General Information

Features

- Companion to the IBM PowerPRS™ 64Gu Packet Routing Switch chip
- CSIX interface attachment to the PowerPRS 64Gu switch core
- CSIX-L1 interface: OC-48 adapter, compliant with the *Common Switch Interface Specification - L1*
- Dual-switch attachment for redundant switch-plane operation, including programmable scheduled switchover (packet lossless) and hot standby switchover
- Redundant switch port attachment at 4 Gbps using 2.5-Gbps serial links compatible with InfiniBand™ physical layer standards
- Supports 8 × 8, 16 × 16, 32 × 32, and 64 × 64 switch fabrics
- PowerPRS 64Gu interface: 16- to 20-byte logical unit (LU) packet processing
- Shared buffer capacity of up to 1024 ingress packets (512 per switch plane) and up to 256 egress packets (shared between switch planes)
- Configurable number of traffic priorities (from one 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 service packets
- Eight-bit parallel processor interface to access all registers for control and error reporting
- Internal loopback support for both the CSIX interface and switch interface
- Internal logic built-in self-test (BIST) and memory BIST
- IEEE® Standard 1149.1 boundary scan to facilitate circuit-board testing
- CMOS 7SF (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 the CSIX-L1 and microprocessor interfaces
- 25-mm, 360-ball ceramic ball grid array (CBGA) package

Description

The IBM PowerPRS C48 Common Switch Interface is a companion device to the IBM PowerPRS 64Gu Packet Routing Switch. It functions as the switch core access layer between the protocol engine's OC-48 CSIX interface and the switch core.

The PowerPRS C48 switch interface is comprised of two 2.5-Gbps high-speed serializer/deserializer (HSS) pairs that provide a total throughput of 5 Gbps. The switch port payload throughput is only 4 Gbps because of the Fibre Channel Standard 8b/10b encoding on the HSS links.

The PowerPRS C48 attaches directly to the HSS-compatible 64Gu. When connected to the PowerPRS 64Gu, the C48 packet length is programmable from 64 to 80 bytes, in 4-byte increments. Ingress and egress packets are divided into four LUs of 16 to 20 bytes each.

The PowerPRS C48 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 C48 hardware-assist functions perform scheduled switchover without packet loss as well as asynchronous (or hot standby) switchover.

The PowerPRS C48 processes data traffic using one to four priorities, depending on register configuration. 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. When PowerPRS C48 ingress traffic congestion occurs, flow control information is propagated through the CSIX interface to the protocol engine on the egress path. When protocol engine egress traffic congestion occurs, flow control information is transmitted in band to the PowerPRS C48 CSIX interface according to the *CSIX-L1 Specification*. The PowerPRS C48 also features an optional out-of-band flow control mechanism. Activation of either the in-band or out-of-band flow control mechanism is selected during PowerPRS C48 configuration.

Ordering Information

Part Number	Description	Throughput
IBM3247P4448	IBM Common Switch Interface	2.5 Gbps

Architecture

Figure 1 illustrates the integration of a PowerPRS C48 in a 32-port 64Gu redundant switching system. An integrated high-speed SerDes (HSS) interface enables the PowerPRS C48 to directly connect to the 64Gu switch core. The PowerPRS C48 supports the 64Gu packet length of 64, 72, or 80 bytes. When attached to a PowerPRS 64Gu master/slave pair, the C48 is used to build a 32-port switching system. With a 64-port destination capability, the PowerPRS C48 will be able to accommodate next-generation PowerPRS switching systems.

Figure 1. System View of the PowerPRS C48 with the 64Gu (configured with redundant 128-Gbps switch planes)

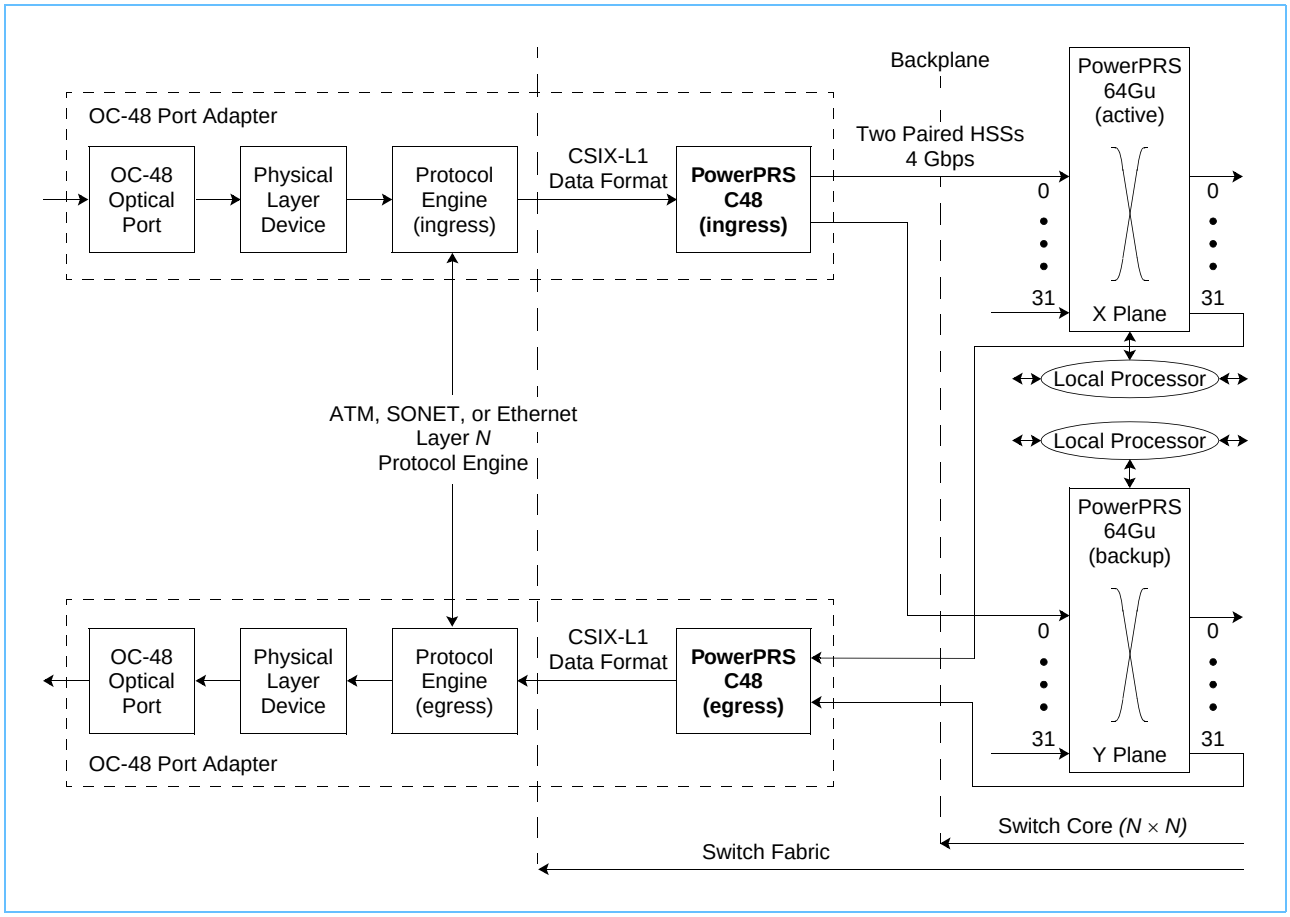

The attachment of a PowerPRS C48 to a redundant switch core allows the balancing of traffic loads between two switch planes. In addition, the PowerPRS C48 and the 64Gu can jointly execute scheduled switchover without packet loss.

The PowerPRS C48 CSIX-L1 interface enables direct attachment to a 2.5-Gbps protocol engine (32 bits wide). PowerPRS C48 ingress *and* egress data paths provide end-to-end flow control as well as access to a 64-K entry multicast table. The PowerPRS C48 multicast table is updated by either the switch local processor (in 64Gu applications) or the C48 eight-bit parallel processor.

The internal structure of the PowerPRS C48 is presented in *Figure 2* on page 4.

Figure 2. PowerPRS C48 Block Diagram

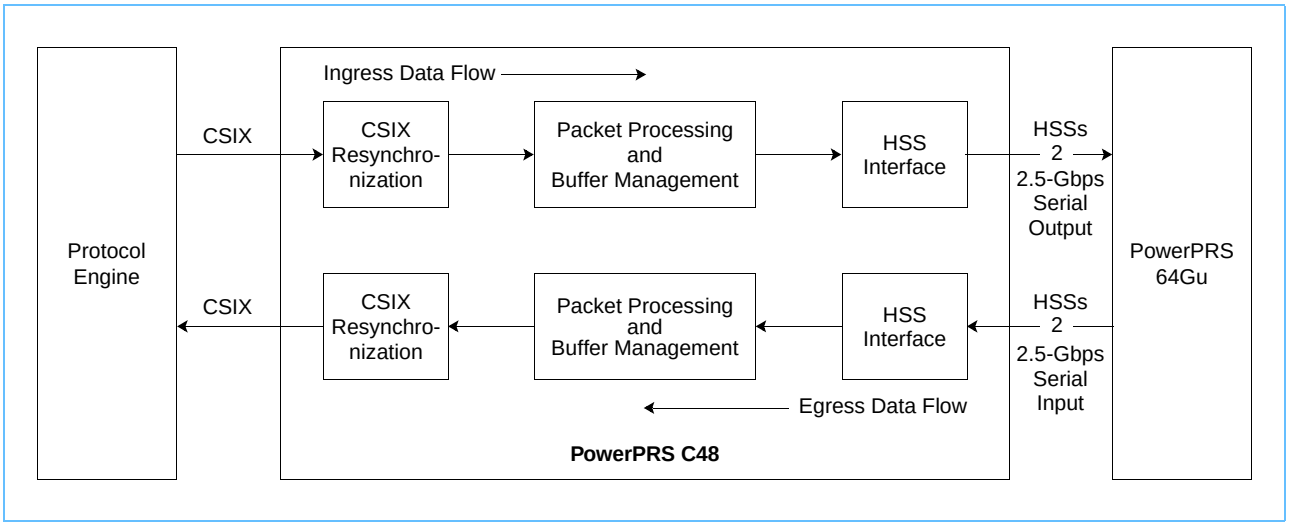

Ingress Data Flow

Ingress packets received from the protocol engine are color-coded and queued for transmission on either the X or Y switch path. The PowerPRS C48 ingress buffer stores up to 1024 incoming packets (512 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 can 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 eight output queues (4 priorities $\times$ 2 switch planes) for transmission to the CSIX interface. The PowerPRS C48 egress buffer is shared between the X and Y switch paths, and can store up to 256 outgoing packets. The egress queue status, packet priority, and CSIX port destination flow control information are reported to the egress scheduler. The egress scheduler uses a flywheel mechanism to select the next packet for transmission to the CSIX interface. The highest-priority packets of granted destinations are generally transmitted first, unless the egress credit table has been programmed to guarantee minimum bandwidth to low-priority packets.

Programming Interface and Registers

The PowerPRS C48 employs an eight-bit parallel processor programming interface. This interface provides read/write access to all PowerPRS C48 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 C48 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
HSS Synchronization 1 Register	x‘C2’	x‘E2’	Read Only
HSS 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
HSS 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 Register	x‘42’	x‘82’	Read/Write
HSS Debug Control Register	x‘75’	x‘B5’	Read/Write
HSS Test Register	x‘76’	x‘B6’	Read/Write
HSS Error 1 Register	x‘C7’	x‘E7’	Read/Write
HSS 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
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 64Gu 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 64Gu 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
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

Table 1. Register Summary (Page 3 of 3)

Register Name	Address		Access
	X Plane	Y Plane	
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 64Gu 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
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} (2.5 V)	Power supply voltage for LVCMOS-level signals		2.5	2.75	V
T_A	Operating ambient temperature	-40		100	°C
T_J	Operating junction temperature	0		125	°C
T_S	Storage temperature	-65		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. Recommended Operating Conditions

Symbol	Parameter	Rating			Units
		Minimum	Typical	Maximum	
V_{DD} (1.8 V)	Power supply voltage	1.71	1.8	1.89	V
V_{DD} (2.5 V)	Power supply voltage for LVCMOS-level signals	2.375	2.5	2.625	V

Table 4. Total Power Requirements

Core Clock Frequency (MHz)	Power (W)		Current (A)			
			1.8 V		2.5 V (LVCMOS)	
	Typical	Maximum	Typical	Maximum	Typical	Maximum
166	10.3	12.0	4.95	5.76	0.52	0.60

Table 5. Thermal Performance

Thermal Resistance (°C/W)		Thermal Resistance θ_{JA} (°C/W) at Air Flow Rate					
θ_{JC}	θ_{JB}	0 LFPM	100 LFPM	200 LFPM	300 LFPM	400 LFPM	600 LFPM
0.51	3.5	14.1	12.7	11.5	10.4	9.6	8.5

Notes: θ_{JA} = Junction-to-ambient thermal resistance
θ_{JB} = Junction-to-board thermal resistance
θ_{JC} = Junction-to-case thermal resistance
LFPM = Linear feet per minute

Figure 3. Pinout (360-ball CBGA package, top view)

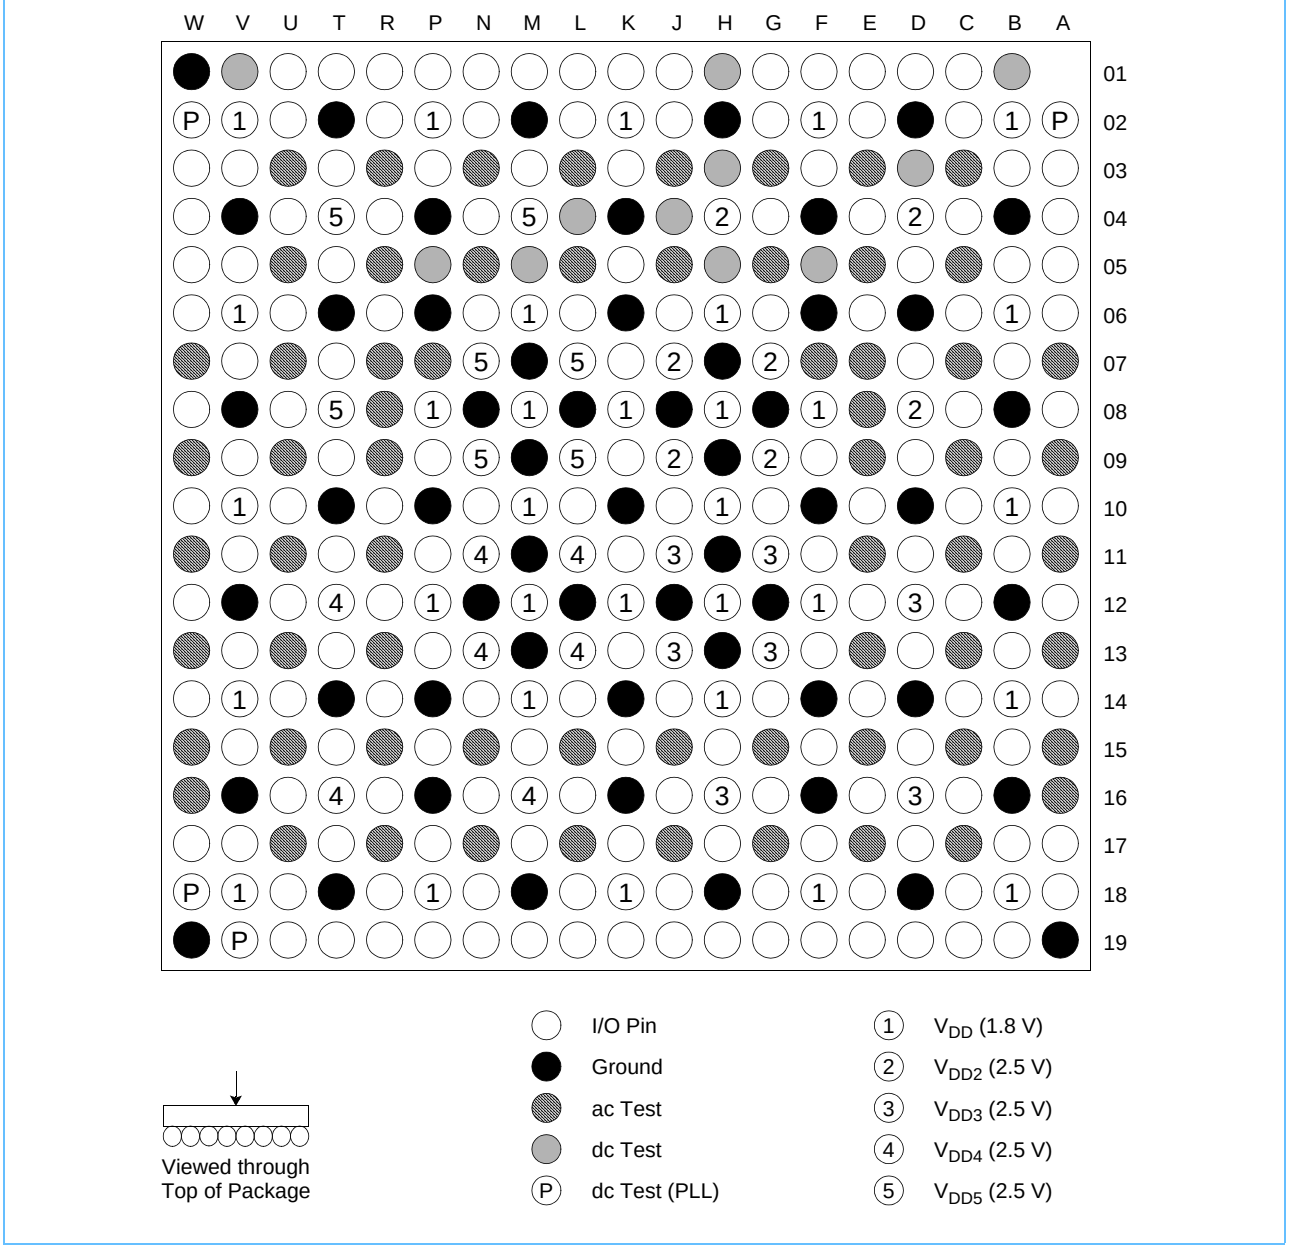

Mechanical Information

Figure 4. Package Mechanical

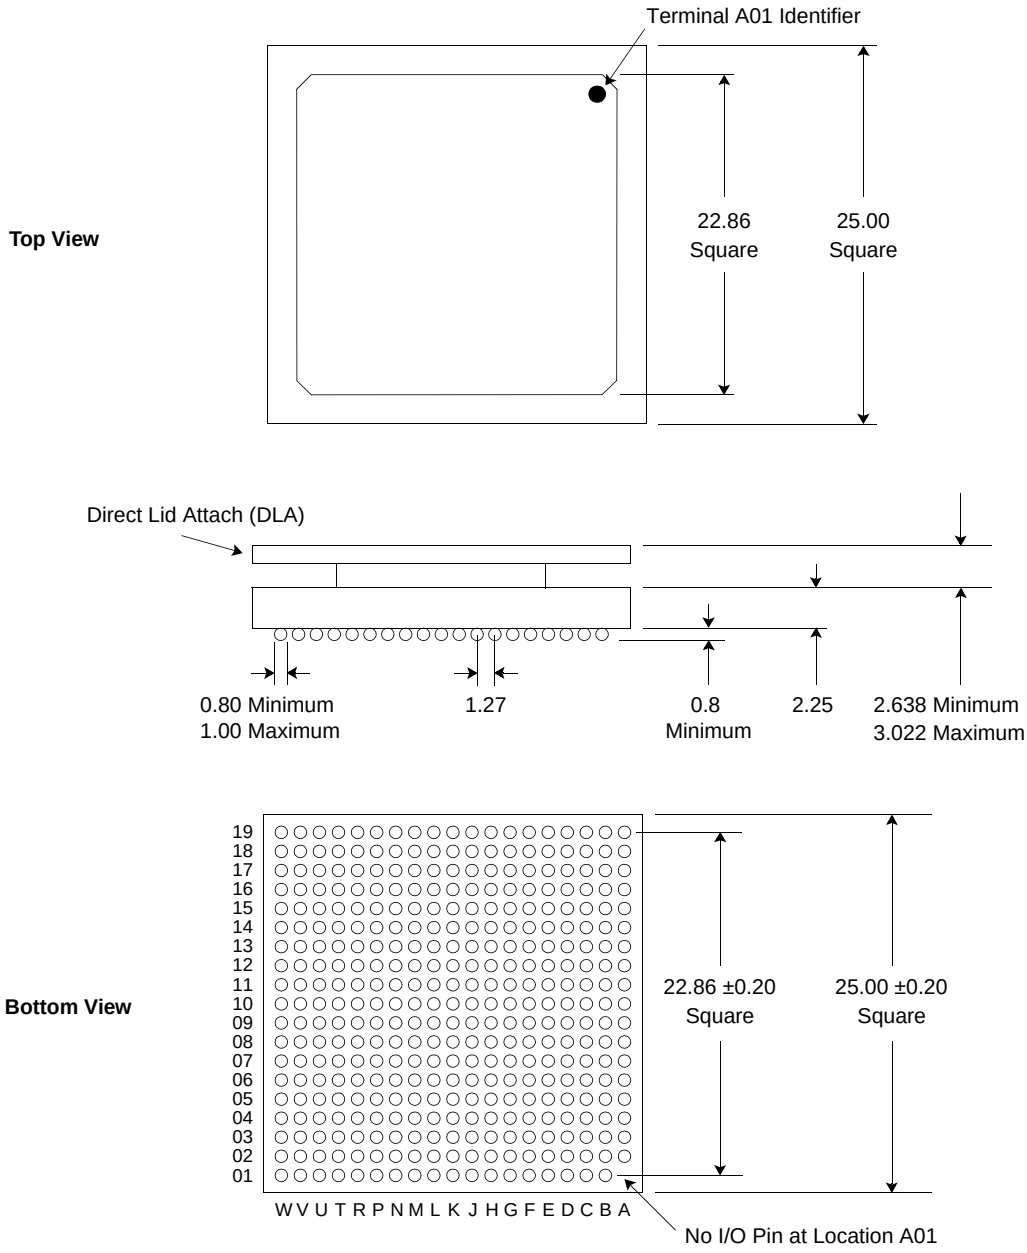

Package type is 360-ball CBGA with DLA, package size is 25 × 25 mm, and array size is 19 × 19. All dimensions are in millimeters.

Note: This document contains information on products in the design, 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. 9, 2002	Initial release (00).

Printed in the United States of America September 2002

Note: This document contains information on products in the design, 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.

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.ibm.com/chips

prsC48sds.00.fm
September 9, 2002