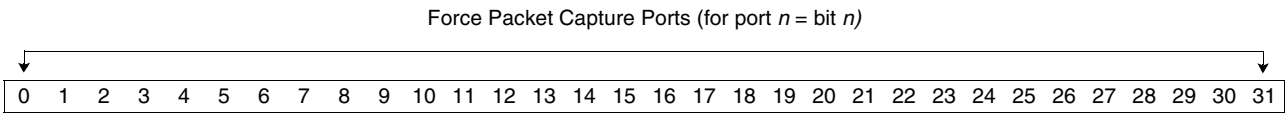

Bit(s)	Field Name	Description
0:31	Force Packet Capture Ports (for port *n* = bit *n)*	When set to ‘1’, enables the port input controller to capture an ingress packet that meets the requirements established by the *Force Packet Capture Header Register* (page 151) and the *Force Packet Capture Mask Register* (page 152).

5.6.12 Force Packet Capture Header Register

This register specifies the header value of the ingress packet to be captured.

Address x‘79’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

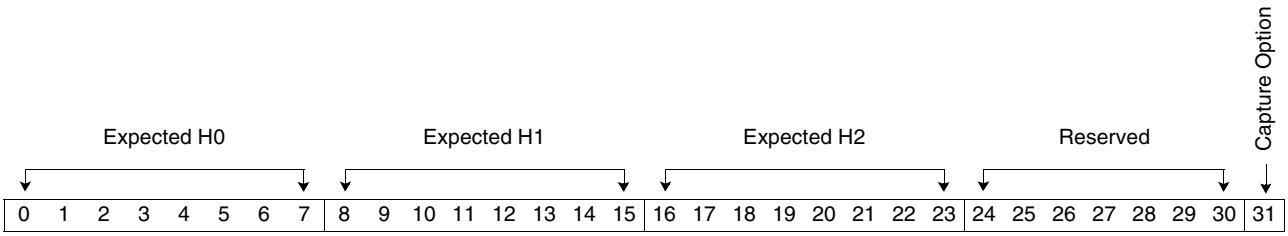

Bit(s)	Field Name	Description
0:7	Expected H0	To be taken into account for packet capture, must be set to ‘1’ or ‘0’ to specify the expected value of the ingress packet H0 byte.
8:15	Expected H1	To be taken into account for packet capture, must be set to ‘1’ or ‘0’ to specify the expected value of the ingress packet H1 byte.
16:23	Expected H2	To be taken into account for packet capture, must be set to ‘1’ or ‘0’ to specify the expected value of the ingress packet H2 byte.
24:30	Reserved	Reserved.
31	Capture Option	If the captured packet is a data packet, specifies whether the packet is propagated to the shared memory: 0 The captured data packet is stored in the *Ingress Control Packet and Service Packet Payload Registers* (page 138) and is propagated to the shared memory. This setting allows port supervision without impacting data traffic. 1 The captured data packet is stored in the *Ingress Control Packet and Service Packet Payload Registers* and is not propagated to the shared memory.

5.6.13 Force Packet Capture Mask Register

This register specifies the header bits that must be taken into account for ingress packet capture.

Address x‘7A’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

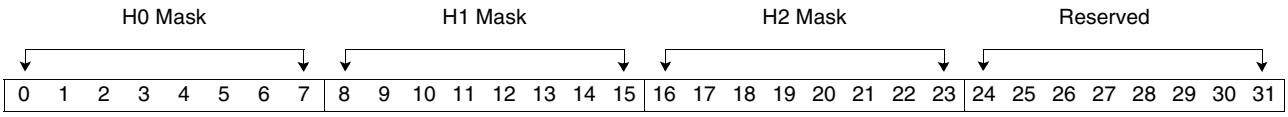

Bit(s)	Field Name	Description
0:7	H0 Mask	When set to ‘1’, the corresponding bit of the ingress packet H0 byte must be taken into account for packet capture.
8:15	H1 Mask	When set to ‘1’, the corresponding bit of the ingress packet H1 byte must be taken into account for packet capture.
16:23	H2 Mask	When set to ‘1’, the corresponding bit of the ingress packet H2 byte must be taken into account for packet capture.
24:31	Reserved	Reserved.

5.6.14 Packet Captured Registers

These registers report that the ingress packet capture requested by the *Force Packet Capture Ports Register* (page 150) has been stored in the *Ingress Control Packet and Service Packet Payload Registers* (page 138). The bit is automatically cleared when the packet payload has been read.

Address x‘7B’ Packet captured on high channel
x‘7C’ Packet captured on low channel

Access Type Read Only

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

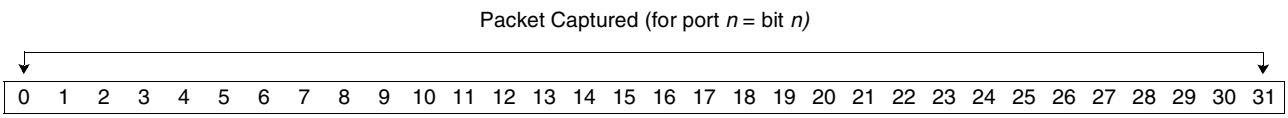

Bit(s)	Field Name	Description
0:31	Packet Captured (for port *n* = bit *n)*	When set to ‘1’, an ingress packet has been captured on the port.

5.6.15 Unilink Debug Control Register

The phase rotator embedded in the Unilink receiver architecture is designed to perform a critical clock-recovery function. This register provides indirect access to the Unilink debug facilities that monitor and control phase-rotator performance. The five phase-rotator control commands are issued using register bits 11:15. These bits must be kept inactive during normal operation.

Write access to the *Unilink Debug Control Register* requires one SHI command:

1. Write the register with the write bit set to ‘1’, the UL RxPort/UL RxSpex bus bit specified, the port number or speed-expansion bus link number field specified, and the phase rotor command specified.

Read access to the *Unilink Debug Control Register* requires two SHI commands:

1. Write the register with the write bit cleared to ‘0’, the UL RxPort/UL RxSpex bus bit specified, and the port number or speed-expansion bus link number field specified.
2. Read the register to return the debug control information.

Address x‘7D’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 uuuu uuuu uuuu uuuu’, where ‘u’ = undefined

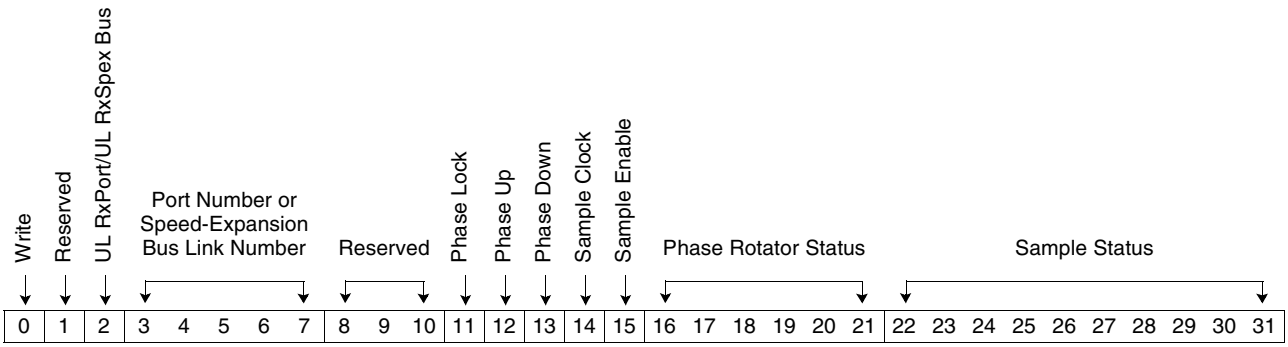

Bit(s)	Field Name	Description
0	Write	1 Specifies a write to the *Unilink Debug Control Register*. 0 Specifies a read from the *Unilink Debug Control Register*.
1	Reserved	Reserved.
2	UL RxPort/ UL RxSpex Bus	1 Specifies a Unilink receive port. 0 Specifies a Unilink receive speed-expansion bus.
3:7	Port Number or Speed-Expansion Bus Link Number	Specifies either the port number (0 to 31) or the speed-expansion bus link number (0 to 15), depending on the UL RxPort/UL RxSpex bus bit setting. Speed-expansion bus links 0 to 7 address ingress links and 8 to 15 address egress links.
8:10	Reserved	Reserved.
11	Phase Lock	When set to ‘1’, enables phase rotator external control.
12	Phase Up	When set to ‘1’, advances the sampling clock phase by one step.
13	Phase Down	When set to ‘1’, retards the sampling clock phase by one step.
14	Sample Clock	When set to ‘1’, enables the observation latch clock.

Bit(s)	Field Name	Description
15	Sample Enable	When set to ‘1’, enables the observation latch.
16:21	Phase Rotator Status (read only)	Provides the current state of the phase rotator. These bits are connected to the Unilink receive macro input pins PHSD[0:5]A, PHSD[0:5]B, PHSD[0:5]C, and PHSD[0:5]D—the phase rotator state buses for channels A, B, C, and D.
22:31	Sample Status (read only)	Unilink receive macro instantaneous sample output.

5.6.16 Unilink Force Error Register

This register is used to generate an error on Unilink receive internal logic.

Address x‘7E’

Access Type Read/Write

Reset Value ‘0000 0000 0000 0000 0000 0000 0000 0000’

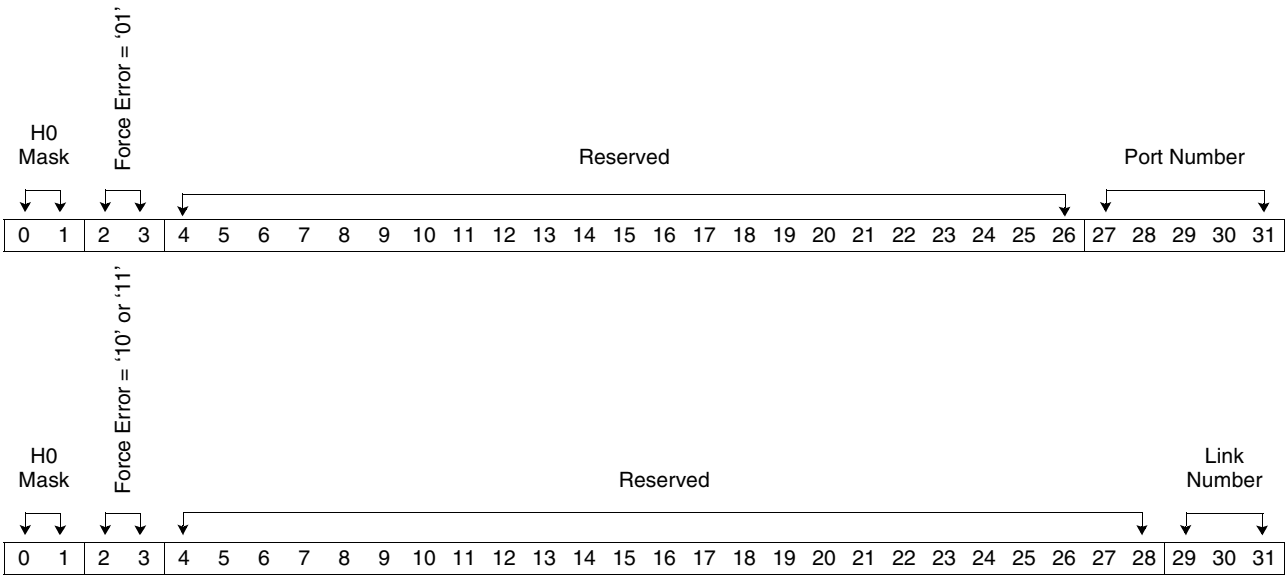

Bit(s)	Field Name	Description
0:1	H0 Mask	Reserved.
2:3	Force Error	Enables a forced error on the specified port or link: 00 Does not force an error. 01 Forces an error on the port specified by the port number field. 10 Forces an error on the ingress speed-expansion bus link specified by the link number field. 11 Forces an error on the egress speed-expansion bus link specified by the link number field.
4:26 *or* 4:28	Reserved	Reserved.
27:31	Port Number	If the force error field is set to ‘01’, specifies the port number.
29:31	Link Number	If the force error field is set to ‘10’ or ‘11’, specifies the link number.

6. Reset, Initialization, and Operation

6.1 Reset Sequence

This sequence must be executed after system power-up.

1. Activate the PowerOnResetIn# input pin and start the serial host interface (SHI) clock.
2. Deactivate the PowerOnResetIn# input pin after at least three SHI clock cycles.
3. Write the range A, range B, multiplier, and tune fields in the *Internal PLL Programming Register* (page 67) with the required values, and then release the phase-locked loop (PLL) reset bit.
4. Write the range A, range B, and multiplier fields in the *Unilink PLL Programming Register* (page 68) with the required values, and then release the PLL reset bit.
5. Wait 0.5 ms.
6. Verify that both PLLs are locked by checking the PLL locked bit settings in the *Internal PLL Status Register* (page 67) and the *Unilink PLL Status Register* (page 68).
7. Enable the Unilink macros by setting the *UL Global Register* (page 72) to x‘F000 0000’.
8. Release the Unilink macro flush bit in the *Reset Register* (page 69).
9. Wait 1 ms.
10. Verify that the Unilink macro internal PLLs are locked by checking the UL RxPorts PLL unlock, UL TxPorts PLL unlock, UL RxSpex bus PLL unlock, and UL TxSpex bus PLL unlock bit settings in the *UL Errors Register* (page 74).
11. Release the flush control and flush bits, and set the output driver enable bit in the *Reset Register*.
12. Wait 10 ms.
13. For slave devices, set the flush control bit in the *Reset Register*. This bit stops the clock on the slave device control logic to reduce power consumption.
14. Write the *Configuration 1 Register* (page 112) to specify part of the device configuration.
15. Set the shared memory, multicast, and output queue thresholds using the *Threshold Access Register* (page 114). (Thresholds can also be modified dynamically while the switch is active.)
16. Write the *Configuration 0 Register* (page 110) to specify the rest of the device configuration, and then release the standby bit.
17. Read the *Status Register* (page 120) to clear all interrupts.
18. If necessary, release the global interrupt mask bit in the *Reset Register.*
19. Set the Unilink clock deskew buffer offset to eight for ports and to six for speed-expansion buses by setting the *UL Global Register* to x‘F000 88C4’.

6.2 Speed-Expansion Bus Initialization

After the reset sequence is complete, the speed-expansion bus must be initialized. Speed-expansion bus initialization includes two steps:

1. Synchronize the speed-expansion bus links (see *Section 6.2.1* on page 156).
2. Set the speed-expansion bus data latency (see *Section 6.2.2* on page 156).

6.2.1 Synchronizing the Speed-Expansion Bus Links

To synchronize the speed-expansion bus links:

1. For each device except the last slave device, set the appropriate parameters in the *UL TxSpex Bus Parameters Register* (page 91). The parameter values depend on the signal characteristics.
2. For each device except the last slave device, set the *UL TxSpex Bus Attachment Enable Register* (page 90) to x‘FFFF 0000’.
3. For each device except the last slave device, set the *UL TxSpex Bus Driver Enable Register* (page 89) to x‘FFFF 0000’.
4. For each slave device, verify that the *UL RxSpex Bus Signal Lost Register* (page 102) reports x‘0000 0000’.
5. For the first slave device, set the *UL RxSpex Bus Attachment Enable Register* (page 96) to x‘FFFF 0000’, and wait 10 µs.
6. Repeat step 5 for each additional slave device in sequential order.
7. For each slave device, verify that both the *UL RxSpex Bus Byte Alignment Done Register* (page 97) and the *UL RxSpex Bus K28.5 Spacing OK Register* (page 98) report x‘FFFF 0000’.
8. For each slave device, set the *UL RxSpex Bus Data Mode Register* (page 101) to x‘FFFF 0000’.
9. For each slave device, read the *UL RxSpex Bus Code Violation Register* (page 105) to clear the errors generated during synchronization.
10. Wait 100 µs.
11. For each slave device, verify that the *UL RxSpex Bus Invalid K Character Register* (page 103), *UL RxSpex Bus Synchronization Lost Register* (page 104), and *UL RxSpex Bus Code Violation Register* (page 105) report x‘0000 0000’.

6.2.2 Setting the Speed-Expansion Bus Data Latency

6.2.2.1 The Source of Data Latency

For ingress packets, the master device provides store addresses to the first slave device, the first slave device propagates the store addresses to the next slave device, and so forth, until the last slave device receives a store address. Because the time required to propagate store addresses to the slave devices is greater than the time required to receive an ingress packet (that is, 64 ns for an eight-byte LU), slave device data path delay is required to guarantee that the input controller is still processing the data when the last slave device receives a store address from the master device (see *Figure 6-1* on page 157*).*

For egress packets, the master device provides retrieve addresses to the first slave device, the first slave device propagates the retrieve addresses to the next slave device, and so forth, until the last slave device receives a retrieve address. Because the time required to propagate retrieve addresses to the slave devices is greater than the time required to transmit an egress packet (that is, 64 ns for an eight-byte LU), a delay is required on the data path of the master device and all but the last slave device to guarantee that all the devices send the data over the Unilink ports at the same time (see *Figure 6-2* on page 157*).*

Figure 6-1. Ingress Speed-Expansion Bus Data Latency

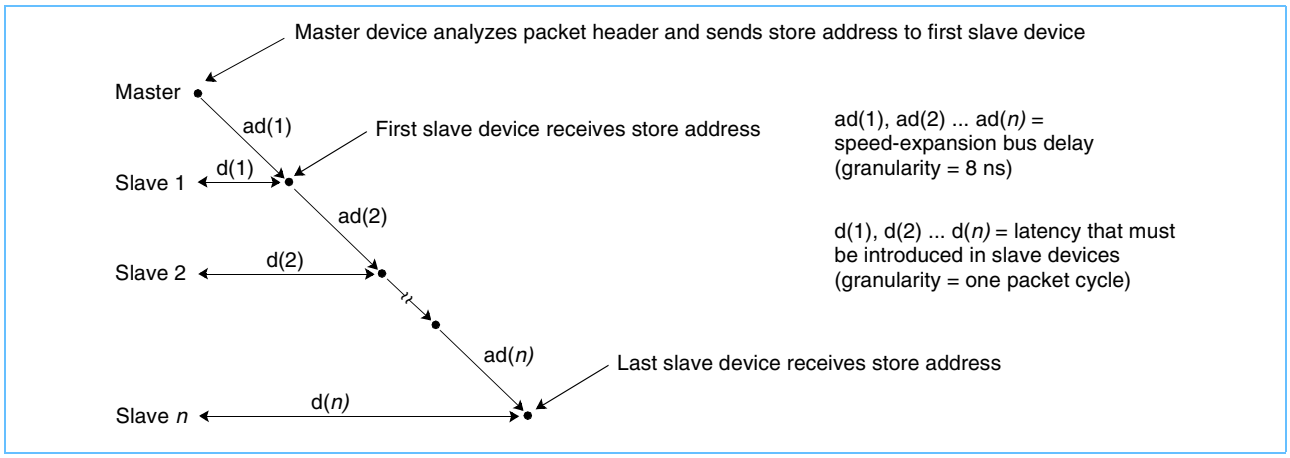

Figure 6-2. Egress Speed-Expansion Bus Data Latency

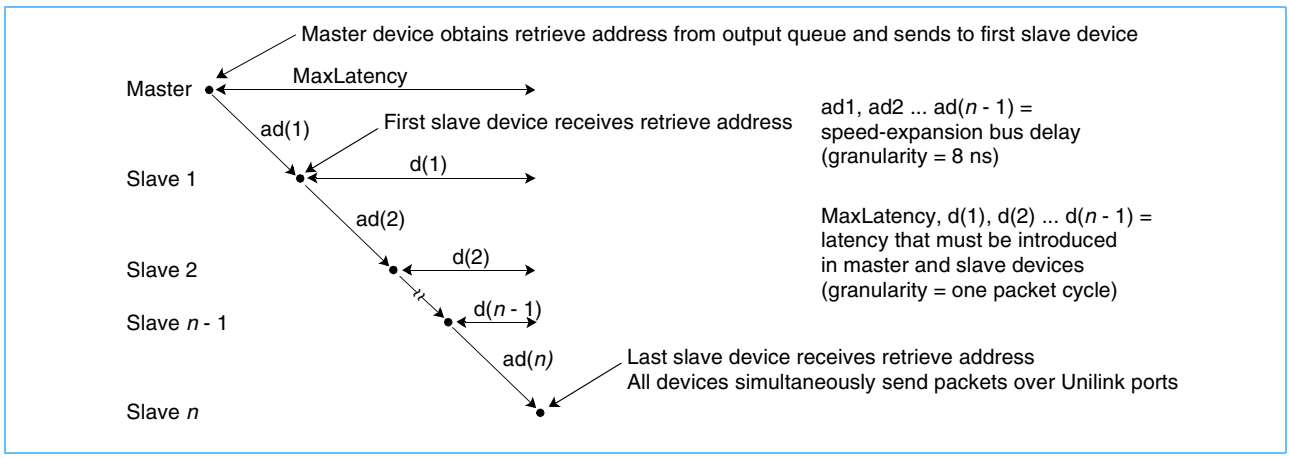

6.2.2.2 Setting Data Latency

The software sets the data latency using the *UL RxSpex Bus Latency Programming Register* (page 99).

The required data latency for all the master device links, when the links are processed independently, is:

- Zero on the ingress path
- MaxLatency (see *Figure 6-2)* on the egress path

The required data latency for the last slave device on the egress path is zero. The required data latency for the last slave device on the ingress path and for the other slave devices on both the ingress and egress paths depends on the speed-expansion bus delay (see *Table 6-1* on page 158*).*

Table 6-1. Required Data Latency

LU length	Ingress Path		Egress Path		
	Speed-Expansion Bus Delay (8-ns cycles)	Required Data Latency (packet cycles)	Speed-Expansion Bus Delay (8-ns cycles)	MaxLatency (packet cycles)	Required Data Latency (packet cycles)
8	0 to 10	1	0 to 10	1	MaxLatency - 1
	11 to 18	2	11 to 18	2	MaxLatency - 2
	19 to 26	3	19 to 26	3	MaxLatency - 3
	27 to 34	4	27 to 34	4	MaxLatency - 4
	35 to 42	5	35 to 42	5	MaxLatency - 5
	43 to 50	6	43 to 50	6	MaxLatency - 6
	> 50	Speed-expansion bus too slow (error)	> 50	Speed-expansion bus too slow (error)	
9	0 to 11	1	0 to 11	1	MaxLatency - 1
	12 to 20	2	12 to 20	2	MaxLatency - 2
	21 to 29	3	21 to 29	3	MaxLatency - 3
	30 to 38	4	30 to 38	4	MaxLatency - 4
	39 to 47	5	39 to 47	5	MaxLatency - 5
	48 to 56	6	48 to 56	6	MaxLatency - 6
	> 56	Speed-expansion bus too slow (error)	> 56	Speed-expansion bus too slow (error)	
10	0 to 12	1	0 to 12	1	MaxLatency - 1
	13 to 22	2	13 to 22	2	MaxLatency - 2
	23 to 32	3	23 to 32	3	MaxLatency - 3
	33 to 42	4	33 to 42	4	MaxLatency - 4
	43 to 52	5	43 to 52	5	MaxLatency - 5
	53 to 62	6	53 to 62	6	MaxLatency - 6
	> 62	Speed-expansion bus too slow (error)	> 62	Speed-expansion bus too slow (error)	

For each link on each of the slave paths, the software must use the *UL RxSpex Bus Latency Programming Register* (page 99) to:

1. Retrieve the speed-expansion bus delay value.
2. Set the data latency value.

When data latency programming is completed, the software sets the speed-expansion bus programming done bit in the *UL RxSpex Bus Latency Programming Register*. This bit must be set first in the master device, then in each slave device.

6.3 Port Initialization and Operation

After the PowerPRS Q-64G has been fully configured, the Unilink port interfaces between the PowerPRS Q-64G and the attached devices must be initialized to provide bit phase alignment and packet delineation. Data traffic cannot be exchanged between the PowerPRS Q-64G and the attached devices until synchronization is complete.

Port synchronization is controlled by the system control processor, which coordinates operations between the switch core and the attached devices. For a single PowerPRS Q-64G, port synchronization is handled by the local processor (which is connected to the PowerPRS Q-64G via the serial host interface). Port synchronization for an attached device is handled by the local processor of the attached device.

6.3.1 Initializing Unilink Ports

To initialize the receive and transmit Unilinks for port *n:*

1. For each PowerPRS Q-64G, set the appropriate Unilink parameters using the *UL TxPort Parameters Register* (page 79).
2. For each PowerPRS Q-64G, set bit *n* of the *UL TxPort Attachment Enable Register* (page 78) to ‘1’.
3. For each PowerPRS Q-64G, set bit *n* of the *UL TxPort Driver Enable Register* (page 78) to ‘1’. The transmit links are now ready.
4. For each PowerPRS Q-64G, poll the *UL RxPort Signal Lost Register* (page 85) until bit *n* equals ‘0’. When bit *n* equals ‘0’, the transmit links on the attached device are enabled and PowerPRS Q-64G receive link synchronization can begin.
5. For each PowerPRS Q-64G, set bit *n* of the *UL RxPort Attachment Enable Register* (page 81) to ‘1’.
6. For each PowerPRS Q-64G, poll the *UL RxPort Byte Alignment Done Register* (page 82) and the *UL RxPort K28.5 Spacing OK Register* (page 82) until both registers report a value of ‘1’ for port *n.* (A short delay [less than 1 ms] will occur while the software implements a timeout.) Bit phase alignment and packet delineation are now complete for each link, and PowerPRS Q-64G receive link LU deskew can begin.
7. For each PowerPRS Q-64G, read the “internal sequencer position on K reception” field of the *UL RxPort LU Deskew Register* (page 83) after the software inputs the PowerPRS Q-64G internal sequencer position from the idle packet K character received on the port. The internal sequencer position will be between 0 and 15 for an 8-byte LU, between 0 and 17 for a 9-byte LU, or between 0 and 19 for a 10-byte LU.

   Note: Because Unilink bit or byte phase synchronization and Unilink clock deskew can introduce multiple bytes of skew between port links, LU deskew must be performed to guarantee that port input controllers (master and slaves) receive data at the same time. LU deskew inserts a pipeline between the Unilink receive logic and the input controller. The pipeline depth is programmed using the LU deskew command field (default value equals zero), which can be set to up to eight stages (each stage is a 4-ns cycle). The pipeline depth must be programmed in each device so that each device reports the same value for the “internal sequencer position on K reception” field. The LU deskew algorithm is as follows:

   a. For each port link, KPOS(link) = internal sequencer position on K reception (link)
   b. Maximum KPOS value = KMAX
   c. Minimum KPOS value = KMIN
   d. KSKEW = KMAX - KMIN
   e. If KSKEW < 8, then go to step h
   f. For all KPOS < 7, KPOS = KPOS + (2 $\times$ LU length)

g. New maximum KPOS value = KMAX
h. For each link, LU deskew command(link) = KMAX - KPOS(link)

8. For each PowerPRS Q-64G, set bit *n* of the *UL RxPort Data Mode Register* (page 84) to ‘1’.
9. For each PowerPRS Q-64G, read the *UL RxPort Code Violation Register* (page 86) to clear any errors generated during synchronization.
10. Wait 10 µs.
11. For each PowerPRS Q-64G, verify that bit *n* of the *UL RxPort Data Valid Register* (page 84) is set to ‘1’.
12. For each PowerPRS Q-64G, verify that bit *n* is *not* set in the *UL RxPort Code Violation Register*, *UL RxPort Invalid K Character Register* (page 85), or *UL RxPort Synchronization Lost Register* (page 86).
13. For the master device, set bit *n* of the *Output Queue Enable Register* (page 124) to ‘1’.
14. For the master device, set bit *n* of the *Input Controller Enable Register* (page 124) to ‘1’. The port is ready to receive and transmit packets.

6.3.2 Deactivating Unilink Ports

To deactivate the Unilink ports for port *n:*

1. For the master device, set bit *n* of the *Input Controller Enable Register* to ‘0’.
2. For the master device, set bit *n* of the *Output Queue Enable Register* to ‘0’.
3. For each PowerPRS Q-64G, set bit *n* of the *UL RxPort Data Mode Register* to ‘0’.
4. For each PowerPRS Q-64G, set bit *n* of the *UL RxPort Attachment Enable Register* (page 81) to ‘0’.
5. For each PowerPRS Q-64G, set bit *n* of the *UL TxPort Driver Enable Register* (page 78) to ‘0’.
6. For each PowerPRS Q-64G, set bit *n* of the *UL TxPort Attachment Enable Register* (page 78) to ‘0’.
7. For each PowerPRS Q-64G, set the LU deskew command field of the *UL RxPort LU Deskew Register* (page 83) o ‘0’ for port *n.*

6.4 Logic BIST Execution Sequence

1. Activate the PowerOnResetIn# pin and start the SHI clock.
2. Deactivate the PowerOnResetIn# pin after at least three SHI clock cycles.
3. Write the range A, range B, multiplier, and tune fields in the *Internal PLL Programming Register* (page 67) with the required values, and then release the PLL reset bit.
4. Write the range A, range B, multiplier, and tune fields in the *Unilink PLL Programming Register* (page 68) with the required values, and then release the PLL reset bit.
5. Wait 0.5 ms.
6. Verify that both PLLs are locked by checking the PLL locked bit settings in the *Internal PLL Status Register* (page 67) and the *Unilink PLL Status Register* (page 68).
7. Enable the Unilink macros by setting the *UL Global Register* (page 72) to x‘F000 0000’.
8. Release the Unilink macro flush bit in the *Reset Register* (page 69).
9. Wait 1 ms.

10. Verify that the Unilink macro internal PLLs are locked by checking the UL RxPorts PLL unlock, UL TxPorts PLL unlock, UL RxSpex bus PLL unlock, and UL TxSpex bus PLL unlock bit settings in the *UL Errors Register* (page 74).
11. Set the logic BIST-requested bit and the Unilink macro flush bit in the *Reset Register* (page 69).
12. Set the PRPG/MISR data field in the *BIST Data Register* (page 70) with a defined value (see example below).
13. Load the *BIST Counter Register* (page 70) (see example below).
14. Set the shift speed to 8 ns and the scan chain length to 2048 in the *BIST Select Register* (page 71). Set the *BIST Select Register* to x‘0000 0800’.
15. Release the flush control and flush bits in the *Reset Register*.
16. Poll the *Reset Register* until the logic BIST done bit is set (see example below).
17. Read the MISR result in the PRPG/MISR data field of the *BIST Data Register* (see example below).
18. Set the flush and flush control bits, and release the logic BIST-requested bit in the *Reset Register*.

An example of logic BIST settings and results follows:

Step *12*. Preset the *BIST Data Register* to x‘007F FFF1’.
Step *13*. Load the *BIST Counter Register* to x‘0002 0000’ (131072 cycles).
Step *16*. Logic BIST duration is about 5 seconds.
Step *17*. The PRPG register values must be x‘0003 A200’.
The MISR0 register value must be x‘006B C885’.
The MISR1 register value must be x‘002B B388’.
The MISR2 register value must be x‘0052 38B3’.
The MISR3 register value must be x‘0026 FD0B’.
If any of the resulting *BIST Data Register* values do not match the values above, then the module is considered invalid.

6.5 Memory BIST Execution Sequence

1. Activate the PowerOnResetIn# pin and start the SHI clock.
2. Deactivate the PowerOnResetIn# pin after at least three SHI clock cycles.
3. Write the range A, range B, multiplier, and tune fields in the *Internal PLL Programming Register* (page 67) with the required values, and then release the PLL reset bit.
4. Write the range A, range B, multiplier, and tune fields in the *Unilink PLL Programming Register* (page 68) with the required values, and then release the PLL reset bit.
5. Wait 0.5 ms.
6. Verify that both PLLs are locked by checking the PLL locked bit settings in the *Internal PLL Status Register* (page 67) and the *Unilink PLL Status Register* (page 68).
7. Enable the Unilink macros by setting the *UL Global Register* (page 72) to x‘F000 0000’.
8. Release the unilink macro flush bit in the *Reset Register*.
9. Wait 1 ms.

10. Verify that the Unilink macro internal PLLs are locked by checking the UL RxPorts PLL unlock, UL TxPorts PLL unlock, UL RxSpex bus PLL unlock, and UL TxSpex bus PLL unlock bit settings in the *UL Errors Register*.
11. Set the memory BIST-requested bit in the *Reset Register* (page 69).
12. Release the flush control and flush bits in the *Reset Register*.
13. Poll the *Reset Register* until the memory BIST done bit is set.
14. Verify that the memory BIST fail bit in the *Reset Register* is off.
15. Set the flush and flush control bits, and release the memory BIST-requested bit in the *Reset Register*.

7. I/O Definitions and I/O Timing

7.1 I/O Definitions

Note: Nondifferential signals are active high unless there is a # symbol at the end of the signal name, in which case the signal is active low. Differential pairs are designated by an _P for the positive signal and an _N for the negative signal at the end of the signal name.

Table 7-1. Signal Definitions (Page 1 of 2)

Signal Name	Type	I/O Cell Name	Description
Reference Clocks and Master/Slave Synchronization Signals			
RefClockIn_P RefClockIn_N	Input Differential	ILVDST_A	System clock used for the internal clock generation network. RefClockIn frequency is 62.5 MHz with a duty cycle of 50 percent. The skew between the RefClockIn signals of all the PowerPRS Q-64Gs in a multiple-device configuration must be less than 1 ns.
UnilinkClockIn_P UnilinkClockIn_N	Input Differential	ILVDST_A	Unilink clock used for all Unilink internal macros. UnilinkClockIn frequency is 125 MHz with a duty cycle of 50 percent.
SyncIn_P SyncIn_N	Input Differential	ILVDS_A	Slave device input used to synchronize the slave device sequencer to the master device sequencer. This differential signal is connected to the master device SyncOut signal, but is not used by the master device. This signal has a 16-ns cycle. **Note:** The maximum wire length between a master device and a slave device is 1 meter.
SyncOut[0:6]_P SyncOut[0:6]_N	Output Differential	OLVDS18_A	Master device outputs used to synchronize the slave device sequencer to the master device sequencer. In the 512-Gbps configuration, each of the seven signals is connected to a slave device SyncIn signal. In the 256-Gbps configuration, only three of the seven signals are used; the other signals are disconnected. This bus is kept at high impedance for the slave devices, and has a 16-ns cycle.
Port Signals			
PortDataOut[0:31]_P PortDataOut[0:31]_N	Output Differential	Unilink	PortDataOut[*n*]_P and PortDataOut[*n*]_N form the 2.5-Gbps differential signal that connects a PowerPRS Q-64G port to the attached device.
PortDataIn[0:31]_P PortDataIn[0:31]_N	Input Differential	Unilink	PortDataIn[*n*]_P and PortDataIn[*n*]_N form the 2.5-Gbps differential signal that connects an attached device to the PowerPRS Q-64G.
Speed-Expansion Bus Signals			
SpexDataOut[0:15]_P SpexDataOut[0:15]_N	Output Differential	Unilink	SpexDataOut[*n*]_P and SpexDataOut[*n*]_N form the 2.5-Gbps differential signal that connects two PowerPRS Q-64Gs. The master device SpexDataOut bus is connected to the first slave device SpexDataIn bus, the first slave device SpexDataOut bus is connected to the second slave device SpexDataIn bus, and so forth, until the last slave device is connected. The total length of the speed expansion bus chain should not exceed 3 meters. The last slave device SpexDataOut bus is kept at high impedance. The master device SpexDataIn bus is disconnected.
SpexDataIn[0:15]_P SpexDataIn[0:15]_N	Input Differential	Unilink	SpexDataIn[*n*]_P and SpexDataIn[*n*]_N form the 2.5-Gbps differential signal that connects two PowerPRS Q-64Gs.
Serial Host Interface Signals			
PowerOnResetIn#	Input	BC1850_A	Asserting this pin forces the two internal PLLs to reset and keeps the internal logic in a flush state. This signal is internally processed as an asynchronous signal. It must be active for at least five SHI clock cycles.
SHIClockIn	Input	BC1850_C	Free-running clock that generates the SHI clock.

Table 7-1. Signal Definitions (Page 2 of 2)

Signal Name	Type	I/O Cell Name	Description
InterruptOut#	Output	BC1850_A	Used to generate interrupts to the local processor. The InterruptOut# signal remains asserted until the *Status Register* (page 120) is no longer empty. To support a wired-OR configuration, InterruptOut# uses an open-drain driver and is in the high-impedance state when inactive. **Note:** InterruptOut# is asserted only if the interrupt is not masked with either the corresponding bit in the *Interrupt Mask Register* (page 122) or the global interrupt mask bit in the *Reset Register* (page 69).
SHISerialDataIn	Input	BC1850_C	Serial data line that shifts into the *SHI Instruction Register* (page 61).
SHISerialDataOut	Output	BC1850_C	Serial data line that shifts out of the *SHI Instruction Register*. SHISerial-DataOut is placed in a high-impedance state when the SHI is not in shift state. The SHI is in shift state one SHI clock cycle after SHISelectIn# becomes inactive.
SHISelectIn#	Input	BC1850_C	Enables SHI operation. One SHI clock cycle after the SHISelect-In# signal becomes active, the instruction is serially shifted into the *SHI Instruction Register*.
Other Signals			
SCCIn[0:3]	Input	BC1850_A	The side communication channel (SCC) signal line that allows communication between the attached devices and the PowerPRS Q-64G. This bus is processed asynchronously to the internal clock and only by the master device. An external pullup or pulldown is required for the unused bits.
DebugBusOut[0:15]	Output	BC1850_C	Sixteen-bit bus that provides direct I/O access (logic analyzer) to the debug bus specified by the *Debug Bus Select Register* (page 140). This bus is enabled by the debug bus enable bit of the *Debug Bus Select Register*.
FullyInsertedIn#	Input	BC1850_A	Used to force the Unilink port drivers to high impedance until the board housing the PowerPRS Q-64G is fully inserted. This ensures that both ends of the board are fully inserted. An external pullup resistor is required to force the inactive state when the board is correctly inserted.
DelayIn	Input	BC1850_A	Internal delay element input used for process measurement. The internal delay element is built with a chain of 300 INVERT_O gates.
DelayOut	Output	BC1850_A	Internal delay element output used for process measurement: • 3.81 ns minimum (process = -3 sigma, temperature = 0°C, and voltage = 1.95 V) • 7.39 ns maximum (process = +3 sigma, temperature = 125°C, voltage = 1.65 V) • 4.25 ns minimum (process = -3 sigma, temperature = 50°C, and voltage = 1.8 V) • 6.4 ns maximum (process = +3 sigma, temperature = 50°C, voltage = 1.8 V)
Osc625TestPointOut_P Osc625TestPointOut_N	Output Differential	OLVDS18_A	Driver connected to the end of the 625-MHz Unilink clock tree. This driver is enabled by the debug bus enable bit of the *Debug Bus Select Register* (page 140).

Table 7-2. Power Signals (Page 1 of 2)

Signal Name	Description
CorePLLVddaIn	Internal logic PLL analog V_{DD}.
CorePLLGndaIn	Internal logic PLL analog ground.
UnilinkPLLVddaIn	Unilink PLL analog V_{DD}.
UnilinkPLLGndaIn	Unilink PLL analog ground.
rulp_av25In[0:7]	Unilink port receive macro 2.5-V analog voltage supply.

Table 7-2. Power Signals (Page 2 of 2)

Signal Name	Description
rulp_avttIn[0:7]	Unilink port receive macro 1.8-V analog voltage supply.
rulp_avregOut[0:7]	Unilink port receive macro observation point for the internal voltage regulator.
xulp_av25In[0:7]	Unilink port transmit macro 2.5-V analog voltage supply.
xulp_avttIn[0:7]	Unilink port transmit macro 1.8-V analog voltage supply.
xulp_avregOut[0:7]	Unilink port transmit macro observation point for the internal voltage regulator.
ruls_av25In[0:3]	Unilink speed expansion bus receive macro 2.5-V analog voltage supply.
ruls_avttIn[0:3]	Unilink speed expansion bus receive macro 1.8-V analog voltage supply.
ruls_avregOut[0:3]	Unilink speed expansion bus receive macro observation point for the internal voltage regulator.
xuls_av25In[0:3]	Unilink speed expansion bus transmit macro 2.5-V analog voltage supply.
xuls_avttIn[0:3]	Unilink speed expansion bus transmit macro 1.8-V analog voltage supply.
xuls_avregOut[0:3]	Unilink speed expansion bus transmit macro observation point for the internal voltage regulator.

Table 7-3. Test Signals

Signal Name	I/O Cell Name	Description
Lssd_A_ClkIn	IC18PDT_A	Used as an external source for the internal set/reset latch (SRL) A clock.
Lssd_B_ClkIn	IC18PUT_A	Used as an external source for the internal SRL B clock.
Lssd_C1_ClkIn	IC18PUT_A	Used as an external source for the internal SRL C clock.
Lssd_C2_ClkIn	IC18PDT_A	Used as an external source for the internal static random access memory (SRAM)-4 clock.
Lssd_C3_ClkIn	IC18PDT_A	Used as an external source for the internal global register array (GRA) C clock.
TestM3In	IC18PUT_A	Used to handle the internal memory built-in self-test (BIST) controllers.
LssdTestEnableIn	IC18TEPDT_A	Allows all clocks to be controlled from the primary inputs, and connects all scan chains.
IOTestIn	IC18PDT_A	Used for reduced pin count testing.
ScanIn[0:23]	IC18PUT_A	Scan chain inputs. Some ScanIn I/Os are shared (JTAG pins): • ScanIn(0) = TDI • ScanIn(1) = TMS • ScanIn(2) = TCK
ScanOut[0:23]	BP1850T_A	Scan chain outputs. ScanOut(6) is shared with the JTAG TDO pin.
LeakageTestIn	IC18LTPUT_A	Used during the leakage test.
ThermalIn/ThermalOut	THERMAL_A	Used to measure the device junction temperature. A THERMAL_A cell is connected between these two I/Os (see the *ASIC SA-27E Databook* under *Related Documents* [page 197] for details).
UnilinkLtestIn	IC18PDT_A	Connected to the Unilink macros LTEST (logic test mode) pin (normal operation = 0, logic test = 1).
UnilinkMtestIn	IC18PDT_A	Connected to the Unilink macro MTEST (macro test mode) pin (normal operation = 0, macro test = 1).
DI1#	IC18D1PUT_A	Inhibits the driver for all nontest outputs. Active low.
DI2#	IC18D2PUT_A	Inhibits the driver for all test outputs. Active low. This pin is also used as the JTAG TRST (test reset) pin and requires an external pulldown.
RI#	IC18RIT_A	Inhibits the receiver for all inputs. An external pullup resistor to 1.8 V is required on this pin. Active low.

7.2 I/O Timing

7.2.1 Unilink Signals

Unilink-signal skew requirements are presented in *Table 7-4.*

Table 7-4. Unilink Interface Skew

Parameter	Rating	Units
Maximum skew between the two lines of a differential pair.	±20	ps
Maximum skew between any 2.5-Gbps links on the same port.	4	ns

7.2.2 SHI Signals

Figure 7-1. SHI Signal Timing Diagram

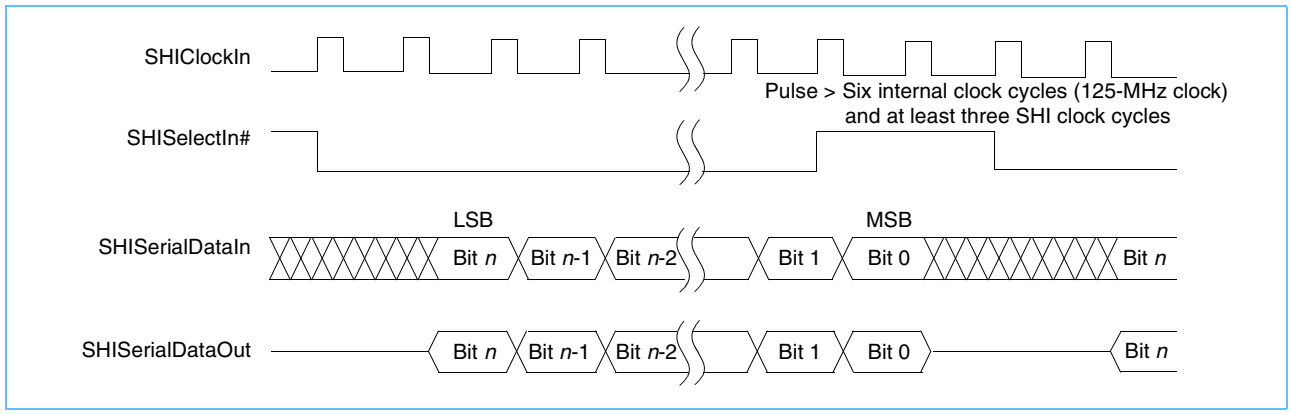

Figure 7-2. SHI Signal-to-Clock Timing Diagram

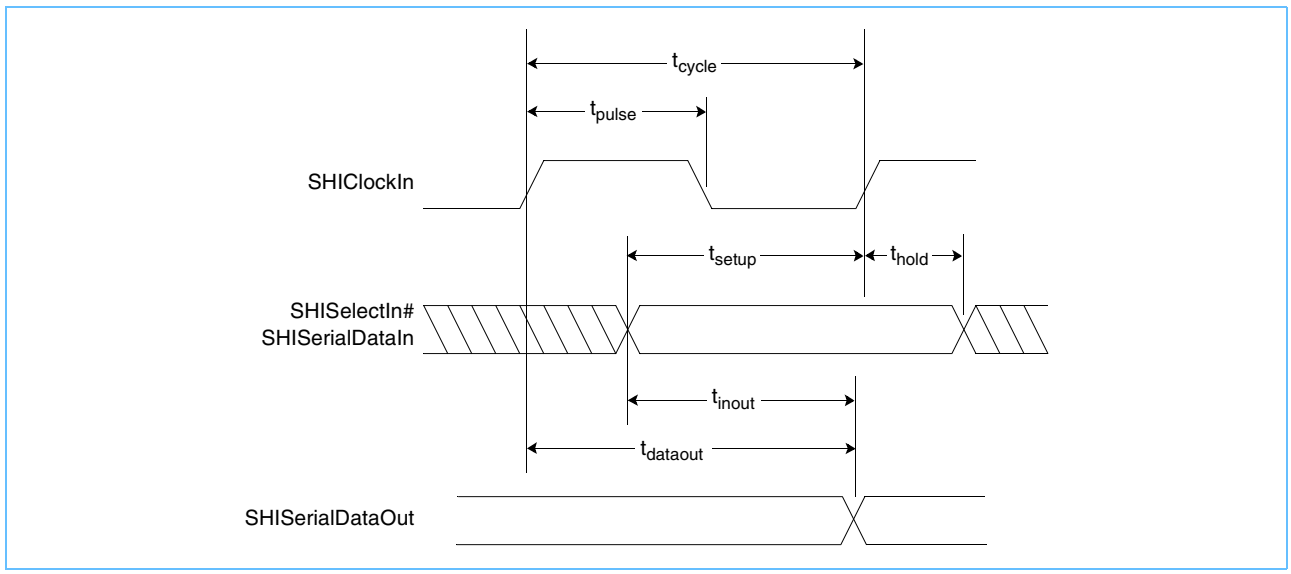

Table 7-5. SHI Signal Timing Values

Symbol	Parameter	Rating		Units
		Minimum	Maximum	
t_{cycle}	Cycle time	16		ns
t_{pulse}	Pulse width	4		ns
t_{setup}	Setup time	4		ns
t_{hold}	Hold time	6		ns
t_{inout}	SHISelectIn# to SHISerialDataOut	2	6	ns
$t_{dataout}$	SHIClockIn to SHISerialDataOut	4	9	ns

7.2.3 SyncIn/SyncOut Signals

Figure 7-3. SyncOut/SyncIn Signal Timing Diagram

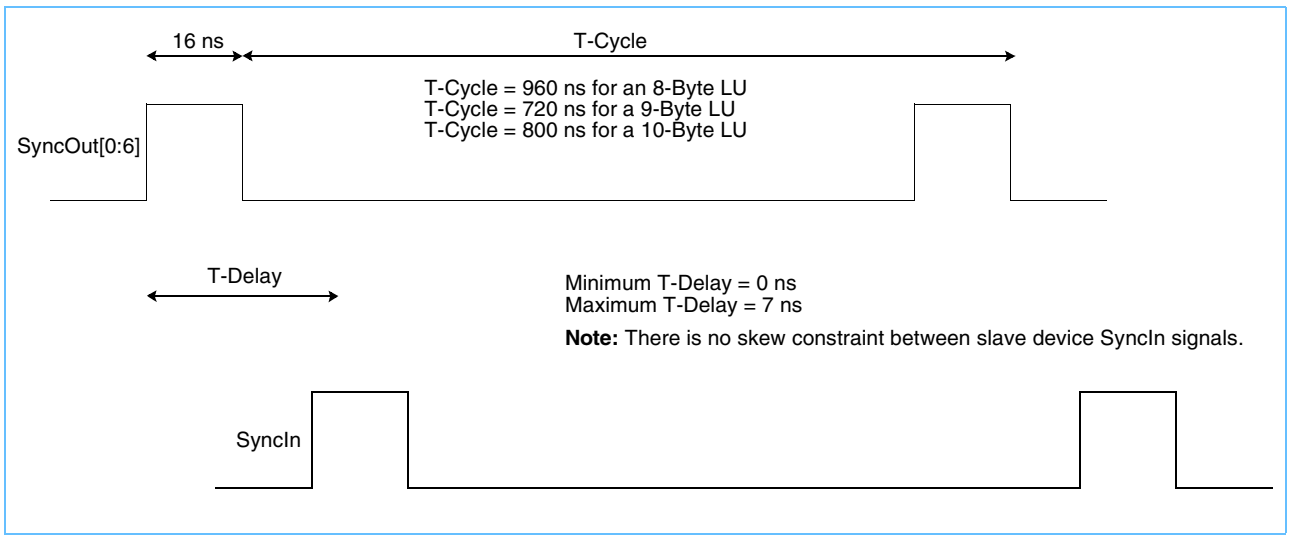

8. Data and Flow Control Latencies

This section discusses PowerPRS Q-64G data packet transmission time, including the time required for grant status changes to take effect after the status-changing events occur. Reported times were calculated from port-byte interface to port-byte interface and, consequently, time must be added to the values reported in the tables below to account for:

- Transmission over the Unilink. The required Unilink transmission time is 200 ns (minimum) to 280 ns (maximum); 100 to 140 ns are necessary for 8b/10b code transmission, clock deskew transmission, serialization, driver propagation, receiver propagation, deserialization, clock deskew reception, 8b/10b code reception, and LU deskew. This time must be calculated twice, once for the Unilink on the ingress path (attached device to PowerPRS Q-64G) and once for the Unilink on the egress path (PowerPRS Q-64G to attached device).
- Propagation through the card(s) and backplane. The required time depends on the application.

8.1 Data Packet Transmission

Table 8-1 shows the difference between the time an input controller receives the first byte of a packet and the time the output controller sends the first byte of the same packet, assuming the packet is not enqueued for flow control. *Table 8-1* does not account for the egress path delay caused by egress speed-expansion bus latency. To account for this delay, add three LU cycles for the 256-Gbps device configuration or six LU cycles for the 512-Gbps device configuration to the minimum and maximum values reported in the table.

Table 8-1. Data Packet Transmission

LU Size (bytes)	Time (LU cycles)		Time (nanoseconds)	
	Minimum	Maximum	Minimum	Maximum
8	4.50	6.88	288	440
9	4.22	6.56	304	472
10	4.00	6.30	320	504

Note: Reported times apply to the PowerPRS Q-64G configurations that feature ports speeds of 16 Gbps; reported times for other port speeds will be different.

8.2 Send Grant Off to Egress Idle Packet

Table 8-2 shows the difference between the time an attached device removes the send grant and the time the first byte of the first idle packet exits the PowerPRS Q-64G due to this flow control mechanism.

Table 8-2. Send Grant Off to Egress Idle Packet

LU Size (bytes)	Time (LU cycles)		Time (nanoseconds)	
	Minimum	Maximum	Minimum	Maximum
8	2.88	3.75	184	240
9	2.56	3.44	184	248
10	2.30	3.20	184	256

Note: Reported times apply to the PowerPRS Q-64G configurations that feature ports speeds of 16 Gbps; reported times for other port speeds will be different.

The values reported in *Table 8-2* are those for the PowerPRS Q-64G with only one enabled priority. As discussed in *Section 3.3.2 Flow Control Flywheels for Grants Carried in Ingress Packets* on page 29, there is one send grant per priority, and only one send grant is transmitted per ingress data packet or control packet cycle. The grant priority flywheel determines the priority for which a send grant is transmitted during a packet cycle and can advance to the priority of a send grant being removed. *Table 8-3* shows the additional latency associated with the longest possible flywheel delay for each number of enabled priorities. To account for grant priority flywheel cycling, the appropriate value in *Table 8-3* should be added to the maximum values in *Table 8-2*.

Table 8-3. Grant Priority Flywheel Cycling

Number of Enabled Priorities	Maximum Latency Addition (LU cycles) (add to maximum values in *Table 8-2)*
1	0
2	1
3	2
4	3

8.3 Ingress Data Packet Received to Output Queue Grant Off

Table 8-4 shows the difference between the time an input controller receives the first byte of an ingress data packet that causes the output queue occupancy to exceed the output queue threshold (and requires the output queue grant to be turned off) and the time the attached device receives the first byte of the packet that contains the updated output queue grant status (which turns off the output queue grant).

Table 8-4. Ingress Data Packet Received to Output Queue Grant Off

LU Size	Time (LU cycles)		Time (nanoseconds)	
	Minimum	Maximum	Minimum	Maximum
8	4.00	6.63	256	424
9	3.67	6.44	264	464
10	3.20	6.20	256	496

Note: Reported times apply to the PowerPRS Q-64G configurations that feature ports speeds of 16 Gbps; reported times for other port speeds will be different.

The values reported in *Table 8-4* are those for the PowerPRS Q-64G with only one enabled priority. As discussed in *Section 3.3.6 Flow Control Flywheels for Grants Carried in Egress Packets* on page 36, two packet cycles are required to convey the output queue grants for each enabled priority. The output queue grant priority flywheel determines the priority for which output queue grants are transmitted during a packet cycle and can advance to the priority of the output queue grant being removed. *Table 8-5* shows the additional latency associated with the longest possible flywheel delay for each number of enabled priorities. To account for output queue grant priority flywheel cycling, the appropriate value in *Table 8-5* should be added to the maximum values in *Table 8-4.*

Table 8-5. Output Queue Grant Priority Flywheel Cycling

Number of Enabled Priorities	Maximum Latency Addition (LU cycles) (add to maximum values in *Table 8-4)*
1	0
2	2
3	4
4	6

8.4 Ingress Data Packet Received to Memory Grant Off

Table 8-6 shows the difference between the time an input controller receives the first byte of an ingress data packet that causes the shared memory occupancy to exceed the shared memory threshold (and requires the memory grant to be turned off) and the time the attached device receives the first byte of the packet that contains the updated memory grant status (which turns off the memory grant).

Table 8-6. Ingress Data Packet Received to Memory Grant Off

LU Size	Time (LU cycles)		Time (nanoseconds)	
	Minimum	Maximum	Minimum	Maximum
8	2.88	5.75	184	368
9	2.67	5.44	192	392
10	2.30	5.20	184	416

Note: Reported times apply to the PowerPRS Q-64G configurations that feature ports speeds of 16 Gbps; reported times for other port speeds will be different.

8.5 Ingress Data Packet Received to Multicast Grant Off

Table 8-7 shows the difference between the time an input controller receives the first byte of an ingress data packet that causes the multicast packet count to exceed the multicast high threshold (and requires the multicast grant to be turned off) and the time the attached device receives the first byte of a packet that contains the updated multicast grant status (which turns off the multicast grant).

Table 8-7. Ingress Data Packet Received to Multicast Grant Off

LU Size	Time (LU cycles)		Time (nanoseconds)	
	Minimum	Maximum	Minimum	Maximum
8	4.75	8.50	304	544
9	3.78	6.78	272	488
10	3.70	6.50	296	520

Note: Reported times apply to the PowerPRS Q-64G configurations that feature ports speeds of 16 Gbps; reported times for other port speeds will be different.

9. Pin Information

Figure 9-1. Pinout (624-ball HyperBGA package, bottom view)

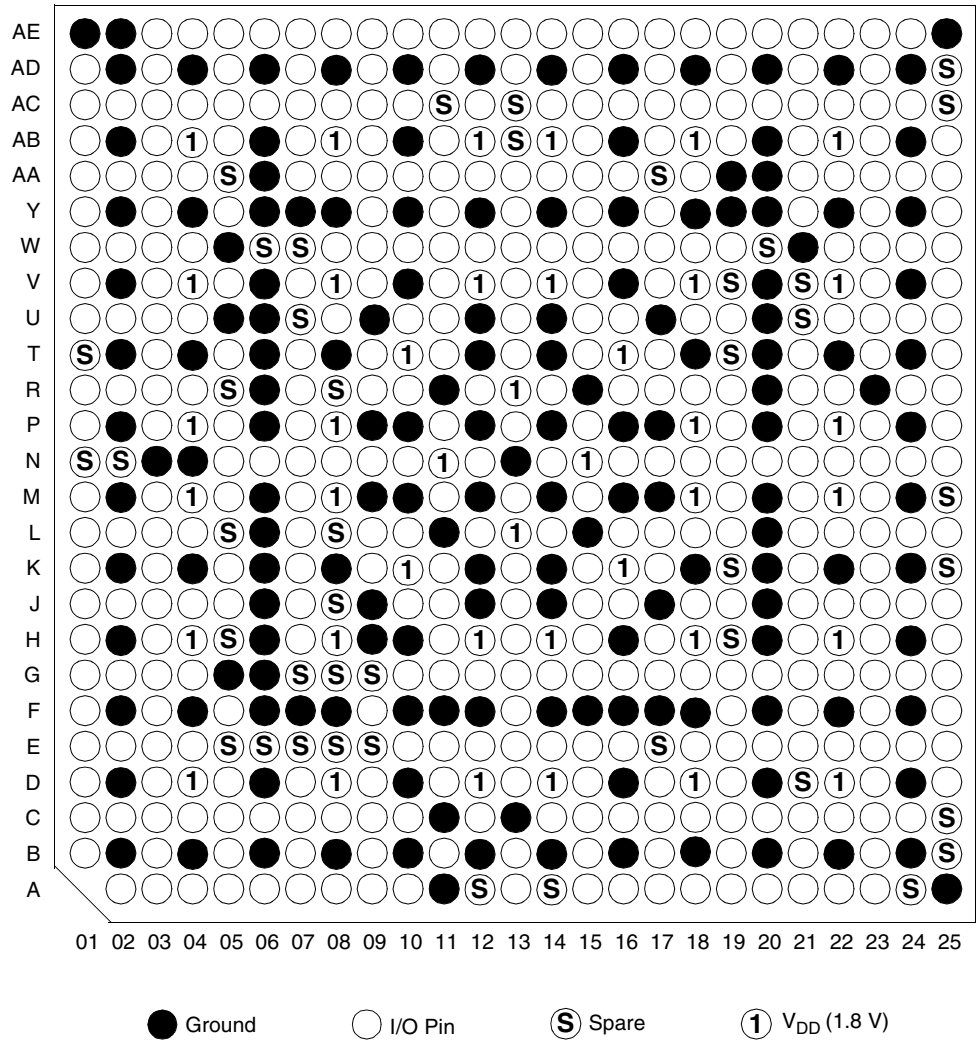

Note: Spare pins must be connected to a ground. However, spare pins should *not* be connected directly to the card's ground plane because they may be needed in a future PowerPRS Q-64G release.

Table 9-1. Ground, V_{DD} and Spare Pin Locations

Pin Function	Pin Locations
Ground	A11, A25, B02, B04, B06, B08, B10, B12, B14, B16, B18, B20, B22, B24, C11, C13, D02, D06, D10, D16, D20, D24, F02, F04, F06, F07, F08, F10, F11, F12, F14, F15, F16, F17, F18, F20, F22, F24, G05, G06, H02, H06, H09, H10, H16, H20, H24, J06, J09, J12, J14, J17, J20, K02, K04, K06, K08, K12, K14, K18, K20, K22, K24, L06, L11, L15, L20, M02, M06, M09, M10, M12, M14, M16, M17, M20, M24, N03, N04, N13, P02, P06, P09, P10, P12, P14, P16, P17, P20, P24, R06, R11, R15, R20, R23, T02, T04, T06, T08, T12, T14, T18, T20, T22, T24, U05, U06, U09, U12, U14, U17, U20, V02, V06, V10, V16, V20, V24, W05, W21, Y02, Y04, Y06, Y07, Y08, Y10, Y12, Y14, Y16, Y18, Y19, Y20, Y22, Y24, AA06, AA19, AA20, AB02, AB06, AB10, AB16, AB20, AB24, AD02, AD04, AD06, AD08, AD10, AD12, AD14, AD16, AD18, AD20, AD22, AD24, AE01, AE02, AE25
V_{DD}	D04, D08, D12, D14, D18, D22, H04, H08, H12, H14, H18, H22, K10, K16, L13, M04, M08, M18, M22, N11, N15, P04, P08, P18, P22, R13, T10, T16, V04, V08, V12, V14, V18, V22, AB04, AB08, AB12, AB14, AB18, AB22
Spare	A12, A14, A24, B25, C25, D21, E05, E06, E07, E08, E09, E17, G07, G08, G09, H05, H19, J08, K19, K25, L05, L08, M25, N01, N02, R05, R08, T01, T19, U07, U21, V19, V21, W06, W07, W20, AA05, AA17, AB13, AC11, AC13, AC25, AD25
Note: Spare pins must be connected to a ground. However, spare pins should *not* be connected directly to the card's ground plane because they may be needed in a future PowerPRS Q-64G release.	

Table 9-2. I/O Signal List, Sorted by Signal Name (Page 1 of 5)

Signal Name	Pin Location	Signal Name	Pin Location	Signal Name	Pin Location
CorePLLGndaIn	AD01	Lssd_C3_ClkIn	A03	PortDataIn[14]_N	A22
CorePLLVddaIn	AC01	Osc625TestPointOut_N	P15	PortDataIn[14]_P	B23
DebugBusOut[0]	V09	Osc625TestPointOut_P	R14	PortDataIn[15]_N	C22
DebugBusOut[1]	W09	PortDataIn[0]_N	A04	PortDataIn[15]_P	C23
DebugBusOut[2]	AA08	PortDataIn[0]_P	B03	PortDataIn[16]_N	AE22
DebugBusOut[3]	Y11	PortDataIn[1]_N	C03	PortDataIn[16]_P	AD23
DebugBusOut[4]	W08	PortDataIn[1]_P	C04	PortDataIn[17]_N	AC23
DebugBusOut[5]	Y09	PortDataIn[2]_N	C05	PortDataIn[17]_P	AC22
DebugBusOut[6]	W10	PortDataIn[2]_P	B05	PortDataIn[18]_N	AC21
DebugBusOut[7]	W11	PortDataIn[3]_N	C06	PortDataIn[18]_P	AD21
DebugBusOut[8]	W18	PortDataIn[3]_P	C07	PortDataIn[19]_N	AC20
DebugBusOut[9]	W16	PortDataIn[4]_N	B07	PortDataIn[19]_P	AC19
DebugBusOut[10]	AA18	PortDataIn[4]_P	A06	PortDataIn[20]_N	AE20
DebugBusOut[11]	Y15	PortDataIn[5]_N	A08	PortDataIn[20]_P	AD19
DebugBusOut[12]	V17	PortDataIn[5]_P	A07	PortDataIn[21]_N	AE18
DebugBusOut[13]	W17	PortDataIn[6]_N	C08	PortDataIn[21]_P	AE19
DebugBusOut[14]	Y17	PortDataIn[6]_P	B09	PortDataIn[22]_N	AC18
DebugBusOut[15]	W15	PortDataIn[7]_N	D09	PortDataIn[22]_P	AD17
DelayIn	A02	PortDataIn[7]_P	C09	PortDataIn[23]_N	AB17
DelayOut	AE24	PortDataIn[8]_N	D17	PortDataIn[23]_P	AC17
DI1#	A23	PortDataIn[8]_P	C17	PortDataIn[24]_N	AB09
DI2#	E10	PortDataIn[9]_N	C18	PortDataIn[24]_P	AC09
FullyInsertedIn#	AE13	PortDataIn[9]_P	B17	PortDataIn[25]_N	AC08
InterruptOut#	AD13	PortDataIn[10]_N	A18	PortDataIn[25]_P	AD09
IOTestIn	AE03	PortDataIn[10]_P	A19	PortDataIn[26]_N	AE08
LeakageTestIn	AE11	PortDataIn[11]_N	B19	PortDataIn[26]_P	AE07
LssdTestEnableIn	AE15	PortDataIn[11]_P	A20	PortDataIn[27]_N	AD07
Lssd_A_ClkIn	AA23	PortDataIn[12]_N	C20	PortDataIn[27]_P	AE06
Lssd_B_ClkIn	AA25	PortDataIn[12]_P	C19	PortDataIn[28]_N	AC06
Lssd_C1_ClkIn	AE21	PortDataIn[13]_N	C21	PortDataIn[28]_P	AC07
Lssd_C2_ClkIn	AE23	PortDataIn[13]_P	B21	PortDataIn[29]_N	AC05

Table 9-2. I/O Signal List, Sorted by Signal Name (Page 2 of 5)

Signal Name	Pin Location	Signal Name	Pin Location	Signal Name	Pin Location
PortDataIn[29]_P	AD05	PortDataOut[13]_N	F05	PortDataOut[28]_P	Y25
PortDataIn[30]_N	AE04	PortDataOut[13]_P	E04	PortDataOut[29]_N	AA22
PortDataIn[30]_P	AD03	PortDataOut[14]_N	D03	PortDataOut[29]_P	Y21
PortDataIn[31]_N	AC04	PortDataOut[14]_P	E02	PortDataOut[30]_N	AB23
PortDataIn[31]_P	AC03	PortDataOut[15]_N	C02	PortDataOut[30]_P	AA24
PortDataOut[0]_N	AB01	PortDataOut[15]_P	D01	PortDataOut[31]_N	AC24
PortDataOut[0]_P	AC02	PortDataOut[16]_N	D25	PortDataOut[31]_P	AB25
PortDataOut[1]_N	AA02	PortDataOut[16]_P	C24	PowerOnResetIn#	R12
PortDataOut[1]_P	AB03	PortDataOut[17]_N	E24	RefClockIn_N	R04
PortDataOut[2]_N	AA04	PortDataOut[17]_P	D23	RefClockIn_P	T03
PortDataOut[2]_P	Y03	PortDataOut[18]_N	E22	RI#	E23
PortDataOut[3]_N	W01	PortDataOut[18]_P	F23	rulp_av25In[0]	A09
PortDataOut[3]_P	Y01	PortDataOut[19]_N	F25	rulp_av25In[1]	A10
PortDataOut[4]_N	W03	PortDataOut[19]_P	G25	rulp_av25In[2]	A15
PortDataOut[4]_P	W04	PortDataOut[20]_N	G23	rulp_av25In[3]	A17
PortDataOut[5]_N	V03	PortDataOut[20]_P	G24	rulp_av25In[4]	AE17
PortDataOut[5]_P	W02	PortDataOut[21]_N	G22	rulp_av25In[5]	AE14
PortDataOut[6]_N	V01	PortDataOut[21]_P	H23	rulp_av25In[6]	AE12
PortDataOut[6]_P	U02	PortDataOut[22]_N	H25	rulp_av25In[7]	AE09
PortDataOut[7]_N	U04	PortDataOut[22]_P	J24	rulp_avregOut[0]	D05
PortDataOut[7]_P	U03	PortDataOut[23]_N	J23	rulp_avregOut[1]	E11
PortDataOut[8]_N	J04	PortDataOut[23]_P	J22	rulp_avregOut[2]	A16
PortDataOut[8]_P	J03	PortDataOut[24]_N	U22	rulp_avregOut[3]	E21
PortDataOut[9]_N	J02	PortDataOut[24]_P	U23	rulp_avregOut[4]	AA21
PortDataOut[9]_P	H01	PortDataOut[25]_N	U24	rulp_avregOut[5]	AA15
PortDataOut[10]_N	H03	PortDataOut[25]_P	V25	rulp_avregOut[6]	AA09
PortDataOut[10]_P	G03	PortDataOut[26]_N	W23	rulp_avregOut[7]	AA07
PortDataOut[11]_N	F03	PortDataOut[26]_P	V23	rulp_avttIn[0]	D07
PortDataOut[11]_P	G04	PortDataOut[27]_N	Y23	rulp_avttIn[1]	D13
PortDataOut[12]_N	G01	PortDataOut[27]_P	W22	rulp_avttIn[2]	C15
PortDataOut[12]_P	F01	PortDataOut[28]_N	W25	rulp_avttIn[3]	D19

Table 9-2. I/O Signal List, Sorted by Signal Name (Page 3 of 5)

Signal Name	Pin Location	Signal Name	Pin Location	Signal Name	Pin Location
rulp_avttIn[4]	AB21	ScanIn[15]	M21	ScanOut[22]	K05
rulp_avttIn[5]	AC15	ScanIn[16]	L22	ScanOut[23]	E03
rulp_avttIn[6]	AA11	ScanIn[17]	K21	SCCIn[0]	F09
rulp_avttIn[7]	AB05	ScanIn[18]	T23	SCCIn[1]	A13
ruls_av25In[0]	M07	ScanIn[19]	R24	SCCIn[2]	B13
ruls_av25In[1]	G12	ScanIn[20]	P23	SCCIn[3]	E15
ruls_av25In[2]	G14	ScanIn[21]	M23	SHIClockIn	AB07
ruls_av25In[3]	M19	ScanIn[22]	L24	SHISelectIn#	AE10
ruls_avregOut[0]	L07	ScanIn[23]	K23	SHISerialDataIn	AE16
ruls_avregOut[1]	J07	ScanOut[0]	AA10	SHISerialDataOut	AB19
ruls_avregOut[2]	J18	ScanOut[1]	AB11	SpexDataIn[0]_N	K09
ruls_avregOut[3]	L19	ScanOut[2]	AA12	SpexDataIn[0]_P	L09
ruls_avttIn[0]	K07	ScanOut[3]	AA14	SpexDataIn[1]_N	N07
ruls_avttIn[1]	H07	ScanOut[4]	AB15	SpexDataIn[1]_P	N08
ruls_avttIn[2]	H17	ScanOut[5]	AA16	SpexDataIn[2]_N	L10
ruls_avttIn[3]	J19	ScanOut[6]	AC10	SpexDataIn[2]_P	M11
ScanIn[0]	C10	ScanOut[7]	AD11	SpexDataIn[3]_N	N12
ScanIn[1]	B11	ScanOut[8]	AC12	SpexDataIn[3]_P	M13
ScanIn[2]	C12	ScanOut[9]	AC14	SpexDataIn[4]_N	G10
ScanIn[3]	C14	ScanOut[10]	AD15	SpexDataIn[4]_P	G11
ScanIn[4]	B15	ScanOut[11]	AC16	SpexDataIn[5]_N	J13
ScanIn[5]	C16	ScanOut[12]	AA01	SpexDataIn[5]_P	K13
ScanIn[6]	D11	ScanOut[13]	R02	SpexDataIn[6]_N	K11
ScanIn[7]	E12	ScanOut[14]	P03	SpexDataIn[6]_P	J10
ScanIn[8]	E14	ScanOut[15]	M03	SpexDataIn[7]_N	H11
ScanIn[9]	D15	ScanOut[16]	L02	SpexDataIn[7]_P	J11
ScanIn[10]	E16	ScanOut[17]	E01	SpexDataIn[8]_N	G13
ScanIn[11]	A21	ScanOut[18]	AA03	SpexDataIn[8]_P	H13
ScanIn[12]	T21	ScanOut[19]	T05	SpexDataIn[9]_N	L14
ScanIn[13]	R22	ScanOut[20]	P05	SpexDataIn[9]_P	M15
ScanIn[14]	P21	ScanOut[21]	M05	SpexDataIn[10]_N	F13

Table 9-2. I/O Signal List, Sorted by Signal Name (Page 4 of 5)

Signal Name	Pin Location	Signal Name	Pin Location	Signal Name	Pin Location
SpexDataIn[10]_P	E13	SpexDataOut[10]_N	T15	UnilinkClockIn_N	L04
SpexDataIn[11]_N	N16	SpexDataOut[10]_P	U16	UnilinkClockIn_P	K03
SpexDataIn[11]_P	N17	SpexDataOut[11]_N	N14	UnilinkPLLGndaIn	C01
SpexDataIn[12]_N	K15	SpexDataOut[11]_P	P13	UnilinkPLLVddaIn	B01
SpexDataIn[12]_P	J16	SpexDataOut[12]_N	T17	UnilinkLtestIn	A05
SpexDataIn[13]_N	K17	SpexDataOut[12]_P	R18	UnilinkMtestIn	E25
SpexDataIn[13]_P	L18	SpexDataOut[13]_N	N21	xulp_av25In[0]	U01
SpexDataIn[14]_N	H15	SpexDataOut[13]_P	N20	xulp_av25In[1]	P01
SpexDataIn[14]_P	J15	SpexDataOut[14]_N	R17	xulp_av25In[2]	M01
SpexDataIn[15]_N	L17	SpexDataOut[14]_P	R16	xulp_av25In[3]	K01
SpexDataIn[15]_P	L16	SpexDataOut[15]_N	N18	xulp_av25In[4]	J25
SpexDataOut[0]_N	T09	SpexDataOut[15]_P	N19	xulp_av25In[5]	L25
SpexDataOut[0]_P	R09	SyncIn_N	G16	xulp_av25In[6]	P25
SpexDataOut[1]_N	N05	SyncIn_P	G15	xulp_av25In[7]	T25
SpexDataOut[1]_P	N06	SyncOut[0]_N	F19	xulp_avregOut[0]	Y05
SpexDataOut[2]_N	R10	SyncOut[0]_P	G19	xulp_avregOut[1]	R03
SpexDataOut[2]_P	P11	SyncOut[1]_N	E19	xulp_avregOut[2]	L01
SpexDataOut[3]_N	N09	SyncOut[1]_P	E18	xulp_avregOut[3]	J01
SpexDataOut[3]_P	N10	SyncOut[2]_N	G18	xulp_avregOut[4]	H21
SpexDataOut[4]_N	AA13	SyncOut[2]_P	G17	xulp_avregOut[5]	L23
SpexDataOut[4]_P	Y13	SyncOut[3]_N	F21	xulp_avregOut[6]	R25
SpexDataOut[5]_N	W13	SyncOut[3]_P	E20	xulp_avregOut[7]	U25
SpexDataOut[5]_P	V13	SyncOut[4]_N	N24	xulp_avttIn[0]	V05
SpexDataOut[6]_N	U11	SyncOut[4]_P	N25	xulp_avttIn[1]	R01
SpexDataOut[6]_P	V11	SyncOut[5]_N	G20	xulp_avttIn[2]	L03
SpexDataOut[7]_N	T11	SyncOut[5]_P	G21	xulp_avttIn[3]	J05
SpexDataOut[7]_P	U10	SyncOut[6]_N	N23	xulp_avttIn[4]	J21
SpexDataOut[8]_N	U13	SyncOut[6]_P	N22	xulp_avttIn[5]	L21
SpexDataOut[8]_P	T13	TestM3In	AE05	xulp_avttIn[6]	R21
SpexDataOut[9]_N	V15	ThermalIn	L12	xulp_avttIn[7]	W24
SpexDataOut[9]_P	U15	ThermalOut	G02	xuls_av25In[0]	P07

Table 9-2. I/O Signal List, Sorted by Signal Name (Page 5 of 5)

Signal Name	Pin Location	Signal Name	Pin Location	Signal Name	Pin Location
xuls_av25In[1]	W12	xuls_avregOut[1]	U08	xuls_avttIn[1]	V07
xuls_av25In[2]	W14	xuls_avregOut[2]	U18	xuls_avttIn[2]	W19
xuls_av25In[3]	P19	xuls_avregOut[3]	R19	xuls_avttIn[3]	U19
xuls_avregOut[0]	R07	xuls_avttIn[0]	T07		

Table 9-3. I/O Signal List, Sorted by Grid Position (Page 1 of 5)

Pin Location	Signal Name	Pin Location	Signal Name	Pin Location	Signal Name
A02	DelayIn	C01	UnilinkPLLGndaIn	D19	rulp_avttIn[3]
A03	Lssd_C3_ClkIn	C02	PortDataOut[15]_N	D23	PortDataOut[17]_P
A04	PortDataIn[0]_N	C03	PortDataIn[1]_N	D25	PortDataOut[16]_N
A05	UnilinkLtestIn	C04	PortDataIn[1]_P	E01	ScanOut[17]
A06	PortDataIn[4]_P	C05	PortDataIn[2]_N	E02	PortDataOut[14]_P
A07	PortDataIn[5]_P	C06	PortDataIn[3]_N	E03	ScanOut[23]
A08	PortDataIn[5]_N	C07	PortDataIn[3]_P	E04	PortDataOut[13]_P
A09	rulp_av25In[0]	C08	PortDataIn[6]_N	E10	DI2#
A10	rulp_av25In[1]	C09	PortDataIn[7]_P	E11	rulp_avregOut[1]
A13	SCCIn[1]	C10	ScanIn[0]	E12	ScanIn[7]
A15	rulp_av25In[2]	C12	ScanIn[2]	E13	SpexDataIn[10]_P
A16	rulp_avregOut[2]	C14	ScanIn[3]	E14	ScanIn[8]
A17	rulp_av25In[3]	C15	rulp_avttIn[2]	E15	SCCIn[3]
A18	PortDataIn[10]_N	C16	ScanIn[5]	E16	ScanIn[10]
A19	PortDataIn[10]_P	C17	PortDataIn[8]_P	E18	SyncOut[1]_P
A20	PortDataIn[11]_P	C18	PortDataIn[9]_N	E19	SyncOut[1]_N
A21	ScanIn[11]	C19	PortDataIn[12]_P	E20	SyncOut[3]_P
A22	PortDataIn[14]_N	C20	PortDataIn[12]_N	E21	rulp_avregOut[3]
A23	DI1#	C21	PortDataIn[13]_N	E22	PortDataOut[18]_N
B01	UnilinkPLLVddaIn	C22	PortDataIn[15]_N	E23	RI#
B03	PortDataIn[0]_P	C23	PortDataIn[15]_P	E24	PortDataOut[17]_N
B05	PortDataIn[2]_P	C24	PortDataOut[16]_P	E25	UnilinkMtestIn
B07	PortDataIn[4]_N	D01	PortDataOut[15]_P	F01	PortDataOut[12]_P
B09	PortDataIn[6]_P	D03	PortDataOut[14]_N	F03	PortDataOut[11]_N
B11	ScanIn[1]	D05	rulp_avregOut[0]	F05	PortDataOut[13]_N
B13	SCCIn[2]	D07	rulp_avttIn[0]	F09	SCCIn[0]
B15	ScanIn[4]	D09	PortDataIn[7]_N	F13	SpexDataIn[10]_N
B17	PortDataIn[9]_P	D11	ScanIn[6]	F19	SyncOut[0]_N
B19	PortDataIn[11]_N	D13	rulp_avttIn[1]	F21	SyncOut[3]_N
B21	PortDataIn[13]_P	D15	ScanIn[9]	F23	PortDataOut[18]_P
B23	PortDataIn[14]_P	D17	PortDataIn[8]_N	F25	PortDataOut[19]_N

Table 9-3. I/O Signal List, Sorted by Grid Position (Page 2 of 5)

Pin Location	Signal Name	Pin Location	Signal Name	Pin Location	Signal Name
G01	PortDataOut[12]_N	J02	PortDataOut[9]_N	L04	UnilinkClockIn_N
G02	ThermalOut	J03	PortDataOut[8]_P	L07	ruls_avregOut[0]
G03	PortDataOut[10]_P	J04	PortDataOut[8]_N	L09	SpexDataIn[0]_P
G04	PortDataOut[11]_P	J05	xulp_avttIn[3]	L10	SpexDataIn[2]_N
G10	SpexDataIn[4]_N	J07	ruls_avregOut[1]	L12	ThermalIn
G11	SpexDataIn[4]_P	J10	SpexDataIn[6]_P	L14	SpexDataIn[9]_N
G12	ruls_av25In[1]	J11	SpexDataIn[7]_P	L16	SpexDataIn[15]_P
G13	SpexDataIn[8]_N	J13	SpexDataIn[5]_N	L17	SpexDataIn[15]_N
G14	ruls_av25In[2]	J15	SpexDataIn[14]_P	L18	SpexDataIn[13]_P
G15	SyncIn_P	J16	SpexDataIn[12]_P	L19	ruls_avregOut[3]
G16	SyncIn_N	J18	ruls_avregOut[2]	L21	xulp_avttIn[5]
G17	SyncOut[2]_P	J19	ruls_avttIn[3]	L22	ScanIn[16]
G18	SyncOut[2]_N	J21	xulp_avttIn[4]	L23	xulp_avregOut[5]
G19	SyncOut[0]_P	J22	PortDataOut[23]_P	L24	ScanIn[22]
G20	SyncOut[5]_N	J23	PortDataOut[23]_N	L25	xulp_av25In[5]
G21	SyncOut[5]_P	J24	PortDataOut[22]_P	M01	xulp_av25In[2]
G22	PortDataOut[21]_N	J25	xulp_av25In[4]	M03	ScanOut[15]
G23	PortDataOut[20]_N	K01	xulp_av25In[3]	M05	ScanOut[21]
G24	PortDataOut[20]_P	K03	UnilinkClockIn_P	M07	ruls_av25In[0]
G25	PortDataOut[19]_P	K05	ScanOut[22]	M11	SpexDataIn[2]_P
H01	PortDataOut[9]_P	K07	ruls_avttIn[0]	M13	SpexDataIn[3]_P
H03	PortDataOut[10]_N	K09	SpexDataIn[0]_N	M15	SpexDataIn[9]_P
H07	ruls_avttIn[1]	K11	SpexDataIn[6]_N	M19	ruls_av25In[3]
H11	SpexDataIn[7]_N	K13	SpexDataIn[5]_P	M21	ScanIn[15]
H13	SpexDataIn[8]_P	K15	SpexDataIn[12]_N	M23	ScanIn[21]
H15	SpexDataIn[14]_N	K17	SpexDataIn[13]_N	N05	SpexDataOut[1]_N
H17	ruls_avttIn[2]	K21	ScanIn[17]	N06	SpexDataOut[1]_P
H21	xulp_avregOut[4]	K23	ScanIn[23]	N07	SpexDataIn[1]_N
H23	PortDataOut[21]_P	L01	xulp_avregOut[2]	N08	SpexDataIn[1]_P
H25	PortDataOut[22]_N	L02	ScanOut[16]	N09	SpexDataOut[3]_N
J01	xulp_avregOut[3]	L03	xulp_avttIn[2]	N10	SpexDataOut[3]_P

Table 9-3. I/O Signal List, Sorted by Grid Position (Page 3 of 5)

Pin Location	Signal Name	Pin Location	Signal Name	Pin Location	Signal Name
N12	SpexDataIn[3]_N	R14	Osc625TestPointOut_P	U19	xuls_avttIn[3]
N14	SpexDataOut[11]_N	R16	SpexDataOut[14]_P	U22	PortDataOut[24]_N
N16	SpexDataIn[11]_N	R17	SpexDataOut[14]_N	U23	PortDataOut[24]_P
N17	SpexDataIn[11]_P	R18	SpexDataOut[12]_P	U24	PortDataOut[25]_N
N18	SpexDataOut[15]_N	R19	xuls_avregOut[3]	U25	xulp_avregOut[7]
N19	SpexDataOut[15]_P	R21	xulp_avttIn[6]	V01	PortDataOut[6]_N
N20	SpexDataOut[13]_P	R22	ScanIn[13]	V03	PortDataOut[5]_N
N21	SpexDataOut[13]_N	R24	ScanIn[19]	V05	xulp_avttIn[0]
N22	SyncOut[6]_P	R25	xulp_avregOut[6]	V07	xuls_avttIn[1]
N23	SyncOut[6]_N	T03	RefClockIn_P	V09	DebugBusOut[0]
N24	SyncOut[4]_N	T05	ScanOut[19]	V11	SpexDataOut[6]_P
N25	SyncOut[4]_P	T07	xuls_avttIn[0]	V13	SpexDataOut[5]_P
P01	xulp_av25In[1]	T09	SpexDataOut[0]_N	V15	SpexDataOut[9]_N
P03	ScanOut[14]	T11	SpexDataOut[7]_N	V17	DebugBusOut[12]
P05	ScanOut[20]	T13	SpexDataOut[8]_P	V23	PortDataOut[26]_P
P07	xuls_av25In[0]	T15	SpexDataOut[10]_N	V25	PortDataOut[25]_P
P11	SpexDataOut[2]_P	T17	SpexDataOut[12]_N	W01	PortDataOut[3]_N
P13	SpexDataOut[11]_P	T21	ScanIn[12]	W02	PortDataOut[5]_P
P15	Osc625TestPointOut_N	T23	ScanIn[18]	W03	PortDataOut[4]_N
P19	xuls_av25In[3]	T25	xulp_av25In[7]	W04	PortDataOut[4]_P
P21	ScanIn[14]	U01	xulp_av25In[0]	W08	DebugBusOut[4]
P23	ScanIn[20]	U02	PortDataOut[6]_P	W09	DebugBusOut[1]
P25	xulp_av25In[6]	U03	PortDataOut[7]_P	W10	DebugBusOut[6]
R01	xulp_avttIn[1]	U04	PortDataOut[7]_N	W11	DebugBusOut[7]
R02	ScanOut[13]	U08	xuls_avregOut[1]	W12	xuls_av25In[1]
R03	xulp_avregOut[1]	U10	SpexDataOut[7]_P	W13	SpexDataOut[5]_N
R04	RefClockIn_N	U11	SpexDataOut[6]_N	W14	xuls_av25In[2]
R07	xuls_avregOut[0]	U13	SpexDataOut[8]_N	W15	DebugBusOut[15]
R09	SpexDataOut[0]_P	U15	SpexDataOut[9]_P	W16	DebugBusOut[9]
R10	SpexDataOut[2]_N	U16	SpexDataOut[10]_P	W17	DebugBusOut[13]
R12	PowerOnResetIn#	U18	xuls_avregOut[2]	W18	DebugBusOut[8]

Table 9-3. I/O Signal List, Sorted by Grid Position (Page 4 of 5)

Pin Location	Signal Name	Pin Location	Signal Name	Pin Location	Signal Name
W19	xuls_avttIn[2]	AA21	rulp_avregOut[4]	AC17	PortDataIn[23]_P
W22	PortDataOut[27]_P	AA22	PortDataOut[29]_N	AC18	PortDataIn[22]_N
W23	PortDataOut[26]_N	AA23	Lssd_A_ClkIn	AC19	PortDataIn[19]_P
W24	xulp_avttIn[7]	AA24	PortDataOut[30]_P	AC20	PortDataIn[19]_N
W25	PortDataOut[28]_N	AA25	Lssd_B_ClkIn	AC21	PortDataIn[18]_N
Y01	PortDataOut[3]_P	AB01	PortDataOut[0]_N	AC22	PortDataIn[17]_P
Y03	PortDataOut[2]_P	AB03	PortDataOut[1]_P	AC23	PortDataIn[17]_N
Y05	xulp_avregOut[0]	AB05	rulp_avttIn[7]	AC24	PortDataOut[31]_N
Y09	DebugBusOut[5]	AB07	SHIClockIn	AD01	CorePLLGndaIn
Y11	DebugBusOut[3]	AB09	PortDataIn[24]_N	AD03	PortDataIn[30]_P
Y13	SpexDataOut[4]_P	AB11	ScanOut[1]	AD05	PortDataIn[29]_P
Y15	DebugBusOut[11]	AB15	ScanOut[4]	AD07	PortDataIn[27]_N
Y17	DebugBusOut[14]	AB17	PortDataIn[23]_N	AD09	PortDataIn[25]_P
Y21	PortDataOut[29]_P	AB19	SHISerialDataOut	AD11	ScanOut[7]
Y23	PortDataOut[27]_N	AB21	rulp_avttIn[4]	AD13	InterruptOut#
Y25	PortDataOut[28]_P	AB23	PortDataOut[30]_N	AD15	ScanOut[10]
AA01	ScanOut[12]	AB25	PortDataOut[31]_P	AD17	PortDataIn[22]_P
AA02	PortDataOut[1]_N	AC01	CorePLLVddaIn	AD19	PortDataIn[20]_P
AA03	ScanOut[18]	AC02	PortDataOut[0]_P	AD21	PortDataIn[18]_P
AA04	PortDataOut[2]_N	AC03	PortDataIn[31]_P	AD23	PortDataIn[16]_P
AA07	rulp_avregOut[7]	AC04	PortDataIn[31]_N	AE03	IOTtestIn
AA08	DebugBusOut[2]	AC05	PortDataIn[29]_N	AE04	PortDataIn[30]_N
AA09	rulp_avregOut[6]	AC06	PortDataIn[28]_N	AE05	TestM3In
AA10	ScanOut[0]	AC07	PortDataIn[28]_P	AE06	PortDataIn[27]_P
AA11	rulp_avttIn[6]	AC08	PortDataIn[25]_N	AE07	PortDataIn[26]_P
AA12	ScanOut[2]	AC09	PortDataIn[24]_P	AE08	PortDataIn[26]_N
AA13	SpexDataOut[4]_N	AC10	ScanOut[6]	AE09	rulp_av25In[7]
AA14	ScanOut[3]	AC12	ScanOut[8]	AE10	SHISelectIn#
AA15	rulp_avregOut[5]	AC14	ScanOut[9]	AE11	LeakageTestIn
AA16	ScanOut[5]	AC15	rulp_avttIn[5]	AE12	rulp_av25In[6]
AA18	DebugBusOut[10]	AC16	ScanOut[11]	AE13	FullyInsertedIn#

Table 9-3. I/O Signal List, Sorted by Grid Position (Page 5 of 5)

Pin Location	Signal Name	Pin Location	Signal Name	Pin Location	Signal Name
AE14	rulp_av25In[5]	AE18	PortDataIn[21]_N	AE22	PortDataIn[16]_N
AE15	LssdTestEnableIn	AE19	PortDataIn[21]_P	AE23	Lssd_C2_ClkIn
AE16	SHISerialDataIn	AE20	PortDataIn[20]_N	AE24	DelayOut
AE17	rulp_av25In[4]	AE21	Lssd_C1_ClkIn		

10. Electrical Characteristics

10.1 General Information

Table 10-1. Absolute Maximum Ratings

Symbol	Parameter	Rating			Units	Notes
		Minimum	Typical	Maximum		
V_{DD}	Supply voltage	1.65	1.8	1.95	V	1
V_{IN}	Input voltage (LVDS)	-0.2		V_{DD} + 0.2	V	1, 2
V_{icm}	Receiver common mode range (ILVDS)	0	1.25	V_{DD}	V	
Θ_{JA}	Thermal impedance junction to ambient package, airflow = 0 feet per minute		10.8		°C/W	3
Θ_{JA}	Thermal impedance junction to ambient package, airflow = 200 feet per minute		8.4		°C/W	3
Θ_{JA}	Thermal impedance junction to ambient package, airflow = 400 feet per minute		6.6		°C/W	3
Θ_{JC}	Thermal impedance junction to case package		0.20		°C/W	4
T_A	Operating junction temperature	0	70	125	°C	1
T_S	Storage temperature	-65		150	°C	1
	Electrostatic discharge	-3000	6000	3000	V	1

1. Permanent device damage can occur if the above absolute maximum ratings are exceeded. Extended exposure to absolute maximum rating conditions can affect device reliability. Normal operation should be restricted to the conditions listed in *Table 10-2.*
2. For $V_{DD} \leq 1.95$ V.
3. For devices mounted to a 2S2P card (1-ounce copper, size 63.5 mm $\times$ 76.2 mm), with flow on both sides of the card in a vertical orientation.
4. Θ_{JC} represents the temperature difference between the junction and the top center of the outside surface of the component package, divided by the power applied to the component mounted to the test card.

Table 10-2. Recommended Operating Conditions

Symbol	Parameter (for TTL-compatible I/Os)		Rating			Units	Notes
			Minimum	Typical	Maximum		
V_{DD}	Supply voltage		1.71	1.8	1.89	V	1, 2
V_{IH}	Input up level		V_{REF} + 0.1		V_{DD} + 0.3	V	3
V_{IL}	Input down level		-0.3		V_{REF} - 0.1	V	3
V_{OH}	High-level output voltage (I_{OH} = -8 mA)		V_{DD} - 0.4			V	
V_{OL}	Low-level output voltage (I_{OL} = 8 mA)				0.4	V	
I_{IL}	Receiver maximum input leakage, low-level input current at least-positive down level	No pullup or pulldown			0	µA	4
		With pullup			0	µA	
		With pulldown			-150	µA	
I_{IH}	Receiver maximum input leakage, high-level input current at most-positive up level	No pullup or pulldown			0	µA	4
		With pullup			200	µA	
		With pulldown			0	µA	
C_I	Input capacitance (V_{DD} = nominal)				5	pF	

1. For power-up, the +2.5 V supply is activated either prior to or concurrently with the +1.8 V supply.
2. For power-down, the +2.5 V supply is deactivated either after or concurrently with the +1.8 V supply.
3. $V_{REF} = V_{DD} \div 2$
4. Applies to bidirectional receivers without pullup or pulldown resistors.

Table 10-3. Power Dissipation for Device Configured as Master

Supply	Power (W)		Current (A)	
	Nominal	Maximum	Nominal	Maximum
Unilink AV25 (2.5 V)	1	1.35	0.4	0.54
1.8 V	24	27	13.4	15

Table 10-4. Power Dissipation for Device Configured as Slave

Supply	Power (W)		Current A)	
	Nominal	Maximum	Nominal	Maximum
Unilink AV25 (2.5 V)	1	1.35	0.4	0.54
1.8 V	17	19	9.5	10.6

Table 10-5. Core Clock Characteristics

Parameter		Rating		Units	Notes
		Minimum	Maximum		
Internal clock frequency 1		220	250	MHz	
Internal clock frequency 2		110	125	MHz	
RefClockIn frequency	Internal PLL	55	62.5	MHz	1, 2

1. The RefClockIn input must have a tolerance of ±100 ppm (0.01%), a duty cycle of 40 to 60 percent, and a phase jitter of ±150 ps (cycle to cycle) maximum.
2. The skew between all the PowerPRS Q-64G RefClockIn signals in a multiple-device configuration must be less than 1 ns.

Table 10-6. Unilink Clock Characteristics

Parameter		Rating		Units	Notes
		Minimum	Maximum		
Internal clock frequency		550	625	MHz	
UnilinkClockIn frequency	Internal PLL	110	125	MHz	1, 2

1. The UnilinkClockIn input must have a tolerance of ±100 ppm (0.01%), a duty cycle of 40 to 60 percent, and a phase jitter of ±40 ps (cycle to cycle) maximum.
2. There is no timing constraint on the skew between the master and slave UnilinkClockIn input pins.

10.2 Internal PLL AVDD and AGND Pins

The AVDD and AGND pins are the voltage supply pins to the analog circuits in the internal phase-locked loop (PLL). The AVDD voltage supply must be between 1.65 and 1.95 V. Noise on these two signals causes phase jitter at the PLL output. To isolate the PLL from the noisy internal digital V_{DD} and ground signals, the AVDD and AGND are connected internally to dedicated package pins. If limited noise is expected at the board level, the AVDD pin can be connected directly to the digital V_{DD} plane; however, it is often prudent to place a filter circuit on the AVDD pin (see *Figure 10-1).* The AGND pin should be brought out from the package and connected to the digital ground plane at the AVDD capacitor. All wire lengths should be as short as possible to minimize coupling from other signals.

Figure 10-1. Internal PLL AVDD and AGND Signals

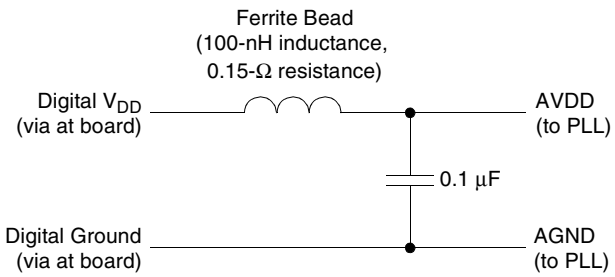

The impedance of the ferrite bead should be much higher than that of the capacitor at frequencies where noise is expected. In many applications, a resistor is better than a ferrite bead at reducing jitter. The resistor value should be less than 10 Ω. Experimentation is the best way to determine the optimal filter design for a specific application.

10.3 InfiniBand Compatibility

The PowerPRS Q-64G is compatible, but not completely compliant, with the InfiniBand physical layer standards. The PowerPRS Q-64G core logic, PLL, and phase rotator are identical to the InfiniBand standard. Differences between the PowerPRS Q-64G and the InfiniBand standard are summarized in *Table 10-7*. The PowerPRS Q-64G can operate with InfiniBand links as described in *InfiniBand Architecture, Volume 2: Physical Specifications* (see *Related Documents* on page 197) for high-speed (2.5-Gbps) electrical signaling.

Table 10-7. PowerPRS Q-64G versus InfiniBand

Parameter	PowerPRS Q-64G	InfiniBand	Units
Transmitter differential signal level (normalized driver setting)	665 nominal	1300 nominal	mV_{PP}
Receiver differential voltage required at center of receive eye	200	150	mV_{PP}
Power (per link at 2.5 Gbps)	200	235	mW
Signal detection level (SIGLEVx)	0 > 200 mV guarantees signal < 50 mV guarantees noise 1 > 400 mV guarantees signal < 2000 mV guarantees noise	0 > 175 mV guarantees signal < 85 mV guarantees noise	N/A (binary)

11. Mechanical Information

Figure 11-1. Package Mechanical

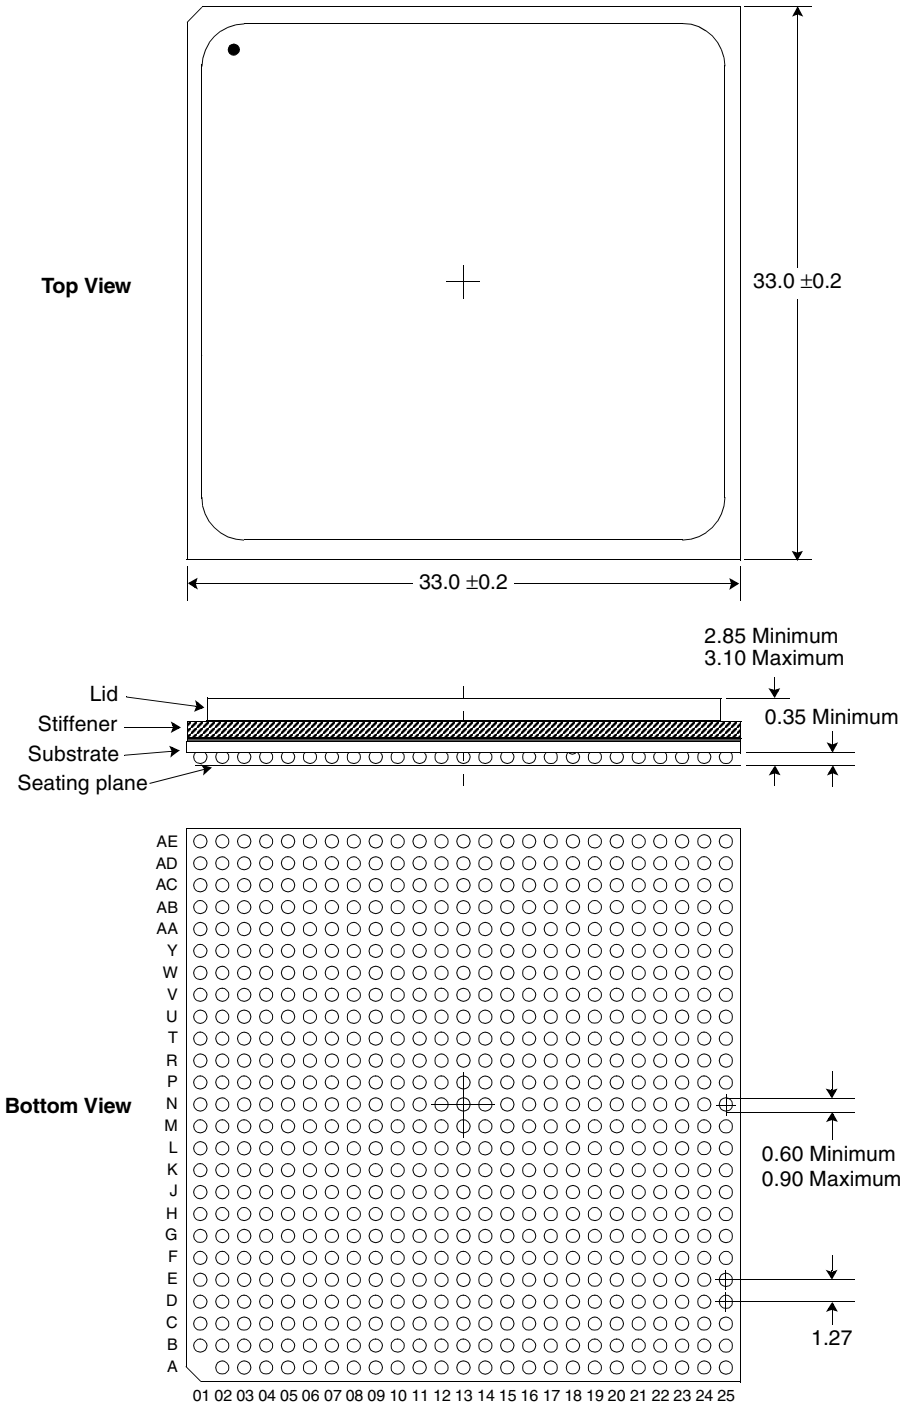

Package type is 624-ball HyperBGA, package size is 33 × 33 mm, and array size is 25 × 25. All dimensions are in millimeters.

12. Glossary

absolute maximum rating The highest value a quantity can have before malfunction or damage occurs.

address A number designating a particular memory location.

array An ordered arrangement of data elements.

ATM asynchronous transfer mode

b bit

B byte

bandwidth The minimum capacity required for effective transmission and reception of a data packet.

best-effort delivery The delivery of packets with no bandwidth guarantee and an unspecified quality of service. Packets can be discarded during periods of traffic congestion.

BGA ball grid array

bidirectional The ability to transmit in both directions.

BIST built-in self-test

bitmap A binary representation in which a bit or set of bits corresponds to an assigned value or condition.

buffer A memory bank used for temporary storage.

bus A common pathway over which input and output signals are routed.

clock An internal timing device that synchronizes the data pulses between the transmitter and receiver.

clock cycle One “tick” of the clock. For example, a 100-MHz clock has 100 million ticks per second.

clock frequency The reciprocal of the time period of a single clock cycle.

CMOS complementary metal-oxide semiconductor

configuration The arrangement of speed expansion, port expansion, and port-paralleling options that create a custom switching device.

control packet The packet that carries the communications between the local processor and the protocol engine. Control packets do not have a specific priority.

counter A circuit that counts pulses and generates an output at a specified time.

CPU central processing unit

CRC cyclic redundancy check. Code used to validate a block of data.

credit table	A weighted cycling mechanism that can be programmed to guarantee minimum bandwidth for low-priority packets. The credit table includes 256 entries, or credits, per port.
CSIX	common switch interface
data packet	See “packet.”
differential pair	Two wires of opposite polarity configured as a pair to reduce signal noise and crosstalk.
driver	Also called a “device driver.” A routine that links a peripheral device to the operating system and performs internal functions, or a functional unit that increases the output current, power, or voltage of another functional unit.
egress flow	Data flow from a PowerPRS Q-64G to an attached device.
FIFO	first-in-first-out
filter	A pattern or mask through which only selected data is passed (for example, bitmap filter, best-effort discard filter, and so forth).
FIR	finite impulse response
Gbps	gigabits per second
GRA	global register array
hysteresis	The lag between making a change and the response or effect of the change.
I/O	input/output
idle packet	The packet that is transmitted when the other packet types are either unavailable for transmission or prevented from transmission. Idle packets carry flow control and link synchronization information, but they do not carry user data.
IEEE	Institute of Electrical and Electronics Engineers
ingress flow	Data flow from an attached device to a PowerPRS Q-64G.
interrupt	A signal that gains the attention of the CPU and is usually generated when input or output is required.
jitter	A flicker or fluctuation in a transmission signal caused by a bit arriving either ahead or behind the standard clock cycle or, more generally, the variable arrival of packets.
JTAG	Joint Test Action Group. IEEE Standard 1149.1 (regarding boundary-scan architecture) is also referred to as the JTAG standard, after the group that developed it.
junction	The region of contact between opposite types of semiconductor materials.

K character An 8b/10b code control character used for link synchronization and supervision.

latency The lag between initiating a data request and starting the actual data transfer.

line An electronic communications channel such as a wire.

load balancing The fine tuning of the network to more evenly distribute the data and/or processing across all available resources.

local processor The microprocessor connected to the PowerPRS Q-64G with the serial host interface (SHI) and used to control device configuration.

look-up table A table that allows the byte transmission sequence of egress packets to be rearranged.

LSB least significant bit

LSSD level-sensitive scan design

LU logical unit. The part of a packet processed by one PowerPRS Q-64G input controller or one output controller.

LVCMOS low-voltage complementary metal-oxide semiconductor

mask A bit pattern used to change (reject) or extract (accept) bit positions in another bit pattern. For example, when the Boolean AND operation is used to match a mask of ‘0’s and ‘1’s with a string of data bits and a ‘1’ occurs in both the mask and the data, the resulting bit will contain a ‘1’ in that position. Hardware interrupts are enabled and disabled in this manner, with each interrupt assigned a bit position in the *Interrupt Mask Register* (page 122).

Mbps megabits per second

MBps megabytes per second

memory bank A physical section of a device memory.

MISR multiple-input signature register

mode An operational state of at least two possible conditions to which a system can be switched.

MSB most significant bit

multicast A one-to-many transmission, where “many” is specifically defined (contrast with broadcast).

multiplexer A device that merges several low-speed transmissions into one high-speed transmission and vice versa.

OC-192 The synchronous optical network (SONET) transmission rate of 9953.28 Mbps.

OC-48	The SONET transmission rate of 2488.32 Mbps.
packet	The user data element that the PowerPRS Q-64G processes in equal lengths called logical units (LUs).
packet header	The first two to three bytes in a packet that contain destination bitmap, packet priority, and flow control information, all protected by a parity bit.
packet priority	Four levels of data packet priority provide quality-of-service support. Data packets are prioritized from zero to three, with zero being the highest priority.
packet qualifier byte	The first byte (H0) of the packet header. This byte contains important information about the packet, such as packet type, packet priority, and so forth.
parity	The number of bits (or the number of similar bits) are even or odd, as intended.
payload	The part of the packet that carries the message data (contrast with packet header).
pin	The male lead on a chip that is plugged into its female counterpart to complete the circuit. The number of pins reflects the number of wires, or pathways, that can carry signals.
pinout	A diagram of the integrated circuit that shows the locations of the pins for various functions.
pipeline processing	Parallel processing that involves overlapping operations by moving data or instructions into a conceptual pipe, where all stages of the pipe are processed simultaneously.
PLL	phase-locked loop
port paralleling	A PowerPRS Q-64G configuration option for reducing the number of device ports. Multiple ports are grouped to form one link.
PRPG	pseudorandom pattern generator
PowerPRS C192	The IBM-approved product nickname for the IBM PowerPRS C192 Common Switch Interface.
PowerPRS Q-64G	The IBM-approved product nickname for the IBM PowerPRS Q-64G Packet Routing Switch.
pulldown	A circuit that lowers the value of a connected device.
pullup	A circuit that raises the value of a connected device.
pulse	A transient signal of short duration, constant amplitude, and one polarity.
quality of service	The ability to define a level of performance.
queue	A temporary holding place for egress data.

RAM	random access memory
read/clear	A register field in which the value is cleared immediately after a read.
receiver	An electronic device that accepts signals, and processes or converts them for internal use.
register	A small, high-speed circuit that stores internal operation values, such as the address of the instruction being executed and the data being processed.
resistor	An electronic component that resists the flow of current in a device.
SCC	side communication channel
sequencer	The PowerPRS Q-64G component that controls the internal data flow by granting shared memory access to the input and output ports.
service packet	The packet that carries the communications between the local processor and the attached devices, and the packet that tests link liveness. There are event-1, event-2, and command service packets.
SHI	serial host interface
skew	A timing change in a transmission signal.
SRAM	static random access memory
SRL	set/reset latch
stage	A complete functional unit of a system (see “pipeline processing”).
standby	The state in which the PowerPRS Q-64G is not in operation, but can be immediately activated.
stream	A contiguous flow of bits, bytes, or data from one place to another.
subswitch element	One of four 16 × 16-port PowerPRS Q-64G components that house a shared memory bank and a control section. The four subswitch elements are designated A, B, C, and D.
switch fabric	The PowerPRS Q-64G internal interconnect architecture that redirects the ingress and egress data flow.
switchover	The process of redirecting the data flow between redundant switch planes.
TDM	time-division multiplexing
timeout	The expiration of an allotted time period for a given operation.
tolerance	The amount of error allowed in a value, rating, dimension, and so forth.
traffic	Data crossing the network.
transmitter	An electronic device that generates signals.

UL	See “Unilink.”
UL RxPort	Unilink port receiver
UL RxSpex Bus	Unilink speed-expansion bus receiver
UL TxPort	Unilink port transmitter
UL TxSpex Bus	Unilink speed-expansion bus transmitter
unicast	A single transmission.
Unilink	The high-speed (2.125 to 2.6 Gbps) serial link that provides backplane point-to-point connectivity between the PowerPRS Q-64G and attached devices. There is one Unilink per PowerPRS Q-64G device port. Each Unilink is comprised of two pairs of differential lines; one differential pair carries ingress flow and the other differential pair carries egress flow.
via	In the printed circuit board, a conducting pathway between two or more substrates (layers).
WAN	wide-area network

13. Related Documents

ASIC SA-27E Databook, Part 1: Base Library and I/Os and Part 2: Macros, IBM Corporation, June 2001 (contact your IBM representative for access to this document).

IBM Packet Routing Switch PRS64G Datasheet, IBM Corporation, July 2001 (see http://www-3.ibm.com/chips/techlib/techlib.nsf/products/Packet_Routing_Switch_PRS64G).

IBM PowerPRS C192 Common Switch Interface Summary Datasheet, Preliminary, IBM Corporation, September 2001 (see http://www-3.ibm.com/chips/techlib/techlib.nsf/products/PowerPRS_C192_Common_Switch_Interface).

IEEE Standard Test Access Port and Boundary-Scan Architecture, IEEE Standard 1491.1-1990, IEEE Standards Association, 1990.

InfiniBand Architecture Specification, Volume 2: Physical Specifications, Release 1.0.a, InfiniBand Trade Association, June 2001 (see http://www.infinibandta.org/estore.html).

Revision Log

Revision Date	Contents of Modification
Sept. 7, 2001	Initial release (00).
Dec. 20, 2001	First revision (01). Added IEEE trademark information to legal page. Added a note about H2 to *Figure 3-4*. In *Table 5-1* and *Section 5.2.2,* changed *UL Errors Register* access type to read only. In *Section 5.1.3 Unilink PLL Programming Register,* updated description of bits 22:31. In *Section 5.2.1 UL Global Register,* updated descriptions of bits 10, 13, 24:27, 28:31. In *Section 5.2.6 UL TxPort Parameters Register,* updated descriptions of bits 11:13, 17:19. In *Section 5.2.26 UL TxSpex Bus Parameters Register,* updated descriptions of bits 11:13, 17:19. In *Section 5.2.33 UL RxSpex Bus Latency Programming Register,* updated description of bit 4. In *Section 5.3.3 Threshold Access Register,* updated description of shared memory threshold. In *Section 5.3.6 Status Register,* updated reset value. In *Section 5.6.15 Unilink Debug Control Register,* updated description of bits 3:7. In *Section 5.6.16 Unilink Force Error Register,* updated description of bits 27:31. In *Section 6.1 Reset Sequence,* updated step 19. In *Figures 6-1* and *6-2,* changed “step” to “granularity”. In *Table 6-1,* adjusted Speed-Expansion Bus Delay column heading. In *Section 6.3.1 Initializing Unilink Ports,* added LU deskew algorithm to step 7. In *Section 6.3.2 Deactivating Unilink Ports,* added step 7. In *Section 6.4 Logic BIST Execution Sequence,* updated steps 11-18. In *Table 7-1,* updated descriptions of SpexDataOut[0:15]_N/P, SCCIn[0:3], and DelayOut. In *Table 7-3,* updated descriptions of ScanIn[0:23], ScanOut[0:23], and ThermalIn/Out. In *Table 7-4,* updated the second parameter. Updated *Figure 7-1.* Added *Section 7.2.3 SyncIn/SyncOut Signals* and *Figure 7-3*.