

PM2329

ClassiPI

Datasheet.Live

Network Classification Processor

Datasheet

Proprietary and Confidential

Issue 4, November 2001

Copyright 2001 PMC-Sierra, Inc. All rights reserved

PMC-Sierra

8555 Baxter Place

Burnaby, BC Canada V5A4V7

Phone 604.415.6000, Fax 604.415.6200

The information is proprietary and confidential to PMC-Sierra, Inc., and for its customers' internal use. In any event, no part of this document may be reproduced in any form without the express consent of PMC-Sierra, Inc.

Document ID: PMC-2010146, Issue 4

Legal Information

Copyright

Copyright 2001 PMC-Sierra, Inc.

The information is proprietary and confidential to PMC-Sierra, Inc., and for its customers' internal use. In any event, you cannot reproduce any part of this document, in any form, without the express written consent of PMC-Sierra, Inc.

PMC-2010146 (R4)

Disclaimer

None of the information contained in this document constitutes an express or implied warranty by PMC-Sierra, Inc. as to the sufficiency, fitness or suitability for a particular purpose of any such information or the fitness, or suitability for a particular purpose, merchantability, performance, compatibility with other parts or systems, of any of the products of PMC-Sierra, Inc., or any portion thereof, referred to in this document. PMC-Sierra, Inc. expressly disclaims all representations and warranties of any kind regarding the contents or use of the information, including, but not limited to, express and implied warranties of accuracy, completeness, merchantability, fitness for a particular use, or non-infringement.

In no event will PMC-Sierra, Inc. be liable for any direct, indirect, special, incidental or consequential damages, including, but not limited to, lost profits, lost business or lost data resulting from any use of or reliance upon the information, whether or not PMC-Sierra, Inc. has been advised of the possibility of such damage.

Patents

Relevant patent applications and other patents may also exist.

Contacting PMC-Sierra

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

Tel: (604) 415-6000

Fax: (604) 415-6200

Document Information: document@pmc-sierra.com

Corporate Information: info@pmc-sierra.com

Technical Support: apps@pmc-sierra.com

Web Site: <http://www.pmc-sierra.com>

Contents

1	PM2329 Overview	14
1.1	Introduction	14
1.2	Features	14
1.3	Functional Overview	15
1.4	Architectural Overview	16
1.4.1	Algorithmic Support	16
1.4.2	Architectural Support	17
1.4.3	Software Model Support	17
1.4.4	Policy Search Engine PSE	18
1.4.5	Field Extraction Engine FEE	18
1.4.6	Operation Control Logic	18
1.4.6.1	Operation Cycles, Descriptors and E-RAM	19
1.5	System Diagram	19
2	Signal Description	20
2.1	Introduction	20
2.2	Signals Listed by Function	23
2.3	Signals Listed by Ball	33
2.4	Interface Description	41
2.4.1	System Interface	41
2.4.1.1	Clock Frequency	41
2.4.1.2	Processor Bus Cycles	42
2.4.1.2.1	SyncBurst Bus Cycles	42
2.4.1.2.2	ZBT Bus Cycles	44
2.4.1.3	Packet Source Interface Signals	45
2.4.1.4	64-Bit Mode	48
2.4.1.5	32-Bit Mode	49
2.4.1.6	Byte Ordering	49
2.4.1.7	Cascade Mode Addressing	49
2.4.1.8	Local and Global Register Access	49
2.4.1.9	Multiple Context Support	50
2.4.1.10	Reset and Interrupts	50
2.4.2	Extended RAM (E-RAM) Interface	51
2.4.3	Cascade Interface	53
2.5	System Configurations	53
2.5.1	Stand Alone Configuration	53
2.5.2	Cascaded Configuration	53
2.5.3	Operation with Extended RAM	54
3	Functional Description	57
3.1	Internal Organization and Data Flow	57

3.1.1	Basic Blocks	58
3.1.2	Data Flow	59
3.1.3	Context Support	61
3.1.4	Channel Input, Output and Status Mechanism	62
3.2	Field Extraction Engine (FEE)	63
3.2.1	Introduction	63
3.2.2	Supported Packet Formats	64
3.3	Control Unit	66
3.4	Policy Search Engine (PSE)	66
3.4.1	Rule Memory	67
3.4.2	Cell Organization	67
3.4.3	Priority of Rules	68
3.4.4	Rule Partitions	68
3.5	E-RAM Operation	69
3.5.1	Organization of E-RAM Words	69
3.6	Cascade Operation	69
4	Registers	71
4.1	PM2329 Access Modes	71
4.1.1	Address Space	71
4.1.2	Channels	72
4.1.3	Channel Register Blocks	73
4.1.4	Direct and Indirect Access	73
4.2	Register Interface	74
4.2.1	Programmable Register Overview	74
4.2.2	Register Description	79
4.2.2.1	Local Configuration Register (LCR; n000h)	80
4.2.2.2	Rule Indirect Command Register (RICR; n008h)	82
4.2.2.3	Rule Indirect Address Register (RIAR; n010h)	84
4.2.2.4	Rule Indirect Data Register Set (RIDR0; n018h) (RIDR2; n020h) (RIDR4; n028h)	85
4.2.2.5	OC Descriptors (OCD; n400h, n408h...n7F0h, n7F8h)	88
4.2.2.6	E-RAM Indirect Data Register Set (EIDR0; 8200h) (EIDR2; 8208h) (EIDR2; 8208h) (EIDR4; 8210h) (EIDR6; 8218h)	91
4.2.2.7	E-RAM Indirect Command Register (EICR; 8220h)	94
4.2.2.8	E-RAM Indirect Address Register (EIAR; 8228h)	96
4.2.2.9	E-RAM Configuration Register (ECR; 8230h)	97
4.2.2.10	Interrupt Enable Register (IER; 8238h)	99

4.2.2.11	Status Register (STSR; 8240h)	100
4.2.2.12	Operation Control Register (OPCR; 8248h)	102
4.2.2.13	Channel Assignment Register (CAR; 8250h)	104
4.2.2.14	OC Conductor Register (OCCR; 8258h)	105
4.2.2.15	Packet Information Register (PIR; 8260h)	107
4.2.2.16	Timer Register (TMR; 8268h)	111
4.2.2.17	Alternate OCC Register (AOCC; 8270h)	113
4.2.2.18	Packet Buffer Input Register (PBIR; Base 3 +00h, +08h,... +0E8h, +0F0h) (PBIR, EOPD0; Base 3 +0F8h) (PBIR, EOPD1; Base 0 +08h)	115
4.2.2.19	Channel Status Register (CSR; Base 0 +00h)	118
4.2.2.20	OC Results FIFO Output Reg (OCRf; Base 1 +00h)	120
4.2.2.21	Data Results FIFO Output Register	125
4.3	Indirectly Addressable Locations	126
4.3.1	Rule Memory Cells	126
4.3.2	E-RAM Words	127
4.4	Register Summary	130
5	Rule Formats and OC Sequencing	135
5.1	Rule Formats	135
5.1.0.1	Rule Data Field	136
5.1.0.2	Rule Control Field	136
5.1.1	Rule Operations	138
5.1.2	Masking	138
5.1.3	Composite Rules	139
5.1.4	Rule Attribute Bit	140
5.1.5	Rule Negation	140
5.2	Operation Cycles	140
5.3	OC Sequencing	141
5.3.1	OC Conductor	141
5.3.2	OCC Sequencing	142
5.3.2.1	Trace Feature	143
5.3.2.2	Sequencing Control Modes	144
5.3.2.2.1	Automated Sequencing	144
5.3.2.2.2	Processor Controlled Sequencing	145
5.3.2.3	OCC Sequencing	146
5.3.2.4	E-RAM Sequencing	147
5.3.3	E-Word Association with Cells	150
6	Electrical and Timing Characteristics	152
6.1	DC Characteristics	153

6.2	Switching Characteristics	154
6.2.1	Reset Timing Parameters	154
6.2.1.1	VDD Power On Sequence	154
6.2.2	Clock Timing Parameters	155
6.2.3	System Interface Timing	156
6.2.4	E-RAM Interface Timing	159
6.2.5	Cascade Interface Timing	160
6.2.6	JTAG Interface Timing	161
7	Package Details	163
7.1	Package Type, Characteristics and Mechanical Drawing	163
7.1.1	Package Type	163
7.1.2	Thermal Characteristics	163
7.1.3	Package Mechanical Drawing	165

List of Figures

Figure 1	System Block Diagram	19
Figure 2	PM2329 External Signals	21
Figure 3	PM2329 Ball Number Assignments	22
Figure 4	On-Chip PLL Bypass Mechanism	42
Figure 5	System Interface Register Timing Diagram (SyncBurst)	43
Figure 6	System Interface Register Timing Diagram (ZBT)	45
Figure 7	System Interface DMA Timing (SyncBurst Mode)	47
Figure 8	System Interface DMA Timing (ZBT Mode)	48
Figure 9	E-RAM Interface Timing	52
Figure 10	E-RAM Connectivity	52
Figure 11	Cascaded Bus Connectivity	54
Figure 12	Single PM2329 with 32-bit/64-bit E-RAM	55
Figure 13	Two Cascaded PM2329 devices with 64 or 96 bits of ERAM	56
Figure 14	Maximum Cascade Configuration	56
Figure 15	PM2329 Block Diagram	57
Figure 16	Simplified Data Flow	59
Figure 17	Simplified Data Flow--OC Processing and Results Posting	61
Figure 18	Packet Input, Result Output, and associated Status Handshake	63
Figure 19	Packet Formats Supported by PM2329	65
Figure 20	Headers Formats within Packet	66
Figure 21	Organization of Rules	68
Figure 22	Local vs. Global Register Space; Conceptual View A	71
Figure 23	PM2329 Packet Input Buffer	72
Figure 24	Channel Register Blocks	73
Figure 25	PM2329 Address Space	76
Figure 26	Access to Rule Memory Cells via RIDR0-4 Registers	82
Figure 27	Timer Logic	112
Figure 28	Rule Control and Data Field	135
Figure 29	Rule Control Fields	137
Figure 30	Trace Feature	144
Figure 31	Processor controlled Sequencing	145
Figure 32	General Overview of OCC Sequencing	147
Figure 33	General Overview of ERAM Sequencing	150
Figure 34	Recommended VDD Power On Sequence	154
Figure 35	SCLK to ECLKOUT Skew	155
Figure 36	System Interface Timing Parameters	156
Figure 37	Load Equivalent	158
Figure 38	E-RAM Interface Timing	159
Figure 39	Cascade Interface Timing	160
Figure 40	JTAG Interface Timing	161

Figure 41	JTAG IDCode Register	162
Figure 42	Device Compact Model	163
Figure 43	Package Mechanical Drawing	165

List of Tables

Table 1	Timing and Common Control Signals	23
Table 2	System Interface Signals	26
Table 3	ERAM Interface Signals	29
Table 4	Cascade Interface Signals	30
Table 5	Test Signals	31
Table 6	VDD and VSS Signals	32
Table 7	Signals Listed by Ball Assignment; Rows A through C	33
Table 8	Signals listed by Ball Assignment; Rows D through K	34
Table 9	Signals listed by Ball Assignment; Rows L through W	35
Table 10	Signals listed by Ball Assignment; Rows Y through AD	36
Table 11	Signals listed by Ball Assignment; Rows AE and AF	37
Table 12	Signals Listed by Name (Alphabetically)	38
Table 13	PSPBA Deassertion Delay	46
Table 14	System Bus 64-bit	48
Table 15	System Bus 32-bit	49
Table 16	Cascade Size vs. Maximum Physical E-RAM Width	54
Table 17	PM2329 Register Memory Map	74
Table 18	Channel Register Block Base Addresses	77
Table 19	Channel Registers	78
Table 20	E-Word Depth Chart	98
Table 21	EMA[18:17] Usage	99
Table 22	Direction Specifier Bit	103
Table 23	OC Conductor Register format	105
Table 24	Processor Controlled OC Sequencing & Trace OC Execution	109
Table 25	Timestamp Increment Interval Example (SCLK 66.67 MHz)	111
Table 26	Data Results FIFO Output Register (64-bit mode)	125
Table 27	Data Results FIFO Output Register (32-bit mode)	126
Table 28	Common Control Rule - CCR Bits	136
Table 29	Operations Supported for Rule Data Sub-fields	138
Table 30	Masked sub-field and Associated Mask Source	139
Table 31	OCC Sequencing Control Bits	142
Table 32	OCC Sequencing Status Bits	142
Table 33	Registers Applicable to OCC Sequencing	142
Table 34	Absolute Maximum Ratings	152
Table 35	Recommended Operating Conditions	152
Table 36	Terminal Capacitance	152
Table 37	DC Characteristics	153
Table 38	Reset Timing	154
Table 39	Clock Timing	155
Table 40	System Interface Timing Parameters	157

Table 41	E-RAM Interface Timing Parameters	160
Table 42	Cascade Interface Timing Parameters	161
Table 43	JTAG Interface Timing Parameters	162
Table 44	Thermal Characteristics	164

1 PM2329 Overview

1.1 Introduction

The PM2329 is a member of the ClassiPI family of sophisticated Network Classification Processors capable of supporting Gigabit/OC-48 interfaces. It is optimized for network environments--network equipment can use the PM2329's classification and analysis capability to implement wire-speed routing, QoS, firewall and other functionality such as NAT and network monitoring that requires packet inspection and classification. The PM2329 can be used to implement a mixed L2, L3, L4 and payload data (L5 to L7) based search. With a peak throughput of up to OC-48 IPv4 packets per second, the PM2329 is an ideal choice for all classification requirements.

The PM2329's patented architecture has been designed to work efficiently in a wide range of system configurations such as per-port, customized hardware, gigabit environment, or a centralized search engine shared by a number of ports. The PM2329 architecture minimizes the bandwidth and latency incurred due to multiple packet data transfers within the equipment. It can be programmed to perform multiple pattern searches sequentially, and/or conditionally without host processor intervention, providing the high throughput required for complex classification applications.

1.2 Features

- Packet Header, Packet Data based or user-defined data based classification at Gigabit wire-speed.
- Single PM2329 can store up to 16K policy rules.
- Up to eight PM2329 devices can be cascaded to appear as a large PM2329 with support for up to 128K rules.
- Supports external Synchronous RAM to extend capabilities such as programmed operation cycle sequencing, per rule statistical information, and aging.
- On-chip Policy Database or Rule Memory that can be partitioned to implement multiple classification partitions.
- Mechanism to allow support for up to 32 independent tasks running on the external processor.
- Supports multiple classification lookups per input request. Each lookup extracts its unique key from input data stream and applies its dedicated classification rule table to produce match results. Classification lookups can be invoked sequentially and/or conditionally using sophisticated sequencing mechanisms.
- Register, External RAM, or Processor controlled lookup sequencing mechanism
- Powerful assist for Routing, QoS, NAT, Firewall and Load Balancing applications.
- High-speed synchronous packet data input and result output interface.

1.3 Functional Overview

The PM2329 is designed for use in switches, routers, access concentrators and other telecommunications equipment such as traffic shapers, firewalls and network address translators. These equipment implement one or more of software modules in the PM2329, such as the Routing, QoS, NAT, Firewall, Load Balancing and Network monitoring modules. These modules implement functions such as address lookup, flow classification, and connections cache identification. The PM2329 accelerates these functions, extending the performance of such equipment in Gigabit/OC-48 wirespeed environments.

The PM2329 is a high performance search and classification processor. During initialization the external packet processor and application code will configure the description of various searches that would be performed by the PM2329. These descriptions consist of key extraction procedures, rule table identities, search methodologies and termination criteria. The external packet processor will then load the PM2329's on-chip rule memory with a set of initial rules for every search. These rules can also be updated at a later time. Once the PM2329 has been configured, the external packet processor can submit packet data together with classification requests. Classification requests can trigger multiple searches to be performed on the same packet. Multiple searches are typically useful when multiple applications are supported on the network equipment - such as routing, QoS, NAT etc. For example, a route lookup table can be stored in one table, while another table can store a QoS flow classification policy. The searches can be specified to be conditional upon the results of the previous searches. The ability to partition the rule memory combined with the ability to sequence multiple searches enable the system designer to offload complex packet processing tasks to the PM2329.

The PM2329 returns a series of results for every packet it analyzes. These can be used by software either directly or to index into a user data structure. The PM2329 provides the capability of attaching an external SRAM, which can store user programmable data corresponding to the indexes returned upon a match. In addition to searches, the PM2329 can gather statistics--counts and time stamps--against every rule match. The PM2329 gathers these statistics in the external SRAM.

Before performing a search operation on a packet, the PM2329 performs key extraction. The PM2329 has an on-chip Field Extraction Engine that can extract Layer 3 and Layer 4 header fields from IPv4 packets as the packet data is presented to the chip. Using these fields, it constructs a header key that be used during searches. Further, it can also extract keys (short and long) at arbitrary offsets into the input data stream. Each extracted key is made up of a number of fields. Every rule in the rule tables specifies values and conditions for all fields in a key.

The PM2329 search engine can perform tasks such as:

- Single or multiple match identification
- Prioritized match selection on multiple match
- Layer-N searches: patterns and signature composed of strings, numbers, etc.
- Longest prefix (LP) search for IP Routing

1.4 Architectural Overview

The PM2329 architecture supports the requirements of networking protocols pertaining to high-speed search and classification. It supports the implementation of several protocols in the same equipment to achieve wire-speed performance at Gigabit/OC-48 rates. This architecture permits sequential, parallel and conditional operations. The PM2329 device can also be cascaded to allow searches consisting of a large number of rules. The PM2329 architecture is tuned to high-speed yet sophisticated classification operation.

Networking protocol processing requires performing searches and lookups in data structures based on information contained in the header of a packet. In IPv4 packets, this is typically the Source and Destination IP address (Layer 3), the IP Protocol Field, and the Source and Destination Port Numbers (Layer 4). A lookup operation would consist of extracting the header of a packet and conducting a search in a data structure using this and some specific search criterion. As the sophistication of network protocols has increased, data lookups using information contained in the data payload of the packet (Layer 5 to Layer 7) has also become necessary. The PM2329 supports high-speed lookups using this (Layer 5 to Layer 7) data.

To perform these lookups at wirespeed packet rates, the accelerator engine must have algorithmic support, architectural support and software model support. The PM2329 provides this support in an optimized manner as outlined below.

1.4.1 Algorithmic Support

The PM2329 provides algorithmic support for searches and examination of packet contents. Based on the information provided by the PM2329, the associated packet content manipulation is performed by an external Packet Processor.

Key capabilities of the PM2329 for algorithmic support are as listed below:

- Flexible key extraction
- Search Prioritization
- Single or Multiple Match identification
- Range Searches
- Counters/Statistics (Packet and Byte Count) and Timestamps
- Longest Prefix Search
- Aging support

1.4.2 Architectural Support

The PM2329 architecture supports the following to improve system performance:

- Scalable Policy Search Engine can be scaled up to accommodate a wide variation in the number of rules as well as number of search partitions (to store multiple search rules).
- Supports search sequencing to minimize Packet Processor or CPU intervention.
- Supports complex field extraction from input packet for further classification.
- Lookup/Search model that is easy to integrate into software implementations of protocols.
- Supports statistics collection.

1.4.3 Software Model Support

From the software programmers perspective, the PM2329 presents a flexible Lookup/Search model. It implements a generalized Lookup/Search operation, which is atomic in nature. This operation has two main inputs: the key data using which the lookup is to be performed and the set of policy rules, which have to be applied to the key data to perform this lookup. The former is selected from the packet and can be the packet header L3/L4 information, other fields from the packet header or fields from the packet data payload. The latter is a set of rules, which are to be applied to this data to obtain a result. The result is in the form of an occurrence of match, in other words, the identification of a rule, which the input data satisfies. Since it is possible for more than one rule to match the input data, the rules are prioritized. In case of multiple matches, multiple results can be returned and the higher priority matching rule is presented first.

Multiple such Lookup/Search functions can be stored within a single PM2329. Each such function is said to occupy a partition within the chip. For example, a route lookup table can be stored in one partition, while another partition can store a QoS flow classification policy. The PM2329 also has the ability to perform a sequence of different lookups or classifications on the same packet data.

The Search/Lookup operations programmed into the PM2329 can be automatically applied to every packet presented to the PM2329's high-speed wide synchronous data transfer bus. Thus, the sequence of operations is performed on every packet and the result is made available, on a per packet basis, to the Packet Processor.

To assist in high-speed packet classification, the PM2329 is designed to be Ethernet and IPv4-aware. It incorporates a Field Extraction Engine (FEE) that understands Ethernet II, 802.3 & 802.1p/q and determines where the Layer 3 payload starts from in the packet. It can extract key IP, TCP and UDP header fields as well as payload data at any offset. The FEE is tolerant of the various idiosyncrasies of the IP and TCP headers. It is also possible to bypass the FEE on a per-packet basis.

Operating at a high clock rate, the 64-bit wide synchronous bus provides adequate bandwidth to support Gigabit/OC-48 rate transfer of data, commands and results. The PM2329's flexible data interface architecture allows complete hardware controlled data transfer for high-speed applications, as well as a processor driven software controlled data transfer for systems where performance is less critical. The PM2329 can be configured to receive streaming commands and data through its data interface. Data can consist of arbitrary bytes, partial IP packets (IP header, TCP header and partial data), or complete IP packets. In all cases the data could be optionally preceded by command words that identify the nature of data following immediately as well as the kinds of classification to be performed on the data. This method

of operation ensures high-speed command and data input.

Architecturally the PM2329 consists of three major functional blocks:

- Policy Search Engine PSE
- Field Extraction Engine FEE
- Operation Control Logic: Operation Cycles, Descriptors and E-RAM

Each of these blocks is introduced in the following sections.

1.4.4 Policy Search Engine PSE

Operating at up to 232 MHz, the PM2329 Policy Search Engine performs multiple search operations on input data against a set of pre-loaded rules. Lookup criteria or Search policies called rules are stored in the Policy Database or Rule Memory. A rule is made up of an operation code and corresponding operand data. These operation codes and operand data specify match criteria for multiple fields that make up the key data. Rules related to a single Search/Lookup operation are stored in a Partition within the Rule memory.

A single search/lookup operation is termed an Operation Cycle (OC). It consists of applying the rules within a specified partition to the key data extracted from the packet by the FEE. The result of an OC is in the form of identification of the rule that caused a Match with the data presented by the FEE to the Policy Search Engine. More complex rules can be created with the help of the Composite rule facility using which up to four consecutive instructions can be combined to form a composite instruction.

1.4.5 Field Extraction Engine FEE

The FEE performs IP, TCP and UDP header analysis of every packet presented to the PM2329. It extracts relevant Layer 3 addresses and Layer 4 port numbers together with protocol type from these headers to form a header key. The FEE can also extract arbitrary data at any offset in the input data stream, thereby enabling classification based on the packet data other than Layer 3 and Layer 4 header fields. The FEE can also maintain some amount of TCP state information in the E-RAM if so desired. The FEE is Ethernet (II, 802.3 & 802.1p/q) aware and hence it can automatically identify offsets for the Layer 3 header. Likewise it can correctly locate Layer 4 header offsets by correct interpretation of the IP options fields.

For applications where it is more efficient to carry out the IP/TCP/UDP header extraction externally, IP/TCP/UDP header processing can be disabled in the FEE.

1.4.6 Operation Control Logic

The Operation Control Logic orchestrates the primary search and classification operations of the PM2329 device.

1.4.6.1 Operation Cycles, Descriptors and E-RAM

The Operation Control Logic controls the execution of the operation cycles. An OC is the atomic action of looking up a packet against a class of rules and returning a match index. In order to run an Operation Cycle, the rules that will participate in the OC must be specified. This information is provided by using an Operation Cycle Descriptor (OCD). Each descriptor stores adequate information to describe the start and end of a rule partition. Rule partitions can be uniquely described in each cascaded PM2329 device. The OCDs are arranged in the form of a table within each device, and are specified by an index into this table.

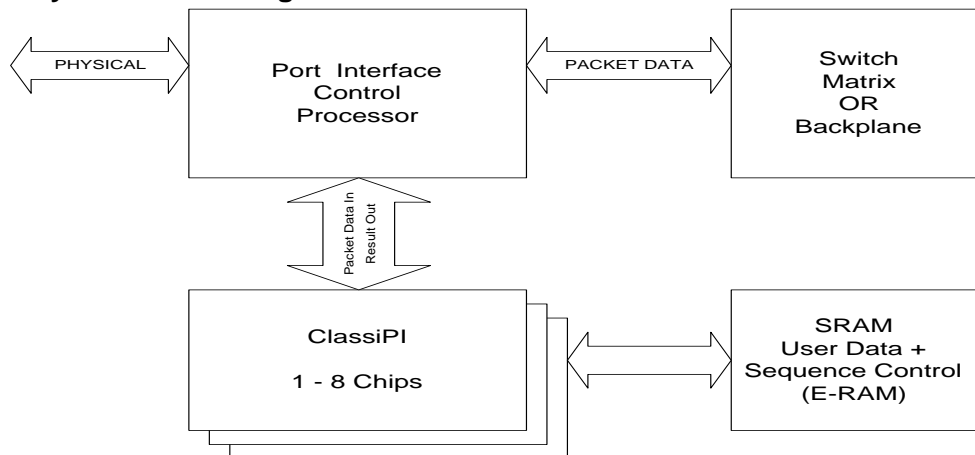
In the minimal mode, a standalone PM2329 device can execute a sequence of up to 4 fixed operation cycles.

To support enhanced functionality including conditional sequencing and user data storage, the PM2329 supports an external synchronous SRAM (E-RAM) that contains control and data information. The control information is stored in Control Words (C-Words). The Operation Control Logic can execute operation cycles based on a flow control mechanism driven by these C-Words stored in the external SRAM. Each control word contains an index into the OC Descriptor Table and identifies a single OC Descriptor. Control words also supply control information that drives the Policy Search Engine. Control words can also specify branch conditions that can start the execution of another OC based on the result of the current OC. Besides C-Words, the E-RAM also contains Data Words (D-Words). These D-Words contain statistical and state information such as Packet Count, Byte Count, TCP State and Time-Stamp. D-Words are typically accessed and updated by the device after a match or when an OC execution is complete.

1.5 System Diagram

Figure 1 illustrates one possible application of the PM2329 in a line interface card of a high performance switch.

Figure 1 System Block Diagram



2 Signal Description

2.1 Introduction

The PM2329 has three external interfaces to connect external devices:

- System interface
- Extended RAM (E-RAM) interface
- Cascade interface

The System interface is a general purpose synchronous 64-bit interface over which the PM2329 can connect to a processor and optionally a Packet Source or DMA device. This interface is designed to work like a synchronous SRAM interface so that it can be connected to an external search machine interface or memory interface on common network processors. On 32-bit platforms, this interface can be configured to work in 32-bit mode.

The Extended RAM (E-RAM) interface is used in the Extended (single device or a cascaded set of devices) mode of operation. This interface drives all the signals required for an external SRAM memory array.

The Cascade interface is used when multiple PM2329 devices are used in the Cascade mode. The cascade mode of operation allows up to eight PM2329 devices to work together as one large logical PM2329.

Extended and Cascade modes are described in Sections 2.4.2 and 2.4.3, respectively.

Figure 3 PM2329 Ball Number Assignments

A26	A25	A24	A23	A22	A21	A20	A19	A18	A17	A16	A15	A14	A13	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1
B26	B25	B24	B23	B22	B21	B20	B19	B18	B17	B16	B15	B14	B13	B12	B11	B10	B9	B8	B7	B6	B5	B4	B3	B2	B1
C26	C25	C24	C23	C22	C21	C20	C19	C18	C17	C16	C15	C14	C13	C12	C11	C10	C9	C8	C7	C6	C5	C4	C3	C2	C1
D26	D25	D24	D23	D22	D21	D20	D19	D18	D17	D16	D15	D14	D13	D12	D11	D10	D9	D8	D7	D6	D5	D4	D3	D2	D1
E26	E25	E24	E23																			E4	E3	E2	E1
F26	F25	F24	F23																			F4	F3	F2	F1
G26	G25	G24	G23																			G4	G3	G2	G1
H26	H25	H24	H23																			H4	H3	H2	H1
J26	J25	J24	J23																			J4	J3	J2	J1
K26	K25	K24	K23																			K4	K3	K2	K1
L26	L25	L24	L23																			L4	L3	L2	L1
M26	M25	M24	M23																			M4	M3	M2	M1
N26	N25	N24	N23																			N4	N3	N2	N1
P26	P25	P24	P23																			P4	P3	P2	P1
R26	R25	R24	R23																			R4	R3	R2	R1
T26	T25	T24	T23																			T4	T3	T2	T1
U26	U25	U24	U23																			U4	U3	U2	U1
V26	V25	V24	V23																			V4	V3	V2	V1
W26	W25	W24	W23																			W4	W3	W2	W1
Y26	Y25	Y24	Y23																			Y4	Y3	Y2	Y1
AA26	AA25	AA24	AA23																			AA4	AA3	AA2	AA1
AB26	AB25	AB24	AB23																			AB4	AB3	AB2	AB1
AC26	AC25	AC24	AC23	AC22	AC21	AC20	AC19	AC18	AC17	AC16	AC15	AC14	AC13	AC12	AC11	AC10	AC9	AC8	AC7	AC6	AC5	AC4	AC3	AC2	AC1
AD26	AD25	AD24	AD23	AD22	AD21	AD20	AD19	AD18	AD17	AD16	AD15	AD14	AD13	AD12	AD11	AD10	AD9	AD8	AD7	AD6	AD5	AD4	AD3	AD2	AD1
AE26	AE25	AE24	AE23	AE22	AE21	AE20	AE19	AE18	AE17	AE16	AE15	AE14	AE13	AE12	AE11	AE10	AE9	AE8	AE7	AE6	AE5	AE4	AE3	AE2	AE1
AF26	AF25	AF24	AF23	AF22	AF21	AF20	AF19	AF18	AF17	AF16	AF15	AF14	AF13	AF12	AF11	AF10	AF9	AF8	AF7	AF6	AF5	AF4	AF3	AF2	AF1

Bottom View

2.2 Signals Listed by Function

Table 1 Timing and Common Control Signals

Signal Name	Ball #	Size	I/O	Description
RESET*	Y3	1	I	Reset Input (Active Low) This is the asynchronous reset input, it must be active for a minimum of 100 SCLK cycles. When this signal is asserted, PM2329 is forced into its reset state.
CID [2:0]	AB3 Y4 AB2	3	I	PM2329 ID Number These signals are sensed at reset to determine the 3-bit PM2329 ID number.
SDWIDTH64	AA3	1	I	SD Width Select 64 This signal is sampled at reset to determine the width of the System Interface Data bus. 0: 32-bit (Data transfers on SD[63:32] only) 1: 64-bit
ZBTMODE	AB1	1	I	ZBT Mode Select This signal is sampled at reset to indicate the System Bus Interface type 0: SyncBurst Pipelined Synchronous SCD SRAM Mode 1: ZBT Pipelined Synchronous SRAM Mode
ECLKOUT	N23	1	O	E-RAM Clock Output Regenerated SCLK is output on this pin. When the device is operating in the single device configuration, this signal should be used to drive the external SSRAM clock input. When the device is operating in cascade mode, it is recommended that an external PLL be used to generate the clock input for the Extended SSRAM array. The external PLL must guarantee a minimum skew to meet the SSRAM timing requirement.
ACLKIN	P4	1	I	Internal Clock Input During normal device operation, this input should be tied low. When PLLA is bypassed, this pin drives the internal ACLK signal.

Table 1 Timing and Common Control Signals

Signal Name	Ball #	Size	I/O	Description						
ACLKOUT	R2	1	O	<p>Internal ACLK Output</p> <p>During normal device operation (TESTMODE is low), this signal is driven low.</p> <p>During test mode operation (TESTMODE is high), this pin outputs a divided by 4 version of the qualified ACLK. The source of ACLKOUT can either be:</p> <p>Internal ACLK signal generated by the on-chip PLLA (when PLLABYPS is low) - or - External ACLKIN input signal (when PLLABYPS is high)</p> <p>If the ACLK frequency is 232MHz, the corresponding ACLKOUT will be 50 MHz.</p>						
SCLKOUT	M2	1	O	<p>SCLK Output</p> <p>During normal device operation (TESTMODE is low), this signal is driven low.</p> <p>During test mode operation (TESTMODE is high), this pin outputs the qualified SCLK. The source of SCLKOUT can either be:</p> <p>Internal SCLK signal generated by the on-chip PLLS (when PLLSBYPS is low) - or - External SCLK input signal (when PLLSBYPS is high)</p>						
SCLK	N4	1	I	<p>System Clock Input</p> <p>This is the main timing clock input to the PM2329. It must be active at all times. The maximum clock input frequency is dictated by the CVDD voltage input level.</p> <table border="0"> <tr> <td>Nominal CVDD</td> <td>Max SCLK input</td> </tr> <tr> <td>1.5V</td> <td>100MHz</td> </tr> <tr> <td>1.6V</td> <td>116MHz</td> </tr> </table>	Nominal CVDD	Max SCLK input	1.5V	100MHz	1.6V	116MHz
Nominal CVDD	Max SCLK input									
1.5V	100MHz									
1.6V	116MHz									
PLLSBYP	N3	1	I	<p>PLLS Bypass</p> <p>During normal device operation, this signal must be grounded and its state must not be changed during operation of the PM2329.</p> <p>In order to bypass the internal PLLS, this signal must be forced high and the SCLK clock signal supplied on the SCLK pin is used to drive the clock internally.</p>						

Table 1 Timing and Common Control Signals

Signal Name	Ball #	Size	I/O	Description
PLLACTL[1:0]	P2 P3	2	I	<p>PLLA Control</p> <p>These pins select an internal ACLK clock of 4xSCLK, 3xSCLK, 2xSCLK or 1xSCLK. These control signals must be tied high or low as required and must not change during the operation of the PM2329.</p> <p>If Nominal CVDD = 1.5V 00: 4x for SCLK < 50 MHz 01: 3x for SCLK < 66 MHz 10: 2x for SCLK < 100 MHz 11: 1x for SCLK < 100 MHz (for reduced power application)</p> <p>If Nominal CVDD = 1.6V 00: 4x for SCLK < 58 MHz 01: 3x for SCLK < 77 MHz 10: 2x for SCLK < 116 MHz 11: 1x for SCLK < 116 MHz (for reduced power application)</p> <p>These control signals must be tied appropriately, even when the PLLA is bypassed.</p>
PLLABYPS	N2	1	I	<p>PLLA Bypass</p> <p>During normal device operation, this signal must be grounded and its state must not be changed during operation of the PM2329.</p> <p>In order to bypass the internal PLLA, this signal must be forced high and the ACLK clock signal supplied on the ACLKIN pin from an external source.</p>

Table 2 System Interface Signals

Signal Name	Ball #	Size	I/O	Description
SCE0*	V1	1	I	System Chip Enable 0 (Active Low) This is the Chip Enable signal for the PM2329. It should be driven active low for every read or write cycle to the device.
SCE1*	U2	1	I	System Chip Enable 1 (Active Low) This is the Chip Enable signal for the PM2329. It should be driven active low for every read or write cycle to the device.
SCE2	G2	1	I	System Chip Enable 2 (Active High) This is the Chip Enable signal for the PM2329. It should be driven active high for every read or write cycle to the device.
SA[15:3]	Refer to section 2.3 tables 2.3-2.4	13	I	System Address Bus Address bus to address the internal locations in the PM2329. SA[11:3] : These signals are used to address the internal 64-bit registers of the PM2329 SA[14:12] : These 3 address lines must match the CID# of the PM2329 for its internal registers to be addressed as local registers. SA[15] : If this address bit is asserted high during a write cycle, the register addressed by SA[11:3] is accessed irrespective of the value on SA[14:12]. A Global write is performed in this manner. Note: When the PM2329 is configured in 32-bit mode, SA[2] is driven on the SWHE* signal.
SD[63:0]	Refer to section 2.3 tables 2.2-2.3	64	I/O	System Data Bus All data transfers between the PM2329 and the external system components take place on this bus. When PM2329 is configured for 64-bit data bus width, SD[63:0] should be tied to the external Packet Processor. When PM2329 is configured for 32-bit data bus width, signals SD[63:32] are used, SD[31:0] should be left open.
SINT*	U1	1	O	System Interrupt (Active Low) Interrupt signal from PM2329 to the processor. In cascaded applications, SINT* from PM2329 Zero (CID=0) is connected to the processor, SINT* signals from all other PM2329 devices are left open and unconnected.

Table 2 System Interface Signals

Signal Name	Ball #	Size	I/O	Description
SOE*	U3	1	I	System Output Enable (Active Low) This signal is driven active low by the Packet Processor to enable PM2329 data onto the external bus when a read operation is performed. This signal must be held high during write cycle.
SRW*	V2	1	I	System Write Enable (Active Low) If SyncBurst mode is selected, this signal can be tied permanently low. If ZBT mode is selected, this signal should be driven low during a write cycle and held high during a read cycle.
SWHE* {SA[2]}	U4	1	I	System Write High Enable (Active Low) or System Address bit 2 In 64-bit mode, this signal is SWHE* and should be asserted during write cycles to indicate write data transfers on the upper 32 bits of the System Data Bus (SD[63:32]). In 32-bit mode, this signal is connected to the System Address line #2 of the Packet Processor, and functions as the SA[2] signal.
SWLE* {SWE*}	W1	1	I	System Write Low Enable (Active Low) In 64-bit mode, this signal should be asserted during write cycles to indicate write data transfers on the lower 32 bits of the System Data Bus (SD[31:0]). In 32-bit mode, this signal should be asserted during write cycle.
PSPD	E24	1	I	Packet Source Packet Direction This signal can be used by the PS to indicate the direction of packet movement (upstream/downstream) to the PM2329. This signal should indicate the packet direction during each cycle of packet data transfer after PSPBA has been asserted. This signal is used in DMA mode only and should be tied low otherwise.
PSEOP	E26	1	I	Packet Source End of Packet PS asserts this signal to indicate that the current packet transfer to the PM2329 is complete. The PM2329 can now process the completed packet. This signal is typically connected to the Terminal Count pin of the DMA device. This signal is used in DMA mode only and should be tied low otherwise.

Table 2 System Interface Signals

Signal Name	Ball #	Size	I/O	Description
PSCC	G25	1	I	<p>Packet Source Current Cycle</p> <p>The PS asserts this signal to indicate to the PM2329 that the current bus cycle is packet data download. This signal is typically connected to the DMA Acknowledge pin of a DMA controller. When this signal is asserted, the other System control lines SA[15:3] and SCE* are ignored.</p> <p>This signal is used in DMA mode only and should be tied low otherwise.</p>
PSPBA	H24	1	O	<p>Packet Source Packet Buffer Available</p> <p>This signal when asserted indicates to the PS that at least one Packet Buffer is currently available to receive packet data. This signal is typically connected to the DMA Request pin of the DMA device.</p>
SCHNUM[4:0]	E25 F24 G24 F25 F26	5	O	<p>System Channel Number</p> <p>In multi-channel mode, when the SCHSTB signal is active, these signals indicate the channel number for which the result is available.</p>
SCHSTB	G23	1	O	<p>System Channel Strobe</p> <p>This signal is asserted for one SCLK cycle when a result for a packet becomes available. If multiple results are generated for a packet, then this signal is asserted every time a result becomes available.</p>

Table 3 ERAM Interface Signals

Signal Name	Ball #	Size	I/O	Description
EMCE*	P24	1	O	<p>Extended Memory Chip Enable (Active Low)</p> <p>This signal when active selects the entire bank of Extended Memory. Only the primary PM2329 device drives this signal. The EMCE* signal from the primary device should be tied to the the Chip Enable inputs of all the E-RAMs (C-Word and all D-Words) in the E-RAM array. The EMCE* signal of all secondary PM2329 devices should be left unconnected.</p>
EMOE*	R25	1	O	<p>Extended Memory Output Enable (Active Low)</p> <p>Read enable to the entire bank of Extended Memory. Only the primary PM2329 device drives this signal. The EMOE* signal from the primary device should be tied to the the Output Enable inputs of all the E-RAMs (C-Word and all D-Words) in the E-RAM array. The EMOE* signal of all secondary PM2329 devices should be left unconnected.</p>
SPARE4 SPARE3	T24 M26	2	O	<p>Spare signals for future use</p> <p>These signals are reserved for future use, in the current device they can be left unconnected.</p> <p>Current proposed assignment on future devices, subject to change, is shown below. The current default output state of the signals is shown in parentheses.</p> <p>SPARE4: ECTL (1) SPARE3: EMA[19] (0)</p>
EMA[18:0]	Refer to section 2.3 tables 2.4-2.5	19	O	<p>Extended Memory Address Bus</p> <p>Address bus for the Extended Memory. Any of the cascaded PM2329 devices can drive this bus.</p> <p>Address lines EMA[16:0] are driven from the E-RAM Address field whereas EMA[18:17] are driven from the depth control field.</p> <p>For EMA[16:0] signals, the number of address bits that need to be connected to the external E-RAM device is a function of the total number of E-Words required.</p> <p>EMA[18:17] can be left unconnected if depth is 1. For depth of 2, only EMA[17] needs to be connected. For depth of 4, EMA[18:17] should be connected.</p>
EDD[31:0]	Refer to section 2.3 tables 2.5-2.6	32	I/O	<p>Extended Data Memory Data Bus</p> <p>Data bus for the data memory assigned to each of the PM2329 devices.</p>

Table 3 ERAM Interface Signals

Signal Name	Ball #	Size	I/O	Description
EDWE*	U24	1	O	Extended Data Memory Write Enable (Active Low) This is the write enable for the data memory assigned to the PM2329. In the cascade mode, each PM2329 device drives the write enable of the corresponding D-Word ERAM.
ECD[31:0]	Refer to section 2.3 tables 2.5-2.6	32	I/O	Extended C-Word Data Bus Shared Data Bus for the C-Words in the Extended memory. This bus transfers the C-Words to all the PM2329 devices in the cascade. Only the primary PM2329 writes C-Word data over this bus.
ECWE*	V24	1	O	Extended C-Word Write Enable (Active Low) This signal is driven by the primary PM2329 to write C-Words to the E-RAM. In the cascade mode, this signal from all secondary devices is left unconnected.

Table 4 Cascade Interface Signals

Signal Name	Ball #	Size	I/O	Description
COCS	H26	1	I or O	Cascade OC Start Output for primary and input for all others. Asserted by Primary PM2329 to indicate start of a new OC to all others. De-asserted when all the PM2329 devices in the cascade complete their OCs.
COCDIN[7:0]	K24 J26 K25 K26 L24 L25 L26 M25	8	I	Cascade OC Done Inputs Connected to the corresponding COCDOUT pins of the other PM2329 devices in the Cascade. Unused COCDIN pins should be tied high.
COCDOUT	J24	1	O	Cascade OC Done Output Connected to the appropriate COCDIN[n] pin on the other PM2329 devices in the cascade.
COCM*	K23	1	I/O	Cascade OC Match (Active Low) This signal is asserted low by a PM2329 if and only if its own OC has resulted in a Match AND all the higher order PM2329 devices participating in that OC have asserted their COCDOUT signals. This signal is wire-OR connected to all PM2329 devices and must be pulled high by an external pull up.
CEMRQ	G26	1	I/O	Extended Memory Bus Request This signal is asserted by the Primary PM2329 to indicate to the other PM2329 devices in the cascade that it has a pending Extended Memory data transfer operation. It is used for Extended Memory bus arbitration.

Table 5 Test Signals

Signal Name	Ball #	Size	I/O	Description
JTRST	V3	1	I	JTAG Reset (tie LOW during normal operation)
JTMS	W2	1	I	JTAG Mode (tie HIGH during normal operation)
JTDI	Y1	1	I	JTAG Data In (tie HIGH during normal operation)
JTDO	W3	1	O	JTAG Data Out (No Connect during normal operation)
JTCK	Y2	1	I	JTAG Clock Input (tie LOW during normal operation)
TRISTATE	AA1	1	I	Tristate Mode Select During normal operation, this signal should be tied permanently low. If this signal is asserted high all outputs of the PM2329 are tristated.
TESTCTL[3:0]	H25 T2 T3 J25	3	I	Test Control signals These test signals are for test and debug purpose only. During normal device operation, these TESTCTL[3] should be tied permanently low and TESTCTL[2:0] should be tied permanently high. For debug purpose, a provision should be made on the board to drive these signals low or high.
TESTOUT[7:0]	AD5 AC7 AC20 AD22 C22 D20 D7 C5	8	O	Test Output signals These test signals are for test and debug purpose only. During normal device operation, these signals can be left unconnected. For debug purpose, these signals should be routed to test points for observation.
TESTMODE	AA2	1	I	Test signal This test signal is for internal use only. During normal device operation this signal should be tied low.

Table 6 VDD and VSS Signals

Signal Name	Ball #	Size	Description
AVSS1 AVSS2	M1 R1	2	PLL analog ground Each AVSS signal should be tied to its corresponding AVDD signal using two low-inductance bypass capacitors (1-10 uF and 0.01-0.001uF). These signals should not be tied to the general ground plane. In order to ensure that the digital and analog sections of the on-chip PLLs have a common reference, they are tied to a ground reference internally on-chip.
AVDD1 AVDD2	N1 P1	2	PLL analog power 1.5V All signals must be connected to an appropriate power supply for correct device operation.
CVSS	Refer to section 2.3 tables 2.3-2.7	16	Core ground All signals must be connected to an appropriate ground plane for correct device operation.
CVDD	Refer to section 2.3 tables 2.3-2.7	16	Core power 1.5V All signals must be connected to an appropriate power supply for correct device operation.
VSS	Refer to section 2.3 tables 2.3-2.7	64	I/O ground All signals must be connected to an appropriate ground plane for correct device operation.
VDD	Refer to section 2.3 tables 2.3-2.7	16	I/O power 3.3V All signals must be connected to an appropriate power supply for correct device operation.

2.3 Signals Listed by Ball

Table 7 Signals Listed by Ball Assignment; Rows A through C

Ball	Signal Name	Ball	Signal Name	Ball	Signal Name
A1	VSS	B1	VSS	C1	VSS
A2	VSS	B2	VSS	C2	VSS
A3	VSS	B3	VSS	C3	VSS
A4	VSS	B4	VSS	C4	VSS
A5	SD10	B5	SD8	C5	TESTOUT0
A6	SD13	B6	SD12	C6	SD9
A7	SD16	B7	SD14	C7	SD11
A8	SD19	B8	SD17	C8	SD15
A9	SD23	B9	SD21	C9	SD18
A10	SD25	B10	SD24	C10	SD22
A11	SD28	B11	SD27	C11	SD26
A12	SD31	B12	SD30	C12	SD29
A13	SD32	B13	SD35	C13	SD34
A14	SD39	B14	SD36	C14	SD37
A15	SD40	B15	SD41	C15	SD42
A16	SD43	B16	SD44	C16	SD45
A17	SD46	B17	SD47	C17	SD49
A18	SD48	B18	SD50	C18	SD53
A19	SD52	B19	SD54	C19	SD56
A20	SD55	B20	SD57	C20	SD60
A21	SD58	B21	SD59	C21	SD62
A22	SD61	B22	SD63	C22	TESTOUT3
A23	VSS	B23	VSS	C23	VSS
A24	VSS	B24	VSS	C24	VSS
A25	VSS	B25	VSS	C25	VSS
A26	VSS	B26	VSS	C26	VSS

Table 8 Signals listed by Ball Assignment; Rows D through K

Ball	Signal Name	Ball	Signal Name	Ball	Signal Name
D1	VSS	E1	SD3	H1	SA11
D2	VSS	E2	SD5	H2	SA13
D3	VSS	E3	SD7	H3	SA15
D4	VSS	E4	CVDD	H4	CVSS
D5	CVDD	E23	CVDD	H23	CVSS
D6	CVDD	E24	PSPD	H24	PSPBA
D7	TESTOUT1	E25	SCHNUM4	H25	TESTCTL3
D8	CVSS	E26	PSEOP	H26	COCS
D9	CVSS	F1	SD0	J1	SA7
D10	SD20	F2	SD1	J2	SA9
D11	VDD	F3	SD4	J3	SA12
D12	VDD	F4	CVDD	J4	CVSS
D13	SD33	F23	CVDD	J23	CVSS
D14	SD38	F24	SCHNUM3	J24	COCDOUT
D15	VDD	F25	SCHNUM1	J25	TESTCTL0
D16	VDD	F26	SCHNUM0	J26	COCDIN6
D17	SD51	G1	SA14	K1	SA5
D18	CVSS	G2	SCE2	K2	SA6
D19	CVSS	G3	SD2	K3	SA8
D20	TESTOUT2	G4	SD6	K4	SA10
D21	CVDD	G23	SCHSTB	K23	COCM*
D22	CVDD	G24	SCHNUM2	K24	COCDIN7
D23	VSS	G25	PSCC	K25	COCDIN5
D24	VSS	G26	CEMRQ	K26	COCDIN4
D25	VSS				
D26	VSS				

Table 9 Signals listed by Ball Assignment; Rows L through W

Ball	Signal Name	Ball	Signal Name	Ball	Signal Name
L1	N.C.	P1	AVDD2	U1	SINT*
L2	SA3	P2	PLLACTL1	U2	SCE1*
L3	SA4	P3	PLLACTL0	U3	SOE*
L4	VDD	P4	ACLKIN	U4	SWHE*
L23	VDD	P23	EMA14	U23	EMA4
L24	COCDIN3	P24	EMCE*	U24	EDWE*
L25	COCDIN2	P25	EMA15	U25	EMA7
L26	COCDIN1	P26	EMA13	U26	EMA8
M1	AVSS1	R1	AVSS2	V1	SCE0*
M2	SCLKOUT	R2	ACLKOUT	V2	SRW*
M3	N.C.	R3	N.C.	V3	JTRST
M4	VDD	R4	VDD	V4	CVSS
M23	VDD	R23	VDD	V23	CVSS
M24	N.C.	R24	EMA11	V24	ECWE*
M25	COCDIN0	R25	EMOE*	V25	EMA5
M26	SPARE3	R26	EMA12	V26	EMA6
N1	AVDD1	T1	N.C.	W1	SWLE*
N2	PLLBYPS	T2	TESTCTL2	W2	JTMS
N3	PLLSBYP	T3	TESTCTL1	W3	JTDO
N4	SCLK	T4	VDD	W4	CVSS
N23	ECLKOUT	T23	VDD	W23	CVSS
N24	EMA17	T24	SPARE4	W24	EMA0
N25	EMA16	T25	EMA9	W25	EMA2
N26	EMA18	T26	EMA10	W26	EMA3

Table 10 Signals listed by Ball Assignment; Rows Y through AD

Ball	Signal Name	Ball	Signal Name	Ball	Signal Name
Y1	JTDI	AC1	VSS	AD1	VSS
Y2	JTCK	AC2	VSS	AD2	VSS
Y3	RESET*	AC3	VSS	AD3	VSS
Y4	CID1	AC4	VSS	AD4	VSS
Y23	EDD25	AC5	CVDD	AD5	TESTOUT7
Y24	EDD28	AC6	CVDD	AD6	ECD1
Y25	EDD31	AC7	TESTOUT6	AD7	ECD3
Y26	EMA1	AC8	CVSS	AD8	ECD7
AA1	TRISTATE	AC9	CVSS	AD9	ECD10
AA2	TESTMODE	AC10	ECD12	AD10	ECD14
AA3	SDWIDTH64	AC11	VDD	AD11	ECD18
AA4	CVDD	AC12	VDD	AD12	ECD21
AA23	CVDD	AC13	ECD25	AD13	ECD26
AA24	EDD27	AC14	ECD30	AD14	ECD29
AA25	EDD29	AC15	VDD	AD15	EDD2
AA26	EDD30	AC16	VDD	AD16	EDD5
AB1	ZBTMODE	AC17	EDD11	AD17	EDD9
AB2	CID0	AC18	CVSS	AD18	EDD13
AB3	CID2	AC19	CVSS	AD19	EDD16
AB4	CVDD	AC20	TESTOUT5	AD20	EDD20
AB23	CVDD	AC21	CVDD	AD21	EDD22
AB24	EDD24	AC22	CVDD	AD22	TESTOUT4
AB25	EDD26	AC23	VSS	AD23	VSS
AB26	N.C.	AC24	VSS	AD24	VSS
		AC25	VSS	AD25	VSS
		AC26	VSS	AD26	VSS

Table 11 Signals listed by Ball Assignment; Rows AE and AF

Ball	Signal Name	Ball	Signal Name
AE1	VSS	AF1	VSS
AE2	VSS	AF2	VSS
AE3	VSS	AF3	VSS
AE4	VSS	AF4	VSS
AE5	ECD0	AF5	ECD2
AE6	ECD4	AF6	ECD5
AE7	ECD6	AF7	ECD8
AE8	ECD9	AF8	ECD11
AE9	ECD13	AF9	ECD15
AE10	ECD16	AF10	ECD17
AE11	ECD19	AF11	ECD20
AE12	ECD22	AF12	ECD23
AE13	ECD27	AF13	ECD24
AE14	ECD28	AF14	ECD31
AE15	EDD1	AF15	EDD0
AE16	EDD4	AF16	EDD3
AE17	EDD7	AF17	EDD6
AE18	EDD10	AF18	EDD8
AE19	EDD14	AF19	EDD12
AE20	EDD17	AF20	EDD15
AE21	EDD19	AF21	EDD18
AE22	EDD23	AF22	EDD21
AE23	VSS	AF23	VSS
AE24	VSS	AF24	VSS
AE25	VSS	AF25	VSS
AE26	VSS	AF26	VSS

Table 12 Signals Listed by Name (Alphabetically)

Signal Name	Ball	Signal Name	Ball	Signal Name	Ball
ACLKIN	P4	CVSS	D18	ECD27	AE13
ACLKOUT	R2	CVSS	D19	ECD28	AE14
AVDD1	N1	CVSS	V23	ECD29	AD14
AVDD2	P1	CVSS	AC18	ECD30	AC14
AVSS1	M1	CVSS	AC19	ECD31	AF14
AVSS2	R1	CVSS	AC8	ECLKOUT	N23
CEMRQ	G26	CVSS	AC9	ECWE*	V24
CID0	AB2	CVSS	H23	EDD0	AF15
CID1	Y4	CVSS	H4	EDD1	AE15
CID2	AB3	CVSS	J23	EDD2	AD15
COCDIN0	M25	CVSS	J4	EDD3	AF16
COCDIN1	L26	CVSS	W23	EDD4	AE16
COCDIN2	L25	CVSS	W4	EDD5	AD16
COCDIN3	L24	ECD0	AE5	EDD6	AF17
COCDIN4	K26	ECD1	AD6	EDD7	AE17
COCDIN5	K25	ECD2	AF5	EDD8	AF18
COCDIN6	J26	ECD3	AD7	EDD9	AD17
COCDIN7	K24	ECD4	AE6	EDD10	AE18
COCDOUT	J24	ECD5	AF6	EDD11	AC17
COCM*	K23	ECD6	AE7	EDD12	AF19
COCS	H26	ECD7	AD8	EDD13	AD18
CVDD	D5	ECD8	AF7	EDD14	AE19
CVDD	D6	ECD9	AE8	EDD15	AF20
CVDD	D21	ECD10	AD9	EDD16	AD19
CVDD	D22	ECD11	AF8	EDD17	AE20
CVDD	AA23	ECD12	AC10	EDD18	AF21
CVDD	AA4	ECD13	AE9	EDD19	AE21
CVDD	AB23	ECD14	AD10	EDD20	AD20
CVDD	AB4	ECD15	AF9	EDD21	AF22
CVDD	AC21	ECD16	AE10	EDD22	AD21
CVDD	AC22	ECD17	AF10	EDD23	AE22
CVDD	AC5	ECD18	AD11	EDD24	AB24
CVDD	AC6	ECD19	AE11	EDD25	Y23
CVDD	E23	ECD20	AF11	EDD26	AB25
CVDD	E4	ECD21	AD12	EDD27	AA24
CVDD	F23	ECD22	AE12	EDD28	Y24
CVDD	F4	ECD23	AF12	EDD29	AA25
CVSS	V4	ECD24	AF13	EDD30	AA26
CVSS	D8	ECD25	AC13	EDD31	Y25
CVSS	D9	ECD26	AD13	EDWE*	U24
EMA0	W24	SA4	L3	SD19	A8

Table 12 Signals Listed by Name (Alphabetically)

Signal Name	Ball	Signal Name	Ball	Signal Name	Ball
EMA1	Y26	SA5	K1	SD20	D10
EMA2	W25	SA6	K2	SD21	B9
EMA3	W26	SA7	J1	SD22	C10
EMA4	U23	SA8	K3	SD23	A9
EMA5	V25	SA9	J2	SD24	B10
EMA6	V26	SA10	K4	SD25	A10
EMA7	U25	SA11	H1	SD26	C11
EMA8	U26	SA12	J3	SD27	B11
EMA9	T25	SA13	H2	SD28	A11
EMA10	T26	SA14	G1	SD29	C12
EMA11	R24	SA15	H3	SD30	B12
EMA12	R26	SCE0*	V1	SD31	A12
EMA13	P26	SCE1*	U2	SD32	A13
EMA14	P23	SCE2	G2	SD33	D13
EMA15	P25	SCHNUM0	F26	SD34	C13
EMA16	N25	SCHNUM1	F25	SD35	B13
EMA17	N24	SCHNUM2	G24	SD36	B14
EMA18	N26	SCHNUM3	F24	SD37	C14
EMCE*	P24	SCHNUM4	E25	SD38	D14
EMOE*	R25	SCHSTB	G23	SD39	A14
JTCK	Y2	SCLK	N4	SD40	A15
JTDI	Y1	SCLKOUT	M2	SD41	B15
JTDO	W3	SD0	F1	SD42	C15
JTMS	W2	SD1	F2	SD43	A16
JTRST	V3	SD2	G3	SD44	B16
N.C.	AB26	SD3	E1	SD45	C16
N.C.	L1	SD4	F3	SD46	A17
N.C.	M24	SD5	E2	SD47	B17
N.C.	M3	SD6	G4	SD48	A18
N.C.	R3	SD7	E3	SD49	C17
N.C.	T1	SD8	B5	SD50	B18
PLLBYPS	N2	SD9	C6	SD51	D17
PLLACTL0	P3	SD10	A5	SD52	A19
PLLACTL1	P2	SD11	C7	SD53	C18
PLLSBYPS	N3	SD12	B6	SD54	B19
PSCC	G25	SD13	A6	SD55	A20
PSEOP	E26	SD14	B7	SD56	C19
PSPBA	H24	SD15	C8	SD57	B20
PSPD	E24	SD16	A7	SD58	A21
RESET*	Y3	SD17	B8	SD59	B21
SA3	L2	SD18	C9	SD60	C20
SD61	A22	VDD	M4	VSS	AE4

Table 12 Signals Listed by Name (Alphabetically)

Signal Name	Ball	Signal Name	Ball	Signal Name	Ball
SD62	C21	VDD	R23	VSS	AF1
SD63	B22	VDD	R4	VSS	AF2
SDWIDTH64	AA3	VDD	T23	VSS	AF23
SINT*	U1	VDD	T4	VSS	AF24
SOE*	U3	VSS	A1	VSS	AF25
SPARE3	M26	VSS	A2	VSS	AF26
SPARE4	T24	VSS	A23	VSS	AF3
SRW*	V2	VSS	A24	VSS	AF4
SWHE*	U4	VSS	A25	VSS	B1
SWLE*	W1	VSS	A26	VSS	B2
TESTCTL0	J25	VSS	A3	VSS	B23
TESTCTL1	T3	VSS	A4	VSS	B24
TESTCTL2	T2	VSS	AC1	VSS	B25
TESTCTL3	H25	VSS	AC2	VSS	B26
TESTMODE	AA2	VSS	AC23	VSS	B3
TESTOUT0	C5	VSS	AC24	VSS	B4
TESTOUT1	D7	VSS	AC25	VSS	C1
TESTOUT2	D20	VSS	AC26	VSS	C2
TESTOUT3	C22	VSS	AC3	VSS	C23
TESTOUT4	AD22	VSS	AC4	VSS	C24
TESTOUT5	AC20	VSS	AD1	VSS	C25
TESTOUT6	AC7	VSS	AD2	VSS	C26
TESTOUT7	AD5	VSS	AD23	VSS	C3
TRISTATE	AA1	VSS	AD24	VSS	C4
VDD	AC11	VSS	AD25	VSS	D1
VDD	AC12	VSS	AD26	VSS	D2
VDD	AC15	VSS	AD3	VSS	D23
VDD	AC16	VSS	AD4	VSS	D24
VDD	D11	VSS	AE1	VSS	D25
VDD	D12	VSS	AE2	VSS	D26
VDD	D15	VSS	AE23	VSS	D3
VDD	D16	VSS	AE24	VSS	D4
VDD	L23	VSS	AE25	ZBTMODE	AB1
VDD	L4	VSS	AE26		
VDD	M23	VSS	AE3		

2.4 Interface Description

2.4.1 System Interface

The System Interface is a general-purpose, 64-bit synchronous bus interface, which can optionally be configured as a 32-bit interface by holding the SDWIDTH64 pin low during reset. The PM2329's ability to perform wire-speed operations at Gigabit rates is dependent on the speed of data transfer over this interface. In general, data transfer speed will be halved when the system interface is configured for 32-bit wide data.

The System Interface is used by:

1. The processor to access the registers and Rules Memory of the PM2329 and the Extended RAM locations.
2. The packet source (or DMA source) to load packets (or packet data) to the PM2329.

The PM2329 is designed to interface to the processor as well as the packet source (configured as a DMA device) for packet transfer without any additional logic. In case a packet source (DMA) is not present, the processor can perform packet transfer.

2.4.1.1 Clock Frequency

The system clock input signal (SCLK) controls the timing of the synchronous system bus signals; address, data and control signal timings are all referenced to the system clock.

The Cascade and E-RAM interfaces also run at the SCLK frequency and are synchronous to SCLK. However, in order to provide improved timing on the E-RAM interface, the PM2329 generates the ECLKOUT signal for the external E-RAM devices.

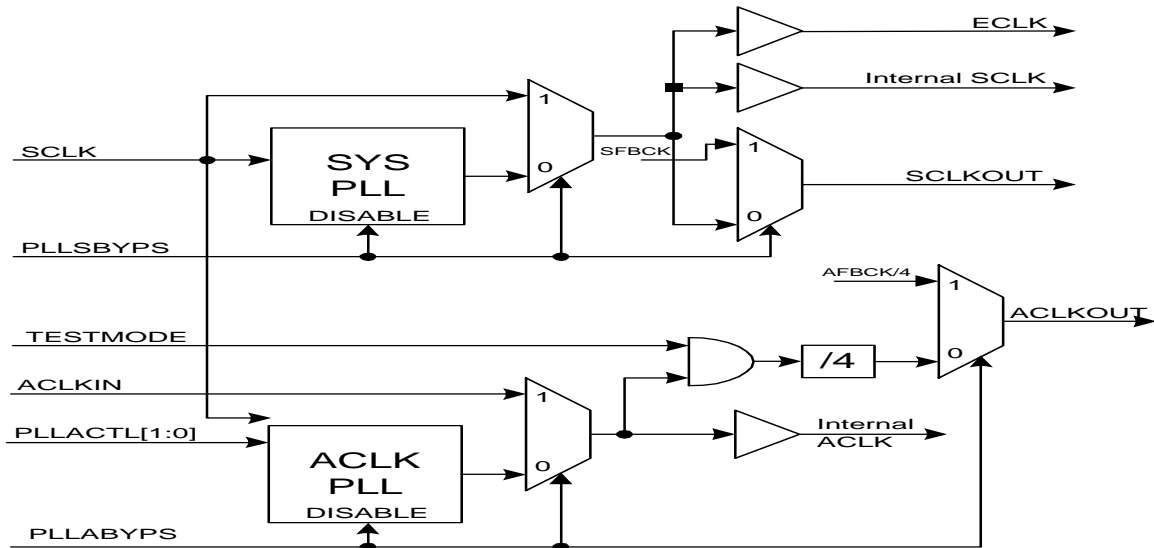
In a lightly loaded system bus with up to two cascaded PM2329 devices and associated E-RAMs, the SCLK frequency may be up to 100MHz. When running at 116MHz SCLK frequency, only a single PM2329 device operation is recommended. Factors such as bus loading must be considered in determining the SCLK frequency. When operating the device at 100 MHz or 116MHz, it is recommended that all high-speed signals should be source- and destination-terminated, to minimize problems due to ground bounce or ringing. Low inductance terminating resistors in the range of 17 to 33 ohms are recommended.

In applications with 3 or more cascaded PM2329 devices and associated E-RAMs, the SCLK frequency should be less than the rated maximum clock input. The actual frequency used in these systems should be appropriately selected in order to achieve a balance of interface speed and the maximum throughput from the PSE logic.

The Policy Search Engine PSE of the PM2329 can run at a frequency of up to 232 MHz regardless of the number of devices in the cascade. It is clocked by an internally generated PSE clock, which is derived from the externally supplied SCLK using an on-chip PLL. The internal PSE clock is derived by multiplying the SCLK by a multiplier of 1x, 2x, 3x or 4x. The multiplier value is sampled at reset by the PM2329 on the PLLCTL[1:0] pins.

The diagram below shows the bypass mechanism for on-chip PLLs. When PLLA is bypassed, PLLACTL[1:0] signals must still be tied appropriately to specify the desired SCLK vs. ACLK clock ratio.

Figure 4 On-Chip PLL Bypass Mechanism



TESTMODE	PLLABYPS	PLLSBYPS	Function
0	0	0	Normal
0	0	1	Bypass SCLK
0	1	0	Bypass ACLK
0	1	1	Bypass both
1	X	X	Test use

2.4.1.2 Processor Bus Cycles

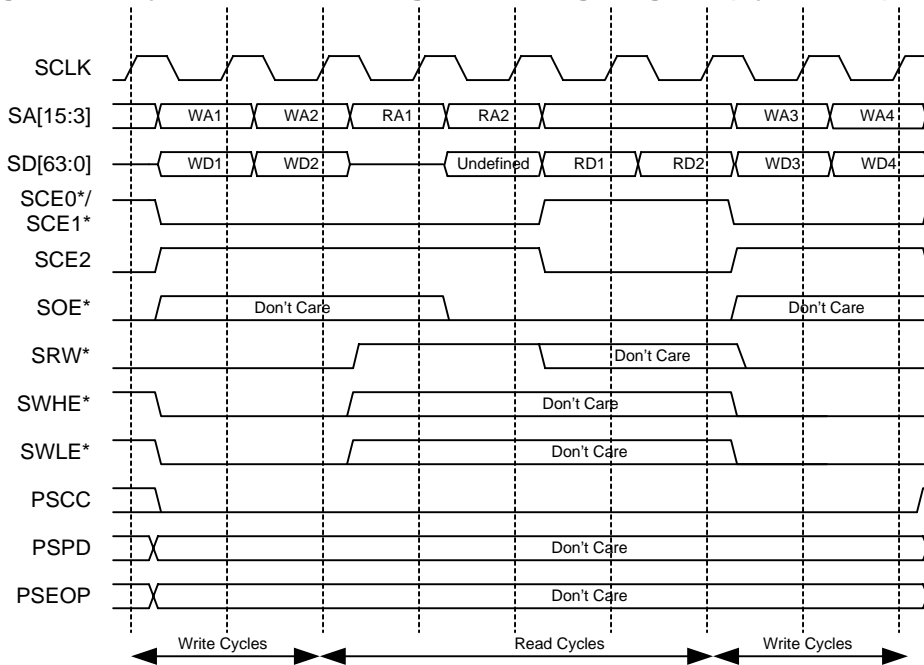
The PM2329 interface to the processor has been designed to appear like a synchronous pipelined SRAM. Two types of bus timings are supported: SyncBurst and ZBT. The block read or block write throughput is the same for SyncBurst and ZBT timings. However, ZBT timings allow greater throughput for random read and write operations. The basic signals involved in a processor access are: SD[63:0], SA[15:3], SWE*, SOE*, SWHE*, SWLE*, SCE0*, SCE1* and SCE2.

In addition to the SCE0* signal, two additional chip enable signals--SCE1* and SCE2--are provided, for a total of two active low and one active high chip enable signals. This allows the PM2329 to interface with external processor at high speed, without the need of any external glue logic for address decode. For systems that do not require multiple chip enable signals, the (one or two) unused SCEn signals can be permanently tied to their active state.

2.4.1.2.1 SyncBurst Bus Cycles

Processor accesses may be either write cycles or read cycles. The write cycle is conducted in a single clock period, whereas the read takes 3 clock periods. The PM2329 supports contiguous or back-to-back transfers for both read and write operations.

Figure 5 System Interface Register Timing Diagram (SyncBurst)



Register Read and Write Timing, SyncBurst Mode

For Reads (address and control cycle), either...
 1) SRW* should be high and SWHE* and SWLE* are don't care (shown), or
 2) SRW* should be low and SWHE* and SWLE* should be high (not shown).

A Write cycle occurs as follows:

1. SA[15:3], and SD[63:32] and SD[31:0] (as required) are driven by the processor. SCE0* and SCE1* are asserted low and SCE2 is asserted high by the processor. One or both of SWHE* and SWLE* are asserted low to indicate the 32-bit lanes on which the write should occur.
2. On the next rising edge of SCLK, the write cycle occurs.
3. In case of back-to-back writes, the next write occurs on the next clock edge.

A Read cycle occurs as follows:

1. SA[15:3] is driven by the processor. SCE0* and SCE1* are asserted low and SCE2 is asserted high by the processor. PM2329 latches this state on the first rising edge of clock and starts the read cycle internally. (Note: If SRW* is driven low and SWHE* and SWLE* are high, then the PM2329 treats this as a read cycle in SyncBurst mode to be compatible with SyncBurst SRAM specification. In ZBT mode, this case is treated as no operation.)
2. On the second rising edge, the PM2329 internally fetches the data to be read. The processor must drive SOE* before the second rising edge. This causes the PM2329 output buffers to be enabled and drive the read data onto the System bus. Keeping SOE* low through the third rising edge will allow the processor to latch the data on this edge.
3. In case of back-to-back reads, the next cycle address and control signals can be asserted after the first rising edge and PM2329 will latch this state on the second rising edge in a pipelined fashion. Thus, read data can be transferred on every cycle.

Note: SRW* should be tied low for SyncBurst mode.

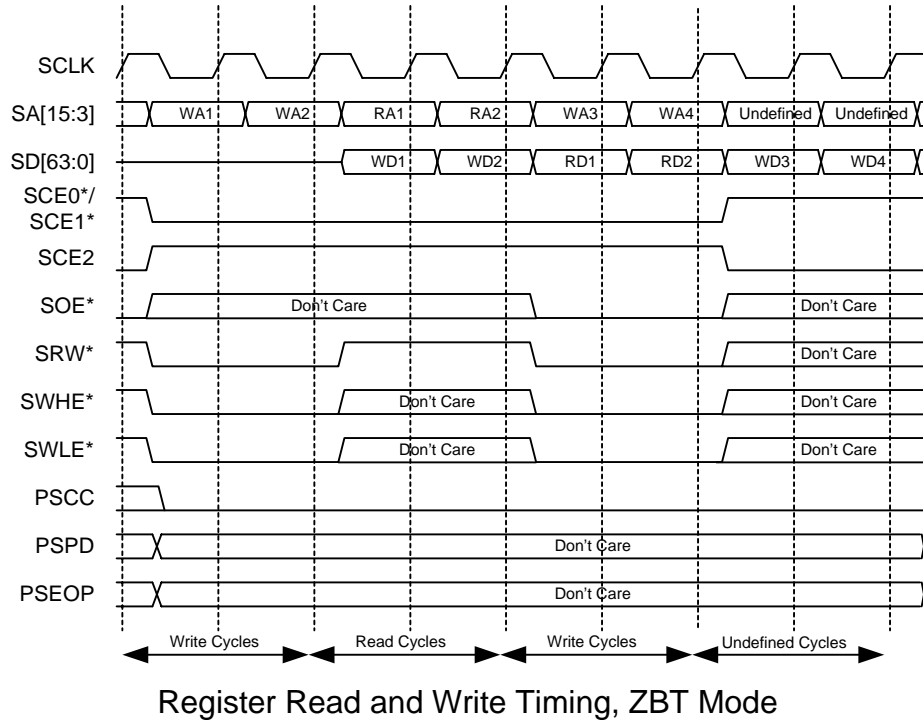
2.4.1.2.2 ZBT Bus Cycles

In the ZBT mode, both the read and write cycles are pipelined. The read cycle is the same as the normal (SyncBurst) mode. (Except for the case stated above: if SRW* is driven low and SWHE* and SWLE* are high in ZBT mode, the PM2329 treats the case as a 'no operation'.). The write cycle is conducted in 3 clock periods, just like the read. The ZBT mode thus supports contiguous or back-to-back transfers on both read and write or interleaved read/write/read cycles. The ZBT operation results in greater available bus bandwidth when random reads/writes are performed.

In a write cycle in ZBT mode, the processor drives the SA[15:3] on the first rising edge. It also drives SRW* low to indicate a write cycle and asserts SCE0*, SCE1* and SCE2 signals on the first rising edge. One or both of SWHE* and SWLE* are asserted low on the first rising edge to indicate the 32-bit lanes on which the write should occur. Finally, the processor asserts write data SD[63:32] and SD[31:0] (as required) on the third rising edge.

There is an important consideration when the PM2329 is used in ZBT mode: If a register is written in cycle 1, its value will only be available for read during bus cycle 4. That is, read cycles performed in bus cycles 2 and 3 will return the old register value. Similarly, if a write to register Q is expected to cause a change in the value of register R, then the new value of register R will only be visible when read at least 2 clock cycles after register Q has been written.

Figure 6 System Interface Register Timing Diagram (ZBT)



2.4.1.3 Packet Source Interface Signals

The PM2329 System interface has four control signals designed to support a dedicated packet processor configured as a DMA controller. These signals are PSPBA, PSEOP, PSPD and PSCC. This DMA interface is a write-only interface, and active in single channel mode only.

When the PM2329 has space in the packet buffer to accept packet data, PSPBA signal is asserted. The signal is deasserted when:

1. The last available 256-byte segment is being written and either the PSEOP signal is asserted or the EOP command terminates the current packet transfer, or
2. The Packet Buffer limit is reached.

The PSPBA signal is deasserted low corresponding to the 27th (64-bit) word written into the last 32-locations by 64-bits segment (that is, the last 256-byte segment) of the packet input buffer. Due to the pipelined nature of the device, there will be a certain delay between the external write cycle for the 27th word and deassertion of PSPBA signal, as noted in Table 13 below.

Table 13 PSPBA Deassertion Delay

Access Mode	PSPBA Deassertion Delay	
	64-bit Mode	32-bit Mode
ZBT	4 clocks	4 clocks
SyncBurst	2 clocks	6 clocks

Depending on the write cycle timing, up to 4 additional 64-bit writes can be performed without causing packet input buffer overflow.

If the DMA source performs a burst write into the PM2329 with data transfers on every clock cycle, it can perform a number of additional data transfers without causing a data overrun, as shown in Figure 7. The number of additional transfers varies depending on transfer width and mode. In 64-bit SyncBurst mode, it can perform two additional data transfers after the cycle in which it detects PSPBA deasserted low (since the internal pipeline is 2 deep). In 64-bit ZBT mode, it can perform data transfers until the cycle in which it detects PSPBA deasserted low (since the internal and external pipeline is 4 deep). The corresponding numbers in 32-bit mode are six for SyncBurst mode and four for ZBT mode.

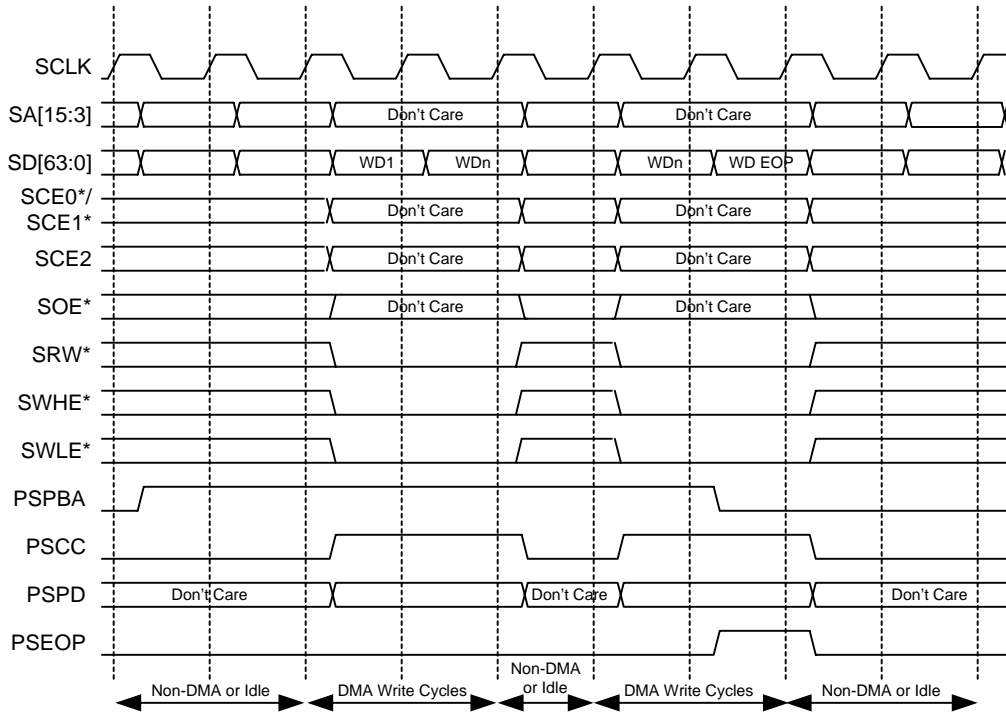
When PSPBA is asserted, the Packet Source acquires the bus and asserts PSCC to indicate to the PM2329 that the present transfer is from the packet source. Assertion of the SWE* in conjunction with SWLE* and/or SWHE* lines causes the current cycle to be treated as a packet data write cycle and data present on SD[63:0] is written to the PM2329 on the next SCLK rising edge. A number of such packet data write cycles could occur in a back-to-back cycle burst. Alternatively, PSCC can be de-asserted while PSPBA is asserted and the processor can execute read or write cycles interleaved with DMA cycles. PSCC is generally tied to the DMA Acknowledge signal.

The PM2329 provides an additional signal, PSPD, to allow the DMA Controller to communicate to the PM2329 the direction of each packet. Packets can be associated with a Direction bit and the PM2329 rules can be programmed to inspect this Direction bit in header lookups. This signal (if used) must be asserted during each cycle of packet transfer. If PSPD is not used, packet direction can be signalled via the packet attribute field.

In order to complete the packet transfer, the PS must assert the PSEOP signal or issue an EOP command for the last word of the packet to be transferred. This indicates to the PM2329 that the current packet transfer is complete and the PM2329 can begin processing the packet data. PSEOP would generally be connected to the Terminal Count signal in a DMA based system.

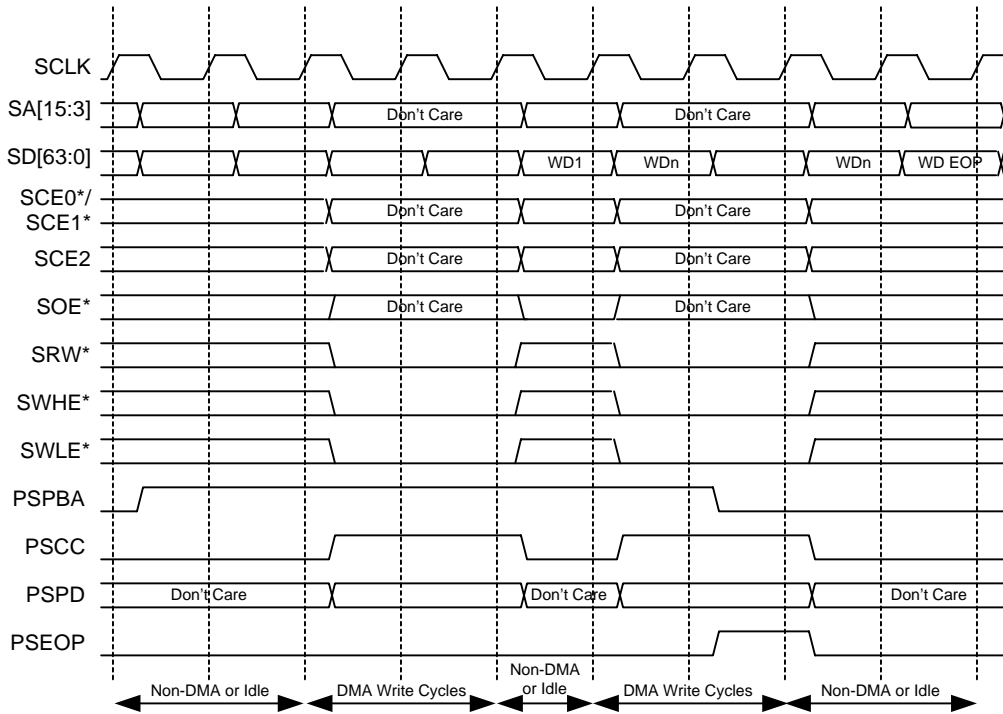
Systems that do not use this DMA mechanism to perform packet transfers should leave the PSPBA output signal unconnected and tie the PSEOP, PSPD and PSCC input signals inactive. In such systems, a processor write to the Packet Buffer Input register can be used to load the packet data into the device. The Packet Buffer Availability status is readable in a PM2329 status register, and End of Packet can be indicated to the PM2329 by a write to the alternate address of the Packet Buffer Input register.

Figure 7 System Interface DMA Timing (SyncBurst Mode)



DMA Timing, SyncBurst Mode
(Supported in single-channel mode only)

Figure 8 System Interface DMA Timing (ZBT Mode)



DMA Timing, ZBT Mode
(Supported in single-channel mode only)

2.4.1.4 64-Bit Mode

If SDWIDTH64 is pulled high during reset, the System bus works in 64-bit mode. In this mode, the bus supports 64-bit read and 64/32 bit write operations. All bus data is assumed to be aligned on an 8 byte boundary; thus SA2 is assumed to be 0 and is not connected to the processor address bus.

Packet data transfers over the system interface follow the network byte order for packet transfer. All other data transfers are assumed to be Big Endian (that is, the most significant byte transfers at the lowest address).

The processor must take care when performing 32-bit read accesses when it is configured for a 64-bit bus width. It can perform 64-bit read accesses without any side effects; however, if a 32-bit read access is performed to an address that is part of a 64-bit register and which contains self clearing bits or which, if read, can trigger other actions, then undesirable side effects can occur. Similarly, if the processor code performs a misaligned read, indeterminate side effects can occur.

Table 14 System Bus 64-bit

	System Data Bus (SD)							
Bit Range	63:56	55:46	47:40	39:32	31:24	23:16	15:8	7:0
Byte Address	0	1	2	3	4	5	6	7

2.4.1.5 32-Bit Mode

When the SDWIDTH64 pin is tied low during reset, the system interface is configured to work in 32-bit mode. In this case, all data transfer takes place over SD[63:32] only, and SA[2] must be connected to address bit 2 of the external processor (the PM2329 interprets it as SA[2]).

Table 15 System Bus 32-bit

	System Data Bus (SD)							
Bit Range	63:56	55:46	47:40	39:32	31:24	23:16	15:8	7:0
Byte Address	0	1	2	3	na	na	na	na

2.4.1.6 Byte Ordering

PM2329 byte organization convention assumes that lower addresses contain the more significant bytes. For example, 32-bit mode, the MSB lies at address 0 and is passed on SD [63:56], while the LSB lies at address 3 and is passed over SD [39:32].

2.4.1.7 Cascade Mode Addressing

All the SA[15:3] address lines, and the SCE0*, SCE1*, and SCE2 signals, are connected in parallel to the corresponding pins of all PM2329 devices in the cascade.

When the processor accesses the PM2329 register space with SA[15] = 0, each device in the cascade will compare address lines SA[14:12] against the assigned Cascade ID (CID#) of the device. Access to a particular device in the cascade will be enabled if SA[14:12] matches the CID# of that device.

When the processor accesses the PM2329 register space with SA[15] equal to 1, SA[14:12] lines are ignored and cascade bus mechanism determines the device that responds to the access.

2.4.1.8 Local and Global Register Access

The PM2329 register space supports two access mechanisms: Local and Global. Some registers are accessed in Local access mode while others are accessible via Global access mode.

This access mechanism permits all system interface signals of the devices in the cascade (except SINT*) to be wired in parallel to the system bus, including the SCE0*, SCE1* and SCE2 system chip enable signals. While all devices in the cascade are selected by the processor at the same time (since all system chip enables are simultaneously asserted active to all devices), arbitration logic inside the PM2329 devices ensure that only one of them will respond during processor read cycles. Only the primary device's SINT* is tied to the processor; all other SINT* signals are left unconnected.

For Local access, the following conditions must be met:

- SA[15] should be driven to 0,
- SA[14:12] must match the assigned CID# of the device, and
- The lower SA signals must specify the address of the desired register.

Reads or Writes to the individual registers accessible via Local access mechanism can be performed this way in both non-cascaded or cascaded configurations.

For Global access, the following conditions must be met:

- SA[15] should be driven to 1,
- SA[14:12] are ignored by the device, and
- The lower SA signals must specify the address of the desired register.

Writes to registers accessible via Global access mode occur in the specified registers of all devices in the cascade simultaneously. Reads from registers accessible via Global access mode are either from the primary device (CID 0) or arbitrated via the cascade bus. Note that the non-cascaded configuration is a simplified case of the cascade configuration.

Further details of Local and Global register access mechanism are provided in Chapter 4, Registers and Programming.

2.4.1.9 Multiple Context Support

When the PM2329 is operating in multi-channel mode, it can support multiple contexts running on the external packet processor (see Operation Control Register, Chapter 4). In this mode, the PM2329 supports up to 32 internal channels for processing packets from up to 32 different contexts. Each context of the packet processor can access the assigned PM2329 channel, sending packets to it and obtaining the associated results, without conflict with other contexts. Channels are numbered 0 to 31.

The SCHNUM[4:0] and SCHSTB signals are provided on the PM2329 System Interface. Whenever a particular PM2329 channel has a result available, the channel number is output on the SCHNUM[4:0] pins, and SCHSTB is driven active for one SCLK cycle. This provides a direct hardware mechanism to signal the particular context of the packet processor to access the PM2329 Result FIFO for its results.

In cascade mode, the SCHNUM[4:0] and SCHSTB signals of the primary device are used. SCHNUM[4:0] and SCHSTB signals from secondary devices in the cascade should be left unconnected and ignored.

For Packet processors that do not support this type of interface, the PM2329 channels can be polled by individual contexts, or the PM2329 SINT* pin can be set up to provide a common interrupt whenever a packet's processing is complete.

In single-channel mode, SCHNUM[4:0] signals may be left unconnected. SCHSTB is asserted whenever the packet processing is complete.

2.4.1.10 Reset and Interrupts

RESET* is the asynchronous reset input to the PM2329. RESET* must be asserted for a minimum of 100 SCLK cycles after SCLK has become active. When asserted, it forces the device into the power-on/reset state. The device enters the normal operating mode after this signal is deasserted.

SINT* is the interrupt from the PM2329 to the rest of the system. Masking, and programming the conditions for assertion, are accomplished via the Interrupt Enable Register.

2.4.2 Extended RAM (E-RAM) Interface

The capabilities of the PM2329 can be enhanced significantly by the addition of the Extended RAM (E-RAM). The PM2329 E-RAM interface is designed to connect directly to a bank of synchronous pipelined SRAMs. The E-RAM stores Control Words (C-Words) and Data Words (D-Words) corresponding to each on-chip Rule.

With the E-RAM connected, the PM2329 (or a cascade of PM2329 devices) works in the extended mode of operation. This mode allows powerful OC sequencing capabilities in addition to storage of statistics and aging information on a per-OC or per-connection basis. The E-RAM data can be automatically updated by the PM2329 based upon the results of OC processing. The E-RAM can also provide a fast user-defined lookup based on the results of an OC match.

The E-RAM interface memory address bus EMA[16:0] is connected to the C-Word and D-Word E-RAM memory devices' address bus. The E-RAM interface control word data bus (ECD[31:0]) is connected to the C-Word memory device's data bus. In the cascade configuration, the EMA bus from all PM2329 devices are connected in parallel to all the memory buses of the C-Word and D-Word memory devices. The ECD bus from all PM2329 devices are tied in parallel together and to the C-Word memory data bus.

The E-RAM interface data word data bus (EDD[31:0]) is connected to the D-Word memory device's data bus. In the cascade configuration, the EDD bus from each PM2329 device is connected to the data bus of the corresponding D-Word memory device.

The E-RAM interface chip enable (EMCE*) and output enable (EMOE*) signals are always driven by the primary device. Regardless of the configuration (single or cascade), the EMCE* and EMOE* signals of the primary PM2329 device are tied to the corresponding chip enable signals and output enable signals of the C-Word and D-Word E-RAM devices. In the cascade configuration, the EMCE* and EMOE* signals of the secondary PM2329 devices are left unconnected.

The E-RAM interface write enable for the C-Word memory device (ECWE*) is connected from the primary PM2329 device only to the write enable of the C-Word memory device. In the cascade configuration, the ECWE* signals of all secondary PM2329 devices are left unconnected.

The E-RAM interface write enable for the D-Word memory device (EDWE*) is connected to the write enable of the D-Word memory device. In the cascade configuration, the EDWE* signals of each PM2329 device is connected to the write enable signal of the corresponding D-Word memory device.

In the non-cascade configuration, the Bus Request (CEMRQ) signal can be left unconnected. In the cascade configuration, the CEMRQ signals of every PM2329 device in the array are tied together.

The E-RAM width, depth and organization are all highly flexible through programmable registers. The interface can handle a maximum E-RAM size of 128K words by 256 bits/word. The access times of the E-RAM devices depends on the SCLK frequency as well as the number of PM2329 devices in cascade, since they load the memory data bus. Moreover, the loading of the C-Word SRAM is not the same as that of the D-Word SRAM, as explained in Section 2.5. The table below gives typical access times expected, but these may change depending on board design and operating conditions.

Figure 9 E-RAM Interface Timing

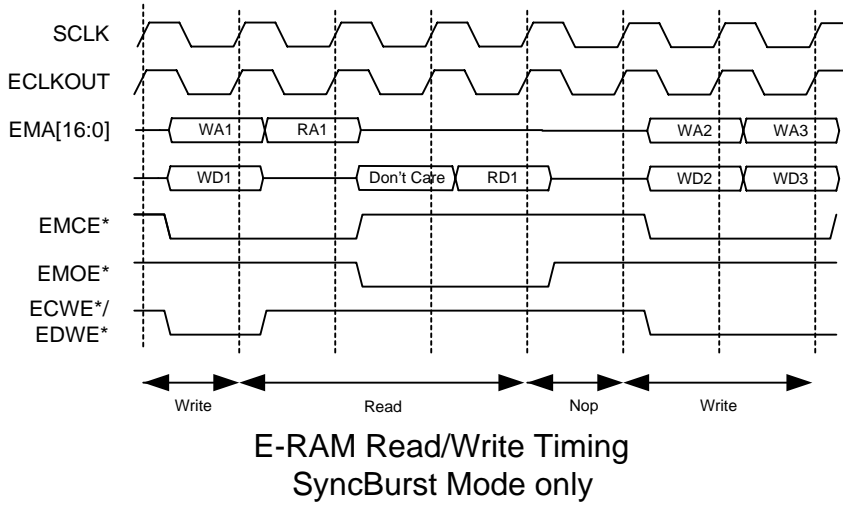
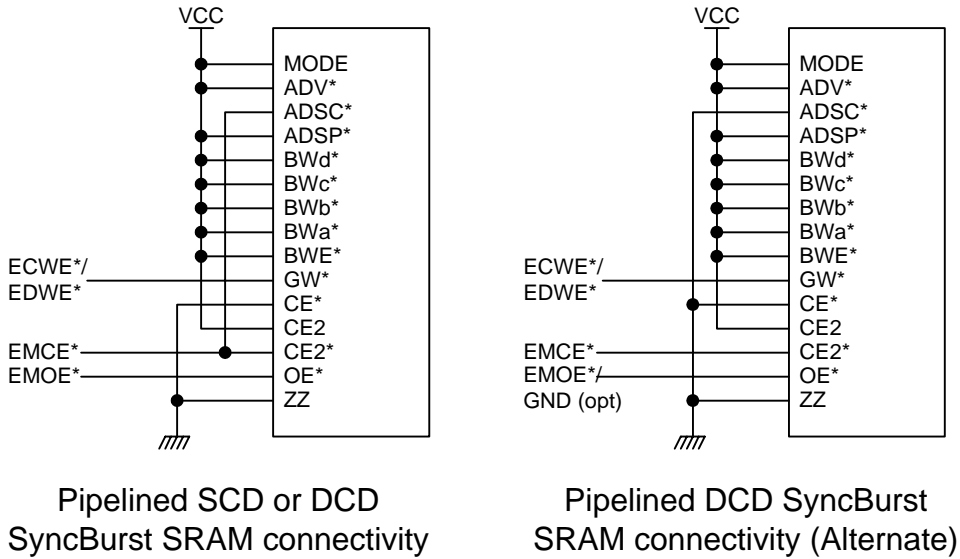


Figure 10 E-RAM Connectivity



2.4.3 Cascade Interface

The Cascade Interface allows up to 8 PM2329 devices to be connected in a cascade configuration. The interface is designed so that the entire cascade appears to the processor as a single PM2329 having 128K rules.

In the cascade configuration, the primary PM2329 should have its CID# set to 0. The remaining PM2329 devices should have their CID numbers set from 1 to n ($n < 8$) by strapping the CID[2:0] signals of each device appropriately. The primary PM2329 has the highest priority, with priority of each PM2329 decreasing with an increase in its CID number.

CID numbers must be assigned serially starting from 0 with no gaps. For a non-cascaded configuration with a single PM2329 device, the CID[2:0] pins of the device should be tied low during reset; this assigns CID# 0 to this device.

The primary PM2329 device asserts the COCS signal to start packet processing in itself and all the secondary PM2329 devices. (Packet processing and sequencing is described in the next chapter.) After processing, each PM2329 drives its COCDOUT output pin, which, when active, indicates the completion of the current operation by a particular PM2329. Since the COCDOUT pin of each PM2329 is connected to the appropriate COCDIN[n] pins of all the remaining PM2329 devices, each one is informed of the completion status of all the others.

The highest priority PM2329 chip with a match, asserts the COCM* signal. Based on a proprietary arbitration mechanism, it gains mastership of the E-RAM bus and accesses a word based on the result of the current operation.

2.5 System Configurations

2.5.1 Stand Alone Configuration

Stand Alone is a minimal configuration, where a single PM2329 can store up to 16K rules and perform packet classification and other functions based on that set of rules. Stand Alone configuration affords the smallest system component count; there is no Extended RAM and no cascaded PM2329 devices in this configuration.

2.5.2 Cascaded Configuration

A Cascaded configuration is the most powerful configuration for packet processing using a set of PM2329 devices. Cascading allows up to 8 PM2329 devices to be cascaded, thus increasing the rule space to a maximum of 128K rules.

PM2329 interfaces are designed to allow cascading without any glue logic. A cascaded set of devices is designed to appear as a single device to the external processor and other system components. CID#s are configured to uniquely identify each PM2329 device. The CID#s also define the address space (as seen by an external processor) for each PM2329 devices' local registers, and determine the participation of the PM2329 in the packet processing operations. The CID# also determines the priority of the PM2329 devices; PM2329 #0 has the highest priority, and PM2329 #7 has the lowest priority.

Figure 11 shows the Cascade Bus connectivity between cascaded PM2329 devices.

Figure 12 Single PM2329 with 32-bit/64-bit E-RAM


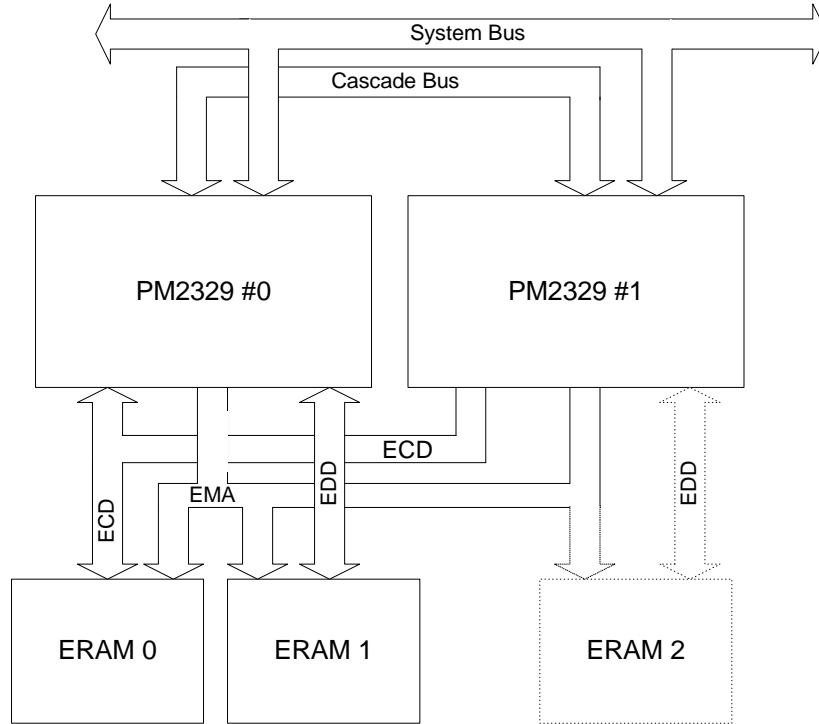
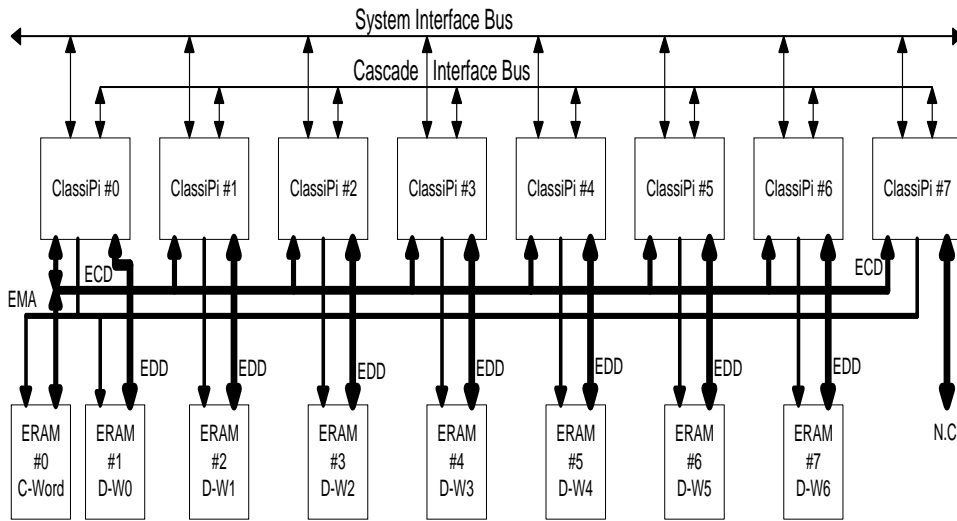


Figure 13 Two Cascaded PM2329 devices with 64 or 96 bits of ERAM



Two PM2329 Devices in Cascade with 64-bit or 96-bit ERAM

Figure 14 Maximum Cascade Configuration



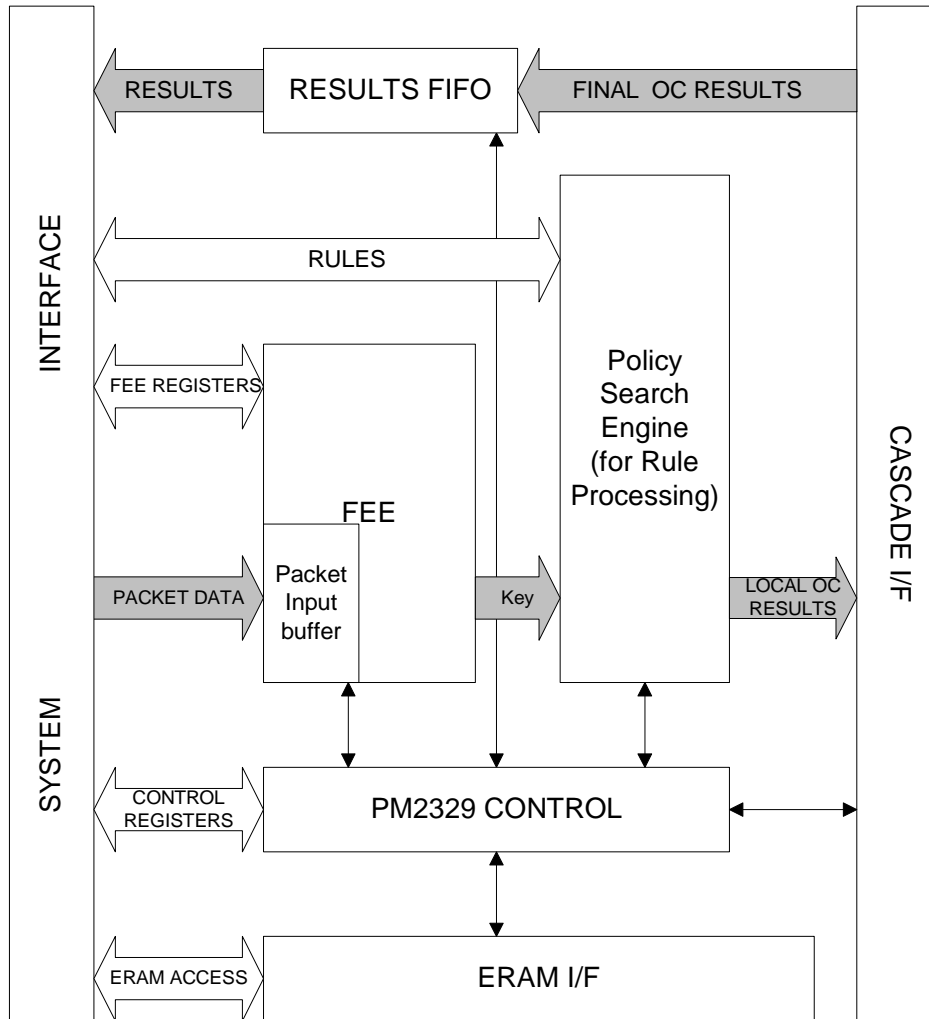
8 Cascaded PM2329 devices with a 256-bit wide E-RAM

3 Functional Description

3.1 Internal Organization and Data Flow

This section provides an overview of the major functional blocks of the PM2329 and illustrates the flow of data through the device. Later sections explain the various blocks in greater detail.

Figure 15 PM2329 Block Diagram



3.1.1 Basic Blocks

As illustrated in Figure 15, the PM2329 consists of 7 major internal blocks:

1. System Interface

This block allows access to the PM2329 registers, and implements specific registers for programming the Rule Memory, transferring packet data to the PM2329, and reading Results. Additionally, indirect memory access registers in this block provide access to the E-RAM memory.

2. Field Extraction Engine (FEE)

This block includes the 8 KB Packet Input Buffer which can store multiple packets in the single channel mode or one for each channel, in the multi channel mode. This block performs data extraction on the packet and transfers this data to the Policy Search Engine for classification. The extracted data can be either from the header or pattern data from the user data field.

3. Control Logic

This block controls the overall operation of the device.

4. Policy Search Engine (PSE)

This is the core of the PM2329 classification engine and holds the policy database (also known as the Rule Memory). This block executes Operation Cycles (OCs) and returns results.

5. Results FIFO

This block contains the Results FIFO logic. The results of packet processing are stored in the Results FIFO. In multi-channel mode, results are organized on a per channel basis in the FIFO.

6. E-RAM Interface

This block interfaces with the External RAM (E-RAM), both Control RAM and Data RAM. The E-RAM is useful for conditional sequencing of OCs and for maintaining statistics and user data.

7. Cascade Logic

This block provides the interface between multiple PM2329 devices in a cascade configuration. It resolves priorities between devices and enables processing across the cascade.

3.1.2 Data Flow

This section presents a brief overview of the data flow of a packet as it is processed. Details of packet processing are discussed in later sections of this chapter.

Figure 15 illustrates the data flow for the PM2329 during normal operation. The data flow in the device can be split into three phases:

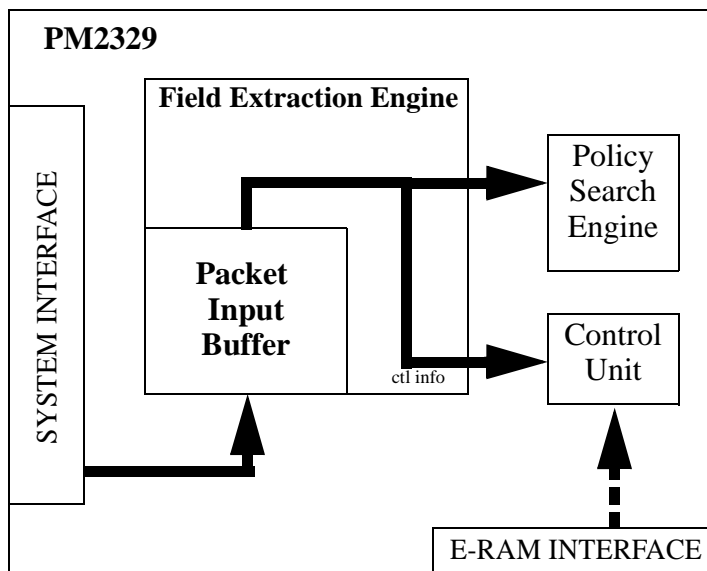
- Packet Transfer Phase
- OC Processing Phase
- Result Phase

Packet Transfer Phase

During the Packet Transfer Phase, the packet data and associated control information is transferred to the device, over the PM2329's system interface. It can be loaded into the PM2329 by a Packet Source (PS), or directly by the external processor.

The packets are stored in the PM2329 Packet Input Buffer--an input FIFO which is the first element in the data path inboard from the system interface.

Figure 16 Simplified Data Flow



The Packet Input Buffer's physical size is 256 bytes deep times the total number of channels supported by the PM2329 (256 x 32 bytes). A Packet Source can load the Packet Input Buffer using either a single write port (the Packet Buffer Input Register, or PBIR), or alternatively with SRAM-like addressing, wherein a block memory transfer mechanism is utilized in which the destination address increments.

In single-channel mode, the Packet Input Buffer acts like a FIFO for multiple packets. In multi-channel mode, several packets may be transferred collectively into the Packet Input Buffer. If a cascade

configuration is used, a transfer of packet data or packet control information to the Packet Input Buffer loads all PM2329 devices in the cascade simultaneously.

For further details regarding the Packet Input Buffer size and functionality, refer to the Channel Assignment Register and Packet Buffer Input Register descriptions in Chapter 4.

The Packet Transfer Phase of processing is not complete until key data is extracted from the newly-acquired packet(s), including IP/TCP/UDP fields (header fields), and data within the packet at given offsets which must undergo pattern searches. To extract such data for processing, the Field Extraction Engine (FEE) accesses the packets in the order in which packet transfers are completed. The FEE transfers the extracted data fields to the Policy Search Engine (PSE) for classification with respect to a predetermined subset of the defined Rules.

In addition to the needs of the FEE, the PM2329 Control Unit needs the control information associated with every packet, in order to start the packet processing operation. This control information defines the search space in the Rule Memory, the sequence of packet processing operations, the data upon which to act, and the action to be taken after each step in the packet processing. The control information may come from the Packet Source (PS) or from an external Packet Processor (PP), through the System Interface. The control information can also come from the E-RAM, if the OC execution is controlled by E-RAM sequencing.

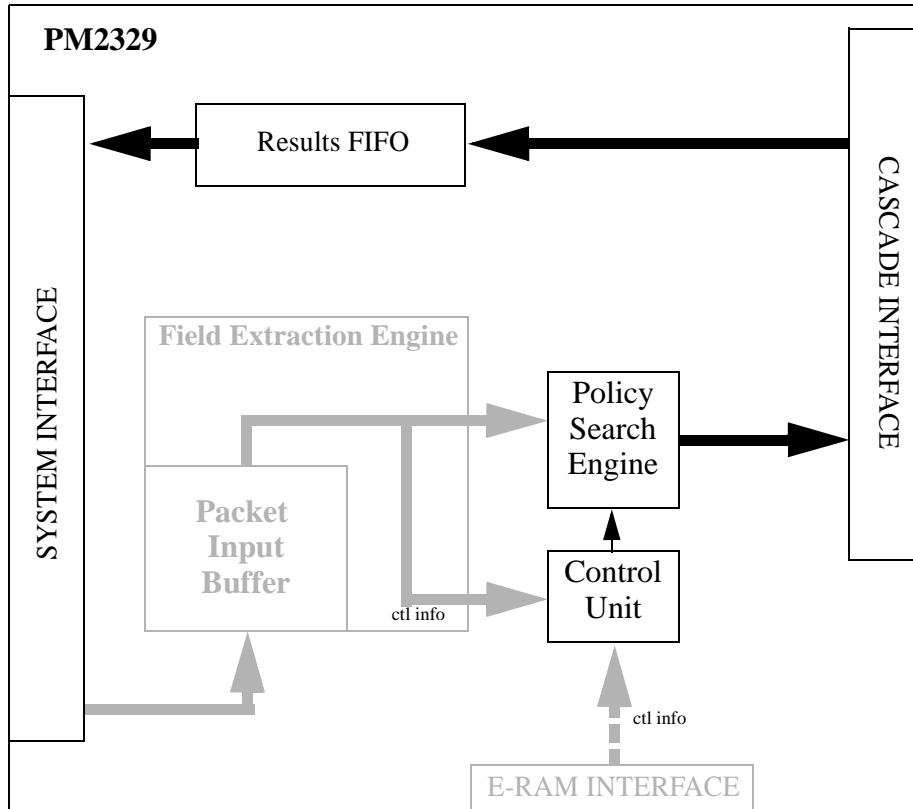
OC Processing Phase

Once the packet data has been received, the relevant fields have been delivered to the Policy Search Engine, and the control and sequencing information is in the hands of the Control Unit, the PM2329 starts the OC Processing Phase. The Policy Search Engine applies select rules from a master Rules Memory in the order specified by the Control Unit, to packet header/data it received from the FEE. Application of a rule is essentially a pattern matching operation. (Section 3.2 describes the Field Extraction Engine in detail.)

Result Phase

During the Result Phase, the results of successful pattern searches are transferred from the PSE into the OC Results FIFO. If an applied rule yields a match, the PM2329 uses the Cascade Interface signals and bus to assert an OC Match signal and return the results of that comparison, with statistical data relating to that rule, in the Results FIFO. The System Interface allows the external processor access to the results.

Figure 17 Simplified Data Flow--OC Processing and Results Posting



Additionally, if the device is extended with External RAM, the Control Unit can retrieve from this E-RAM and deliver to the Results FIFO the D-Words associated with the rule that produced a match. The Control Unit can also update the Data portion of the E-RAM word (E-Word), based on update instructions in the control information for that packet.

The cascade bus is in operation during the OC Processing and Results phases, in order to co-ordinate result processing activity between the cascade of devices. The System Interface allows access to the Results FIFO by the external processor.

3.1.3 Context Support

The PM2329 has been designed to work with high speed network processing applications where multiple processors or multiple tasks running on the same processor can send packets for processing and then retrieve only their own results. The PM2329 supports this through the use of channels. The concept of a channel is a path through the PM2329 from packet entry to result buffering. Thirty two independent channels are supported, allowing for up to 32 independent contexts to utilize the PM2329. Interleaved writing of packets by the contexts, as well as interleaved access to the context-based results, are possible; this can be accomplished without the use of any semaphore mechanism between the contexts.

The multi-channel operation of the PM2329 is enabled by setting a bit in the Operation Control Register to enable the multi-channel mode. Note that packet data transfer when multi-channel mode is enabled must be done using processor cycles. Packet Source (DMA) transfers are not supported in this case.

In the multi-channel mode, the 8 KB input buffer is organized as 32 segments of 256 bytes each. Packet Input Buffer Registers for each segment are provided so that the external processor(s) may write to any segment (based on the context that is writing the packet). Similarly, the 128 entry results FIFO gets organized into 32 segments of 4 entries each. The combination of the input buffer and the associated results segment defines a channel. These channels are numbered 0 through 31. Consequently, 32 simultaneous channels can co-exist and be used by 32 independent tasks/contexts. The registers used to load the packet, retrieve results, and check the packet processing status are organized in a per-channel fashion into channel-oriented register groups. This allows multiple contexts on the external processor to access their channel registers simply by using the appropriate base address for that channel group.

The PM2329 processes the packets in the order in which packet transfers complete. Results from packet processing are then sent to the appropriate segment in the Results FIFO. The results for a channel are made available through an OC Results FIFO Read Port for each channel.

The SCHNUM[4:0] and SCHSTB pins on the PM2329 provide the required support to accomplish a hardware-based handshake between the PM2329 and the processor. For processors equipped to handle such handshake signals, this interface awakens a context whenever its packet processing results are ready. If the processor does not implement this interface, a Channel Status Register can be read by any context to determine the status of the Packet Input Buffers and result FIFOs for that channel.

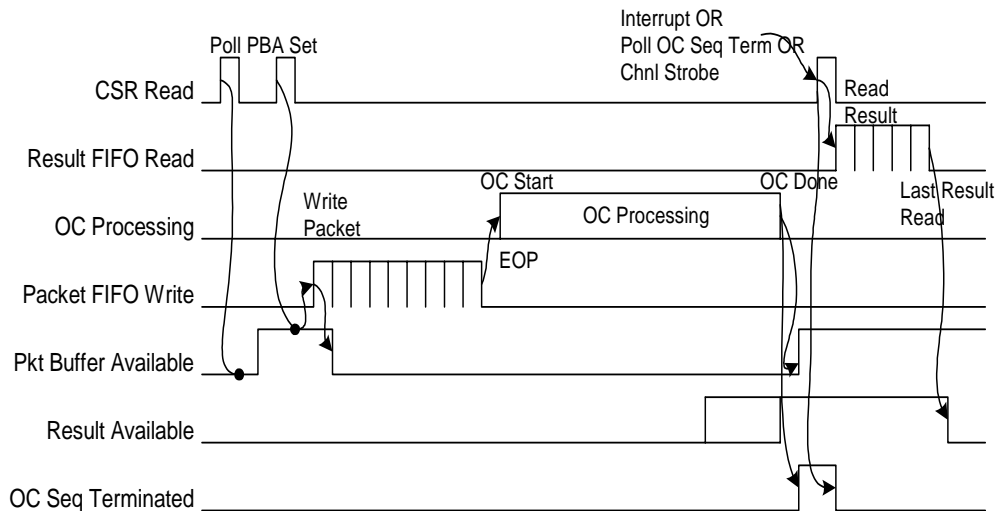
Concatenation of segments is supported to allow any channel to transfer a packet larger than 256 bytes or to accumulate more than 4 results. This is accomplished using the Channel Assignment Register. This allows any number of adjacent channels to be concatenated to work as a single larger channel. When 4 channels are concatenated, for example, the Packet Input Buffer will appear to be 1 KB in length and the OC Results will appear to be a 16 deep FIFO.

3.1.4 Channel Input, Output and Status Mechanism

Each channel (0 through 31) has its pre-assigned packet input buffer available for loading after reset and after the previously loaded packet has been processed. Since the write ports associated with the packet input buffers share the processor bus, only one of them can be written to at a time; however, they can be loaded in an interleaved fashion. Once a packet has been loaded by a context, the context must wait until the processing for this packet is complete as indicated by OC Sequence Terminated bit, before loading the next packet for its assigned channel.

The diagram below shows the operation of status bits used for Packet input and Result output.

Figure 18 Packet Input, Result Output, and associated Status Handshake



(Status handshake associated with Result Output involves PBA, RA and OCST signals)

3.2 Field Extraction Engine (FEE)

3.2.1 Introduction

The Field Extraction Engine (FEE) represents the first stage of processing for all packet classification operations in the PM2329.

Packet Input Buffer

The FEE includes an 8 KB Packet Input Buffer which can store up to 32 packets simultaneously. Whenever the external device (processor or other packet source) sends packet data to the PM2329, the packet is first deposited in the Packet Input Buffer.

- Packets may transfer to the PM2329 while earlier packets are still being processed.
- Packet bytes may begin to transfer to the PM2329 before other packets have completed transfer.
- Multiple packets may be stored concurrently in the Packet Input Buffer.

Key Extraction

Packet transfer for any packet completes when the PM2329 receives an EOP indication for that packet. Packets are then taken for further processing in the order in which the packet transfers complete. Since the FEE's task is to glean one or more "keys" from the packet which will be compared against a series of rules, it extracts the relevant header or data key from the packet, constructs the key(s), and transfers the key(s) to the PSE for processing.

A key can be one of three main types:

PI Field Attribute Key: The Attribute segment of the packet's PI field (or the Attribute segment of the pre-programmed PI register) becomes the key to be compared, according to the sequence of rules identified for the OC.

Header Key: The IP, TCP and UDP information from within the packet's header are concatenated to construct a 108-bit Packet Header key, which is compared (with appropriate masking if desired) in the Rules Processing of the OC. The FEE parses and extracts header fields from Ethernet at Layer 2, IP at layer 3, and TCP/UDP at layer 4. Within Ethernet, it can decode Ethernet II, Ethernet 802.3 and Ethernet 802.1p/q.

Packet Data Key: The FEE also has the capability of extracting a given length of packet data starting at a given byte offset, and constructing short or long keys. This feature can be used for searching for a set of specific strings within a packet, whereby the FEE repeatedly extracts data using an offset shifted by one additional byte for each iteration, in the forward or reverse direction, thereby allowing pattern searches in either direction. A Packet Data field can include payload bytes, header bytes, or both.


The FEE can extract multiple fields for any packet, and submit them all to the PSE for processing. There can be one unique PI Field Attribute key, one unique 108-bit header key, and many unique Packet Data keys of varying lengths and beginning at varying offsets within the header/payload parts of the packet.

3.2.2 Supported Packet Formats

The FEE performs extraction of fields and key construction as data is being received into its buffers. The FEE uses the information in the PI field (or the PI register if no PI field exists in the packet) to direct its extract efforts. Figure 19 and Figure 20 show the packet formats recognized by the PM2329.


Figure 19 Packet Formats Supported by PM2329

L3 Header Extraction Enable = 0



L3 Header Extraction Enable = 1





Ethernet Framing Enable = 0



* If L4 Extraction Enable = 1 & Protocol = TCP/UDP

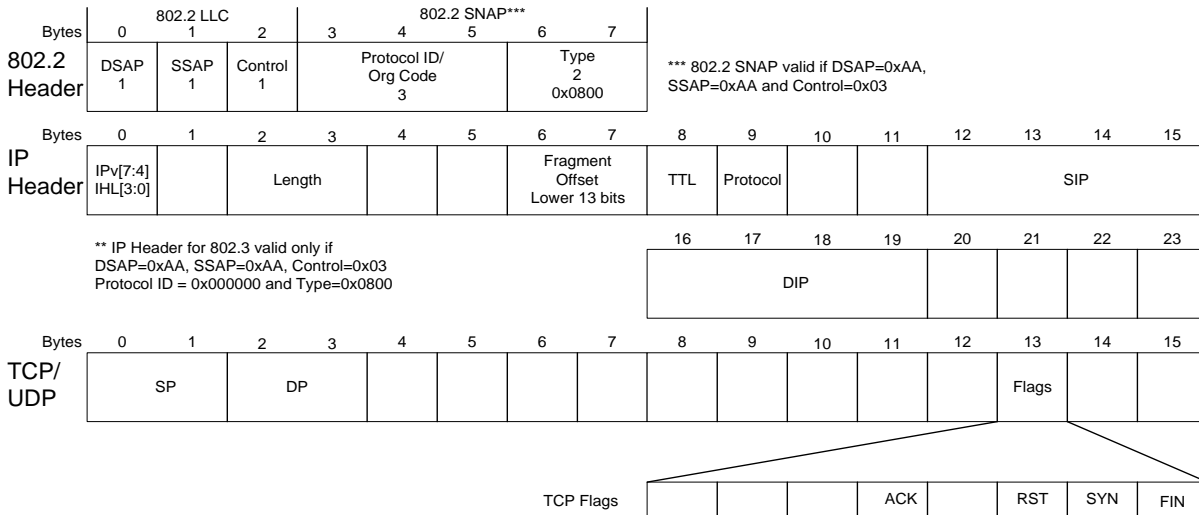
L3 Header Extraction Enable = 1

Ethernet Framing Enable = 1



* If L4 Extraction Enable = 1 & Protocol = TCP/UDP

Figure 20 Headers Formats within Packet



3.3 Control Unit

The Control Unit (or Control Logic Block) consists of registers required to control the operation of the PM2329 and its various internal blocks, and state machines to perform the control functions.

The registers of the Control Logic block include various configuration and status registers for controlling the OC processing, as well as the 64-entry OC Descriptor table. A brief description of some of these key registers follows:

- **Operation Control Register** - This register controls the operation modes of the PM2329.
- **Interrupt Enable Register** - This register specifies the mask bits for the various conditions that can cause the PM2329 to assert an external interrupt.
- **Status Register** - A register that provides common information regarding the state of the single or cascaded PM2329 devices.
- **OC Descriptors** - Registers that define the processing of the OC.

Details of register definitions, as well as OC Descriptor format and usage, is available in 4.

3.4 Policy Search Engine (PSE)

The PM2329 Policy Search Engine (PSE) is the work-horse of Rules Processing. It can perform powerful policy-based search operation sequences, applying rules from the 16K-deep Rule Memory and comparing a specified sequence of rules with the extracted key information from the FEE. It generates detailed results of those searches, and returns those results (under control of the Control Unit) through the Cascade Interface to the Results FIFO.

Operation Cycle (OC)

The PSE performs Operation Cycles (referred to throughout this data book as OCs) on packets in the order that their field extraction is complete. The processing of any one search using one key constitutes a single OC. As there could be multiple keys extracted from a packet, and as each key can be subjected to multiple searches, a packet can utilize many OCs for full processing.

An OC is complete when the key is compared against a specific rule, and the results are queued in the Results FIFO. The time duration of an OC depends therefore on the availability of space in the Results FIFO; if the OC must wait for space, it remains active until space is available.

OC Sequencing

The order in which searches are performed for a particular key (that is, the order in which OCs execute) is called OC Sequencing. There are 2 basic types of OC Sequencing: OCC Sequencing and E-RAM Sequencing. Within each type, sequencing can be Automated or Processor Controlled.

Chapter 5 describes the sequencing modes in detail.

3.4.1 Rule Memory

The Rule Memory can store up to 16,384 rules in on-chip memory. This Rule Memory is composed of cells arranged into rows and columns. Each rule is stored in a location known as a cell. The PM2329 can perform multiple and simultaneous data operations on the extracted fields of the input packet and the specified set of rules (cells) in the Rule Memory.

Refer to Chapter 5 for a description of the rules and OC sequencing.

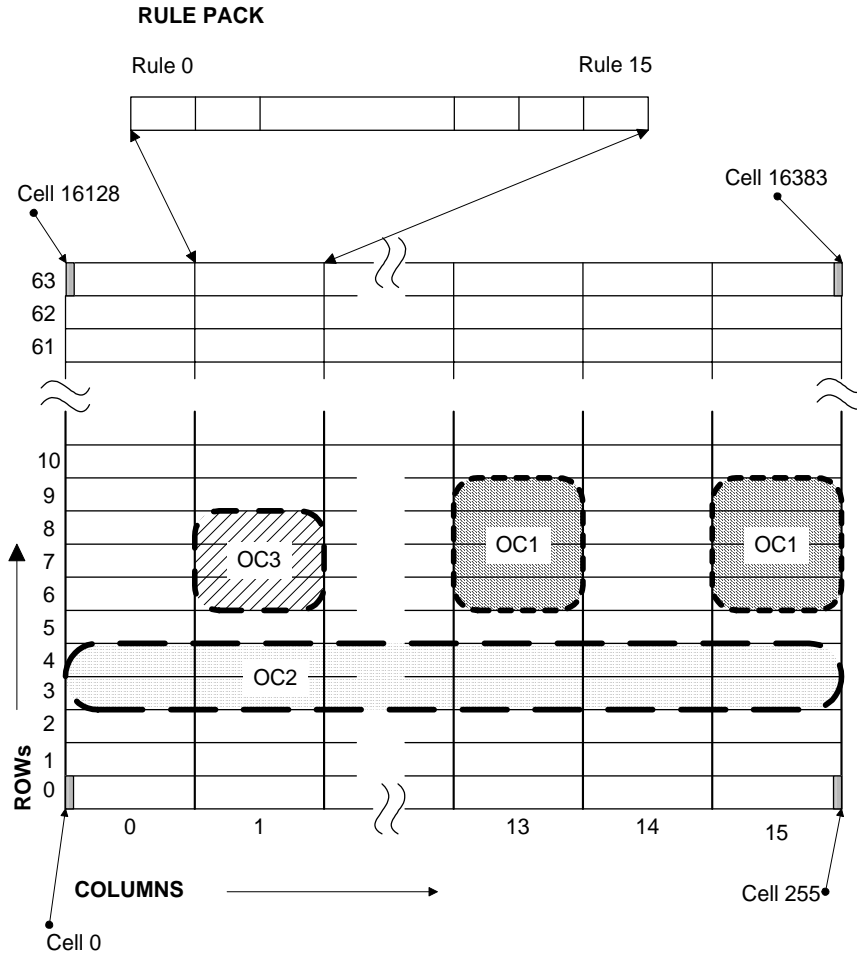
3.4.2 Cell Organization

The organization of the cells in the Rule Memory is shown in **Figure 21**. The smallest element is a cell, which contains a single rule. Rules are grouped into "rule packs" of 16 cells (that is, up to 16 rules) each. The rule pack represents the smallest number of cells that can participate in any OC. Therefore, partitioning of rules among different OCs is done in "rule pack" units.

Rule packs themselves are further organized into a two dimensional array consisting of 64 rows and 16 columns. Each row contains 16 rule packs, and each column contains 64 rule packs. **Figure 21** illustrates.

Each cell can be uniquely identified by the row number (0-63) and column number (0-15) of the rule pack, and finally by its position (0-15) within the rule pack to which it belongs. The cell address is a 14-bit number where the most significant 6 bits represent its row number, the middle 4 bits indicate its column number, and the least significant 4 bits describe its position within the rule pack. The unique cell number of each cell ranges from 0 to 16,383 for a single PM2329 device.

Figure 21 Organization of Rules



3.4.3 Priority of Rules

There is an inherent priority structure built into the Rule Memory, with each of the 16K rules having a distinct priority. The priority of a rule is a function of the cell number. Rules at lower cell numbers have higher priority. Thus, priority is first determined by its row number, then its column number, and lastly by its position within the rule pack. Rule 0 has the highest priority and Rule 16,383 has the lowest priority.

The priority structure allows prioritization of match results when multiple rules match the packet key data. In the case of multiple matches in the PSE, the matched rule (cell) numbers are returned in the order of their priority.

3.4.4 Rule Partitions

A partition is defined as a collection of rule packs. The selection of rule packs that belong to a partition is decided by specifying row numbers and column numbers. The row numbers are specified as a range, while the column numbers are defined by an enable bit for each column. For example, Rows 6 to 9 and Columns 13 and 15 can be defined as a partition. Separately, Rows 3 to 4 and Columns 0, 1, 2... 15 can be defined as

a second partition. Both these partitions are shown in Figure 21 as "OC1" and "OC2" respectively. Partitions are used when running operation cycles (OCs). An OC executes over a specific partition, and typically all rules in a partition participate in the same kind of operation. From a software point of view, a partition will consist of cells storing rules relating to a specific networking operation. For example, routing rules, firewall rules, QoS rules, etc.

In general, partitions can overlap; for example, two partitions can point to the same rule space. This is useful when multiple software contexts in a network processor access Rule Memory; partitions can then be associated with the appropriate context. Another situation where Rule space sharing is beneficial is in MAC address tables where the source and destination MAC addresses can be maintained in one table.

3.5 E-RAM Operation

The basic functionality of a single PM2329 or a cascade of multiple PM2329 devices can be greatly enhanced by the addition of the Extended RAM. The E-RAM may be used for storing statistical counts, aging, and per-connection state information. It can also be used to implement conditional sequencing of multiple OCs. OC Sequencing is a powerful way of selectively running multiple classification operations on a single packet.

The PM2329's E-RAM interface is designed to work with or without cascading. It allows external synchronous pipelined SRAM to be gluelessly interfaced to the PM2329. The PM2329 supports up to a maximum of 128K word deep E-RAM, allowing up to a word per rule when using 8 cascaded PM2329 devices (128K rules). The width of the E-RAM can be decided based upon the system requirements, with the maximum width specified by the number of cascaded PM2329 devices. A single PM2329 can interface to a 64-bit-wide E-RAM, and each additional PM2329 in the cascade allows an additional 32 bits to be added to the E-RAM word width, up to a maximum of 256 bits with 7 PM2329 devices in cascade. PM2329 chip #7 (the eighth chip of a cascade) does not support an additional 32 bits of E-RAM width. Figures in Chapter 2 show some typical configurations using E-RAM.

3.5.1 Organization of E-RAM Words

Every E-RAM Word (E-Word) is composed of the following fields:

Extended Memory word (E-Word)				
Control Word (32 - bit) C-Word	Data Words (D-Words)			
	Byte Count (32-bit)	Packet Count (32-bit)	Timestamp/State (24/8- bit)	User Defined (32-bit)

Further details of E-Word format and E-RAM based OC Sequencing is provided in Chapter 5.

3.6 Cascade Operation

The Cascade bus allows up to 8 PM2329 devices to work in parallel, appearing to the external processor as a single large PM2329 chip. The 16K-rule memory size of a single PM2329 is thus expandable up to 128K rules on 8 PM2329 devices. The rule set of an Operation Cycle (OC) can be spread across several PM2329 devices allowing greater parallelism in operation across a larger rule memory space.

Each PM2329 in the cascade has its CID pins uniquely wired to the supply rails. These CID pins are sampled on reset, thereby assigning a unique CID# (0 to 7) to each PM2329. Rules participating in an OC running across multiple PM2329 devices are now prioritized by CID# of the PM2329 to which they belong. Thus, for any OC, all the rules within PM2329 #0 will have a higher priority than rules for that OC within PM2329 #1. Priority decreases with increasing CID#; the priority of a rule match is automatically determined among the PM2329 devices by communication over the Cascade interface.

OCs can be processed as "single hit" or "multiple hit" OCs, referring to the number of matches that should be processed before execution is terminated and a result is returned. In the case of a single-hit OC, a single result is generated--the result of the first match found. In case of a multi-hit OC, a result is generated for each match found. A single-hit OC execution is terminated as soon as the first (highest priority) match is found. In case of a multi-hit OC, all the rules in the OC are executed and all matches are returned (in order of priority) The external processor need not determine in which PM2329 a match has occurred, since the appropriate information is returned when it reads the globally accessible Results FIFO.

4 Registers

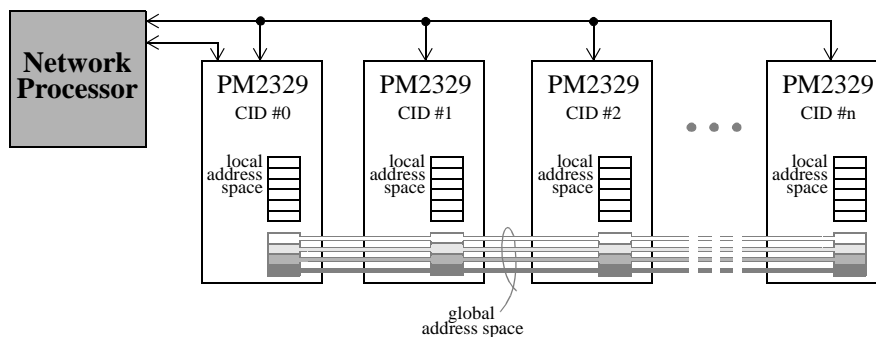
4.1 PM2329 Access Modes

The mechanism to access the internal status and programmable locations within the PM2329 device permits simple and flexible access in a variety of configurations. The access mechanism provides a consistent interface to the external processor in both single-chip or cascade (multi-chip) configurations. The PM2329 can be connected to either a 64-bit or a 32-bit interface on the packet processor. The access mechanism allows multiple contexts on the packet processor to communicate easily with an assigned channel within the PM2329.

4.1.1 Address Space

The PM2329 address space is divided into Local and Global spaces. The concept of local vs. global addressing is relevant to the cascade (multi-chip) configuration. Local space in general represents registers involved with processing functions that are specific to a particular PM2329 device. Global space in general represents registers involved with processing functions that span all PM2329 devices in the cascade. A conceptual view of Global and Local register space is considered in Figure 22.

Figure 22 Local vs. Global Register Space; Conceptual View A



Each register within the PM2329 is mapped either to Local or Global address space. The access mode is determined by the state of the SA[15] address line in any given access cycle. If SA[15] is low, the access is to the Local space. If SA[15] is high, the access is to the Global space. The other address lines dictate which specific PM2329 device, and which specific register within the prescribed space of that device, is accessed.

The PM2329 has three pins (CID[2:0]) which are sampled at reset to assign a unique PM2329 ID Number (CID #N) to each chip in the cascade. The value on these pins during reset must be set up to assign a serial number from 0 to N to each PM2329 device, where N varies from 0 to 7. In single chip configuration, the PM2329 is assigned Cascade ID #0 (or CID #0).

When the Local address space is accessed (SA[15] is low), each PM2329 compares the address lines SA[14:12] with its CID #N. If they match, then the particular PM2329 device is accessed.

When the Global address space is accessed (SA[15] is high), all PM2329 devices participate in the operation. When a Global Write is performed, all PM2329 devices are written. When a Global Read is performed, all devices in the cascade respond in turn, while an internal address resolution mechanism within the PM2329 devices avoids contention.

Once initialization is complete and local PM2329 registers are configured, the network processor thereafter treats the bank of cascaded PM2329 devices as a single PM2329 device; whether a single device or a cascade of devices is implemented, there is no difference in the interface protocol. (Thus, in discussion throughout this data book, references to “the PM2329 device” can mean a single device or multiple cascaded devices.)


4.1.2 Channels

The PM2329 implements a set of thirty two channels to support independent contexts of the external network processor (a network processor divides packets among its various processing cores; each stream of packets from any given core represents a context to the PM2329 device[s]). PM2329 channels can be assigned to these contexts; each context can send a packet to PM2329 in an independent manner and then fetch its own classification results. Packet input activities, as well as access to results, can be interleaved between the various contexts.

To facilitate packet transfers, on board each PM2329 device is a Packet Input Buffer--an input FIFO which is the first element in the data path inboard from the network processor (system) interface. The Packet Input Buffer’s physical size is 256 bytes times the maximum number of channels supported in hardware, or 256 x 32 bytes. The Packet Input Buffer can be accessed using either a single write port (called the Packet Buffer Input Register, or PBIR), or with SRAM-like addressing.

The Packet Input Buffer resides in Global address space. Thus, when packets are written to a Packet Input Buffer, they are written simultaneously to all Packet Input Buffers of all cascaded PM2329 devices. At any given time, then, all Packet Input Buffers are identical, and all PM2329 devices in the cascade have the opportunity to utilize or classify each packet.

Figure 23 PM2329 Packet Input Buffer




To facilitate the return of processing information back to the network processor’s context, each PM2329 device has a Results FIFO that can hold 8 entries per channel, or $(32 \times 8) = 256$ entries total. Like the Packet Input Buffer, the Results FIFO can be accessed using either a single write port, or with SRAM-like addressing. The Results FIFO also resides in Global address space.

The PM2329 channels are implemented through segmentation of the Packet Buffer and the Results FIFO. Channel segments may be concatenated to allow larger packet or result storage for a given channel (with correspondingly fewer channels supported).

4.1.3 Channel Register Blocks

PM2329 registers involved in packet transfer, packet processing status and packet results are organized into Channel Register blocks. There is one block of Channel registers for each of the 32 possible channels. These registers reside in Global address space. Thus, every PM2329 in a cascade has them, and a write to them by the network processor affects all devices in the cascade.

Figure 24 Channel Register Blocks



As illustrated in Figure 24, each channel's register block is spread over four address regions identified by base addresses 0 to 3. However, the specific registers of each channel have a unique offset relative to base addresses. That is, for any given channel, the offset distinguishes the channel's registers from corresponding registers of other channels. Thus, each context on the packet processor can access a particular channel easily, and perform its own packet processing independently. (The four base address spaces differ in size, so a channel's offset from Base 0 will differ from its offset from Base 1, or 2, or 3.)

Figure 25 and Table 18 show the complete Register Block address space.

4.1.4 Direct and Indirect Access

The registers of the PM2329 are accessed using direct addressing by directly reading or writing to these registers. Memory locations controlled by the PM2329 (the E-RAM and PM2329 Rule Memory), however, are accessed using Indirect Addressing. This is done by writing and reading a set of Address, Data and Command Registers. The following registers provide indirect access to the Rule Memory:

- Rule Indirect Address Register,
- Rule Indirect Data Register, and
- Rule Indirect Command Register.

The following registers provide indirect access to the E-RAM:

- E-RAM Indirect Address Register,
- E-RAM Indirect Data Register, and
- E-RAM Indirect Command Register.

Both sets of indirect access registers support an Auto-Increment mode which is optimized for writing or reading a block of data. Enabling auto-increment avoids repeated writes to the address and command registers. Also, if initializing a block of locations with the same data, multiple writes to the data register can be avoided.

4.2 Register Interface

4.2.1 Programmable Register Overview

The table below lists the registers of the PM2329 and their addresses. These are all addressable using Direct Addressing. The table shows the register address for a 64-bit packet processor. The Source and

Table 17 PM2329 Register Memory Map

#	Register Name	Mode	Local or Global	Source (Processor Read)	Destination (Processor Write)	Address
1	Local Configuration Register	R/W	Local	CID n	CID n	n000h
2	Rule Indirect Command Register	R/W	Local	CID n	CID n	n008h
3	Rule Indirect Address Register	R/W	Local	CID n	CID n	n010h
4	Rule Indirect Data Register Set Reg 0 Reg 2 Reg 4	R/W	Local	CID n	CID n	n018h n020h n028h
5	OC Descriptors Upper OCD 0 Lower OCD 0 Upper OCD 1 Lower OCD 1 ... Upper OCD 63 Lower OCD 63	R/W	Local	CID n	CID n	n400h n408h n410h n418h .. n7F0h n7F8h
6	E-RAM Indirect Data Register Set C-Word : D-Word 0 D-Word 1: D-Word 2 D-Word 3 : D-Word 4 D-Word 5 : D-Word 6	R/W	Global	CID 0 CID 1: CID 2 CID 3 : CID 4 CID 5 : CID 6	CID 0 CID 1 : CID 2 CID 3 : CID 4 CID 5 : CID 6	8200h 8208h 8210h 8218h
7	E-RAM Indirect Command Register	R/W	Global	CID 0	CID 0~7	8220h
8	E-RAM Indirect Address Register	R/W	Global	CID 0	CID 0~7	8228h
9	E-RAM Configuration Register	R/W	Global	CID 0	CID 0~7	8230h
10	Interrupt Enable Register	R/W	Global	CID 0	CID 0~7	8238h
11	Status Register	RO	Global	CID 0	NA	8240h
12	Operation Control Register	R/W	Global	CID 0	CID 0~7	8248h
13	Channel Assignment Register	R/W	Global	CID 0	CID 0~7	8250h
14	OC Conductor Register	R/W	Global	CID 0	CID 0~7	8258h
15	Packet Information Register	R/W	Global	CID 0	CID 0~7	8260h
16	Timer Register	R/W	Global	CID 0	CID 0~7	8268h

Table 17 PM2329 Register Memory Map

#	Register Name	Mode	Local or Global	Source (Processor Read)	Destination (Processor Write)	Address
17	Alternate OC Conductor Register	R/W	Global	CID 0	CID 0~7	8270h
18	Channel Register Block Base Addresses See Table 4.2 for further details	See Table 19	Global	See Table 19	See Table 19	8400h - 8600h, C000h - CFFFh, D000h - D7FFh, E000h - FFFFh

Destination columns show the PM2329 devices participating in the read or write operations, respectively. CID n indicates the address device n responds to the access, CID # or #-# indicates the device with the specified ID responds to the access. In the address column, the addresses for local registers are shown as 'nNNN' (for example, 'n000'). To access the register of a specific PM2329 device in the cascade, substitute a number from 0 to 7 for 'n', depending on the device to be accessed. For example, in a single PM2329 system or to access the Primary PM2329 in a cascade, 'n' would equal '0'.

PM2329 registers generally fall into one of three categories:

- Setup and Control,
- Packet Input, and
- Result Output.

They must be written to or read from in specified sequences and at appropriate times to ensure proper device operation. For example, improperly sequenced or ill-timed modification of setup and control registers when a packet is being processed by the device can result in operation failure, the new value being ignored, or some other unpredictable behavior.

All PM2329 registers are 64-bit locations; however, they can be accessed either in 64-bit or 32-bit mode. If the processor executes a 64-bit access using a register address specified in Table 17, it will access the entire register. If the processor executes a 32-bit access using the same address, it will access the upper half of the register (bits 63:32). To access the lower half (bits 31:0) of the register using 32-bit access, the value '04h' should be added to the address specified in Table 17. Bits in every register are aligned toward bit 0 (that is, right-aligned), so for all registers with bits in the lower half only, 32-bit accesses should be at the address+04h location.

Figure 25 provides a graphical representation of the address space.

Figure 25 PM2329 Address Space

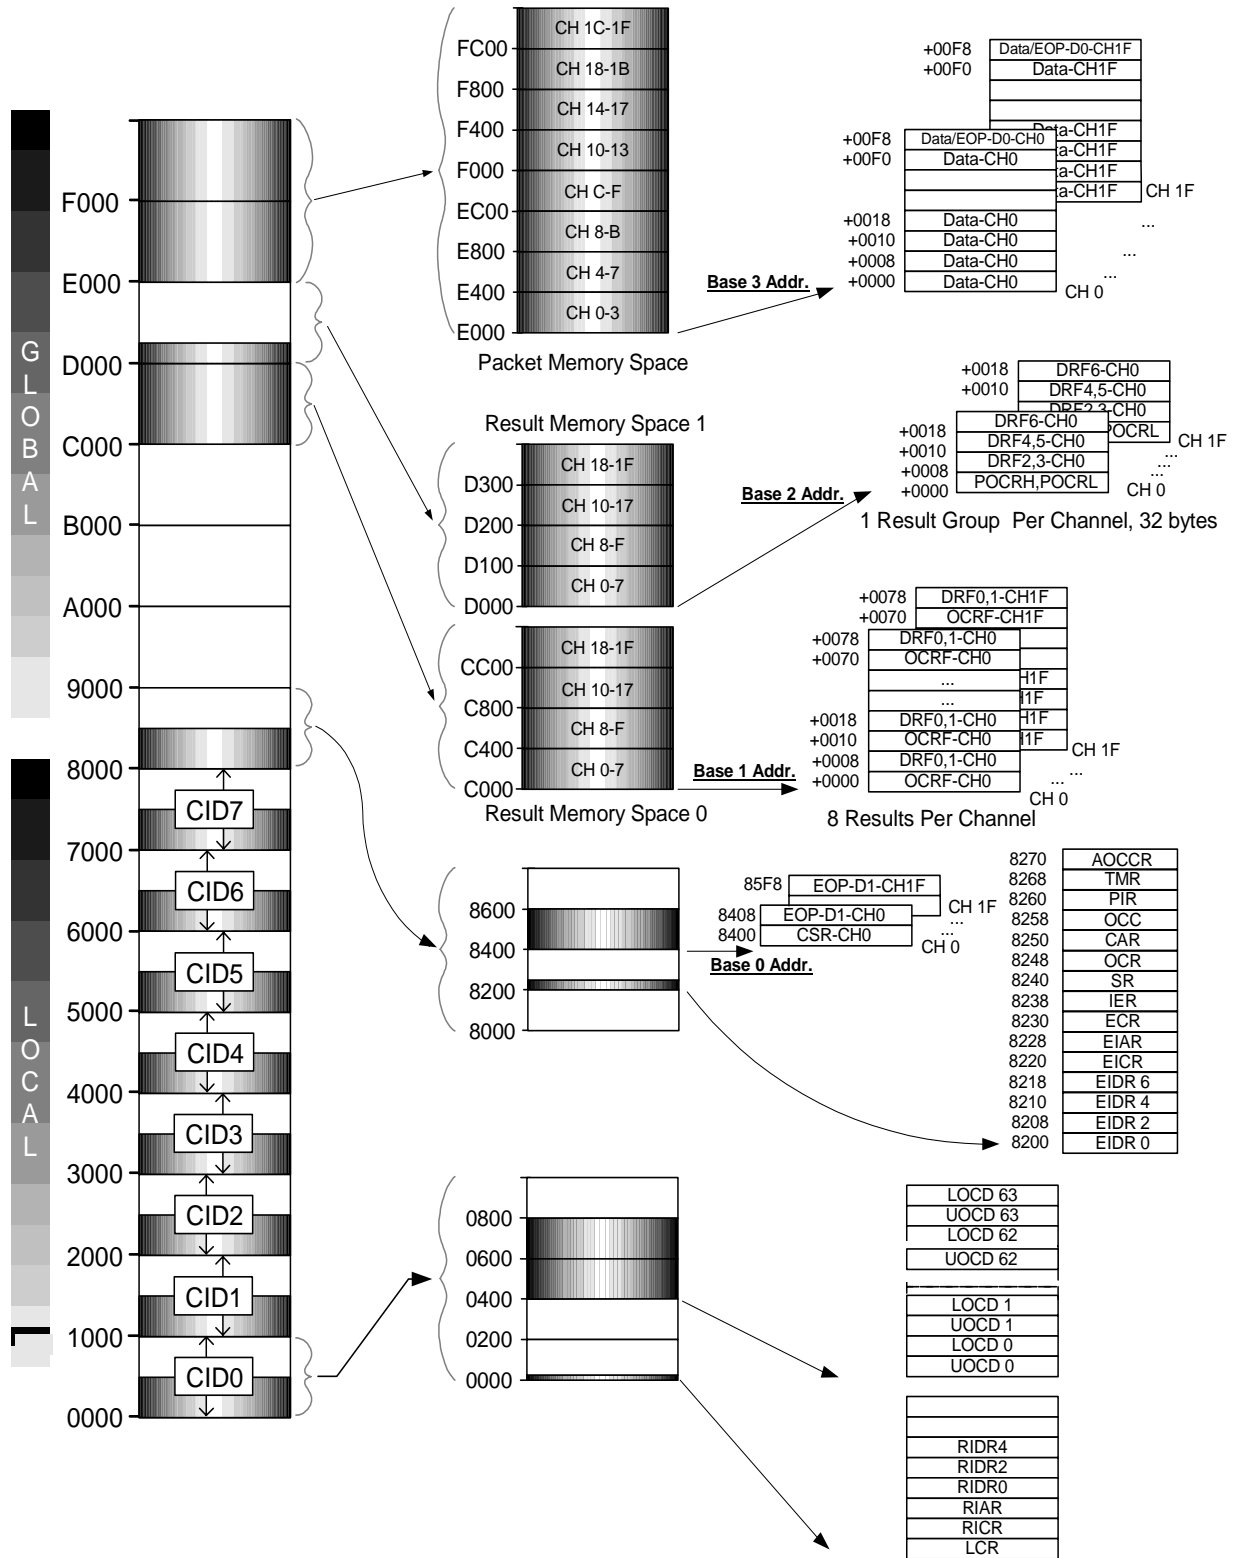


Table 18 summarizes the address assignments. Note that all reserved bits must be written with zeroes. When read, reserved bit values returned will be invalid and must be masked out by the processor.

Table 18 Channel Register Block Base Addresses

Channel Number	<u>Base 0</u>	<u>Base 1</u>	<u>Base 2</u>	<u>Base 3</u>
	CSR and EOP-D1	OC Result and Data Result 0/1	Data Result 2/3, Data Result 4/5, and Data Result 6	Packet Input and EOP D0
0	8400h	C000h	D000h	E000h
1	8410h	C080h	D020h	E100h
2	8420h	C100h	D040h	E200h
3	8430h	C180h	D060h	E300h
4	8440h	C200h	D080h	E400h
5	8450h	C280h	D0A0h	E500h
6	8460h	C300h	D0C0h	E600h
7	8470h	C380h	D0E0h	E700h
8	8480h	C400h	D100h	E800h
9	8490h	C480h	D120h	E900h
10	84A0h	C500h	D140h	EA00h
11	84B0h	C580h	D160h	EB00h
12	84C0h	C600h	D180h	EC00h
13	84D0h	C680h	D1A0h	ED00h
14	84E0h	C700h	D1C0h	EE00h
15	84F0h	C780h	D1E0h	EF00h
16	8500h	C800h	D200h	F000h
17	8510h	C880h	D220h	F100h
18	8520h	C900h	D240h	F200h
19	8530h	C980h	D260h	F300h
20	8540h	CA00h	D280h	F400h
21	8550h	CA80h	D2A0h	F500h
22	8560h	CB00h	D2C0h	F600h
23	8570h	CB80h	D2E0h	F700h
24	8580h	CC00h	D300h	F800h
25	8590h	CC80h	D320h	F900h
26	85A0h	CD00h	D340h	FA00h
27	85B0h	CD80h	D360h	FB00h
28	85C0h	CE00h	D380h	FC00h
29	85D0h	CE80h	D3A0h	FD00h
30	85E0h	CF00h	D3C0h	FE00h
31	85F0h	CF80h	D3E0h	FF00h

Table 18 shows the Channel Register Block base addresses for each of the 32 channels. Note that each channel has four blocks of registers at four separate base addresses. Relative to the base addresses, individual channel registers can be accessed as shown in Table 19 below.

Table 19 Channel Registers

#	Register Name	Mode	Local or Global	Source (Processor Read)	Destination (Processor Write)	Address
1	Channel Status Register	RO	Global	CID 0	NA	Base 0 +00h
2	End of Packet Data Direction1	WO	Global	NA	CID n	Base 0 +08h
3	OC Result FIFO/ OC Result 0	RO	Global	CID n ^a	NA	Base 1 +00h
4	Data Result 0/1 FIFO/ Data Result 0/1 0	RO	Global	CID n	NA	Base 1 +08h
5	OC Result 1	RO	Global	CID n ^a	NA	Base 1 +10h
6	Data Result 0/1 1	RO	Global	CID n	NA	Base 1 +18h
7	OC Result 2	RO	Global	CID n ^a	NA	Base 1 +20h
8	Data Result 0/1 2	RO	Global	CID n	NA	Base 1 +28h
9	OC Result 3	RO	Global	CID n ^a	NA	Base 1 +30h
10	Data Result 0/1 3	RO	Global	CID n	NA	Base 1 +38h
11	OC Result 4	RO	Global	CID n ^a	NA	Base 1 +40h
12	Data Result 0/1 4	RO	Global	CID n	NA	Base 1 +48h
13	OC Result 5	RO	Global	CID n ^a	NA	Base 1 +50h
14	Data Result 0/1 5	RO	Global	CID n	NA	Base 1 +58h
15	OC Result 6	RO	Global	CID n ^a	NA	Base 1 +60h
16	Data Result 0/1 6	RO	Global	CID n	NA	Base 1 +68h
17	OC Result 7	RO	Global	CID n ^a	NA	Base 1 +70h
18	Data Result 0/1 7	RO	Global	CID n	NA	Base 1 +78h
19	Previous OC Result	RO	Global	CID n ^a	NA	Base 2 +00h
20	Data Result 2/3	RO	Global	CID n	NA	Base 2 +08h
21	Data Result 4/5	RO	Global	CID n	NA	Base 2 +10h
22	Data Result 6	RO	Global	CID n	NA	Base 2 +18h

Table 19 Channel Registers

#	Register Name	Mode	Local or Global	Source (Processor Read)	Destination (Processor Write)	Address
23	Packet Buffer Input	WO	Global	NA	CID 0~7	Base 3 +00h +08h ... +F0h
24	End of Packet Data Direction0	WO	Global	NA	CID 0~7	Base 3 +F8h

Note:

- a. CID n or CID 0 if no valid result is available.

4.2.2 Register Description

This section provides a description of the registers in the PM2329. This section also specifies the access mode: Local vs.Global, and the access type: Read Only, Write Only, or Read/Write, for each register.

In each Register bit description table, the reset values shown in the “Value after Reset” column applies to both soft and hard reset conditions.

For local addresses shown, substitute CID numbers 0, 1, 2,... or 7 for “n” depending on the PM2329 to be accessed.

4.2.2.1 Local Configuration Register (LCR; n000h)

Access Mode: Read/Write, Local

Bit Range	Size	Name	Value after Reset
63	1	BIST Result	0
62	1	BIST Enable/Status	0
61:59	3	(Reserved)	Undefined
58	1	SCH Timing Mode Select	0
57:53	5	(Reserved)	Undefined
52:32	21	(Reserved)	00 0000h
31:24	8	Device Revision Number	04h
23:18	6	(Reserved)	Undefined
17:16	2	PLL Multiplier	PLL Multiplier
15:8	8	Cascade Mask	Cascade Mask
7	1	(Reserved)	Undefined
6	1	ZBT Enable	ZBT Mode
5	1	System I/F Bus Width	Sys I/F Bus Width
4:2	3	CID Number	CID#
1:0	2	(Reserved)	Undefined

This register returns bits that indicate the hardware configuration of the PM2329. All bits, except bit 62, are read-only.

- **BIST Result**

This bit indicates the result of the last BIST run. It should be read after the BIST operation has been enabled (BIST Enable/Status LCR[62] is written with 1) and the BIST execution is complete (indicated by BIST Enable/Status LCR[62] = 0). When sampled, if this bit LCR[63] is 0, then the BIST failed; if this bit is 1, then the BIST passed.

This bit is set to 0 when BIST Enable/Status LCR[62] is written with 1.

This bit is 0 after reset.

- **BIST Enable/Status**

This bit controls the operation of the on-chip BIST facility for internal RAM blocks. Writing a 1 to this bit activates the BIST operation inside the device. This bit will stay 1 as long as the BIST operation is in progress, and it will become 0 when the BIST execution is complete.

After the BIST sequence is completed successfully, the following on-chip resources are initialized:

Rule Memory:NO MATCH

OCDs: INVALID OC

Since all rules are loaded with NO MATCH and all OCDs are loaded with INVALID OCs, control software need not initialize unused rules and OCDs.

This bit is 0 after reset.

- **SCH Timing Mode Select**

This bit controls the clock frequency at which the SCHSTB and SCHNUM operate. This bit can be configured as follows:

- 0: SCHSTB and SCHNUM[] operate at the same frequency as SCLK.
- 1: SCHSTB and SCHNUM[] operate at half the frequency of SCLK.

- **Device Revision Number**

This field indicates the revision number of the device. For the PM2329-A1, this field is hardwired to 04h.

- **PLLA Multiplier**

These two bits indicate the value of the PLLA multiplier sensed at reset. For a description of this field's value and the corresponding PLLA multiplier, see the PLLACTRL[1:0] description in Chapter 2, Section 2.2 Pin Description Table.

- **Cascade Mask**

These bits specify which of the chips in the cascade were detected via handshake using COCDOUT lines during reset. The external processor can read this field to determine the number of PM2329 devices in the cascade.

- **ZBT Enable**

This bit indicates the value of ZBT_MODE signal sampled at reset. If this bit is '0', a low level was sensed at reset and the PM2329 system interface is configured for SyncBurst mode. If this bit is '1', a high level was sensed at reset and the PM2329 system interface is configured for ZBT operation.

- **System Interface Bus Width**

This bit indicates the value of SD_WIDTH signal sampled at reset. If this bit is '0', a low level was sensed at reset and the PM2329 is configured for 32-bit mode. If this bit is '1', a high level was sensed at reset and the PM2329 is configured for 64-bit mode.

- **CID Number**

These 3 bits return the PM2329 ID number for that particular PM2329 as set up by the CID[2:0] inputs sensed at reset.

4.2.2.2 Rule Indirect Command Register (RICR; n008h)

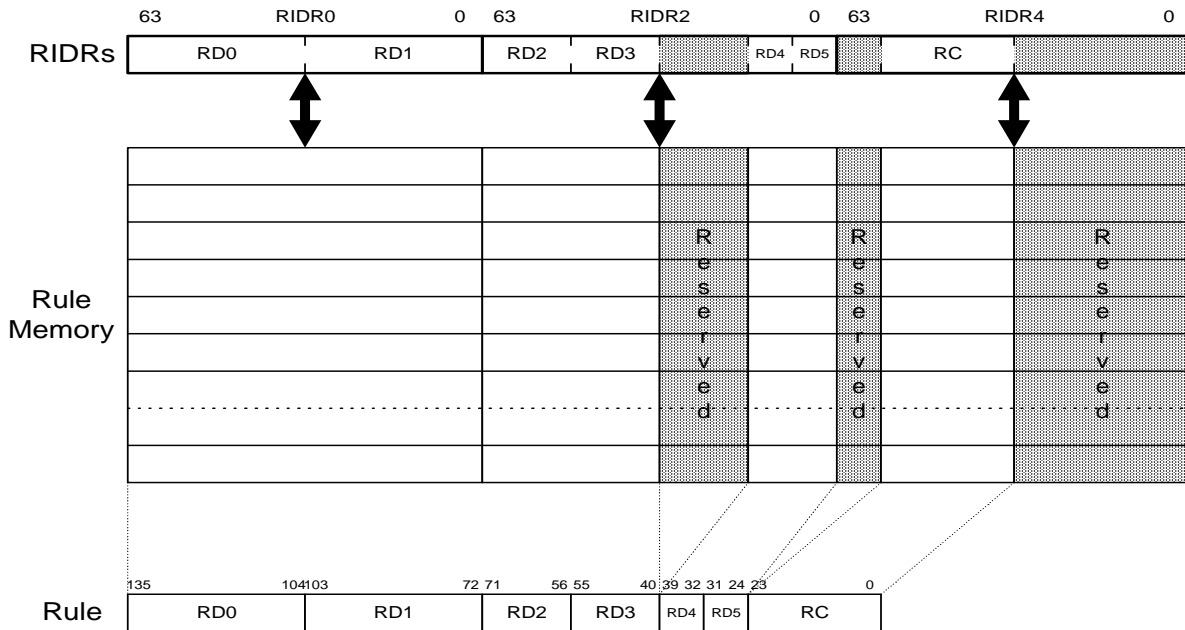
Access Mode: Read/Write, Local

Bit Range	Size	Name	Value after Reset
63:13	51	(Reserved)	Undefined
12:8	5	Read/Write Data Enable	00h
7:6	2	(Reserved)	Undefined
5	1	Auto-Increment	0
4:2	3	Trigger	000
1	1	(Reserved)	Undefined
0	1	Register Set Ready (RO)	1

This register controls the indirect access mechanism to the Policy Database Rule Memory cells in the PM2329.

The values in Rule Memory cells are accessed using the RIDR registers (explained in Section 4.2.2.4). The diagram below shows the Policy Database Rule memory cell mapping to these (RIDR0~4) registers.

Figure 26 Access to Rule Memory Cells via RIDR0-4 Registers



- **Read/Write Data Enable**

These bits enable or disable reading or writing to each of the 32-bit words comprising the 136-bit rule memory. When a bit is '1', read or write of the corresponding word is enabled. During write, the value in the Rule Indirect Data Register will be written to the addressed rule memory location. When a bit is '0', the corresponding word is disabled. The higher bit (bit 12) is for Rule Indirect Data 0, and the lower bit

(bit 8) is for Rule Indirect Data 4. For example, programming a value of 01100 will enable writes to words 1 and 2 only. During read, only the enabled words are read.

For a detailed description of the Rule Memory read and write operation, refer to the Rule Indirect Data Register set description.

- **Auto-Increment**

If this bit is '1', then after the read or write operation to the Rule Memory, the Rule Indirect Address Register is automatically incremented to point to the next cell (post-increment). This bit will normally be written with '1' when a block of data needs to be transferred to or from the Rule Memory.

If this bit is '0', the Rule Indirect Address Register is unchanged after a read or write operation to the Rule Memory.

- **Trigger**

These bits specify which of the five 32-bit data words will serve as the Trigger for the PM2329 to perform the internal read or write operation to the Rule Memory. After the Trigger word is written to or read from, the device performs Read or Write operation on the rule memory. These bits are programmed as follows:

000	Rule Indirect Data Register 0
001	Rule Indirect Data Register 1
010	Rule Indirect Data Register 2
011	Rule Indirect Data Register 3
100	Rule Indirect Data Register 4
101	(Reserved)
110	(Reserved)
111	(Reserved)

- **Register Set Ready**

This bit when '1' indicates that the Rule Indirect registers are currently not in use. It is cleared when a read or write command is issued and remains '0' as long as the device is busy performing this operation. It is set when the device has internally executed the read or write operation. For correct operation, the processor must check that this bit is set before it writes to the RICR, RIAR or RIDR. When the RSR bit is 0, write cycles to any of these (RICR, RIAR or RIDR) registers are ignored. This is a read-only bit and is set to '1' after reset.

4.2.2.3 Rule Indirect Address Register (RIAR; n010h)

Access Mode: Read/Write, Local

Bit Range	Size	Name	Value after Reset
63:14	50	(Reserved)	Undefined
13:0	14	Cell Number	0000h

This register is used to specify the address of the Rule Memory cell to be accessed.

- **Cell Number**

This is the address of the Rule to be accessed. This is an absolute cell number within the PM2329. If the auto-increment bit in Rule Indirect Command Register is set to '1', then after a read or write operation is performed to the Rule memory, the Rule Indirect Address Register is automatically incremented to point to the next cell (post-increment).

For a detailed description of the Rule Memory read and write operation, refer to the Rule Indirect Data Register set description.

4.2.2.4 Rule Indirect Data Register Set (RIDR0; n018h) (RIDR2; n020h) (RIDR4; n028h)

Access Mode: Read/Write, Local

This is a set of five 32-bit registers or three 64-bit registers. They are used by the processor to read or write data from the Rule Memory cells of the PM2329. The registers are organized as follows.

Register 0

Bit Range	Size	Name	Value after Reset
63:32	32	Rule Bits (135:104) or Rule Data Field #0 (SIP)	0000 0000h
31:0	32	Rule Bits (103:72) or Rule Data Field #1 (DIP)	0000 0000h

Register 2

Bit Range	Size	Name	Value after Reset
63:48	16	Rule Bits (71:56) or Rule Data Field #2 (SP)	0000h
47:32	16	Rule Bits (55:40) or Rule Data Field #3 (DP)	0000h
31:16	16	(Reserved)	Undefined
15:8	8	Rule Bits (39:32) or Rule Data Field #4 (Protocol)	00h
7:0	8	Rule Bits (31:24) or Rule Data Field #5 (Flags (7:4) & Mask (3:0))	00h

Register 4

Bit Range	Size	Name	Value after Reset
63:56	8	(Reserved)	Undefined
55:32	24	Rule Bits (23:0) or Rule Control Field	00 0000h
31:0	32	(Reserved)	Undefined

Note: The field names (e.g., SIP, DIP, SP and DP) shown in the table above are for user convenience only and generally correspond to header related fields described in the Field Extraction Engine description. They are used for other functions depending on the usage of the rule (e.g., pattern searches use these fields to store patterns or mask values).

Rule memory write or read operation can be performed either in random access mode or in sequential access mode. In general, the processor should set up the Rule Indirect Command Register to specify the auto-increment mode, etc. It should then set up the address of the location or the starting address of the block of locations to be accessed. It can then access the addressed location or access sequential location depending on the programming of the auto-increment bit.

Note that read or write operations are performed internally when the trigger word is accessed as explained below. However, if the RIAR, RICR or RIDR have not been written to since the last rule memory access, an internal read to the rule memory is not performed since the current RIDR content correctly reflects the address rule memory content.

During random reads, the processor must perform a dummy read (and discard the value read back) to trigger the actual read internally, and then perform a second read to get the real data. For sequential reads, only one (the first) dummy read is required for each block of data read sequentially. In either case, the first dummy read needs to be performed to the trigger word only.

During writes, processor can simply write to the Indirect Data Register and the device will execute the internal write cycle when the trigger word is written. There are no dummy writes to be performed, however, the processor must ensure that all the enabled data words are written to first, before the trigger word is written.

As long as the processor manages rule memory Write and Read operations using the RSR handshake, the processor can perform these operations at any time even when OC sequencing is in progress--the read or write operation will take place arbitrated by PM2329 control logic. The software must ensure coherency of OC operation vs. its access operation. Also, it must manage multiple rule word updates using the RSR bit.

Note that all the rules in the rule memory are initialized to NO MATCH after the BIST sequence is completed successfully; control software need not initialize unused rules.

Rule Updates when PM2329 is processing an OC

When an OC is in progress, rule update for all columns that are in the partition of the current OC is deferred until the last result of the OC has been transferred out of the processing engine to the result FIFO. This update deferral includes times when the processing engine has stalled if the Result FIFO is full, in other words, the update will not occur during processing engine stalls.

Note that these columns can contain rules that belong to the current partition (as defined by the RSA and REA fields) and those that do not belong to the current partition--none of the rules in these columns will be updated until the condition stated above occurs.

Columns that are not participating in the current OC can be updated at all times. However, note if a rule update is started that results in update deferral (since it belongs to a column of the current partition), the RSR handshake control will prevent loading of a new rule to be updated that could otherwise be updated (one that does not belong to a column of the current partition).

Read and write sequences are explained in greater detail below.

For random read operation, the processor must perform the following steps.

1. Check the RSR bit is set to '1' to ensure the previous memory transaction is complete.
2. Set up RICR with the Data Enable and Trigger fields, and the auto-increment bit reset to '0'.
3. Set up the RIAR with the address of the location to be read.
4. Read the trigger word and discard the value read back.
5. Wait for RSR bit to be set to '1'
6. Read the content of the addressed location by reading the RIDR (the RIDR set as trigger word should be read last).

To read a set of random locations, steps 1 through 6 outlined above can be repeated. However, a set of random reads can be optimized (only one dummy read for a set of random reads), if (step 5a) the address of the next random location to be read is loaded in the RIAR after step 5, before reading the previous locations content in step 6. In this case, the processor can simply repeat steps 5, 5a, and 6 until all the random location have been accessed.

For random write operation, the processor must perform the following steps.

1. Check the RSR bit is set to '1' to ensure the previous memory transaction is complete.
2. Set up RICR with the Data Enable and Trigger fields, and the auto-increment bit reset to '0'.
3. Set up the RIAR with the address of the location to be written.
4. Write the RIDR, the trigger word should be written last.

For sequential read operation, the processor must perform the following steps.

1. Check the RSR bit is set to '1' to ensure the previous memory transaction is complete.
2. Set up RICR with the Data Enable and Trigger fields, and the auto-increment bit set to '1'.
3. Set up the RIAR with the start address of the locations to be read.
4. Read the trigger word and discard the value read back.
5. Wait for RSR bit to be set to '1'
6. Read the content of the addressed location by reading the RIDR (the trigger word should be read last).

Repeat steps 5 and 6 to read the rest of the memory block.

For sequential write operation, the processor must perform the following steps.

1. Check the RSR bit is set to '1' to ensure the previous memory transaction is complete.
2. Set up RICR with the Data Enable and Trigger fields, and the auto-increment bit set to '1'.
3. Set up the RIAR with the start address of the locations to be written.
4. Write the RIDR, the trigger word should be written last.
5. Wait for RSR bit to be set to '1'

Repeat steps 5 and 6 to write the rest of the memory block.

4.2.2.5 OC Descriptors (OCD; n400h, n408h...n7F0h, n7F8h)

Access Mode: Read/Write, Local

OC Descriptor

Each OCD occupies two consecutive 64-bit locations: Upper OCD [n] and Lower OCD [n], where n is the OC Descriptor Index. OCDI. Since each full OC Descriptor is 96 bits wide spanning two 64-bit register locations or three 32-bit register locations, the user must exercise caution when updating the OC Descriptor while an OC that uses this OC Descriptor is executing. OC Descriptors can be updated without any side effects when:

- 1) no OC execution is in progress, or
- 2) if the OC Descriptor to be updated will not be used by the current OC that is executing.

Upper OC Descriptor

Bit Range	Size	Name	Value after Reset
63	1	OC Descriptor Valid	undefined
62:61	2	OC Type	undefined
60	1	Enable Multi-hit	undefined
59:54	6	Row Start Number	undefined
53:48	6	Row End Number	undefined
47:32	16	Column Select Enable	undefined
31:29	3	(Reserved)	undefined
28:16	13	Pattern Search Start Offset	undefined
15	1	Pattern Search Direction	undefined
14:11	2	(Reserved)	undefined
12:0	13	Pattern Search Count	undefined

Lower OC Descriptor

Bit Range	Size	Name	Value after Reset
63:32	32	(Reserved)	undefined
31:29	3	(Reserved)	undefined
28:16	13	E-Word Segment Base Offset	undefined
15:13	3	(Reserved)	undefined
12:0	13	Partition Start Offset	undefined

The PM2329 contains a total of 64 OC Descriptors (OCD), which are used to describe the OC partition to be executed. The OC to be executed is specified by a 6-bit OCD Index programmed into the OC Conductor or into the E-RAM Control Word.

Note that all the OCDs are initialized to INVALID OC after the BIST sequence is completed successfully; control software need not initialize unused OCDs.

- **OC Descriptor Valid**

The OCD is executed only if this bit is set.

If an OCD Index points to an OCD with this bit cleared, it will terminate with OC Done without match. No error condition is flagged in this case.

- **OC Type**

The field specifies the type of OC to be executed, valid values are:

- 00 Header OC
- 01 Attribute OC
- 10 Pattern Search Short (up to 12 bytes)
- 11 Pattern Search Long (up to 192 bytes)

Note that the Header OC and Pattern Search OCs (Short or Long) use packet data information and their execution starts when the EOP is detected during packet transfer. In case of Attribute OC, if the attribute data is supplied during packet transfer, the same technique (EOP) can be used to start OC execution, however, if the attribute data is in the Packet Information register, the processor must execute a dummy write cycle to the EOP port in order to start the OC execution.

- **Enable Multi-hit**

If this bit is '0', the OC execution returns the highest priority result only (single hit OC). If the bit is '1', then all the match results are returned in a prioritized order (Multi-hit OC). The number of results generated by a Multi-hit OC are not predictable.

- **Row Start Number**

This field specifies the Rule Memory Row for the start of the partition which will participate in the OC.

- **Row End Number**

This field specifies the last Rule Memory Row which will participate in the OC. This field along with the Row Start field gives the range of rows for the partition.

Note that Row Start Number and Row End Number specification wraps around. In other words, the following applies:

1. If Row Start Number and Row End Number are equal, the partition spans one row.
2. If Row Start Number is less than Row End Number, the partition spans from Row Start Number from Row End Number
3. If Row Start Number is greater than Row End Number, the partition spans from Row Start Number to Row 63 and then from Row 0 up to Row End Number.

- **Column Select Enable**

This field specifies which columns are included in the OC partition. This field assign one bit per column, the highest bit specifies column 15 and the lowest bit specifies column 0. A column participates in the OC only if the corresponding bit is '1'.

- **Pattern Search Start Offset**

This field is only valid for a Pattern Search (Short or Long) OC.

This field specifies the starting offset in the packet data which will be scanned to search for a set of preloaded patterns. Bytes are taken starting from this offset (and incrementing or decrementing based on Search Direction) to form the Data Source for the OC.

The first byte of the packet is at offset 0.

- **Pattern Search Direction**

This field is only valid for a Pattern Search (Short or Long) OC.

If this bit is '0', the pattern search is in the forward direction (incrementing offsets). If this bit is '1', the pattern search is in the reverse direction (decrementing offsets).

- **Pattern Search Count**

This field is only valid for a Pattern Search (Short or Long) OC.

This field specifies the window (in bytes) in the packet data which will be scanned to search for a set of preloaded patterns. Note that this window specifies the starting offset (0) of the string.

- **E-Word Segment Base Offset**

This field defines the starting location of the E-Word segment corresponding to this OC. An E-RAM segment associated with an OC can start on a 16 location boundary, so the value in this field is multiplied by 16 for calculating the location of the E-Word to access.

- **Partition Start Offset**

For any OC partition spanning multiple PM2329 devices, the relative cell offset of the first (highest priority) cell of the partition is defined by this field. Since OC partitions can be defined only on 16 cell boundaries, this field is multiplied by 16 to get the actual absolute partition start address.

4.2.2.6 E-RAM Indirect Data Register Set (EIDR0; 8200h) (EIDR2; 8208h) (EIDR2; 8208h) (EIDR4; 8210h) (EIDR6; 8218h)

Access Mode: Read/Write, Global

The processor views these registers as a set of 32-bit registers which are used to write or read data from locations in the E-RAM devices attached to the PM2329.

Recall that each PM2329 device has a 32-bit E-RAM control bus and a 32-bit E-RAM data bus. Further, the E-RAM control bus from all devices is connected to a shared memory device, and the E-RAM data bus from each device is connected to a separate memory device. To access all these memory devices over the E-RAM bus, there are essentially two E-RAM access registers implemented in the PM2329--one that can access the memory connected to the control bus and the other that can access the memory connected to the data bus.

These two E-RAM access registers are mapped to the C-Word bus and D-Word bus of the PM2329 cascade depending on the CID# of the PM2329 as shown in below.

PM2329 0	C-Word Bus
PM2329 0	D-Word 0 Bus
PM2329 1	D-Word 1 Bus
PM2329 2	D-Word 2 Bus
PM2329 3	D-Word 3 Bus
PM2329 4	D-Word 4 Bus
PM2329 5	D-Word 5 Bus
PM2329 6	D-Word 6 Bus
PM2329 7	No D-Word Bus supported

Note that both the C-Word and the D-Word 0 buses are accessed using CID#0 and no D-Word is supported on CID#7. In other words, both the E-RAM access registers in the PM2329 with CID#0 are used, only one of the two E-RAM access registers in CIDs #1 through #6 are used and neither of the E-RAM access registers in CID#7 are used.

This register set contains as many 32-bit registers as the physical width of the connected E-RAM. Based on the Read/Write Data Enable in the E-RAM Indirect Address Register, the cascaded PM2329 devices can read from or write to multiple 32-bit E-RAM words in a single cycle.

The E-RAM Configuration Register defines how the E-RAM is configured. Based on the configuration, the E-RAM contains a C-Word (mandatory) and a set of (optional) D-Words.

Since these registers are distributed across the cascaded PM2329 devices, their validity depends on the number of PM2329 devices in the cascade. Shown below is the association of these registers with the CID# of the PM2329 devices:

EIDR0 High CID#0

EIDR0 Low CID#0

EIDR2 High CID#1

EIDR2 Low CID#2

EIDR4 High CID#3

EIDR4 Low CID#4

EIDR6 High CID#5

EIDR6 Low CID#6

The structure of the E-RAM Indirect Data Register Set is as follows:

Control Word/D-Word 0

Bit Range	Size	Name	Value after Reset
63:32	32	Control Word Bus	Undefined
31:0	32	Data Word 0 Bus	Undefined

D-Word 1/D-Word 2

Bit Range	Size	Name	Value after Reset
63:32	32	Data Word 1 Bus	Undefined
31:0	32	Data Word 2 Bus	Undefined

D-Word 3/D-Word 4

Bit Range	Size	Name	Value after Reset
63:32	32	Data Word 3 Bus	Undefined
31:0	32	Data Word 4 Bus	Undefined

D-Word 5/D-Word 6

Bit Range	Size	Name	Value after Reset
63:32	32	Data Word 5 Bus	Undefined
31:0	32	Data Word 6 Bus	Undefined

For a description of the fields and associated values in the E-RAM, see Section 4.3.2

For a detailed description of the E-RAM access mechanism, refer to the Rule Memory Indirect Data Register section. Processor access mechanism to E-RAM memory through the E-RAM Indirect Register set (Data, Command and Address) is similar to the Rule Memory access mechanism with the following differences.

Rule memory for each PM2329 device is local to it and contained on-chip, i.e., the processor sees as many rule memory blocks as the number of PM2329 devices in the cascade with their own associated address pointers and indirect data registers for access. While various PM2329 devices have separate E-RAMs attached to them, they present a unified view (a common block of memory) to the processor accessed via a common address pointer and a distributed yet common set of indirect data register set.

Additionally, the E-RAM access mechanism supports the Auto Increment Mode Select bit as explained in the E-RAM Indirect Command Register section. It also has a depth Level field (part of the E-RAM address) as explained in the E-RAM Indirect Address Register section.

As long as the processor manages E-RAM Write and Read operations using the RSR handshake, the processor can perform these operations at any time even when OC sequencing is in progress--the read or write operation will take place arbitrated by PM2329 control logic. The software must ensure coherency of OC operation vs. the access operation. Also, it must manage multiple E-RAM updates using the RSR bit.

Unlike RIDR (Rule Memory access mechanism) read operation, reading the EIDR set as the trigger word will always cause an E-RAM read operation regardless of whether the EICR, EIAR or EIDR were written to or not since the last ERAM read operation.

4.2.2.7 E-RAM Indirect Command Register (EICR; 8220h)

Access Mode: Read/Write, Global

Bit Range	Size	Name	Value after Reset
63:16	48	(Reserved)	Undefined
15:8	8	Read/Write Data Enable	00h
7	1	(Reserved)	Undefined
6	1	Auto Increment Mode Select	0
5	1	Auto Increment Enable	0
4:2	3	Trigger	000
1	1	(Reserved)	Undefined
0	1	Register Set Ready (RO)	1

- **Read/Write Data Enable**

This field enables or disables reading or writing to each of the 32-bit physical E-RAMs, which may be up to (32 x8) 256 bits wide. When a bit is '1', it enables the corresponding E-RAM access. When a bit is '0', the corresponding E-RAM access is disabled. The higher bit (bit 15) is for E-RAM Indirect Data Register 0, and the lowest bit (bit 8) is for E-RAM Indirect Data Register 7. For example, programming a value of '01101011' will enable writes to E-RAMs 1, 2, 4, 6 and 7 only. See Figure 2.8 for ERAM # reference description.

Read/Write Data Enable bits in the E-RAM Indirect Command Register have a specific assignment for each device as determined by its CID #, it controls the C-Word and/or D-Word connected to that device regardless of the type of D-Word programming. This field is independent of D-Word type or depth, in other words if a Read/Write Data Enable bit is disabled, all corresponding read/write access associated with that E-RAM will be disabled. This field is applicable for E-RAM indirect operation only, OC Sequencing controlled accesses are not qualified by this field.

- **Auto Increment Mode Select**

If this bit is '0', increment operation is performed on the E-RAM address value only. If this bit is '1', the Level value is also incremented in addition to the E-RAM address value. See Level field description for further explanation.

- **Auto Increment Enable**

This bit is written with '1' when a block of data needs to be transferred to or from the E-RAM. After the read or write operation, the E-RAM Indirect Address Register is automatically incremented to point to the next memory location (post increment).

If this bit is '0', the E-RAM Indirect Address Register is unchanged after a read or write operation to the E-RAM Memory.

- **Trigger**

This field specifies which of the (up to) eight 32-bit data words will serve as the Trigger for the PM2329 to perform the read or write operation to the E-RAM. After the Trigger word is written to or read from, the device performs Read or Write operation on the E-RAM. These bits are programmed as follows:

000	E-RAM Indirect Data Register 0 or C-Word Bus
001	E-RAM Indirect Data Register 1 or D-Word 0 Bus
010	E-RAM Indirect Data Register 2 or D-Word 1 Bus
011	E-RAM Indirect Data Register 3 or D-Word 2 Bus
100	E-RAM Indirect Data Register 4 or D-Word 3 Bus
101	E-RAM Indirect Data Register 5 or D-Word 4 Bus
110	E-RAM Indirect Data Register 6 or D-Word 5 Bus
111	E-RAM Indirect Data Register 7 or D-Word 6 Bus

- **Register Set Ready**

This bit when '1' indicates that the E-RAM Indirect registers are currently not in use. It gets cleared when a read or write command is issued and remains '0' as long as the device is busy performing this operation. It is set again as soon as the device has executed the read or write operation to the E-RAM. For correct operation, the processor must check that this bit is set before it writes to the EICR, EIAR or EIDR. When the RSR bit is 0, write cycles to any of these (RICR, RIAR or RIDR) registers are ignored.

This is a read-only bit and is set to '1' after reset.

4.2.2.8 E-RAM Indirect Address Register (EIAR; 8228h)

Access Mode: Read/Write, Global

Bit Range	Size	Name	Value after Reset
63:24	40	(Reserved)	Undefined
23:22	2	Level	00
21:17	5	(Reserved)	Undefined
16:0	17	E-Word Address	0 0000h

This register is used to specify the indirect read/write address of the E-RAM location to be accessed.

- **Level**

In case the E-Word is larger than the physical width of the E-RAM, in one indirect read or write operation of the E-RAM, the number of 32-bit words read or written will be determined by the E-RAM Width. The Level field thus provides a means to access the other parts of the E-Word.

The Level is programmed as

- 00 Level 0
- 01 Level 1
- 10 Level 2
- 11 Level 3

Note that the results returned in the Data FIFO correspond to level 0 only. Exception to this rule is in the case of a single E-RAM device where both the ECD and EDD buses are connected in parallel to the same memory device. In this case, D-Word 0 (at Level 1) will be returned in the Data FIFO. Higher levels in all cases must be accessed by the processor using E-RAM indirect addressing mechanism. See E-RAM Configuration Register for more details.

When the Auto-increment mode select bit enables increment of the level field, the PM2329 will increment the level field depending on the depth of E-Word in the E-RAM. If depth is 1, Level field is not incremented. If depth is 2 Level field is incremented from 0 to 1 and then rolls over. If depth is 4, Level field is incremented from 0 to 3 and then rolls over. If Level is programmed to an invalid value, e.g., Level 2 when depth is 2, it will count up to 3 and then roll over.

- **E-Word Address**

The location of the E-Word in the E-RAM to be written or read.

4.2.2.9 E-RAM Configuration Register (ECR; 8230h)

Access Mode: Read/Write, Global

Bit Range	Size	Name	Value after Reset
63:32	32	(Reserved)	Undefined
31	1	E-RAM Enable	0
30:28	3	D-Word #0 Definition	000
27:25	3	D-Word #1 Definition	000
24:22	3	D-Word #2 Definition	000
21:19	3	D-Word #3 Definition	000
18:16	3	D-Word #4 Definition	000
15:13	3	D-Word #5 Definition	000
12:10	3	D-Word #6 Definition	000
9:3	7	(Reserved)	Undefined
2:0	3	E-RAM Width	00h

The bits in this register are used to configure the logical E-Word. This register must be programmed before any E-RAM accesses are performed. This register should not be updated when E-RAM operations are in progress, normally this register will be written during initialization only. In cascade mode, global write will configure all the devices together.

- **E-RAM Enable**

When this bit is reset to 0, all E-RAM operations are disabled. No C-Word will be supported in this case.

When this bit is set to 1, E-RAM operations are enabled. C-Word is always assumed to be present in this case. D-Words are defined using the D-Word Definition fields as explained below.

- **D-Word Definition**

If E-RAM is present, up to 7 D-Words (depending on the physical width of the E-RAM) can be configured. These 21 bits define the D-Word type for these 7 D-Words. the PM2329 also interprets this register to obtain the logical E-Word Width. Three bits define each of the D-Words as follows:

000	D-Word Absent
001	Packet Count
010	Byte Count
011	Timestamp/State
100	User Defined
101	(Reserved)
110	(Reserved)
111	(Reserved)

The user can force gaps in the D-Words by programming them to be user defined.

The logical E-Word Width is determined by scanning the D-Word Definition bits from DW6 down to DW0, until it encounters the first D-Word field that is non-zero, i.e., D-Word that is not absent. The E-Word width is C-Word plus the number of D-Words present.

• **E-RAM Width**

Indicates the physical width of attached E-RAM. The PM2329 devices in cascade assume that the first 32 bits (i.e., C-Word) are connected to PM2329 #0, next 32 bits are also on CID #0, next 32 bits on CID #1 and so on. This sequence must be maintained and gaps in the physical connection are not allowed. The bits are written as follows:

001	32-bit	010	64-bit
011	96-bit	100	128-bit
101	160-bit	110	192-bit
111	224-bit	000	256-bit

Note that the PM2329 cascade allows a larger E-Word to reside in an E-RAM whose physical width is smaller than the logical E-Word. This will happen whenever the D-Word Definition field indicates a larger E-Word than what the E-RAM Width field specifies. Consequently, the number of locations accessed to access the entire E-Word (E-Word Depth) varies. Given the E-Word width and E-RAM width, the PM2329 automatically enforces a depth of 1, 2 or 4 as shown in the table below.

Table 20 E-Word Depth Chart

E-Word Depth for E-Word Logical Width & E-RAM Physical Width Combinations									
Physical E-RAM Width (bits)	Logical E-Word Width (bits), C-Word plus D-Word								Number of devices required
	32	64	96	128	160	192	224	256	
32	1	2	4	4	NA	NA	NA	NA	1
64	1	1	2	4	4	NA	NA	NA	1
96	1	1	1	2	2	4	4	4	2
128	1	1	1	1	2	2	2	4	3
160	1	1	1	1	1	2	2	2	4
192	1	1	1	1	1	1	2	2	5
224	1	1	1	1	1	1	1	2	6
256	1	1	1	1	1	1	1	1	7

EMA[18:17] lines carry the E-RAM depth field. For most efficient E-RAM utilization, EMA[18] and EMA[17] signals should be connected to the E-RAM devices depending on the depth of the E-Word as shown in the table below.

Table 21 EMA[18:17] Usage

Depth	EMA[18]	EMA[17]
1	NC	NC
2	NC	Connected
4	Connected	Connected

4.2.2.10 Interrupt Enable Register (IER; 8238h)

Access Mode: Read/Write, Global

Bit Range	Size	Name	Value after Reset
63:8	56	(Reserved)	Undefined
7	1	Enable Result FIFO Full Interrupt (RFOF)	0
6	1	(Reserved)	Undefined
5	1	Enable OC Sequence Halted Interrupt (OCSH)	0
4	1	Enable Idle Interrupt (IDLE)	0
3	1	Enable Packet Buffer Available Interrupt (PBA)	0
2	1	Enable Result Available Interrupt (RAV)	0
1	1	Enable OC Sequence Terminated Interrupt (OCST)	0
0	1	Interrupt Enable	0

This register specifies the mask bits for the various conditions that can cause the PM2329 to assert an external interrupt. If a bit is set, an occurrence of the corresponding condition causes the interrupt signal to be asserted. The OCST status bit is always cleared on reading the STSR. The interrupt signal will continue to be asserted until either the interrupting condition is removed, or the corresponding enable bit in this register is reset.

The bits in this register are defined below.

- **Enable Result FIFO Full Interrupt (RFOF)**

If this bit is set, an interrupt is generated when a result FIFO is full.

- **Enable OC Sequence Halted Interrupt (OCSH)**

If this bit is set, an interrupt is generated when the OC Sequence is halted due to a “Break” or “Wait” condition (see Status register for definitions of these conditions). OC Sequence Halted (Status register bit 5) will be set when this interrupt is asserted.

- **Enable Idle Interrupt (IDLE)**

If this bit is set, an interrupt is generated when all Packet Buffers are empty and all result FIFOs are empty.

- **Enable Packet Buffer Available Interrupt (PBA)**

If this bit is set, an interrupt is generated when a Packet Buffer is available for receiving new packet data from the external processor.

- **Enable Result Available Interrupt (RAV)**

If this bit is set, an interrupt is generated when there is at least one result in the Result FIFO.

- **Enable OC Sequence Terminated Interrupt (OCST)**

If this bit is set, an interrupt is generated when an entire sequence of OCs has terminated, i.e., processing for the current packet is over. Upon completion of an OC sequence, the current packet is discarded by the PM2329.

- **Interrupt Enable**

SINT* is never asserted if this bit is set to '0'. When '1', an interrupt is generated when an interrupting condition exists and the interrupt enable bit for that condition is set to '1' as shown below.

4.2.2.11 Status Register (STSR; 8240h)

Access Mode: Read Only, Global

Bit Range	Size	Name	Value after Reset
63:8	56	(Reserved)	Undefined
7	1	Result FIFO Full	0
6	1	OC Sequence Halt Condition	Undefined
5	1	OC Sequence Halted	0
4	1	Idle	1
3	1	Packet Buffer Available	1
2	1	Result Available	0
1	1	OC Sequence Terminated	0
0	1	Status Valid	1

The Status Register provides common information regarding the state of the single or cascaded PM2329 devices. This is a Global Read register, only the Primary PM2329 responds to the read request. The bits defined hold true for all devices in a cascade of PM2329 devices.

In single channel mode when interrupts are not used, it is possible to process the status using either this register or the appropriate channel status register.

In multi-channel mode, this register provides summary status of all the channels. Channel specific information can be retrieved by reading the channel status registers.

When interrupts are used, this register should be read to clear the OC Sequence Terminated bit.

Bit definitions for this register are as follows:

- **Result FIFO Full**

This bit is set when any one of the result FIFOs becomes full. If this condition is not serviced by the processor, the internal processing engine will eventually stall since any results that are generated can not be transferred to the Result FIFO.

The OC Processing Halt State bit indicates the reason for the halt condition.

- **OC Sequence Halt Condition**

If the OC Sequence Halted bit (bit 5) is 0, then this bit is a don't care.

If the OC Sequence Halted bit is 1, and this bit is 0, the PM2329 has reached a “**Break**” condition--it has completed execution of the previous OC, and that previous OC was not the last OC in the specified sequence.

If the OC Sequence Halted bit is 1, and this bit is 1, the PM2329 has reached a “**Wait**” condition--it has completed execution of the previous OC, and that previous OC was the last OC in the specified sequence.

- **OC Sequence Halted**

This bit is set when the PM2329 is halted awaiting a command from the network processor, and either:

1. OC Trace Enable bit is 0 and OC Sequence Mode bit is 1 (Processor controlled OC sequencing), or
2. OC Trace Enable bit is 1 and OC Sequence Mode bit is 0 (automated OC sequencing with Trace).

The OC Sequence Halt Condition bit (bit 6) indicates the reason for the halt condition.

If the processor writes to the AOCCR register, this bit is reset.

- **Idle**

This bit is set when all Packet Buffers are empty and all result FIFOs are empty. When this bit is set, the PM2329 is in Idle condition and the PM2329 operating modes can be reprogrammed without losing coherency.

This bit is cleared when any of the Packet Buffers is written to or at least one result is present in the result FIFO. If the PM2329 operating modes are reprogrammed when this bit is clear, data or results can be lost.

- **Packet Buffer Available**

This bit when set indicates that the PM2329 is ready to receive another packet into one or more of the Packet Buffers. When the PM2329 is operating in multi-channel mode, this bit when set indicates that at least one of the Packet Buffers is available. The Channel Packet Buffer Available bit in the Channel Status Registers can be used to determine channel(s) for which the Packet Buffer(s) are available. For further information about Packet Buffer Available status, see the Channel Status Register description.

This bit is cleared when all the individual Channel Packet Buffer Available bits are clear.

- **Result Available**

This bit is set when there is at least one result in the Result FIFO. When the PM2329 is operating in multi-channel mode, this bit when set indicates that at least one of the Result FIFOs have results available. The Channel Result Available bit in the Channel Status Registers can be used to determine channel(s) for which the Result(s) are available.

This bit is cleared when all the individual Channel Result Available bits are clear.

- **OC Sequence Terminated**

This bit is set when the packet processing on the current packet is complete. This bit is cleared when this register is read. This also causes the OC Sequence Interrupt, if enabled, to be deasserted.

- **Status Valid**

This bit is set to indicate the status registers has at least one status bit set. This bit is cleared when all other status bits in this register are clear.

4.2.2.12 Operation Control Register (OPCR; 8248h)

Access Mode: Read/Write, Global

Bit Range	Size	Name	Value after Reset
63	1	Soft Reset	See Note (a)
62	1	Hard Reset Status (RO)	See Note (a)
61:4	58	(Reserved)	000 0000 0000 0000h
3	1	Enable Multi-Channel Mode	1
2	1	Enable PI Field	1
1	1	Enable OCC Field	1
0	1	Direction Specifier	0

Notes:

- This bit is set to '1' when the reset signal RESET* is asserted, it is reset to '0' eight SCLK cycles after reset signal is deasserted.

This register controls the operation modes of the PM2329. This register should be written by the processor when it detects the PM2329 is idle (STSR[4]= 1) and prior to supplying the packet data to the device and it should not be updated by the processor until after the EOP for the packet has been transferred to the device. The bits in this register are as follows:

- **Soft Reset**

Writing a 1 to this bit causes the chip to perform a soft reset internally. On reading the register, the bit indicates 1 if a soft reset or hard reset is in progress. The bit automatically resets to 0 eight SCLK cycles when the internal reset operation is done. Soft reset forces all the registers to their Reset state and the Packet Buffer and Result FIFOs are cleared. Soft reset does not affect strap option pins, they are sampled on hard reset only.

- **Hard Reset Status**

This is a read-only bit. On reading the register, this bit is 1 if a hard reset is in progress. This occurs immediately after a hardware reset is applied by asserting the RESET* signal. After the RESET* pin is deasserted, the PM2329 will keep this bit set for eight SCLK cycles until it has completed its internal initializations. The bit automatically resets to 0 when the internal reset operation is complete. The external software should check for this bit '0' before issuing any other accesses to the PM2329.

- **Enable Multi Channel Mode**

Writing '0' to this bit causes the PM2329 to act as a single channel device. The Packet Buffer is configured to work like a single FIFO for multiple incoming packets. The Results Buffer is similarly also configured to work as a single FIFO for storing the results of multiple packets. Setting this bit to '1' causes the PM2329 to work as a multi-channel device, where multiple contexts on the external processor can control individual channels within the PM2329.

- **Enable PI Field**

If this bit is '1', the PM2329 interprets the first 64 bits that it receives as packet data, as the Packet Information (PI) field. If this bit is '0' the Packet Information is taken from the Packet Information Register. The PI contains the Packet Attribute as well as fields which inform the FEE in the PM2329 how the header is to be extracted. Setting this bit to '0' saves one write cycle during packet transfer but causes all packet headers to be processed similarly.

- **Enable OCC Field**

If this bit is set to '1', the PM2329 assumes that the OCC is contained within the first 64-bits of packet data that it receives after the PI Field (if present). If '0' then the OC Conductor (OCC) is taken from the OCC Register. The OCC contains instructions for packet processing in terms of the sequence of OCs to be executed. Setting this bit to '0' saves a write cycle during packet transfer but causes all packets to be processed in a like manner.

- **Direction Specifier**

There are several mechanisms by which packet direction can be indicated to the PM2329. This bit specifies to the PM2329 what source it will use to determine the direction sense of the packet as per the table below.

Table 22 Direction Specifier Bit

Direction Specifier Bit	Packet data from Packet Source (DMA)	Packet data from Processor	Comment
0	PI indicates direction (Enable PI Field = 1)	PI indicates direction (Enable PI Field = 0)	PI controlled direction
1	PSPD Pin indicates direction (Sampled on the last word [EOP] transferred)	Address of Packet Buffer Input Register which processor writes to indicates direction (Sampled on the last word [EOP] transferred)	Hardware Controlled direction

4.2.2.13 Channel Assignment Register (CAR; 8250h)

Access Mode: Read/Write, Global

Bit Range	Size	Name	Value after Reset
63:32	32	(Reserved)	Undefined
31:0	32	Channel [31:0] Assignment	FFFF FFFFh

This register is only valid in the multi-channel mode (see Operation Control Register).

In multi-channel mode, the PM2329 can support up to 32 simultaneous channels. For this reason, the Packet Buffer is comprised of 32 segments, one dedicated to each channel 0 through 31. Similarly, the Results FIFO is also split into 32 segments corresponding to 32 input channels. In this way up to 32 external contexts can use the 32 PM2329 channels simultaneously, in an interleaved manner, without conflict. See Context Support description in Chapter 3 for further information.

Each Packet Buffer segment is 256 bytes long and is associated with a Results FIFO of 8 entries.

If a larger Packet Buffer or a larger Results FIFO is required, the PM2329 allows multiple adjacent segments to be concatenated to generate a larger segment. Packet Buffer and Result segments are both concatenated in this case. With concatenation, the total number of channels is reduced accordingly.

Segments can be concatenated by writing a '1' to the bit of this register associated with the first channel, and a '0' to all subsequent channels, to be concatenated. Thus as an example, segments 5 through 8 can be concatenated by writing '1' to the Assignment bit for channel 5, and writing '0' to the Assignment bits for channels 6, 7, and 8. Provided the bit for channel 9 is '1', this will assign a 1 KB Packet Buffer and a 32 deep OC Result FIFO to channel 5. Channels 6, 7 and 8 will now become unavailable.

As explained in the example above, a set of (one or more) higher numbered channels can be concatenated with a lower numbered channel to create a bigger (lowest numbered) input channel (note that channel concatenation does not wrap around from channel 31 to channel 0, consequently Channel 0 assignment bit[0] is ignored and tied high permanently).

4.2.2.14 OC Conductor Register (OCCR; 8258h)

Access Mode: Read/Write, Global

Bit Range	Size	Name	Value after Reset
63:0	64	OC Conductor	0000 0000 0000 0000h

This register contains the OC Conductor to specify the OCs to be executed. The OC Conductor can be supplied either using this register or as part of the packet stream (preceding the packet header).

The format of the OCC Field, which comes in preceding the packet header, is same as the OC Conductor (OCC) register format. If the OCC Field is received from the Packet Source, the content of this register is not used and its content are not changed. If the OCC Field is not supplied with the packet, this register should be written by the processor prior to supplying the current packet and it should not be updated by the processor until after the EOP for the associated packet has been transferred to the device.

If Trace OC Execution Enable or Processor controlled OC Sequencing Enable bits are set, the initial OC processing starts using the content of this register or the OCC Field received from the Packet Source depending on the state of the Enable OCC Field in the OPCR register. Once the wait condition is reached, further OC processing continues using the content of the Alternate OCC register supplied by the processor.

The register can be read at any time by the processor to find out which OCC is currently active.

The OCC register supports two formats to specify the OCs to be executed as shown below. It can contain:

1. a set of up to 4 OC Identifiers OCIDs, or
2. the address of an E-RAM location that contains a valid C-Word.

Table 23 OC Conductor Register format

Format	Bits 63:48	Bit 47:32	Bits 31:16	Bits 15:0
0	0:OCID ₁ [14:0]	0:OCID ₂ [14:0]	0:OCID ₃ [14:0]	0:OCID ₄ [14:0]
1	1xxx xxxx xxxx xxxx	xxxx xxxx xxxx xxxx	xxxx xxxx xxxx xxxE ₁₆	E[15:0]

If bit [63] of the OCC register is '0' then the OCC register contains up to four OC Identifiers (OCIDs). The first OCID is located at bits [62:48], the following OCIDs at bits [46:32], [30:16] and [14:0], respectively. The OD Identifier format is shown below.

If bit [63] of the OCC register is '1' then the OCC register contains the address of the first C-Word to be executed from the E-RAM, OCC [16:0] supply the address.

- **OC Identifier Format**

Each OCID has 15-bits, these fields are defined as follows:

- **Bit [14] D-Word Present:** If this bit is set to '1', it indicates that D-Words corresponding to this OC are present in the E-RAM. If this OC results in a match condition, the D-Words will be updated depending on the setting of the D-Word Update Control field in the C-Word. The C-Word and D-Word are fetched from the E-Word corresponding to the matched cell.

When performing OCC controlled sequencing, fields other than the D-Word Update Control field in the C-Word are ignored.

- **Bits [13:8] OC Descriptor Index:** This field specifies the index of the descriptor in the OC Descriptor table which will be used for executing this OC.
- **Bits [7:0] Cascade OC Enable:** This field specifies which of the eight PM2329 devices in the cascade will participate in the OC. When the enable bit is set, the corresponding device will execute the OC. If the enable bit is reset, the corresponding device will not execute the OC. Bit 0 corresponds to PM2329 device 0, bit 7 corresponds to PM2329 device 7 in the Cascade.

If this field is set to 00h (i.e., all devices are disabled), it signals an invalid OCID and the current OC execution will be terminated and the following OCIDs, if any, will be ignored. Bits [14:8] of the OCID in this case are not used and should be set to zeroes.

4.2.2.15 Packet Information Register (PIR; 8260h)

Access Mode: Read/Write, Global

Bit Range	Size	Name	Value after Reset
63	1	Direction Bit	0
62	1	L3 Header Extraction Enable	0
61	1	Ethernet Framing Enable	0
60	1	Layer 4 Extraction Enable	0
59	1	OC Sequence Control Mode	0
58	1	OC Trace Enable	0
57:48	10	Layer 3 Header Offset	00 0000 0000
47:0	48	User Defined Packet Attribute	0000 0000 0000h

This register contains fields (control information) that define how the packet header is to be extracted. Additionally, this register also contains a 6-byte user defined packet attribute field. Note that the control information used to extract values from the header can either come from the Packet Information Register, or can be supplied as part of the packet stream. If this control information is supplied as part of the packet stream, this register is not used and remains unchanged. If the Packet Information is not supplied with the packet, this register should be written by the processor prior to supplying the current packet and it should not be updated by the processor until after the EOP for the associated packet has been transferred to the device.

- **Direction Bit**

This bit specifies the value of the Direction Bit associated with the packet. This bit is used only if the Direction Specifier bit in the Operation Control Register is reset to '0'.

- **L3 Header Extraction Enable**

If this bit is '0', then 108 bits from the first two 64-bit words of the packet data are taken and used in place of the extracted header. This allows the PM2329 to accept header information from pre-extracted packets.

If this bit is '1', then Layer 3 header extraction is enabled, and the header extraction is carried out based on the setting of the Ethernet Framing Enable and Layer 4 Extraction Enable control bits, explained below.

- **Ethernet Framing Enable**

If this bit is '0', then the Layer 3 Header Offset field is used as supplied in this PI word.

If this bit is '1', then the PM2329 Field Extraction Engine assumes that the packet is an Ethernet frame starting at offset 0, and the Layer 3 header offset is calculated. SIP, DIP and Protocol fields are then extracted.

- **Layer 4 Extraction Enable**

If this bit is '0', then the SP, DP and Flag (SYN, FIN and ACK) fields are loaded with the default values.

If this bit is '1', then the PM2329 Field Extraction Engine (FEE) identifies whether the L4 header is TCP or UDP. If it is TCP, then the FEE extracts the SP, DP and Flag (SYN, FIN and ACK) fields. The RST bit in the TCP Flags field is also extracted; however this bit is used for updating the TCP state D-Word only. If the L4 header is UDP, the FEE extracts the SP and DP fields only.

The following pseudo-code shows the header extraction flow based on the setting of these three control bits.

```
If (L3 Header Extraction Enable [62] == 0)

    Pre extracted header

else ([62] == 1)

    if (Ethernet Framing Enable [61] == 0)

        Use supplied L3 offset to extract SIP, DIP and Protocol

        if (L4 Header Extraction Enable [60] == 1) AND (Protocol == TCP/UDP)

            Compute L4 offset and extract SP, DP and Flags

        else

            Load SP, DP and Flags with default values

        endif

    else ([61] == 1)

        If (EtherType == IP 0x800)

            Parse packet to determine L3 offset and extract SIP, DIP and Prot

            Carry out L4 extraction shown above

        else

            Load SIP, DIP, Protocol, SP, DP and Flags with default values

        endif

    endif

endif
```

If a field cannot be extracted due to an extraction error, or if the extraction was disabled, then the corresponding fields in the extracted header are loaded with the default values shown below:

SIP:0FFFF FFFFh

DIP:0000 0000h

Protocol: 0000h

SP: 0000h

DP: 0000h

Flags (SYN FIN and ACK): 0

- **OC Sequence Mode**

This bit is a don't care if the OC Trace Enable bit is '1'.

If the OC Trace Enable bit is '0' and this bit is '0', OC Sequencing is automatic and PM2329 terminates packet processing at the end of the OC sequence.

If the OC Trace Enable bit is '0' and this bit is '1', OC Sequencing is under processor control. At the end of the current OC Sequence specified in the OCC word, the PM2329 retains the current packet and enters a wait condition. It permits the processor to control the next sequence of OCs to be executed. For further description of this bit, see OC Sequencing description in Chapter 5.

- **OC Trace Enable**

When this bit is 1, the PM2329 executes an OC and enters a break condition (if the OC just executed was not the last OC in the sequence) or a wait condition (if the OC just executed was the last OC in the sequence). For further description of break and wait conditions, see the Alternate OCC register description.

The table below shows the operation of the device based on the setting of the OC Sequencing Mode and OC Trace Enable bits.

Table 24 Processor Controlled OC Sequencing & Trace OC Execution

OC Trace Enable	OC Sequence Mode	Operation
0	0	Automated OC sequencing operation (original PM2328 compatible operation)
0	1	Processor controlled OC sequencing (wait after each OC sequence)
1	X	Trace OC sequencing--break after each OC execution and wait after each OC sequence

- **Layer 3 Header Offset**

These 10 bits specify the start of the Layer 3 header with respect to the start of the packet. This field is used only if Auto L3 Header Extraction is enabled and the Ethernet Framing bit is '0'.

- **User Defined Packet Attributes**

This field contains 6 bytes of user defined attributes. These can serve as one of the Data Sources for an OC (see OC Descriptors). This is useful for running an initial lookup for the packet to determine which OC sequence to execute.

4.2.2.16 Timer Register (TMR; 8268h)

Access Mode: Read/Write, Global

Bit Range	Size	Name	Value after Reset
63	1	(Reserved)	Undefined
62:48	15	Prescaler	0000h
47:44	4	(Reserved)	Undefined
43:32	12	Divider	000h
31:24	8	(Reserved)	Undefined
23:0	24	Timestamp	00 0000h

The contents of this register are used to update the aging information or Timestamp D-Word in the E-RAM. The Timestamp field of this register is used to update the specific D-Word defined as Timestamp. The Divider field of this register specifies the frequency at which the Timestamp field should be updated. This frequency is derived by dividing the system clock (SCLK) by a 15 bit prescaler and then by the divider field specified in this register. The register can be reprogrammed with a new Divider or Timestamp by writing to it. Whenever this register is written, the internal prescaler is loaded with the newly written value.

The Timestamp field is incremented at the following frequency:

$$(F_{SCLK} / (2^{15} - \text{Prescaler} + 1)) / \text{Divider}$$

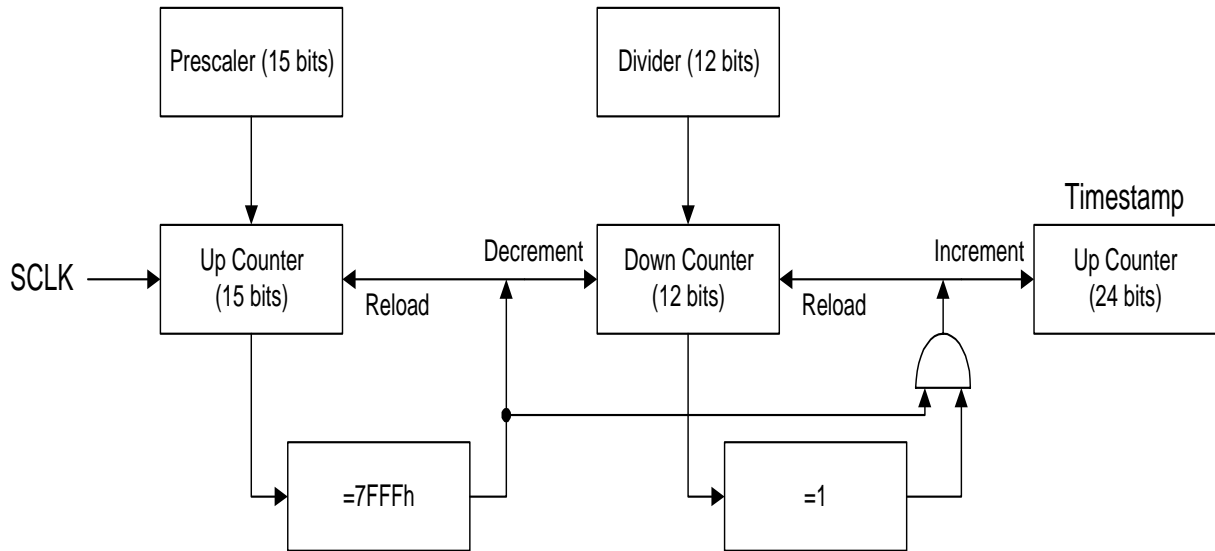
As an example, if the SCLK frequency is 66 MHz, the prescaler output period is 0.4915 msec. As an example, the table below shows the rate at which the Timestamp will be incremented for some example settings of Divider value assuming the prescaler was loaded with 0h.

Table 25 Timestamp Increment Interval Example (SCLK 66.67 MHz)

Divider	Timestamp Increment Interval	Unit
0	2.0133	sec
1	0.4915	ms
50	24.576	ms

Figure 27 shows an overview of the Timer logic.

Figure 27 Timer Logic



4.2.2.17 Alternate OCC Register (AOCC; 8270h)

Access Mode: Read/Write, Global

Bit Range	Size	Name	Value after Reset
63:0	64	OC Conductor	0000 0000 0000 0000h

Writes to this register are recognized by the device only if the current OC execution has reached halt state--either a break or a wait condition--as a result of one of the following two control settings:

1. OC Trace Enable is set (Trace On), or
2. OC Trace Enable is reset and OC Sequence Mode is set (Processor controlled sequencing On).

When a break condition (end of OC execution) occurs, the processor can either Terminate Sequence or Continue Sequence. When a wait condition (end of OC sequence) occurs, the processor can either Terminate Sequence or Start New Sequence.

The processor writes to the AOCC register to issue Terminate Sequence, Continue Sequence or Start New Sequence commands in the following manner.

1. To issue a Terminate Sequence command that terminates the current packet processing, the processor writes an EPP (End of Packet Processing) word (0000 XXXX XXXX XXXXh) to this register.
2. To issue a Continue Sequence command that continues the execution of remaining OCs in the sequence, the processor can write any non-EPP word to this register.
3. To issue a Start New Sequence command that starts executing a new OCC sequence on the current packet, the processor can write a new OCC word to this register. The format of this register is identical to the OCC Register format described earlier.

Note: To issue the Continue Sequence or Start New Sequence commands, the processor writes a non-EPP word to the AOCC. This word is interpreted as "Continue Sequence" if a break condition has occurred (that is, the OC just executed is not the the last OC of the programmed sequence). It is interpreted as "Start New Sequence" if a wait condition has occurred (that is, the OC just executed is the last OC of the programmed sequence).

In the automated OC sequencing operations (Trace OC Enable is reset and OC Sequence Mode bit is reset), this register is ignored.

Switching Between Sequencing Types

Sequencing can switch from OCC Sequencing to E-RAM Sequencing or vice versa, when a new sequence is initiated following a Halt condition. The processor writes to the AOCC register to initiate a sequence, selecting the type of sequencing desired by choice of AOCC bit 63, regardless of the previous state of bit 63.

When the hardware is in the "WAIT" state (that is, after completing an OC sequence), the PM2329 interprets a processor write to the AOCC as follows:

- If AOCC[63] = '1', the PM2329 initiates E-RAM sequencing, jumping to the specified C-word address.
- If AOCC[63] = '0' and AOCC[55:48] = '00000000', the PM2329 terminates sequencing.
- If AOCC[63] = '0' and AOCC[55:48] Not= '00000000', the PM2329 initiates OCC sequencing as specified by the AOCC[55:48] value.

When the hardware is in the "BREAK" state (that is, after any OC except the last one), the PM2329 interprets a processor write to the AOCC as follows:

- If AOCC[63] = '1', the PM2329 continues the next OC.
- If AOCC[63] = '0' and AOCC[55:48] = '00000000', the PM2329 terminates sequencing.
- If AOCC[63] = '0' and AOCC[55:48] Not= '00000000', the PM2329 continues with the next OC.

4.2.2.18 Packet Buffer Input Register

(PBIR; Base 3 +00h, +08h,... +0E8h, +0F0h)

(PBIR, EOPD0; Base 3 +0F8h)

(PBIR, EOPD1; Base 0 +08h)

Channel Register

Access Mode: Write Only, Global

The external processor or the Packet Source device (or DMA Source) writes to this register to transfer the packet data to the PM2329. Data written to this register fills the corresponding Packet Input Buffer.

The PM2329 supports two mechanisms to load the Packet Input Buffer. A FIFO-like load mechanism similar to the PM2328, and an SRAM-like write mechanism to support DMA capabilities of some network processors. While the Packet Input Buffer itself is a FIFO (in the data path between the system interface and the Field Extraction Engine) that supports a single write port type addressing, the FIFO address logic also supports an SRAM like addressing mechanism. The last transfer of the packet that indicates the end of packet and also packet direction information must always be done to one of two separate EOP addresses. When utilizing SRAM like addressing mechanism, the last transfer address is arranged to indicate EOP-Direction 0 in an efficient manner.

For each channel (as determined by the base address), multiple address offsets are assigned to this register. Depending on the offset, direction and end-of-packet information are communicated to the PM2329. This offset assignment is as follows:

Register	Address (64-bit Write)	Address (32-bit Write)	Direction	EOP
PBIR	Base 3 +00H, +08H,..., +0E8H, +0F0H	Base 3 +00H, +04H, +08H,..., +0F0H, +0F4H	X	No
PBIR EOPD0	Base 3 +0F8H	Base 3 +0FCH	0	Yes
PBIR EOPD1	Base 0 +08H	Base 0 +0CH	1	Yes

All data for the packet must be written to the PBIR address except the last write, which is written to the appropriate EOP (EOPD0 or EOPD1) address. For the last packet data write, valid data must be left justified and the rest of the word padded by nulls.

The tables below shows the writes to be performed to transfer 64-, 96- and 128-bit packets in 32- or 64-bit modes. This can be used as a guideline for other packet sizes. Also, this table shows the write cycles

required assuming FIFO addressing (all writes at the same address, except the EOP word) is to be used.

FIFO Addressing		64-bit		32-bit	
Packet Size	Direction	0	1	0	1
64 bits	Not EOP	None		Base 3 +00h	
	EOP	Base 3 +0F8h	Base 0 +08h	Base 3 +0FCh	Base 0 +0Ch
96 bits	Not EOP	Base 3 +00h		2x (Base 3 +00h)	
	EOP (Note1)	Base 3 +0F8h	Base 0 +08h	Base 3 +0FCh	Base 0 +0Ch
128 bits	Not EOP	Base 3 +00h		3x (Base 3 +00h)	
	EOP	Base 3 +0F8h	Base 0 +08h	Base 3 +0FCh	Base 0 +0Ch

Note: For 64-bit mode, the last 32 bits must be left justified (63:32) and the lower 32 bits (31:0) should be padded with zeroes.

The next table shows the write cycles required when SRAM-like addressing is utilized.

SRAM Addressing		64-bit		32-bit	
Packet Size	Direction	0	1	0	1
64 bits	Not EOP	None		Base 3 +any offset other than 0FCh	
	EOP	Base 3 +0F8h	Base 0 +08h	Base 3 +0FCh	Base 0 +0Ch
96 bits	Not EOP	Base 3 +any offset other than 0F8h		Base 3 +any offset other than 0FCh	
	EOP (Note1)	Base 3 +0F8h	Base 0 +08h	Base 3 +0FCh	Base 0 +0Ch
128 bits	Not EOP	Base 3 +any offset other than 0F8h		Base 3 +any offset other than 0FCh	
	EOP	Base 3 +0F8h	Base 0 +08h	Base 3 +0FCh	Base 0 +0Ch

Note that when using the SRAM like addressing to input the packet data (non EOP words), the lower address bits (SA[7:3], or SA[7:2]) are don't care whereas the EOP word address is fixed depending on the Direction to be indicated (EOP D0 is Base 3 +0F8h; EOP D1 is Base 0 +08h [64 bit mode]). This allows the packet data to be input using a block memory transfer mechanism where the destination address increments. Packet data to be transferred with Direction set to 0 can be transferred efficiently using a single block transfer--depending on the number of words to be transferred, the transfer can be started at the appropriate starting offset such that the EOP word gets the last packet word.

For example, when operating in 64-bit mode, in order to transfer 256-bit packet (four 64-bit words) with Direction set to 0, block transfer can be started at Base 3 +0E0h with a transfer count of 4. The first three

words will be written to Base 3 +0E0h, +0E8h and +0F0h and the EOP word will be written to +0F8h, as required.

In case the transfer Direction is 1, block moves can still be used; however, the EOP D1 must be written at a separate non-contiguous address using a separate write cycle.

When channels are concatenated, the EOP address of the concatenated channel is the EOP address of the highest channel in the concatenated set. For example, if channels 0 through 3 are concatenated to form a 1K deep channel, address +000 through +3F0 act as non-EOP addresses and +3F8 will be the EOP D0 address.

If automatic header extraction is disabled for this packet, then the first two 64-bit words (or first four 32-bit words) in the packet can contain pre-extracted header data for use with policies which inspect the packet header. If the pre-extracted header is to be compatible with the FEE extracted header (so that the same classification rules

may be applied to this packet), these words must be formatted as follows.

1st Word	63.....32 SIP		31.....0 DIP		
2nd Word	63.....48 SP	47.....32 DP	31.....24 Protocol	23...20 Flags	19.....0 (Reserved)

The Flags field is further defined as follows:

Bit Position	Flag
23	Reserved; set to 0 ^a
22	ACK
21	SYN
20	FIN

a. Regardless of the source of the header information (pre-extracted or on-chip extraction), the DIR bit is always derived as described in the Direction Specifier control bit description in Operation Control Register and inserted into this bit position in the internal header holding register.

The Packet Input Buffer is organized as 32 segments of 256 bytes each. Regardless of the state of Enable Multi-Channel Mode bit, packets are stored starting at segment boundaries.

In the single channel mode, up to 32 packets of up to 256 bytes each can be transferred into the Packet Input Buffer. Note that if a packet exceeds 256 bytes, the next byte is placed in the next segment and the full 256-byte segment is also assigned to this packet. As an example, if the processor downloads a 264-byte packet, only 30 additional packets up to 256 bytes each can be input.

In the multi-channel mode, up to 32 packets of 256 bytes each can be input. For further information regarding packet input in multi-channel mode, refer to the Channel Assignment register description.

4.2.2.19 Channel Status Register (CSR; Base 0 +00h)

Channel Register

Access Mode: Read Only, Global

Bit Range	Size	Name	Value after Reset
63:8	56	(Reserved)	Undefined
7	1	Channel Result FIFO Full	0
6	1	Channel OC Sequence Halted	0
5	1	Channel OC Sequence Halt Condition	Undefined
4	1	(Reserved)	Undefined
3	1	Channel Packet Buffer Available	1
2	1	Channel Result Available	0
1	1	Channel OC Sequence Terminated	0
0	1	(Reserved)	Undefined

This register indicates the status of the corresponding channel. In single-channel mode, only CSR0 is valid. Since this is a Channel Register, each channel has a corresponding Channel Status register when the PM2329 is in the multi-channel mode. When channels are concatenated, the status for the concatenated channel are all available by reading the Channel Status register corresponding to the lowest numbered concatenated channel. Reading a Channel Status register corresponding to an unassigned channel in the Channel Assignment Register will return invalid value.

- **Channel Result FIFO Full**

This bit is set when the result FIFO of the corresponding channel becomes full. If this condition is not serviced by the processor, the internal processing engine will eventually stall since any results that are generated can not be transferred to the Result FIFO.

- **Channel OC Sequence Halted**

This bit is set when the PM2329 has completed the current OC processing for the corresponding channel and...

1. OC Trace Enable bit is 0 and OC Sequence Mode bit is 1 (Processor controlled OC sequencing), or
2. OC Trace Enable bit is 1 and OC Sequence Mode bit is 0 (automated OC sequencing with Trace)

The OC Sequence Halt State bit indicates the reason for the halt condition.

- **Channel OC Sequence Halt Condition**

If the OC Sequence Halted bit is 0 then this bit is a don't care.

If the OC Sequence Halted bit is 1 and this bit is 0, the PM2329 has reached a break condition for the corresponding channel--it has completed execution of the previous OC and this just completed OC was not the last OC in the specified sequence.

If the OC Sequence Halted bit is 1 and this bit is 1, the PM2329 has reached a wait condition for the corresponding channel--it has completed execution of the previous OC and this just completed OC was the last OC in the specified sequence.

- **Channel Packet Buffer Available**

This bit is set to indicate that the corresponding Packet Buffer is empty. A Packet Buffer is empty as soon as the OC processing for that channel is complete. Depending on the single- vs. multi-channel mode setting, the behavior of the PBA bit is different as explained below.

In single-channel mode, the PBA bit is set as long as a 256 byte block is available in the Packet Buffer. This bit is cleared when the first packet data word is written to the last available 256 byte block in the Packet Buffer.

In multi-channel mode, the PBA bit is set as long as the 256 byte block corresponding to this channel is available. This bit is cleared when the first packet data word is written to the 256 byte block in the Packet Buffer. In multi channel mode with concatenation enabled, this bit (for the lowest numbered channel in the concatenated chain, all other concatenated channels in this chain are not used) behaves the same as the multi channel mode except the block size now refers to (N x 256) bytes where N is the number of channels concatenated together.

- **Channel Result Available**

This bit is set whenever one or more results are available in the corresponding Results FIFO. This bit is cleared when all of the results are read out by the processor.

- **Channel OC Sequence Terminated**

This bit is set whenever the packet processing is completed and all the results have been transferred into the Results FIFO corresponding to that channel. This bit is cleared when the Channel Status Register is read by the processor or the EOP for the corresponding channel is detected.

The STSR contains summary status of all CSRs. Note that the OC Sequence Terminated bit in the STSR is set when any of the Channel OC Sequence Terminated bit are set. Each OC Sequence Terminated bit is cleared when the corresponding register is read. If in response to an interrupt generated due to the OC Sequence Terminated condition, the processor reads the STSR (thereby clearing the OC Sequence Terminated bit in the STSR) without reading the CSRs (thus leaving the OC Sequence Terminated bit set in the CSRs) and another OC sequence is terminated, the PM2329 will set the OC Sequence Terminated bit in the STSR.

4.2.2.20 OC Results FIFO Output Reg (OCRF; Base 1 +00h)

Channel Register

Access Mode: Read Only, Global

Bit Range	Size	Name	Value after Reset
63	1	OC Match	Undefined
62	1	OC Result Valid	0
61	1	OC Seq Control Mode	Undefined
60	1	Packet First Result	Undefined
59	1	Packet Last Result / Sequence Terminate	Undefined
58	1	OC Done	Undefined
57	1	Trace Enable	Undefined
56	1	OCD Single or Multi-hit OC	Undefined
55:54	2	OCD OC Type	Undefined
53:48	6	OCD Index	Undefined
47:45	3	PM2329 ID Number	Undefined

If OC Type is 'Pattern Search OC':			
44:32	13	Pattern Search Match Offset	Undefined

If OC Type is 'Header OC':			
44:36	9	(Reserved)	Undefined
35:33	3	Extraction Error Code	Undefined
32	1	TTL Equal to 0	Undefined

31:24	8	Cascade OC Enable	Undefined
23	1	OC Match	Undefined
22	1	OC Result Valid	0
21	1	Data Result Available	Undefined
20	1	Match Rule Attribute	Undefined
19:17	3	(Reserved)	Undefined
16:0	17	Match Cell Number	Undefined

The results of packet processing are stored in the Results FIFO, which can be accessed by reading this register (read port). In 32-bit mode, the upper word (bits 63:32) of this register should be read first.

Regardless of the result of the processing (match, no match, invalid submission, etc.) at least one result is generated for each OC submitted.

This is a Channel Register; each channel has a corresponding Results FIFO Read Port when the PM2329 is in the multi-channel mode. When channels are concatenated, the result segments are also concatenated. Results for the concatenated channel will all be available by reading a single read port corresponding to the lowest numbered concatenated channel. Reading a Result FIFO Read Port register corresponding to an

unassigned channel in the Channel Assignment Register will return invalid value.

Every time the OC Results FIFO Read Port register is read, the OC Results FIFO advances and the next result is returned in the next read of the register. In case the PP performs 32-bit reads, it should read the higher word first and then the lower word since the FIFO advances when the lower word (bits 31:0) is read. Also see Data Result Valid bit below regarding sequence of reads to the OC Results FIFO and Data Results FIFO.

In the single-channel mode, only Channel 0 is valid.

- **OC Match**

This bit is set to '1' when an OC results in a match condition. Other fields of the result word contain further information about the result.

This bit is reset to '0' when no match was found.

This function appears in both bit 63 and bit 23, for compatibility with some network processors and the original PM2328, while still facilitating easy access in 32-bit mode.

- **OC Result Valid**

This bit is set to '1' for a valid result. All other bits in the result register are valid only if this bit is set. This bit is reset to '0' for an invalid result to indicate no result is available (other bits in this case are invalid).

Upon power-on reset, this bit is reset to 0.

This function appears in both bit 62 and bit 22, for compatibility with some network processors and the original PM2328, while still facilitating easy access in 32-bit mode.

- **OC Sequence Control Mode**

This bit is interpreted only if Trace Enable (bit 57) is '0'. Its description is as follows:

If this bit is '0' it indicates Automated OC Sequencing (AOS). The PM2329 performs OC sequencing without processor intervention.

If this bit is '1' it indicates Processor Controlled OC Sequencing (PCOS). The PM2329 stops after completion of an OC Sequence, and then waits for the Processor to issue a command. The command could be to start another OC Sequence or to End Packet Processing. Note that in this mode, the PM2329 does not stop after every OC.

- **Packet First Result**

If this bit is set, it indicates that this result is the first result of the first OC run on the packet.

- **Packet Last Result / Sequence Terminate**

If this bit is set, it specifies that the last result for a packet is in the FIFO, and indicates termination of the OC Sequence.

- **OC Done**

If this bit is set, it indicates that this result was the last result for the specified OC. Every OC always generates at least one result and a single-hit OC generates only one result.

- **Trace Enable**

If this bit is ‘1’ then the PM2329 is in “Trace Mode”. It will stop after every OC and wait for the processor to issue a command. In this mode, OC Sequence Control Mode (bit 61) is ignored.

If this bit is ‘0’, the PM2329 mode is decided by OC Sequence Control Mode (bit 61).

- **OCD Single or Multi-hit OC**

This bit is valid only if the OC Match bit is set.

This bit is set to ‘1’ if this was a Multi-hit OC. It is reset to ‘0’ if this was a single-hit OC.

- **OCD OC Type**

This bit is valid only if the OC Match bit is set.

This field has the same value as the OC-Type field in the OCD which was used to run this OC. This field indicates a Header, Attribute, Pattern Search Short or Pattern Search Long OC.

- **OCD Index**

This bit is valid only if the OC Match bit is set.

This field contains the 6-bit index which gives the position of the Descriptor within the OCD Table.

- **PM2329 ID Number**

This field is valid only if the OC Match bit is set to ‘1’.

These bits specify which of the PM2329 devices has generated this result.

- **Pattern Search Match Offset**

This field is valid only if the OC was a Short or Long Pattern Search OC and the OC Match bit is set to ‘1’.

This field indicates the position at which a match occurred. The value returned is the absolute offset of the first byte for forward searches or last byte for reverse searches of the packet data which matched.

- **Extraction Error Code**

This field is valid only if the OC Type was Header.

This field indicate the cause of extraction error. They are defined as follows.

000	No Extraction Error
001	Incomplete Header Extraction due to Offset Error
010	Intermediate IP Fragment when L4 bit in PI was set
011	IHL field has a value less than 5 (20 Bytes)
100	IP version is not 4
101	(Reserved)
110	(Reserved)
111	(Reserved)

This field returns 000 for Attribute OCs.

- **TTL Equal to 0**

This bit is valid only if the OC Type was Header.

This bit is set to '1' to indicate that the Time To Live field within the IP header is '0'. This field is valid only if no extraction error occurred and it was not a pre-extracted header packet.

- **Cascade OC Enable**

This field contains a copy of the Cascade OC Enable field from the OCC of the packet corresponding to this result.

- **Data Result Available**

The condition when this bit is set are as follows:

1. OCC controlled sequencing is selected,
2. D-Word Update Enable was set for this OC,
3. ECR defined at least one D-Word, and
4. the OC results in a match condition.

or

1. E-RAM controlled sequencing is selected,
2. D-Word Update Enable was set for this OC,
3. ECR defined at least one D-Word,
4. the OC results in a match condition, and
5. the specified branch opcode was not Terminate.

If this bit is set, then some E-RAM data fields were either read or updated and are present in the Data Results FIFO. The operation of the Data Results FIFO is linked to the reading out of the OC Results FIFO. The Data Results, if required, should be read out before the next read access to the OC Results FIFO.

- **Match Rule Attribute**

This bit is valid only if the OC Match bit is set.

This bit is a copy of the Attribute bit within the rule for the cell which matched.

- **Match Cell Number**

This field is valid only if the OC Match bit is set to '1'.

This field returns the 17-bit cell number of the rule within the partition that resulted in the match. This cell number corresponds to a relative cell number, i.e., the cell number within the OC partition across multiple PM2329 devices regardless of the physical partition fragmentation within a device or across multiple devices.

For the following three conditions, the PM2329 will not execute an OC.

1. The OC Descriptor is invalid.
2. The OC Type is Pattern Search and the search range lies outside the packet.
3. The Column Select Mask of the specified partition is 00h.

Under these conditions, the following result is generated:

OC Match = 0, OC Done = 1

For the following three error conditions, the PM2329 cannot execute an OC, since it cannot identify a valid OC Descriptor:

1. With OCC sequencing, the Cascade ID Mask for the first OCID is 00h.
2. OCC specifies E-RAM sequencing, but the ECR specifies no E-RAM is present.
3. With E-RAM sequencing, each C-Word has the Cascade ID Mask set to 00h and the last C-Word specifies the Terminate condition.

Under the above three conditions, the result shown below is generated:

- OC Match = 0
- Result Valid = 1
- Data Result = X
- Packet First = 1
- Packet Last = 1
- OC Done = 0

4.2.2.21 Data Results FIFO Output Register

(DRFO 0/1; Base 1 +08h)
(DRFO 2/3; Base 2 +08h)
(DRFO 4/5; Base 2 +10h)
(DRFO 6; Base 2 +18h)

Channel Register

Access Mode: Read Only, Global

In the Extended mode of operation, the PM2329 supports external E-RAM. Besides Control Words, this E-RAM array can also contain Data Words. These Data Words can contain information such as Packet Count, Byte Count, Timestamp, TCP State or User Defined Results Field. If a match is found while running an OC then a result is generated and placed in the OC Results FIFO. Additionally, the PM2329 also accesses the Data Words in E-RAM and updates them as needed. Note that the PM2329 will update the D-Words regardless of the depth level. However, see the Note below regarding depth and D-Words returned in the Data Results FIFO.

Typically the processor might want information about the Data Words associated with the match result. To support this, the PM2329 contains a 32-bit Data Results FIFO that is synchronized with the OC Results FIFO. For each OC Results FIFO entry which has the Match bit set, the PM2329 will copy the information from the Data Word E-RAM, corresponding to the matched cell location, into the Data Results FIFO. This is only done by the PM2329 that is configured to connect to External Data RAM.

This makes it possible for the processor to first get information about a match condition, and then to read Data Results FIFO Read Port to get information such as Packet Count, Byte Count, Timestamp, State and User Defined results. This saves the processor explicit read requests to the E-RAM to get the same information. Providing User Defined Results also saves the processor translation time which it would otherwise spend in translating the cell number to final results. The information read from the Data Results FIFO will always correspond to the last OC Result read from the OC Results FIFO.

The Data Results FIFO is organized at 4 consecutive addresses so that a long E-Word can be transferred easily out of the PM2329 for either 32- or 64-bit accesses.

Table 26 Data Results FIFO Output Register (64-bit mode)

Bit Range	Register Name	Size	Responding Device	Address
63:32, 31:0	Data Result FIFO Output Register 0 (DW0 / DW1)	64	CID 0, CID 1	Base 1 + 08h
63:32, 31:0	Data Result FIFO Output Register 2 (DW2 / DW3)	64	CID 2, CID 3	Base 2 + 08h
63:32, 31:0	Data Result FIFO Output Register 4 (DW4 / DW5)	64	CID 4, CID 5	Base 2 + 10h
63:32	Data Result FIFO Output Register 6 (DW6)	32	CID 6	Base 2 + 18h

Table 27 Data Results FIFO Output Register (32-bit mode)

Register Name	Size	Responding Device	Address
Data Result FIFO Output Register 0 (DW0)	32	CID 0	Base 1 + 08h
Data Result FIFO Output Register 1 (DW1)	32	CID 1	Base 1 + 0Ch
Data Result FIFO Output Register 2 (DW2)	32	CID 2	Base 2 + 08h
Data Result FIFO Output Register 3 (DW3)	32	CID 3	Base 2 + 0Ch
Data Result FIFO Output Register 4 (DW4)	32	CID 4	Base 2 + 10h
Data Result FIFO Output Register 5 (DW5)	32	CID 5	Base 2 + 14h
Data Result FIFO Output Register 6 (DW6)	32	CID 6	Base 2 + 18h

Note that the Data Results FIFO returns values associated with depth level 0 only. Words at other levels, if defined, must be accessed using E-RAM indirect addressing mechanism.

(The only exception is in case of a single device connected to a single external E-RAM memory device with both ECD and EDD buses connected to the same physical memory device. In this case, D-Word 0 at the next sequential address is accessed and returned in the Data Results FIFO.)

4.3 Indirectly Addressable Locations

This section describes indirectly addressable locations in the PM2329 and in the E-RAM, which fall in one of the following two classes:

- Rule Memory Cells
- E-Words

4.3.1 Rule Memory Cells

Indirect Access using: Rule Indirect Registers

Access Mode: Local

Rule Memory Cells are accessed by indirect addressing using the Rule Indirect Address, Command and Data Register Set. The Rule Memory Cells store the Rule Data and the Rule Control fields used when an OC is executed. The Rule Memory Cells are 136 bits in size.

The format of the Rule Memory Cells is same as the Rule Indirect Data Register Set.

Detailed rule operation is described in Chapter 5.

4.3.2 E-RAM Words

Indirect Access using: E-RAM Indirect Registers

Access Mode: Global

E-RAM is accessed by indirect addressing using the E-RAM Indirect Address, Command and Data Register Set. A single location in the E-RAM is referred to as an E-Word. The E-Word is made up of several 32-bit words, which are either a C-Word or D-Words. The size of the E-Word depends on the number of PM2329 devices in the cascade (See E-RAM Indirect Data Register Set).

The format of the C-Word is given below. This must be written to the upper half of the C-Word/D-Word Register of the E-RAM Indirect Data Register Set.

Control Word

Bit Range	Size	Name	Value after Reset
31	1	Reserved	Undefined
30	1	D-Word Present	Undefined
29:24	6	OC Descriptor Index	Undefined
23:16	8	Cascade OC Enable	Undefined
15:12	4	D-Word Update Control	Undefined
11:8	4	Branch Opcode	Undefined
7:0	8	Immediate Address	Undefined

- **D-Word Present**

If this bit is set to '1', it indicates that D-Words corresponding to this OC are present in the E-RAM. If this OC ends in a match condition, the D-Words will be updated depending on the setting of the D-Word Update Control field.

If this bit is reset to '0', no D-Word updates are performed based on the result of the current OC.

Note, when OCs are executed from E-RAM, the D-Word fields corresponding to the first C-Words fetched (to determine the OC to be executed) are updated regardless of the state of this bit provided the updates are enabled using the Update Control bits in the C-Word.

- **OC Descriptor Index**

This field specifies the index of the descriptor in the OC Descriptor table which will be used for executing this OC.

- **Cascade OC Enable**

This field specifies which of the eight PM2329 devices in the cascade will participate in the OC. When the enable bit is set, the corresponding device will execute the OC. If the enable bit is reset, the corresponding device will not execute the OC. Bit 16 corresponds to PM2329 device 0, bit 23 corresponds to PM2329 device 7 in the Cascade.

If this field is set to 00h (i.e., all devices are disabled), it signals an invalid OCID and current OC execution will be skipped and based on the branch code in the current C-Word, the next C-Word will be fetched and processed.

- **D-Word Update Control**

This field specifies which of the D-Word fields (if defined) need to be updated. The 4 bits are defined as follows.

Bit 15	Update Byte Count
Bit 14	Update Packet Count
Bit 13	Update Timestamp
Bit 12	Update TCP State

- **Branch Opcode**

This 4 bit field specifies the following opcodes.

0000	Continue
0001	Return
	Return address is popped from an internal single-level stack.
0010	Goto Immediate
0011	Call Immediate
	Return address is pushed into an internal single-level stack.
0100	Goto immediate address on match, else continue
0101	Call immediate address on match, else continue
0110	Goto immediate address on match else terminate
0111	Call immediate address on match, else terminate
1000	Goto cell number address on match, else continue
1001	Call cell number address on match, else continue
1010	Goto cell number address on match, else terminate
1011	Call cell number address on match, else terminate
1100	Goto cell number address on match, else goto immediate
1101	Call cell number address on match, else call immediate
1110	(Reserved)
1111	Terminate

The cell number address is the relative address of the rule within the OC partition that generated the match. The address is used as the 17-bit (target) address in the E-RAM.

- **Immediate Address**

This is an 8-bit field which specifies one of 256 C-Words to branch to using a goto or a call opcode. Hence, the target for any immediate address must always lie within the first 256 locations of the E-RAM C-Words.

Byte Count or Packet Count D-Word

Bit Range	Size	Name	Value after Reset
31:0	32	Byte Count or Packet Count	Undefined

Timestamp and State D-Word

Bit Range	Size	Name	Value after Reset
31:28	4	(Reserved)	Undefined
27:24	4	TCP State	Undefined
23:0	24	Timestamp	Undefined

User Defined Data D-Word

Bit Range	Size	Name	Value after Reset
31:0	32	User Defined Data	Undefined

- **Byte Count**

The D-Word containing Byte Count is incremented by the Byte Count for the current packet. This is determined by the FEE header extraction. The 32-bit value will wraparound unless it is cleared periodically by the processor. If the FEE cannot determine the Byte Count field, it is assumed to be '0'.

- **Packet Count**

The Packet Count is incremented by 1 every time this D-Word is updated. The 32-bit value will wraparound unless it is cleared periodically by the processor.

- **TCP State**

This D-Word field is updated to reflect the current state of the packet.

- **Timestamp**

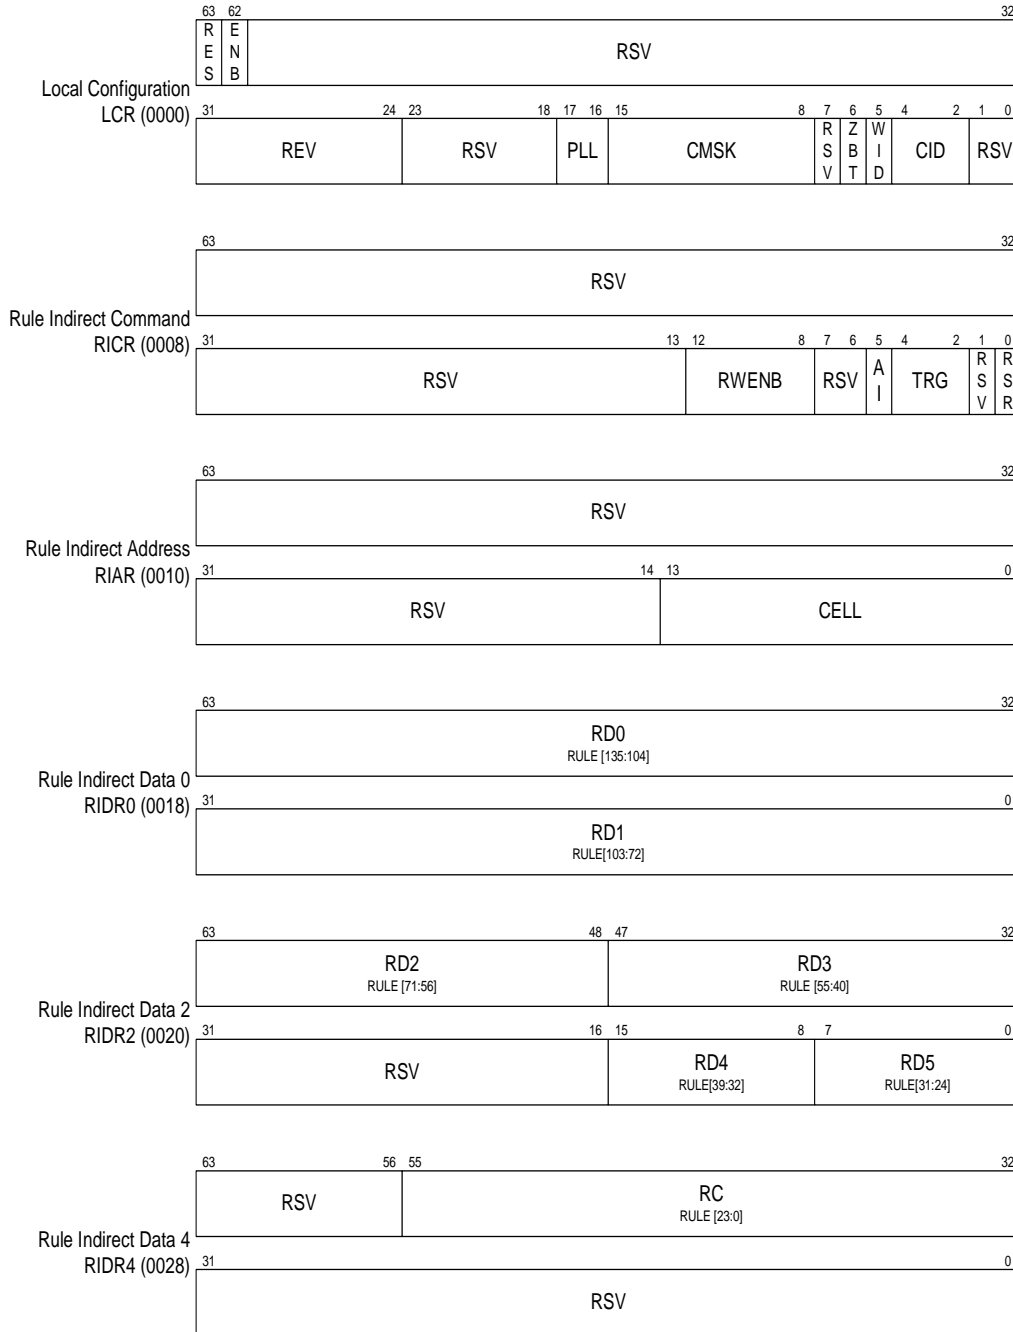
This D-Word is updated with the current Timestamp value from the Timer Register in the PM2329.







- **User Defined Data**







This field can be set to any 32-bit value. It provides a lookup mechanism for OC matches.






Note: Updates of Byte Count and TCP State fields are valid only if a prior OC using header has been executed.

4.4 Register Summary









5 Rule Formats and OC Sequencing

5.1 Rule Formats

The PM2329 can perform powerful policy-based search operations. These operations involve comparing the packet data fields with a pre-loaded set of policy rules in the policy database of the device. Results of search operations are returned in the Results FIFO. Additionally, based on the result of the operation, further operations can be performed on the same packet.

PM2329 rules have been designed to handle a wide range of packet processing applications. Key features of the rule format are listed below:

- Supports multi-operand (up to 6) operations within an instruction
- EQ, GE, or LE (unsigned integers) operations
- Bit-wise mask capability
- Force Match function
- Rule Attribute bit
- Composite Rules made of up to 4 rules
- Rule Negation function


Rule Structure

All PM2329 policy rules have a uniform structure, consisting of two parts: a Rule Control field (RC) and a Rule Data field (RD). The Rule Control field determines the operation to be performed on the packet data and policy contents of the Rule Data field.

A policy rule is 136 bits wide, comprised of the 24-bit Rule Control field and the 112-bit Rule Data field. Rule Control and Rule Data fields are further divided into six sub-fields. This allows different operations to be performed on the Rule Data sub-fields, with the operations being specified by the corresponding sub-fields in the Rule Control sub-fields.

Figure 28 shows the organization of the 136-bit Rule memory word.

Figure 28 Rule Control and Data Field



5.1.0.1 Rule Data Field

The Rule data field is 112 bits wide and is equal to the width of the PM2329 policy rules which are compared with the packet data field. These 112 bits are divided into sub-fields, each of which can be used independently of the others. There are a maximum of 6 sub-fields: two 32-bit fields (RD0 and RD1), two 16-bit fields (RD2 and RD3) and two 8-bit fields (RD4 and RD5). Sub-fields RD2 and RD3 can be used separately as two 16-bit sub-fields, or they can be merged into a single 32-bit sub-field.

By programming the Rule Control field appropriately, the individual Rule Data sub-fields can effectively be combined into wider fields to perform Equal operation.

Some Rule data sub-fields can act as mask fields for other rule data sub-fields. A rule data sub-field and the corresponding packet data sub-field are masked with the associated mask field (if enabled) and are then compared. This enables the PM2329 to implement power-of-2 (binary) range compare operations. Further details about masking are provided in Section 5.1.3.

5.1.0.2 Rule Control Field

The Rule Control Field controls the selection of packet data fields, operations to be performed on the rule and packet data fields, and other actions related to the rule. Each rule has a set of Common Control (CC) bits and Field Control (FC) bits associated with each rule data sub-field.

Since the rule data is divided into a number of sub-fields, the packet data must also be divided into a compatible set of sub-fields. The sub-fields are then compared with the corresponding rule data sub-fields.

Common Control Bits

There are three common control (CC) bits used to specify global properties of the rule, or for indicating the action to be performed on a rule match. These bits are described in Table 28:

Table 28 Common Control Rule - CCR Bits

Rule Control Bits	No of Bits	Description
Rule Attribute	1	This bit is returned as part of the result and can be interpreted by the user as desired
Composite Rule Bit	1	Specifies if the rule is part of a composite rule. Up to 4 rules can be combined to form a composite rule. This bit cannot be used with a Long Pattern Search OC and must be reset to '0' in that case.
Negation Bit	1	This bit allows negation of the match result for the entire rule. In case the rule is part of a Composite Rule, result negation is applied for each Rule, and then the results from all appropriate rules are combined.

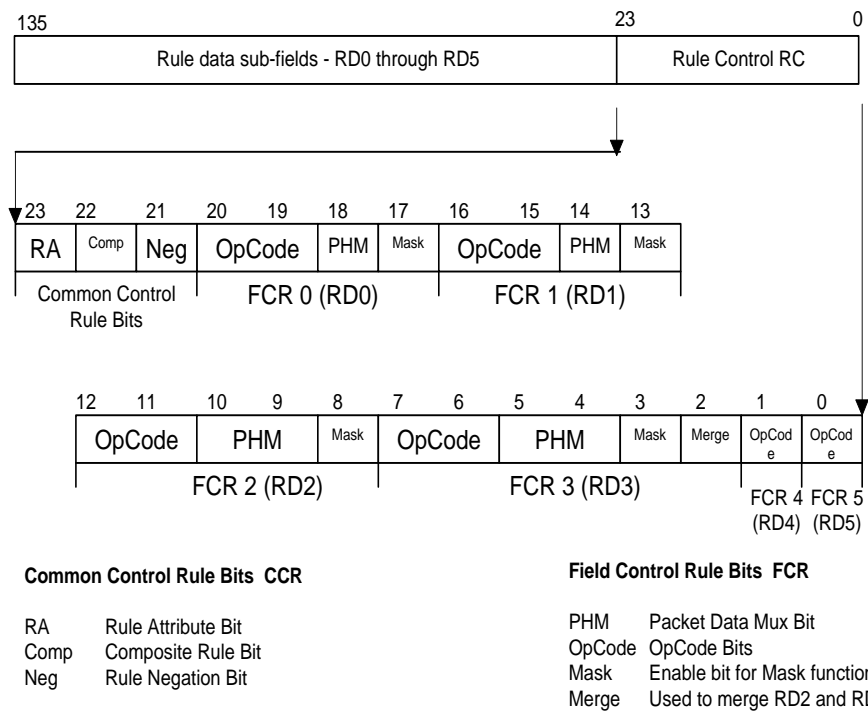
Field Control Bits

The field control (FC) bits corresponding to each rule data sub-field specify the following:

- **OpCode** bits specify the operation on the corresponding data sub-field and the packet data.
- The **Packet Data Mux (PDM)** selection bit specifies the association of fields in the packet data with the rule data sub-fields.
- The corresponding **Mask** bits specify whether the mask field for this rule data sub-field should be applied.
- The **Merge** bit (specific to RD2 and RD3 only) specifies that the two rule data sub-fields are to be combined into a wider 32-bit field.

A brief overview of the Rule Control fields is presented below. For further details, contact SwitchOn Networks.

Figure 29 Rule Control Fields



5.1.1 Rule Operations

The PM2329 supports different operations that can be performed on the rule data fields. The six rule data fields can be controlled separately to perform different operations on each of these fields. Operations supported by the PM2329 are as follows:

- Equal To (with or without masks)
- Greater Than or Equal To (with or without masks)
- Less Than or Equal To (with or without masks)
- Match

Equal To, Greater Than or Equal, and Less Than or Equal operations compare the corresponding packet and rule data fields for the specified condition using unsigned integer arithmetic. The Match operation is used to force a TRUE result regardless of the input operands.

Each of the six rule data sub-fields supports operations as shown in the table below.

Table 29 Operations Supported for Rule Data Sub-fields

Operation	RD0	RD1	RD2	RD3	RD4	RD5
Match (00 or 0)	Yes	Yes	Yes	Yes	Yes	Yes
EQ (01 or 1)	Yes	Yes	Yes	Yes	Yes	Yes
GE (10)	Yes	Yes	Yes	Yes	-	-
LE (11)	Yes	Yes	Yes	Yes	-	-

Using the above operations, it is possible for the PM2329 to perform the following operations:

- Binary range compares
- Non-Binary range compares
- Longest prefix (LP) matches
- Pattern or String Searches

5.1.2 Masking

Comparison operations between rule and packet data fields can be qualified by bit mask fields. Using bit masks, it is possible to store a "don't care" (as opposed to a '1' or a '0') in the PM2329 rule data bits, for comparison against packet data bits.

Both the rule data sub-fields and their associated mask bits are stored in the same PM2329 rule where some of the rule data sub-fields are used as the mask fields for the non-mask rule data sub-fields. This sub-field correspondence (data vs. mask) is fixed, as shown in Table 30 below. This imposes some restrictions on the number of mask fields, but improves the memory efficiency of the PM2329 policy rule set.

Table 30 Masked sub-field and Associated Mask Source

Rule and Packet Data Field (Size)	Bits Masked	Mask Source
RD0 (32 Bits)	All 32 bits	RD1 (32 bits)
RD1 (32 Bits)	All 32 bits	RD2 (16 Bits) & RD3 (16 Bits)
RD2 (16 Bits)	All 16 bits	RD3 (16 Bits)
RD3 (16 Bits)	Lower 8 bits only	RD4 (8 Bits)
RD4 (8 Bits)	None	None
RD5 (8 Bits)	Upper 4 bits	Lower 4 Bits of RD5

In general, each rule data sub-field can be masked from the next rule data field. However, the following special cases should be noted.

1. The first Rule data sub-field RD0 cannot act as the mask field for any other sub-field.
2. The upper 8 bits of RD3 cannot be masked using this mechanism; however, RD3 (all 16 bits) can be ignored using appropriate rule programming.
3. RD4 can never be masked; however, RD4 can be ignored using appropriate rule programming.
4. The last Rule data sub-field, RD5, is divided into a 4-bit data and a 4-bit mask field. The 4-bit mask field cannot be used as a data sub-field.

Also note that while a sub-field has been enabled as a mask field, it is still used to generate the result as specified in the corresponding FCR sub-field. It is up to the user to program the rules and interpret the results appropriately. Alternatively, if a field is used as a mask field, its corresponding Opcode field could be set to Match, to effectively ignore the result of the processing of the corresponding processing unit.

5.1.3 Composite Rules

The PM2329 rule set supports creation of Composite rules using the Composite Rule enable bit. Composite rules combine up to 4 (or less) rules into one wide-policy rule. This composite rule would result in a compare condition, which is a combination of all the rules that constitute the composite rule. A composite rule set must lie within the same rule pack.

Note that the same packet data bits are compared against the composite rule. A composite rule generates a match only if all members of the composite rule match the packet data bits. This makes it possible for the PM2329 to impose highly complex policies on the packet data.

For all other purposes, a composite rule behaves in the same way as any other single rule. The packet data bits are compared with each of the member rules, as specified by the individual rule control bits. The results of the individual compares are logically ANDed for generating the match result of the composite rule.

In case of a Long Pattern Search OC, the Composite Rule enable bit should be reset to '0', since this OC uses multiple rules (16 total) by default.

5.1.4 Rule Attribute Bit

Each PM2329 Rule contains a Rule Attribute Bit. This bit does not participate in Rule processing operations; it is intended for use by the external processor. When a particular rule matches, the match result will contain the Attribute Bit taken from the rule that matched. This provides an additional degree of flexibility for the external processor software when handling match results.

5.1.5 Rule Negation

The PM2329 performs operations between various data sub-fields of the rule and packet data as described earlier, and returns a combined result of this operation. The Rule Negation bit allows the returned result to be negated. For example, if this bit is set to '0' and an equality check is performed between rule and data sub-fields and the packet and rule data are equal, the result will indicate a match condition. However, if this bit is set to '1' for the same example, the result will indicate a "no match" condition.

5.2 Operation Cycles

The **Operation Cycle (OC)** is the basic classification operation performed by the PM2329 on the data extracted from a packet. It consists of comparing a string of extracted data against a set of rules and returning the relative offset of the rule that matched (single-hit OC), or a set of relative offsets of rules that matched (multi-hit OC).

OCs specify the set of rules on which the OC will be executed (also called the Rule Partition) and the type of extracted packet data that will be classified (also known as the Data Source).

The Rule Partition defines a subset of the Policy Database that will be used during the execution of an OC. The Policy Database is organized with a total of 16 columns and 64 rows. Each column is further made up of 16 rule cells. A Rule Partition is defined by two sets of values:

- a) The set of columns containing the rules, and
- b) The set of rows containing the rules within the partition.

For a given Rule Partition, individual columns (between 1 and 16, in any combination) can be specified. The rows of a Rule Partition must be defined with a start and end row (both inclusive). Given this definition scheme, note that the columns of the partition need not be contiguous, whereas the rows within a column will always be contiguous. The minimum partition granularity permitted by this scheme is 16 rule cells (a single row of a single column) and the maximum is the entire Policy Database rule memory in the device. The number of rule cells in a Rule Partition will always be a multiple of 16.

The Data Source can be one of the following types:

- a) **Extracted Packet Header**, which can contain the extracted IP and also TCP/UDP fields, or alternatively can be pre-constructed by the external processor.
- b) **Packet Attribute**, which is the 48-bit field obtained from the Packet Information field.
- c) **Packet Data**, for pattern search OCs which can be taken from any starting byte offset within the packet. Pattern Search OCs could be applied to search for up to 12-byte (short) strings or up to 192-byte (long) strings.

Additionally, OCs can be specified to be single-hit or multi-hit OCs. When the OC is executed, it can either result in no matches, or yield one or more matches. In case one or more matches occur, a single-hit OC will return the only match or the highest priority match, whereas a multi-hit OC will return all the matches in the order of priority.

5.3 OC Sequencing

5.3.1 OC Conductor

The PM2329 is capable of running different OCs on the same packet. The entire set of OCs that are applied to the same packet is called an **OC Sequence**.

The OC Sequence is specified by the **OC Conductor (OCC)** word associated with each packet. The OCC can be either specified one time in the OCC Register, or it can be supplied preceding each packet when the packet is written into the PM2329.

The OCC format can be one of two types:

- a) The OCC contains up to 4 OC Identifiers, or
- b) The OCC contains the address of an E-RAM location.

Format (a) results in OCC controlled sequencing, where a fixed set of up to 4 OCs are executed on the corresponding packet. No conditional OC sequencing is possible, and the E-RAM is not used for sequencing, since the sequence is fixed. However, D-Word updates are supported and performed, if enabled.

Format (b) results in E-RAM controlled sequencing, where the next OC to be executed is defined by a Control Word stored in the E-RAM. The OC sequence can now be variable where the next OC depends on the result of the previous OC. D-Word updates are performed, if enabled.

The PM2329 contains a table of 64 OC Descriptors. Each entry in this table contains an OC Descriptor that can specify up to 64 different types of OCs. Additional information including the OC descriptor format is available in Chapter 4.

5.3.2 OCC Sequencing

The PM2329 supports Automated (PM2328 compatible) OC Sequencing or Processor controlled OC Sequencing. Additionally, the PM2329 can be programmed to Trace OC execution for debug.

Various control bits, status bits, and registers determine the sequencing operation. They are summarized below:

Table 31 OCC Sequencing Control Bits

Location	Control Bits
In PI Word	OC Sequence Control Mode
	OC Trace Enable
In Interrupt Enable Register	OC Processing Halted Enable

Table 32 OCC Sequencing Status Bits

Location	Status Bits
In Status Register	OC Processing Halted
	OC Processing Halt Condition

Table 33 Registers Applicable to OCC Sequencing

Registers	Action	Function
OCC		Determine OCs Executed (OCID word or ERAM pointer)
AOCC		Determine OCs Executed (OCID word or ERAM pointer)
AOCC	write to: -- Terminate -- Continue -- Start New Sequence	Determines OC sequence executed in Processor Controlled mode or with Trace Enabled

5.3.2.1 Trace Feature

OC Trace Enable can be set to debug the OC processing regardless of the OC Sequencing employed (Automated or Processor Controlled). With the OC Trace Enable bit set, the PM2329 will enter a Halt state at the end of each OC execution. The Status register will indicate the Halted state and the Halt condition (Break or Wait). The processor can examine the state of the device between OC execution by reading various registers, rule memory or ERAM memory locations. The OC execution can be terminated or continued, as explained below.

Note that processor intervention is required if the Trace feature is enabled, regardless of the OC Sequencing Mode employed.

When the PM2329 enters the Halt condition (Break or Wait), the processor can use the AOCC register to determine how the OC processing is to continue. The operations and conditions for their employ are listed below.

1. Terminate Sequence

When a Break or Wait condition occurs, the processor can write an EPP word to the AOCC register. This terminates the current packet processing, and packet is discarded.

2. Continue Sequence

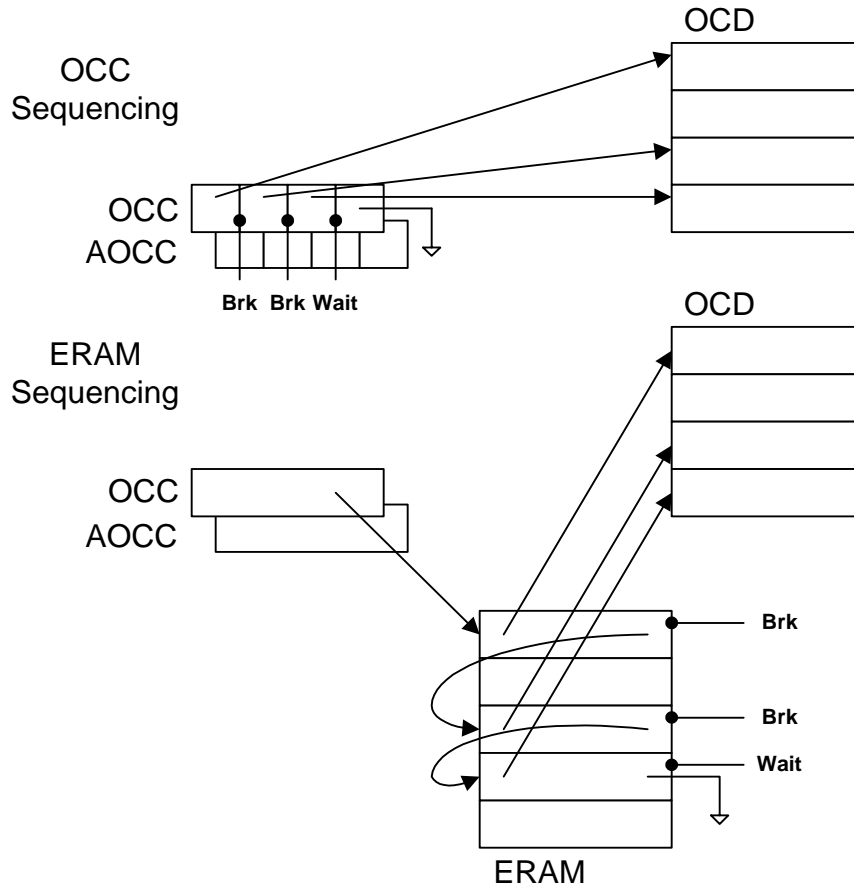
When a Break condition (end of OC execution) occurs, the processor can issue a Continue Sequence command by writing a non-EPP word to the AOCC register. This causes the PM2329 to continue processing the current packet with the next specified OC.

3. Start New Sequence

When a wait condition (end of OC sequence) occurs, the processor can Start New Sequence by writing a non-EPP, OCC format-compatible word to the AOCC register.

Figure 30 shows the Halt state (Break or Wait) entered by the PM2329 as the OC execution proceeds with the Trace feature enabled.

Figure 30 Trace Feature



Note 1. When a Break condition occurs, the processor can Continue Sequence or Terminate Sequence using the AOC Register.

Note 2. When a Wait condition occurs, the processor can Start New Sequence or Terminate Sequence using the AOC Register

5.3.2.2 Sequencing Control Modes

OC Sequence Control Mode bit determines if Automated or Processor controlled Sequencing is enabled.

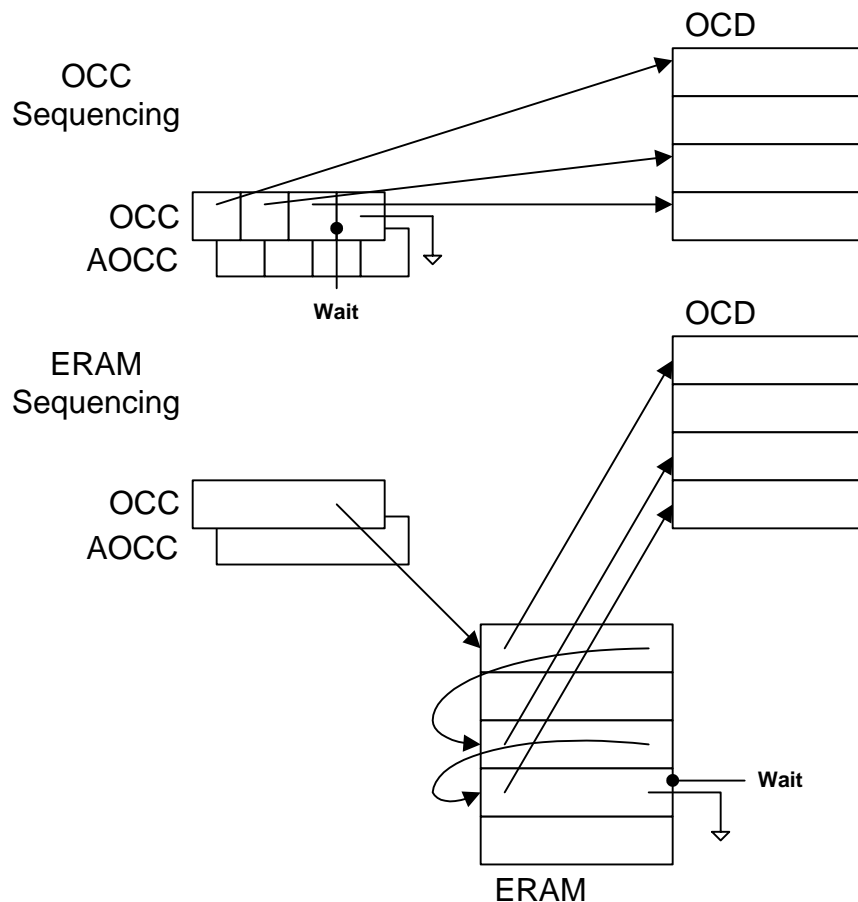
5.3.2.2.1 Automated Sequencing

When the OC Sequence Control Mode bit is reset and OC Trace Enable bit is reset, the PM2329 operates in **Automated Sequencing** (PM2328 compatible) mode. Refer to the OCC or ERAM Sequencing description below for device operation. In this mode, after the specified sequence has been executed, the packet is discarded.

5.3.2.2.2 Processor Controlled Sequencing

When the OC Sequence Control Mode bit is set and OC Trace Enable bit is reset, the PM2329 operates in **Processor Controlled Sequencing** mode. In this case, the device will enter a Halt state at the end of OC sequence, and set the Halt status to indicate a Wait condition. Refer to Figure 31 and the Trace feature description above regarding the options available to the processor in this case (Terminate Sequence or Start New Sequence). In this mode, the packet is retained until the processor issues a Terminate Sequence command.

Figure 31 Processor controlled Sequencing



Note 1. When a Wait condition occurs, the processor can Start New Sequence or Terminate Sequence using the AOCC Register

When Processor Controlled Sequencing is enabled, it is possible to switch between OCC Sequencing and ERAM Sequencing by writing the appropriate control word to the AOCC when the Start New Sequence command is issued.

5.3.2.3 OCC Sequencing

If the OCC contains an OCC word containing OC Identifiers (the structure of the OCC word is as described for the OCC Register in Chapter 4 along with the format of the OC Identifiers) then OCC Sequencing is employed.

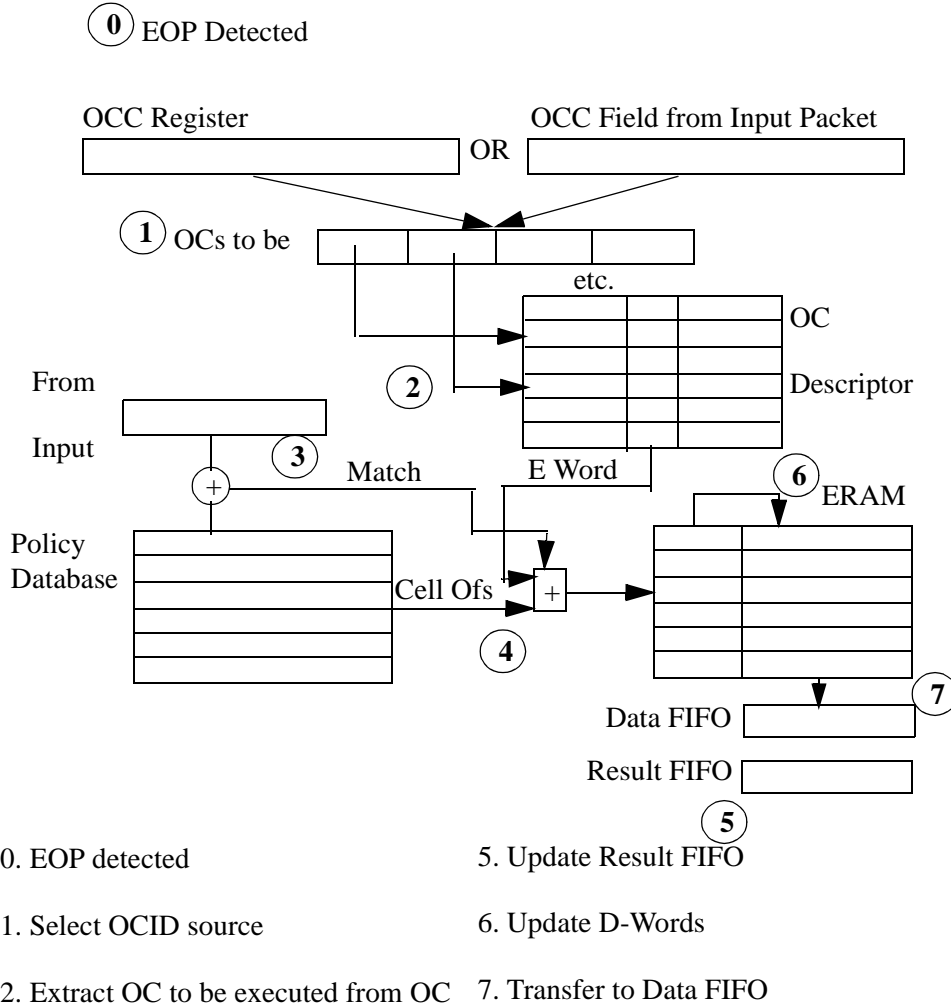
OC execution starts when the EOP for the packet is detected. The OCC can specify up to 4 OCs. This allows up to 4 OCs to be run in a fixed sequence. Packet processing terminates when an OC with the Cascade OC Enable field value set to 00 is encountered, or when the last (fourth) OC is executed.

When operating in the cascaded mode, each OC Identifier also identifies the set of devices in the cascade which will participate in the OC. The OC Descriptor Index across all devices in the cascade executing this OC must be identical.

While ERAM sequencing using C-Words (described in the next section) is not supported in this case, E-RAM D-Words can be updated, if enabled. In other words, all fields of the C-Word except the D-Word Update Control field are ignored, and D-Words are updated as specified in the D-Word Update Control field.

Figure 32 shows the sequence of events during OC execution when OCC sequencing is performed.

Figure 32 General Overview of OCC Sequencing



5.3.2.4 E-RAM Sequencing

If the OCC does not contain OC Identifiers, it must contain a pointer to an E-RAM address that contains a valid C-Word. In this case, an OC execution sequence is initiated with E-RAM sequencing.

In E-RAM sequencing, each C-Word contains an OC Identifier, a D-Word Update Control field and a Branch Condition.

For every C-Word executed, the following steps are performed by the PM2329 in the order shown:

1. Update the D-Words based on the D-Word Update Control Field
2. Execute the OC specified by the OC Identifier.
3. Upon completion of the OC and based on its results, execute the branch condition.

The branch conditions supported can be grouped into three categories, as follows:

Unconditional Branches

1. Terminate
No further OCs are executed on the current packet. The current packet processing is terminated.
2. Continue (with the next C-Word)
The next C-Word is fetched from the next E-RAM location and a new OC execution is started.
3. Goto immediate address
The next C-Word is fetched from the specified address in the immediate address field and a new OC execution is started.
4. Call immediate address
The address of the current C-Word is incremented and pushed onto a single-level stack. The next C-Word is fetched from the specified address and a new OC execution is started.
5. Return
Fetch the next C-Word after popping the single location stack.

2-Way Conditional Branches

These are conditional branches. Each opcode specifies an address to call or goto if a match occurs, and a default action in case of no match.

6. On match, goto immediate address else continue
7. On match, call immediate address else continue
8. On match, goto immediate address else terminate
9. On match, call immediate address else terminate

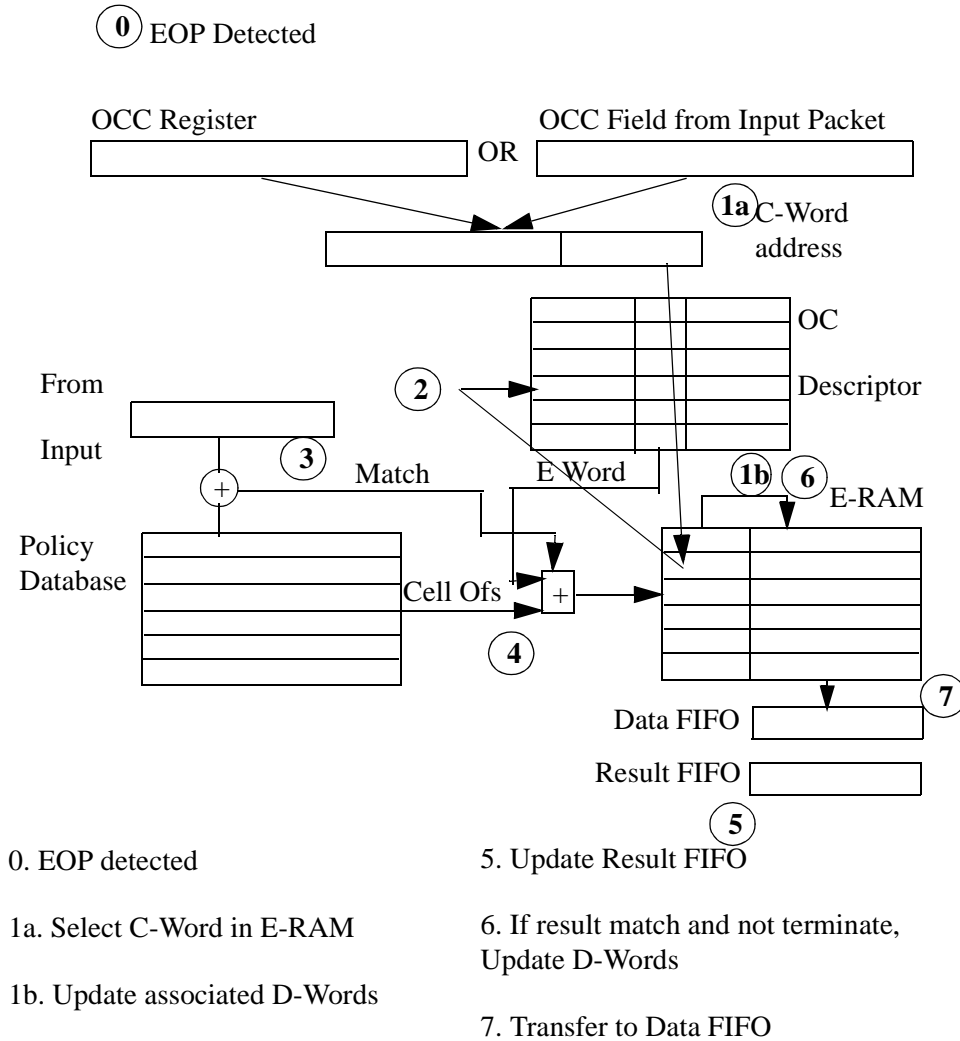
Multi-Way Conditional Branches

These are conditional branches with more than two possible outcomes. If the OC results in a match, then the branch taken depends on the cell number that matched. This depends on the E-RAM Offset and the Partition Offset specified for the OC executed. Each branch opcode also has a default action which is taken in case the OC does not result in a match.

- 10. On match goto cell number which matched, else continue
- 11. On match goto cell number which matched, else terminate
- 12. On match goto cell number which matched, else goto immediate
- 13. On match call cell number which matched, else continue
- 14. On match call cell number which matched, else terminate
- 15. On match call cell number which matched, else call immediate

Figure 33 shows the sequence of events during OC execution when ERAM sequencing is performed:

Figure 33 General Overview of E-RAM Sequencing



5.3.3 E-Word Association with Cells

Every OC (which requires E-RAM) is mapped to a set of E-Words. The mapping between a cell in an OC and the corresponding address in the E-RAM is done as follows:

- Each OC is assigned a block of E-Words in the E-RAM. Generally, the number of E-Words in this block would equal the number of cells in the OC (this is, however, not mandatory). An E-Word Segment Start Offset is programmed within each PM2329 to correspond to the start address of its E-Word block.

- A Partition Start Offset is further associated with each descriptor to specify the cell number offset of the OC partition within the entire OC. This is useful in cascade applications. For example, if an OC has 4K cells and has four partitions across four PM2329 devices where each partition is 1K cells in length, then the partition offsets within the four PM2329 devices should be programmed to be 0, 1K, 2K and 3K respectively. In a single PM2329 configuration, this field will be normally programmed to 0.

When a match occurs, the location of the E-Word to be fetched is calculated by adding the E-Word Segment Start Offset, the Partition Start offset, and the relative offset of the matching cell within that partition. This computation is performed for a Multi-way branch condition only. In case of a two way or unconditional branch, the E-Word address to be fetched is not dependent on the cell number.

6 Electrical and Timing Characteristics

Table 34 Absolute Maximum^a Ratings

Item	Parameter	Symbol	Conditions	Value	Unit
1	Power Supply Voltage (I/O)	V_{DD}	V _{SS} =0, CV _{SS} =0, AV _{SS} =0,	-0.5 to 4.0	V
2	Power Supply Voltage (Core)	CV_{DD}		-0.5 to 2.5	V
3	Input Voltage	V_{in}		-0.5 to $V_{DD} + 0.5$	V
4	Short Circuit Output Current	I_o		±20	mA
5	Electro-Static Discharge Voltage	V_{ESD}		±1000	V
6	Latch-Up Current	I_{LU}		±100	mA
7	Junction Temperature	T_{jmax}		125	°C
8	Storage Temperature	T_{stg}		-40 to 150	°C

a. Maximum ratings are those values beyond which damage to the device may occur.

Table 35 Recommended Operating Conditions

Item	Parameter	Symbol	Conditions	Min.	Typ.	Max	Unit
1	Power Supply Voltage (I/O)	V_{DD}	V _{SS} =0V	3.14	3.3	3.46	V
2	Power Supply Voltage (Core)	CV_{DD}	CV _{SS} =0V (see Note ^a)	1.44	1.5	1.56	V
			CV _{SS} =0V (see Note ^b)	1.54	1.6	1.66	V
3	Analog Power Supply Voltage	AV_{DD}	AV _{SS} =0V (see Note ^c)	1.44	1.5	1.56	V
			AV _{SS} =0V (see Note ^d)	1.54	1.6	1.66	V
4	Ambient Temperature	T_a		0	25	70	°C

- a. The required core power supply voltage when running at SCLK = 100MHz or below
 b. The required core power supply voltage when running at SCLK = 116MHz
 c. The required analog power supply when using core power supply voltage of 1.5V (nominal)
 d. The required analog power supply when using core power supply voltage of 1.6V (nominal)

General Notes:

- All parameters are characterized for DC conditions after thermal equilibrium has been established.
- Unused inputs must always be tied to appropriate logic level (either V_{SS} or V_{DD}).
- The input pin contains circuitry to protect the inputs against damage due to high static voltages or electric fields; however, it is advised that normal precautions be taken to avoid applying any voltage higher than the maximum rated voltages. For proper operation it is recommended that V_{in} be constrained to the range V_{SS} < V_{in} < V_{DD}.
- The device requires two digital power supplies: V_{DD} = 3.3V and CV_{DD} = 1.5V or 1.6V, and one analog power supply: AV_{DD} = CV_{DD}.
- Correct association between the SCLK frequency and the CV_{DD} supply voltage must be followed.

Table 36 Terminal Capacitance

Item	Parameter	Symbol	Conditions	Min.	Typ	Max	Unit
1	Input	C_I	Measured at V _{DD} =V _{IN} =V _{OUT} =V _{SS} f=1MHZ T _a =25°C	-	5	-	pF
2	Output	C_O		-	5	-	pF
3	Bidirectional	C_{IO}		-	5	-	pF

6.1 DC Characteristics

Table 37 DC Characteristics

Item	Parameter	Symbol	Conditions	Min.	Max	Unit
1	High-Level Input Voltage	V_{IH}	TTL Input	2.0		V
2	Low-Level Input Voltage	V_{IL}	TTL Input		0.8	V
3	Input Leakage Current (for Input-only pins without pull-up resistor)	I_{IH}	$V_{IN}=V_{DD}$		10	μA
		I_{IL}	$V_{IN}=V_{SS}$		10	μA
4	Input Leakage Current (for bi-directional pins with pull-up resistor)	I_{UPIH}	$V_{IN}=V_{DD}$		15	μA
		I_{UPIL}	$V_{IN}=V_{SS}$	20	100	μA
5	High-Level Output Voltage	V_{OH}	$I_{OH} = -400\mu A$	2.4		V
6	Low-Level Output Voltage	V_{OL}	$I_{OL} = 4mA$		0.4	V
7	Output Leakage Current Tri-State Output (for output-only pins without pull-up resistor)	I_{OZH}	$V_{OUT} = V_{DD}$		15	μA
		I_{OZL}	$V_{OUT} = V_{SS}$		15	μA
8	Max Current at V_{DD} Power Supply	I_{DDOP}	SCLK = 100MHz, SD, EDD load = 15pF EMA, ECD load = 20pF		see Note ^a	A
			SCLK = 116MHz, SD, EDD load = 15pF EMA, ECD load = 20pF		see Note ^b	A
9	Max Current at CV_{DD} Power Supply	CI_{DDOP}	$CV_{DD} = 1.5V + 4\%$, SCLK = 100MHz, ACLK = 200MHz		2.0	A
			$CV_{DD} = 1.6V + 4\%$, SCLK = 116MHz, ACLK = 232MHz		2.5	A

a. The maximum current at V_{DD} power supply (I_{DDOP}) is based on worst case simulation data. The following specification applies, Max = 0.5A (condition: SCLK = 100MHz).

b. The maximum current at V_{DD} power supply (I_{DDOP}) is based on worst case simulation data. The following specification applies, Max = 0.6A (condition: SCLK = 116MHz)

General Note:

1. All DC parameters are guaranteed across recommended operating conditions.

6.2 Switching Characteristics

6.2.1 Reset Timing Parameters

Table 38 Reset Timing


Item	Parameter	Symbol	Conditions	Min.	Max	Unit
1	V _{DD} rise time	T_{rVDD}		1		ms
2	CV _{DD} rise time	T_{rCVDD}		1		ms
3	AV _{DD} rise time	T_{rAVDD}		1		ms
4	Reset low time	$Trst$	After V _{DD} and CV _{DD} are stable ^a	10 100		ms Valid SCLK Cycles

a. The minimum reset low time should be the larger of the two given values.

6.2.1.1 VDD Power On Sequence

V_{DD} can be applied concurrently with, or prior to CV_{DD}. The recommended sequence is to apply V_{DD} prior to CV_{DD} and AV_{DD}. The recommended power-on sequence is illustrated in Figure 34.

Figure 34 Recommended V_{DD} Power On Sequence



6.2.2 Clock Timing Parameters

Table 39 Clock Timing

Item	Parameter	Symbol	Conditions	Min.	Max	Unit
1	SCLK period	T_{sck}	$CV_{DD} = 1.5V \pm 4\%$	10.0	100.0	ns
			$CV_{DD} = 1.6V \pm 4\%$	8.6	100.0	ns
2	SCLK rise time	T_{rsck}	10% to 90%	see Note ^a	see Note ^a	ns
3	SCLK fall time	T_{fsck}	10% to 90%	see Note ^b	see Note ^b	ns
4	SCLK high time	T_{hsck}	$CV_{DD} = 1.5V \pm 4\%$	4.0		ns
			$CV_{DD} = 1.6V \pm 4\%$	3.6		ns
5	SCLK low time	T_{lsck}	$CV_{DD} = 1.5V \pm 4\%$	4.0		ns
			$CV_{DD} = 1.6V \pm 4\%$	3.6		ns
6	SCLK to ECLKOUT skew	T_{skew}		-0.5	1.0	ns


a.The SCLK rise time is based on simulation data. The following specification applies, Min = 0.5ns, Max = 1.5ns.

b.The SCLK fall time is based on simulation data. The following specification applies, Min = 0.5ns, Max = 1.5ns.

General Notes:

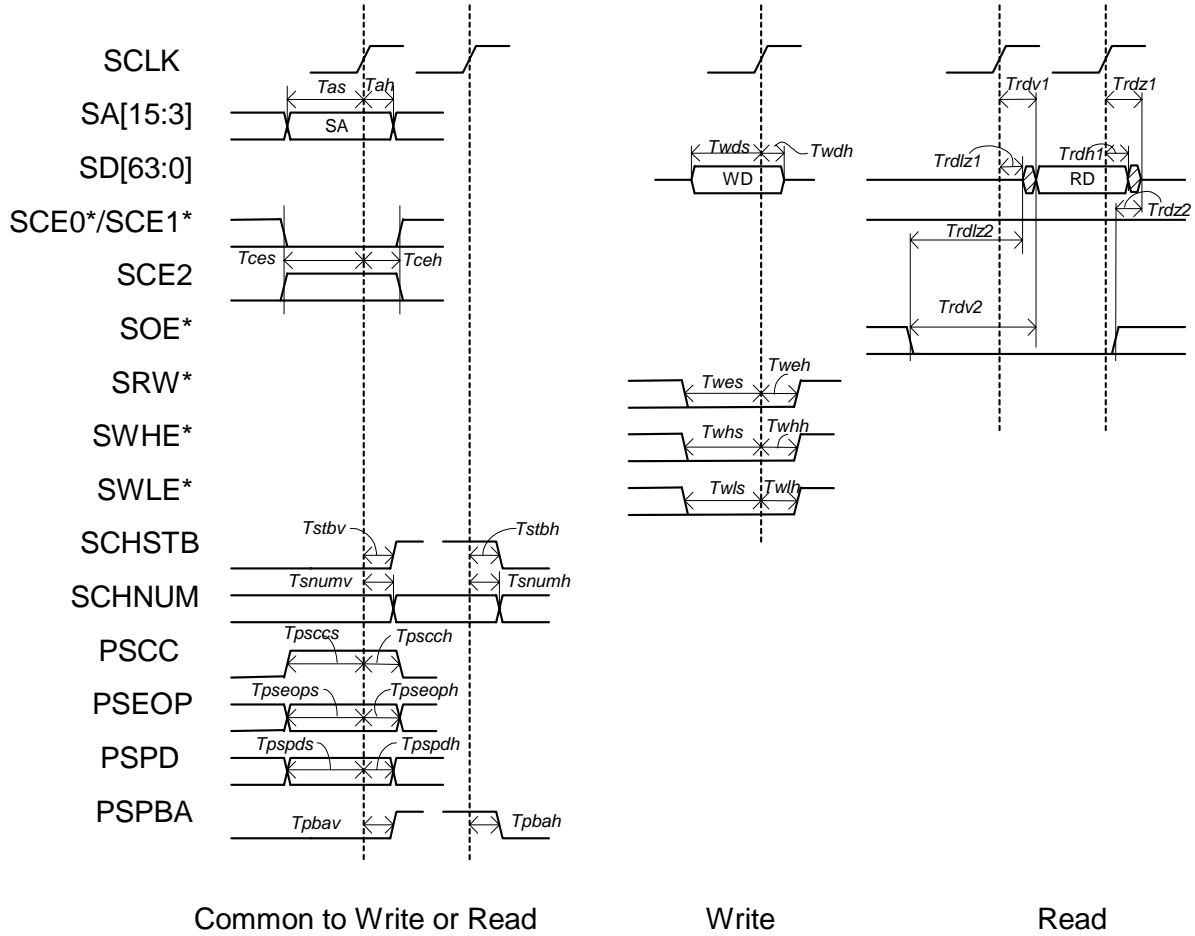
1. The specified input clock swing of 10% to 90% refers to percentages of nominal rail-to-rail supply voltage (0V to 3.3V).
2. Equivalent slew rates are 5.28 V/ns (max) and 1.76 V/ns (min).
3. SCLK duty cycle must ensure that the specified high and low times are met.
4. Correct association between the SCLK frequency and the CV_{DD} supply voltage must be followed.

Figure 35 SCLK to ECLKOUT Skew



6.2.3 System Interface Timing

Figure 36 System Interface Timing Parameters



This diagram shows timing relationships with respect to SCLK.
For functional timing, see Chapter 2


Table 40 System Interface Timing Parameters

Item	Parameter	Symbol	Conditions	Min.	Max	Unit
1	SA setup time	T_{as}		2.00		ns
2	SA hold time	T_{ah}		0.50		ns
3	SCLK to SD Lo-Z (Read)	$Trdlz1$		1.50		ns
4	SCLK to SD Valid (Read)	$Trdv1$	SCLK = 100MHz, $CV_{DD} = 1.5V \pm 4\%$		5.00	ns
			SCLK = 116MHz, $CV_{DD} = 1.6V \pm 4\%$		4.50	ns
5	SCLK to SD Invalid (Read)	$Trdh1$		1.15		ns
6	SCLK to Hi-Z (Read)	$Trdz1$			5.00	ns
7	SOE* to Lo-Z (Read)	$Trdlz2$		1.20		ns
8	SOE* to Data Valid	$Trdv2$			5.00	ns
9	SOE* to Hi-Z (Read)	$Trdz2$			4.50	ns
10	SCE0*, SCE1*, SCE2 setup time	T_{ces}		2.00		ns
11	SCE0*, SCE1*, SCE2 hold time	T_{ceh}		0.50		ns
12	SD setup time (write)	T_{wds}		2.00		ns
13	SD hold time (write)	T_{wdh}	SCLK = 100MHz, $CV_{DD} = 1.5V \pm 4\%$	1.00		ns
			SCLK = 116MHz, $CV_{DD} = 1.6V \pm 4\%$	0.90		ns
14	SRW* setup time (write)	T_{wes}		2.00		ns
15	SRW* hold time (write)	T_{weh}		0.50		ns
16	SWHE* setup time (write)	T_{whs}		2.00		ns
17	SWHE* hold time (write)	T_{whh}		0.50		ns
18	SWLE* setup time (write)	T_{wls}		2.00		ns
19	SWLE* hold time (write)	T_{wlh}		0.50		ns
20	SINT* Valid	T_{sintv}			5.00	ns
21	SINT* Invalid	T_{sintv}		1.50		ns
22	SCHSTB Valid	T_{stbv}			5.00	ns
23	SCHSTB Invalid	T_{stbh}		1.50		ns
24	SCHNUM Valid	T_{snumv}			5.00	ns
25	SCHNUM Invalid	T_{snumh}		1.50		ns
26	PSCC setup time	T_{pscscs}		2.00		ns
27	PSCC hold time	T_{pscch}		0.60		ns
28	PSEOP setup time	T_{pseops}		2.00		ns
29	PSEOP hold time	T_{pseoph}		0.50		ns
30	PSPD setup time	T_{pspds}		2.00		ns
31	PSPD hold time	T_{pspdh}		0.50		ns
32	PSPBA Valid	T_{pbav}			5.00	ns
33	PSPBA Invalid	T_{pbah}		1.50		ns

General Notes:

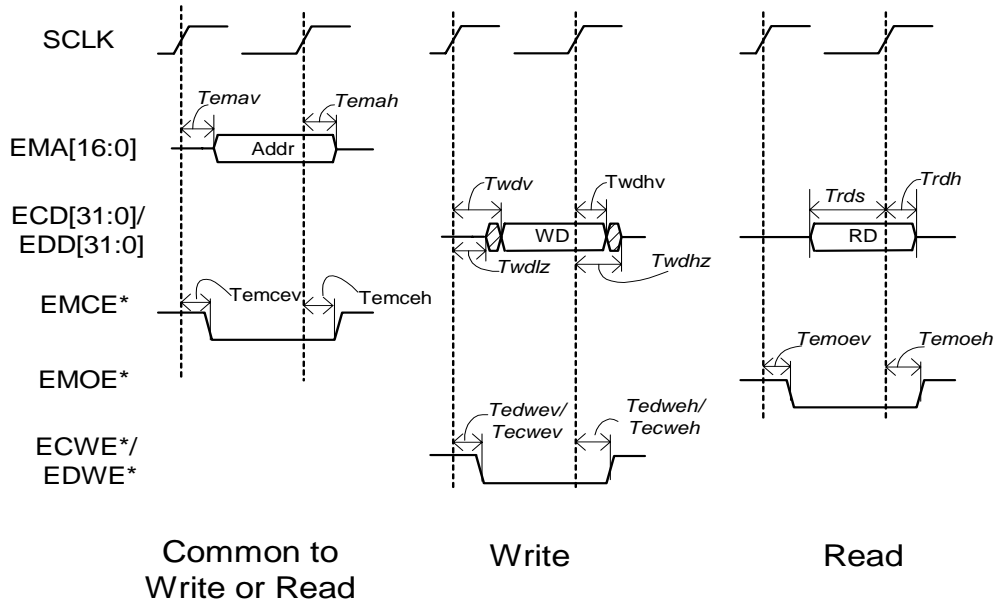
1. Depending on the clock to output delay of other devices on the system bus and the bus protocol characteristics (back-to-back read and write cycles), bus contention may occur.
2. The output load used in the timing measurement is shown in Figure 37.
3. Correct association between the SCLK frequency and the V_{DD} supply voltage must be followed.

Figure 37 Load Equivalents



6.2.4 E-RAM Interface Timing

Figure 38 E-RAM Interface Timing



This diagram shows timing relationships with respect to SCLK.
For functional timing, see Chapter 2

Table 41 E-RAM Interface Timing Parameters

Item	Parameter	Symbol	Conditions	Min.	Max	Unit
1	SCLK to EMA Valid	<i>Temav</i>	Max $C_L=20\text{pF}$		5.0	ns
2	SCLK to EMA Invalid	<i>Temah</i>	Max $C_L=20\text{pF}$	1.5		
3	SCLK to EMOE* Valid	<i>Temoev</i>	Max $C_L=20\text{pF}$		5.0	ns
4	SCLK to EMOE* Invalid	<i>Temoeh</i>	Max $C_L=20\text{pF}$	1.5		ns
5	SCLK to EMCE* Valid	<i>Temcev</i>	Max $C_L=20\text{pF}$		5.0	ns
6	SCLK to EMCE* Invalid	<i>Temceh</i>	Max $C_L=20\text{pF}$	1.5		ns
7	SCLK to EDWE*/ ECWE* Valid	<i>Tedwev/ Tecwev</i>	Max $C_L=15\text{pF}$		5.0	ns
8	SCLK to EDWE*/ ECWE* Invalid	<i>Tedweh/ Tecweh</i>	Max $C_L=15\text{pF}$	1.5		ns
9	SCLK to ECD/EDD Low-Z	<i>Twdlz</i>	C_L (see Note ^a)	1.5		ns
10	SCLK to ECD/EDD Valid	<i>Twdv</i>	C_L (see Note ^a)		5.0	ns
11	SCLK to ECD/EDD Invalid	<i>Twdhv</i>	C_L (see Note ^a)	1.0		ns
12	SCLK to ECD/EDD Hi-Z	<i>Twdhz</i>	C_L (see Note ^a)		5.5	ns
13	EDD/ECD setup time	<i>Trds</i>	C_L (see Note ^a)	2.0		ns
14	EDD/ECD hold time	<i>Trdh</i>	C_L (see Note ^a)	0.5		ns

a. The shown data bus timing parameters are based on ECD load of 20 pF and EDD load of 15 pF.

General Note:

1. The ERAM interface timing parameters are specified with reference to SCLK.
2. Correct ERAM device should be chosen according to clock speed.

6.2.5 Cascade Interface Timing

Figure 39 Cascade Interface Timing

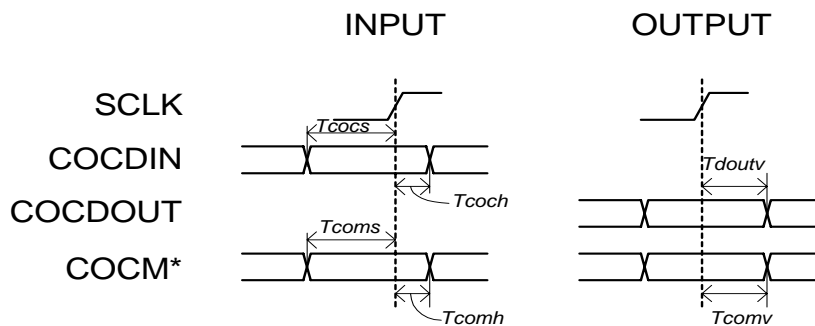


Table 42 Cascade Interface Timing Parameters

Item	Parameter	Symbol	Conditions	Min.	Max	Unit
1	COCS (input mode) setup time	T_{cocss}		1.2		ns
2	COCS (input mode) hold time	T_{cocsh}		0.5		ns
3	SCLK to COCS (output mode) valid	T_{cocsv}			7.8	ns
4	COCDIN[n] setup time	T_{cocds}		3.0		ns
5	COCDIN[n] hold time	T_{cocdh}		0.5		ns
6	SCLK to COCDOUT valid	T_{cocdv}			5.0	ns
7	COCM* (input mode) setup time	T_{cocms}		1.5		ns
8	COCM* (input mode) hold time	T_{cocmh}		0.5		ns
9	SCLK to COCM* (output mode) valid	T_{cocmv}			8.0	ns
10	CEMRQ (input mode) setup time	T_{cemss}		1.8		ns
11	CEMRQ (input mode) hold time	T_{cemsh}		0.5		ns
12	SCLK to CEMRQ (output mode) valid	T_{cemsv}			7.7	ns

General Note:

1. The cascade interface timing parameters are specified with reference to SCLK.

2. Number of devices in cascade is limited by the SCLK frequency as follows: 116MHz : 1 device, 100MHz : 2 devices, 77MHz: 3 devices, 66MHz : 4 devices, 50MHz : 8 devices. Note that the number of devices in cascade is also determined by the capacitive loading due to board layout and the presence of other devices on the system bus.

6.2.6 JTAG Interface Timing

Figure 40 JTAG Interface Timing

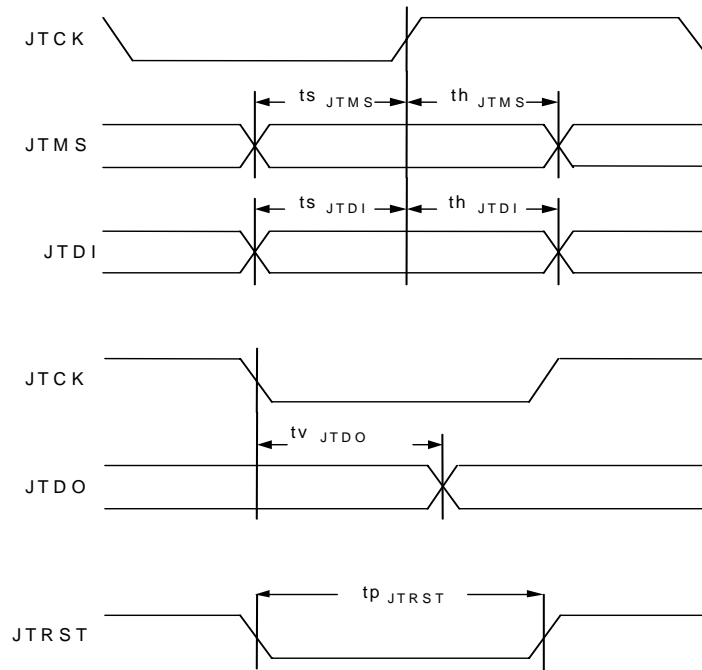
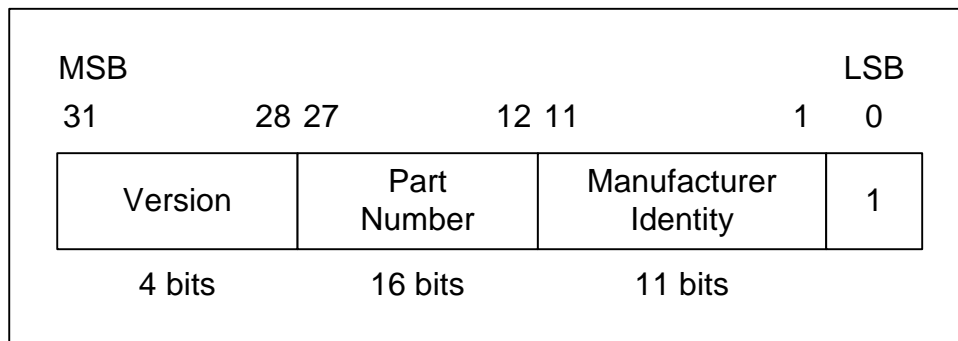


Table 43 JTAG Interface Timing Parameters

Item	Parameter	Symbol	Conditions	Min	Typ	Max	Unit
1	JTCK Frequency					10	MHz
2	JTCK Duty Cycle			40		60	%
3	JTMS Set-up time to JTCK	T_{sJTMS}		4.0			ns
4	JTMS Hold time to JTCK	T_{hJTMS}		1.0			ns
5	JTDI Set-up time to JTCK	T_{sJTDI}		4.0			ns
6	JTDI Hold time to JTCK	T_{hJTDI}		1.0			ns
7	JTCK to JTDO Valid	T_{VJTDO}		2.0		10.0	ns
8	JTRST Pulse Width	T_{pJTRST}		200.0			ns

Figure 41 JTAG IDCode Register



Version	0h
Part Number	0040h
Manufacturer Identity	038h

7 Package Details

7.1 Package Type, Characteristics and Mechanical Drawing

7.1.1 Package Type

352 pin TEBGA package, 35x35 mm (1.378" x 1.378")

7.1.2 Thermal Characteristics

Figure 42 Device Compact Model




Table 44 Thermal Characteristics

Item	Parameter	Symbol	Conditions ^a		Value	Unit
1	Junction to Ambient Thermal Resistance (Theta Ja)	θ_{JA}	SCLK = 100MHz, ACLK = 200MHz, $V_{DD} = 3.3V+5\%$, $CV_{DD} = 1.5V+4\%$	Natural Convection	12.50	°C/W
				Airflow = 1 m/s (200 LFPM)	10.00	°C/W
				Airflow = 2 m/s (400 LFPM)	7.80	°C/W
			SCLK = 116MHz, ACLK = 232MHz, $V_{DD} = 3.3V+5\%$, $CV_{DD} = 1.6V+4\%$	Natural Convection	12.20	°C/W
				Airflow = 1 m/s (200 LFPM)	9.50	°C/W
				Airflow = 2 m/s (400 LFPM)	7.70	°C/W
2	Junction to Top Thermal Resistance (Theta Jt)	θ_{JT}			0.43	°C/W
3	Junction to Bottom Thermal Resistance (Theta Jb)	θ_{JB}	SCLK = 100MHz, ACLK = 200MHz, $V_{DD} = 3.3V+5\%$, $CV_{DD} = 1.5V+4\%$		9.00	°C/W
			SCLK = 116MHz, ACLK = 232MHz, $V_{DD} = 3.3V+5\%$, $CV_{DD} = 1.6V+4\%$		7.20	°C/W
4	Maximum Operating Power Dissipation	P_{op-max}	SCLK = 100MHz, ACLK = 200MHz, $V_{DD} = 3.3V+5\%$, $CV_{DD} = 1.5V+4\%$	SD, EDD load = 15pF EMA, ECD load = 20pF (see Note ^b)	4.0	W
			SCLK = 116MHz, ACLK = 232MHz, $V_{DD} = 3.3V+5\%$, $CV_{DD} = 1.6V+4\%$	SD, EDD load = 15pF EMA, ECD load = 20pF (see Note ^b)	5.5	W
5	Maximum long-term operating junction temperature	T_{j-op}	(see Note ^c)		105	°C


a. Specified parameter values vary depending on the given device operation conditions.

b. The shown maximum operating power dissipation is based on simulated data.

c. The maximum long-term operating junction temperature is to ensure adequate long-term life. The absolute maximum junction temperature of 125 °C is for short-term excursions with guaranteed continued functional performance. Short-term is understood as the definition stated in Bellcore Generic Requirements GR-63-Core.


7.1.3 Package Mechanical Drawing

Figure 43 Package Mechanical Drawing



NOTES:

1. ALL DIMENSION ARE IN MM.
2. TERMINAL POSITIONS DESIGNATION PER JESD 95-1, SPP-010.
3. CORNER DETAILS PER STATS OPTION.
4. PIN 1 IDENTIFIER CAN BE CHAMFER, INK MARK, METALLIZED MARK, BUT LOCATED WITHIN ZONE INDICATED.
5. REFLOW BALL DIAMETER.
6. JEDEC CODE MO-151 VARIATION BAR-2.



DIMENSION	MINIMUM	NOMINAL	MAXIMUM
A	-	-	2.10
A1	0.50	0.60	0.70
A2	0.80	1.10	1.40
b	0.60	0.75	0.90
NUMBER OF BALLS 352			