General Information

Features

- Companion to the IBM PowerPRS™ 64G and PowerPRS Q-64G Packet Routing Switch chips
- Direct attachment to the IBM PowerPRS Switch Core Interface Chip (SCIC), the Unilink to PowerPRS 64G-compatible data-aligned synchronous link (DASL) converter
- CSIX interface attachment to PowerPRS 64G and PowerPRS Q-64G switch cores
- CSIX-L1 interface: OC-48 and OC-192 adapter, compliant with the Common Switch Interface - Level 1 specification
- Redundant switch port attachment at 16 Gbps using 2.5-Gbps serial links compatible with InfiniBand™ physical layer standards
- Dual-switch attachment for redundant switch-plane operation, including programmable scheduled switchover (packet lossless) and hot standby switchover
- PowerPRS 64G interface: 16- to 20-byte logical unit (LU) packet processing with four-way port paralleling
- PowerPRS Q-64G interface: 8- to 10-byte LU packet processing with two-way port paralleling
- Shared buffer capacity of up to 2048 ingress packets (1024 per switch plane) and up to 1024 egress packets (shared between switch planes)
- Configurable number of traffic priorities (from two to four)
- Packet header parity generation and checking
- End-to-end packet payload protection, with optional cyclic redundancy check (CRC) insertion
- Programmable generation and detection of link liveness messages in yellow packets
- Eight-bit parallel processor interface to access all registers for control and error reporting
- Internal loopback support for both the switch interface and CSIX adapter interface
- Internal logic built-in self-test (BIST) and memory BIST
- SA-27E technology (L_{drawn} = 0.18 µm, L_{eff} = 0.11 µm):
  - 1.8-V core voltage
  - 2.5-V LVCMOS-compatible (3.3-V tolerant) I/Os for microprocessor interface
  - 2.5-V LVCMOS- or 1.5-V HSTL-compatible I/Os for CSIX-L1 interface
- IEEE Standard 1149.1 boundary scan to facilitate circuit-board testing
- 840-ball IBM HyperBGA™ package

Description

The IBM PowerPRS C192 Common Switch Interface is a companion device to the IBM PowerPRS 64G and PowerPRS Q-64G Packet Routing Switches. It functions as the switch core access layer between the protocol engine's four OC-48 or OC-192 CSIX interfaces and the switch core. The PowerPRS C192 switch interface is comprised of eight 2.5-Gbps Unilink pairs that provide a total throughput of 20 Gbps. The switch port payload throughput is 16 Gbps because the Unilinks use the same 8b/10b line coding scheme as the fiber channel standard.

In PowerPRS 64G applications, the PowerPRS C192 is attached to the PowerPRS SCIC, which converts the 8 Unilinks into 32 PowerPRS 64G-compatible data-aligned synchronous links (DASLs). The packet length in the PowerPRS 64G, SCIC, and C192 is programmable from 64 to 80 bytes, in 4-byte increments. Packets flowing to (ingress) and from (egress) the PowerPRS 64G are divided into four logical units (LUs) of 16 to 20 bytes each, depending on the packet length. In PowerPRS 64G applications, the PowerPRS C192 processes four packets in parallel.

The PowerPRS C192 attaches directly to the Unilink-compatible PowerPRS Q-64G. When connected to the PowerPRS Q-64G, the PowerPRS C192 packet length is programmable from 64 to 80 bytes, in 8-byte increments. Ingress and egress packets are divided into eight LUs of 8 to 10 bytes each. In PowerPRS Q-64G applications, the PowerPRS C192 processes two packets in parallel.

The PowerPRS C192 can also be used to attach four OC-48 subports to the PowerPRS Q-64G. Ingress traffic received from the four OC-48 protocol engines is multiplexed by the PowerPRS C192 16-Gbps switch interface, and egress traffic received from the 16-Gbps switch interface is demultiplexed by the PowerPRS C192 to its OC-48 destination.

The PowerPRS C192 provides attachment to a redundant switch fabric. Two independent data paths (X and Y) can be clocked, reset, and controlled separately to activate or deactivate each switch plane independently. PowerPRS C192 hardware-assist functions perform scheduled switchover without packet loss as well as hot standby switchover.

The PowerPRS C192 processes data traffic using two to four priorities (configured in the registers). An in-band flow control mechanism, carried in the packet header, controls the traffic flow. In-band flow control is performed per priority and destination, and end-to-end flow control is performed per PowerPRS C192 OC-48 subport (PowerPRS Q-64G configuration only) and propagated to the CSIX protocol engine. During PowerPRS C192 ingress traffic congestion, flow control information is propagated through the PowerPRS C192 CSIX interface to the egress protocol engine. If the egress protocol engine traffic is congested, the flow control information is transmitted in band through the ingress interface according to CSIX specification. An optional out-of-band flow control interface is reserved for OC-192 applications (no subports). Activation of either the in-band or out-of-band flow control interface is selected during PowerPRS C192 initialization.

Ordering Information

Part Number	Description	Throughput
IBM3254P2537	IBM Common Switch Interface	10 Gbps

Architecture

Figure 1 illustrates the integration of a PowerPRS C192 in a PowerPRS 64G switching system. The PowerPRS C192 connects to the switch core through the PowerPRS SCIC, which converts the eight 2.5-Gbps Unilinks (2 Gbps of switch payload) into 500-Mbps data-aligned synchronous links (DASLs) for compatibility with the four PowerPRS 64G 4-Gbps switch ports configured for port paralleling. A 128-Gbps switch core can accommodate up to eight OC-192 optical ports.

Figure 1. System View of the PowerPRS C192 with the Network Processor and Packet Routing Switches

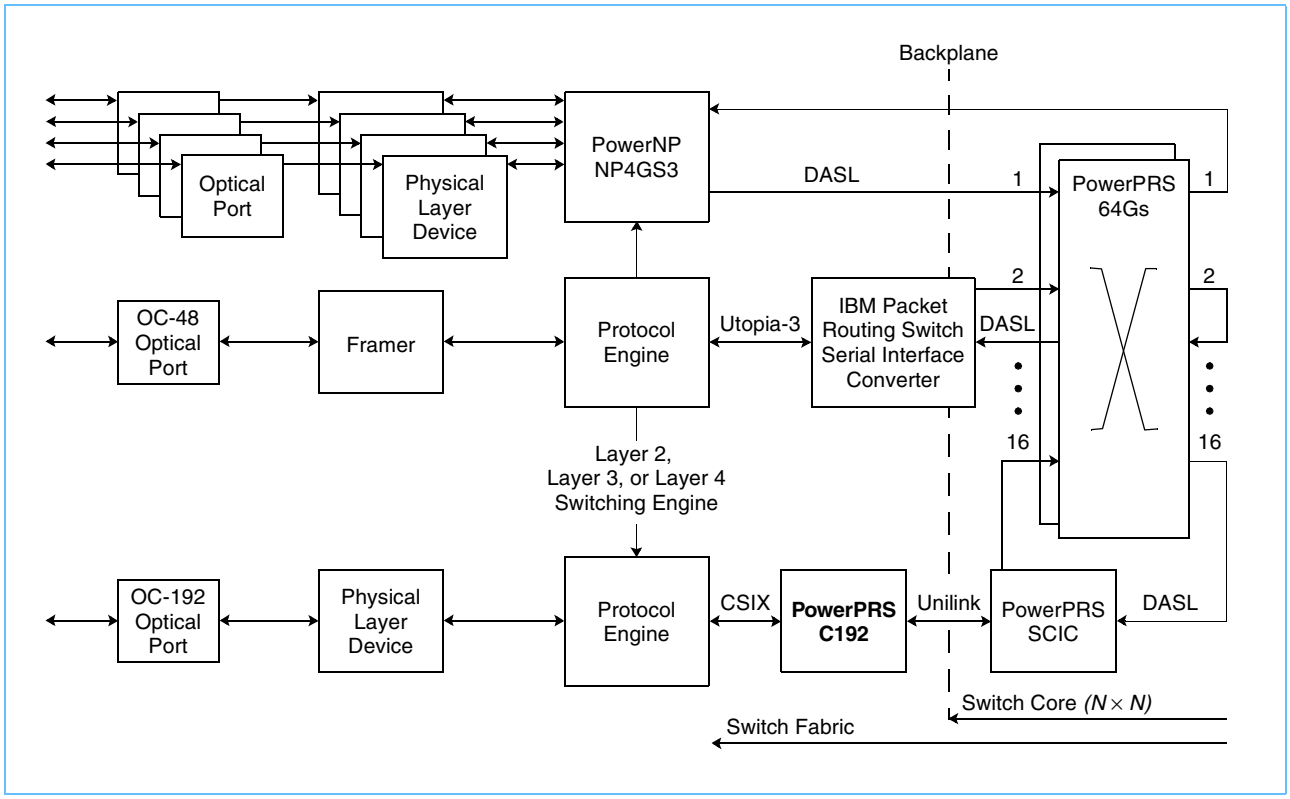

In a system architecture comprised of 4-Gbps and 16-Gbps protocol engines, the PowerPRS C192 operates with both the IBM PowerNP™ NP4GS3 Network Processor and the IBM Packet Routing Switch Serial Interface Converter for a smooth migration to new high-speed optical ports, such as OC-192 SONET or 10-Gbps Ethernet. To facilitate migration to a higher port speed, the PowerPRS C192's ingress and egress byte-shuffling tables can be configured via corresponding registers to adapt to any packet format. The PowerPRS C192 can operate in a switching system with as few as 4 or as many as 32 logical ports.

Figure 2 illustrates the integration of a PowerPRS C192 in a 32-port PowerPRS Q-64G redundant switching system. An integrated serializer/deserializer enables the PowerPRS C192 and PowerPRS Q-64G to directly connect to each other. The PowerPRS C192 supports the PowerPRS Q-64G packet length of 64, 72, or 80 bytes. When attached to four or eight PowerPRS Q-64Gs, the PowerPRS C192 can be used to build a 16- or 32-port switching system, respectively, or extended to build 64-port systems in the future. The PowerPRS C192 can be used to attach either a single 10-Gbps protocol engine (64- or 128-bits wide) or up to four 2.5-Gbps protocol engines (32-bits wide each) using the CSIX-L1 protocol and interface specifications. To support the four CSIX subports, ingress *and* egress PowerPRS C192s provide end-to-end flow control as well

as 64K-entry multicast tables for separate subport addressing. A central control located in the switch local processor can update the multicast tables or updates can be distributed through the PowerPRS C192's eight-bit parallel processor interface.

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

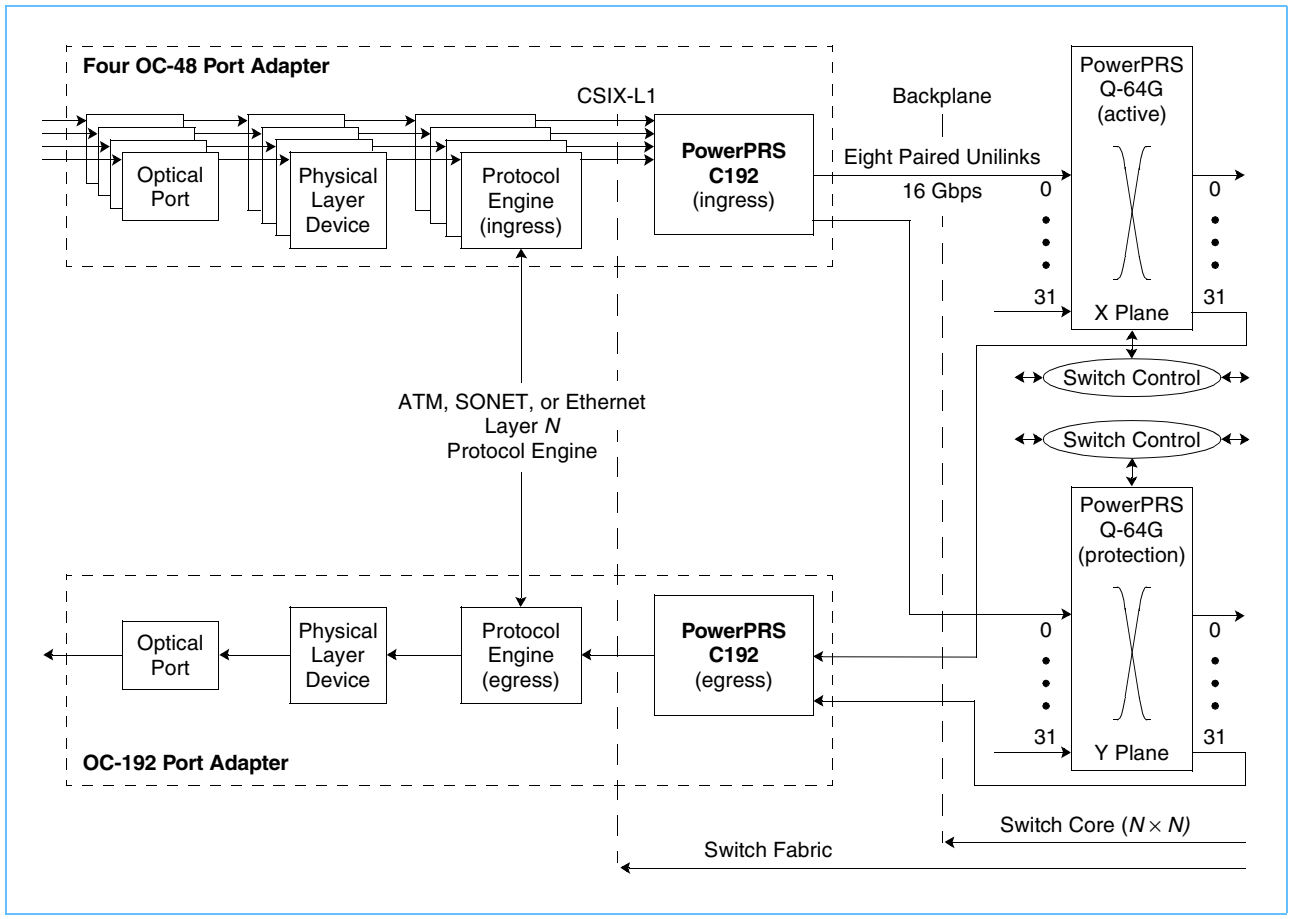

PowerPRS C192 attachment to a redundant switch core allows the balancing of traffic loads between two switch planes. In addition, the PowerPRS C192 and the PowerPRS Q-64G can jointly execute scheduled switchover without packet loss. The internal structure of the PowerPRS C192 is presented in *Figure 3*.

Figure 3. PowerPRS C192 Block Diagram

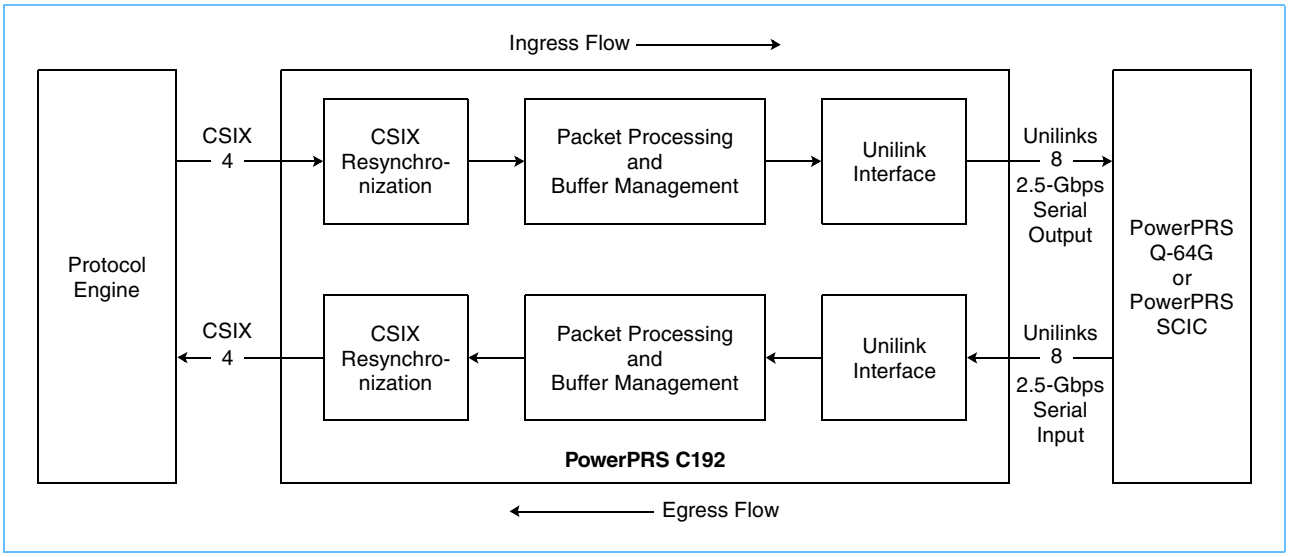

Ingress Data Flow

Ingress packets received from the protocol engine are color-coded and queued for transmission on either the X or Y switch path. The PowerPRS C192 ingress buffer stores up to 2048 incoming packets (1024 per switch plane), and implements programmable filtering to prevent packet duplication and wasted buffer space. Packets are queued (or dequeued) using a first-in-first-out (FIFO) mechanism per priority. The destination queue status (empty or occupied), packet priority, and target switch output queue status flow control information are reported to the ingress scheduler. The ingress scheduler uses a flywheel mechanism to select the next packet for transmission to the switch interface. The highest-priority packets of granted destinations are generally transmitted first. However, an ingress credit table can be programmed at system initialization to guarantee minimum bandwidth to low-priority packets. When activated via the corresponding register, the credit table alters the scheduler selection. Alternatively, a flywheel weighted in favor of the low-priority packets may be used to offset the transmission of higher-priority packets.

Egress Data Flow

Egress packets received from the switch are queued by switch plane in up to 32 output queues (4 ports $\times$ 4 priorities $\times$ 2 switch planes) for transmission to the CSIX interface. The PowerPRS C192 egress buffer is shared between the X and Y switch paths, and stores up to 1024 outgoing packets. The egress queue status and packet priority and CSIX subport destination flow control information are reported to the egress scheduler. The egress scheduler uses four flywheels (one per subport) to select the next packet for transmission to the CSIX interface. A subport service flywheel provides each subport with equal access to the egress buffer. The subport flywheels toggle between the occupied queues of the same priority on the X and Y switch paths to select packets. The highest-priority packets of granted destinations are generally transmitted first. However, the egress credit table can be programmed to guarantee minimum bandwidth to low-priority packets. Each subport flywheel has access to the egress credit table.

When a multicast packet is received from the switch, it is stored (one time only) in the egress buffer and queued in the designated number of subport output queues. The buffer space allocated to the multicast packet is available again only after the required number of packets (up to four) have been transmitted.

Programming Interface and Registers

The PowerPRS C192 employs an eight-bit parallel processor programming interface. This parallel processor interface provides the read/write access to all the PowerPRS C192 internal registers and diagnostic functions, such as online error detection and reporting, and built-in self-test (BIST). *Table 1* summarizes the registers that provide the mechanism for PowerPRS C192 configuration specification and status reporting.

Table 1. Register Summary (Page 1 of 3)

Register Name	Address		Access
	X Plane	Y Plane	
CSIX Interface Control Registers			
CSIX Mode Control Register	x‘00’		Read/Write
CSIX Checking Enable Register	x‘01’		Read/Write
CSIX Interface Error/Status Register	x‘02’		Read/Write
CSIX Interface Error/Status Interrupt Register	x‘06’		Read/Write
Switch Interface Configuration Registers			
Switch Interface System Configuration 1 Register	x‘74’	x‘B4’	Read/Write
Switch Interface System Configuration 2 Register	x‘64’	x‘A4’	Read/Write
Switchover Control Register	x‘09’		Read/Write
Unilink Synchronization 1 Register	x‘C2’	x‘E2’	Read/Write
Unilink Synchronization 2 Register	x‘C3’	x‘E3’	Read/Write
Unilink Synchronization 3 Register	x‘C4’	x‘E4’	Read/Write
Ingress Data Count Register	x‘62’	x‘A2’	Read/Clear
Egress Data Count Register	x‘63’	x‘A3’	Read/Clear
Unilink Control Register	x‘C0’	x‘E0’	Read/Write
Switch Interface Event/Error Register	x‘61’	x‘A1’	Read/Clear
Switch Interface Interrupt Register	x‘5A’	x‘9A’	Read/Write
Switch Interface Checking Enable Register	x‘5B’	x‘9B’	Read/Write
Payload CRC Error Counter Register	x‘40’	x‘80’	Read/Clear
Yellow Packet Transmit Counter Register	x‘41’	x‘81’	Read/Write
Yellow Packet Receive Counter Registers	x‘42’ to x‘43’	x‘82’ to x‘83’	Read/Clear
Yellow Packet Receive Counter Registers (PowerPRS 64G only)	x‘5C’ to x‘5D’	x‘9C’ to x‘9D’	Read/Clear
Unilink Debug Control Register	x‘75’	x‘B5’	Read/Write
Unilink Test Register	x‘76’	x‘B6’	Read/Write
Unilink Error 1 Register	x‘C7’	x‘E7’	Read/Write
Unilink Error 2 Register	x‘C8’	x‘E8’	Read/Write

Table 1. Register Summary (Page 2 of 3)

Register Name	Address		Access
	X Plane	Y Plane	
Ingress Byte Shuffling Table Register	x‘1C’		Read/Write
Ingress Byte Shuffling Table Byte Location Register	x‘1D’		Read/Write
Egress Byte Shuffling Table Register	x‘44’	x‘84’	Read/Write
Egress Byte Shuffling Table Byte Location Register	x‘45’	x‘85’	Read/Write
Flow Control and Packet Scheduling Control Registers			
Ingress Credit Table Access Register	x‘46’	x‘86’	Read/Write
Egress Credit Table Access Register	x‘03’		Read/Write
Ingress Buffer Flow Control High Threshold Registers	x‘04’ to x‘05’		Read/Write
Ingress Buffer Flow Control Low Threshold Registers	x‘07’ to x‘08’		Read/Write
Ingress VOQ Flow Control High Threshold Registers	x‘1E’ to x‘20’		Read/Write
Ingress VOQ Flow Control Low Threshold Registers	x‘21’ to x‘23’		Read/Write
Egress Buffer Flow Control Threshold Registers	x‘0A’ to x‘0B’		Read/Write
Egress Subport Flow Control Threshold Registers (PowerPRS Q-64G only)	x‘0C’ to x‘0D’		Read/Write
Egress Multicast Flow Control Threshold 1 Register (PowerPRS Q-64G only)	x‘0E’		Read/Write
Egress Multicast Flow Control Threshold 2 Register	x‘0F’		Read/Write
Ingress Filter 1 Registers	x‘47’ to x‘48’	x‘87’ to x‘88’	Read/Write
Ingress Filter 2 Registers	x‘49’ to x‘4A’	x‘89’ to x‘8A’	Read/Write
Ingress Filter Command Register	x‘4B’	x‘8B’	Read/Write
Internal Status Registers			
Ingress Queue Status 1 Register	x‘4C’	x‘8C’	Read/Write
Ingress Queue Status 2 Register (PowerPRS Q-64G only)	x‘4D’	x‘8D’	Read/Write
Ingress Queue Status Selection Register	x’4E’	x’8E’	Read/Write
Egress Queue Status Register	x‘4F’	x‘8F’	Read/Write
Local Multicast Table Access Registers			
Multicast Table Access 1 Register	x‘10’		Read/Write
Multicast Table Access 2 Register	x‘11’		Read/Write
Multicast Table Access 3 Register (PowerPRS Q-64G only)	x‘12’		Read/Write
Internal Resource Monitoring Registers			
Ingress Free Buffer List Register	x‘50’	x‘90’	Read/Write
Egress Free Buffer List Register	x‘13’		Read/Write

Table 1. Register Summary (Page 3 of 3)

Register Name	Address		Access
	X Plane	Y Plane	
Ingress Link List Register	x‘51’	x‘91’	Read/Write
Ingress First-Last Table Access Register	x‘52’	x‘92’	Read/Write
Egress Link List 1 Register	x‘14’		Read/Write
Egress Link List 2 Register	x‘15’		Read/Write
Egress First-Last Table Access Register	x‘16’		Read/Write
Ingress Flow Control Register	x‘17’		Read/Write
Egress Flow Control Register	x‘18’		Read/Write
Switch Fabric Environment Status Registers			
Card/Slot ID Register	x‘70’	x‘B0’	Read/Write
Remote Card Availability 1 Register	x‘71’	x‘B1’	Read Only
Remote Card Availability 2 Register (PowerPRS Q-64G only)	x‘72’	x‘B2’	Read Only
Clock Configuration Registers			
Switch Clock PLL Register	x‘73’	x‘B3’	Read/Write
Switch Clock PLL Observe Register	x‘CA’	x‘EA’	Read Only
Local Clock PLL Register	x‘78’		Read/Write
Local Clock PLL Observe Register	x‘1A’		Read Only
Reset and Test Registers			
Reset Control Register	x‘54’	x‘94’	Read/Write
Memory BIST Status Register	x‘26’		Read/Write
Chip ID Register	x‘24’		Read Only
Logic BIST 1 Register	x‘79’		Read/Write
Logic BIST 2 Register	x‘7A’		Read/Write
Test Configuration Register	x‘55’	x‘95’	Read/Write
Internal Hardware Checking Registers			
Event 1 Register	x‘56’	x‘96’	Read Only
Event 1 Mask Register	x‘57’	x‘97’	Read/Write
Event 1 Interrupt Enable Register	x‘5E’	x‘9E’	Read/Write
Event 2 Register	x‘58’	x‘98’	Read Only
Event 2 Mask Register	x‘59’	x‘99’	Read/Write
Event 2 Interrupt Enable Register	x‘5F’	x‘9F’	Read/Write

Electrical Information

Table 2. Absolute Maximum Ratings

Symbol	Parameter	Rating			Units
		Minimum	Typical	Maximum	
V_{DD} (1.8 V)	Power supply voltage		1.8	1.95	V
V_{DD} (1.5 V)	Power supply voltage for HSTL-level signals		1.5	1.65	V
V_{DD} (2.5 V)	Power supply voltage for LVCMOS-level signals		2.5	2.75	V
T_A	Operating junction temperature	0		100	°C
T_S	Storage temperature	-55		150	°C
	Electrostatic discharge		3000		V

Note: Permanent device damage may occur if the above absolute maximum ratings are exceeded. Extended exposure to absolute maximum rating conditions may affect device reliability.

Table 3. Power Requirements

CSIX Driver	Subport	Format	Core Clock (MHz)	Maximum Power (W)	Maximum Current (A)		
					1.5 V	1.8 V	2.5 V
2.5-V LVCMOS at 125 MHz	No	Fixed	130	14.5		7.44	0.44
	Yes	Fixed	130	14.5		7.44	0.44
	Yes	Variable	166	16.6		8.61	0.44
1.5-V HSTL at 250 MHz	No	Fixed	130	16.0	1.53	7.44	0.12
	Yes	Fixed	130	16.0	1.53	7.44	0.12
	Yes	Variable	166	18.2	1.53	8.67	0.12

Figure 4. Pinout (840-ball HyperBGA package, bottom view)

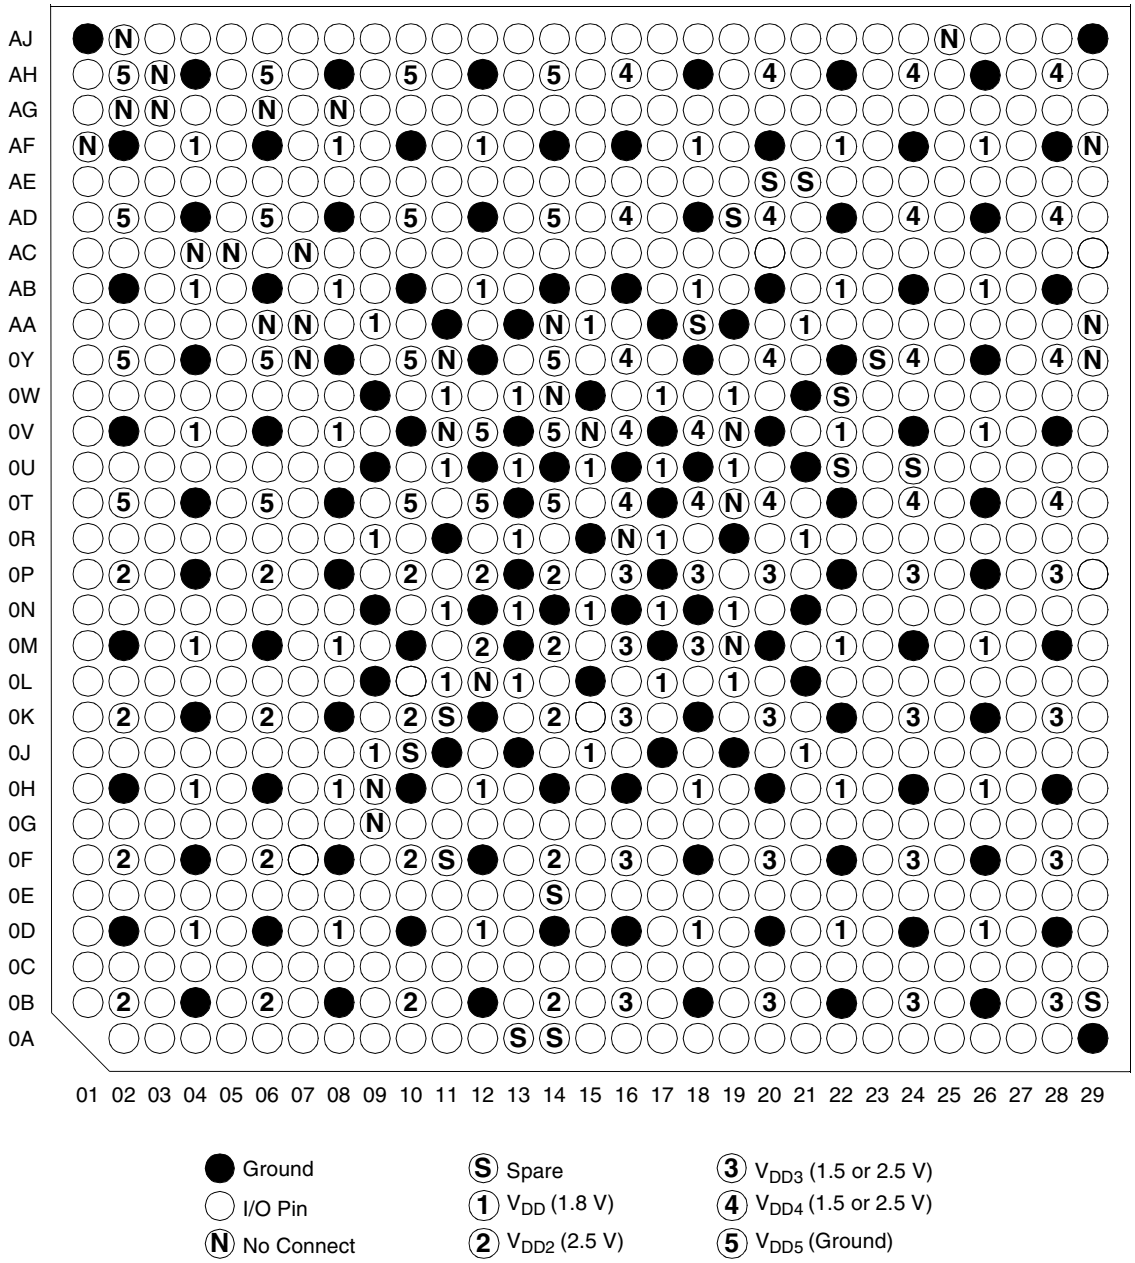

Note: The value of V_{DD3} and V_{DD4} is either 1.5 V for HSTL-level signals or 2.5 V for LVCMOS-level signals.

Note: The “no connect” I/Os must be connected to ground to improve noise immunity. Spare pins must be connected to ground through a via. Spare pins should *not* be connected directly to the card’s ground plane because they may be needed in a future PowerPRS C192 release.

Mechanical Information

Figure 5. Package Mechanical

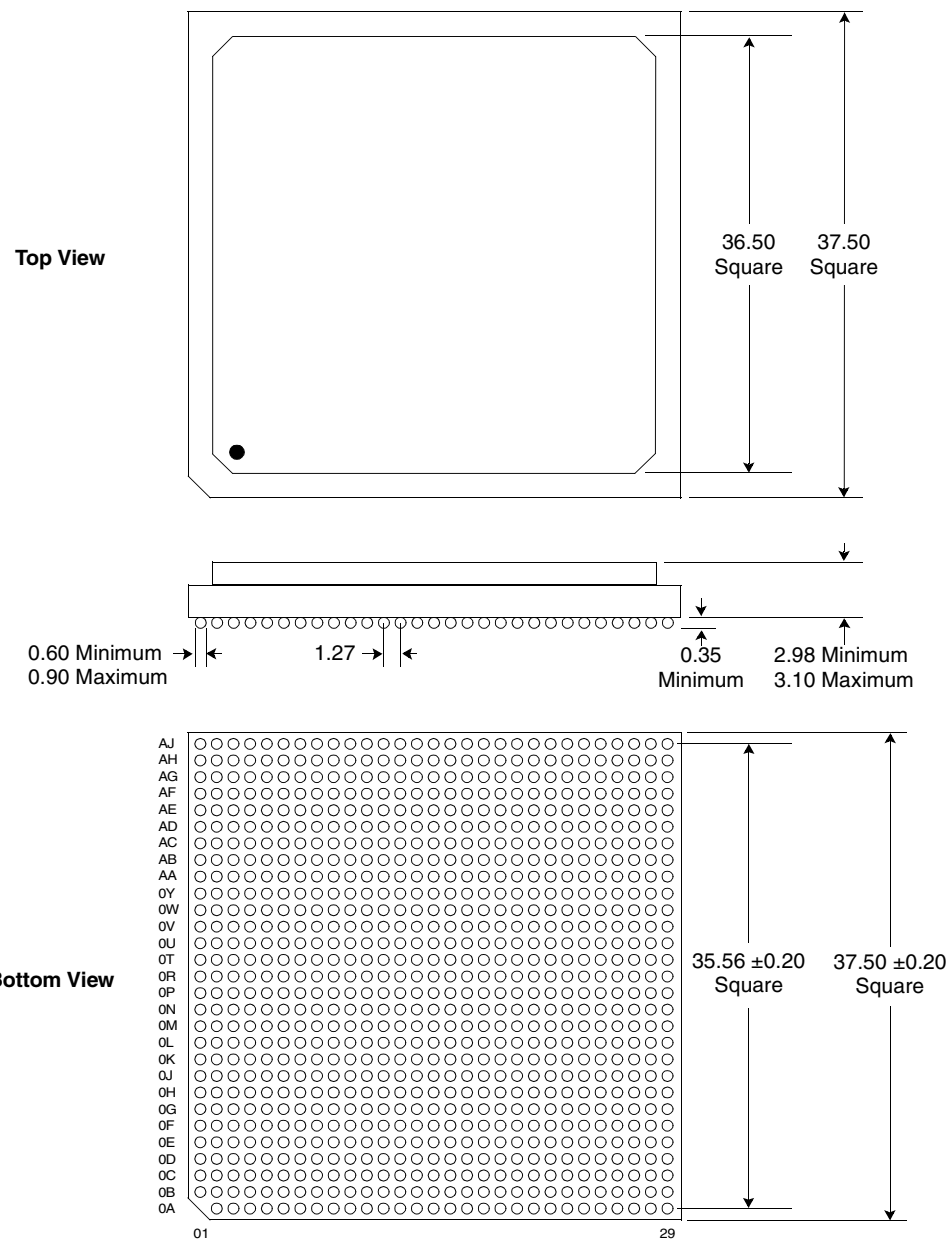

Package type is 840-ball HyperBGA, package size is 37.5 × 37.5 mm, and array size is 29 × 29. All dimensions are in millimeters.

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.

Revision Log

Revision Date	Contents of Modification
Sept. 7, 2001	Initial release (00).

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

HyperBGA	PowerNP
IBM	PowerPRS
IBM Logo	

InfiniBand is a trademark of the InfiniBand Trade Association.

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

C192sds.00.fm
September 7, 2001