PM7324

S/UNI-ATLAS

DATA SHEET ERRATA

ISSUE 10: MARCH 2001

PUBLIC REVISION HISTORY

Issue No.	Issue Date	Details of Change
1	Jan 1999	Creation of document
2	Apr 1999	Included new errata items as of Apr. 13, 1999
3	May 1999	Final errata list for Rev A
4	Jun 1999	Aligns with Rev B
5	Sep 1999	Aligns with Rev C
6	Oct 1999	Aligns with Rev C Production Release
7	Nov 1999	Adding min CLK Frequency of 24.5 MHz @ -20°C to 85°C
8	June 2000	Aligns with Rev D. Removed all errata from issue 7. Added new items under sections 2.1-2.6, 3.1-3.2.
9	Mar 2001	Added datasheet and functional discrepancies: 2.7 - Quad UL1 mode, 2.8 – Maximum ISYSCLK Frequency, 2.9 -Maximum ESYSCLK Frequency and 3.3 – Non-Optimal Ingress Throughput Observed With Certain Combinations of Features, 3.4 Non-Optimal Ingress Throughput Observed Due to “Stuttering” Effect
10	Mar 2001	Modified Sections 2.8 and 3.4 from the previous issue of this document. Under section 2.8, the maximum ISYSCLK frequency is increased to 59.5MHz from 59MHz. Under section 3.4, further clarification on the worst-case condition is presented.

CONTENTS

1 ISSUE 10 ERRATA

This document is the errata notice for Revision D of the S/UNI-ATLAS (PM7324-BI) and Issue 7 S/UNI-ATLAS datasheet. The Issue 7 S/UNI-ATLAS datasheet (PMC-1971154) and Issue 10 errata supersede all prior editions and versions of the datasheet.

1.1 Device Identification

The information contains in this document applies to the PM7324 S/UNI-ATLAS Revision D device only. The device revision code is marked at the end of the Wafer Batch Code on the face of the device (as shown in Figure 1). PM7324 S/UNI-ATLAS Revision D is packaged in a 432 pin Super BGA package.

Figure 1: PM7324 S/UNI-ATLAS Branding Format

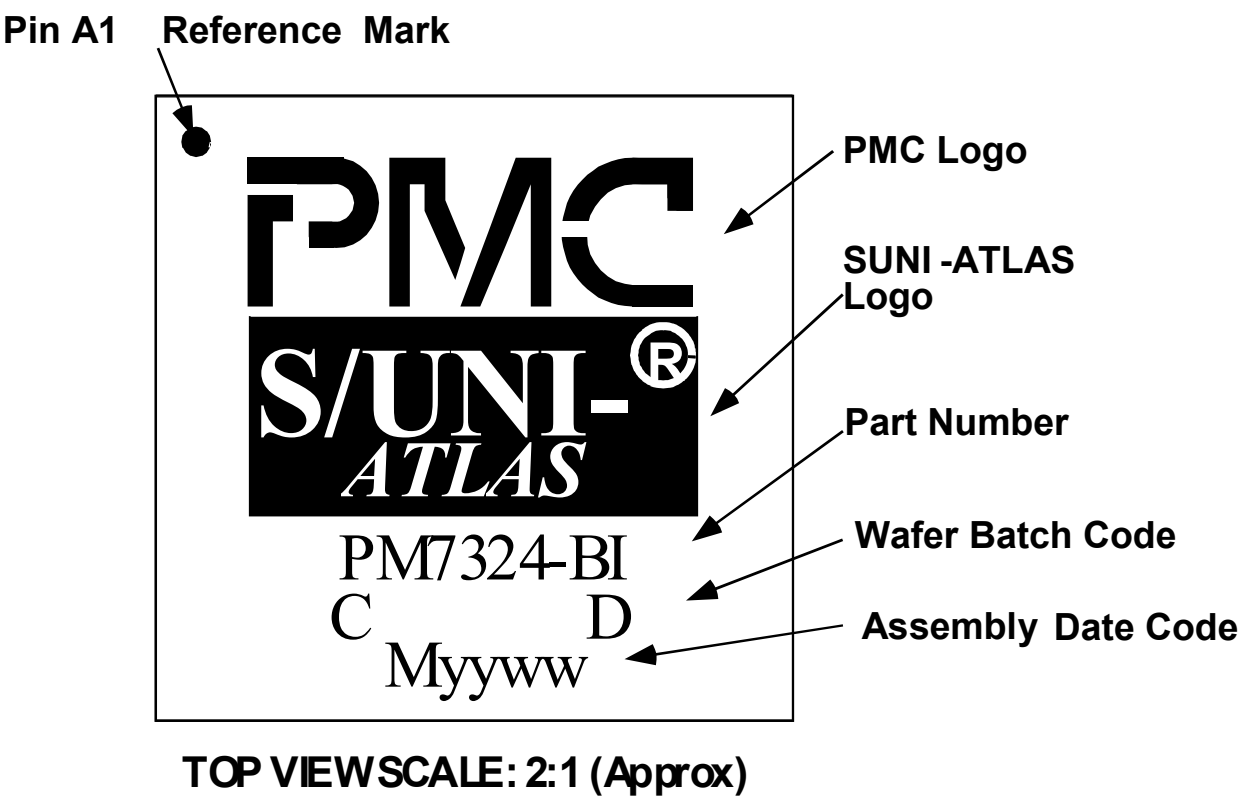

2 DATASHEET DISCREPANCIES

Legend

1. unaltered text is unchanged to add context to changes

2. new material is bold and Italicized

~~3. obsolete material is struck out~~

4. comments specific to this document are in italics

5. A vertical bar in left margin indicates that this is a new item which was not present in the previous issue of this document.

NOTE: All items in Section 2 are documentation changes only.

2.1 CMOS input high DC level

In Table 37 – D.C. Characteristics, the Min value for Input High Voltage (CMOS Only) should read as follows:

Symbol	Parameter	Min	Typ	Max	Units	Conditions
V_{IH}	Input High Voltage (CMOS Only)	~~$0.3V_{dd}$~~ ***$0.7V_{dd}$***			Volts	Guaranteed Input HIGH Voltage Notes 7

2.2 TFIFODP[1:0] additional register bit description

The TFIFODP[1:0] registers bits are found in Register 0x40 – Egress Output Cell Configuration #1. The register bits' description with the additional text is shown below.

TFIFODP[1:0]:

The TFIFODP[1:0] register field determines the apparent cell depth of all FIFOs for the Egress output interface. The FIFO depth control may be important in systems where the cell latency through the Egress portion of the ATLAS. The apparent FIFO cell depth is configured as shown below:

TFIFODP[1]	TFIFODP[0]	APPARENT DEPTH
0	0	4 cells
0	1	3 cells
1	0	2 cells
1	1	1 cell

In single PHY OC-12 mode, TFIFODP[1:0] must be set to 00B to guarantee full OC-12 throughput.

2.3 Clarification in the use of CLRONRD and WM register bits

The CLRONRD and WM register bits are found in Register 0x182 - Ingress VC Table External RAM Access Control and Register 0x184 - Ingress VC Table Write Mask respectively. The register bit description of CLRONRD with the additional text is shown below.

CLRONRD:

If CLRONRD is logic 1, then after a read access of Row 5 of the Ingress VC Table specified in Register 0x181, a write with data bits [47:0] set to all '0' is automatically initiated. ~~Other bits in the words~~ ***Data bits [63:48] and data bits with WM from register 0x184 set to 1*** are preserved in the write back.

If CLRONRD is logic 1, then after a read access of Row 6 and Row 9 of the Ingress VC Table specified in Register 0x181, a write with data bits all '0' is automatically initiated. ***Data bits with WM from register 0x184 set to 1 are preserved in the write back***.

If CLRONRD = '0', no write back to clear the data bits is initiated.

2.4 COSDATA[25:0] Bit Descriptions

The following changes are made to the COSDATA[25:16] bit descriptions in Registers 0x23b:

COSDATA[25:16]:

The COS***DATA***[25:16] field contains the End-point information and Status field of a connection whose address is identified by the COS***DATA***[15:0] register. COSDATA[25] is the Segment End-Point bit (if this bit is logic 1, the connection is configured as a segment end-point), COSDATA[24] is the End-to-End Point bit (if this bit is logic 1, the connection is configured as an end-to-end point), COSDATA[23:16] is the Status field (LSB justified with COSDATA***[23]***~~[23:21]~~ set to logic 0) of the connection. The presence of data in this register indicates that the connection has undergone a change in connection state. ~~This is the most significant 8 bits of the Ingress Change of State FIFO.~~

The following changes are made to the COSDATA[25:16] bit descriptions in Registers 0x2d7:

COSDATA[25:16]:

The COS***DATA***[25:16] field contains the Status field of a connection whose address is identified by the COS***DATA***[15:0] register. COSDATA[25] is the Segment End-Point bit (if this bit is logic 1, the connection is configured as a segment end-point), COSDATA***[24]***~~[25]~~ is the End-to-End point bit (if this bit is logic 1, the connection is configured as an end-to-end point), COSDATA[23:16] is the Status field (LSB justified with COSDATA***[23:22]***~~[23:15]~~ set to logic 0) of the connection. The presence of data in this register indicates that the connection has undergone a change in connection state. ~~This is the most significant 8 bits of the Egress Change of State FIFO.~~

2.5 ITU-I.610 Recommendation Version

The reference to ITU-T Recommendation I.610 in Section 3 of the datasheet is changed to reflect the most recent version as shown below:

- ITU-T Recommendation I.610 – "B-ISDN Operation and Maintenance Principles and Functions", ***February, 1999*** ~~June, 1997 (Rapporteur's edition).~~

2.6 Device ID Revision Number

The device ID revision number is changed from 0010 (Rev C) to 0011 (Rev D) in Register 0x000: S/UNI-ATLAS Master Reset and Identity / Load Performance Meters. The following note is added to the ID[3:0] bit description:

ID[3:0]:
The ID bits can be read to provide a binary number indicating the S/UNI-ATLAS feature version.

NOTE 1:
Rev A ID[3:0] = 0000
Rev B ID[3:0] = 0001
Rev C ID[3:0] = 0010
Rev D ID[3:0] = 0011

2.7 Quad UL1 Mode

The RPOLL, IPOLL and TPOLL pin description notes found in section 7 of the datasheet states that "the 4-PHY configuration is not recommended." This is changed to "the 4-PHY configuration should not be used" as shown below:

Note: In direct addressing mode, the 4-PHY configuration ~~is not recommended~~ **should not be used**. Instead the 4-PHY address-polling mode should be used. This does not apply to the Single or Dual-PHY configurations.

2.8 Maximum ISYSCLK Frequency

In Table 44 – Ingress SRAM Interface, the Max value for ISYSCLK frequency (including the reference to note 1) are changed as follows:

Symbol	Description	Min	Max	Units	Notes
	ISYSCLK Frequency	24.5	~~52~~ ***59.5***	MHz	1
	ISYSCLK Duty Cycle	40	60	%	

Notes:

1. Over the ISYSCLK frequency range 25 MHz to ~~52~~ ***59.5*** Mhz T_j = -40°C to +120°C. Over the ISYSCLK frequency range 24.5 to 25 MHz T_j = -20°C to +120°C.

2.9 Maximum ESYSCLK Frequency

In Table 45 – Egress SRAM Interface, the Max value for ESYSCLK frequency (including the reference to note 1) are changed as follows:

Symbol	Description	Min	Max	Units	Notes
	ESYSCLK Frequency	24.5	~~52~~ ***57.0***	MHz	1
	ESYSCLK Duty Cycle	40	60	%	

Notes:

1. Over the ESYSCLK frequency range 25 MHz to ~~52~~ ***57.0*** Mhz T_j = -40°C to +120°C. Over the ESYSCLK frequency range 24.5 to 25 MHz T_j = -20°C to +120°C.

3 FUNCTIONAL DISCREPANCIES

3.1 Egress Header Translation Not Performed under Certain Conditions

3.1.1 Description

In the Egress direction of the ATLAS, the Egress Cell Processor (ECP) does not perform header translation on the incoming egress traffic when cells are also being inserted to the ECP from the microprocessor **without** header translation (i.e. E_UPHDRX bit from Register 0x061 is set to 0). This phenomenon has been observed at microprocessor cell insertion rates of 1 cell every 20ms or higher.

Note: Egress header translation is performed properly on the incoming egress traffic when cells that are inserted to the ECP from the microprocessor are header translated.

3.1.2 Software Workaround

The software workaround is implemented by creating a dummy Virtual Connection (VC) with header translation enabled. Before each microprocessor cell insertion, copy the inserting cell's VPI/VCI address to the VPI/VCI address field in the dummy VC Table. Subsequently, copy the dummy VC index into the inserting cell. Finally, write the cell into the Egress insert FIFO with header translation enabled.

The following steps describe the above procedure in more detail:

1. Select an Egress Search Key (SRAM address) to be assigned to the dummy VC.
2. Create and Initialize a dummy Egress VC Table at the SRAM address specified by the Egress Search Key from step 1.
3. Set EAD[15:0] from Register 0x2AB to the Egress Search Key (SRAM address) of the chosen dummy VC.
4. Set ROW[0] from Register 0x2AD to 1 to allow write access to row 0 of the Egress VC Table.
5. Set the VPI/VCI address field located at ESA[19:16] = 0x0000 of the dummy VC Table by configuring ROW0[15:0] from Register 0x2AE and ROW0[31:16] from Register 0X2AF to the VPI/VCI address of the inserting cell.
6. Program the cell to be inserted so that its Egress Search Key will resolve to the SRAM address of the dummy VC
7. Set the E_UPHDRX bit from Register 0x061 to 1 to enable header translation.
8. Insert the cell to the Egress insert FIFO through Register 0x062.
9. Repeat Steps 3 to 8 for each microprocessor cell insertion.

Note: Only one dummy VC is needed for all connections.

3.1.3 Performance with Workaround

The implementation of the software workaround requires the user to dedicate only one Virtual Connection for all microprocessor cell insertions.

Also, the active bit for this connection will be set even though it is not really an active connection. This may present a problem in cases where the system needs to go through the records in sequence and then verify if the active bit is set before doing any action on that particular connection.

3.2 Reserved Bit in the Status Field of the Egress VC Table Erroneously Set when COS FIFO is Full

3.2.1 Description

When the COS FIFO is full and a change of state occurs for a VC, the Egress Cell Processor (ECP) is supposed to set the Reserved bit in the Status Field of the Egress VC table (bit 31 of ESA[19:16] = 0x0010) to 1. This indicates that once there is space in the COS FIFO, the status of this VC must be copied to the COS FIFO.

Instead, at the instance the COS FIFO becomes full, the ECP sets the Reserved bits of all connections that receive cells regardless of the change of state condition. This situation only occurs for connections with the Active and COS_Enable bits set to 1. As the Reserved bits are set erroneously, the COS FIFO will contain vectors to a mix of connections that have no change of state and those that have a change of state.

This phenomenon can be detected when an interrupt generated by the setting of the E_COSFULLI bit in Register 0x006 is received.

3.2.2 Software Workaround

When the COS FIFO is full, E_COSFULLI in Register 0x006 will be set to 1 and will generate an interrupt. The software should poll the interrupt and then implement the following routine:

1. Clear the COS bit in Register 0x280. (This will disable the COS FIFO)
2. Empty the COS FIFO by reading Registers 0x2D6 and 0x2D7.
3. Initiate a wait state routine set to 1.0s. (This will ensure that the background processor has cleared all the Reserved bits in the Status field of the Egress VC Table)

4. Set COS bit in Register 0x280. (This will enable the COS FIFO)
5. Read the Status field from ESA[19:16] = 0010 for all connections in the Egress VC Table.
6. Check the value of E_COSFULLI in Register 0x006. If E_COSFULLI = 1, then go back to step 1, else go to step 7.
7. Empty the COS FIFO. At this point, the size of the COS FIFO should be less than 256.

Note: Step 5 can be modified or ignored if the microprocessor has limited allocated time to read all the connections.

3.2.3 Performance with Workaround

When the COS FIFO is not full, the performance is unaffected. When the COS FIFO overflows, it takes greater than 1s for the COS FIFO to stabilize. Furthermore, the entire VC context table must be read to ensure that the current state of the connections is in the microprocessor memory.

Note: Software convergence rates will vary according to the microprocessor speed and the number of active connections.

3.3 Non-Optimal Ingress Throughput Observed With Certain Combinations of Features

3.3.1 Description

The two major parameters that determine the ingress throughput performance of the device are the ISYSCLK (Ingress System Clock) frequency and the average number of clock cycles spent to process an ATM cell in each cell period.

For a 1xOC-12 or 4xOC-3 system, the maximum cell throughput is 1.413×10^6 cells/s. Therefore, in order to process the maximum cell throughput at an ISYSCLK frequency of 50MHz (most OC-12 systems run the ISYSCLK at this rate), the maximum number of clock cycles spent per cell period must be:

$$\begin{aligned}\text{Max Number of Clock Cycles} &= \text{ISYSCLK frequency} / \text{Max cell throughput} \\ &= 50\text{MHz} / 1.413 \times 10^6 \text{ cells/s} \\ &= 35.39 \text{ clock cycles} \\ &= 35 \text{ clock cycles}\end{aligned}$$

Under certain conditions, which will be described below, the average number of clock cycles spent in each cell period exceeds 35 clock cycles.

As a rule of thumb, the average number of clock cycles spent in each cell period is dependent on the rate at which the microprocessor accesses the SRAM, the

binary search tree depth and the amount of cell processing required by each cell. These factors are now described in further detail.

MICROPROCESSOR SRAM ACCESS

The device permits one microprocessor SRAM access per cell period. In the worst-case, the microprocessor may request a clear-on-read or a byte-masked write access, which takes 4 clock cycles to complete.

CPU Access = 4 clock cycles

By comparison, read-only or write-only operations take only 1 clock cycle.

BINARY SEARCH ALGORITHM

The searching algorithm consists of the following steps:

1. Primary Lookup
2. Secondary Branch Lookups, and
3. Verification

As specified in section 8.2.1 of the datasheet, the maximum number of secondary lookups that can be performed to guarantee full throughput is 16. Therefore, the maximum possible number of clock cycles used up for the search is:

Search = 1 (Primary Lookup) + 16 (Secondary Branches) + 1 (Verification)
= 18 clock cycles

CELL PROCESSING

The number of clock cycles used by the cell processor is highly dependent upon the type of cell received and upon the configuration of the device. Accordingly, the cell processor performs the necessary reads and writes to the appropriate rows in the Ingress VC Table outlined in Table 1 of the datasheet. Three cases are presented below.

Case A: User Cells with a Cell Length = 27 or 28 Words

For the case of a user cell, in which policing, PM and header translation are enabled, and the user cell is part of an F4 termination, the cell processing is as follows (assume that the output cell length is either 27 or 28 words):

1. Read Row2 (status, configuration, VPC pointer)
2. Read Row4 (policing parameters)
3. Read Row3 (policing parameters)
4. Read Row5 (policing parameters)
5. Read F4 Row2 (assumes user cell is part of a VPC flow)
6. Read Row7 (translated header, UDF, prepend, postpend)
7. Read Row6 (ingress cell count 1, 2)
8. Bus Turnaround
9. Bus Turnaround
10. Write Row 6* (ingress cell count 1, 2)
11. Write Row 3 (update TAT1, TAT2)
12. Write Row 5* (update non-compliant cell count or GFR state)
13. Write Row 2 (update status)

*Not always required.

Cell Processing = 13 ISYSCLK cycles

In the example above, if the associated F4 connection was in AIS alarm or CC alarm (or in any case once per half-second per F4), an additional write would be performed to Row 2 of that F4 Connection, and thus, the total number of processing cycles would be 14.

Therefore, the worst-case total number of clock cycles required for all processing (searching, microprocessor access and cell processing) in this scenario is:

Total cycles = CPU Access + Worst-Case Search + Cell Processing
= 4 + 18 + 14
= 36 cycles

NOTE: Segment and End-to-End F4 Must be Set Up as a Single Connection

In some applications, segment and end-to-end F4 OAM connections are configured as two separate connections. This should not be done, as it will cause one additional read to Row 2 of the second F4 OAM connection per cell access, thereby increasing the maximum processing cycle by 1.

In the case where the additional F4 OAM connection is in AIS or CC alarm, or if the background process has just decremented its CC alarm counters (which occurs once per half second), an additional write to Row 2 of that OAM connection may also occur. It is possible that this additional write could coincide with a write to the first F4 OAM connection.

Instead, **a single connection should be used for both Segment and End-To-End F4 OAM, and cells from both VCI= 0003 and VCI = 0004 should resolve to this connection.** This combined OAM connection handles both the End-to-End and Segment OAM simultaneously in much the same way that any other connection does. This avoids the additional read or write operations to Row 2 of the second F4 OAM connection. This is shown in the figure below

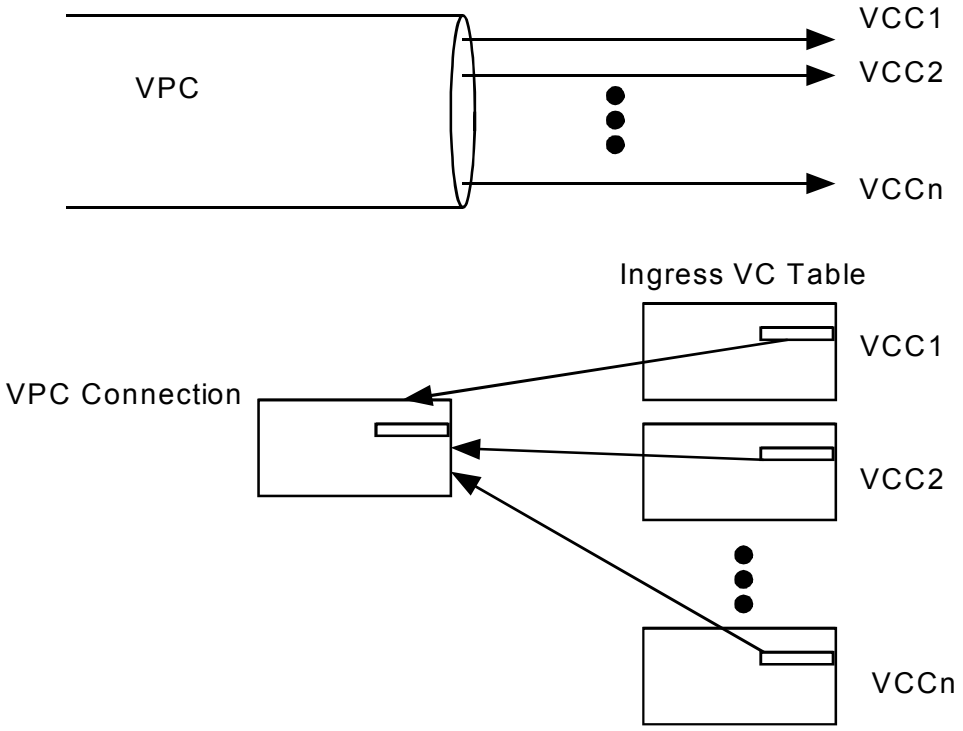

Case B: User Cells with a Cell Length > 28 words

For output cell lengths greater than 28 words (GPREPO bit in Register 0x200 = 1), an additional read from Row 8 is performed thereby increasing the maximum cell processing cycle by 1.

Case C: AIS Cells

For AIS cell processing with IAISCOPY in Register 0x238 set to 1, an additional read to Row A, and writes to Row A, B (or D) and C (or E) are performed thereby increasing the maximum cell processing cycle by 4. Please note, however, that these read and writes are only performed if IAISCOPY = 1. If all features are enabled for these AIS cells as outlined above, then the maximum worst-case cell processing cycles required is 19 ISYSCLK cycles.

Therefore, the worst-case total number of clock cycles required for all processing in this scenario is:

Total cycles = CPU Access + Worst-Case Search + Cell Processing
= 4 + 18 + 19
= 41 cycles

This is the worst-case average number of clock cycles that can be encountered in a given system.

3.3.2 Workaround

In general, the workaround involves either reducing the average number of clock cycles spent per cell period or increasing the ISYSCLK frequency. Most system configurations are different and therefore will require a different workaround. Examples for common configurations are provided for each workaround.

Pace Microprocessor SRAM Access

The microprocessor SRAM access can be paced appropriately to ensure that the maximum average number of clock cycles required to guarantee full throughput is not exceeded.

As an example, consider the case where GPREPO = 0 and IAISCOPY = 0, this results in a maximum average number of clock cycle per cell period of 36. With the clock at 50MHz, the maximum average number of clock cycle per cell period must not exceed 35.39. By pacing the microprocessor accesses to 1 every 2 cell period at most, the average number of clock cycles is reduced to 34. This guarantees that full throughput is achieved.

Total cycles = Ave CPU Access + Search + Maximum Cell Processing
= 2 + 18 + 14
= 34 cycles

Avoid the use of Clear-On-Read or Masked-Write operations

If microprocessor accesses are limited to simple read and simple write operations, then each access requires only 1 cycle.

As an example, consider the case where GPREPO = 0 and IAISCOPY = 0, this results in a maximum average number of clock cycle per cell period of 36. With the clock at 50MHz, the maximum average number of clock cycle per cell period must not exceed 35.39. By eliminating the use of Clear-On-Read or Masked-Write operations, the average number of clock cycles is reduced to 33. This guarantees that full throughput is achieved.

Total cycles = 1 read or write + Search + Maximum Cell Processing
= 1 + 18 + 14
= 33 cycles

Reduce Search Tree Depth

The search tree depth can be reduced appropriately to ensure that the maximum average number of clock cycles required to guarantee full throughput is not exceeded.

Using the same example as above, where GPREPO = 0 and IAISCOPY = 0, the search tree depth can be reduced to 15 (instead of 16) to get a maximum average number of clock cycles of 35. This guarantees that full throughput is achieved.

Total cycles = CPU Access + Search + Maximum Cell Processing
= 4 + 17 + 14
= 35 cycles

Increase ISYSCLK Frequency

The ISYSCLK frequency can be increased appropriately to ensure that the device can process the worst-case average number of clock cycles per cell period while guaranteeing full throughput.

Using the same example as above, where GPREPO = 0 and IAISCOPY = 0, the maximum average number of clock cycles is 36. In order to ensure full throughput the ISYSCLK frequency must be run at:

ISYSCLK frequency = 36 cycles x 1.413×10^6 cells/s
= 50.87 MHz

In this case, a clock of 52 MHz should be used, to permit the use of the internally-generated half-second clock.

In the absolute worst-case scenario as in case C (section 3.3.1) above, the maximum average number of clock cycles is 41. In order to ensure full throughput the ISYSCLK frequency must be run at:

ISYSCLK frequency = 41 cycles x 1.413×10^6 cells/s
= 57.94 MHz = 58 MHz

If a 58 MHz clock is used, then the internally-generated half-second clock (which supports ISYSCLK frequencies of 25, 50, or 52 MHz) will be slightly fast. As a result, an external half-second clock must be provided on the HALFSECCLK pin, to provide accurate timing to the background processes that implement the CC, AIS, and RDI functions.

3.3.3 Performance with Workaround

The device is able to process a full OC-12 worth of traffic after the workaround is implemented.

3.4 Non-Optimal Ingress Throughput Observed Due to "Stuttering" Effect

3.4.1 Description

The Ingress SRAM bus is shared between two main blocks, the Ingress Search Engine (ISE) and Ingress Cell Processor (ICP). Under certain conditions, the way these blocks share the bus can enter an inefficient state, resulting in an inability to process cells at full throughput. In this inefficient state, the bandwidth is limited by the depth of the search tree and the length of the output cell in the ingress direction.

The ISE performs searching on 3 cells in parallel. As soon as the next cell finishes its search, and is ready for transfer into the ICP, the SRAM bus is handed over. When the ICP is done processing cells, the SRAM bus is returned to the control of ISE for searching. Thus, searching is interleaved between each cell-processing step. Meanwhile, as soon as policing is complete, and the data for header translation has been read from the SRAM, the ICP begins streaming the cell from the ISE into the OCIF, modifying it as needed. This cell streaming continues in parallel with searching once the bus is turned back over to the ISE. The diagram below shows a 32-word cell case where the interleaving is occurring correctly.

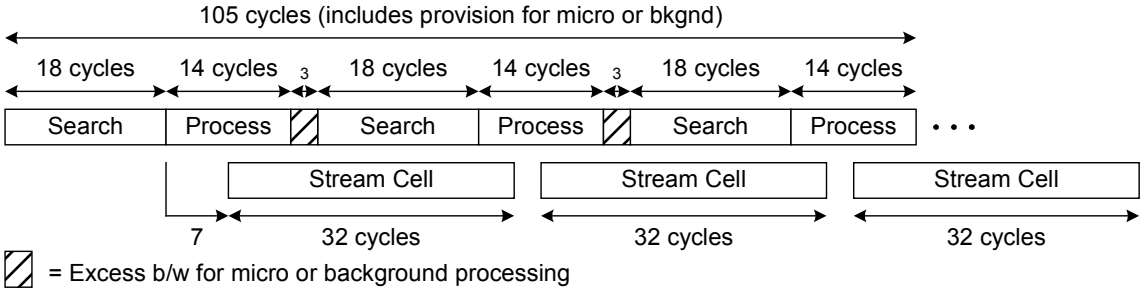

Under certain conditions, the device can enter a state where all the searching for three cells is done at once, followed by cell processing of all three cells, without interleaving. Under this condition, the cell processing is throttled by the length of time it takes to stream the previous cell. No searching is done while the cell processor is waiting for the previous cell to stream. The diagram below illustrates this condition.

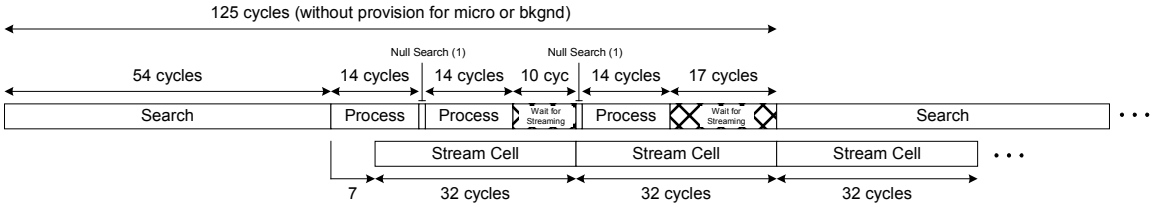

The process of getting into the stuttering state is probabilistic. The mechanism for entering this state is a combination of disturbing factors such as uneven search-tree depths, reception of unprovisioned cells, background processing, and other events. While it is unlikely that the worst version of this state will be entered, it cannot be ruled out. Once in this state, the throughput equation is as follows:

Time for 3 cells = 2(Out_Cell_Len) + 3(Sec_Search_Depth + 2) + Fixed_Latency

Out_Cell_Len is the length of cells sent through the Ingress OCIF
Sec_Search_Depth is the worst-case depth of the secondary search
Fixed_Latency is 6 for cell lengths of 27 or 28, and 7 for greater cell lengths.

In the worst-case scenario, the microprocessor accesses can be additive to the total clock cycles consumed per 3 cell period. This occurs when the microprocessor access to SRAM is requested during the cell processing time before the long search in the stuttering case. This occurs less than 1/3 of the time, but could happen many times in succession resulting in cell loss. For a clear-on-read or a byte-masked write operation an additional of 4 clock cycles is added to the throughput equation as follows:

Time for 3 cells = 2(Out_Cell_Len) + 3(Sec_Search_Depth + 2) + Fixed_Latency + 4

For a 1xOC-12 or 4xOC-3 system, the maximum cell throughput is 1.413×10^6 cells/s. Therefore the frequency at which ISYSCLK must be run is:

Min ISYSCLK Freq = (1.413×10^6) (Time for 3 cells)/3.

As long as the throughput is sufficient to temporarily handle the worst-case stuttering condition, the Input Cell Interface will be emptied, and the device will exit the stuttering condition.

NOTE: Background processes have no impact, as they are automatically paused while bandwidth needs are high, and resumed once the inefficiency has passed.

3.4.2 Workaround

In general, the workaround involves a combination of reducing the maximum number of clock cycles spent during the worst-case stuttering condition, or increasing the ISYSCLK frequency. Most system configurations are different and therefore will require a different workaround. Examples for common configurations are provided for each workaround.

Decrease the Output Cell Length

Reducing the output cell length will reduce the length of time spent waiting for cells to stream into the OCIF. As per the equation, 2 cycles are saved from the 3-cell repeating pattern for every reduction of 1 word in the cell length.

The following are some examples of the required throughput:

Output Cell Length 30, Search Depth 16:
(30x2 + 18x3 + 7 + 4)(1/3)(1.413) = 58.88 MHz

Output Cell Length 29, Search Depth 16:
(29x2 + 18x3 + 7 + 4)(1/3)(1.413) = 57.93 MHz

Output Cell Length 28, Search Depth 16:
(28x2 + 18x3 + 6 + 4)(1/3)(1.413) = 56.52 MHz

Output Cell Length 27, Search Depth 16:
(27x2 + 18x3 + 6 + 4)(1/3)(1.413) = 55.58 MHz

Reduce Search Tree Depth

Reducing the search tree depth reduces the time to search each cell. As per the equation, 3 cycles are saved from the 3-cell repeating pattern for every reduction by 1 in the search tree.

Output Cell Length 28, Search Depth 12:
(28x2 + 14x3 + 6 + 4)(1/3)(1.413) = 50.87 MHz

Output Cell Length 30, Search Depth 11:
(30x2 + 13x3 + 7 + 4)(1/3)(1.413) = 51.81 MHz

Output Cell Length 32, Search Depth 9:
(32x2 + 11x3 + 7 + 4)(1/3)(1.413) = 50.87 MHz

Reducing the depth of the secondary search can often be achieved by more fully utilizing the primary search. For instance, a 4-PHY UNI application might concatenate the 2-bit PHYID with the 8-bit UNI VPI and the first 6 bits of the VCI into a 16-bit index to the primary search table. Most of the entries in that table would not have any associated connections, and thus would be coded to zero. Those entries in the table that did have associated connections would point to the base of secondary search trees. Because only 10 bits of the VCI remain to be resolved, the maximum search depth in the secondary tree will be 10. Thus, the total search will consume 12 cycles (1 primary + 10 secondary + 1 confirmation).

Further reductions in the VPI or VCI range allowable permit further reductions in the secondary search depth.

Increase ISYSCLK Frequency

To the extent that a combination of search tree reduction and cell length reduction is insufficient to bring the required frequency below 50 or 52 MHz, ISYSCLK may be increased up to a maximum of 59.5 MHz.
However:

- The ESYSCLK must not be raised above 57 MHz, and
- If a frequency other than 50 MHz or 52 MHz is used, then the HALFSECCLK input pin must be driven with an external half-second clock reference

Pace Microprocessor SRAM Access

The microprocessor SRAM access can be paced appropriately to ensure that the maximum average number of clock cycles required to guarantee full throughput is not exceeded. In the worst-case scenario, where the output cell length is 32 words and the search tree depth is 16, it is suggested that the microprocessor accesses to the SRAM be paced to 1 every 48 cell periods at most in order to guarantee full throughput at a maximum ISYSCLK rate of 59.5MHz.

In the worst-case scenario, the time for 3 cells (excluding microprocessor access) is 2x32 + 3x18 + 7 = 125 cycles. At 59.5MHz, this translates to 2101ns. Hence, at a rate of 1 access per 48-cell period (or 16 x 3 cell periods), the time interval between each access must be 16*2101ns = 33.614us.

Therefore, for every 3 cell period, an additional of 1/16*4 = 0.25 cycles must be added to the total number of clock cycles. At this rate, the maximum throughput will be 125.25/3 * 1.413MHz = 58.993MHz.

If the microprocessor access is not paced as described above, then either the output cell length or search tree depth must be reduced in order to guarantee full throughput at 59.5MHz.

At an output cell length of 32 words, the search tree depth must be reduced to 15 as shown in the calculation below:

(3x32 + 2x(15+2) + 7 +4)(1/3)(1.413) = 59.35MHz

At a search tree depth of 16, the output cell length must be reduced to 30 as shown in the calculation below:

(3x30 + 2x(16+2) + 7 +4)(1/3)(1.413) = 58.88MHz

This guarantees that full throughput is achieved.

Avoid the use of Clear-On-Read or Masked-Write operations

If microprocessor accesses are limited to simple read and simple write operations, then each access requires only 1 cycle.

As an example, consider the case of an output length of 32 with a secondary search depth of 16, and constant clear-on-read or masked-write operations that all fall in the worst-case position so that they affect throughput. In this case, microprocessor accesses must be paced to no more than 1 in 16x3 = 1 in 48 cell periods.

(3x32 + 2x(16+2) + 7 + 4/16)(1/3)(1.413) = 58.993 MHz

By comparison, if only simple reads and writes were involved, the pacing can be reduced to 1 in 4x3 = 1 in 12 cell periods.

(3x32 + 2x(16+2) + 7 + 1/4)(1/3)(1.413) = 58.993 MHz

3.4.3 Performance with Workaround

The device is able to process a full OC-12 worth of traffic so long as the

ISYSCLK Freq > (1.413 x10^6) (Time for 3 cells)/3,

where

Time for 3 cells = 2(Out_Cell_Len) + 3(Sec_Search_Depth + 2) + Fixed_Latency + Provision for Microprocessor Access

NOTES

CONTACTING PMC-SIERRA, INC.

PMC-Sierra, Inc.
105-8555 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
Application Information: apps@pmc-sierra.com
Web Site: http://www.pmc-sierra.com

PMC-1981505(R10) ref PMC-1971154(R7) Issue date: Mar 2001