

PM7326



S/UNI APEX

ATM/PACKET TRAFFIC MANAGER AND SWITCH

DATA SHEET

ISSUE 6: APRIL 2000

REVISION HISTORY

Issue No.	Issue Date	Details of Change
Issue 6	April, 2000	Final update for production release. Changes from Issue 5 marked with change bars
Issue 5	December, 1999	Removed Applications Examples and Operations sections, replaced with dedicated documents. IDDOP (operating current) value inserted.
Issue 4	August, 1999	Datasheet re-written to incorporate extensive updates and clarifications.
Issue 3	June, 1999	No material change from Issue 2, formatted for web site.
Issue 2	February, 1999	Updates and clarifications throughout
Issue 1	November, 1998	Document created.

CONTENTS

1 DEFINITIONS 1

2 FEATURES 3

3 APPLICATIONS 7

4 REFERENCES..... 8

5 APPLICATION EXAMPLES 9

6 BLOCK DIAGRAM 10

7 DESCRIPTION 12

8 PIN DIAGRAM 16

9 PIN DESCRIPTION..... 17

9.1 LOOP ANY-PHY RECEIVE MASTER/TRANSMIT SLAVE
INTERFACE (28 SIGNALS) 17

9.2 LOOP ANY-PHY TRANSMIT MASTER/RECEIVE SLAVE
INTERFACE (34 SIGNALS) 22

9.3 WAN ANY-PHY RECEIVE MASTER/TRANSMIT SLAVE
INTERFACE (25 SIGNALS) 26

9.4 WAN ANY-PHY TRANSMIT MASTER/RECEIVE SLAVE
INTERFACE (25 SIGNALS) 31

9.5 CONTEXT MEMORY SYNCHRONOUS SSRAM INTERFACE (60
SIGNALS)..... 36

9.6 CELL BUFFER SDRAM INTERFACE (52 SIGNALS) 38

9.7 MICROPROCESSOR INTERFACE (44 SIGNALS)..... 40

9.8 GENERAL (9 SIGNALS) 44

9.9 JTAG & SCAN INTERFACE (7 SIGNALS)..... 45

9.10 POWER..... 46

10 FUNCTIONAL DESCRIPTION..... 48

10.1	ANY-PHY INTERFACES	48
	10.1.1 RECEIVE INTERFACE	48
	10.1.2 TRANSMIT INTERFACE	50
10.2	LOOP PORT SCHEDULER	53
10.3	WAN PORT SCHEDULER	54
10.4	WAN PORT ALIASING	56
10.5	WAN AND LOOP ICI SELECTION	57
10.6	MICROPROCESSOR INTERFACE	57
10.7	MEMORY PORT	61
10.8	SAR ASSIST	62
	10.8.1 TRANSMIT	62
	10.8.2 RECEIVE	63
10.9	QUEUE ENGINE	64
	10.9.1 SERVICE ARBITRATION	65
	10.9.2 CELL QUEUING	66
	10.9.3 CLASS SCHEDULING	73
	10.9.4 CONGESTION CONTROL	75
	10.9.5 STATISTICS	82
	10.9.6 MICROPROCESSOR QUEUE BUFFER RE- ALLOCATION/TEAR DOWN	84
10.10	CONTEXT MEMORY SSRAM INTERFACE	84
10.11	CELL BUFFER SDRAM INTERFACE	89
10.12	JTAG TEST ACCESS PORT	93
11	PERFORMANCE	94
	11.1 THROUGHPUT	94

11.2	LATENCY	96
11.3	CDV.....	96
12	REGISTER.....	97
12.1	GENERAL CONFIGURATION AND STATUS.....	98
12.2	LOOP CELL INTERFACE	107
12.3	WAN CELL INTERFACE.....	113
12.4	MEMORY PORT	119
12.5	SAR.....	125
	12.5.1 RECEIVE.....	125
	12.5.2 TRANSMIT	127
	12.5.3 CELL BUFFER DIAGNOSTIC ACCESS.....	128
12.6	QUEUE ENGINE.....	129
12.7	MEMORY INTERFACE	144
12.8	TEST INTERFACE.....	145
12.9	CBI INTERFACE	149
13	CBI REGISTER PORT MAPPING	151
14	MEMORY PORT MAPPING.....	157
14.1	CONTEXT SIZE AND LOCATION.....	157
14.2	QUEUE CONTEXT DEFINITION	160
	14.2.1 VC CONTEXT RECORDS.....	161
	14.2.2 PORT CONTEXT RECORDS.....	169
	14.2.3 CLASS CONTEXT RECORDS.....	173
	14.2.4 SHAPING CONTEXT RECORDS.....	178
	14.2.5 CELL CONTEXT RECORD	180

14.2.6	MISC CONTEXT	180
14.3	WAN PORT SCHEDULER CONTEXT	184
14.3.1	WAN TRANSMIT PORT POLLING WEIGHT RECORD	184
14.3.2	WAN TRANSMIT CLASS STATUS RECORD	185
14.4	LOOP PORT SCHEDULER CONTEXT	186
14.4.1	LOOP TRANSMIT PORT POLLING SEQUENCE RECORD	186
14.4.2	LOOP TRANSMIT PORT POLLING WEIGHT RECORD	187
14.4.3	LOOP TRANSMIT CLASS STATUS RECORD	188
15	TEST FEATURES DESCRIPTION	190
15.1	JTAG TEST PORT	190
16	OPERATION	194
17	FUNCTIONAL TIMING	195
17.1	MICROPROCESSOR INTERFACE	195
17.2	SDRAM INTERFACE	197
17.3	ZBT SSRAM INTERFACE	199
17.4	LATE WRITE SSRAM INTERFACE	200
17.5	ANY-PHY/UTOPIA INTERFACES	201
17.5.1	RECEIVE MASTER/TRANSMIT SLAVE INTERFACES	201
17.5.2	TRANSMIT MASTER/RECEIVE SLAVE INTERFACES	204
18	ABSOLUTE MAXIMUM RATINGS	209
19	D.C. CHARACTERISTICS	210
20	A.C. TIMING CHARACTERISTICS	212
20.1	JTAG INTERFACE	217
21	ORDERING AND THERMAL INFORMATION	219

22	MECHANICAL INFORMATION	220
----	------------------------------	-----

LIST OF REGISTERS

REGISTER 0X00: RESET AND IDENTITY 98

REGISTER 0X10: HI PRIORITY INTERRUPT STATUS REGISTER 99

REGISTER 0X14: HIGH PRIORITY INTERRUPT MASK..... 101

REGISTER 0X18: LOW PRIORITY INTERRUPT ERROR REGISTER 102

REGISTER 0X1C: LOW PRIORITY INTERRUPT ERROR MASK..... 104

REGISTER 0X20: LOW PRIORITY INTERRUPT STATUS REGISTER 105

REGISTER 0X24: LOW PRIORITY INTERRUPT STATUS MASK..... 106

REGISTER 0X100: LOOP CELL RX INTERFACE CONFIGURATION..... 107

REGISTER 0X104: LOOP CELL TX INTERFACE CONFIGURATION..... 110

REGISTER 0X200: WAN CELL RX INTERFACE CONFIGURATION 113

REGISTER 0X204: WAN CELL TX INTERFACE CONFIGURATION 116

REGISTER 0X300: MEMORY PORT CONTROL..... 119

REGISTER 0X340-0X34C: MEMORY WRITE DATA (BURSTABLE)..... 121

REGISTER 0X350: MEMORY WRITE DATA OVERFLOW (BURSTABLE) ... 122

REGISTER 0X380-0X38C: MEMORY READ DATA (BURSTABLE)..... 123

REGISTER 0X390: MEMORY READ DATA OVERFLOW (BURSTABLE) 124

REGISTER 0X400-0X43C: SAR RECEIVE DATA (BURSTABLE)..... 125

REGISTER 0X500-0X53C: SAR TRANSMIT DATA, CLASS 0 (BURSTABLE)127

REGISTER 0X540-0X57C: SAR TRANSMIT DATA, CLASS 1 (BURSTABLE)127

REGISTER 0X580-0X5BC: SAR TRANSMIT DATA, CLASS 2 (BURSTABLE)127

REGISTER 0X5C0-0X5FC: SAR TRANSMIT DATA, CLASS 3 (BURSTABLE)127

REGISTER 0X600: CELL BUFFER DIAGNOSTIC CONTROL 128

REGISTER 0X700: QUEUE CONTEXT CONFIGURATION	129
REGISTER 0X704: RECEIVE AND TRANSMIT CONTROL	132
REGISTER 0X710: MAX DIRECTION CONGESTION THRESHOLDS	134
REGISTER 0X714: CLP0 DIRECTION CONGESTION THRESHOLDS	135
REGISTER 0X718: CLP1 DIRECTION CONGESTION THRESHOLDS	136
REGISTER 0X71C: RE-ASSEMBLY MAXIMUM LENGTH.....	137
REGISTER 0X720: WATCH DOG ICI PATROL RANGE	138
REGISTER 0X724: TEAR DOWN QUEUE ID	139
REGISTER 0X728: WATCH DOG / TEAR DOWN STATUS	140
REGISTER 0X730: SHAPER 0 CONFIGURATION (N = 0).....	141
REGISTER 0X734: SHAPER 1 CONFIGURATION (N = 1).....	141
REGISTER 0X738: SHAPER 2 CONFIGURATION (N = 2).....	141
REGISTER 0X73C: SHAPER 3 CONFIGURATION (N = 3).....	141
REGISTER 0X800: SDRAM/SSRAM CONFIGURATION.....	144
REGISTER 0X900: BIST_OK.....	145
REGISTER 0X904: BIST_MODE	146
REGISTER 0X908: BIST_RESULTS A.....	147
REGISTER 0X90C: BIST_RESULTS B.....	148
REGISTER 0XA00: CBI REGISTER PORT	149
CBI REGISTER 0X00: CONFIGURATION	151
CBI REGISTER 0X01: VERNIER CONTROL.....	153
CBI REGISTER 0X02: DELAY TAP STATUS	154
CBI REGISTER 0X03: CONTROL STATUS	155

LIST OF FIGURES

FIGURE 1 - S/UNI APEX BLOCK WITH DATAPATH 11

FIGURE 2 - S/UNI APEX BOTTOM VIEW PIN OUT 16

FIGURE 3 - 16BIT RECEIVE CELL TRANSFER FORMAT 48

FIGURE 4 - 8-BIT RECEIVE CELL TRANSFER FORMAT 49

FIGURE 5 - 16-BIT TRANSMIT CELL TRANSFER FORMAT 51

FIGURE 6 - 8-BIT TRANSMIT CELL TRANSFER FORMAT 52

FIGURE 7 - I960 (80960CF) INTERFACE..... 60

FIGURE 8 - POWERPC (MPC860) INTERFACE 60

FIGURE 9 - SAR ASSIST TRANSMIT CELL TRANSFER FORMAT 63

FIGURE 10- SAR ASSIST RECEIVE CELL TRANSFER FORMAT 64

FIGURE 11- SERVICE ARBITRATION HIERARCHY 66

FIGURE 12- QUEUE LINKED LIST STRUCTURE 67

FIGURE 13- TRAFFIC SHAPING ON THE WAN PORT 71

FIGURE 14- NON-INTEGERSHPINCR..... 72

FIGURE 15- THRESHOLDS AND COUNT DEFINITIONS..... 76

FIGURE 16- EPD/PPD CONGESTION DISCARD RULES 79

FIGURE 17 CELL CONGESTION DISCARD RULES 80

FIGURE 18 FCQ DISCARD RULES 81

FIGURE 19- 1 BANK CONFIGURATION FOR 1MB OF ZBT SSRAM 85

FIGURE 20- 1 BANK CONFIGURATION FOR 1MB OF LATE WRITE SSRAM86

FIGURE 21- 2 BANK CONFIGURATION FOR 2MB OF ZBT SSRAM 87

FIGURE 22- 2 BANK CONFIGURATION FOR 2MB OF LATE WRITE SSRAM88

FIGURE 23- 4 BANK CONFIGURATION FOR 4MB OF ZBT SSRAM	89
FIGURE 24- CELL STORAGE MAP.....	90
FIGURE 25- 4 MB – 64K CELLS.....	91
FIGURE 26- 8 MB – 128K CELLS.....	91
FIGURE 27- 16 MB – 256K CELLS.....	92
FIGURE 28- CONTEXT LOCATION.....	157
FIGURE 29- INPUT OBSERVATION CELL (IN_CELL).....	191
FIGURE 30- OUTPUT CELL (OUT_CELL)	192
FIGURE 31- BI-DIRECTIONAL CELL (IO_CELL)	192
FIGURE 32- LAYOUT OF OUTPUT ENABLE AND BIDIRECTIONAL CELLS	193
FIGURE 33- SINGLE WORD READ AND WRITE	195
FIGURE 34- BURST READ AND WRITE.....	196
FIGURE 35- CONSECUTIVE WRITE ACCESSES USING WRDONEB.....	197
FIGURE 36- READ TIMING	198
FIGURE 37- WRITE TIMING.....	198
FIGURE 38- REFRESH.....	199
FIGURE 39- POWER UP AND INITIALIZATION SEQUENCE.....	199
FIGURE 40- READ FOLLOWED BY WRITE TIMING.....	200
FIGURE 41- READ FOLLOWED BY WRITE TIMING.....	201
FIGURE 42- UTOPIA L2 TRANSMIT SLAVE (LOOP & WAN)	202
FIGURE 43- UTOPIA L1 RECEIVE MASTER (LOOP & WAN).....	202
FIGURE 44- UTOPIA L2 RECEIVE MASTER (LOOP & WAN).....	203
FIGURE 45- ANY-PHY RECEIVE MASTER (LOOP & WAN).....	204
FIGURE 46- UTOPIA L2 RECEIVE SLAVE (LOOP & WAN).....	205

FIGURE 47- WAN UTOPIA L1 TRANSMIT MASTER	205
FIGURE 48- LOOP UTOPIA L1 TRANSMIT MASTER.....	206
FIGURE 49- WAN UTOPIA L2 TRANSMIT MASTER	206
FIGURE 50- LOOP UTOPIA L2 TRANSMIT MASTER.....	207
FIGURE 51- WAN ANY-PHY TRANSMIT MASTER	207
FIGURE 52- LOOP ANY-PHY TRANSMIT MASTER	208
FIGURE 53- RSTB TIMING.....	212
FIGURE 54- SYNCHRONOUS I/O TIMING	213
FIGURE 55- JTAG PORT INTERFACE TIMING	217
FIGURE 56- MECHANICAL DRAWING 352 PIN BALL GRID ARRAY (SBGA)220	

LIST OF TABLES

TABLE 1 - TERMINOLOGY 1

TABLE 2 - SAMPLE FEATURE SET AS A FUNCTION OF MEMORY
CAPACITY 15

TABLE 3 - PIN TYPE DEFINITION 17

TABLE 4 - NUMBER OF PORTS SUPPORTED, RECEIVE INTERFACE 50

TABLE 5 - NUMBER OF PORTS SUPPORTED, TRANSMIT INTERFACE.. 53

TABLE 6 - EXAMPLE WIRR TRANSMISSION SEQUENCE 56

TABLE 7 - AVAILABLE QUEUING PROCEDURES 68

TABLE 8 - OAM & RRM CELL IDENTIFICATION 73

TABLE 9 - CONGESTION ERROR FLAGS 77

TABLE 10 - CONGESTION DISCARD RULES SELECTION 78

TABLE 11 - STATISTICAL COUNTS 82

TABLE 12 - IN/OUT BOUND CLP STATE FOR STATISTICAL COUNTS 83

TABLE 13 - CONGESTION RULE & COUNT SUMMARY 83

TABLE 14 - RECEIVE INTERFACE THROUGHPUT, MCELLS/SEC 94

TABLE 15 - QUEUE ENGINE THROUGHPUT, MCELLS/SEC 95

TABLE 16 - TRANSMIT INTERFACE THROUGHPUT, MCELLS/SEC 95

TABLE 17 - EXTERNAL QUEUE CONTEXT MEMORY MAP 158

TABLE 18 - INTERNAL QUEUE CONTEXT MEMORY MAP 158

TABLE 19 - INTERNAL WAN PORT SCHEDULER CONTEXT MEMORY MAP
..... 159

TABLE 20 - INTERNAL LOOP PORT SCHEDULER CONTEXT MEMORY MAP
..... 159

TABLE 21 - 2 BIT LOGARITHMIC, 2 BIT FRACTIONAL 160

TABLE 22 - 4 BIT LOGARITHMIC, 2 BIT FRACTIONAL	160
TABLE 23 - 4 BIT LOGARITHMIC, 4 BIT FRACTIONAL	160
TABLE 24 - VC CONTEXT RECORD STRUCTURE	161
TABLE 25 - VC STATISTICS RECORD STRUCTURE	168
TABLE 26 - VC ADDRESS MAP RECORD STRUCTURE	168
TABLE 27 - PORT THRESHOLD CONTEXT RECORD STRUCTURE	170
TABLE 28 - PORT COUNT CONTEXT RECORD STRUCTURE.....	171
TABLE 29 - CLASS SCHEDULER RECORD STRUCTURE	173
TABLE 30 - CLASS CONTEXT RECORD STRUCTURE	176
TABLE 31 - SHAPE TXSLOT CONTEXT RECORD STRUCTURE	178
TABLE 32 - SHAPE RATE CONTEXT RECORD STRUCTURE.....	179
TABLE 33 - CELL CONTEXT RECORD STRUCTURE	180
TABLE 34 - FREE COUNT CONTEXT STRUCTURE	181
TABLE 35 - OVERALL COUNT CONTEXT STRUCTURE.....	181
TABLE 36 - CONGESTION DISCARD CONTEXT STRUCTURE	182
TABLE 37 - MAXIMUM CONGESTION ID CONTEXT STRUCTURE.....	183
TABLE 38 - MISC ERROR CONTEXT STRUCTURE.....	183
TABLE 39 - WAN TRANSMIT PORT POLLING WEIGHT	184
TABLE 40 - WAN POLL WEIGHT FORMAT	185
TABLE 41 - WAN CLASS STATUS.....	185
TABLE 42 - LOOP TRANSMIT PORT POLLING SEQUENCE	186
TABLE 43 - LOOP TRANSMIT PORT POLLING WEIGHT.....	187
TABLE 44 - LOOP CLASS STATUS	188
TABLE 45 - INSTRUCTION REGISTER.....	190

TABLE 46 - IDENTIFICATION REGISTER	190
TABLE 47 - BOUNDARY SCAN REGISTER	190
TABLE 48 - ABSOLUTE MAXIMUM RATINGS	209
TABLE 49 - D.C. CHARACTERISTICS	210
TABLE 50 - RTSB TIMING	212
TABLE 51 - SYSCLK TIMING	213
TABLE 52 - CELL BUFFER SDRAM INTERFACE	213
TABLE 53 - CONTEXT MEMORY ZBT & LATE WRITE SSRAM INTERFACE	213
TABLE 54 - MICROPROCESSOR INTERFACE	214
TABLE 55 - LOOP ANY-PHY TRANSMIT INTERFACE	214
TABLE 56 - WAN ANY-PHY TRANSMIT INTERFACE	215
TABLE 57 - LOOP ANY-PHY RECEIVE INTERFACE	215
TABLE 58 - WAN ANY-PHY RECEIVE INTERFACE	216
TABLE 59 - JTAG PORT INTERFACE	217

This page left blank intentionally.

1 DEFINITIONS

Table 1 - Terminology

Term	Definition
AAL5	ATM Adaptation Layer
ABR	Available Bit Rate
Any-PHY	Interoperable version of UTOPIA and SCI-PHY, with inband addressing.
ATLAS	PMC's OAM and Address Resolution device
ATM	Asynchronous Transfer Mode
BOM	Beginning of Message
CBI	Common Bus Interface
CBR	Constant Bit Rate
CDV	Cell Delay Variation
CDVT	Cell Delay Variation Tolerance
CES	Circuit Emulation Service
CLP	Cell Loss Priority
COM	Continuation of Message
COS	Class of Service
CTD	Cell Transfer Delay
DLL	Delay Locked Loop
DSL	Digital Subscriber Loop
DSLAM	DSL access Multiplexer
DUPLEX	PMC UTOPIA deserializer
ECI	Egress Connection Identifier
EFCI	Early forward congestion indicator
EOM	End of Message
EPD	Early Packet Discard
FIFO	First-In-First-Out
GCRA	Generic Cell Rate Algorithm

GFR	Guaranteed Frame Rate
IBT	Intrinsic Burst Tolerance
ICI	Ingress Connection Identifier
MBS	Maximum Burst Size
MCR	Minimum Cell Rate
OAM	Operation, Administration and Maintenance
PCR	Peak Cell Rate
PDU	Packet Data Unit
PHY	Physical Layer Device
PPD	Partial Packet Discard
PTI	Payload Type Indicator
QOS	Quality of Service
QRT	PMC's traffic management device
QSE	PMC's switch fabric device
RRM	Reserved or Resource Management
SAR	Segmentation and Re-assembly
SCI-PHY	PMC-Sierra enhanced UTOPIA bus
SCR	Sustained Cell Rate
UBR	Unspecified Bit Rate
UTOPIA	Universal Test & Operations PHY Interface for ATM
VBR	Variable Bit Rate
VCC	Virtual Channel Connection
VORTEX	PMC UTOPIA/Any-PHY slave serializer
VPC	Virtual Path Connection
WAN	Wide Area Network
WIRR	Weighted Interleaved Round Robin
WRR	Weighted Round Robin
ZBT	Zero Bus Turnaround

2 FEATURES

- Monolithic single chip ATM traffic manager providing VC queuing/shaping and VC, Class Of Service(COS), and Port scheduling, congestion management, and switching across 2048 ports.
- Targeted at systems where many low speed ATM data ports are multiplexed onto few high speed ports.
- 869 Kcells/s non shaped throughput in full duplex.
- 1.73 Mcells/s non shaped throughput in half duplex.
- 1.42 Mcells/s shaped throughput (aggregate of the four shapers)
- Supports four WAN uplink ports, with port aliasing
- Supports 2048 loop ports. Loop port can support an uncongested rate up to 230Kcells/sec.
- Provides 4 Classes of Service per port with configurable traffic parameters enabling support for a mix of CBR, VBR, GFR, and UBR classes.
- Provides 64k per-VC queues individually assignable to any COS in any port.
- Provides support of up to 256k cells of shared buffer
- Provides 2 independent cell emission schedulers, 1 for the WAN ports, and 1 for the Loop ports. The schedulers have the following features: Three level hierarchical cell emission scheduling at the port, class, and VC levels.
 - WAN Port Scheduling
 - Weighted Interleaved Round Robin WAN port scheduling.
 - Per port Priority Fair Queued class scheduling with port independence.
 - Per Class
 - Weighted Fair Queued VC scheduling with class independence or

- Shaped Fair Queued VC scheduling applying rate based per VC shaping or
- Frame Continuous Queued VC scheduling for VC Merge and packet re-assembly.
- Loop Port Scheduling
 - Weighted Interleaved Round Robin Loop port scheduling.
 - Per port Priority Fair Queued class scheduling with port independence.
 - Per Class
 - Weighted Fair Queued VC scheduling with class independence or
 - Frame Continuous Queued scheduling for VC Merge and packet re-assembly
- Congestion Control applied per-VC, per-class, per-port and per-direction.
 - Flexible, progressive hierarchical throttling of buffer consumption. Provides sharing of resources during low congestion, memory reservation during high congestion.
 - Applies EPD and PPD on a per-VC, per-class, per-port, and per-direction basis with CLP differentiation, following emerging GFR standards.
 - Provides EFCI marking on a per VC basis.
 - Provides interrupts and indication of most recent VC/Class/Port that exceeded maximum thresholds.
- Provides flexible VPC or VCC switching selectable on a per VC basis as follows
 - Any WAN port to any WAN port
 - Any WAN port to any Loop port
 - Any Loop port to any WAN port.
 - Any Loop port to any Loop port.

- Microprocessor port to any loop or WAN port.
- Any loop or WAN port to microprocessor port
- VP Termination (in conjunction with ATLAS)
- VPI or VPI/VCI header mapping
- VC merge
- Provides flexible signaling and control capabilities
 - Provides 4 independent uP transmit queues
 - Provides simultaneous AAL5 SAR assistance for traffic to/from the uP on up to 64k VCs.
 - Supports uP cell injection into any queue.
 - Provides per VC selectable OAM cell pass through or switching to microprocessor port.
 - Supports CRC10 calculation for OAM cells destined for/originating from the microprocessor.
- Diagnostic access provided to context memory and cell buffer memory via the microprocessor.
- Provides per VC CLP0/1 transmit counts.
- Provide global per CLP0/1 discard counts
- Provides various error statistics accumulation.
- Determines the ingress connection identifier from one of several locations: the cell prepend, the VPI/VCI field, or the HEC/UDF field.
- Interface support
 - Provides a 8/16-bit Any-PHY compliant master/slave Loop side interface supporting up to 2048 ports (logical PHYs).
 - Provides an 8/16-bit Any-PHY compliant master/slave WAN side interface supporting up to 4 ports (PHYs).

- Provides a 32-bit multiplexed microprocessor bus interface for signaling, control, and cell message extraction and insertion, context memory access, control and status monitoring, and configuration of the IC.
- Provides a 32-bit SDRAM interface for cell buffering.
- Provides a 36-bit pipelined ZBT or register to register late write SSRAM interface for context storage.
- Packaging
 - Provides a standard 5 signal P1149.1 JTAG test port for boundary scan board test purposes.
 - Implemented in low power, 0.25 micron, +2.5/3.3V CMOS technology with CMOS compatible inputs and outputs.
 - 352-pin high-performance ball grid array (SBGA) package.

3 APPLICATIONS

- DSL Access Multiplexers
- ATM Switches
- Multiservice Access Multiplexers
- 3rd generation wireless base stations and base station controllers
- OC-12 ingress congestion and traffic management
- OC-12 egress traffic manager and shaper

4 REFERENCES

1. PMC-Sierra; "Saturn Compatible Interface For ATM PHY Layer And ATM Layer Devices, Level 2"; PMC-940212; Dec. 8, 1995
2. PMC-Sierra; VORTEX engineering document
3. PMC-Sierra; DSLAM engineering document
4. "Traffic Management And Switching With The Vortex Chip Set: S/UNI-APEX Technical Overview", PMC-981024
5. PMC-Sierra, "S/UNI APEX Test Bench Engineering Document", PMC-981029
6. ATM Forum, "Universal Test & Operations PHY Interface for ATM (UTOPIA), Level 2", Version 1.0, af-phy-0039.000, June 1995
7. ITU-T Recommendation I.432.1, "B-ISDN user-network interface – Physical layer specification: General characteristics", 08/96
8. ITU-T Recommendation I.363, "B-ISDN ATM Adaptation Layer (AAL) Specification", March 1993
9. AF Traffic Management Specification Version 4.1 AF-TM-0121.000, March 1999.
10. AF Traffic Management Baseline Text Document BTM-TM-01.01, April 1998
11. I.610 OAM
12. PMC Sierra, "Saturn Interface Specification and Interoperability Framework for Packet and Cell Transfer Between Physical Layer and Link Layer Devices", PMC980902
13. PMC Sierra, "S/UNI APEX H/W Programmer's Guide", PMC-991454

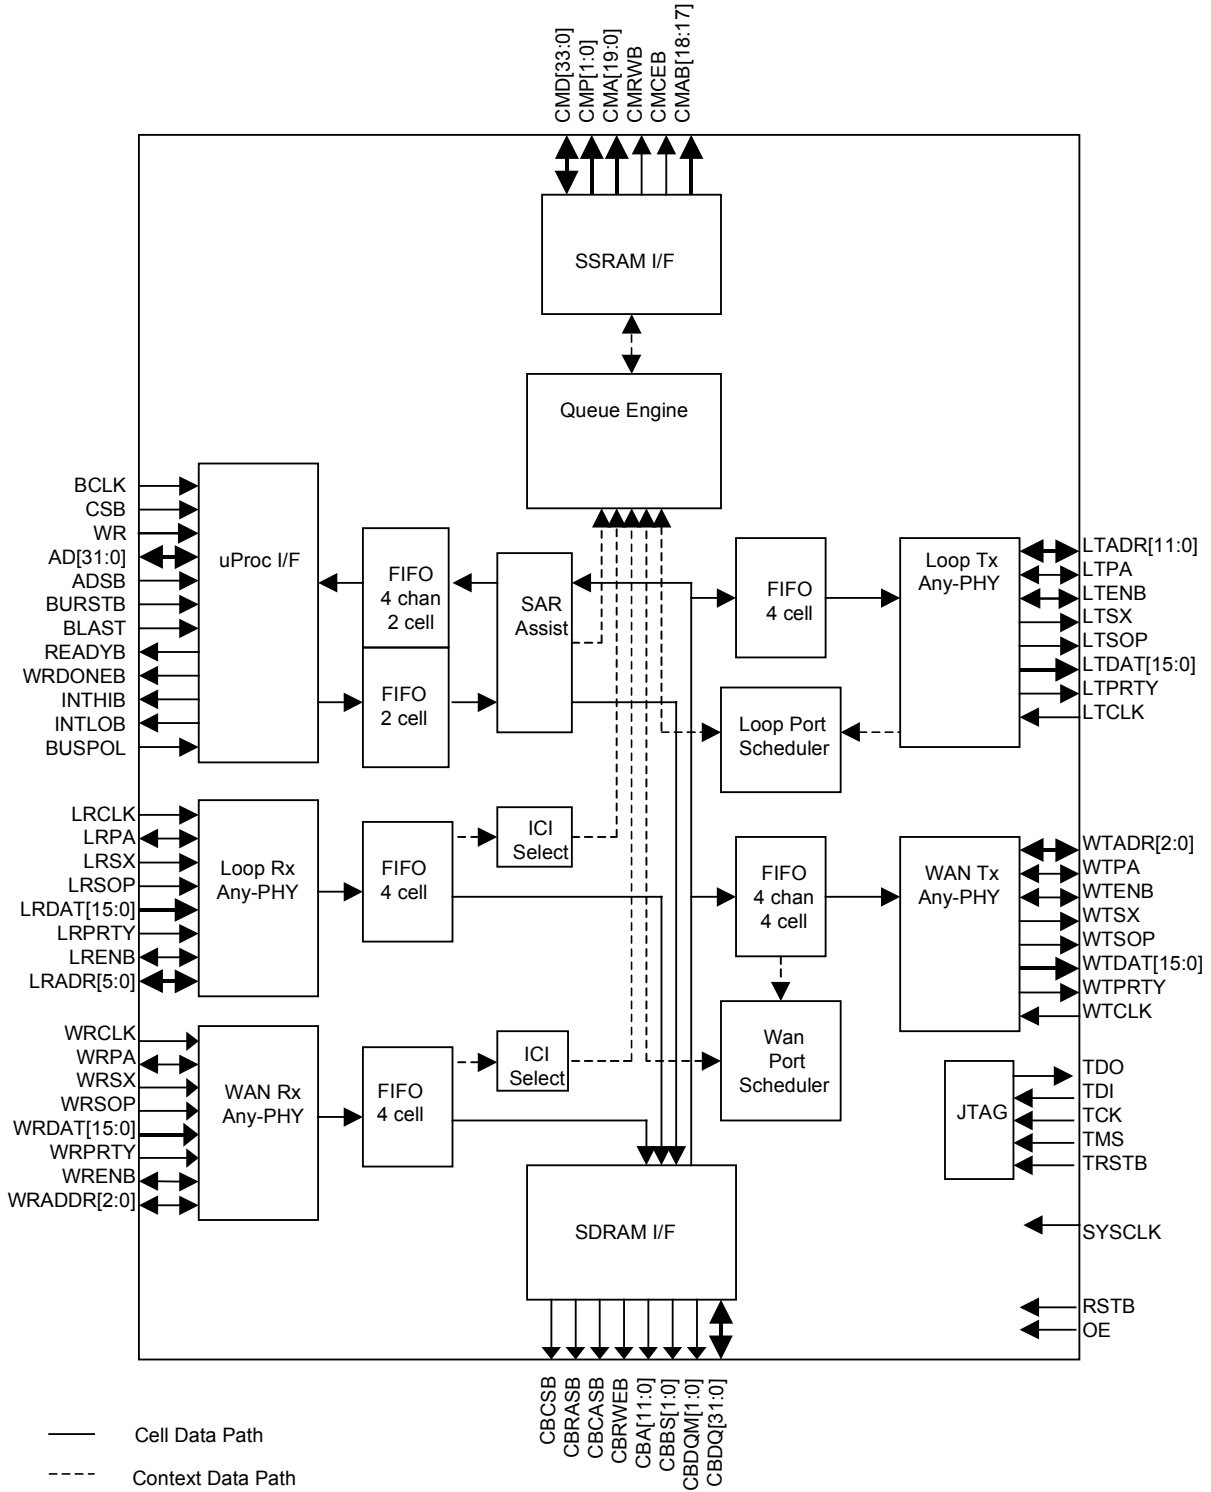
5 APPLICATION EXAMPLES

Please refer to the document "Traffic Management And Switching With The Vortex Chip Set: S/UNI-APEX Technical Overview", PMC-981024

6 **BLOCK DIAGRAM**

Figure 1 shows the function block diagram of the S/UNI APEX ATM traffic manager. The functional diagram is arranged such that cell traffic flows through the S/UNI APEX from left to right.

Figure 1 - S/UNI APEX Block with Datapath



7 DESCRIPTION

The PM7326 S/UNI APEX is a full duplex ATM traffic management device, providing cell switching, per VC queuing, traffic shaping, congestion management, and hierarchical scheduling to up to 2048 loop ports and up to 4 WAN ports.

The S/UNI APEX provides per-VC queuing for 64K VCs. A per-VC queue may be allocated to any Class of Service (COS), within any port, in either direction (ingress or egress path). Per-VC queuing enables PCR or SCR per-VC shaping on WAN ports and greater fairness of bandwidth allocation between VCs within a COS.

The S/UNI APEX provides three level hierarchical scheduling for port, COS, and VC level scheduling. There are two, three level schedulers; one for the loop ports and one for the WAN ports. The three level scheduler for the WAN ports provides

- Weighted Interleaved Round Robin (WIRR) scheduling across the 4 WAN ports enabling selectability of bandwidth allocation between the ports.
- Priority Fair scheduling across the 4 COS's within each port. This class scheduler is a modified priority scheduler allowing minimum bandwidth allocations to lower priority classes within the port. Class scheduling within a port is independent of activity on all other ports.
- There are three types of VC schedulers. VC scheduling within a class is independent of activity on all other classes
 - Shaped fair queuing is available for 4 classes. If the COS is shaped, each VC within the class is scheduled for emission based on its VCs shaping rate. During class congestion, the VC scheduler may lower a VCs rate in proportion to a normalization factor calculated as a function of the VCs rate and the aggregate rate of all active VCs within the class.
 - Weighted Interleaved Round Robin scheduling in which weights are used to provide fairness between the VCs within a class.
 - Frame continuous scheduling where an entire packet is accumulated prior to transferring to a class queue.

The three level scheduler for the loop ports provides

- Weighted Interleaved Round Robin (WIRR) scheduling across the 2048 loop ports enabling selectability of bandwidth allocation between the ports

and ensuring minimal PHY layer FIFOing is required to support a wide range of port bandwidths.

- Priority scheduling across the 4 COS's within each port. Class scheduling within a port is independent of activity on all other ports.
- VCs within a class are scheduled with a Round Robin scheduler or Frame Continuous scheduling. VC scheduling within a class is independent of activity on all other classes. Shaping is not supported on loop ports.

The S/UNI APEX forwards cells via tail of queue enqueueing and head of queue dequeuing (emission) where tail of queue enqueueing is controlled by the VC context record and subject to congestion control, and head of queue dequeuing is controlled by the three level hierarchical schedulers. The VC context record allows for enqueueing to any queue associated with any port, thus full switching is supported, any port to any port.

The S/UNI APEX supports up to 256k cells of shared buffering in a 32-bit wide SDRAM. Memory protection is provided via an inband CRC on a cell by cell basis. Buffering is shared across direction, port, class, and VC levels. The congestion control mechanism provides guaranteed resources to all active VCs, allows sharing of available resources to VCs with excess bandwidth, and restricts buffer allocation on a per-VC, per-class, per-port, and per-direction basis. The congestion control mechanism supports PPD and EPD on a CLP0 and CLP1 basis across per-VC, per-class, per-port, and per-direction structures. EFCI marking is supported on a per-VC basis. Congestion thresholds and packet awareness is selectable on a per connection basis.

The S/UNI APEX provides flexible capabilities for signaling, management, and control traffic. There are 4 independent uP receive queues to which both cell and AAL5 frame traffic may be en-queued for termination by the uP. A staging buffer is also provided enabling the uP to en-queue both cell and AAL5 frame traffic to any outgoing queue. AAL5 SAR assistance is provided for AAL5 frame traffic to and from the uP. AAL5 SAR assistance includes the generation and checking of the 32-bit CRC field and the ability to reassemble all the cells from a frame in the VC queue prior to placement on the uP queues. Any or all of the 64k VCs may be configured to be routed to/from the uP port. Any or all of the VCs configured to be routed to/from the uP port may also be configured for AAL5 SAR assistance simultaneously. OAM cells may optionally (per-VC selectable) be routed to a uP receive queue or switched with the user traffic. CRC10 generation and checking is optionally provided on OAM cells to/from the uP.

The S/UNI APEX maintains cell counts of CLP0 and CLP1 cell transmits on a per-VC basis. Global CLP0 and CLP1 congestion discards are also maintained. Various error monitoring conditions and statistics are accumulated or flagged.

The uP has access to both internal S/UNI APEX registers and the context memory as well as diagnostic access to the cell buffer memory.

The S/UNI APEX provides a 8/16-bit Any-PHY compliant loop side master/slave interface supporting up to 2048 ports. Egress cell transfers across the interface are identified via an inband port identifier prepended to the cell. The slave devices must match the inband port identifier with their own port ID or port ID range in order to accept the cell. Per port egress flow control is effected via an 12-bit address polling bus to which the appropriate slave device responds with out of band per port flow control status. Ingress cell transfers across the interface are effected via a combination of UTOPIA L2 flow control polling and device selection for up to 32 slave devices. The Any-PHY loop side interface may be reconfigured as a standard single port UTOPIA L2 compliant slave interface. 16-bit preponds are optionally supported on both ingress and egress for cell flow identification enabling use with external address resolution devices, switch fabric interfaces, or other layer devices.

The S/UNI APEX provides an 8/16-bit Any-PHY or UTOPIA L2 compliant WAN side master/slave interface supporting up to 4 ports. 16-bit preponds are optionally supported on both ingress and egress for cell flow identification enabling use with external address resolution devices, switch fabric interfaces, or other layer devices. The WAN port has port aliasing on the egress, providing in service re-direction without requiring re-programming the context of active VCs.

The S/UNI APEX provides a 32-bit microprocessor bus interface for signaling, control, cell and frame message extraction and insertion, VC. Class and port context access, control and status monitoring, and configuration of the IC. Microprocessor burst access for registers, cell and frame traffic is supported.

The S/UNI APEX provides a 36-bit ZBT or late write SSRAM interface for context storage supporting up to 4MB of context for up to 64kVCs and up to 256k cell buffer pointer storage. Context Memory protection is provided via 2 bits of parity over each 34-bit word.

The total number of cells, the total number of VCs, support for address mapping and shaped fair queuing is limited to the amount of context and cell buffer memory available. Below is a table illustrating the most common combinations of memory/features.

Table 2 - Sample feature set as a function of memory capacity

Context Memory Size SSRAM	Cell Buffer Memory Size SDRAM	# VC	# Cell Buffers	Address Mapping Support	Shaping Support
1 MB	4MB	16 K	64 K	Yes	No
2 MB	4MB	16 K	64 K	Yes	Yes
2 MB	4MB	64 K	64 K	No	No
4 MB	16MB	64 K	256 K	Yes	Yes

The S/UNI APEX provides facilities to enable sparing capability with another S/UNI APEX device. The facilities enable a 'warm standby' capability in which connection setup between the two devices can be maintained identically but some cell loss will occur at the point of device swapping. The facilities do not include a cell by cell lock step between the two S/UNI APEX devices. To avoid any cell replication, queues in the 'spare' S/UNI APEX will be kept empty, thus causing all queued traffic in the 'active' S/UNI APEX to be lost at the point of switch over. However, since connection setup is maintained identically between the two S/UNI APEX devices, switch over can happen instantaneously, thus avoiding any connection timeout or tear down issues.

The S/UNI APEX facilities provided are the disable and filter control bits in the Receive and Transmit Control register. These control bits are asserted in the spare S/UNI APEX to ensure the queues remain empty until swapping is initiated. Alternatively, asserting only the filter enable bits allow signalling and control traffic continuity to be maintained to the spare S/UNI APEX to enable datapath integrity testing on the spare plane and to ensure control communications paths to the spare plane are usable.

8 PIN DIAGRAM

The S/UNI APEX is packaged in a 352-pin ball grid array (SBGA) package having a body size of 35 mm by 35 mm.

Figure 2 - S/UNI APEX Bottom View Pin out

	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	
A	vss5	vss4	CMD [0]	CMD [4]	CMD [7]	CMD [11]	CMD [14]	CMD [18]	CMD [21]	PCH	CMD [27]	CMD [30]	vss3	vss2	AD [2]	AD [4]	PCH	AD [10]	AD [13]	AD [17]	AD [20]	AD [24]	AD [27]	AD [31]	vss1	vss0	A
B	vss9	vdd10	vss8	CMD [1]	CMD [5]	CMD [8]	CMD [12]	CMD [15]	CMD [19]	CMD [23]	CMD [25]	CMD [29]	CMD [32]	CMD [33]	SYSCLK	AD [6]	AD [8]	AD [12]	AD [16]	AD [19]	AD [23]	AD [26]	AD [30]	vss7	vdd9	vss6	B
C	CMRWB	vss11	vdd12	CMP [1]	CMD [2]	CMD [6]	CMD [9]	CMD [13]	CMD [16]	CMD [20]	CMD [24]	CMD [28]	CMD [31]	AD [0]	AD [3]	AD [7]	AD [11]	AD [15]	AD [18]	AD [22]	AD [25]	AD [29]	INTHB	vdd11	vss10	BCLK	C
D	CMAB [1]	CMCEB	CMP [0]	vdd17	nc	CMD [3]	PCH	CMD [10]	vdd16	CMD [17]	CMD [22]	CMD [26]	vdd15	AD [1]	AD [5]	AD [9]	AD [14]	vdd14	AD [21]	PCH	AD [28]	nc	vdd13	INTLOB	WRDONE	BLAST	D
E	CMA [16]	CMA [17]	CMA [19]	nc																			BJSPOL	BTERMB	BURSTB	CSB	E
F	CMA [12]	CMA [15]	CMAB [0]	CMA [18]																			READYB	WR	ADSB	LRADR [2]	F
G	CMA [9]	CMA [11]	CMA [14]	PCH																			PCH	LRADR [0]	LRADR [3]	LRADR [5]	G
H	CMA [5]	CMA [8]	CMA [10]	CMA [13]																			LRADR [1]	LRADR [4]	LRFB	LRFRY	H
J	CMA [2]	CMA [4]	CMA [7]	vdd18																			vdd19	LRPA	LRSDP	LRDAT [0]	J
K	PCH	CMA [9]	CMA [3]	CMA [6]																			LRCLK	LRBK	LRDAT [1]	LRDAT [3]	K
L	LRDAT [12]	LRDAT [14]	LRDAT [15]	CMA [1]																			PCH	LRDAT [2]	LRDAT [4]	LRDAT [6]	L
M	LRDAT [9]	LRDAT [10]	LRDAT [11]	LRDAT [13]																			LRDAT [5]	LRDAT [7]	LRDAT [8]	LRDAT [9]	M
N	vss3	LRDAT [6]	LRDAT [7]	LRDAT [8]																			vdd20	LRDAT [10]	LRDAT [11]	vss12	N
P	vss5	LRDAT [5]	LRDAT [4]	vdd21																			LRDAT [14]	LRDAT [13]	LRDAT [12]	vss14	P
R	LRDAT [3]	LRDAT [2]	LRDAT [1]	LRPA																			PCH	WRADR [1]	WRADR [9]	LRDAT [15]	R
T	LRDAT [0]	LRCLK	LRFB	LRSDP																			WRSDP	WRCLK	WRFB	WRADR [2]	T
U	PCH	LRBK	LRADR [11]	LRADR [8]																			WRDAT [3]	WRDAT [0]	WRFRY	WRPA	U
V	LRFRY	LRADR [10]	LRADR [7]	vdd23																			vdd22	WRDAT [4]	WRDAT [1]	WRBK	V
W	LRADR [9]	LRADR [6]	LRADR [4]	LRADR [11]																			PCH	WRDAT [7]	WRDAT [5]	WRDAT [2]	W
Y	LRADR [5]	LRADR [3]	LRADR [0]	WRDAT [14]																			WRDAT [13]	WRDAT [10]	WRDAT [8]	WRDAT [6]	Y
AA	LRADR [2]	WRDAT [15]	WRDAT [13]	WRDAT [10]																			RSTB	WRDAT [14]	WRDAT [11]	WRDAT [9]	AA
AB	PCH	WRDAT [12]	WRDAT [9]	nc																			nc	TDI	WRDAT [15]	WRDAT [12]	AB
AC	WRDAT [11]	WRDAT [8]	WRDAT [6]	vdd4	SCANMB	WRDAT [2]	WRFB	WRSDP	vdd3	CBA [9]	CBA [4]	CBA [0]	CBCASB	vdd2	CBDQ [26]	CBDQ [22]	CBDQ [17]	vdd1	CBDQ [10]	CBDQ [7]	CBDQ [3]	nc	vdd0	RSTB	TMS	OE	AC
AD	WRDAT [7]	vss17	vdd6	SCANEN	WRDAT [3]	WRPA	WRCLK	WRADR [1]	CBA [10]	CBA [6]	CBA [2]	CBBES [0]	CBBEB	CBDQ [31]	CBDQ [28]	CBDQ [24]	CBDQ [20]	CBDQ [16]	CBDQ [13]	CBDQ [9]	CBDQ [5]	CBDQ [2]	TDO	vdd5	vss16	TCK	AD
AE	vss21	vdd8	vss19	WRDAT [4]	WRDAT [0]	WRBK	WRADR [0]	CBA [11]	CBA [7]	CBA [3]	CBA [1]	CBCSB	CBDQM [1]	CBDQM [0]	CBDQ [29]	CBDQ [25]	CBDQ [23]	CBDQ [19]	CBDQ [15]	CBDQ [12]	CBDQ [8]	CBDQ [5]	CBDQ [1]	vss20	vdd7	vss18	AE
AF	vss27	vss26	WRDAT [5]	WRDAT [1]	PCH	WRFRY	WRADR [2]	CBA [8]	CBA [5]	PCH	CBBES [1]	CBBASB	vss25	vss24	CBDQ [30]	CBDQ [27]	PCH	CBDQ [21]	CBDQ [18]	CBDQ [14]	CBDQ [11]	PCH	CBDQ [4]	CBDQ [3]	vss23	vss22	AF

9 PIN DESCRIPTION

Notes on Pin Description:

1. All S/UNI APEX inputs and bi-directionals present minimum capacitive loading
2. LVCMOS, LVTTTL compatible logic levels.
3. All pins are 5V tolerant.
4. Inputs RSTB, OE, TMS, TDI and TRSTB have internal pull-up resistors.
5. The recommended power supply sequencing is as follows:
 - 3.1 VDD power must be supplied either before or simultaneously with PCH.
 - 3.2 The VDD power must be applied before input pins are driven or the input current per pin be limited to less than the maximum DC input current specification. (20 mA)
 - 3.3 Power down the device in the reverse sequence.

Table 3 - Pin Type Definition

Type	Definition
Input	Input
Output	Pin is always driven
Tri-State	Pin is either driven, or held in Hi-Z
BiDi	Bidirectional
OD	Open drain. Either driven low or held in Hi-Z.

9.1 Loop Any-PHY Receive Master/Transmit Slave Interface (28 Signals)

Pin Name	Type	Pin No.	Function
LRCLK	Input	K4	Loop Receive Clock. LRCLK is used to transfer data blocks in the receive directions across the Any-PHY interface. LRCLK must cycle at a 52 MHz or lower instantaneous rate.

Pin Name	Type	Pin No.	Function
LRPA	Input (Master) Tri-state (Slave)	J3	<p>Loop Receive Packet Available. LRPA indicates whether at least one cell is queued for transfer in the selected PHY device.</p> <p>This pin is in Hi-Z when the loop receive interface is not enabled.</p> <p>If receive master mode is selected, this signal is an input. The selected PHY device drives LRPA with the cell availability status N LRCLK cycles after LRADR[5:0] matches the PHY device address. If the PHY device is a UTOPIA device, N=1. If the PHY device is an Any-PHY device, N=2. Assertion of LRPA indicates that at least one entire cell is available.</p> <p>If transmit slave mode is selected, this signal is a tri-state output. The S/UNI APEX drives LRPA high 1 LRCLK after LRADR[5:0] matches the programmed LoopRxSlaveAddr register. A logical high indicates that the S/UNI APEX is capable of accepting at least one cell.</p> <p>LRPA is sampled/updated/Hi-Z'd on the rising edge of LRCLK.</p>

Pin Name	Type	Pin No.	Function
LRENB	Input (Slave) Output (Master)	H2	<p>Loop Receive Enable. The active low receive enable (LRENB) signal is used to initiate the transfer of a data block from the selected Physical layer device to the S/UNI APEX.</p> <p>This pin is in Hi-Z when the loop receive interface is not enabled.</p> <p>If receive master mode is selected, this signal is an output and the start of block transfer must occur 1 or 2 LRCLK cycles after device selection occurs. If the PHY device is a UTOPIA device, N=1. If the PHY device is an Any-PHY device, N=2. Device selection occurs when the selected device address is placed on LRADR[5:0] with LRENB held high followed by LRENB low in the next LRCLK period. LRENB is held low for M cycles where M is the number of 8 or 16-bit words in the block transfer.</p> <p>If transmit slave mode is selected, this signal is an input and LRDAT[15:0] word is accepted coincident with LRENB being sampled.</p> <p>LRENB is sampled/updated on the rising edge of LRCLK.</p>

Pin Name	Type	Pin No.	Function
LRADR[0] .. LRADR[5]	Input (Slave) Output (Master)	G3 H4 F1 G2 H3 G1	<p>Loop Receive Address. The LRADR[5:0] signals are used to address up to thirty two Physical layer devices for the purposes of polling and device selection.</p> <p>This pin is in Hi-Z when the loop receive interface is not enabled.</p> <p>If UL2 or Any-PHY receive master mode is selected, these signals are outputs. LRADR[5:0] selects a device for polling by applying the device address N LRCLK cycles prior to sampling LRPA. If the PHY device is a UTOPIA device, N=1. If the PHY device is an Any-PHY device, N=2. LRADR will insert 1 NULL address between address changes.</p> <p>If UL1 master mode is selected, this bus is driven to a high NULL address.</p> <p>LRADR[5:0] selects a device to transfer a data block when the LRENB is last sampled high. The start of data block transfer must occur 1 or 2 LRCLK cycles after device selection occurs.</p> <p>LRADR[5:0] = 3F hex is used as the NULL address. No PHY device can match the NULL address.</p> <p>If transmit slave mode is selected, these signals are inputs. The S/UNI drives the LRPA 1 LRCLK after the LRADR[4:0] matches the programmed LoopRxSlaveAddr register, and LRADR[5] is zero.</p> <p>LRADR[5:0] is sampled/updated or on the rising edge of LRCLK.</p>
LRSX	Input	K3	<p>Loop Receive Start of Transfer. LRSX is asserted by the selected PHY device during the first cycle of a data block transfer coinciding with the port address prepend. Required only during Any-PHY mode.</p> <p>For UTOPIA modes, this signal should be tied low.</p> <p>LRSX is sampled on the rising edge of LRCLK.</p>

Pin Name	Type	Pin No.	Function
LRSOP	Input	J2	<p>Loop Receive Start of Packet . LRSOP marks the start of the cell on the LRDAT[15:0] bus. When LRSOP is high, the first data word of the cell is present on the LRDAT[15:0] stream. If the selected device is an Any-PHY device, the LRSOP cycle will be preceded by the LRSX cycle marking the Any-PHY port address transfer cycle.</p> <p>LRSOP considered valid only when the LRENB signal is low. LRSOP becomes high impedance upon sampling LRENB high or if no physical layer device was selected for transfer.</p> <p>LRSOP is sampled on the rising edge of LRCLK.</p>
LRDAT[0] .. LRDAT[15]	Input	J1 K2 L3 K1 L2 M4 L1 M3 M2 M1 N3 N2 P2 P3 P4 R1	<p>Loop Receive Data. LRDAT[15:0] carries the transfer block words that have been read from the physical layer device to the S/UNI APEX internal cell buffers.</p> <p>LRDAT bus is considered valid only when the LRENB signal was low N cycles previous. LRDAT is expected to become high impedance N LRCLK cycles after sampling LRENB high or upon completion of a data block transfer. If the PHY device is a UTOPIA device, N=1. If the PHY device is an Any-PHY device, N=2.</p> <p>All 16 bits are used in 16-bit mode. In 8 bit mode, LRDAT[15:8] should either be tied high or low, as only the first 8 bits LRDAT[7:0] are valid.</p> <p>LRDAT[15:0] is sampled on the rising edge of LRCLK.</p>

Pin Name	Type	Pin No.	Function
LRPRTY	Input	H1	<p>Loop Receive Parity. LRPRTY provides programmable odd/even parity of the LRDAT[15:0] bus.</p> <p>LRPRTY is considered valid only when the LRENB signal was low N cycles previous. If the PHY device is a UTOPIA device, N=1. If the PHY device is an Any-PHY device, N=2. LRPRTY is expected to become high impedance N LRCLK cycles after sampling LRENB high.</p> <p>A parity error is indicated by a status bit and a maskable interrupt.</p> <p>LRPRTY is sampled on the rising edge of LRCLK.</p>

9.2 Loop Any-PHY Transmit Master/Receive Slave Interface (34 Signals)

Pin Name	Type	Pin No.	Function
LTCLK	Input	T25	<p>Loop Transmit Clock. LTCLK is used to transfer data blocks in the transmit direction across the Any-PHY interface. LTCLK must cycle at a 52 MHz or lower instantaneous rate.</p>

Pin Name	Type	Pin No.	Function
LTADR[0] .. LTADR[11]	Output (Master) Input (Slave)	Y24 W23 AA26 Y25 W24 Y26 W25 V24 U23 W26 V25 U24	<p>Loop Transmit Address. The LTADR[11:0] signals are used to address up to 2048 logical channels for the purposes of polling on the LTPA signal. 1 or more PHY devices can share the LTPA signal</p> <p>This pin is in Hi-Z when the loop transmit interface is not enabled.</p> <p>If transmit master mode is selected, these signals are outputs. LTADR[11:0] selects a logical channel for polling by applying the logical channel address N LTCLK cycles prior to sampling LTPA. If the PHY device is a UTOPIA device, N=1. If the PHY device is an Any-PHY device, N=2. LTADR inserts NULL cycles between addresses.</p> <p>For Any-PHY transmit master, LTADR[11:0] corresponds to the PORTID[11:0] fields in the Any-PHY address word prepend format.</p> <p>For UTOPIA L2 transmit master, LTADR[5:0] is also used to select a UTOPIA device to transfer a cell to, when LTENB is last sampled high. LTADR[11:6] should be left unconnected.</p> <p>For UTOPIA L1 transmit master, LTADR[11:0] is unused and should be left unconnected.</p> <p>If UTOPIA L2 receive slave mode is selected, these signals are inputs. The S/UNI APEX drives LTPA high 1 LTCLK after the LTADR[5:0] matches the programmed LoopTxSlaveAddr register. LTADR[11:6] are unused and should be tied either high or low.</p> <p>LTADR[11:0] is sampled/updated on the rising edge of LTCLK.</p>

Pin Name	Type	Pin No.	Function
LTPA	Input (Master) Tri-state (Slave)	R23	<p>Loop Transmit Packet Available. LTPA indicates the availability of space in the selected polled port when polled using the LTADR[11:0] signals.</p> <p>This pin is in Hi-Z when the loop transmit interface is not enabled.</p> <p>If transmit master mode is selected, this signal is an input. The PHY device whose address or address range matches LTADR[11:0] drives the LTPA signal with the transmit FIFO availability status of the selected logical channel N LTCLK cycles after the match. If the PHY device is a UTOPIA device, N=1. If the PHY device is an Any-PHY device, N=2. Assertion of LTPA indicates that at least K entire cell buffer is available in that logical channel. K = 1 if the register LoopTxTwoCellEn = 0. K = 2 if the register LoopTxTwoCellEn = 1.</p> <p>If receive slave mode is selected, this signal is a tri-state output. The S/UNI APEX drives LTPA 1 LTCLK after LTADR[5:0] matches the programmed LoopTxSlaveAddr register. A logical high indicates that at least one cell is available for transmission.</p> <p>LTPA is sampled/updated/Hi-Z'd on the rising edge of LTCLK.</p>
LTENB	Output (Master) Input (Slave)	T24	<p>Loop Transmit Enable. LTENB indicates cell transfers to UTOPIA and SCI-PHY devices. The device is selected via a match on LTADR[11:0] when LTENB is last sampled high.</p> <p>This pin is in Hi-Z when the loop transmit interface is not enabled.</p> <p>If transmit master mode is selected, this signal is an output. LTENB is held low for the duration of the cell transfer.</p> <p>If receive slave mode is selected, this signal is an input.</p> <p>LTENB is sampled/updated on the rising edge of LTCLK.</p>

Pin Name	Type	Pin No.	Function
LTSX	Output	U25	<p>Loop Transmit Start of Transfer. LTSX is asserted by the S/UNI APEX during the first cycle of a data block transfer. LTSX assertion will coincide with the port address prepend, if the cell being transferred has a prepended port address. Required only during Any-PHY mode. Should be left unconnected during UTOPIA modes.</p> <p>LTSX is updated on the rising edge of LTCLK.</p>
LTSOP	Output (Master) Tri-state (Slave)	T23	<p>Loop Transmit Start of Cell. LTSOP marks the start of cell on the LTDAT[15:0] data bus. LTSOP is driven high when the first word of the cell (excluding address prepend) is present on the LTDAT[15:0] stream. LTSOP is asserted for each cell.</p> <p>In transmit master mode, the signal is always driven.</p> <p>In receive slave mode, this signal is driven 1 LTCLK after LTENB is asserted.</p> <p>LTSOP is updated/Hi-Z'd on the rising edge of LTCLK.</p>
LTDAT[0] .. LTDAT[15]	Output (Master) Tri-state (Slave)	T26 R24 R25 R26 P24 P25 N25 N24 N23 M26 M25 M24 L26 M23 L25 L24	<p>Loop Transmit Data. LTDAT[15:0] carries the data block transfers to the physical layer devices.</p> <p>In 8 bit mode, only LTDAT[7:0] are valid.</p> <p>In transmit master mode, the entire bus is always driven.</p> <p>In receive slave mode, this bus is driven 1 LTCLK after LTENB is asserted. Pull up/downs are required for the entire bus, regardless of whether the bus is in 8 or 16 bit mode.</p> <p>LTDAT[15:0] is updated/Hi-Z'd on the rising edge of LTCLK.</p>

Pin Name	Type	Pin No.	Function
LTPRTY	Output (Master) Tri-state (Slave)	V26	<p>Loop Transmit Parity. This signal provides programmable odd/even parity of the LTDAT[15:0] bus.</p> <p>In transmit master mode, the signal is always driven.</p> <p>In receive slave mode, this signal is driven 1 LTCLK after LTENB is asserted.</p> <p>LTPRTY is updated/Hi-Z'd on the rising edge of LTCLK.</p>

9.3 WAN Any-PHY Receive Master/Transmit Slave Interface (25 Signals)

Pin Name	Type	Pin No.	Function
WRCLK	Input	T3	<p>WAN Receive Clock. WRCLK is used to transfer data blocks in the receive direction across the Any-PHY interface. WRCLK must cycle at a 52 MHz or lower instantaneous rate.</p>

Pin Name	Type	Pin No.	Function
WRPA	Input (Master) Tri-state (Slave)	U1	<p>WAN Receive Packet Available. WRPA indicates cell availability.</p> <p>This pin is in Hi-Z when the WAN receive interface is not enabled.</p> <p>If master mode is selected, the selected PHY device drives WRPA with the cell availability status N WRCLK cycles after WRADR[2:0] matches the PHY device address. If the PHY device is a UTOPIA device, N=1. If the PHY device is an Any-PHY device, N=2. Assertion of WRPA indicates that at least one entire cell is available.</p> <p>If slave mode is selected, this signal is an output and the S/UNI APEX plays the roll of a single port UTOPIA L2 slave device driving the WRPA when the WRADR matches the programmed WANRxSlaveAddr register. A logical high indicates that the S/UNI APEX is capable of accepting at least one cell.</p> <p>WRPA is sampled/updated/Hi-Z'd on the rising edge of WRCLK.</p>

Pin Name	Type	Pin No.	Function
WRENB	Output (Master) Input (Slave)	T2	<p>WAN Receive Enable. The active low receive enable (WRENB) output is used to initiate the transfer of a data block from the selected Physical layer device to the S/UNI APEX.</p> <p>This pin is in Hi-Z when the WAN receive interface is not enabled.</p> <p>If master mode is selected, this signal is an output and the start of block transfer must occur 1 or 2 WRCLK cycles after device selection occurs. Device selection occurs when the selected device address is placed on WRADR[2:0] with WRENB held high followed by WRENB low in the next WRCLK period.</p> <p>WRENB is held low for M cycles where M is the number of 8 or 16-bit words in the block transfer.</p> <p>If slave mode is selected, this signal is an input and WRDAT[15:0] word is accepted coincident with WRENB being sampled.</p> <p>WRENB is sampled/updated on the rising edge of WRCLK.</p>

Pin Name	Type	Pin No.	Function
WRADR[0] .. WRADR[2]	Output (Master) Input (Slave)	R2 R3 T1	<p>WAN Receive Address. The WRADR[2:0] signals are used to address up to four Physical layer devices for the purposes of polling and device selection.</p> <p>This pin is in Hi-Z when the WAN receive interface is not enabled.</p> <p>If UL2 or Any-PHY receive master mode is selected, this bus is an output. WRADR[2:0] selects a device for polling by applying the device address N WRCLK cycles prior to sampling WRPA. If the PHY device selected is a UTOPIA device, N=1. If the PHY device selected is an Any-PHY device, N=2. When supporting multiple PHYs, WRADR will insert 1 NULL address between address changes.</p> <p>If UL1 master mode is selected, this bus is driven to a high NULL address.</p> <p>WRADR[2:0] selects a device to transfer a data block when the WRENB is last sampled high. The start of data block transfer must occur 1 or 2 WRCLK cycles after device selection occurs.</p> <p>WRADR[2:0] = 7 hex is used as the NULL address. No PHY device can match the NULL address.</p> <p>If slave mode is selected, this signal is an input and the S/UNI APEX plays the roll of a single port UTOPIA L2 slave device driving the WRPA 1 WRCLK after the WRADR[1:0] matches the programmed WANRxSlaveAddr register, and WRADR[2] is zero.</p> <p>WRADR[2:0] is sampled/updated on the rising edge of WRCLK.</p>
WRSX	Input	V1	<p>WAN Receive Start of Transfer. WRSX is asserted by the selected PHY device during the first cycle of a data block transfer coinciding with the port address prepend. WRSX is ignored during cell transfers from UTOPIA or SCI-PHY devices.</p> <p>WRSX is updated on the rising edge of WRCLK.</p>

Pin Name	Type	Pin No.	Function
WRSOP	Input	T4	<p>WAN Receive Start of Packet . WRSOP marks the start of the cell on the WRDAT[15:0] bus. When WRSOP is high, the first data word of the cell is present on the WRDAT[15:0] stream. If the selected device is an Any-PHY device, the WRSOP cycle will be preceded by the WRSX cycle marking the Any-PHY port address transfer cycle.</p> <p>WRSOP is considered valid only when the WRENB signal is low. WRSOP becomes high impedance upon sampling WRENB high or if no physical layer device was selected for transfer.</p> <p>WRSOP is sampled on the rising edge of WRCLK.</p>
WRDAT[0] .. WRDAT[15]	Input	U3 V2 W1 U4 V3 W2 Y1 W3 Y2 AA1 Y3 AA2 AB1 Y4 AA3 AB2	<p>WAN Receive Data. WRDAT[15:0] carries the transfer block words that have been read from the physical layer device to the S/UNI APEX internal cell buffers. All 16 bits are used in 16-bit mode, only the first 8 bits WRDAT[7:0] are valid in 8-bit mode.</p> <p>The WRDAT bus is considered valid only when the WRENB signal was low N cycles previous. WRDAT is expected to become high impedance N WRCLK cycles after sampling WRENB high or upon completion of a data block transfer. If the PHY device selected is a UTOPIA device, N=1. If the PHY device selected is an Any-PHY device, N=2.</p> <p>WRDAT[15:0] is sampled on the rising edge of WRCLK.</p>

Pin Name	Type	Pin No.	Function
WRPRTY	Input	U2	<p>WAN Receive Parity. WRPRTY provides programmable odd/even parity of the WRDAT[15:0] bus.</p> <p>The WRPRTY signal is considered valid only when the WRENB signal was low N cycles previous. If the PHY device selected is a UTOPIA device, N=1. If the PHY device selected is an Any-PHY device, N=2. WRPRTY is expected to become high impedance N WRCLK cycles after sampling WRENB high.</p> <p>A parity error is indicated by a status bit and a maskable interrupt.</p> <p>WRPRTY is sampled on the rising edge of WRCLK.</p>

9.4 WAN Any-PHY Transmit Master/Receive Slave Interface (25 Signals)

Pin Name	Type	Pin No.	Function
WTCLK	Input	AD20	<p>WAN Transmit Clock. WTCLK is used to transfer data blocks in the transmit direction across the Any-PHY interface. WTCLK must cycle at a 52 MHz or lower instantaneous rate.</p>

Pin Name	Type	Pin No.	Function
WTADR[0] .. WTADR[2]	Output (Master) Input (Slave)	AE20 AD19 AF20	<p>WAN Transmit Address. The WTADR[2:0] signals are used to address up to four logical channels for the purposes of polling.</p> <p>This pin is in Hi-Z when the WAN transmit interface is not enabled.</p> <p>If master mode is selected, these signals are outputs. WTADR[2:0] selects a logical channel for polling by applying the logical channel address N WTCLK cycles prior to sampling WTPA. If the PHY devices are UTOPIA devices, N=1. If the PHY devices are Any-PHY devices, N=2. , WTADR will insert 1 NULL address between address changes</p> <p>For Any-PHY transmit master, WTADR[1:0] corresponds to the PORTID[1:0] fields in the Any-PHY address word prepend format.</p> <p>WTADR[2:0] = 7 hex is used as the NULL address. No PHY device can match the NULL address.</p> <p>For UTOPIA L2 transmit master, WTADR[2:0] signals are also used for cell transfer PHY selection to UTOPIA compliant PHY devices. WTADR[2:0] selects a device to transfer a data block to when the WRENB is last sampled high.</p> <p>For UTOPIA L1 transmit master, WTADR[1:0] contains the value of the WANTxSlaveAddr register. WTADR[2] is held low.</p> <p>If UTOPIA L2 receive slave mode is selected, these signals are inputs and the S/UNI APEX plays the roll of a single port UTOPIA L2 slave device driving the WTPA 1 WTCLK after the WTADR matches the programmed WANTxSlaveAddr register.</p> <p>WTADR[2:0] is sampled/updated on the rising edge of WTCLK.</p>

Pin Name	Type	Pin No.	Function
WTPA	Input (Master) Tri-state (Slave)	AD21	<p>WAN Transmit Packet Available. WTPA indicates cell availability.</p> <p>This pin is in Hi-Z when the WAN transmit interface is not enabled.</p> <p>If master mode is selected, this signal is an input. The PHY device whose address or address range matches WTADR[2:0] drives the WTPA signal with the transmit FIFO availability status of the selected logical channel N WTCLK cycles after the match. If the PHY devices are UTOPIA devices, N=1. If the PHY devices are Any-PHY devices, N=2. Assertion of WTPA indicates that at least one entire cell buffer is available in that logical channel.</p> <p>If slave mode is selected, this signal is a tri-state output and the S/UNI APEX plays the roll of a single port UTOPIA L2 slave device driving the WTPA when the WTADR matches the programmed WANTxSlaveAddr register. A logical high indicates that at least one cell is available for transmission.</p> <p>WTPA is sampled/updated/Hi-Z'd on the rising edge of WTCLK.</p>
WTENB	Output (Master) Input (Slave)	AC20	<p>WAN Transmit Enable. WTENB indicates cell transfers to UTOPIA and SCI-PHY devices. The device is selected via a match on WTADR[2:0] when WTENB is last sampled high.</p> <p>This pin is in Hi-Z when the WAN transmit interface is not enabled.</p> <p>If master mode is selected, this signal is an output.</p> <p>If slave mode is selected, this signal is an input.</p> <p>WTENB is held low for the duration of the cell transfer.</p> <p>WTENB is sampled/updated on the rising edge of WTCLK.</p>

Pin Name	Type	Pin No.	Function
WTSX	Output	AE21	<p>WAN Transmit Start of Transfer. WTSX is asserted by the S/UNI APEX during the first cycle of a data block transfer. WTSX assertion will coincide with the port address prepend, if the cell being transferred has a prepended port address. Required only during Any-PHY mode.</p> <p>WTSX is updated on the rising edge of WTCLK.</p>
WTSOP	Output (Master) Tri-state (Slave)	AC19	<p>WAN Transmit Start of Packet. WTSOP marks the start of cell on the WTDAT[15:0] data bus. WTSOP is driven high when the first word of the cell (excluding address prepend) is present on the WTDAT[15:0] stream. WTSOP is asserted for each cell.</p> <p>In transmit master mode, the signal is always driven.</p> <p>In receive slave mode, this signal is driven 1 WTCLK after WTENB is asserted.</p> <p>WTSOP is updated/Hi-Z'd on the rising edge of WTCLK.</p>
WTDAT[0] .. WTDAT[15]	Output (Master) Tri-state (Slave)	AE22 AF23 AC21 AD22 AE23 AF24 AC24 AD26 AC25 AB24 AA23 AC26 AB25 AA24 Y23 AA25	<p>WAN Transmit Data. WTDAT[15:0] carries the data block transfers to the physical layer devices.</p> <p>In 8 bit mode, only WTDAT[7:0] are valid.</p> <p>In 8/16bit transmit master mode, the entire bus is always driven.</p> <p>In receive slave mode, this bus is driven 1 WTCLK after WTENB is asserted Pull up/downs are required for the entire bus, regardless of whether the bus is in 8 or 16 bit mode.</p> <p>WTDAT[15:0] is updated/Hi-Z'd on the rising edge of WTCLK.</p>

Pin Name	Type	Pin No.	Function
WTPRTY	Output (Master) Tri-state (Slave)	AF21	<p>WAN Transmit Parity. This signal provides programmable odd/even parity of the WTDAT[15:0] bus.</p> <p>In transmit master mode, the signal is always driven.</p> <p>In receive slave mode, this signal is driven 1 WTCLK after WTENB is asserted.</p> <p>WTPRTY is updated/Hi-Z'd on the rising edge of WTCLK.</p>

9.5 Context Memory Synchronous SSRAM Interface (60 Signals)

Pin Name	Type	Pin No.	Function
CMD[0] .. CMD[33]	BiDi	A24 B23 C22 D21 A23 B22 C21 A22 B21 C20 D19 A21 B20 C19 A20 B19 C18 D17 A19 B18 C17 A18 D16 B17 C16 B16 D15 A16 C15 B15 A15 C14 B14 B13	<p>Context Memory SSRAM Data. The bi-directional SSRAM data bus pins interface directly with the synchronous SSRAM data ports.</p> <p>The S/UNI APEX presents valid data on the CMD[33:0] pins upon the rising edge of SYSCLK during write cycles. CMD[33:0] is Hi-Z'd on the rising edge of SYSCLK for read cycles.</p> <p>CMD[33:0] is sampled/updated/Hi-Z'd on the rising edge of SYSCLK.</p>

Pin Name	Type	Pin No.	Function
CMP[0] .. CMP[1]	BiDi	D24 C23	<p>Context Memory SSRAM Data Parity. The SSRAM parity pins provide parity protection over the CMD[33:0] data bus.</p> <p>CMP[0] completes the odd parity for CMD[16:0] CMP[1] completes the odd parity for CMD[33:17] CMP[1:0] has the same timing as CMD[33:0]. The CMP[1:0] may be unconnected if parity protection is not required.</p> <p>CMP[1:0] is sampled/updated/Hi-Z'd on the rising edge of SYSCLK.</p>
CMA[0] .. CMA[19]	Output	K25 L23 J26 K24 J25 H26 K23 J24 H25 G26 H24 G25 F26 H23 G24 F25 E26 E25 F23 E24	<p>Context Memory SSRAM Address. The SSRAM address outputs identify the SSRAM locations accessed.</p> <p>The maximum size of the SSRAM is 4M bytes but smaller configurations are possible depending on the number of ports and VCs supported by the particular application. In the case where smaller SSRAM sizes are used the most significant bits of CMA[19:0] may be left unconnected.</p> <p>CMA[19:0] is updated on the rising edge of SYSCLK.</p>

Pin Name	Type	Pin No.	Function
CMAB[17] .. CMAB[18]	Output	F24 D26	<p>Context Memory SSRAM Address Bar. These active low address outputs are provided to enable glueless connection to 4 banks of ZBT SSRAM, or 2 banks of Late Write SSRAM.</p> <p>In ZBT SSRAM mode, these bits are the inverse of CMA[18:17].</p> <p>In Late Write SSRAM mode, CMAB[17] is the chip enable bar for even addresses, CMAB[18] is the chip enable bar for odd addresses.</p> <p>CMAB[18:17] is updated on the rising edge of SYSCLK.</p>
CMRWB	Output	C26	<p>Context Memory SSRAM Read Write Bar. CMRWB determines the cycle type when CMCEB is asserted low. When CMRWB is asserted high, the cycle type is a read. When CMRWB is asserted low, the cycle type is a write.</p> <p>CMRWB is updated on the rising edge of SYSCLK.</p>
CMCEB	Output	D25	<p>Context Memory SSRAM Chip Enable Bar. CMCEB initiates an access. When CMCEB is asserted low, the external SSRAM samples the address and CMRWB asserted by the S/UNI APEX.</p> <p>CMCEB is updated on the rising edge of SYSCLK.</p>

9.6 Cell Buffer SDRAM Interface (52 Signals)

Pin Name	Type	Pin No.	Function
CBCSB	Output	AE15	<p>Cell Buffer SDRAM Chip Select Bar. CBCSB, CBRASB, CBCASB, and CBWEB define the command being sent to the SDRAM.</p> <p>CBCSB is updated on the rising edge of SYSCLK.</p>
CBRASB	Output	AF15	<p>Cell Buffer SDRAM Row Address Strobe Bar. CBCSB, CBRASB, CBCASB, and CBWEB define the command being sent to the SDRAM.</p> <p>CBRASB is updated on the rising edge of SYSCLK.</p>

Pin Name	Type	Pin No.	Function
CBCASB	Output	AC14	Cell Buffer SDRAM Column Address Strobe Bar. CBCSB, CBRASB, CBCASB, and CBWEB define the command being sent to the SDRAM. CBCASB is updated on the rising edge of SYSCLK.
CBWEB	Output	AD14	Cell Buffer SDRAM Write Enable Bar. CBCSB, CBRASB, CBCASB, and CBWEB define the command being sent to the SDRAM. CBWEB is updated on the rising edge of SYSCLK.
CBA[0] .. CBA[11]	Output	AC15 AE16 AD16 AE17 AC16 AF18 AD17 AE18 AF19 AC17 AD18 AE19	Cell Buffer SDRAM Address. The Cell Buffer SDRAM address outputs identify the row address (CBA[11:0]) and column address (CBA[7:0]) for the locations accessed. CBA[11:0] is updated on the rising edge of SYSCLK.
CBBS[0] .. CBBS[1]	Output	AD15 AF16	Cell Buffer SDRAM Bank Select. The bank select signal determines which bank of a dual/quad bank Cell Buffer SDRAM chip is active. CBBS[1:0] is generated along with the row address when CBRASB is asserted low. CBBS is updated on the rising edge of SYSCLK.
CBDQM[0] .. CBDQM[1]	Output	AE13 AE14	Cell Buffer SDRAM Input/Output Data Mask. The data mask changes state from high to low when the SDRAM is enabled. These pins are held low during normal operation CBDQM is updated on the rising edge of SYSCLK.

Pin Name	Type	Pin No.	Function
CBDQ[0] .. CBDQ[31]	BiDi	AF3 AE4 AD5 AC6 AF4 AE5 AD6 AC7 AE6 AD7 AC8 AF6 AE7 AD8 AF7 AE8 AD9 AC10 AF8 AE9 AD10 AF9 AC11 AE10 AD11 AE11 AC12 AF11 AD12 AE12 AF12 AD13	<p>Cell Buffer SDRAM Data. The bi-directional Cell Buffer SDRAM data bus pins interface directly with the Cell Buffer SDRAM data ports.</p> <p>The Cell Buffer SDRAM is accessed as a burst of 32-bit long words.</p> <p>CBDQ[31:0] is updated/Hi-Z'd on the rising edge of SYSCLK.</p>

9.7 Microprocessor Interface (44 Signals)

Pin Name	Type	Pin No.	Function
BCLK	Input	C1	<p>Bus Clock. This clock is the bus clock for the microprocessor interface. BCLK must cycle at 66 MHz or lower instantaneous rate.</p>

Pin Name	Type	Pin No.	Function
AD[0] .. AD[31]	BiDi	C13 D13 A12 C12 A11 D12 B11 C11 B10 D11 A9 C10 B9 A8 D10 C9 B8 A7 C8 B7 A6 D8 C7 B6 A5 C6 B5 A4 D6 C5 B4 A3	<p>Multiplexed Address Data Bus. The multiplexed address data bi-directional bus AD[31:0] is used to connect the S/UNI APEX to the microprocessor.</p> <p>During the address phase when ADSB = 0, AD[1:0] are ignored as all transfers are 32 bits wide.</p> <p>AD[31:0] is sampled/updated/Hi-Z'd on the rising edge of BCLK.</p>
ADSB	Input	F2	<p>Address Status. This signal is active-low and indicates a long-word address is present on the address/data bus AD[31:2].</p> <p>Address space used is 0->4K. Attempts to access above this address space is prohibited.</p> <p>ADSB is sampled on the rising edge of BCLK.</p>
CSB	Input	E1	<p>Active Low Chip Select. The chip select (CSB) signal is low during the address cycle (as defined by ADSB) of S/UNI APEX register accesses.</p> <p>CSB is sampled on the rising edge of BCLK.</p>

Pin Name	Type	Pin No.	Function
WR	Input	F3	<p>Write/Read. The write/read (WR) signal is evaluated when the ADSB and CSB are sampled active by S/UNI APEX. The BUSPOL input pin controls the polarity of this input.</p> <p>WR is sampled on the rising edge of BCLK.</p>
BURSTB	Input	E2	<p>Burst Bar. This signal is evaluated when the ADSB and CSB are sample active by S/UNI APEX. When low, this signal indicates that the current access is a burst access (and the BLAST input can be used to detect the end of the transaction).</p> <p>BURSTB is sampled on the rising edge of BCLK.</p>
BLAST	Input	D1	<p>Burst Last. This signal indicates the last data access of the transfer. When the BURSTB input is low, the BLAST input is driven active during the last transfer of a transaction (even if the transaction is one word in length). When the BURSTB input is high, the BLAST input is ignored by S/UNI APEX. The BUSPOL input pin controls the polarity of this input.</p> <p>BLAST is sampled on the rising edge of BCLK.</p>
READYB	Tri-state	F4	<p>Ready Bar. This signal is asserted low by S/UNI APEX when the data on the AD[31:0] bus has been accepted (for writes), or when the data on the AD[31:0] is valid (for reads). This signal may be used by S/UNI APEX to delay a data transaction. This output is Hi-Z'd one clock cycle after an S/UNI APEX access, allowing multiple slave device to be tied together in the system. This output should be pulled up externally.</p> <p>READYB is updated on the rising edge of BCLK.</p>

Pin Name	Type	Pin No.	Function
BTERMB	Tri-state	E3	Burst Terminate Bar. This signal is asserted low by S/UNI APEX when a data transfer has reached the address boundary of a burstable range. Attempts to extend the burst transfer after this signal is asserted will be ignored. This output is Hi-Z'd one clock cycle after an S/UNI APEX access, allowing multiple slave device to be tied together in the system. This output should be pulled up externally. BTERMB is updated on the rising edge of BCLK.
WRDONEB	Output	D2	Write Done Bar. This signal is asserted low by S/UNI APEX when the most recent write access to internal registers is complete. This signal may be used by external circuitry to delay the issuance of a write operation address cycle until S/UNI APEX can accept write data. This signal is only needed in systems where the READYB output cannot be used to delay a write data transaction (due to microprocessor restrictions). WRDONEB is updated on the rising edge of BCLK.
INTHIB	OD	C4	Active Low Open-Drain High Priority Interrupt. This signal goes low when an S/UNI APEX high priority interrupt source is active and that source is unmasked. The S/UNI APEX may be enabled to report many alarms or events via interrupts. INTHIB becomes high impedance when the interrupt is acknowledged via an appropriate register access. INTHIB is an asynchronous signal.
INTLOB	OD	D3	Active Low Open-Drain Low Priority Interrupt. This signal goes low when an S/UNI APEX low priority interrupt source is active and that source is unmasked. The S/UNI APEX may be enabled to report many alarms or events via interrupts. INTLOB becomes high impedance when the interrupt is acknowledged via an appropriate register access. INTLOB is an asynchronous signal.

Pin Name	Type	Pin No.	Function
BUSPOL	Input	E4	<p>Bus Control Polarity. This signal indicates the polarity of the WR and BLAST inputs to S/UNI APEX.</p> <p>When high, the BLAST pin is active high (high indicates the last word of the burst) and the WR pin is active low (low indicates write).</p> <p>When low, the BLAST pin is active low (low indicates the last word of the burst) and the WR pin is active high (high indicates write).</p> <p>BUSPOL is sampled on the rising edge of BCLK.</p>

9.8 General (9 signals)

Pin Name	Type	Pin No.	Function
RSTB	Input	AA4	<p>Reset Bar. This signal provides an asynchronous S/UNI APEX reset. RSTB is a Schmitt triggered input with an internal pull-up resistor.</p>
OE	Input	AC1	<p>Output Enable OE is an active high signal, which allows all of the outputs of the device to operate in their functional state. When this signal is low, all outputs of the S/UNI APEX are Hi-Z'd, with the exception of TDO.</p> <p>OE has an internal pull up resistor.</p>
SYSCLK	Input	B12	<p>System Clock. This clock is the master clock for the S/UNI APEX device. All non-Any-PHY or microprocessor interface related internal synchronous logic is timed to this signal. SYSCLK must cycle at a 80 MHz or lower instantaneous rate. External SSRAM and SDRAM devices share this clock and must have clocks aligned within 0.2ns skew of the clock seen by the S/UNI APEX device.</p> <p>This clock must be stable prior to deasserting RSTB 0->1.</p>

Pin Name	Type	Pin No.	Function
NC		AB4 AC5 AB23 E23 D22 D5	No Connect. These balls are not connected to the die.

9.9 JTAG & Scan Interface (7 Signals)

Pin Name	Type	Pin No.	Function
TCK	Input	AD1	Test Clock. This signal provides timing for test operations that are carried out using the IEEE P1149.1 test access port.
TMS	Input	AC2	Test Mode Select. This signal controls the test operations that are carried out using the IEEE P1149.1 test access port. TMS is sampled on the rising edge of TCK. TMS has an integral pull-up resistor.
TDI	Input	AB3	Test Data Input. This signal carries test data into the S/UNI APEX via the IEEE P1149.1 test access port. TDI is sampled on the rising edge of TCK. TDI has an integral pull-up resistor.
TDO	Tri-state	AD4	Test Data Output. This signal carries test data out of the S/UNI APEX via the IEEE P1149.1 test access port. TDO is a tri-state output, which is inactive except when scanning of data is in progress. TDO is updated/Hi-Z'd on the falling edge of TCK.
TRSTB	Input	AC3	Active low Test Reset. This signal provides an asynchronous S/UNI APEX test access port reset via the IEEE P1149.1 test access port. TRSTB is a Schmitt triggered input with an integral pull-up resistor. Note that when not being used, TRSTB must be connected to the RSTB input.

Pin Name	Type	Pin No.	Function
SCANEN	Input	AD23	Scan Enable This signal enables the internal scan logic for production testing. Should be held to its inactive low state.
SCANMB	Input	AC22	Scan Mux This signal is connected directly to the control of the internal scan muxes. Should be held to its inactive high state.

9.10 Power

Pin Name	Type	Pin No.	Function
VDD	Power	AC4 AC9 AC13 AC18 AC23 AD3 AD24 AE2 AE25 B2 B25 C3 C24 D4 D9 D14 D18 D23 J23 J4 N4 P23 V4 V23	The pad ring power pins should be connected to a well de-coupled +3.3 V DC.

Pin Name	Type	Pin No.	Function
PCH	Power	G4 L4 R4 W4 AF5 AF10 AF17 AF22 AB26 U26 K26 G23 D20 A17 A10 D7	The core power pins should be connected to a well-decoupled +2.5 V DC.
VSS	Ground	A1 A2 A13 A14 A25 A26 B1 B3 B24 B26 C2 C25 N1 N26 P1 P26 AD2 AD25 AE1 AE24 AE3 AE26 AF1 AF2 AF13 AF14 AF25 AF26	The pad ring and core ground pins should be connected to GND.

10 FUNCTIONAL DESCRIPTION

This section describes the function of each entity in the S/UNI APEX block diagram. In this document, receive and transmit are used with the S/UNI APEX as the frame of reference. For example, receive is used to describe data paths which are coming into the device.

10.1 Any-PHY Interfaces

The S/UNI APEX Interface are Any-PHY compliant 8/16-bit master/slave interface for both Loop and WAN ports. The loop and WAN interfaces are configured independently. Both interfaces are fully compatible with the following Any-PHY options:

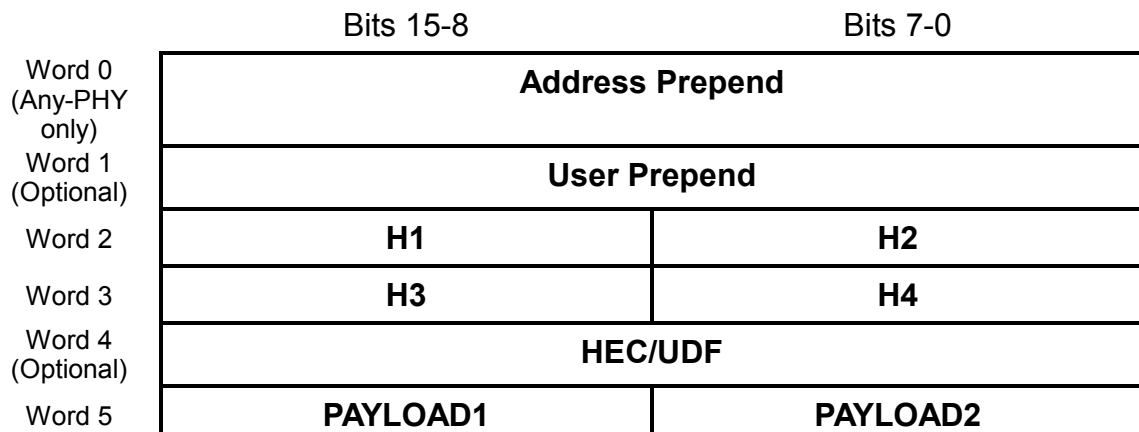
- Any-PHY master.
- UTOPIA L2 master (UL2M).
- UTOPIA L1 master (UL1M).
- UTOPIA L2 slave (UL2S).

10.1.1 Receive Interface

The S/UNI APEX requires a 16-bit Ingress Connection Identifier (ICI) to be received with every cell. The ICI uniquely identifies the VCC or VPC. The ICI can be received within the HEC/UDF field (16bit I/F only), as a user prepend, or encoded within the VPI/VCI field. In Any-PHY mode, an address prepend is expected to be in the first word/byte of every cell. Inclusion of optional words/bytes are statically configured for the interface.

The Receive Cell Transfer Format is shown in Figure 3 and Figure 4.

Figure 3 - 16bit Receive Cell Transfer Format




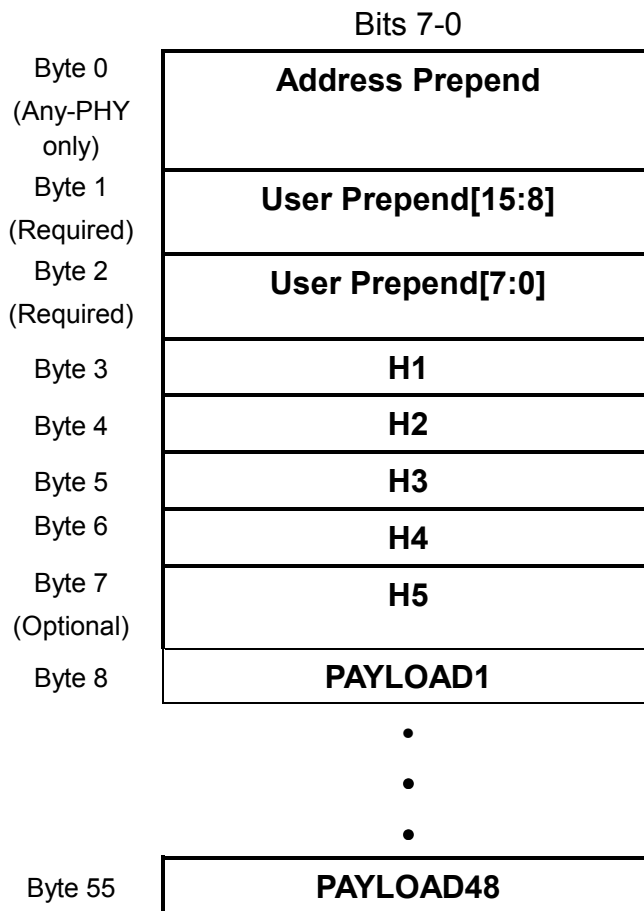


Figure 4 - 8-bit Receive Cell Transfer Format



The Loop and WAN receive master mode interface supports per-device or per-port RPA (Receive Packet Available) status polling via round robin polling address enabling support for up to 32 loop or 4 WAN devices and/or ports. Polling ceases once a device or port has been identified as having a cell available. Polling recommences on the following address that was serviced. Since the S/UNI APEX requires a unique 16-bit ICI with every cell, knowledge of which polling addresses are associated with devices and which are associated with ports is not required.

If UL2M or Any-PHY, and the number of ports connected is less than 32 (loop) or 4 (WAN), there is an option of limiting the polling range; thereby providing optimal polling efficiency.

The UL1M is effectively a UL2M without address polling, but retains the port selection handshake. Hence a single external UL2S may be connected to the S/UNI APEX UL1M transmit interface.

If Any-PHY, the S/UNI APEX expects the Any-PHY slave device to act as a proxy for its internal ports. The S/UNI APEX places no restrictions on the number of internal ports within an Any-PHY slave device. Since the polling is tied to the data transfer, both the WAN and loop Any-PHY receive interface is capable of mixing prepend enabled UL2 and Any-PHY slaves on the same bus with some external glue logic.

If UL2S, the S/UNI APEX operates as a single port UTOPIA L2 transmit slave port. The address pins become inputs and can be configured to respond to any port identifier from 0 to 31 for loop, and 0 to 3 for WAN.

Table 4 - Number of Ports Supported, Receive Interface

Mode	Loop (8/16bit)	WAN (8/16bit)
Any-PHY Master	32	4
UTOPIA L2 Master	32	4
UTOPIA L1 Master	1	1
UTOPIA L2 Slave	1 of 32	1 of 4

10.1.2 Transmit Interface

The Transmit Cell Transfer Format is shown in Figure 5 and Figure 6. Word/byte 0 is required for cell transfers to Any-PHY slaves. The address prepend is the S/UNI APEX port id associated with the transmit queue in which the cell was enqueued. The unused bits in the address prepend are reserved and devices should not rely on the content. Optional word 1 or bytes {1,2} enables the prepending of a 16-bit switch tag. Optional word 2 or bytes {3,4} enables the prepending of a 16-bit Egress Connection Identifier (ECI). Both the Switch tag and the Egress Connection Identifier are sourced on a per-VC basis from VC context. The S/UNI APEX also maps the ECI tag to the HEC/UDF field (word 5) for 16-bit transfer. Word 5 or byte 9 is optional. The S/UNI APEX supports optional VPI and/or VCI mapping, selectable on a per VC basis.


Inclusion of optional words is statically configurable for the interface. Selection of the usage of the included optional words is configurable on a per-VC basis. On a

per-VC basis, either mapping of switch tag and/or ECI or mapping of the switch tag and VPI, or of VPI and VCI is supported. If optional words 2, and/or 5 are included on the interface, they contain the original ICI if ECI remapping is not supported for the VC. If optional word 1 is included on the interface, it is defined as a reserved field for those VCs that are not mapping the switch tag.

Figure 5 - 16-bit Transmit Cell Transfer Format

	Bits 15-8	Bits 7-0
Word 0 (Any-PHY only)	Address Prepend Loop I/F: {4 MSB reserved, PortID[11:0]} WAN I/F: {14 MSB reserved, PortID[1:0]}	
Word 1 (Optional)	Switch Tag Prepend	
Word 2 (Optional)	ECI Prepend	
Word 3	H1	H2
Word 4	H3	H4
Word 5 (Optional)	ECI Prepend	
Word 6	PAYLOAD1	PAYLOAD2
Word 7	PAYLOAD3	PAYLOAD4
	•	•
	•	•
	•	•
Word 29	PAYLOAD47	PAYLOAD48

Figure 6 - 8-bit Transmit Cell Transfer Format



In the loop interface Any-PHY mode, 16bit, per-port status polling is supported via a 12 bit polling address bus and a single transmit packet available input enabling up to 2048 port polling. 8-bit loop interface is limited to an 8-bit polling address, enabling 256 port polling. The loop interface polling is completely independent of the data transfer.

In the WAN interface Any-PHY mode, 8/16bit, per-port status polling is supported via a 3 bit polling address bus and a single transmit packet available input enabling up to 4 port polling. The WAN interface polling ceases once a device or port has been identified as having a cell available. WAN polling recommences on the following address that was serviced. Since the polling is tied to the data

transfer, the WAN transmit interface is capable of mixing prepend enabled UL2 and Any-PHY slaves on the same bus with some external glue logic.

In UL2M, loop interface port selection is done via the 6 lower bits of the 12-bit polling address bus, supporting up to 32 ports. WAN interface port selection is done via the 3-bit polling address bus, supporting up to 4 ports.

Details of the polling algorithm for the loop and WAN interface can be found in the loop port scheduler and WAN port scheduler section respectively.

The UL1M is effectively a UL2M without address polling, but retains the port selection handshake. Hence a single external UL2S may be connected to the S/UNI APEX UL1M transmit interface. Specific only to the WAN UL1M mode, port address is presented with a programmable value, giving the option of port sparing.

In slave mode, the transmit interface operates as a single port UTOPIA L2 receive slave port. The 6 lower bits of the 12-bit loop polling address, or the entire 3 bits of the WAN polling address become inputs. The loop interface can be configured to respond to any port identifier from 0 to 31. The WAN interface can be configured to respond to any port identifier from 0 to 3.

Table 5 - Number of Ports Supported, Transmit Interface

Mode	Loop (8 bit)	Loop (16 bit)	WAN (8 bit)	WAN (16bit)
Any-PHY Master	256	2048	4	4
UTOPIA L2 Master	32	32	4	4
UTOPIA L1 Master	1 (no sparing)	1 (no sparing)	1 (4 sparing)	1 (4 sparing)
UTOPIA L2 Slave	1 of 32	1 of 32	1 of 4	1 of 4

10.2 Loop Port Scheduler

The S/UNI APEX loop port scheduler provides weighted interleaved round robin scheduling of up to 2k Any-PHY addresses. To achieve fairness among the 2k ports and to avoid wasted polling opportunities, the selection of what ports to poll is based on what ports have transmit data queued and have a high probability of being able to accept the cell.

The scheduler has 128 polling sequences and 8 different weighting groups. Software configures the number of polling sequences a port should participate in by assigning a 3-bit logarithmic weight value and a 7-bit sequence number to

each port. The scheduler maintains a 7-bit polling sequence number and increments it after each scheduler polling cycle. During a scheduler polling cycle each of the 2k ports is evaluated. The port will be polled if the following conditions are met:

- the port's transmit data queue is not empty
- the n LSB's of the scheduler poll sequence number match the n LSB's of the port's sequence number (n is equal to the port's weight). For ports with a weight of zero, this compare is ignored. For ports with a weight of one, then only the LSB is compared. For ports with a weight of seven, then the entire seven bits are compared.

To maintain even distribution of ports within the same weight class, software must assign sequence numbers to ports evenly across the 128 polling sequences. This sequence number need only be changed when a port's weight is changed or the distribution ports in a weight group becomes significantly unbalanced due to port deactivations. Sequence numbers and weights may be modified at any time.

The logarithmic weights are set so that lower speed ports are evaluated less often relative to higher speed ports. The following formula show relationship between the 3 bit logarithmic weights (lw) and the assigned relative throughput weight (rw) in the case where the aggregate throughput of all the ports is greater than the available bandwidth:

$$rw = 2^{(7-lw)}$$

The maximum polling rate for any given port is dictated by the number of active ports. In Any-PHY mode, if only one port is active for all 2k ports (port's transmit data queue is not empty), the maximum polling rate is governed by the following formula:

$$\text{Max. polling rate} = f(\text{SYSCLK}) / (64 * 2^{lw})$$

The equivalent equation for UL2M mode is the following:

$$\text{Max. polling rate} = f(\text{SYSCLK}) / 2^{lw}$$

10.3 Wan Port Scheduler

The WAN port scheduler operates between the queue engine and the multi-channel WAN port FIFO. The S/UNI APEX WAN port scheduler provides weighted interleaved round robin scheduling of up to 4 WAN ports. The dynamic range of the weights is 8 to 1.

The scheduler has 8 polling sequences and 4 different weighting groups. Port weighting is achieved by configuring the number of polling sequences a port should participate in. This configuration is done by assigning a 2-bit logarithmic weight value and a 3-bit sequence number to each port. Software assigns the 2-bit weight value, and the hardware always maps the 4 ports to the following sequence numbers: port 0 is assigned 000, port 1 is assigned 010, port 2 is assigned 101, and port 3 is assigned 111. The scheduler maintains a 3-bit polling sequence number and increments it after each scheduler polling cycle. During a scheduler polling cycle each of the 4 ports is evaluated. The port will be selected for transmission if the following conditions are met:

- the port's transmit data queue is not empty
- the n LSB's of the scheduler poll sequence number match the n LSB's of the port's sequence number (n is equal to the port's weight). For ports with a weight of zero, this compare is ignored and assumed successful. For ports with a weight of one, only the LSB is compared. For ports with a weight of two, only the first two LSBs are compared. For ports with a weight of three, all three bits are compared.
- the S/UNI APEX internal WAN FIFO for the port is not full

The logarithmic weights are set so that lower speed ports are evaluated less often relative to higher speed ports. The following formula shows relationship between the logarithmic weights values and the resulting linear relative weight.

$$rw = 2^{(3-lw)}$$

If port 0 were assigned a weight of 0, port 1 a weight of 1, port 2 a weight of 2, and port 3 a weight of 3, and all the ports had data to send, and none of the WAN FIFOs were full, then cells would be transmitted in the following order:

Table 6 - Example WIRR Transmission Sequence

Sequence Number	Ports Transmitted
000	0,1
001	0,2
010	0,1
011	0
100	0,1
101	0,2
110	0,1
111	0,3

The above example was constrained by several conditions under which the queue engine WAN interface bandwidth was the transmission bottleneck. In the WAN transmit datapath, there are actually three places where a transmission bottleneck can occur: the queue engine's WAN interface bandwidth, the Any-PHY bus, and the actual physical lines.

If the queue engine's WAN interface bandwidth is the bottleneck, then the WIRR WAN scheduler will determine the transmission order. In this case, the queue engine's WAN interface does not have enough bandwidth to service all of the physical lines and each physical line will receive a weighted proportion the queue engine's available WAN bandwidth.

If the Any-PHY bus becomes the bottleneck, then a simple round robin scheduler at the Any-PHY interface will determine the transmission order. For this reason, the system designers should ensure that the Any-PHY bus does not become the bottleneck.

Finally, if the physical lines are the bottleneck, then the physical line rates and the WIRR WAN scheduler will determine the transmission order. This last situation is the most desirable one because in this case no transmission opportunities will ever be missed.

10.4 WAN Port Aliasing

For each of the four channels, a port aliasing register is provided to allow for port sparing for the uplinks. These registers map the internal VC's PortID to the external Any-PHY address. By having this layer of indirection, it is possible to re-direct all traffic from one Any-PHY address to another by modifying a single register, and without having to change any per-VC context information.

10.5 WAN and Loop ICI Selection

The S/UNI APEX requires that an ICI be received with every cell. A connection identified by a single ICI may be either a VCC or a VPC connection. The ICI may be prepended to the cell or embedded in the VPI/VCI header for interfacing to devices that cannot add prepends to the cell.

The S/UNI APEX accepts cells from the following sources: WAN ports, loop ports, and the microprocessor port. Each cell is directed to a particular connection, which is identified by an ICI. For cells from the microprocessor port, the ICI is given directly. For the WAN and Loop ports, this ICI may be selected from one of several locations within the cell and is programmable per interface. The ICI may be:

- A two byte user tag prepended to the cell.
- The two byte HEC/UDF field of the cell.
- Embedded in the 12 bit VPI & 16 bit VCI field as defined as follows:

If the VPI < "FFF" then

ICI = "0" & VPI; -- This connection is a VPC connection.

else

ICI = VCI; -- This connection is a VCC connection.

end if;

In an UNI environment, the S/UNI APEX considers the 4 bit GFC field plus the 8 bit VPI field as the VPI field described above.

10.6 Microprocessor Interface

The microprocessor interface supports the following features:

- 32-bit wide multiplexed address data bus.
- Synchronous microprocessor interface supporting linear bursts of up to 16-long words for cell transfers, up to 5 long words for performance sensitive context memory, and single long word accesses for registers and remaining context memory.

- Microprocessor clock independent of the system clock, allowing for easy integration into any host system without altering device performance.
- Addressing:
 - Direct addressing for internal control and status, SAR assist, CBI register port and memory port.
 - Indirect addressing (via the Memory Port) for context memory accesses.
 - Indirect addressing (via the CBI Register Port) for CBI register accesses.
- High and Low Priority Interrupt outputs provided for efficient task management.
- Bus Polarity Select pin provided to allow interconnect between the S/UNI APEX and PowerPC or i960 microprocessors.
- Write Done Indicator output provided to allow interconnect with the IDT MIPS microprocessor (with minimal external logic for system command generation and interpretation).

The microprocessor interface receives a multiplexed address and data bus, where an address strobe input defines the address cycle. During the address cycle, the bus contains the address for the beginning of the transaction. Also during this cycle, the chip select, write indicator, and burst indicator are latched to define the transaction. The interpreted polarity of the write indicator and burst indicator are controlled by a single configuration input pin, for compatibility with multiple microprocessors such as the PowerPC or the i960.

If a read transaction is indicated at the address cycle, then S/UNI APEX will respond with a ready indicator concurrent with each long word of valid data, until the burst is complete. The delay between the address cycle and the first valid long word of read data is variable, depending on the specific register address (not less than 2 clock cycles). If a read transaction is issued to the receive SAR when no data is available, or issued to the memory port when the current command is not yet complete, the first word of valid read data will be delayed until data is available (this can be many clock cycles). If excessive delay for the first word of valid read data cannot be tolerated, then polling (or interrupt processing) must be used for accesses to these regions. The ready indicator may be deasserted by S/UNI APEX in the middle of a burst read operation to allow for read data synchronization delay.

If a write transaction is indicated at the address cycle, then S/UNI APEX will respond with a ready indicator concurrent with each long word of valid data, until

the burst is complete. The delay between the address cycle and the first valid long word of write data is variable, depending on the specific register address (not less than 1 clock cycle). If a write transaction is issued to the transmit SAR when the buffer is full, or issued to the memory port when the current command is not yet complete, the first word of valid write data will be delayed by the ready indicator until buffer space is available (this can be many clock cycles). If excessive delay for the first word of valid write data cannot be tolerated, then polling (or interrupt processing) must be used for accesses to these regions. Once the ready indicator has been asserted, it will remain asserted until the completion of the burst.

An additional output is provided to indicate when the current write operation is complete (write done indicator). Processors which do not allow the ready indicator to be used to delay the advance of write data, but do allow a write operation to be delayed before it is issued (such as the IDT MIPS processor) may use this output. The write done indicator is asserted when S/UNI APEX can accept another write command. Typically, an external circuit may be employed which uses this S/UNI APEX output to determine when to allow the processor to issue another write command. When this output is used prior to the address cycle, the normal ready indicator need not be used for write operations, as S/UNI APEX can accept write data always once the write done indicator is asserted (unless polling of buffer status is disabled). Note that polling of buffer status must be employed when the processor does not allow the ready indicator to be used to delay the advance of write data.

If a burst is indicated at the address cycle, then the transaction will not complete until the processor asserts the burst last indicator. If a burst is not indicated, then the transaction will be completed after the ready indicator is asserted by S/UNI APEX.

The multiplexed address/data bus will be Hi-Z'd immediately following the last word of read data to allow a new address cycle to commence. The microprocessor interface will allow an address cycle to occur with no wait states between the last word of valid data and the new address; however, care must be taken to minimize bus contention in the system design if no wait states are provided by the microprocessor.

The diagrams below illustrate possible connections between the APEX and various microprocessors. For the i960 interface, the two lower order bits of the address may be tied to ground as all accesses to the APEX are 32bits wide.

Figure 7 - i960 (80960CF) Interface


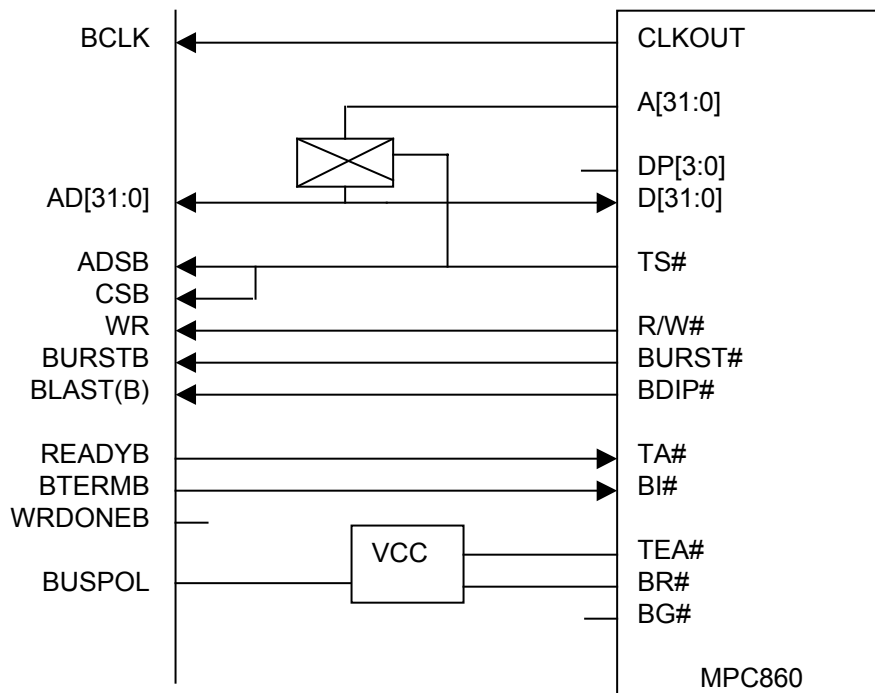


Figure 8 - PowerPC (MPC860) Interface



10.7 Memory Port

Much of the configuration information that S/UNI APEX requires for normal operation is accessed indirectly through the memory port, as the configuration storage is tightly coupled to performance. Register arrays are provided to allow access to the following memory apertures:

- External Queue context
- Internal Queue context
- Internal Loop context

The memory port is primarily used for context setup, but may also be used for diagnostic purposes. Features include

- Control register allows the microprocessor to specify the aperture, address, and length of the burst. Access to the internal loop context are restricted to single long word accesses.
- 4-word burst write buffer with 8-bit overflow register, supporting writes of up to 4 contiguous 34-bit words to valid apertures.
- Masked write mechanism, which can be used to overwrite specific bits of 1 word without affecting other bits.
- 4-word burst read buffer with 8-bit overflow register, supporting reads of up to 4 contiguous 34-bit words from valid apertures.
- Memory port status provided in the low priority interrupt status register, allowing for polling or for interrupt driven accesses to memory.

Memory is accessed using a 4-long word address in the control register, along with 4 long-word enables. This approach allows non-contiguous bursts within a 4-long word section of memory, or to specify which long word is to be accessed in single long word transfer. (For example, the first and third word of a section may be modified without changing the second and fourth).

To compensate for the difference between the 34-bit context memory bus and the 32-bit microprocessor bus, an 8-bit overflow register is provided for both reads and writes. The overflow register represents the most significant 2 bits of up to 4 words in a burst access. In this manner, 4 34-bit words can be accessed using a 5-word burst on the microprocessor bus.

The masked write mechanism is provided to allow the microprocessor to change a field within a word in context memory while traffic is present, without risk of context corruption. The masked write can be performed on one word per operation. In this mode (as indicated in the control register), the second word in the 4-word burst write buffer and the second pair of bits in the overflow register represent a bit mask which will be used by S/UNI APEX to perform a masked write function.

10.8 SAR Assist

The SAR assist module allows cells or AAL5 frames to be transferred to and from the queue engine. Burst transfers from the microprocessor into and out of the SAR staging buffers enable efficient access to the queuing structures. The staging buffers are organized as 64 byte units, including the ICI/ECI, the cell header, the payload, and control or status information. A complete buffer can be written or read in one continuous burst, or the data can be accessed individually or with a series of shorter bursts. Within this structure, both the cell header and the payload are aligned on 32-bit boundaries, to simplify microprocessor access. The SAR assist module can also optionally perform calculation, checking, and insertion of AAL5 CRC32 or CRC10.

One staging buffer is provided for cell or frame injection, while four staging buffers are provided for cell or frame reception (one for each microprocessor class queue).

10.8.1 Transmit

The transmit function of the SAR has the following features:

- Read staging buffer for each of the 4 class queues associated with the microprocessor.
- CRC-32 checking for AAL5 re-assembly.
- Simultaneous re-assembly assist on all 4 class queues.
- CRC-10 checking for OAM.
- Cell header is provided with each PDU, including PTI for end-of-message detection by the microprocessor.

Each read buffer represents a 2-cell pipeline, providing minimum latency for cell retrieval. While a cell is read out, a second cell is retrieved from the queue engine automatically. By having read buffers for each class, the microprocessor can decide which class has the highest priority. The microprocessor can

interrupt the retrieval of a frame from one class to transmit a higher priority packet/cell from another class without impacting the CRC32 calculation.

The SAR will accept another cell from the queue engine when the 14th long word of the transmit buffer has been read. The SAR assist transmit cell transfer format is shown in Figure 9.

Figure 9 - SAR Assist Transmit Cell Transfer Format

Register	Bits 31-24	Bits 23-16	Bits 15-8	Bits 7-0
SarTxData0	CRC Status		SarTxECI	
SarTxData1	H0	H1	H2	H3
SarTxData2	Payload1	Payload2	Payload3	Payload4
	•	•	•	•
	•	•	•	•
	•	•	•	•
SarTxData13	Payload45	Payload46	Payload47	Payload 48

The SAR performs CRC32 error checking over the entire frame. The CRC32 accumulator for a class is automatically reset on frame boundaries or when a non-user or WFQ cell is encountered (see Table 8). A CRC32 status bit is updated as the EOM cell enters the read buffer.

The SAR performs CRC10 error checking over an OAM cell. A CRC10 status bit is updated as the OAM cell enters the read buffer. The processor should verify the cell type (via the cell header) when determining the validity of these status bits.

All the CRC status bits in the buffer are updated prior to indicating data is available. Should a CRC error be detected, the microprocessor can skip reading the cell's entire payload and move on to the next cell by reading the 14th word of the transmit buffer.

10.8.2 Receive

The receive function of the SAR has the following features:

- Single write staging buffer
- ICI (Ingress Connection Identifier) prepended to all cells
- Option to overwrite the end of a cell with AAL5 CRC32 or OAM CRC-10

The single write buffer represents a 2-cell pipeline, allowing the microprocessor to fill one payload while the other one is waiting to be queued. A Not Full status bit is provided, indicating whether the write buffer is capable of accepting at least one cell.

Cell enqueueing is initiated by writing to the 14th word of the receive buffer. The SAR assist receive cell transfer format is shown in Figure 10.

Figure 10 - SAR Assist Receive Cell Transfer Format

Register	Bits 31-24	Bits 23-16	Bits 15-8	Bits 7-0
SarRxLWord0	CRC Control		SarRxICI	
SarRxLWord1	H0	H1	H2	H3
SarRxLWord2	Payload1	Payload2	Payload3	Payload4
	•	•	•	•
	•	•	•	•
	•	•	•	•
SarRxLWord13	Payload45	Payload46	Payload47	Payload 48

Once there are 2 cells in the process of being en-queued, any further attempts to write to the write buffer will be held pending until the first cell has been en-queued.

The CRC Control gives each cell the option of being overwritten with an AAL5 CRC-32 or an OAM CRC-10 trailer. These CRC values cannot be invoked if OAM cells are interspersed within AAL5 packets.

For frame traffic, it is necessary to write SarRxLWord0&1 for the first two cells, SarRxLWord0 for the third cell and SarRxLWord0 for the last cell of the frame. SarRxLWord0&1 write of the first cell is required to reset the CRC, and establish the ICI and header for the first pipe. SarRxLWord0&1 write of the second cell is required to set the CRC for normal operation, and establish the ICI and header for the second pipe. SarRxLWord0 write of the third cell is required to remove the reset of the CRC established in the first cell and set the CRC for normal operation. SarRxLWord0 write of the last cell is required to concatenate the CRC onto the end of the cell. The middle cells of the frame only require the payload to be updated.

10.9 Queue Engine


The queue engine performs the following functions:

- Service Arbitration
- Congestion Control
- Statistics
- Cell Queuing (VC Scheduling)
- Class Scheduling
- Watch Dog: VC time out patrol and re-allocation
- Microprocessor queue buffer re-allocation

10.9.1 Service Arbitration

There are 9 components that request services from the queue engine. Three components (SAR Rx, WAN Rx, and Loop Rx) can request a cell to be en-queued. Another three components (SAR Tx, WAN Tx, and Loop Tx) can request a cell to be de-queued. The shaper, if enabled, can request the transmission slots to be advanced (see section on Shape Fair Queuing). There are two possible requests from the watch dog, one to patrol a range of VC queues to detect a timed out VC, and another request to re-allocate buffers from a VC that has timed out. The uP can request a VC or Class queue to have their buffers re-allocated and removed from service. The queue engine is capable of simultaneously servicing any one or all of the requests from an en-queue component, a de-queue component, watch dog patrol and transmission slot advancement. The queue engine is capable of servicing the re-allocation of buffers from either the uP or watch dog alone. To resolve all these requests, there are four arbitration units. See Figure 11.

Figure 11 - Service Arbitration Hierarchy



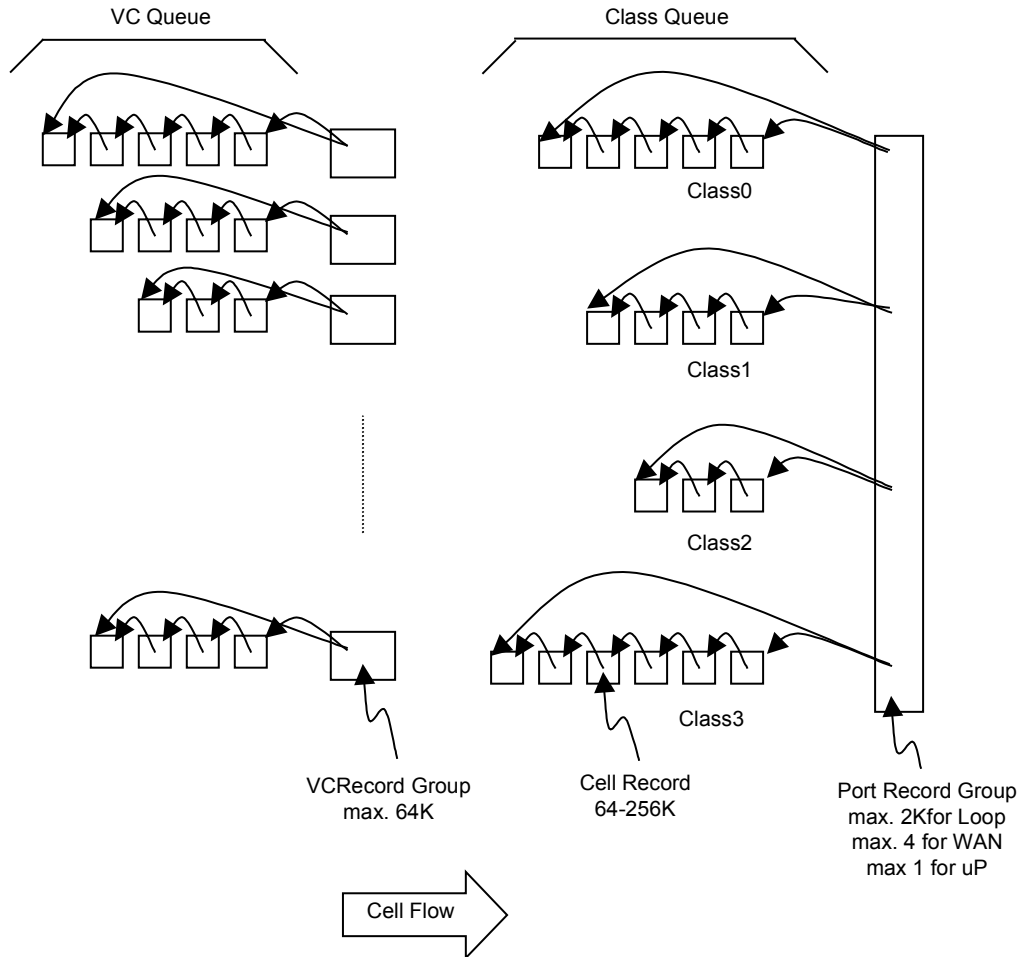
There is a Rx arbiter that receives requests to queue a cell from the SAR, loop and WAN Rx interfaces. There is a Tx arbiter that receives requests to de-queue a cell from the SAR, loop and WAN Tx interfaces. These Rx and Tx arbiters have two options for arbitration. The default option is to have the arbiters use round robin to select between the three interfaces. The alternate option is to have the arbiters use round robin between the loop and the WAN interfaces, with the SAR set to the lowest priority. The results of the Rx and Tx arbiters, along with the request from the watch dog patrol and the shaper transmission slot advancement, are OR'd together to represent a single request from the "queue arbiter".

There is a round robin arbiter that receives re-allocation requests from the watch dog and uP. The results of this arbiter, and the one from the "queue arbiter" goes to the final round robin arbiter.

10.9.2 Cell Queuing

After congestion control, a cell will be queued onto a linked list structure. The structure is made up of context records, on a per-Port, per-Class, and per-VC basis. Context records are stored in both the external SSRAM and internal RAM. Figure 12 below illustrates the structure of the linked lists, and the relationships between the different context records.

Figure 12 - Queue Linked List Structure



Note: The class queue and VC queue as illustrated in the above diagram cannot be directly correlated with the per-Class and per-VC levels as defined in the congestion control.

The rules for queuing, and the way the linked lists are utilized is configured on a per-VC basis. A VC may be configured to one of three mutually exclusive queuing procedures. In addition, the queuing of non user cells may be handled differently. The available queuing procedures as a function of the port destination are outlined in Table 7.

Table 7 - Available Queuing Procedures

	Loop	WAN	uP
Weighted Fair Queuing	x	x	x
Frame Continuous Queuing	x	x	x
Shape Fair Queuing		x	
Non User Cell Queuing	x	x	x

10.9.2.1 Weighted Fair Queuing

Weighted fair queuing is available to cells destined for the loop, uP and WAN ports. It is configured on a per-Class basis. The VC and class queue work together to provide weighted fair queuing. The class queue is a staging area for cells from different VCs to be lined up for their final destination. The VC demographics in the class queue are defined by each VC's scheduled weight. The WFQ maintains N cells from a VC in the class queue, where N is the weight of the VC. If greater than N cells exist, the excess is maintained in the VC queue. Cells are transferred from the VC queue to the Class queue to maintain the VC weight in the class queue.

10.9.2.2 Frame Continuous Queuing

Frame continuous queuing, or VC merge is available to all ports. It is configured on a per-VC basis. The VC queue is transformed into a frame re-assembly area. Frame traffic is assumed to use the AAL5 EOM PTI field indicator to delineate frame boundaries. Frames are completely assembled in the VC queue before being transferred over to the class queue. Non-user cells encountered on FCQ VCs are handled differently. Please refer the section on Non-User Cell Queuing.

The maximum length of the re-assembled frame can be one of two globally defined sizes, selected on a per-VC basis. Should a VC that is in process of re-assembly exceed the maximum length, a frame discard will be invoked. The cells in the VC queue will be discarded, as well as the cells that are about to be received up to and including the EOM. From a statistical count perspective, this frame discard is identical to a frame discard caused by congestion. In addition, a per-VC maskable interrupt is invoked and the ICI is stored in a register that only holds the ICI of the last VC that violated the maximum re-assembly length.

If a frame has a zero length field in the AAL5 trailer, there is a per-VC context parameter VcLenChkEn that will configure the queue engine to perform a frame discard. As with the maximum length frame discard, this zero length frame

discard is identical to a frame discard caused by congestion from a statistical count perspective.

A VC timeout watchdog is provided to protect memory resource should a re-assembly not complete in a timely manner. There are two procedures that are carried out by the watch dog. The first procedure is the patrol, which is performed by the queue engine during regular cell en-queue and de-queue sequence. Within the context record, there is a re-assembly parking state bit, ReasPark. The watch dog has a current re-assembly parking state bit, CurrentReasPark. Whenever a user cell (ie not RRM or OAM cell) arrives, the ReasPark state bit is set to the CurrentReasPark. The watch dog, initiated by the microprocessor, will walk through a programmable range of marked VCs, that are currently being re-assembled, to check and see if ReasPark = CurrentReasPark. If this is true, then the VC is deemed OK. If it finds a valid VC with ReasPark != CurrentReasPark, the VC is deemed dead. The discovery of a dead VC initiates the watch dog re-allocation procedure. When the patrol is complete, the CurrentReasPark bit is automatically inverted to prepare for the next patrol.

The watch dog re-allocation procedure is performed between the cell receive/transmit servicing. All the buffers in the VC queue are reclaimed, the VC Q congestion counters are reset to zero, the general discard count is updated, and VC status is reset to receive the next incoming cell as a BOM. A per-VC maskable interrupt is invoked and the ICI is stored in a register that only holds the ICI of the last timed out VC.

10.9.2.3 Shape Fair Queuing

The S/UNI APEX shaper is a passive dual rate shaper based on a time slot design. It will shape on a per VC basis, to the traffic parameters PCR, SCR & MBS. Traffic shaping is available on the four WAN ports, but not on the loop ports. A maximum of four out of the sixteen WAN port classes (four ports, four classes per port) can have shaping applied to their output. Every VC connected to a shaped class will have shaping applied to it, but each VC can have a unique shape rate. Classes that are not shaped can coexist on the same port as classes that are shaped, and there can be more than one shaped class on a single port.

Each shaper has a fundamental time unit, QShpNRTRate, which defines the minimum time increment between successively scheduled cells. Although each shaper is independent, the aggregate shape rate ($1/QShpNRTRate$) of the active shapers must be less than the device overall cell rate limit (1.42Mc/s @ 80MHz).

The VC's SCR is defined by the number of fundamental time units, ShpIncr, inserted between the VC's cell as they are scheduled by the shaper. The SCR is


proportional to $1/\text{ShpIncr}$. The CDV introduced is within 1 fundamental time unit. When there is neither contention (brief period where multiple VCs are scheduled to transmit at the same time) nor congestion (over-subscription of the port/class), the shaper will always transmit at SCR. See ideal emission of Figure 13.

The PCR and MBS only come into effect when the VC experiences one or more periods of contention or congestion (and hence the term passive dual rate shaper). An internal "late counter" is maintained that represents how late the current cell's scheduled emission time slot is relative to the ideal emission time slot. A non-zero late counter will cause the shaper to attempt to recover the lost opportunities by scheduling the cell with an increment value no smaller than $\text{ShpIncr} - \text{ShpCdvt}$. The ShpCdvt parameter, user defined on a per-VC basis, is in terms of the shaper's fundamental time unit. The difference ($\text{ShpIncr} - \text{ShpCdvt}$) is minimum number of fundamental time units inserted between cells, and is proportional to the VC's $1/\text{PCR}$. Given the opportunity, the shaper will burst at PCR rates until the late counter returns to zero¹. The size of the counter, programmable on a per-VC basis, therefore defines the MBS. See case #1 of Figure 13.

If congestion persists for an extended period, the late counter will continue to accumulate and eventually wrap around once MBS is reached. The resulting emission pattern is one where the duration of bursting is the remainder of the rolled counter. Every time the counter wraps, a CDV, equal to the MBS, is introduced into the emission stream. Recovery of the cumulative CDVs can only occur if the ingress stream pauses long enough for the VC queue to empty entirely. MaxCDV can be imposed by limiting the length of the VC queue via the per-VC max congestion threshold. See case #2 of Figure 13.


¹ Note that the inter-cell transmission times may actually exceed $1/\text{PCR}$. Factors include the number of active WAN ports, the number of active loop ports, and back pressure created by the external WAN port.

Figure 13 - Traffic Shaping on the WAN Port



For VCs that are shaped to rates approaching the fundamental time unit, there is the issue of granularity caused by the nature of time slots. For example, if the fundamental time unit is the equivalent of 100Mb/s, the maximum shaped rate is 100Mb/s ($ShpIncr = 1$), the next possible shaped rate is 50Mb/s ($ShpIncr = 2$). In order to achieve shaped rates between 100Mb/s and 50Mb/s, the $ShpIncr$ may be defined as an integer plus a fractional component. The shaper will schedule a cell to its integer value of $ShpIncr$, while maintaining a remainder count of the fractional portion. Whenever the remainder count exceeds a unit value, the shaper will schedule the next cell to the integer value + 1. The effective SCR rate over time will be the correct rate, but a CDV equal to the fractional value is introduced into the egress stream. If the $ShpIncr$ is an integer value, then there is no additional CDV introduced due to time slot granularity. PCR and MBS parameters are not supported when non-integer $ShpIncr$ is invoked. See Figure 14 where the $ShpIncr$ has been set to 1.5.

Figure 14 - Non-integer ShpIncr



User programming can also define the action of the shaper when overall egress congestion (i.e. too much traffic being sent through the shaper) is causing all VCs on that port/class to experience shaping delay due to congestion. When congestion is detected, the shaper will temporarily increase the fundamental shaping time unit, thereby causing each VC to schedule cells less frequently. This will eventually relieve the congestion, at which point the time unit will be brought back to its previous value. The impact of the congestion is distributed fairly across all VCs on the congested port **because all VCs on the port experience the same relative decrease in scheduling frequency.**

10.9.2.4 Non-User Cell Queuing

When a non-user cell is encountered, it may be queued with special handling. The cases requiring special handling are:

- Cells identified as an end to end OAM may be redirected to the uP's class 0 queue. This can occur independent of the queuing mechanism selected for the VC (WFQ, FCQ, and SFQ).
- Cells identified as a segment OAM may be redirected to the uP's class 0 queue. This can occur independent of the queuing mechanism selected for the VC (WFQ, FCQ, and SFQ).
- During FCQ, a cell identified as an OAM that is not being redirected to the uP will bypass the VC queue re-assembly area and go directly to the class queue.
- During FCQ, a cell identified as an RRM (Reserved or Resource Management) will bypass the VC queue re-assembly area and go directly to the class queue.

The table below lists the rules used to identify OAM and RRM cell types.

Table 8 - OAM & RRM Cell Identification

Type	Parameter	Location
VPC End to End OAM	VcVPC = 1	VC Context
	VCI = 4	Cell header
VPC Segment OAM	VcVPC = 1	VC Context
	VCI = 3	Cell header
VCC End to End OAM	VcVPC = 0	VC Context
	PTI = 5	Cell header
VCC Segment OAM	VcVPC = 0	VC Context
	PTI = 4	Cell header
RRM	VcVPC = 0	VC Context
	PTI = 11x	Cell header

VPC/VCC End to End OAM cells will be redirected to the uP's class 0 queue if the context parameter VcEEOam = 1, independent of the queue method selected. If VcEEOam = 0 and FCQ is selected, then the cell will be queued directly onto the class queue.

VPC/VCC segment OAM cells will be redirected to the uP's class 0 queue if the context parameter VcSegOam = 1, independent of the queue method selected. If VcSegOam = 0 and FCQ is selected, then the cell will be queued directly onto the class queue.

Non user cells not meeting any of the above conditions will not be redirected and will be treated like a normal user cell in terms of queuing.

The re-direction applied on OAM cells will preclude any performance measuring sessions on VCs that are programmed with FCQ.

10.9.3 Class Scheduling

Class scheduling is performed on the loop and WAN ports. There is no class scheduling for the uP ports as all four classes are accessible simultaneously. The class scheduler provides modified priority scheduling with class zero having the highest priority and class three having the lowest. The high priority classes can be utilized for real time services such as CBR and VBR-rt. The lower priority classes can be utilized for VBR-nrt, GFR and UBR services. There are three configurations for class scheduling:

- strict priority, round robin or modified strict priority between classes, evaluated after the transmission of each cell.
- strict priority between classes, evaluated after the transmission of an entire packet, and available only to those VCs configured for FCQ.
- strict priority between classes, evaluated after the transmission of a partial packet of a programmable length, and available only to those VCs configured for FCQ.

10.9.3.1 Cell and Packet Scheduling

In order to ensure that the lower priority classes are not starved when the high priority classes are under heavy utilization, a minimum bandwidth reservation scheme is employed. The user can program the minimum bandwidth requirements of classes one, two, and three and thus avoid starvation. Setting the minimum bandwidth requirements to zero ($\text{ClassXCellLmt} = 0$) on all classes will result in the class scheduler acting as a strict priority scheduler. Setting the minimum bandwidth requirements to three ($\text{ClassXCellLmt} = 3$) on all classes will result in the class scheduler acting as a round robin scheduler.

The mechanism utilized to ensure that a class does not starve is as follows. The class scheduler keeps track of the number of missed transmit opportunities the lower priority classes within a port have had. When a cell is transmitted on a particular class the ClassXCellCnt counters are incremented for all other classes which have missed an opportunity to transmit a cell. Once the ClassXCellCnt for a class reaches a maximum value (as defined by ClassXCellLmt), the class is in starvation. On the next cell transmit opportunity for that port, the starving class will be allowed to transmit one cell. If multiple classes were indicating starvation then the highest priority class would transmit first, then the next class until all starving classes have been serviced.

A starving class is only allowed to transmit one cell at a time. This ensures that the higher priority classes do not experience a large amount of CDV caused by the lower priority classes. When a class has an opportunity to transmit (due to starvation avoidance or otherwise), its ClassXCellCnt is reset and the above procedure is repeated.

A per-Port parameter, ClassPacket , is provided to support continuous packet transmission. In this packet mode, a VC that is configured for FCQ will retain permission to transmit cells for the length of the entire packet, regardless of the starvation states of the other classes, including class 0. This feature enables traffic to be emitted from the S/UNI APEX packet contiguously and thus minimizing the buffering requirements for an external SAR device. Strict priority must be set whenever packet class scheduling is selected.

It is possible to mix VCs that have FCQ and WFQ scheduling within the same class, and ClassPacket enabled. Non FCQ VCs and non-user cells (as defined in Table 8) are treated as single cell packets.

All counters, control and limit fields for the class scheduler are located in the class scheduler context memory. As stated earlier, memory is only allocated for classes one, two, and three. Class zero does not require any class scheduler context information.

10.9.3.2 Partial Packet Scheduling

A per-Port parameter, ClassFragEn, is provided to support packet fragmentation. In this fragmentation mode, classes are selected on a strict priority basis. Once a class is selected, the packet at the head of the class queue is transmitted up to a programmable length or until the EOM is encountered, whichever comes first. Non FCQ VCs and non-user cells (as defined in Table 8) are treated as single cell packets. When the length/EOM is reached, the classes are evaluated once again in a strict priority. The transmission of the original packet will resume once the original class regains transmission rights. Note that by virtue of the strict priority scheduling, Class0 will always have its packets transmitted in their entirety.

10.9.4 Congestion Control

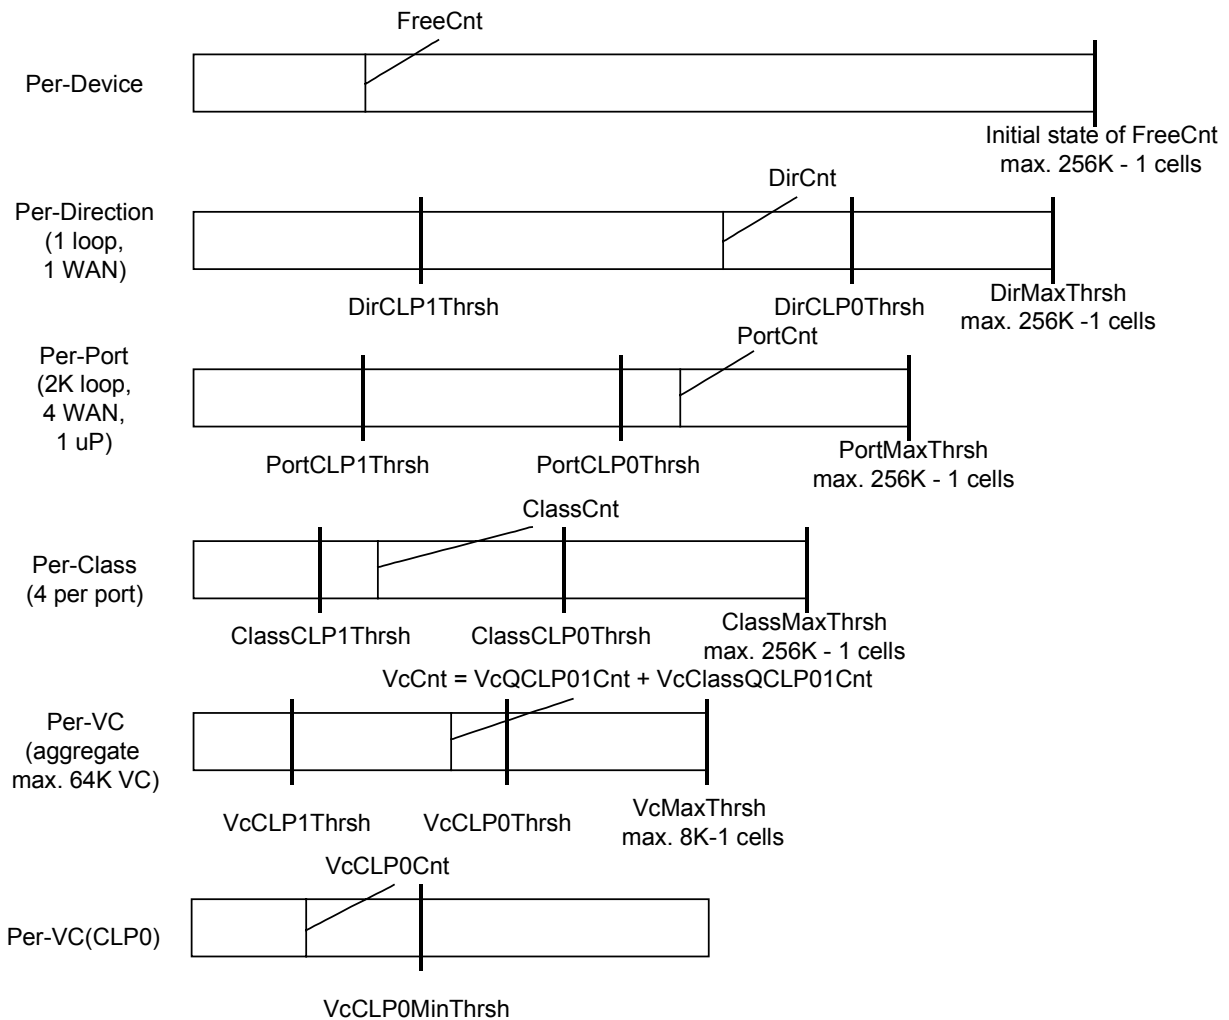
The congestion control decides whether to permit a cell to enter the queue structure. The objective is to provide a minimum reserved buffer allocation to all active VCs and to fairly allocate shared buffer resources to eligible VCs. The algorithm is applied to both frame and non-frame traffic. The objectives of the algorithm are as follows:

- provide guaranteed resources to all active VCs
- share available buffer resources to eligible VCs with excess buffering requirements
- restrict resource allocation on a per-VC, per-Class, per-Port, and per-Direction basis to those levels that have exceeded their allotment of resources.
- avoid global synchronization
- Provide interrupts and ID of the last maximum threshold discard invoked.

These objectives are achieved by having several thresholds and hierarchical count values, at the per-Device, per-Direction, per-Port, per-Class, and per-VC

levels. Figure 15 illustrates the relationship between the hierarchical count values and their associated thresholds.

Figure 15 - Thresholds and Count Definitions



Each hierarchical level has three population zones, each with its own discard rules:

- 1) Plenty of resources available, no discard
- 2) Some resources available, discard all cells with inbound CLP state = 1
- 3) Restricted resources available, discard all cells except cells that have inbound CLP state = 0 and have not met their minimum allocation of resources (**VcCLP0MinThrsh**).

No equivalent per-Direction count and threshold for uP destined cells, since there is only 1 uP port.

All counts represent the number of cells found at the hierarchical level, with the exception of FreeCnt at the per-Device level. The FreeCnt count value represents the number of free buffers remaining in the device. The initial value of FreeCnt is defined by the user.

The congestion algorithm has three possible definitions for CLP:

- 1) cell CLP, the CLP found in each cell's header;
- 2) BOM CLP, the CLP found in the frame BOM cell's header;
- 3) OR CLP, the running OR of all received user cell's CLPs since the BOM of a frame. Non-user cells do not affect the state of the running OR CLP.

Depending on the VC configuration, anyone of these three definitions can be used to increment a congestion count, or to select a threshold when comparing to a count.

When the queue engine receives a cell, the congestion control will apply the discard rules at each hierarchical level. Only when a cell has passed through each hierarchical level without being discarded will it be permitted entry into the queue.

Setting the Max threshold to zero on any given hierarchical level will effectively disable congestion discards at that hierarchical level. Exception to this rule is the VcMaxThrsh, which will always have the 8k-1 limit. The xxxCLP0Thrsh thresholds must always be set greater than or equal to the xxxCLP1Thrsh thresholds.

There are several error flags set whenever a non-zero maximum threshold is exceeded. Table 9 correlates the interrupts and context record identification parameters to the corresponding maximum threshold.

Table 9 - Congestion Error Flags

Threshold	Interrupt	Identification
VcMaxThrsh	QVcMaxThrshErr (Maskable on per-VC basis)	VcMaxThrshErrID
ClassMaxThrsh	QClassMaxThrshErr	ClassMaxThrshErrID ClassMaxThrshErrPortID
PortMaxThrsh	QPortMaxThrshErr	PortMaxThrshErrID

DirMaxThrsh	QDirMaxThrshErr	Check WANCnt or LoopCnt
FreeCnt = 0	QFreeCntZeroErr	N/A

EFCI marking may be performed on cells as they are transmitted out of the queue, based on the state of congestion at the time of transmission. Marking of EFCI is per VC selectable to occur at either the CLP1 thresholds or the CLP0 thresholds. A cell will be marked if hierarchical count values exceed one of the CLP1 thresholds (VcCLP1Thrsh, ClassCLP1Thrsh, PortCLP1Thrsh, DirCLP1Thrsh) or one of the CLP0 thresholds (VcCLP0Thrsh, ClassCLP0Thrsh, PortCLP0Thrsh, DirCLP0Thrsh), and the third bit of the PTI field in the cell is zero (PTI = 0xx).

There are three unique congestion discard rules. The selection of the rule to be applied is based on the cell type (user or non-user), the queuing mechanism, and finally the congestion type. If it is a non-user cell, the congestion mode is always cell discard. If shaping is not enabled for the destination port/class, the discard rule is selected on a per-VC basis, and is a function of the queue mechanism selected (VcQueue), as well as a per-VC congestion context parameter (VcCongMode). If the port/class is shaped, only two of the three rules is available, and is selectable on a per-VC basis. Table 10 below illustrates how the congestion discard rule is selected.


Table 10 - Congestion Discard Rules Selection

Cell	Shaped	VcQueue	VcCongMode	Congestion Mode
User	No	0 (WFQ)	0	EPD/PPD discard
User	No	0 (WFQ)	1	Cell discard
User	No	1 (FCQ)	x	FCQ discard
User	Yes	x	0	EPD/PPD discard
User	Yes	x	1	Cell discard
Non-user	x	x	x	Cell discard

OAM cells that are redirected to the microprocessor are subject to cell discard rules applied to the uP congestion counts at the per-port and per-class levels. There is no congestion control at the VC level for these redirected OAM cells.

10.9.4.1 EPD/PPD Discard

Figure 16 - EPD/PPD Congestion Discard Rules



When EPD/PPD discard is selected, the discard mechanism uses the AAL5 EOM PTI field indicator to delineate frame boundaries

EPD discard is evaluated only when the BOM is received, and is based on the BOM CLP state.

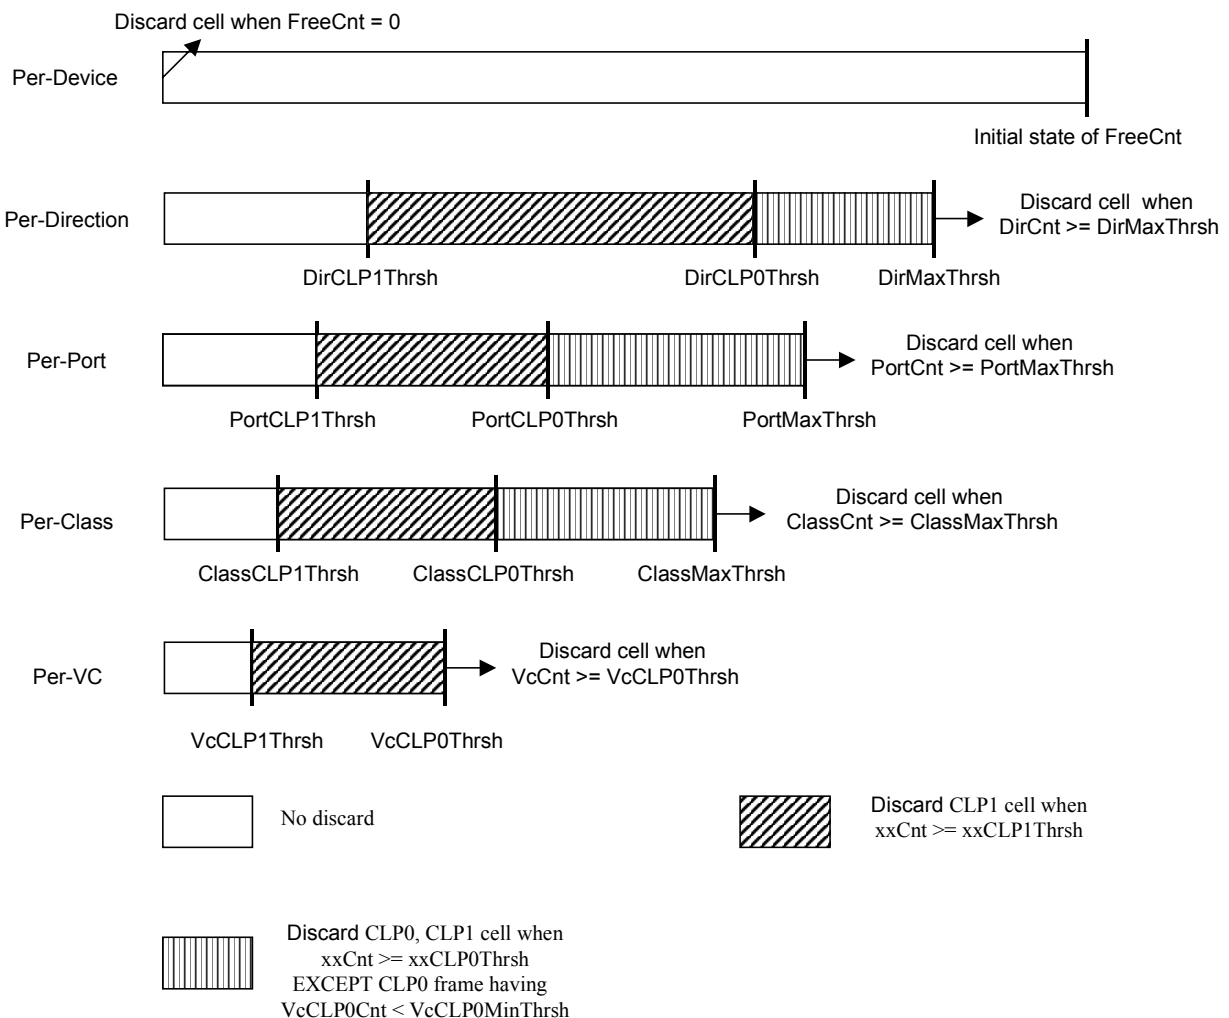
The VcCLP0Cnt increments when a received cell passes congestion and the inbound CLP state is zero. The VcCLP0Cnt decrements when the outbound CLP state is zero. The in/outbound CLP state is defined by the per-VC context parameter, VcGFRMode. When VcGFRMode = 0, the in/outbound CLP is

defined by the receiving/transmitting cell CLP state, respectively. When $VcGFRMode = 1$, the in/outbound CLP is defined by the receiving/transmitting frame's BOM CLP state, respectively.

If a PPD discard is invoked, the EOM will not be discarded unless one or more of the hierarchical count values is greater than or equal to the Dir/Port/Class/Vc MaxThrsh at the time the EOM is received. If the EOM is discarded, the following frames will be discarded and the congestion status will remain in PPD until an EOM is accepted. In the case when $VcGFRMode = 1$, the BOM CLP state of the first frame will be used to define the CLP state of the following discarded frames.

10.9.4.2 Cell Discard

Figure 17 Cell Congestion Discard Rules




When in cell discard mode, the CLP state is defined by each cell, and the decision to discard is evaluated upon receiving each cell. The minimum resource counter VcCLP0Cnt increments/decrements based on the cell CLP received/transmitted.

Non-user cells always have cell discard congestion rules applied, regardless of the original VC's congestion setting. Non-user cells do not have per-VC congestion as the VcQCLP01Cnt is not active when a non-user cell is encountered.

10.9.4.3 FCQ Discard

Figure 18 FCQ Discard Rules



When FCQ discard is selected, the discard mechanism uses the AAL5 EOM PTI field indicator to delineate frame boundaries.

Frame discard is evaluated after receiving each cell. The per-VC context parameter, VcGFRMode, dictate how frame discard is evaluated. When VcGFRMode = 0, frame discard is based on the OR CLP. When VcGFRMode = 1, frame discard is based on the BOM CLP.

The minimum resource counter is incremented/decremented after receiving/transmitting a cell that has the BOM CLP = 0.

When frame discard is invoked, the minimum resource count value will be reduced by the number of cells found in the VC queue if the BOM CLP = 0.

10.9.5 Statistics

There are two transmit counts, and three discard counts. All counts are 32-bits wide. The sum of all counts equals the total number of cells received by the S/UNI APEX. Table 11 gives a summary of the statistical counts.

Table 11 - Statistical Counts

Count	Scope	Description
VcCLP0TxCnt	Per- VC	Per-VC count of all cells transmitted that had an outbound CLP state of zero. OAM cells re-directed to the uP will not be represented by this count.
VcCLP1TxCnt	Per- VC	Per-VC count of all cells transmitted that had an outbound CLP state of one. OAM cells re-directed to the uP will not be represented by this count.
CLP0DiscardCnt	Global	Global count of all inbound CLP0 cells discarded due to congestion, re-assembly maximum length limit, or zero length check. There is an associated register that holds the last ICI that caused this count to increment.
CLP1DiscardCnt	Global	Global count of all inbound CLP1 cells discarded due to congestion, re-assembly maximum length limit, or zero length check. There is an associated register that holds the last ICI that caused this count to increment.
DiscardCnt	Global	Global count of all discards that are not due to congestion. These include cells discarded due re-assembly time outs, cells received on VCs that were not

Count	Scope	Description
		enabled, execution of a VC queue or a class queue tear down.

The definition of the in/out bound CLP state is a function of the congestion discard mode, and the per-VC parameter VcGFRMode. Table 12 illustrates the definition of in/out bound CLP state.

Table 12 - In/out Bound CLP State For Statistical Counts

Congestion Mode	VcGFRMode	Inbound CLP	Outbound CLP
Cell Discard	X	Cell CLP	Cell CLP
EPD/PPD Discard	0	Cell CLP	Cell CLP
	1	BOM CLP	Cell CLP
FCQ Discard	0	OR CLP	Cell CLP
	1	BOM CLP	Cell CLP

The table below give a brief summary of the rules applied for discard, and CLP definition for incrementing various counts as a function of the discard mode and the specific cell encountered.

Table 13 - Congestion Rule & Count Summary

Condition	Discard Mode	cell	EPD/PPD				FCQ			
			VcGFRMode	x	x	x	0	1	x	x
	Cell Type	x	pass thru OAM	redir ect OAM	user	user	pass thru OAM	redir ect OAM	user	user
Rules for discard	when discard decision is made	cell	cell	cell	BOM	BOM	cell	cell	cell	cell
	CLP def'n	cell	cell	cell	BOM	BOM	cell	cell	OR	BOM
CLP definition for various counts	VcCLP0Cnt	cell	cell	n/a	cell	BOM	cell	n/a	BOM	BOM
	VcCLP0TxCnt or VcCLP1TxCnt	cell	cell	cell	cell	cell	cell	cell	cell	cell

Condition	Discard Mode	cell	EPD/PPD				FCQ			
	VcGFRMode	x	x	x	0	1	x	x	0	1
	Cell Type	x	pass thru OAM	redirect OAM	user	user	pass thru OAM	redirect OAM	user	user
	CLP0DiscardCnt or CLP1DiscardCnt	cell	cell	cell	cell	BOM	cell	cell	OR	BOM

One reads the table vertically. Take the last column. A user cell arrives in a connection configured for FCQ, VcGFRMode = 1, will have

- its discard decision made on a cell by cell basis;
- the CLP is defined by the BOM for discard purposes;
- the minimum CLP0 count will be incremented based on the BOM, if the frame is not discarded;
- either VcCLP0TxCnt or VcCLP1TxCnt will be incremented based on the CLP of the cell, if the frame is not discarded;
- either CLP0DiscardCnt or CLP1DiscardCnt will be incremented based on the CLP of the BOM, if the frame is discarded.

10.9.6 Microprocessor Queue Buffer Re-allocation/Tear Down

The microprocessor has the option of engaging one of two macros that provide a fast mechanism to tear down either a VC queue or a Class queue for non-shaped port class. Specified and initiated through registers, the macro will go to the specified queue, reclaim the buffers in the queue, and reset the appropriate congestion counters. The number of cells that were in the queue are added to the general discard count. The VC queue or Class queue remain enabled after the re-allocation. Invoking of these functions may reduce general throughput of the device.

10.10 Context Memory SSRAM Interface

The context memory SSRAM interface stores and retrieves context data from one of two SSRAM devices: pipelined ZBT or register to register late write. Up to 4 banks and 4 SSRAM devices are supported, with 1M addressing capability for a total of 4MB data capacity. 2 parity bits are provided to protect the 34-bit data bus. If a parity error occurs, an interrupt is sent to the microprocessor.

The clock source drawn in Figure 19, Figure 20, Figure 21, Figure 22, Figure 23, must all be skew aligned at between S/UNI APEX, SDRAM & SSRAM clock input pins.

Figure 19 - 1 Bank Configuration for 1MB of ZBT SSRAM

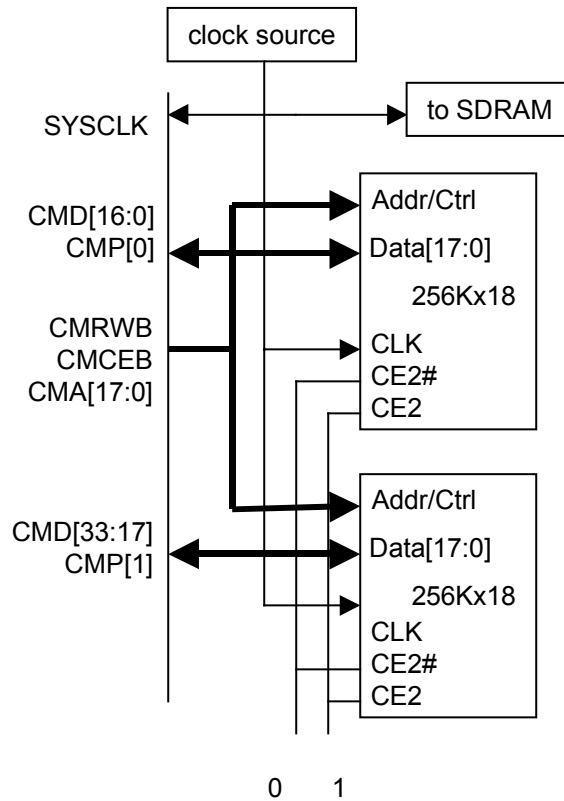


Figure 20 - 1 Bank Configuration for 1MB of Late Write SSRAM

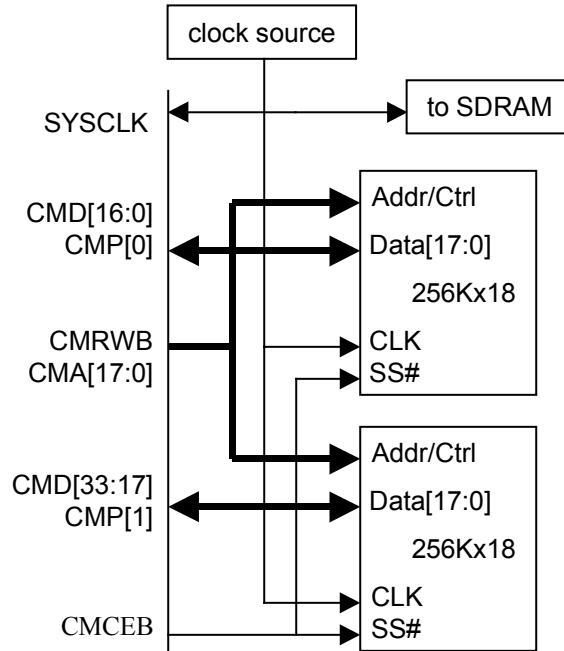


Figure 21 - 2 Bank Configuration for 2MB of ZBT SSRAM




Figure 22 - 2 Bank Configuration for 2MB of Late Write SSRAM

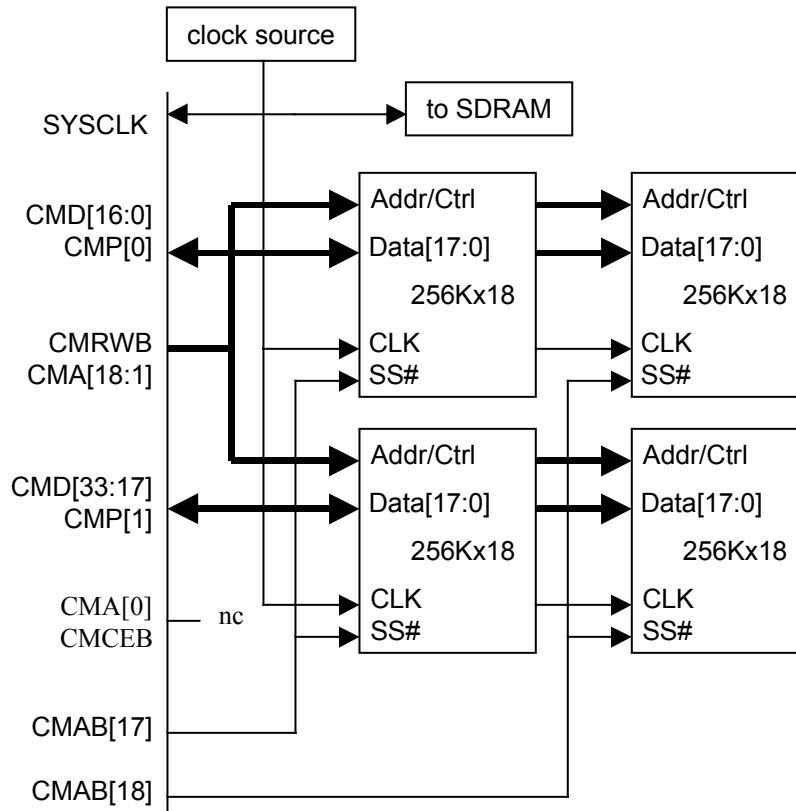
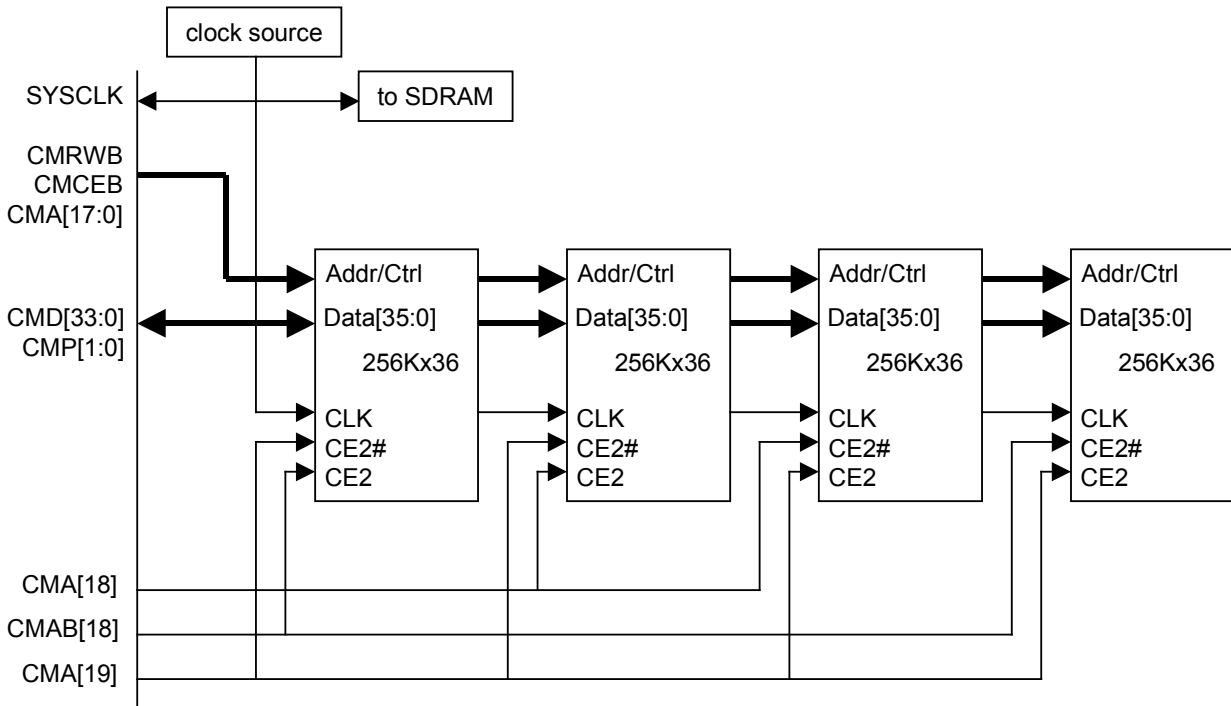


Figure 23 - 4 Bank Configuration for 4MB of ZBT SSRAM



There are two processes, arbitrated by the SSRAM arbiter, that access the context SSRAM:


1. The queue engine, for reading and writing context information, and for executing the re-assembly watchdog, described in the AAL5 Re-assemble Queuing section;
2. The microprocessor interface, for reading or writing context information, including the option of mask writes. See the Memory Port section for a description of the SSRAM access via the Microprocessor Interface;

10.11 Cell Buffer SDRAM Interface

The S/UNI APEX uses the external SDRAM to buffer queued cells. The cell buffer SDRAM interface permits up to 2 devices in parallel, with 4M addressing capability for a total of 16 MB of storage, sufficient for up to 256k cells. It has a 32 bit wide data bus, with CRC-16 checking applied on a per-cell basis. Each cell takes up 64 bytes of memory. The CRC-16 is applied to the first 60 bytes. If an error occurs, an interrupt is sent to the microprocessor.

The following diagram shows the cell storage map with the 64 byte memory boundary.

Figure 24 - Cell Storage Map



The clock source drawn in Figure 25, Figure 26, Figure 27 must all be skew aligned at between S/UNI APEX, SDRAM & SSRAMs clock input pins.

The following diagrams illustrate the various configurations supported:

Figure 25 - 4 MB – 64k Cells

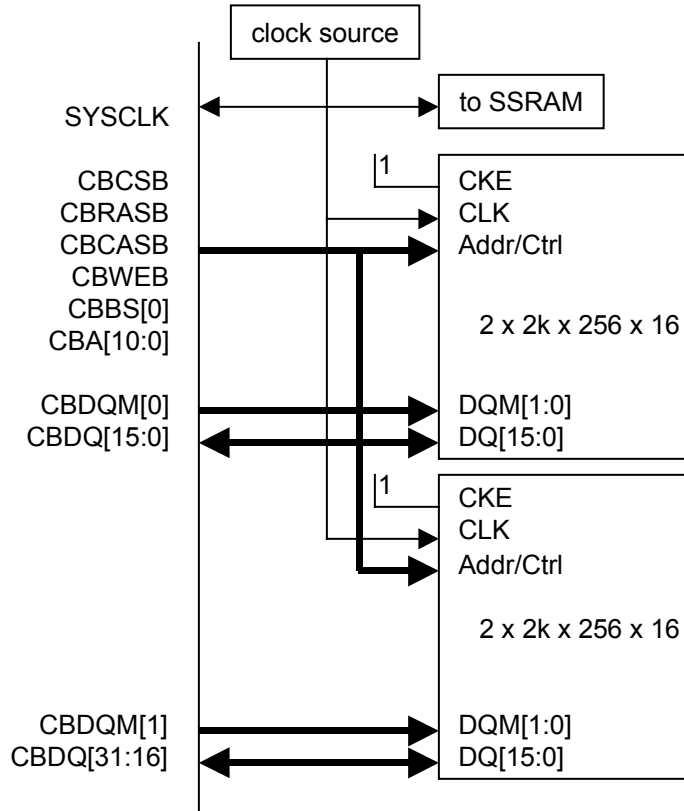


Figure 26 - 8 MB – 128k Cells





Figure 27 - 16 MB – 256k Cells



There are three processes, arbitrated by the SDRAM arbiter, that access the cell buffer SDRAM:

1. The queue engine, for reading and writing cells. The granularity of access by the queue engine is a concatenated 1 cell write - 1 cell read. Either the write or the read may not be performed, depending on the queue engine's requirements;
2. The microprocessor interface, for diagnostic reading or writing of 64 bytes of data. This data is aligned with the cell data. See the Operations section for a description of the data format;
3. The refresh controller, which has a programmable refresh rate.

The SDRAM interface will perform the initialization sequence for the SDRAM. This sequence is triggered by the SDRAM enable bit CBEn. The sequence will program the SDRAM with a CAS latency of 3, sequential access, write burst mode, and a burst length of 8. Application should ensure that sufficient time is provided between SDRAM power up and when this enable bit is set.

The SDRAM interface, under the direction of the queue engine, performs the header remapping function as the cell is read from SDRAM. It also attaches the cell preprends including the Switch Tag, ECI/ICI and Any PHY address preprend.

10.12 JTAG Test Access Port

The JTAG Test Access Port block provides JTAG support for boundary scan. The standard JTAG EXTEST, SAMPLE, BYPASS, IDCODE and STCTEST instructions are supported. The S/UNI APEX identification code is 073260CD hexadecimal.

11 PERFORMANCE

11.1 Throughput

The maximum throughput is governed at 3 potential bottlenecks:

- Receive interface configuration.
- Queue engine mode (shaper on or off).
- Transmit interface configuration.

The lowest throughput of the 3 bottlenecks will dictate the overall throughput for a given cell datapath.

General assumptions:

SYSCLK = 80MHz
no watch dog recovery in progress
SDRAM refresh set to slowest rate

BCLK = 66MHz
no setup/tear downs/context memory access in progress

Any-PHY clocks = 52MHz
16 bit interface for non uP receive and transmit interface
ICI in HEC/UDF field for non uP receive interface
ECI and Switch Tag disabled for non uP transmit interface

Table 14 - Receive Interface Throughput, Mcells/sec

Port/Configuration	Loop	WAN	uP	Assumptions
Any-PHY Master	1.67	1.67	n/a	address polling range optimized, n external port slaves, n >= 1, equal traffic
UTOPIA L2 Master	1.44	1.44	n/a	address polling range set to 2, 1 external port slave
	1.74	1.74	n/a	address polling range optimized, n external port slaves, n >= 2, equal traffic
UTOPIA L1 Master	1.74	1.74	n/a	none
UTOPIA L2 Slave	1.74	1.74	n/a	none

Port/Configuration	Loop	WAN	uP	Assumptions
uP full cell insertion	n/a	n/a	1.1	SarRxData0->13 written
uP partial cell (multi-cast)	n/a	n/a	1.74	SarRxData0, 13 written

Table 15 - Queue Engine Throughput, Mcells/sec

Configuration		Assumptions
Shaper disabled	1.74*	none
Shaper enabled	1.42	none

* Throughput drops down to 1.63 Mcells/sec if only a single WAN port is transmitting and no cells are being received by either the loop, WAN or uP Rx interfaces.

Table 16 - Transmit Interface Throughput, Mcells/sec

Port/Configuration	Loop	WAN	uP	Assumptions
Any-PHY Master	0.56*	n/a	n/a	limit 1 cell/port in FIFO, 1 external port slave
	1.11*	n/a	n/a	limit 2 cell/port in FIFO, 1 external port slave
	1.0*	n/a	n/a	limit 1 cell/port in FIFO, 2 external port slave
	1.6*	n/a	n/a	limit 2 cell/port in FIFO, 2 external port slave
	n/a	1.74	n/a	1 external port slave
	n/a	1.68	n/a	2 external port slave, equal traffic
UTOPIA L2 Master	0.61	n/a	n/a	limit 1 cell/port in FIFO, 1 external port slave
	1.25	n/a	n/a	limit 2 cell/port in FIFO, 1 external port slave
	1.1	n/a	n/a	limit 1 cell/port in FIFO, 2 external port slave
	1.6	n/a	n/a	limit 2 cell/port in FIFO, 2 external port slave

Port/Configuration	Loop	WAN	uP	Assumptions
	n/a	1.44	n/a	1 external port slave
	n/a	1.68	n/a	2 external port slave, equal traffic
UTOPIA L1 Master	1.67	1.74	n/a	none
UTOPIA L2 Slave	1.67	1.74	n/a	none
uP full cell extraction	n/a	n/a	1.61	none

* Throughput reduced when shaper is enabled. Guarantee minimum 0.4 Mcells/sec per loop port when there are 3 or less active loop ports.

11.2 Latency

The latency that a cell assumes an empty queue, SYSCLK = 80MHz, Any-PHY clocks = 52MHz, measure from SOP of the Any-PHY WAN/Loop receive interface to the SOP Any-PHY WAN transmit interface

Minimum Latency = 2340ns.

11.3 CDV

There are many points in the S/UNI APEX where CDV can be introduced.

12 REGISTER

Please use the List of Registers as a reference for the register map.

Notes on Register Bits:

1. Writing values into unused register bits has no effect. However, to ensure software compatibility with future, feature-enhanced versions of the product, unused register bits must be written with logic zero. Reading back unused bits can produce either a logic one or a logic zero; hence, unused register bits should be masked off by software when read.
2. Writable normal mode register bits are cleared to logic zero upon reset unless otherwise noted.
3. Writing into read-only normal mode register bit locations does not affect S/UNI APEX operation unless otherwise noted.
4. Certain register bits are reserved. These bits are associated with megacell functions that are unused in this application. To ensure that the S/UNI APEX operates as intended, reserved register bits must be written with their default value as indicated by the register bit description.
5. S/UNI APEX is addressable on a long-word basis only. Data fields are loaded into S/UNI APEX registers as described in each specific register section. S/UNI APEX does not perform any byte swapping.
6. With the exception of the CBI register port, part of the RAMBIST, and the Reset and Identity register, all registers are inaccessible until the software reset bit in the Reset and Identity register is removed.

12.1 General Configuration and Status

Register 0x00: Reset and Identity

Bit	Type	Function	Default
31:8		Unused	0
7	R/W	Reset	1
6:4	R	Type[2:0]	001
3:0	R	ID[3:0]	0

ID[3:0]

The ID bits can be read to provide a binary number indicating the S/UNI APEX feature version. These bits are incremented only if features are added in a revision of the chip.

Type[2:0]

The TYPE bits can be read to distinguish the S/UNI APEX from the other members of the DSLAM family of devices.

Reset

The RESET bit allows the S/UNI APEX to be reset under software control. If the RESET bit is a logic one, the entire S/UNI APEX except for the microprocessor interface is held in reset. This bit is not self-clearing. Therefore, a logic zero must be written to bring the S/UNI APEX out of reset. Holding the S/UNI APEX in a reset state places it into a low power, stand-by mode. A hardware reset sets the RESET bit, thus negating the software reset.

Notes:

- 1) Software should ensure that the DIIRun in the CBI register port reads back a 1 before releasing the S/UNI APEX from reset.
- 2) Software should wait 2 clock periods of the slowest clock (WTCLK, WRCLK, LTCLK, LRCLK, SYSCLK) before attempting to write to any other register. Exception to this rule is the CBI register port.

Register 0x10: Hi Priority Interrupt Status Register

Bit	Type	Function	Default
31	R/W	Reserved	0
30:15		Unused	
14	R/W	WTCelIXfErr	0
13	R/W	WRRuntCellErr	0
12	R/W	WRParErr	0
11		Unused	0
10	R/W	LTCelIXfErr	0
9	R/W	LRRuntCellErr	0
8	R/W	LRParErr	0
7:5		Unused	0
4	R/W	QFreeCntZeroErr	0
3:2		Unused	0
1	R/W	SSRAMParErr	0
0	R/W	SDRAMCrcErr	0

Each bit in the register is masked with the high priority interrupt mask register. The results are then NOR'd together to produce the state of INTHIB pin. All bits are cleared when this register is read. All may be set to one by the microprocessor for interrupt testing. Note that if the interrupt condition persists, the associated status bit will be reasserted.

SDRAMCrcErr

This bit goes high when a CRC-16 error was detected during a transaction on the SDRAM interface.

SSRAMParErr

This bit goes high when a parity error was detected during a transaction on the SSRAM interface.

QFreeCntZeroErr

This bit goes high when the entire device is completely congested, and that there is no more memory left to accept one more cell.

LRParErr

Loop receive parity error detected.

LRRuntCellErr

Loop receive runt cell error detected. A SOP is detected prior to receiving enough bytes for a cell.

LTCellXfErr

Loop transmit cell transfer error was encountered. This interrupt status is asserted when an external master device selects the Loop transmit interface for a cell transfer when the FIFO is empty.

WRParErr

WAN receive parity error detected.

WRRuntCellErr

WAN receive runt cell error detected. A SOP is detected prior to receiving enough bytes for a cell.

WTCellXfErr

WAN transmit cell transfer error was encountered. This interrupt status is asserted when an external master device selects the WAN transmit interface for a cell transfer when the FIFO is empty.

Register 0x14: High Priority Interrupt Mask

Bit	Type	Function	Default
31:0	R/W	Mask[31:0]	FFFFFFFF

Mask[31:0]

These bits mask the High Priority Interrupt Status Register.

Register 0x18: Low Priority Interrupt Error Register

Bit	Type	Function	Default
31	R/W	Reserved	0
30	R/W	QShp3lctrErr	0
29	R/W	Reserved	0
28	R/W	QShp2lctrErr	0
27	R/W	Reserved	0
26	R/W	QShp1lctrErr	0
25	R/W	Reserved	0
24	R/W	QShp0lctrErr	0
23	R/W	QDirMaxThrshErr	0
22	R/W	QPortMaxThrshErr	0
21	R/W	QClassMaxThrshErr	0
20	R/W	QVcMaxThrshErr	0
19		Unused	0
18	R/W	QCellRxErr	0
17	R/W	QVcReasLenErr	0
16	R/W	QVcReasTimeErr	0
15:0		Unused	0

Each bit in the register is masked with the low priority interrupt mask register. The results are then NOR'd together with the other low priority interrupt register to produce the state of INTLOB pin. All parameters are cleared when this register is read. All parameters may be set to one by the microprocessor for interrupt testing.

QVcReasTimeErr

Watch Dog found a re-assembled VC that timed out. The Misc Error Context Structure VcReasLenErrICl parameter has been updated to indicate the ICI of the frame that encountered this error. All cells in the re-assembly queue have been re-allocated. The next cell to arrive will be considered the BOM.

QVcReasLenErr

Maximum length in the re-assembly queue has been encountered. The Misc Error Context Structure VcReasLenErrICI parameter has been updated to indicate the ICI of the frame that encountered this error. All cells in the re-assembly queue have been re-allocated. PPD has been invoked on the VC.

QCellRxErr

Status bit indicating that a cell was received on a VC that was not enabled. Caused when either VcEn, ClassEn or PortEn are not set in the context record. The Misc Error Context Structure CellRxErrICI parameter has been updated to indicate the ICI of the cell that encountered this error.

QVcMaxThrshErr

Status bit indicating that a VC has reached the VC maximum threshold of VcMaxThrsh. The Maximum Congestion ID Misc Context Structure VcMaxThrshErrICI parameter has been updated to indicate the ICI of the cell that encountered this error.

QClassMaxThrshErr

Status bit indicating that a class queue has reached the class maximum threshold ClassMaxThrsh. The Maximum Congestion ID Misc Context ClassMaxThrshErrID & ClassMaxThrshErrPortID parameters have been updated to indicate the Port and class of the cell that encountered this error.

QPortMaxThrshErr

Status bit indicating that a port queue has reached the port maximum threshold PortMaxThrsh. The Maximum Congestion ID Misc Context Structure PortMaxThrsh parameter has been updated to indicate the PortID of the cell that encountered this error.

QDirMaxThrshErr

Status bit indicating that the per-direction maximum congestion threshold has been encountered. Check LoopCnt and WANCnt in the Overall Count Misc Context to determine whether it was the loop or the WAN ports that reached this threshold.

QShpNlctrErr, N = 0..3

Status bit indicating that the ingress counter has saturated.

Register 0x1C: Low Priority Interrupt Error Mask

Bit	Type	Function	Default
31:0	R/W	Mask[31:0]	FFFFFFFF

Mask[31:0]

These bits mask the Low Priority Interrupt Status Register.

Register 0x20: Low Priority Interrupt Status Register

Bit	Type	Function	Default
31:7		Unused	0
6	R	MPIdleStatus	1
5	R	SarRxEmptyStatus	1
4	R	SarRxRdyStatus	1
3:0	R	SarTxRdyStatus[3:0]	0

Each bit in the register is masked with the low priority interrupt status mask register. The results are then NOR'd together with the other low priority interrupt register to produce the state of INTLOB pin.

SarTxRdyStatus[3:0]

Status bit indicating that one of four SAR read buffers contains a least one cell for reading. **Warning:** Software should not attempt to read this status or clear the associated interrupt mask immediately after extracting a cell from the S/UNI APEX. There is a latency of 3 BCLKs + 2 SYSCLKs between the last word of a cell read out and this signal going inactive. Removing the mask prematurely may generate an unintentional interrupt.

SarRxRdyStatus

Status bit indicating that SAR receive buffer is ready to accept the next cell.

SarRxEmptyStatus

Status bit indicating that SAR receive buffer is empty. Typically used for diagnostic writes.

Warning: Software should not attempt to read this status or clear the associated interrupt mask immediately after injecting a cell into the S/UNI APEX. There is a latency of 5 SYSCLKs between the last word of a cell written out and this signal going inactive. Removing the mask prematurely will generate an unintentional interrupt.

MPIdleStatus

Status bit indicating the memory port is idle and ready to accept a new command. This signal is the inverse of MPBusy found in the Memory Port Control Register.

Register 0x24: Low Priority Interrupt Status Mask

Bit	Type	Function	Default
31:0	R/W	Mask[31:0]	FFFFFFFF

Mask[31:0]

These bits mask the Low Priority Interrupt Status Register.

12.2 Loop Cell Interface

Register 0x100: Loop Cell Rx Interface Configuration

Bit	Type	Function	Default
31:18		Unused	0
17:16	R/W	LoopRxCISel[1:0]	0
15:13		Unused	0
12:8	R/W	LoopRxPollAddr[4:0]	0
7	R/W	LoopRxParPolarity	0
6	R/W	LoopRx8bitEn	0
5	R/W	LoopRxHecDis	0
4		Unused	0
3	R/W	LoopRxCIPreEn	1
2:1	R/W	LoopRxMode[1:0]	0
0	R/W	LoopRxEn	0

All parameters in this register should only be set once at the same time the interface is enabled. Once the interface is set, none of these parameters may be changed.

LoopRxEn

The LoopRxEn enables the loop receive interface. When set to one, the loop receive Any-PHY interface operates normally. Once set, this bit should not be reset to zero.

LoopRxMode[1:0]

Selects the receive interface mode

LoopRxMode[1:0]	Operation
00	UTOPIA L2 master, supports 32 PHYs
01	UTOPIA L1 master, supports 1 PHY
10	Any-PHY master mode inband port notification via the address prepend. Supports 32 PHYs.
11	UTOPIA L2 slave

LoopRxICIPreEn:

When set to one, the default, an ICI prepend (Word 1 in Figure 3, Byte 1&2 in Figure 4) is expected on this interface. When reset to zero an ICI prepend is not expected on this interface. When in 8 bit mode, this bit must be set to 1.

LoopRxHecDis

When reset to zero, the HEC/UDF field (Word 4 in Figure 3, Byte 7 in Figure 4) is expected on this interface. When set to a one an HEC/UDF field is not expected on this interface.

LoopRx8bitEn

When reset to zero, this bit sets the interface bus width to 16 bits. When set to one, this bit sets the interface bus width to 8 bit.

LoopRxParPolarity

When reset to zero, the loop receive parity is odd. When set to one, the loop receive parity is even.

LoopRxPollAddr[4:0]

In UTOPIA L2 slave mode: The five bit UTOPIA Address to which the slave will respond.

In Any-PHY and UTOPIA L2 master mode: These bits represents the polling address range.

LoopRxPollAddr[4:0]	Polling range
0-2	Not valid
3	Address 3->0
...	...
31	Address 31->0

In UTOPIA L1 master mode: These bits are reserved.

LoopRxICISel[1:0]

Indicates which part of the incoming cell, the internal ICI is selected from:

LoopICISel[1:0]	Source
00	User Prepend (Note: LoopRxICIPreEn must be set).
01	VPI/VCI fields: if VPI < "FFF" then

	ICI = "0" & VPI; else ICI = VCI; end if;
10	HEC/UDF fields (Note: Both LoopRxHecDis and LoopRx8bitEn must be clear, and not valid for 8-bit mode).
11	Unused

Register 0x104: Loop Cell Tx Interface Configuration

Bit	Type	Function	Default
31	R	LoopTxPfFullStatus	0
30:29		Unused	0
28:24	R/W	LoopTxPfThres	0
23:22		Unused	0
21	R/W	LoopTxTwoCellEn	0
20	R/W	LoopTxSchEn	0
19:13		Unused	0
12:8	R/W	LoopTxSlaveAddr[4:0]	0
7	R/W	LoopTxParPolarity	0
6	R/W	LoopTx8bitEn	0
5	R/W	LoopTxHecDis	0
4	R/W	LoopTxSwitchPreEn	0
3	R/W	LoopTxECIPreEn	1
2:1	R/W	LoopTxMode[1:0]	0
0	R/W	LoopTxEn	0

All parameters in this register should only be set once at the same time the interface is enabled. Once the interface is set, none of these parameters may be changed.

LoopTxEn

The LoopTxEn enables the loop transmit interface. When set to one, the loop transmit Any-PHY interface operates normally. Once set, this bit should not be reset to zero.

LoopTxMode[1:0]

Selects the transmit interface mode.

LoopTxMode[1:0]	Operation
00	UTOPIA L2 master mode, supports 32 PHYs
01	UTOPIA L1 master mode, supports 1 PHY

LoopTxMode[1:0]	Operation
10	Any-PHY master mode inband port selection, via the address prepend. Supports 2048 PHYs.
11	UTOPIA L2 slave

LoopTxECIPreEn:

When reset to zero, a ECI prepend (Word 2 in Figure 5, Byte 3&4 in Figure 6) is not present on this interface. When set to one a ECI prepend is present on this interface.

LoopTxSwitchPreEn

When reset to zero, a switch tag prepend (Word 1 in Figure 5, Byte 1&2 in Figure 6) is not present for this interface. When set to a one a switch tag is present for this interface.

LoopTxHecDis

When reset to zero, the ECI is generated on the HEC/UDF field (Word 5 in Figure 5, Byte 9 in Figure 6) in this interface. When set to a one an HEC/UDF field is not generated on this interface.

LoopTx8bitEn

When reset to zero, this bit sets the interface bus width to 16 bits. When set to one, this bit sets the interface bus width to 8 bit.

LoopTxParPolarity

When reset to zero, the loop transmit parity is odd. When set to one, the loop transmit parity is even.

LoopTxSlaveAddr[4:0]

The five bit UTOPIA Address to which the slave will respond. Used only when the loop Tx interface is configured for UTOPIA L2 slave.

LoopTxSchEn

The LoopTxSchEn enables the loop port scheduler for normal operation. The LoopTxSchEn enable should be set to a '1' for normal operation after the initialization of the loop's class not empty context memory.

LoopTxTwoCellEn

When set to 0, the loop port scheduler will allow a maximum one cell per port in the transmit pipeline for each LTPA. When set to 1, the loop port scheduler will allow a maximum of two cells per port in the transmit pipeline for each LTPA.

LoopTxPfThres[4:0]

Controls the depth of the poll request FIFO, offset by 1. A value of 0 represents poll depth of 1. Recommended value is 0x1f.

LoopTxPfFullStatus

Read only bit. When read high, it indicates that the polling FIFO has reached the LoopTxPfThres value. Reading this bit will reset the value to zero. Used for diagnostics.

12.3 WAN Cell Interface

Register 0x200: WAN Cell Rx Interface Configuration

Bit	Type	Function	Default
31:18		Unused	0
17:16	R/W	WANRxICISel[1:0]	0
15:13		Unused	0
12:10	R/W	Reserved	0
9:8	R/W	WANRxPollAddr[1:0]	0
7	R/W	WANRxParPolarity	0
6	R/W	WANRx8bitEn	0
5	R/W	WANRxHecDis	0
4		Unused	0
3	R/W	WANRxICIPreEn	1
2:1	R/W	WANRxMode[1:0]	0
0	R/W	WANRxEn	0

All parameters in this register should only be set once at the same time the interface is enabled. Once the interface is set, none of these parameters may be changed.

WANRxEn

The WANRxEn enables the WAN receive interface. When set to one, the WAN receive Any-PHY interface operates normally. Once set, this bit should not be reset to zero.

WANRxMode[1:0]

Controls the port selection mode.

WANRxMode[1:0]	Operation
00	UTOPIA L2 master, supports 4 PHYs

WANRxMode[1:0]	Operation
01	UTOPIA L1 master, supports 1 PHY
10	Any-PHY master mode inband port notification via the address prepend. Supports 4 PHYs.
11	UTOPIA L2 slave

WANRxICIPreEn

When set to one, the default, an ICI prepend (Word 1 in Figure 3, Byte 1&2 in Figure 4) is expected on this interface. When reset to zero an ICI prepend is not expected on this interface. When in 8-bit mode, this bit must be set to 1.

WANRxHecDis

When reset to zero, the HEC/UDF field (Word 4 in Figure 3, Byte 7 in Figure 4) is expected on this interface. When set to a one an HEC/UDF field is not expected on this interface.

WANRx8bitEn

When reset to zero, this bit sets the interface bus width to 16 bits. When set to one, this bit sets the interface bus width to 8 bit.

WANRxParPolarity

When reset to zero, the WAN receive parity is odd. When set to one, the WAN receive parity is even.

WANRxPollAddr[1:0]

In UTOPIA L2 slave mode: The two bit UTOPIA Address to which the slave will respond.

In Any-PHY and UTOPIA L2 master mode: These bits represents the polling address range.

WANRxPollAddr[1:0]	Polling range
0-2	Not valid
3	Address 3->0

In UTOPIA L1 master mode: These bits are reserved.

WANRxICISel[1:0]

Indicates which part of the incoming cell, the internal ICI is selected from:

WANRxICISel[1:0]	Source
00	User Prepend (Note: WANRxICIPreEn must be set).
01	VPI/VCI fields: if VPI < "FFF" then ICI = "0" & VPI; else ICI = VCI; end if;
10	HEC/UDF fields (Note: Both WANRxHecDis and WANRx8bitEn must be clear, and not valid for 8-bit mode).
11	Unused

Register 0x204: WAN Cell Tx Interface Configuration

Bit	Type	Function	Default
31:30	R/W	WANTx3PortID[1:0]	3
29:28	R/W	WANTx2PortID[1:0]	2
27:26	R/W	WANTx1PortID[1:0]	1
25:24	R/W	WANTx0PortID[1:0]	0
23:21		Unused	0
20	R/W	WANTxSchEn	0
19:10		Unused	0
9:8	R/W	WANTxSlaveAddr[1:0]	0
7	R/W	WANTxParPolarity	0
6	R/W	WANTx8bitEn	0
5	R/W	WANTxHecDis	0
4	R/W	WANTxSwitchPreEn	0
3	R/W	WANTxECIPreEn	1
2:1	R/W	WANTxMode[1:0]	0
0	R/W	WANTxEn	0

All parameters in this register should only be set once at the same time the interface is enabled. With the exception of the WANTxXPortID and WANTxSlaveAddr bits, once the interface is set, none of these parameters may be changed.

WANTxEn

The WANTxEn enables the WAN transmit interface. When set to one, the WAN transmit Any-PHY interface operates normally. Once set, this bit should not be reset to zero.

WANTxMode[1:0]

Controls the port selection mode.

WANTxMode[1:0]	Operation
00	UTOPIA L2 master, supports 4 PHYs
01	UTOPIA L1 master , supports 1 PHY

WANTxMode[1:0]	Operation
10	Any-PHY master mode inband port notification via the address prepend. Supports 4 PHYs.
11	UTOPIA L2 slave

WANTxECIPreEn

When reset to zero, an ECI prepend (Word 2 in Figure 5, Byte 3&4 in Figure 6) is not generated on this interface. When set to one an ECI prepend is generated on this interface.

WANTxSwitchPreEn

When reset to zero, a switch tag prepend (Word 1 in Figure 5, Byte 1&2 in Figure 6) is not generated for this interface. When set to a one a switch tag prepend is generated for this interface.

WANTxHecDis

When reset to zero, the ECI is generated on the HEC/UDF field (Word 5 in Figure 5, Byte 9 in Figure 6) in this interface. When set to a one an HEC/UDF field is not generated on this interface.

WANTx8bitEn

When reset to zero, this bit sets the interface bus width to 16 bits. When set to one, this bit sets the interface bus width to 8 bit.

WANTxParPolarity

When reset to zero, the WAN transmit parity is odd. When set to one, the WAN transmit parity is even.

WANTxSlaveAddr[1:0]

In UTOPIA L2 slave mode : The two bit UTOPIA address to which the slave will respond.

In UTOPIA L1 master mode: The two bit UTOPIA address that is presented during the address selection phase. Permits an external UTOPIA L2 slave to be connected when interface is in UTOPIA L1 master mode.

WANTxSchEn

The WANTxSchEn enables the WAN port scheduler for normal operation. The WANTxSchEn enable should be set to a '1' for normal operation after the initialization of the WAN's class not empty context memory.

WANTxXPortID[1:0], X = 0..3 representing channel #.

The value in these register are the alias mapping of the internal WAN PortID to the physical port ID as presented by the Any-PHY/UTOPIA master mode interface. By default, the internal matches the external port IDs.

In slave mode, the WAN port scheduler only operates on Port 0; therefore, the following relationship must always be true:

WANTx0PortID = WANTxSlaveAddr

Warning: the values programmed into these four registers must always be unique. Two or more internally active PortID must never point to the same active physical port ID.

12.4 Memory Port

See section on Memory Port Mapping for specific context parameter definitions.

Register 0x300: Memory Port Control

Bit	Type	Function	Default
31	R/W	MPBusy	0
30:29	R/W	MPCCommand[1:0]	0
28		Unused	
27:24	R/W	MPLWordEn[3:0]	0
23:22	R/W	MPMemSelect[1:0]	0
21:18		Unused	0
17:0	R/W	MPQuadAddr[17:0]	0

Writes to this register will place the microprocessor interface in wait state until the MPBusy bit in the control register is clear.

MPQuadAddr[17:0]

Indicates the beginning quad long word address for the operation in memory. Up to 4 megabytes of memory is supported in each aperture by this address (or 256k 16-byte regions).

MPMemSelect[1:0]

Selects the memory aperture. The aperture is chosen according to the following table.

MPMemSelect[1:0]	Aperture Selected
00	External Queue Context
01	Internal Queue Context
10	Internal WAN Port Context
11	Internal Loop Port Scheduler Context

MPLWordEn[3:0]

Indicates which long words of data will be written to or read from memory. This register is used to resolve the quad long word address MPQuadAddr down to long word resolution.

For masked writes, only one bit should be set, identifying the long word address. If more than one bit in this parameter is set then the least significant active bit will be used to indicate the long word address to be modified.

For normal read/write access, the number of active bits permitted is unrestricted with the exception of the internal WAN port scheduler context memory map and internal loop port scheduler context memory map. These 2 memory maps are restricted to having only 1 active bit per access.

MPCommand[1:0]

Selects the type of access. If a masked write is indicated, a 34-bit mask will be used to determine which bits within one long word will be written.

MPCommand[1:0]	Command Selected
00	Reserved
01	Write
10	Read
11	Masked Write

MPBusy

When set to a one by the microprocessor, the command will be executed. When the command is complete, this bit will be cleared to zero. This signal is the inverse of MPIdleStatus found in the Interrupt Status Register.

Register 0x340-0x34C: Memory Write Data (Burstable)

Bit	Type	Function	Default
31:0	R/W	MPWrDataN[31:0], N = 0..3	0

Writes to this register will be delayed until the MPBusy bit in the control register is clear.

MPWrDataN[31:0], N = 0..3

The least significant 32 bits of write data to be directed to the address and aperture as specified in the memory port control register.

For normal write operations, MPWrDataN corresponds to MPLWordEn[N].

For masked writes, MPWrData0 contains the data, MPWrData1 contains the mask bits: only bits set to a one in this vector will be overwritten with MPWrData0 in memory.

Register 0x350: Memory Write Data Overflow (Burstable)

Bit	Type	Function	Default
31:8		Unused	0
7:6	R/W	MPWrData3[33:32]	0
5:4	R/W	MPWrData2[33:32]	0
3:2	R/W	MPWrData1[33:32]	0
1:0	R/W	MPWrData0[33:32]	0

Writes to this register will be delayed until the MPBusy bit in the control register is clear. This register is only used for writes to the external queue context.

MPWrData0[33:32]

The most significant 2 bits of MPWrData0. See Memory Write Data array for a more detailed description.

MPWrData1[33:32]

The most significant 2 bits of MPWrData1. See Memory Write Data array for a more detailed description.

MPWrData2[33:32]

The most significant 2 bits of MPWrData2. See Memory Write Data array for a more detailed description.

MPWrData3[33:32]

The most significant 2 bits of MPWrData3. See Memory Write Data array for a more detailed description.

Register 0x380-0x38C: Memory Read Data (Burstable)

Bit	Type	Function	Default
31:0	R	MPRdDataN[31:0] , N = 0..3	0

Reads from this register will be delayed until the MPBusy bit in the control register is clear.

MPRdDataN[31:0], N = 0..3

The least significant 32 bits of read data from the address and aperture as specified in the memory port control register. MPRdDataN corresponds to MPLWordEn[N].

Register 0x390: Memory Read Data Overflow (Burstable)

Bit	Type	Function	Default
31:8		Unused	0
7:6	R	MPRdData3[33:32]	0
5:4	R	MPRdData2[33:32]	0
3:2	R	MPRdData1[33:32]	0
1:0	R	MPRdData0[33:32]	0

Reads from this register will be delayed until the MPBusy bit in the control register is cleared. This register is only used for reads from external queue context.

MPRdData0[33:32]

Indicates the most significant 2 bits of the first word of read data from memory. See Memory Read Data array for a more detailed description.

MPRdData1[33:32]

Indicates the most significant 2 bits of the first word of read data from memory. See Memory Read Data for a more detailed description.

MPRdData2[33:32]

Indicates the most significant 2 bits of the second word of read data from memory. See Memory Read Data for a more detailed description.

MPRdData3[33:32]

Indicates the most significant 2 bits of the third word of read data from memory. See Memory Read Data for a more detailed description.

MPRdData4[33:32]

Indicates the most significant 2 bits of the fourth word of read data from memory. See Memory Read Data for a more detailed description.

12.5 SAR

12.5.1 Receive

In normal SAR operation, writing to any SAR receive data register when SarRxRdyStatus = 0 (in the low priority interrupt status register) will cause the microprocessor bus to be held in wait state. Data transfer is initiated when SarRxData13 is written.

When in diagnostic write mode, writing to any SAR receive data register when SarRxEmptyStatus = 0 (in the low priority interrupt status register) will cause the microprocessor bus to be held in wait state. Data transfer is initiated when SarRxData15 is written.

Register 0x400-0x43C: SAR Receive Data (Burstable)

Bit	Type	Function	Default
31:0	W	SarRxDataN[31:0], N = 0..15	X

SarRxData0[15:0]: ICI

In normal SAR operation, this is the Ingress Connection Identifier that identifies the connection in which the cell belongs.

In diagnostic write mode, this is simply a write data register.

SarRxData0[17:16]: CRC Trailer

In normal SAR operation, these bits define whether the next cell received should be over written with a trailer.

SarRxData0[17:16]	Function
00	Do not overwrite the end of the current cell with a trailer
01	Write the current cell with CRC-32 trailer.
10	Write the current cell with CRC-10 trailer.
11	Unused

In diagnostic write mode, this is simply a write data register.

SarRxData0[18]: CRC-32 Init

In normal SAR operation, when set, this bit will set the internal CRC32 to 0xFFFFFFFF. Required when the current cell received is the BOM of an

AAL5 packet. This bit must be reset if the current cell received is not the BOM of an AAL5 packet.

In diagnostic write mode, this is simply a write data register.

SarRxData0[31:19]

In normal SAR operation, these bits are not used.

In diagnostic write modes, this is simply a write data register.

SarRxDataN[31:0], N=1..15:

In normal SAR operation, these registers contain the header and payload. Data stored in SarRxData14 & SarRxData15 are not used. Transfer is initiated when SarRxData13 is written.

In diagnostic write mode, this is simply a write data register. Transfer is initiated when SarRxData15 is written.

12.5.2 Transmit

Reading from any registers within the transmit SAR when SarTxRdyStatus = 0 (in the Low Priority Interrupt Status register), or when SarDiagRdBusy = 1 (in the Cell Buffer Diagnostic Control register) will cause the microprocessor bus to be held in wait state.

All registers in the transmit SAR are read-only; the registers may not be initialized by the microprocessor directly.

All 16 registers in the class 3 transmit SAR buffer are used to report read data when cell buffer diagnostic mode is enabled. In this mode, the entire block of data is cleared from the buffer once the last word is read.

Register 0x500-0x53C: SAR Transmit Data, Class 0 (Burstable)

Register 0x540-0x57C: SAR Transmit Data, Class 1 (Burstable)

Register 0x580-0x5BC: SAR Transmit Data, Class 2 (Burstable)

Register 0x5C0-0x5FC: SAR Transmit Data, Class 3 (Burstable)

Bit	Type	Function	Default
31:0	R	SarTxDataN[31:0], N = 0..15	X

SarTxData0[15:0]: ECI

In normal SAR operation. This is the Egress Connection Identifier of the cell received.

In diagnostic read mode, class 3, this is the read data.

SarTxData0[16]: CRC32 Status

In normal SAR operation, when high, this bit indicates that the AAL5 CRC32 polynomial check failed. This bit is only valid for cells that belong to an AAL5 packet (as indicated by the cell header). Within AAL5 packets, this status will return the actual CRC32 status with the EOM, and all other cells will be accompanied by a status of 0. This status should be ignored for all other cell types.

In diagnostic read mode, class 3, this is the read data.

SarTxData0[17]: CRC10Stat

In normal SAR operation, when high, this bit indicates that the CRC10 polynomial check failed. This status is only valid when the cell in the buffer is an OAM cell (as indicated by the cell header). This status should be ignored for all other cell types.

In diagnostic read mode, class 3, this is the read data.

SarTxData0[31:18]: Unused

In normal SAR operation, not used.

In diagnostic read mode, class 3, this is the read data.

SarTxDataN[31:0], N=1..15:

In normal SAR operation, these registers contain the header and payload. For Data stored in SarTxData14 & SarTxData15 are reserved and are not used. A new transfer is initiated when SarTxData13 is read.

In diagnostic read mode, class 3, these registers are the read data. A new transfer is initiated when SarTxData15 is read.

12.5.3 Cell Buffer Diagnostic Access

Register 0x600: Cell Buffer Diagnostic Control

Bit	Type	Function	Default
31	R/W	SarDiagRdBusy	0
30	R/W	SarDiagRdModeEn	0
29	R	SarDiagRdModeLock	0
28	R/W	SarDiagWrModeEn	0
27:18		Unused	0
17:0	R/W	SarDiagAddr[17:0]	0

SarDiagAddr[17:0]

Indicates the beginning 16-long word address for the operation in cell buffer. Up to 16 megabytes of memory is supported by this address (or 256k 64-byte regions).

SarDiagWrModeEn

When enabled, the SAR receive staging buffer will be used as a port to write data directly to the cell buffer (SDRAM). Receiving of normal cells into the traffic stream is no longer possible while this parameter is enabled. This bit should only be set when SarRxEmptyStatus = 1 in the low interrupt status register.

SarDiagRdModeLock

Read only. Indicates that the SAR is ready to perform a diagnostic reads. A lock will only occur when SarDiagRdModeEn = 1, and all non-diagnostic cells remaining in the class 3 Tx staging buffers have been read by the microprocessor.

SarDiagRdModeEn

When enabled, normal loading of cells into all 4 classes of the SAR Tx staging buffers are withheld. A cell that is in the process of being loaded into a Tx staging buffer when this bit is set will be allowed to complete.

SarDiagRdBusy

Setting this register to one will initiate a diagnostic read from the cell buffer (SDRAM). When the command is complete, this bit will be cleared to zero. This bit should not be set to one until SarDiagRdModeEn = 1, SarDiagRdModeLock = 1, and SarTxRdyStatus = 0.

12.6 Queue Engine

Register 0x700: Queue Context Configuration

Bit	Type	Function	Default
31	R/W	QEngEn	0
30	R	QBusy	0
29	R/W	QSglStep	0
28:27		Unused	0
26	R/W	QRxTxArbSel	0
25		Unused	0
24	R/W	QNumVCSEL	0
23:16	R/W	QCellStartAdr[7:0]	0
15:8	R/W	QLClassStartAdr	0

Bit	Type	Function	Default
		[7:0]	
7:0	R/W	QShpStartAdr[7:0]	0

QShpStartAdr[7:0]

Defines the memory port starting 4 long word address for traffic shaping records in the SSRAM context memory. Units in K values.

Warning: This value must not change while the queue engine is enabled.

QLClassStartAdr[7:0]

Defines the memory port starting 4 long word address for loop class records in the SSRAM context memory. Units in K values.

Warning: This value must not change while the queue engine is enabled.

QCellStartAdr[7:0]

Defines the memory port starting 4 long word address for cell records in the SSRAM context memory. Units in K values

Warning: This value must not change while the queue engine is enabled.

QNumVCSel

Selects the maximum number of VCs. When QNumVCSel = 0, a maximum of 64K VCs is supported. When QNumVCSel = 1, a maximum of 16K VCs is supported

Warning: This value must not change while the queue engine is enabled.

Warning: When only 16K VCs are enabled, only the first 14 LSB of the ICI is used to identify the cells. The remaining 2 MSB are ignored.

QRxTxArbSel

Selects between 2 arbitration schemes. When QRxTxArbSel = 0, RR is used to select between WAN, Loop and SAR. When QRxTxArbSel = 1, RR is used to select between WAN and Loop, and SAR having low priority.

QSglStep

Forces the queue engine to process the winner of the service arbitration and then halt. When invoked, the register QEngEn will be set to one. When

QBusy returns a one, both QEngEn and QSglStep will be reset back to zero. To be used only for diagnostics.

QBusy

Status signal indicating that the queue engine is currently active.

QEngEn

The QEngEn enables the queue engine for normal operation. When QEngEn = 0 and QBusy = 0, all queue engine operations are disabled. The QEngEn should be set to a '1' for normal operation after the initialization of the S/UNI APEX and the S/UNI APEX's Context Memory. This register will reset to zero if QSglStep is invoked.

Register 0x704: Receive and Transmit Control

Bit	Type	Function	Default
31:14		Unused	0
13	R/W	QLoopRxFilter	0
12	R/W	QLoopRxDis	0
11:10		Unused	0
9	R/W	QWANRxFilter	0
8	R/W	QWANRxDis	0
7:5		Unused	0
4	R/W	QLoopTxDis	0
3:1		Unused	0
0	R/W	QWANTxDis	0

QWANTxDis

When setting this bit to one, all cell request by the WAN transmit interface will be ignored.

QLoopTxDis

When setting this bit to one, all cell request by the Loop transmit interface will be ignored.

QWANRxDis

When setting this bit to one, all cell request by the WAN receive interface will be ignored.

QWANRxFilter

When setting this bit to one, only uP destined cells received via the WAN receive interface will be serviced. Cells destined for the WAN or Loop will be discarded and the general discard count incremented. Used when the device is in redundant mode

QLoopRxDis

When setting this bit to one, all cell request by the Loop receive interface will be ignored.

QLoopRxFilter

When setting this bit to one, only uP destined cells received via the Loop receive interface will be serviced. Cells destined for the WAN or Loop will be discarded and the general discard count incremented. Used when the device is in redundant mode

Register 0x710: Max Direction Congestion Thresholds

Bit	Type	Function	Default
31:24		Unused	0
23:16	R/W	QLoopMaxThrsh [7:0]	0
15:8		Unused	0
7:0	R/W	QWANMaxThrsh [7:0]	0

This register contains all the maximum per-Direction congestion threshold values for the WAN and Loop ports. Please refer to section 14.2 for a complete definition of m bit logarithmic, n bit fractional.

QWANMaxThrsh[7:0]

Sets the maximum threshold for cells destined for all 4 WANs combined. 4 bit logarithmic, 4 bit fractional.

QLoopMaxThrsh[7:0]

Sets the maximum threshold for cells destined for all 2K Loops combined. 4 bit logarithmic, 4 bit fractional.

Register 0x714: CLP0 Direction Congestion Thresholds

Bit	Type	Function	Default
31:27		Unused	0
23:16	R/W	QLoopCLP0Thrs h [7:0]	0
15:8		Unused	0
7:0	R/W	QWANCLP0Thr sh [7:0]	0

This register contains all the CLP0 per-Direction congestion threshold values for the WAN and Loop ports. Please refer to section 14.2 for a complete definition of m bit logarithmic, n bit fractional.

QWANCLP0Thrsh[7:0]

Sets the EPD threshold for CLP0 cells destined for all 4 WANs combined. 4 bit logarithmic, 4 bit fractional.

QLoopCLP0Thrsh[7:0]

Sets the EPD threshold for CLP0 cells destined for all 2K Loops combined. 4 bit logarithmic, 4 bit fractional.

Register 0x718: CLP1 Direction Congestion Thresholds

Bit	Type	Function	Default
31:27		Unused	0
23:16	R/W	QLoopCLP1Thrs h [7:0]	0
15:11		Unused	0
7:0	R/W	QWANCLP1Thr sh [7:0]	0

This register contains all the CLP1 per-Direction congestion threshold values for the WAN and Loop ports. Please refer to section 14.2 for a complete definition of m bit logarithmic, n bit fractional.

QWANCLP1Thrsh[7:0]

Sets the EPD threshold for CLP1 cells destined for all 4 WANs combined. 4 bit logarithmic, 4 bit fractional.

QLoopCLP1Thrsh[7:0]

Sets the EPD threshold for CLP1 cells destined for all 2K Loops combined. 4 bit logarithmic, 4 bit fractional.

Register 0x71C: Re-assembly Maximum Length

Bit	Type	Function	Default
31:10		Unused	0
9:0	R/W	QReasMaxLen[9:0]	0

QReasMaxLen[9:0]

Defines the alternate re-assembly maximum length, in cells. When a VC has frame continuous queuing, the VC context parameter VcReasMaxSize = 1, and the VC's frame length exceeds this register, an EPD will be invoked.

Register 0x720: Watch Dog ICI Patrol Range

Bit	Type	Function	Default
31:16	R/W	QWdEndICI[15:0]	0
15:0	R/W	QWdStartICI[15:0]	0

QWdStartICI[15:0]

Identifies the first ICI to be checked for re-assembly time-outs.

Warning: This value must not change while a patrol is active.

QWdEndICI[15:0]

Identifies the last ICI to be checked for re-assembly time-outs.

Warning: This value must not change while a patrol is active.

Register 0x724: Tear Down Queue ID

Bit	Type	Function	Default
31:16	R/W	QTdICI[15:0]	0
15:4	R/W	QTdPortID[11:0]	0
3:2		Unused	0
1:0	R/W	QTdClassID[1:0]	0

QTdClassID[3:0]

Identifies the class to be torn down from service. Used only when tearing down a class.

Warning: This value must not change while a tear down is active.

QTdPortID[11:0]

Identifies the port id of the class queue that is to be torn down from service. Used only when tearing down a class. See VC Context Record for PortID encoding.

Warning: This value must not change while a tear down is active.

QTdICI[15:0]

Identifies the VC queue that is to be torn down from service. Used only when tearing down a VC.

Warning: This value must not change while a tear down is active.

Register 0x728: Watch Dog / Tear Down Status

Bit	Type	Function	Default
31:6		Unused	0
5	R/W	QTdMode	0
4	R/W(1)	QTdActive	0
3:1		Unused	0
0	R/W(1)	QWdActive	0

QWdActive

When setting this bit to one, the watch dog will begin its patrol. When the patrol is over, this bit will reset to zero. This bit cannot be reset to zero by the microprocessor.

QTdActive

When setting this bit to one, the tear down macro, as defined in the QTdMode, will be initiated. When the macro has completed its tear down, this bit will reset to zero. This bit cannot be reset to zero by the microprocessor.

QTdMode

When setting this bit to zero, a tear down macro will remove a VC queue, as identified in the Tear Down Queue ID register, from service, re-allocate the buffers in the queue, and update the general discard count. When setting this bit to a one, a tear down macro will remove a class queue, as identified in the Tear Down Queue ID register, from service, re-allocate the buffers in the queue, and update the general discard count.

Warning: This value must not change while a tear down is active.

Warning:

- 1) Tearing down a class queue should only be done after all the VCs within the class has been torn down first.
- 2) After performing a VC tear down, do not setup another VC until either the class queue has been torn down or drain the class queue until VcClassQCLP01Cnt = 0. Failure to do this will result in anomalies in the new VC, such as inaccurate VC weights (for WFQ VCs) and premature discard at the per-VC hierarchical level.

Register 0x730: Shaper 0 Configuration (N = 0)

Register 0x734: Shaper 1 Configuration (N = 1)

Register 0x738: Shaper 2 Configuration (N = 2)

Register 0x73C: Shaper 3 Configuration (N = 3)

Bit	Type	Function	Default
31	R	QShpNExpCong	0
30:29		unused	
28:20	R/W	QShpNRTRate[8:0], N = 0..3	0
19:18		Unused	0
17:16	R/W	QShpNRedFact[1:0], N = 0..3	0
15:12	R/W	QShpNThrshVal[3:0], N = 0..3	0
11:8	R/W	QShpNMeasInt[3:0], N = 0..3	0
7:6	R/W	QShpNClass[1:0], N = 0..3	0
5:4	R/W	QShpNPort[1:0], N = 0..3	0
3		Unused	0
2	R/W	QShpNThrshEn, N = 0..3	0
1	R/W	QShpNSlowDownEn, N = 0..3	0
0	R/W	QShpNEn,	0

Bit	Type	Function	Default
		N = 0..3	

There are four shapers, N = 0..3. Each shaper can be assigned to the same port, or can be assigned to different ports. Only one shaper may be assigned to any one port/class combination. Only WAN ports may be shaped. If a port/class combination is to be shaped, the VcQueue parameter defined in the VC context record structure will be overridden.

Note: If one or more shaper is enabled, then all 4 registers must be programmed with unique port/class combinations.

QShpNEn, N=0..3

Enable for shaper N.

Warning:

This register can only be modified during a port/class setup, and when the queue engine is disabled.

QShpNSlowDownEn, N=0..3

Enables the slow down of the time reference clock used by shaper N to calculate transmission events. Enabling this feature will provide fair shaping to high speed VCs

QShpNThrshEn, N=0..3

Defines how congestion is to be declared for shaper N. When set to zero, the declaration is based on the comparison between the current class queue length and the number of cells leaving the class queue over the previous measurement period. When set to one, the declaration is based on the comparison between the current class queue length and the QShpThrshVal.

This register has no effect if QShpNSlowDownEn = 0.

QShpNPort[1:0]; N=0..3

WAN Port that shaper N is linked to.

Warning:

This register can only be modified during a port/class setup, and when the queue engine is disabled.

QShpNClass[1:0], N=0..3

Class that shaper N is associated with.

Warning:

This register can only be modified during a port/class setup, and when the queue engine is disabled.

QShpNMeasInt[3:0], N=0..3

Encoded value defining the number of clock cycles over which to measure congestion levels for shaper N. The period defined is not affected by the slow down factor. Encoding formula is $64 * 2^m$.

This register has no effect if QShpNSlowDownEn = 0.

QShpNThrshVal[3:0], N=0..3

Defines a value for the class queue length threshold value. Used to select whether to speed up or slow down the shaper N.

$$\text{Effective value} = 2^{(T+1)} - 1, \text{ where } T = \text{QShpNThrshVal}[3:0]$$

This register has no effect if QShpNSlowDownEn = 0 or QShpNThrshEn = 0.

QShpNRedFact[1:0], N=0..3

Encoded value for the slow down rate reduction factor. Controls how quickly the shaper N speeds up. Low values will produce better fairness at the cost of utilization.

QShpNRedFact[1:0]	Effective Value
0	2
1	4
2	8
3	16

This register has no effect if QShpNSlowDownEn = 0.

QShpNRTRate[8:0], N=0..3

Real time rate for shaper N. Represents the maximum shaped data rate, calculated in the number of clock cycles per timeslot. The sum of all active real time rate must be less than the link rate of 1.42Mcells/s. The table below provides the decimal setting for various configurations:

SYSCLK (MHz)	Max. Shaped Rate	QShpNRTRate (decimal)	Max. Active Shapers
80	1.42 Mcells/s	57	1
80	355 Kcells/s	227	4
40	355 Kcells/s	114	2

where

$$QShpNRTRate = \text{ROUNDUP}(\text{SYSCLK (MHz)} / \text{Effective Data Rate}).$$

QShpNExpCong, N = 0..3

Status bit indicating shaper N has experienced congestion, and that the time slot counter was slowed down. This bit is cleared when this register is read.

12.7 Memory Interface

Register 0x800: SDRAM/SSRAM Configuration

Bit	Type	Function	Default
31:17		Unused	0
16	R/W	CMLateWrite	0
15		Unused	0
14:8	R/W	CBRefDivide[6:0]	0
7:1		Unused	0
0	R/W	CBEn	0

CBEn

The CBEn enables the SDRAM Interface. A transition from 0 to 1 initiates the SDRAM initialization procedures. This enable is provided to ensure that the power-up time before the initialization sequence is applied to the SDRAM is met. When CBEn = '0', no SDRAM accesses will take place and the chip will not operate properly.

CBRefDivide[6:0]

Defines the SYSCLK divide down factor to determine the SDRAM refresh rate. Actual divide down value is the value stored in the register multiplied by 16 decimal, plus 1. For example, a value of 78 decimal will produce a refresh cycle every $78 \times 16 + 1 = 1249$ SYSCLKs. A zero value is not permitted.

CMLateWrite

The CMLateWrite selects the type of SSRAM connected. When set to zero, pipelined ZBT SSRAM is configured. When set to one, register to register late write SSRAM is configured.

When late write is selected, pins CMAB[18] & CMAB[17] change functionality and become the chip enable bar for odd and even addresses, respectively.

12.8 Test Interface

The following registers are for production test only. They should not be invoked while in service.

Register 0x900: BIST_OK

Bit	Type	Function	Default
31:8		Unused	0
7	R	McLRBistOK	0
6	R	McLTBistOK	0
5	R	McWRBistOK	0
4	R	McWTBistOK	0
3	R	McCRBistOK	0
2	R	McCTBistOK	0
1	R	ALBistOK	0
0	R	QEBistOK	0

XXBistOK

A high value indicates that the test is complete and that no error has been encountered. For typical BIST operation, sequence of write should be as follows:

XXBistEnb = 1, BistMode = 0x4

XXBistEnb = 0, BistMode = 0x2

Wait 16*depth SYSCLKs, where depth is the address space of the deepest RAM under test. This register requires an active SYSCLK to access.

Register 0x904: BIST_Mode

Bit	Type	Function	Default
31	R/W	BistSide	0
30:28		Unused	0
27	R/W	McLRBistEnb	1
26	R/W	McLTBistEnb	1
25	R/W	McWRBistEnb	1
24	R/W	McWTBistEnb	1
23	R/W	McCRBistEnb	1
22	R/W	McCTBistEnb	1
21	R/W	ALBistEnb	1
20	R/W	QEBistEnb	1
19:16		Unused	0
15:8	R/W	BistTest[7:0]	0
7:3		Unused	0
2:0	R/W	BistMode[2:0]	0

BistSide

For dual port RAMs, selects the side that is to be tested. When set to 0, PortA will be tested. When set to 1, PortB will be tested.

XXBistEnb

When set to 0, enables the internal RAMBIST for the RAMs within the specific block.

BistTest[7:0]

Seed test pattern bus that is replicated over the RAM data bus

BistMode

Defines the Bist mode of operation.

Register 0x908: BIST_Results A

Bit	Type	Function	Default
31:21		Unused	0
20	R	ALBistResult	1
19	R	ALBistEnd	0
18:9	R	ALBistError[9:0]	0
8:7		Unused	0
6	R	QEBistResult	1
5	R	QEBistEnd	0
4:1	R	QEBistError[3:0]	0
0		Unused	0

XXBistResult

When BISTEnb is enabled for the specific block, a high value indicates that no errors have been encountered.

XXBistEnd

A high value indicates the end of the BIST test.

XXBistError

A high value indicates that the current clock cycle of the RAM undertest has encountered an error.

Register 0x90C: BIST_Results B

Bit	Type	Function	Default
31:24		Unused	0
23	R	McLRBistResult	1
22	R	McLRBistEnd	0
21	R	McLRBistError	0
20		Unused	0
19	R	McLTBistResult	1
18	R	McLTBistEnd	0
17	R	McLTBistError	0
16		Unused	0
15	R	McWRBistResult	1
14	R	McWRBistEnd	0
13	R	McWRBistError	0
12		Unused	0
11	R	McWTBistResult	1
10	R	McWTBistEnd	0
9	R	McWTBistError	0
8		Unused	0
7	R	McCRBistResult	1
6	R	McCRBistEnd	0
5	R	McCRBistError	0
4		Unused	0
3	R	McCTBistResult	1
2	R	McCTBistEnd	0
1	R	McCTBistError	0
0		Unused	0

See register BIST_Result A.

12.9 CBI Interface

Register 0xA00: CBI Register Port

Bit	Type	Function	Default
31	R/W	CBIBusy	0
30	R/W	CBIRdWrb	0
29:16		Unused	0
15:14	R/W	CBIAddr[1:0]	0
13	R/W	CBITrsb	0
12		Unused	0
11	R/W	CBITstb	1
10:8		Unused	0
7:0	R/W	CBIData[7:0]	0

Writes to this register will place the microprocessor interface in wait state until the CBIBusy bit is clear.

CBIData[7:0]

Data value used to write into the CBI port or read from the CBI port. Address specified by CBIAddr[2:0]

CBITstb

Active low signal used during production testing. Normally set high.

CBIAddr[1:0]

Address bus used to select the CBI register.

CBITrsb

Test select bit. Should always be set to 1 during normal mode access.

CBIRdWrb

Defines whether a read or write to the CBI register is to be performed. When set to 0, a write command is defined. When set to 1 a read command is defined. Command invoked when CBIBusy is set to 1.

CBIBusy

When set to a one, the CBIRdWrb command will be executed. When the command is complete, this bit will be cleared to zero.

13 CBI REGISTER PORT MAPPING

CBI Register 0x00: Configuration

Bit	Type	Function	Default
Bit 7		Unused	X
Bit 6		Unused	X
Bit 5	R/W	Reserved	0
Bit 4	R/W	OVERRIDE	0
Bit 3		Unused	X
Bit 2	R/W	ERRORE	X
Bit 1	R/W	VERN_EN	0
Bit 0	R/W	LOCK	0

The DLL Configuration Register controls the basic operation of the DLL.

LOCK:

The LOCK register is used to force the DLL to ignore phase offsets indicated by the phase detector after phase lock has been achieved. When LOCK is set to logic zero, the DLL will track phase offsets measured by the phase detector between the SYSCLK and the REFCLK inputs. When LOCK is set to logic one, the DLL will not change the tap after the phase detector indicates of zero phase offset between the SYSCLK and the REFCLK inputs for the first time.

VERN_EN:

The vernier enable register (VERN_EN) forces the DLL to ignore the phase detector and use the tap number specified by the VERNIER[7:0] register bits. When VERN_EN is set to logic zero, the DLL operates normally adjusting the phase offset based on the phase detector. When VERN_EN is set to logic one, the delay line uses the tap specified by the VERNIER[7:0] register bits.

Used only for diagnostics.

ERRORE:

The ERROR interrupt enable (ERRORE) bit enables the error indication interrupt. When ERRORE is set high, an interrupt is generated upon assertion event of the ERR output and ERROR register. When ERRORE is set low, changes in the ERROR and ERR status do not generate an interrupt.

OVERRIDE:

The override control (OVERRIDE) disables the DLL operation. When OVERRIDE is set low, the DLL generates the DLLCLK by delaying the SYSCLK until the rising edge of REFCLK occurs at the same time as the rising edge of SYSCLK. When OVERRIDE is set high, the DLLCLK output is a buffered version of the SYSCLK input.

Used only for diagnostics.

CBI Register 0x01: Vernier Control

Bit	Type	Function	Default
Bit 7	R/W	VERNIER[7]	0
Bit 6	R/W	VERNIER[6]	0
Bit 5	R/W	VERNIER[5]	0
Bit 4	R/W	VERNIER[4]	0
Bit 3	R/W	VERNIER[3]	0
Bit 2	R/W	VERNIER[2]	0
Bit 1	R/W	VERNIER[1]	0
Bit 0	R/W	VERNIER[0]	0

The Vernier Control Register provides the delay line tap control when using the vernier option.

Used only for diagnostics.

VERNIER[7:0]:

The vernier tap register bits (VERNIER[7:0]) specifies the phase delay through the DLL when using the vernier feature. When VERN_EN is set high, the VERNIER[7:0] registers specify the delay tap used. When VERN_EN is set low, the VERNIER[7:0] register is ignored. A VERNIER[7:0] value of all zeros specifies the delay tap with the minimum delay through the delay line. A VERNIER[7:0] value of 255 specifies the delay tap with the maximum delay through the delay line.

CBI Register 0x02: Delay Tap Status

Bit	Type	Function	Default
Bit 7	R	TAP[7]	X
Bit 6	R	TAP[6]	X
Bit 5	R	TAP[5]	X
Bit 4	R	TAP[4]	X
Bit 3	R	TAP[3]	X
Bit 2	R	TAP[2]	X
Bit 1	R	TAP[1]	X
Bit 0	R	TAP[0]	X

The DLL Delay Tap Status Register indicates the delay tap used by the DLL to generate the outgoing clock.

Writing to this register performs a software reset of the DLL. A software reset requires a maximum of 24*256 SYSCLK cycles for the DLL to regain lock. During this time the DLLCLK phase is adjusting from its current position to delay tap 0 and back to a lock position.

Used only for diagnostics.

TAP[7:0]:

The tap status register bits (TAP[7:0]) specifies the delay line tap the DLL is using to generate the outgoing clock DLLCLK. When TAP[7:0] is logic zero, the DLL is using the delay line tap with minimum phase delay. When TAP[7:0] is equal to 255, the DLL is using the delay line tap with maximum phase delay. TAP[7:0] is invalid when vernier enable VERN_EN is set to one.

CBI Register 0x03: Control Status

Bit	Type	Function	Default
Bit 7	R	SYSCCLKI	X
Bit 6	R	REFCLKI	X
Bit 5	R	ERRORI	X
Bit 4	R	CHANGEI	X
Bit 3		Unused	X
Bit 2	R	ERROR	X
Bit 1	R	CHANGE	0
Bit 0	R	RUN	0

The DLL Control Status Register provides information of the DLL operation. Used only for diagnostics.

RUN:

The DLL lock status register bit (RUN) indicates the DLL found a delay line tap in which the phase difference between the rising edge of REFCLK and the rising edge of SYSCCLK is zero. After system reset, RUN is logic zero until the phase detector indicates an initial lock condition. When the phase detector indicates lock, RUN is set to logic 1. Maximum time for a DLL lock from reset with a stable clock should be under $12 * 256 / f(\text{SYSCCLK})$. The RUN register bit is cleared only by a system reset (CBI[12]) or a software reset (writing to register 2).

CHANGE:

The delay line tap change register bit (CHANGE) indicates the DLL has moved to a new delay line tap. CHANGE is set high for eight SYSCCLK cycles when the DLL moves to a new delay line tap.

ERROR:

The delay line error register bit (ERROR) indicates the DLL has run out of dynamic range. When the DLL attempts to move beyond the end of the delay line, ERROR is set high. When ERROR is high, the DLL cannot generate a DLLCLK phase that causes the rising edge of REFCLK to be aligned to the rising edge of SYSCCLK. ERROR is set low, when the DLL captures lock again.

CHANGEI:

The delay line tap change event register bit (CHANGEI) indicates the CHANGE register bit has changed value. When the CHANGE register changes from a logic zero to a logic one, the CHANGEI register bit is set to logic one. The CHANGEI register bit is cleared immediately after it is read, thus acknowledging the event has been recorded.

ERRORI:

The delay line error event register bit (ERRORI) indicates the ERROR register bit has gone high. When the ERROR register changes from a logic zero to a logic one, the ERRORI register bit is set to logic one. The ERRORI register bit is cleared immediately after it is read, thus acknowledging the event has been recorded. If the ERRERE interrupt enable is high, the INT output is also asserted when ERRORI asserts.

REFCLKI:

The reference clock event register bit REFCLKI provides a method to monitor activity on the reference clock. When the REFCLK primary input changes from a logic zero to a logic one, the REFCLKI register bit is set to logic one. The REFCLKI register bit is cleared immediately after it is read, thus acknowledging the event has been recorded.

SYSCLKI:


The system clock event register bit SYSCLKI provides a method to monitor activity on the system clock. When the SYSCLK primary input changes from a logic zero to a logic one, the SYSCLKI register bit is set to logic one. The SYSCLKI register bit is cleared immediately after it is read, thus acknowledging the event has been recorded.

14 MEMORY PORT MAPPING

14.1 Context Size and Location

The context records for S/UNI APEX are stored in 4 different areas. The majority of the context information is stored in an external SSRAM, the context memory. For performance reasons, some context information is stored in internal memories. There are three internal memories, the Queue context, the WAN Port scheduler context, and the Loop Port scheduler context

Figure 28 - Context Location



The records for the external context memory are partitioned into four groups: VC, Loop Ports, Shape and Cells. Three registers define the starting physical address of each of the Loop class, Shape and Cell groups, with the VC group always starting at address zero.

The starting address, or offset, may be of any order, and may, under certain restriction, overlap one another. For example, if the VC Address Map will never be enabled for a given application, then the loop class offset may be set to the

starting location of the VC Address Map. Should the shaping not be required, then cell offset may follow immediately after the loop class records.

All the Port context records and all the class context records for the WAN and microprocessor are located internally, and are accessed via the appropriate aperture in the context memory port.

Table 17 - External Queue Context Memory Map

MPQuadAddr		Context Record
16K VC	64K VC	
0..16383	0..65535	VC Context
16384..24575	65536..98303	VC Statistic
24576..28671	98304..114687	VC Address Map
QLClassStartAdr * 1024 + {0..2047}		Loop Class0
QLClassStartAdr * 1024 + {2048..4095}		Loop Class1
QLClassStartAdr * 1024 + {4096..6143}		Loop Class2
QLClassStartAdr * 1024 + {6144..8191}		Loop Class3
QShpStartAdr * 1024 + {0..2047}		Shape0 TxSlot
QShpStartAdr * 1024 + {2048..4095}		Shape1 TxSlot
QShpStartAdr * 1024 + {4096..6143}		Shape2 TxSlot
QShpStartAdr * 1024 + {6144..8191}		Shape3 TxSlot
QShpStartAdr * 1024 + {8192..40959}		Shape Rate
QCellStartAdr * 1024 + {0..65535}		Cell Record

Table 18 - Internal Queue Context Memory Map

MPQuadAddr	Record
0..511	Loop Class Scheduler
512..1023	Loop Port Threshold
1024..1535	Loop Port Count
1536..1539	WAN Class0
1540..1543	WAN Class1

MPQuadAddr	Record
1544..1547	WAN Class2
1548..1551	WAN Class3
1552	WAN Class Scheduler
1553	WAN Port Threshold
1554	WAN Port Count
1555	uP Class Scheduler
	uP Port Threshold
	uP Port Count
1556	uP Class0
1557	uP Class1
1558	uP Class2
1559	uP Class3
2048	Free Count
2049	Overall Count
2050	Congestion Discard
2051	Maximum Congestion ID
2052	Misc Error

Table 19 - Internal WAN Port Scheduler Context Memory Map

MPQuadAddr	Record
0	WAN Poll Weight, WAN Class Status

Table 20 - Internal Loop Port Scheduler Context Memory Map

MPQuadAddr	Record
0..127	Loop Poll Sequence
128..191	Loop Poll Weight
192..255	Loop Class Status

Writing values into unused parameter bits has no effect. However, to ensure software compatibility with future, feature-enhanced versions of the product, unused parameter bits must be written with logic zero. Reading back unused bits can produce either a logic one or a logic zero; hence, unused parameter bits should be masked off by software when read.

14.2 Queue Context Definition

Many of the context parameters make references to m bit logarithmic, n bit fractional, where m is the MSB of the parameter field, and n is the LSB of the parameter field. The tables below provide quick references.

Table 21 - 2 Bit Logarithmic, 2 Bit Fractional

		2 bits fractional			
		0	1	2	3
2 bits log	0	0	1	2	3
	1	4	5	6	7
	2	8	10	12	14
	3	16	20	24	28

Table 22 - 4 Bit Logarithmic, 2 Bit Fractional

		2 bits fractional			
		0	1	2	3
4 bits log	0	0	1	2	3
	1	4	5	6	7
	2	8	10	12	14
	3	16	20	24	28
	4	32	40	48	56
	5	64	80	96	112
	6	128	160	192	224
	7	256	320	384	448
	8	512	640	768	896
	9	1024	1280	1536	1792
	10	2048	2560	3072	3584
	11	4096	5120	6144	7168
12	8191				

Table 23 - 4 Bit Logarithmic, 4 Bit Fractional

		4 bits fractional															
		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
4	0	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

1	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
2	32	34	36	38	40	42	44	46	48	50	52	54	56	58	60	62
3	64	68	72	76	80	84	88	92	96	100	104	108	112	116	120	124
4	128	136	144	152	160	168	176	184	192	200	208	216	224	232	240	248
5	256	272	288	304	320	336	352	368	384	400	416	432	448	464	480	496
6	512	544	576	608	640	672	704	736	768	800	832	864	896	928	960	992
7	1024	1088	1152	1216	1280	1344	1408	1472	1536	1600	1664	1728	1792	1856	1920	1984
8	2048	2176	2304	2432	2560	2688	2816	2944	3072	3200	3328	3456	3584	3712	3840	3968
9	4096	4352	4608	4864	5120	5376	5632	5888	6144	6400	6656	6912	7168	7424	7680	7936
10	8192	8704	9216	9728	10240	10752	11264	11776	12288	12800	13312	13824	14336	14848	15360	15872
11	16384	17408	18432	19456	20480	21504	22528	23552	24576	25600	26624	27648	28672	29696	30720	31744
12	32768	34816	36864	38912	40960	43008	45056	47104	49152	51200	53248	55296	57344	59392	61440	63488
13	65536	69632	73728	77824	81920	86016	90112	94208	98304	102400	106496	110592	114688	118784	122880	126976
14	131072	139264	147456	155648	163840	172032	180224	188416	196608	204800	212992	221184	229376	237568	245760	253952
15	262143															

14.2.1 VC Context Records

14.2.1.1 VC Context Record

MPMemSelect = External Queue Context

MPQuadAddr = ICI

Table 24 - VC Context Record Structure

MPLWord En (bit #)	Bits	Parameter	Description
0	33:32	Unused	Reserved
	31	VcEn	Enables the VC. When VcEn = 0, incoming cells of this VC will be discarded, CellRxErrICI updated, DiscardCnt incremented, and the maskable interrupt QCellRxErr will be generated. Any cells remaining in the queue will be transmitted, if possible.
	30	VcIntDis	Disables any maskable interrupts that this VC may generate, including QCellRxErr, QVcReasLenErr, QVcMaxThrsErr.

MPLWord En (bit #)	Bits	Parameter	Description
	29	VcSegOam	Segment OAM re-direct. When VcSegOam = 1, any segment OAM encountered in this VC will be re-directed to the uP. Warning: This bit must be never be set to one if the VC is destined for the uP class 0 port.
	28	VcEEOam	End to end OAM re-direct. When VcEEOam = 1, any end to end OAM encountered in this VC will be re-directed to the uP. Warning: This bit must be never be set to one if the VC is destined for the uP class 0 port.
	27	VcVPC	When VcVPC = 1, identifies this VC as VPC. When VcVPC = 0, identifies this VC as a VCC. Required for OAM re-direction.
	26:24	VcCLP0MinThrsh	Minimum number of CLP0 cells guaranteed to be allowed on a per-VC basis. Values are encoded as follows: 000: 0 001: 24 010: 32 011: 48 100: 64 101: 96 110: 128 111: 256
	23:18	VcMaxThrsh	PPD maximum threshold for cells on a per-VC basis. 4 bits logarithmic/linear, 2 bits fractional. A zero value will effectively disable congestion at the VC level for this VC; however, the 8K-1 maximum limit still remains active.
	17:12	VcCLP0Thrsh	EPD maximum threshold for CLP0 cells on a per-VC basis. 4 bits logarithmic/linear, 2 bits fractional.
	11:6	VcCLP1Thrsh	EPD maximum threshold for CLP1 cells on a per-VC basis. 4 bits logarithmic/linear, 2 bits fractional.

MPLWord En (bit #)	Bits	Parameter	Description
	5:4	VcEFCIMode	<p>Defines EFCI marking.</p> <p>00: No marking 01: Reserved 10: Mark on active CLP0 thresholds 11: Mark on active CLP1 thresholds</p> <p>Note: active thresholds is defined as the hierarchical maximum threshold set to a non-zero value.</p>
	3:2	VcReMapMode	<p>00: ECI=ICI. No VPI/VCI remapping. SwTag available.</p> <p>01: Provide an ECI that is different from the ICI. No VPI/VCI remapping. SwTag available.</p> <p>10: ECI=ICI. VPI header remapping. SwTag available.</p> <p>11: ECI=ICI. VPI/VCI header remapping. SwTag not available</p>
	1	VcCongMode	<p>Defines congestion mode when FCQ is not selected.</p> <p>0: EPD/PPD. Congestion control is set to discard using EPD/PPD.</p> <p>1: Congestion control operates on a cell by cell basis.</p> <p>When FCQ is selected, this bit is reserved.</p> <p>Warning: This bit must only be set during a VC setup. This bit must be set to 1 if the underlying traffic is not AAL5.</p>
	0	VcGFRMode	<p>Selects between I.363 standard definition of a frame's CLP, and the emerging GFR standard definition of a frame's CLP. Zero selects the I.363 standard, a value of one selects the emerging GFR standard.</p> <p>This bit is reserved when cell discard congestion rules is selected.</p>

MPLWord En (bit #)	Bits	Parameter	Description
1	33:27	VcCLP0Cnt[12:6]	<p>Last 7 bits of a 13-bit count of CLP0 cells in both VC and Class. Used for congestion control for minimum resource monitoring.</p> <p>VC tear down and watch dog re-allocation will reduce this count by VcQCLP01Cnt if the VcRxBOMClp = 0 and VcQueue = 1 (FCQ).</p>
	26:15	VcPortID	<p>VcPortID[11] = 0.</p> <p>VC destined for a Loop port.</p> <p>VcPortID[10:0]: Identifies 1 of 2K loop ports as the destination.</p> <p>If the loop port is configured for Any-PHY, all 11 bits are valid.</p> <p>If the loop transmit port is configured for UTOPIA L2 master, only the first 5 bits (ports 31->0) are valid. Cells received in VCs with VcPortID > 31 will never be transmitted.</p> <p>If the loop transmit port is configured for UTOPIA L1 master or UTOPIA L2 slave, VcPortID[10:0] must be set to 0. Cells received in VCs with VcPortID != 0 will never be transmitted.</p>
			<p>VcPortID[11:10] = 10.</p> <p>VC destined for a WAN ports</p> <p>VcPortID[9:2]: Reserved, Must be initialized to zero.</p> <p>VcPortID[1:0]: Identifies 1 of 4 WAN ports as the destination.</p> <p>If the WAN port is configured for UTOPIA L1 master or UTOPIA L2 slave, VcPortID[10:0] must be set to 0. Cells received in VCs with VcPortID != 0 will never be transmitted.</p>
			<p>VcPortID[11:10] = 11.</p> <p>Identifies the uP port as the VC's destination</p> <p>VcPortID[9:0]: Reserved, Must be initialized to zero.</p>
14:13	VcClass	Assigns class to this VC	
	12		Destination Port Class is not shaped

MPLWord En (bit #)	Bits	Parameter	Description
		VcQueue	<p>Defines the queue mechanism to be used.</p> <p>0: Weighted Fair Queuing (WFQ)</p> <p>1: Frame Continuous Queuing (FCQ)</p> <p>Warning: This bit must only be set during a VC setup.</p>
		Destination Port Class is shaped	
		Unused	Reserved
	11:6	VcQueue = 0 (WFQ), and destination Port Class is not shaped	
	11:6	VcWght	<p>Class queuing weight. Linear encoding, multiple of 2 with zero taking a special value of 1. (eg. 1, 2, 4, 6 ... 124, 126)</p> <p>Warning: This bit must only be set during a VC setup.</p>
		VcQueue = 1 (FCQ), and destination Port Class is not shaped	
	11	VcReasPark	<i>Park state bit for re-assembly watch dog. Non-user cells have no effect on this bit.</i>
	10	VcReasActive	<i>Indication that there is at least one cell in the queue. VC tear down and watch dog re-allocation will set this parameter to zero (empty).</i>
	9	VcReasMaxSiz	<p>Defines one of two maximum frame sizes. When this parameter is zero, frames exceeding 1366 cells will have EPD invoked. When this parameter is one, frames exceeding the value programmed in the Re-assembly Maximum Length register will have EPD invoked. EPDs will trigger a maskable interrupt to the microprocessor.</p>
	8	VcWDEn	Marks this VC as one for the watch dog to patrol for re-assembly time outs.
	7	Unused	Reserved

MPLWord En (bit #)	Bits	Parameter	Description
		6 VcLenChkEn	Enables the checking of the AAL5 trailer's length field. 0: No length checking. 1: If the length field is zero, EPD will be invoked.
		Destination Port Class is shaped	
		11:6 Unused	Reserved
	5:0	VcCongMode = 0 (EPD/PPD Congestion)	
		5 VcTxStatus	<i>EOM indication. When VcTxStatus = 0, the next cell transmitted is considered the BOM. Used to indicate when VcTxBOMClp should be re-evaluated.</i>
		4 VcTxBOMClp	<i>State of the BOM's CLP bit during transmit. Used to maintain VcCLP0Cnt during FCQ congestion and also during EPD/PPD congestion when GFR Mode = 1.</i>
		3:2 VcRxStatus	<i>00: Next cell queued is considered the BOM. 01: Next cell queued is considered a COM or EOM. 10: Early packet discard. All cells up to and including the EOM will be discarded. 11: Partial packet discard. All cells but the EOM will be discarded VC tear down and watch dog re-allocation will set this parameter to zero.</i>
		1 VcRxBOMClp	<i>State of the BOM's CLP bit during receive. Used to maintain VcCLP0Cnt during FCQ congestion and also during EPD/PPD congestion when GFR Mode = 1.</i>
		0 VcRxORClp	<i>VcQueue = 1 (FCQ), VcGFRMode = 0, Port Class not shaped: This is a running OR of all the CLPs that have been received since the BOM. Used by congestion control to invoke EPD discard Reserved in all other cases.</i>

MPLWord En (bit #)	Bits	Parameter	Description
		VcCongMode = 1 (Cell by Cell Congestion)	
		5:0 Unused	Reserved
2	33:16	<i>VcQHeadPtr</i>	<i>Points to the first cell in the VC queue. A value of zero indicates that the queue is empty. VC tear down and watch dog re-allocation will set this parameter to zero.</i>
	15:13	<i>VcCLP0Cnt[5:3]</i>	<i>Second 3 bits of a 13-bit count of CLP0 cells in both VC and Class. See MSB for description.</i>
	12:0	<i>VcQCLP01Cnt</i>	<i>Count of CLP01 cells in the VC queue. Used for congestion control. VC tear down and watch dog re-allocation will set this parameter to zero.</i>
3	33:16	<i>VcQTailPtr</i>	<i>Points to the last cell in the VC queue</i>
	15:13	<i>VcCLP0Cnt[2:0]</i>	<i>First 3 bits of a 13-bit count of CLP0 cells in both VC and Class. See MSB for description.</i>
	12:0	<i>VcClassQCLP01Cnt</i>	<i>Count of CLP01 cells in the Class queue for WFQ & FCQ. Count of CLP01 cells in the Class queue & TxSlot table combined for SFQ. Used for congestion control for all queuing mechanisms. Note that during SFQ, the maximum value this parameter can be is 1.</i>

Note

Parameters in italics are active context information. Microprocessor must mask these bits out during in-service context modification. During VC setup, these italicized parameters should be reset to zero.

14.2.1.2 VC Statistics Record

MPMemSelect = External Queue Context

MPQuadAddr = 65536 + ICI/2, if QNumVCSel = 0

MPQuadAddr = 16384 + ICI/2, if QNumVCSel = 1

The statistic record for even numbered ICI's are located in long words 0 & 1. Odd numbered ICI's are located in long words 2 & 3.

OAM cells redirected to the uP are not represented by these counts.

Table 25 - VC Statistics Record Structure

MPLWord En (bit #)	Bits	Parameter	Description
0/2	33:32	Unused	Reserved
	31:0	<i>VcCLP0TxCnt</i>	<i>Free running count of all CLP0 cells that have been transmitted.</i>
1/3	33:32	Unused	Reserved
	31:0	<i>VcCLP1TxCnt</i>	<i>Free running count of all CLP1 cells that have been transmitted.</i>

Note

Parameters in italics are active context information. Microprocessor must mask these bits out during in-service context modification. During VC setup, these italicized parameters should be reset to zero.

14.2.1.3 VC Address Map Record

MPMemSelect = External Queue Context

MPQuadAddr = 98304 + ICI/4, if QNumVCSEL = 0

MPQuadAddr = 24576 + ICI/4, if QNumVCSEL = 1

The address map records are packed together. ICI = 0 would have the address map record located in long word 0. ICI = 3 would have the address map record located in long word 3 etc.

Table 26 - VC Address Map Record Structure

MPLWord En (bit #)	Bits	Parameter	Description
0/1/2/3	33:32	Unused	Reserved
	31:0	VcReMapMode = 00	
		31:16 SwTag	Option to provide Switch Tag for the WAN and LOOP Tx Interfaces. Prepend of the Switch Tag is determined on a port by port basis.
	15:0 Unused	Reserved	

MPLWord En (bit #)	Bits	Parameter	Description
		VcReMapMode = 01	
	31:16	SwTag	Option to provide Switch Tag for the WAN and LOOP Tx Interfaces. Prepend of the Switch Tag is determined on a port by port basis.
	15:0	ECI	ECI value that is applied to the cells upon emission.
		VcReMapMode = 10	
	31:16	SwTag	Option to provide Switch Tag for the WAN and LOOP Tx Interfaces. Prepend of the Switch Tag is determined on a port by port basis.
	15:12	Unused	Reserved
	11:0	VPI	VPI remap
		VcReMapMode = 11	
	31:16	VCI	VCI remap
	15:12	Unused	Reserved
	11:0	VPI	VPI remap

14.2.2 Port Context Records

14.2.2.1 Port Threshold Context Record

For Loop Ports:

MPMemSelect = Internal Queue Context

MPQuadAddr = 512 + Loop#/4

For WAN Ports:

MPMemSelect = Internal Queue Context

MPQuadAddr = 1553

For Microprocessor Port:

MPMemSelect = Internal Queue Context

MPQuadAddr = 1555, MPLWordEn[1] = 1.

The port context records are packed together. Loop/WAN = 0 would have the port context record located in long word 0. Loop/WAN = 3 would have the port context record located in long word 3 etc. The microprocessor port would be located in long word 1.

Table 27 - Port Threshold Context Record Structure

MPLWord En (bit #)	Bits	Parameter	Description
0/1/2/3	33:16	Unused	Reserved
	15:8	PortCLP0Thrs h	Per-Port EPD maximum threshold for CLP0 cells that have met their minimum allocation. 4 bit logarithmic, 4 bit fractional encoding.
	7:0	PortCLP1Thrs h	EPD maximum threshold for CLP1 cells on a per-Port basis. 4 bit logarithmic, 4 bit fractional encoding.

14.2.2.2 Port Count Context Record

For Loop Ports:

MPMemSelect = Internal Queue Context

MPQuadAddr = 1024 + Loop#/4

For WAN Ports:

MPMemSelect = Internal Queue Context

MPQuadAddr = 1554

For Microprocessor Port:

MPMemSelect = Internal Queue Context

MPQuadAddr = 1555, MPLWordEn[2] = 1.

The port context records are packed together. Loop/WAN = 0 would have the port context record located in long word 0. Loop/WAN = 3 would have the port

context record located in long word 3 etc. The microprocessor port would be located in long word 2.

Table 28 - Port Count Context Record Structure

MPLWord En (bit #)	Bits	Parameter	Description
0/1/2/3	33:32	Unused	Reserved
	31	PortEn	<p>Enables the port. When PortEn = 1, it permits cells to be received and transmitted if the class and VCs are setup. The PortEn should be set to zero 0 only after the class and VC have been torn down.</p> <p>If the PortEn is set to zero prior to tearing down a VC, then incoming cells will increment either the DiscardCnt or CLPxDiscardCnt. Cells that exist in the queues when the port is disabled may be transmitted. Details are listed below.</p> <p>If PortEn changes from 1->0, the first cell at the head of each class queue shall be transmitted. Any cells remaining in the queues will not be transmitted. Exception to this case is if the port contains a shaped class, where the transmission of cells in the queues will continue until there is one cell per VC remaining.</p> <p>When PortEn = 0, all incoming cells to this port will be discarded, CellRxErrICI updated, and the maskable interrupt QCellRxErr will be sent.</p> <p>For non FCQ VCs in cell discard mode, the DiscardCnt is incremented when a cell is received.</p> <p>For non FCQ VCs in EPD/PPD discard mode, the DiscardCnt is incremented when the first cell is received, and CLPxDiscardCnt is incremented when subsequent cells are received.</p> <p>For FCQ VCs, the CLPxDiscardCnt is incremented when a cell is received. If a frame is in the process of being re-assembled when PortEn = 1->0, the CLPxDiscardCnt will be incremented by the length of the re-assembly queue plus 1 (cell that arrived).</p>
	30:28	Unused	Reserved

MPLWord En (bit #)	Bits	Parameter	Description
	27:20	PortMaxThrsh	Maximum threshold for cells on a per-Port basis. 4 bit logarithmic, 4 bit fractional encoding. A zero value will effectively disable congestion at the Port level for this port.
	19:18	Unused	Reserved
	17:0	<i>PortCnt</i>	<i>Total count of all cells queued for this port. Used for congestion control.</i> <i>VC tear down and watch dog re-allocation will reduce this count by VcQCLP01Cnt.</i> <i>Class tear down will reduce this count by ClassCnt.</i>

Note

Parameters in italics are active context information. Microprocessor must mask these bits out during in-service context modification. During port setup, these italicized parameters should be reset to zero.

14.2.3 Class Context Records

14.2.3.1 Class Scheduler Context Record

For Loop Ports:

MPMemSelect = Internal Queue Context

MPQuadAddr = Loop#/4

For WAN Ports:

MPMemSelect = Internal Queue Context

MPQuadAddr = 1552

For Microprocessor Port:

MPMemSelect = Internal Queue Context

MPQuadAddr = 1555, MPLWordEn[0] = 1.

The port context records are packed together. Loop/WAN = 0 would have the port context record located in long word 0. Loop/WAN = 3 would have the port context record located in long word 3 etc. The microprocessor port would be located in long word 0.

Table 29 - Class Scheduler Record Structure

MPLWord En (bit #)	Bits	Parameter	Description
0/1/2/3	33:32	Unused	Reserved
	31:20	ClassFragEn = 0	
		31:28 Class1CellLmt	Number of cells that can be transmitted from another class while Class1 is waiting for a transmission opportunity before Class1 enters the starvation condition. 2 bit logarithmic, 2 bit fractional encoding. When ClassPacket = 1, this parameter should be set to 0 to force strict priority.

MPLWord En (bit #)	Bits	Parameter	Description
		27:24 Class2CellLmt	Number of cells that can be transmitted from another class while Class2 is waiting for a transmission opportunity before Class2 enters the starvation condition. 2 bit logarithmic, 2 bit fractional encoding. When ClassPacket = 1, this parameter should be set to 0 to force strict priority.
		23:20 Class3CellLmt	Number of cells that can be transmitted from another class while Class3 is waiting for a transmission opportunity before Class3 enters the starvation condition. 2 bit logarithmic, 2 bit fractional encoding. When ClassPacket = 1, this parameter should be set to 0 to force strict priority.
		ClassFragEn = 1	
		31 Unused	Reserved
		30:20 ClassCellLmt	Maximum number of cells that can be transmitted from a class before the class scheduler selects another class to transmit from. A value of zero is functionally equivalent to a value of one.
19		ClassFragEn	Enable packet fragmentation class scheduling.
18		ClassPacket	Places the Class scheduler into a packet or frame continuous mode. In packet mode, if a VC set to FCQ is at the head of the class queue, the class scheduler will send out the entire packet prior to servicing a different class. When set, ClassNCellLmt must be set to zero to ensure strict priority scheduling. Warning: This bit must only be set during a Class setup. This must be set to 1 when ClassFragEn = 1.
	17:15	<i>ClassPacket = 0</i>	
		Unused	Reserved
		<i>ClassPacket = 1</i>	
		17 ClassStatus	<i>When ClassStatus = 0, indicates that the next cell transmitted is considered the BOM of a packet.</i>
<p>PROPRIETARY AND CONFIDENTIAL TO PMC-SIERRA, INC., AND FOR ITS CUSTOMERS' INTERNAL USE</p>			
			174

MPLWord En (bit #)	Bits	Parameter	Description
		<i>16:15 ClassServiced</i>	<i>Indicates the class that is currently being serviced.</i>
	14:0	<i>ClassFragEn = 0</i>	
		<i>14:10 Class1CellCnt</i>	<i>Current number of cells that have been transmitted while Class1 has remained unserved.</i>
		<i>9:5 Class2CellCnt</i>	<i>Current number of cells that have been transmitted while Class2 has remained unserved.</i>
		<i>4:0 Class3CellCnt</i>	<i>Current number of cells that have been transmitted while Class3 has remained unserved.</i>
		<i>ClassFragEn = 1</i>	
		<i>14:11 Unused</i>	<i>Reserved</i>
		<i>10:0 ClassCellCnt</i>	<i>Current number of cells that have been transmitted from the class indicated by ClassServiced.</i>

Note

Parameters in italics are active context information. Microprocessor must mask these bits out during in-service context modification. During class setup, these italicized parameters should be reset to zero.

14.2.3.2 Class 0 through 3 Context Record

For Loop Ports:

MPMemSelect = External Queue Context

MPQuadAddr = QLClassStartAdr * 1024 + Class * 2048 + Loop#

For WAN Ports:

MPMemSelect = Internal Queue Context

MPQuadAddr = 1536 + Class * 4 + WAN#, Class = 0,1,2,3

For Microprocessor Ports:

MPMemSelect = Internal Queue Context

MPQuadAddr = 1556 + Class * 4, Class = 0,1,2,3

Table 30 - Class Context Record Structure

MPLWord En (bit #)	Bits	Parameter	Description
0	33:32	Unused	Reserved
	31	ClassEn	<p>Enable the class queue. When ClassEn = 1, it permits cells to be received and transmitted if the port and VCs are setup. The ClassEn should be set to zero 0 only after the VC has been torn down.</p> <p>If the ClassEn is set to zero prior to tearing down a VC, then incoming cells will increment either the DiscardCnt or CLPxDiscardCnt. Details are listed below.</p> <p>When ClassEn = 0, all incoming cells to this class will be discarded, CellRxErrICI updated, and the maskable interrupt QCellRxErr will be sent.</p> <p>For non FCQ VCs, any cells remaining in the queue will be transmitted, if possible. If the VC is in cell discard mode, cells received will cause the DiscardCnt to increment. If the VC is in EPD/PPD discard mode, the DiscardCnt is incremented when the first cell is received, and CLPxDiscardCnt is incremented when subsequent cells are received.</p> <p>For FCQ VCs, the CLPxDiscardCnt is incremented when a cell is received. If a frame is in the process of being re-assembled when ClassEn = 1->0, the CLPxDiscardCnt will be incremented by the length of the re-assembly queue plus 1 (cell that arrived).</p>
	30:24	Unused	Reserved
	23:16	ClassMaxThresh	Maximum threshold for cells on a per-Class basis. 4 bit logarithmic, 4 bit fractional encoding. A zero value will effectively disable congestion at the class level for this class.
15:8	ClassCLP0Thresh	Per- Class EPD threshold for CLP0 cells that have met their minimum buffer allocation. 4 bit logarithmic, 4 bit fractional encoding.	

MPLWord En (bit #)	Bits	Parameter	Description
	7:0	ClassCLP1Thresh	EPD maximum threshold for CLP1 cells on a per-Class basis. 4 bit logarithmic, 4 bit fractional encoding.
1	33:26	Unused	Reserved
	25:20	<i>ClassCnt[17:12]</i>	<i>Last 6 bits of an 18 bit count of all cells in both class, TxSlot (SFQ only) & VC queue. Used for congestion control. VC tear down and watch dog re-allocation will reduce this count by VcQCLP01Cnt Class tear down will set this parameters to zero.</i>
	19:18	Unused	Reserved
	17:0	<i>ClassHeadPtr</i>	<i>Points to first Cell in linked list of Cell Records for Class queue. A value of zero indicates that the queue is empty; however, during shaping or VcMerge, it does not indicate that the class is empty. Class tear down will set this parameters to zero.</i>
2	33:32	Unused	Reserved
	31:20	<i>ClassCnt[11:0]</i>	<i>First 12 bits of an 18 bit count of all cells in both class & VC queue. See MSB for description.</i>
	19:18	Unused	Reserved
	17:0	<i>ClassTailPtr</i>	<i>Points to last Cell in linked list of Cell Records for Class queue. Class tear down will set this parameters to zero.</i>
3	33:0	Unused	Reserved

Note

Parameters in italics are active context information. Microprocessor must mask these bits out during in-service context modification. During class setup, these italicized parameters should be reset to zero.

14.2.4 Shaping Context Records

14.2.4.1 Shape TxSlot Context Record

MPMemSelect = External Queue Context

MPQuadAddr = QShpStartAdr * 1024 + Shape# * 2048 + TxSlot/2, TxSlot = 0..4095

The TxSlot record for even numbered TxSlots are located in long words 0 & 1. Odd numbered TxSlots are located in long words 2 & 3.

Table 31 - Shape TxSlot Context Record Structure

MPLWord En (bit #)	Bits	Parameter	Description
0/2	33:16	<i>ShpTxQHeadPtr</i>	<i>Head pointer for the traffic shaped transmission queue. A value of zero indicates that the queue is empty.</i>
	15:0	Unused	Reserved
1/3	33:16	<i>ShpTxQTailPtr</i>	<i>Tail pointer for the traffic shaped transmission queue.</i>
	15:0	<i>ShpTxQCnt</i>	<i>Number of cells en-queued on this time slot.</i>

Note

Parameters in italics are active context information. Microprocessor must mask these bits out during in-service context modification. During VC setup, these italicized parameters should be reset to zero.

14.2.4.2 Shape Rate Context Record

MPMemSelect = External Queue Context

MPQuadAddr = QShpStartAdr * 1024 + 8192 + ICI/2

The shape rate record for even numbered ICI's are located in long words 0 & 1. Odd numbered ICI's are located in long words 2 & 3.

Table 32 - Shape Rate Context Record Structure

MPLWord En (bit #)	Bits	Parameter	Description
0/2	33:32	Unused	Reserved
	31:28	ShpLateBits	Number of bits used to represent ShpTxSlotsLate, the "late counter". In unit of timeslots, $MBS = 2^{\text{ShpLateBits}}$.
	27:16	ShpCdvt	Cell Delay Variance Tolerance for the connection. Format is in integer number of timeslots. It must be less than ShpIncr, and $\text{ShpCdvt} < 2^{\text{ShpLateBits}}$. Relative to the ATMF 4.1 specification terminology, this value is equivalent to $T - T_s$, where $T = 1/PCR$, $T_s = 1/SCR$. This parameter must be set to zero if ShpPrescale = 0.
	15:13	Unused	Reserved
	12	ShpPrescale	Determines the resolution of the Incr field
	11:0	ShpIncr	Increment field for SCR-GCRA. . Format is in integer number of timeslots if Prescale = 1 or in integer + 1/64 fractional timeslots xxxxxx.xxxxxx (6 bit integer part w/ 6-bit fractional part) if ShpPrescale = 0. The ShpIncr must always be greater than or equal to 1.0.
1/3	33:30	Unused	Reserved
	29:18	ShpTxSlotsLate	<i>Indicates the accumulated number of late time slots that the previous cells for this connection were output. ShpLateBits specifies the number of least significant bits used to represent this value.</i>
	17:12	ShpRem	<i>Indicates the remainder when scheduling at fractional timeslots.</i>
	11:0	ShpTxSlot	<i>Indicates the last time slot that this connection was scheduled</i>

Note

Parameters in italics are active context information. Microprocessor must mask these bits out during in-service context modification. During VC setup, these italicized parameters should be reset to zero.

14.2.5 Cell Context Record

MPMemSelect = External Queue Context

MPQuadAddr = QCellStartAdr * 1024 + CellPtr/4

The cell records are packed together. CellPtr = 0 would have the cell record located in long word 0. CellPtr = 3 would have the cell record located in long word 3 etc.

Table 33 - Cell Context Record Structure

MPLWord En (bit #)	Bits	Parameter	Description
0/1/2/3	33:16	<i>CellNextPtr</i>	<i>Points to the next cell in the linked list</i>
	15:0	<i>CellICI</i>	<i>VC Identifier</i>

Note

Parameters in italics are active context information. Microprocessor must mask these bits out during in-service context modification. During freelist setup, CellICI should be reset to zero.

14.2.6 Misc Context

Note:

- 1) Masked writes are not permitted for any of the Misc. context records.
- 2) Parameters in italics are active context information. During chip setup, these italicized parameters should be reset to zero unless otherwise stated.

MPMemSelect = Internal Queue Context

MPQuadAddr = 2048

Table 34 - Free Count Context Structure

MPLWord En (bit #)	Bits	Parameter	Description
0	33:18	Unused	Reserved
	17:0	<i>FreeCnt</i>	<i>Total count of all cells available for buffering. This parameter must be initialized to the cell buffer size minus 1.</i> <i>VC tear down and watch dog re-allocation will increase this count by VcQCLP01Cnt.</i> <i>Class tear down will increase this count by ClassCnt.</i>
1	33:18	Unused	Reserved
	17:0	<i>FreeHeadPtr</i>	<i>Points to the first cell context record of the linked list that holds all the cell records that are free for cell buffering.</i>
2	33:18	Unused	Reserved
	17:0	<i>FreeTailPtr</i>	<i>Points to the last cell context record of the linked list that holds all the cell records that are free for cell buffering.</i>
3	33:0	Unused	Reserved

MPQuadAddr = 2049

Table 35 - Overall Count Context Structure

MPLWord En (bit #)	Bits	Parameter	Description
0	33:18	Unused	Reserved
	17:0	<i>LoopCnt</i>	<i>Total count of all cells queued for all loop ports. This is the "DirCnt" for cells destined to a loop port.</i> <i>VC tear down and watch dog re-allocation will reduce this count by VcQCLP01Cnt if VC is destined for the Loop.</i> <i>Class tear down will reduce this count by ClassCnt if class is destined for the Loop.</i>
1	33:18	Unused	Reserved

MPLWord En (bit #)	Bits	Parameter	Description
	17:0	WANCnt	Total count of all cells queued for all WAN ports. This is the "DirCnt" for cells destined to a WAN port. VC tear down and watch dog re-allocation will reduce this count by VcQCLP01Cnt if VC is destined for the WAN. Class tear down will reduce this count by ClassCnt if class is destined for the WAN.
2	33:0	Unused	Reserved
3	33:0	Unused	Reserved

MPQuadAddr = 2050

Table 36 - Congestion Discard Context Structure

MPLWord En (bit #)	Bits	Parameter	Description
0	33:32	Unused	Reserved
	31:0	CLP0DiscardCnt	Free running count of all inbound CLP0 cells that have been discarded due to congestion, re-assembly maximum length limit, disabled class/port, or zero length check.
1	33:16	Unused	Reserved
	15:0	CLP0DiscardICI	These bits contain the ICI of the last time a CLP0 cell was discarded due to congestion.
2	33:32	Unused	Reserved
	31:0	CLP1DiscardCnt	Free running count of all inbound CLP1 cells that have been discarded due to congestion, re-assembly maximum length limit, disabled class/port, or zero length check.
3	33:0	Unused	Reserved
	15:0	CLP1DiscardICI	These bits contain the ICI of the last time a CLP1 cell was discarded due to congestion.

MPQuadAddr = 2051

Table 37 - Maximum Congestion ID Context Structure

MPLWord En (bit #)	Bits	Parameter	Description
0	33:16	Unused	Reserved
	15:0	<i>VcMaxThrshErrICI</i>	<i>These bits contain the ICI of the last VC queue that has reached the congestion limit of VcMaxThrsh.</i>
1	33:14	Unused	Reserved
	13:12	<i>ClassMaxThrshErrID</i>	<i>These bits identify which class of the last class queue that has reached the ClassMaxThrsh limit.</i>
	11:0	<i>ClassMaxThrshErrPortID</i>	<i>These bits contain the Port ID of the last class queue that has reached the ClassMaxThrsh limit. See VC Context Record for PortID encoding.</i>
2	33:12	Unused	Reserved
	11:0	<i>PortMaxThrshErrPortID</i>	<i>These bits contain the Port ID of the last port queue that has reached the PortMaxThrsh limit. See VC Context Record for PortID encoding.</i>
3	33:0	Unused	Reserved

MPQuadAddr = 2052

Table 38 - Misc Error Context Structure

MPLWord En (bit #)	Bits	Parameter	Description
0	33:32	Unused	Reserved
	31:0	<i>DiscardCnt</i>	<i>Free running general discard count of all the cells that have been discarded due to reasons other than congestion. Possible causes include re-assembly time out, cell encountered on a disable VC/Class/Port, tear down of a VC or class queue..</i>
1	33:16	Unused	Reserved
	15:0	<i>CellRxErrICI</i>	<i>These bits contain the ICI of the last cell that was destined for a VC that was not enabled. This could be caused when any of the following enables are not set to 1: VcEn, ClassEn, PortEn.</i>
2	33:16	Unused	Reserved

MPLWord En (bit #)	Bits	Parameter	Description
	15:0	VcReasLenErr ICI	These bits contain the ICI of the last VC that violated the maximum permitted re-assembly length.
3	33:16	Unused	Reserved
	15:0	VcReasTimeErrICI	These bits contain the ICI of the last timed out VC discovered by the watch dog.

14.3 WAN Port Scheduler Context

14.3.1 WAN Transmit Port Polling Weight Record

MPLMemSelect = WAN Port Scheduler Internal Context

MPQuadAddr = 0

The WAN Transmit Port Polling Weight records are packed together

NOTE: Only single long word accesses are permitted at any one time.

Table 39 - WAN Transmit Port Polling Weight

MPLWord En (bit #)	Bits	Parameter	Description
0	31:14	Unused	Reserved
	13:12	WANPollWght 3	WAN polling weight for port 3
	11:10	Unused	Reserved
	9:8	WANPollWght 2	WAN polling weight for port 2
	7:6	Unused	Reserved
	5:4	WANPollWght 1	WAN polling weight for port 1
	3:2	Unused	Reserved
	1:0	WANPollWght 0	WAN polling weight for port 0

The polling weight determines how frequently the WAN port scheduler will evaluate whether the corresponding WAN port should be polled for transmit packet available. The weights are logarithmic. The weight determines how many of the total polling cycles a port participates in. A weight should be set according to the following table:

Table 40 - WAN Poll Weight Format

WANPollWght[1:0]	Weight ratio
00	1
01	1/2
10	1/4
11	1/8

14.3.2 WAN Transmit Class Status Record

MPMemSelect = WAN Port Scheduler Internal Context

MPQuadAddr = 0

NOTE: Only single long word accesses are permitted at any one time.

Table 41 - WAN Class Status

MPLWord En (bit #)	Bits	Parameter	Description
1	31:16	Unused	Reserved
	15:12	WANClassStat 3	WAN class status for port 3
	11:8	WANClassStat 2	WAN class status for port 2
	7:4	WANClassStat 1	WAN class status for port 1
	3:0	WANClassStat 0	WAN class status for port 0

The bits in each WAN class status indicates whether a cell destined for the given WAN is currently in the queue waiting for transmission. Bit 0 represents class 0, bit 1 represents class 1 etc.

Warning: Masking should always be applied over active classes during writes.

14.4 Loop Port Scheduler Context

14.4.1 Loop Transmit Port Polling Sequence Record

MPMemSelect = Loop Port Scheduler Internal Context

MPQuadAddr = Loop#/16

The Loop Transmit Port Polling Sequence records are packed. Each long word contains sequence numbers for four ports. Each quad word contains entries for 16 ports. There are a total of 128 quad words in the table for a total of 2K loop ports. Long word 0 contains weight for ports 0-3, long word 1 contains weight for ports 4-7 etc.

NOTE: Only single long word accesses are permitted at any one time.

Table 42 - Loop Transmit Port Polling Sequence

MPLWord En (bit #)	Bits	Parameter	Description
0/1/2/3	31	Unused	Reserved
	30:24	LoopPollSeq3	Loop polling sequence for port ((MPLWordEn bit *4) + 3)
	23	Unused	Reserved
	22:16	LoopPollSeq2	Loop polling sequence for port ((MPLWordEn bit *4) + 2)
	15	Unused	Reserved
	15:8	LoopPollSeq1	Loop polling sequence for port ((MPLWordEn bit *4) + 1)
	7	Unused	Reserved
	6:0	LoopPollSeq0	Loop polling sequence for port ((MPLWordEn bit *4) + 0)

The polling sequence determines when within a loop port scheduler weight polling cycle this port will be evaluated to be polled. This allows software to evenly distribute the polling of ports of the same weight. The loop port scheduler will compare the n LSB's of the LoopPollSeq with the current scheduler poll sequence to decide if that port should be polled (n is equal to the port's

LoopPollWght). When these 2 values match and that port has transmit data to be sent, that port is scheduled to be polled.

14.4.2 Loop Transmit Port Polling Weight Record

MPMemSelect = Loop Port Scheduler Internal Context

MPQuadAddr = 128 + Loop#/32

The Loop Transmit Port Polling Weight records are packed together. Each long word contains weights for 8 loop ports. There are a total of 64 quad words in the table for a total of 2K loop ports. Long word 0 contains weight for ports 0-7, long word 1 contains weight for ports 8-15 etc.

NOTE: Only single long word accesses are permitted at any one time.

Table 43 - Loop Transmit Port Polling Weight

MPLWord En (bit #)	Bits	Parameter	Description
0/1/2/3	31	Unused	Reserved
	30:28	LoopPollWght 7	Loop polling weight for port ((MPLWordEn bit *8) + 7)
	27	Unused	Reserved
	26:24	LoopPollWght 6	Loop polling weight for port ((MPLWordEn bit *8) + 6)
	23	Unused	Reserved
	22:20	LoopPollWght 5	Loop polling weight for port ((MPLWordEn bit *8) + 5)
	19	Unused	Reserved
	18:16	LoopPollWght 4	Loop polling weight for port ((MPLWordEn bit *8) + 4)
	15	Unused	Reserved
	14:12	LoopPollWght 3	Loop polling weight for port ((MPLWordEn bit *8) + 3)
	11	Unused	Reserved
	10:8	LoopPollWght 2	Loop polling weight for port ((MPLWordEn bit *8) + 2)

MPLWord En (bit #)	Bits	Parameter	Description
	7	Unused	Reserved
	6:4	LoopPollWght 1	Loop polling weight for port ((MPLWordEn bit *8) + 1)
	3	Unused	Reserved
	2:0	LoopPollWght 0	Loop polling weight for port ((MPLWordEn bit *8) + 0)

The polling weight determines how frequently the loop port scheduler will evaluate whether the corresponding loop port should be polled for transmit packet available.

14.4.3 Loop Transmit Class Status Record

MPMemSelect = Loop Port Scheduler Internal Context

MPQuadAddr = 192 + Loop#/32

The Loop Class Status records are packed together. Each long word contains status for 8 loop ports. There are a total of 64 quad words in the table for a total of 2K loop ports. Long word 0 contains weight for ports 0-7, long word 1 contains weight for ports 8-15 etc.

NOTE: Only single long word accesses are permitted at any one time.

Table 44 - Loop Class Status

MPLWord En (bit #)	Bits	Parameter	Description
0/1/2/3	31:28	LoopClassStat 7	Loop class status for port ((MPLWordEn bit *8) + 7)
	27:24	LoopClassStat 6	Loop class status for port ((MPLWordEn bit *8) + 6)
	23:20	LoopClassStat 5	Loop class status for port ((MPLWordEn bit *8) + 5)
	19:16	LoopClassStat 4	Loop class status for port ((MPLWordEn bit *8) + 4)

MPLWord En (bit #)	Bits	Parameter	Description
	15:12	LoopClassStat 3	Loop class status for port ((MPLWordEn bit *8) + 3)
	11:8	LoopClassStat 2	Loop class status for port ((MPLWordEn bit *8) + 2)
	7:4	LoopClassStat 1	Loop class status for port ((MPLWordEn bit *8) + 1)
	3:0	LoopClassStat 0	Loop class status for port ((MPLWordEn bit *8) + 0)

The bits in each loop class status indicates whether a cell destined for the given loop is currently in the queue waiting for transmission. Bit 0 represents class 0, bit 1 represents class 1 etc.

Warning: Masking should always be applied over active classes during writes.

15 TEST FEATURES DESCRIPTION

15.1 JTAG Test Port

The S/UNI APEX JTAG Test Access Port (TAP) allows access to the TAP controller and the 4 TAP registers: instruction, bypass, device identification and boundary scan. Using the TAP, device input logic levels can be read, device outputs can be forced, the device can be identified and the device scan path can be bypassed.

Table 45 - Instruction Register

Length - 3 bits

Instructions	Selected Register	Instruction Codes, IR[2:0]
EXTEST	Boundary Scan	000
IDCODE	Identification	001
SAMPLE	Boundary Scan	010
BYPASS	Bypass	011
BYPASS	Bypass	100
STCTEST	Boundary Scan	101
BYPASS	Bypass	110
BYPASS	Bypass	111

Table 46 - Identification Register

Length	32 bits
Version number	0H
Part Number	7326H
Manufacturer's identification code	0CDH
Device identification	073260CDH

Table 47 - Boundary Scan Register

Please see file "APEX JTAG Scan Register.xls"

Boundary Scan Cells

In the following diagrams, CLOCK-DR is equal to TCK when the current controller state is SHIFT-DR or CAPTURE-DR, and unchanging otherwise. The multiplexer in the centre of the diagram selects one of four inputs, depending on the status of select lines G1 and G2. The ID Code bit is as listed in the Boundary Scan Register table located above.

Figure 29 - Input Observation Cell (IN_CELL)




Figure 30 - Output Cell (OUT_CELL)

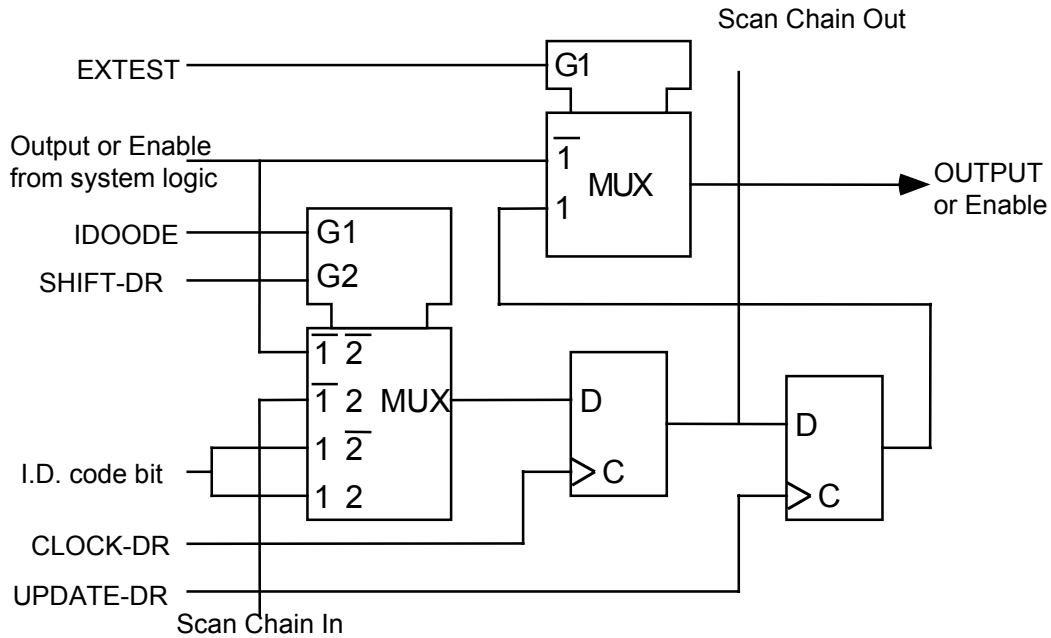


Figure 31 - Bi-directional Cell (IO_CELL)

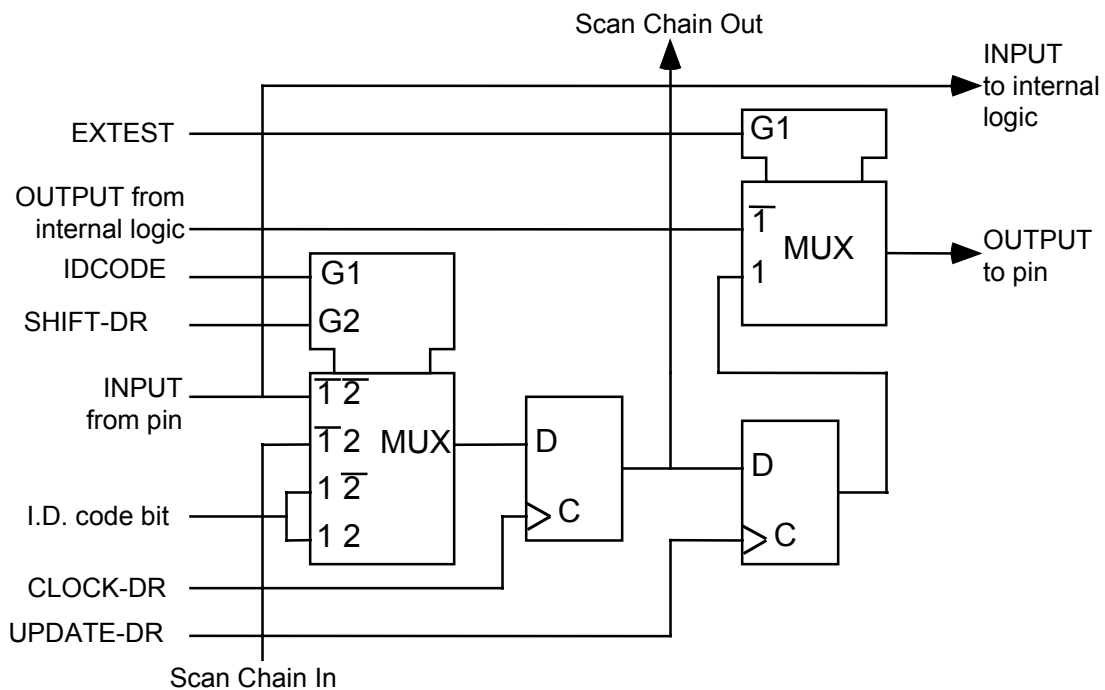



Figure 32 - Layout of Output Enable and Bidirectional Cells



16 OPERATION

Please refer to the document "S/UNI APEX H/W Programmer's Guide", PMC-991454

17 FUNCTIONAL TIMING

17.1 Microprocessor Interface

The following diagrams illustrate the various handshaking required for microprocessor reads and writes.

Figure 33 shows a single read and write operation with bus polarity set to 1. On the first cycle, BURSTB is sampled inactive; therefore, it is expected that the cycle be a single data transfer, and the BLAST signal is of no significance. The subsequent 2 cycles have BURSTB sampled active hence the transfer cycle is terminated when both BLAST and READYB are asserted. Note that between each transfer, there is a turn around cycle provided by the external interface to ensure that there is no bus contention on back to back transfers on the AD bus.

Figure 33 - Single Word Read and Write

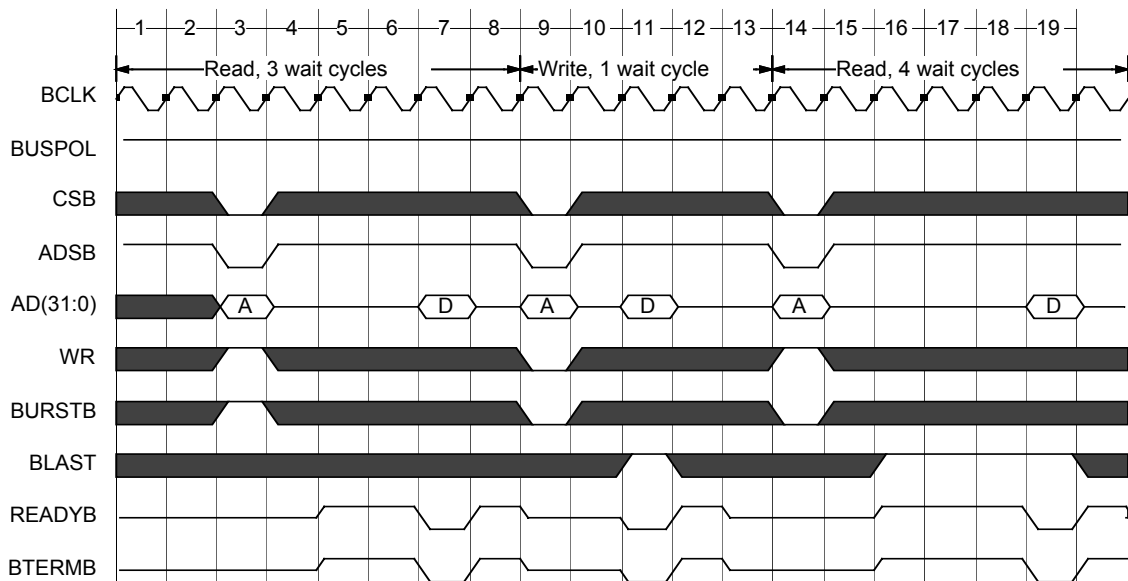


Figure 34 shows a burst read and write operation with bus polarity set to 0. The first and third access illustrate transfers that are terminated by the S/UNI APEX via the assertion of BTERMB. The second and fourth access illustrate transfers that are terminated by the external interface via the assertion of BLAST. Note that between each transfer, there is no turn around cycle. Care must be taken to examine the AC timing to ensure that there is no bus contention on the AD bus between a read followed by a write transfer.

Figure 34 - Burst Read and Write





Figure 35 shows consecutive write operations using the WRDONEB signal without the READYB. Write operations may only begin when WRDONEB is sampled low by the external interface. On the first data transfer, the cycle is terminated normally. Subsequent access does not begin until WRDONEB is sampled low by the external interface. This interface is used when the external processor is incapable of dealing with wait states during write operations.

Figure 35 - Consecutive Write Accesses Using WRDONEB



17.2 SDRAM Interface

The following three diagrams depict the timing for signals destined for the pins of the SDRAM during the Activate-Read (with Auto-precharge), Activate-Write (with Auto-precharge), and Auto-refresh command sequences. The *cbcmd* signal is not an actual signal; it merely represents the memory access command formed by the combination of the individual SDRAM control signals (*cbcsb_o*, *cbrasb_o*, etc.). Another note is that reads/writes are always done in bursts of eight words; the first involves the even banks and the second burst involves the odd banks in SDRAM.

Figure 36 - Read Timing

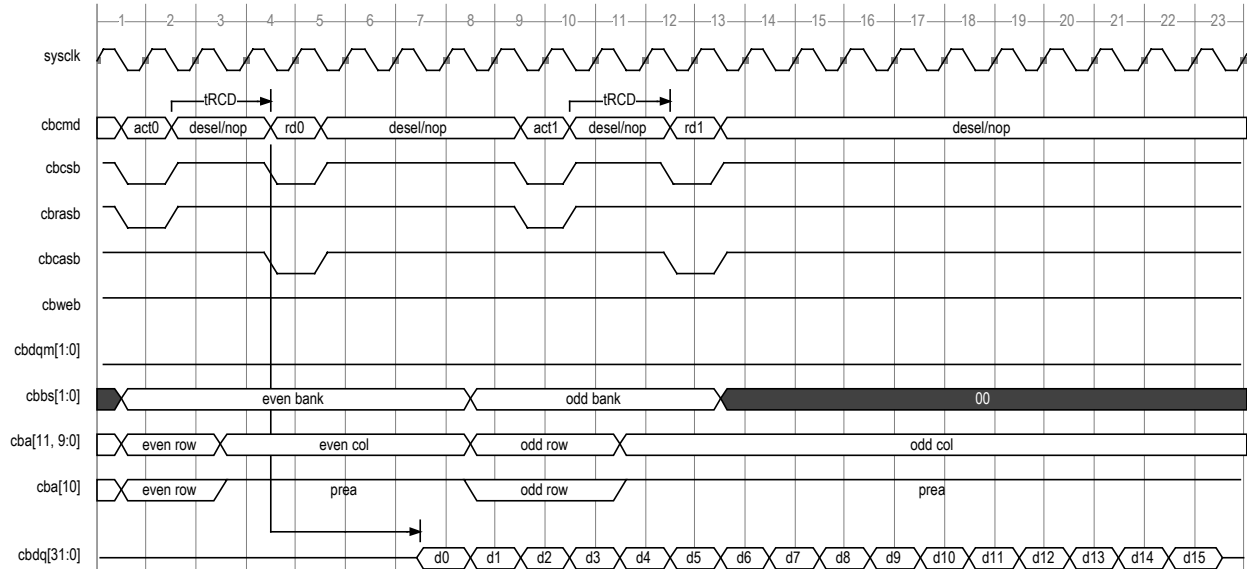


Figure 37 - Write Timing




Figure 38 - Refresh





Figure 39 - Power Up and Initialization Sequence

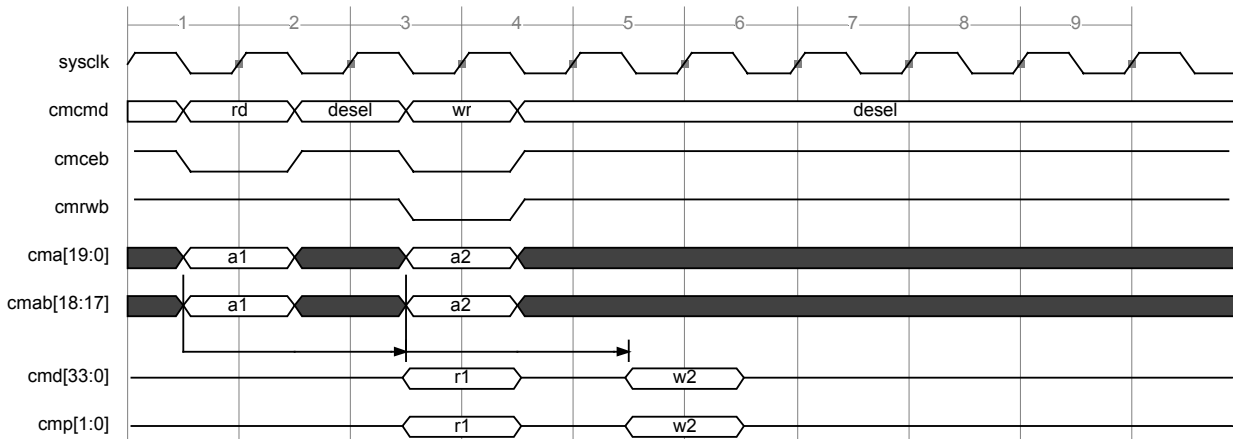


17.3 ZBT SSRAM Interface

The following diagram depicts the timing for signals destined for the pins of the pipelined ZBT SSRAM during a read followed by a write cycle. The cmcmd

signal is not an actual signal; it merely represents the memory access command formed by the combination of the individual SSRAM control.

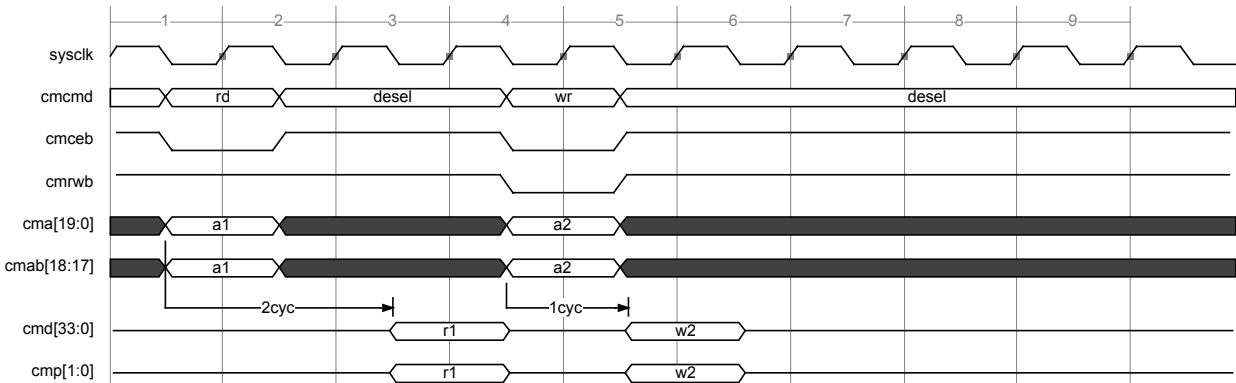
Figure 40 - Read followed by Write Timing



17.4 Late Write SSRAM Interface

The following diagram depicts the timing for signals destined for the pins of the register to register Late Write SSRAM during a read followed by a write cycle. The cmcmd signal is not an actual signal; it merely represents the memory access command formed by the combination of the individual SSRAM control.

Figure 41 - Read followed by Write Timing



17.5 Any-PHY/UTOPIA Interfaces

While the following diagrams present representative waveforms, they are not an attempt to unambiguously describe the interfaces. The Pin Description section is intended to present the detailed pin behavior and constraints on use.

The following parameters apply to all Any-PHY/UTOPIA interface figures:

n = 2 for WAN, 5 for Loop

m = 7 for 8 bit mode, 15 for 16 bit mode

k = function of 8/16 bit mode, and number of prebends selected.

17.5.1 Receive Master/Transmit Slave Interfaces

Figure 42 gives an example of the functional timing of the receive interface when configured as a UTOPIA Level 2 compliant transmit slave. The interface responds to the polling of address “APEX” (which matches the address defined by the register {Loop/WAN}RxSlaveAddr[n:0]) by asserting RPA when it is capable of accepting a complete cell. As a result, the master selects the S/UNI-APEX by presenting “APEX” again during the last cycle RENB is high. Had not the device been selected, RSOP, RDAT[n:0] and RPRTY would have remained high-impedance.

Figure 42 illustrates that a cell transfer may be paused by deasserting RENB. The device is reselected by presenting address "APEX" the last cycle RENB is high to resume the transfer.

Figure 42 - UTOPIA L2 Transmit Slave (Loop & WAN)




Figure 43 gives an example of the functional timing of the receive interface when configured as a UTOPIA Level 1 compliant receive master. When S/UNI APEX is capable of accepting at least one more cell, and it samples RPA high, a transfer cycle is initiated. If the S/UNI APEX is capable of receiving an additional cell and RPA is sampled high, it will de-assert RENB first before initiating the next transfer.

Once transfer is initiated, RENB will remain asserted until the last data is received.

Figure 43 - UTOPIA L1 Receive Master (Loop & WAN)




Figure 44 gives an example of the functional timing of the receive interface when configured as a UTOPIA Level 2 compliant receive master. When S/UNI APEX is capable of accepting at least one more cell, the interface polls addresses until it receives an asserted RPA. As a result, the master re-selects the same RADR again during the last cycle RENB is high to initiate a transfer. If the S/UNI APEX is capable of receiving an additional cell, it will continue to poll for the next available port.

Once transfer is initiated, RENB will remain asserted until the last data is received.

Figure 44 - UTOPIA L2 Receive Master (Loop & WAN)

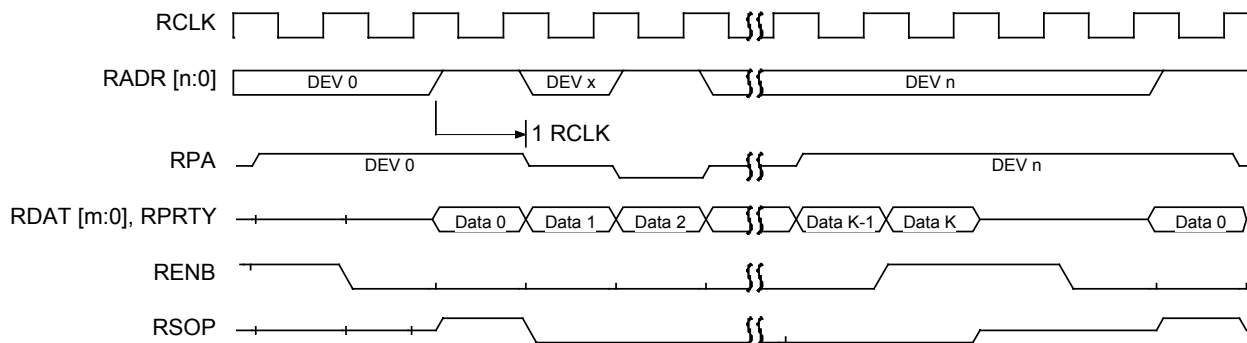



Figure 45 gives an example of the functional timing of the receive interface when configured as a Any-PHY compliant receive master. When S/UNI APEX is capable of accepting at least one more cell, the interface polls addresses until it receives an asserted RPA. As a result, the master re-selects the same RADR again during the last cycle RENB is high to initiate a transfer. If the S/UNI APEX is capable of receiving an additional cell, it will continue to poll for the next available port.

Once transfer is initiated, RENB will remain asserted until the last data is received.

Figure 45 - Any-PHY Receive Master (Loop & WAN)



17.5.2 Transmit Master/Receive Slave Interfaces

Figure 46 gives an example of the functional timing of the transmit interface when configured as a UTOPIA Level 2 compliant receive slave. The interface responds to the polling of address “APEX” (which matches the address defined by the register {Loop/WAN}TxSlaveAddr[n:0]) by asserting TPA when it is capable of transmitting a complete cell. As a result, the master selects the S/UNI-APEX by presenting “APEX” again during the last cycle TENB is high. Had not the device been selected, TSOP, TDAT[n:0] and TPRTY would have remained high-impedance.

Figure 46 illustrates that a cell transfer may be paused by deasserting TENB. The device is reselected by presenting address “APEX” the last cycle TENB is high to resume the transfer.

Figure 46 - UTOPIA L2 Receive Slave (Loop & WAN)




Figure 47 gives an example of the functional timing of the WAN transmit interface when configured as a UTOPIA L1 compliant transmit master. The address presented on the WTADR bus comes from the WANTxSlaveAddr register. When the S/UNI APEX samples WTPA as high and there is a complete cell is available for transfer, a transfer is initiated. If WTPA remains high at the end of the first transfer, and if a cell is available for transfer, another transfer is initiated. Transfers complete without pausing.

Figure 47 - WAN UTOPIA L1 Transmit Master




Figure 48 gives an example of the functional timing of the loop transmit interface when configured as a UTOPIA L1 compliant transmit master. It is nearly identical to the WAN UTOPIA L1 transmit master, with the exception that the address on the LTADR bus is indeterminate and should not be used.

Figure 48 - Loop UTOPIA L1 Transmit Master



Figure 49 gives an example of the functional timing of the WAN transmit interface when configured as a UTOPIA L2 compliant transmit master. The S/UNI APEX polls ports that it has a complete cell available for transfer. The receiving device responds by driving WTPA. When the first port asserts WTPA, the S/UNI APEX will stop polling, and drives WTADR with the port's address until the transfer is initiated. Once the transfer is initiated, as indicated by the assertion of WTENB, polling recommences and continues until the next port asserts WTPA. If the port in transfer is polled at the same time, WTPA is considered valid only for the last 4 clocks from the end of the transfer. The second transfer will not begin until the first transfer is complete. Transfers complete without pausing.

Figure 49 - WAN UTOPIA L2 Transmit Master

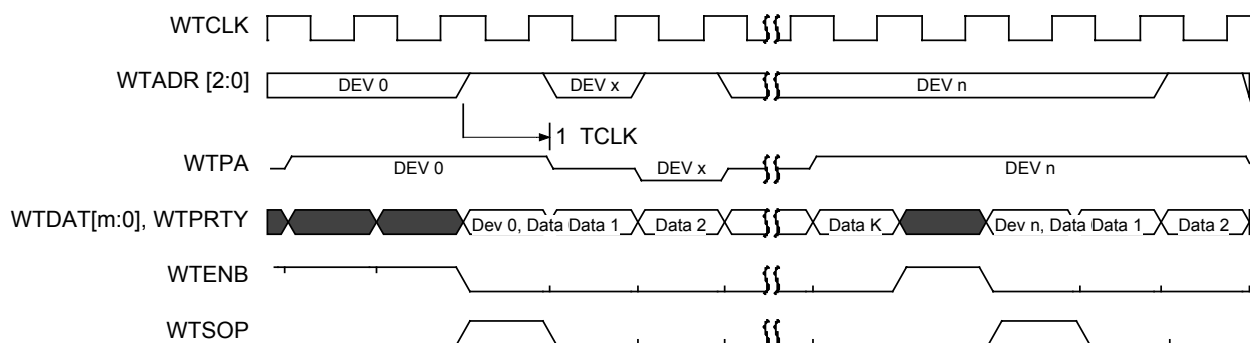


Figure 50 gives an example of the functional timing of the loop transmit interface when configured as a UTOPIA L2 compliant transmit master. It is nearly identical to the WAN UTOPIA L2 transmit master, with the exception of the polling behaviour. Unlike the WAN, the loop interface will assert a NULL address once a

LTPA is sampled high. The NULL address remains on the LTADR bus until the port selection marking the start of the next transfer. Transfers complete without pausing.

Figure 50 - Loop UTOPIA L2 Transmit Master




Figure 51 gives an example of the functional timing of the WAN transmit interface when configured as a Any-PHY compliant transmit master. The S/UNI APEX polls ports that it has a complete cell available for transfer. The receiving device responds by driving WTPA. When the first port asserts WTPA, the S/UNI APEX will stop polling, and drives WTADR with the port's address until the transfer is initiated. Once the transfer is initiated, as indicated by the assertion of WTENB & WTSX, polling recommences and continues until the next port asserts WTPA. Transfers complete without pausing.

Figure 51 - WAN Any-PHY Transmit Master


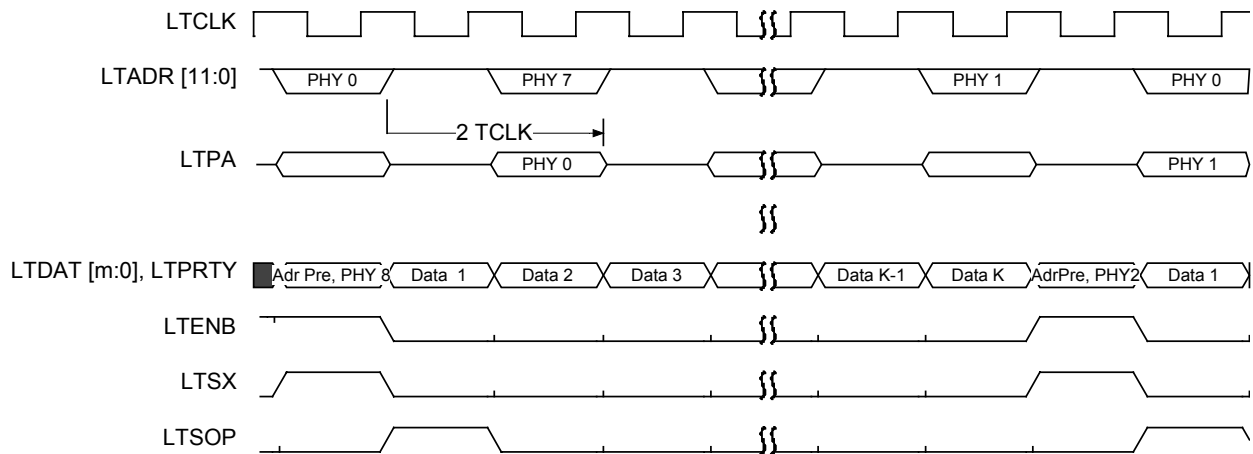


Figure 52 gives an example of the functional timing of the loop transmit interface when configured as a Any-PHY compliant transmit master. The S/UNI APEX

polls ports that it has a complete cell available for transfer. The receiving device responds by driving LTPA. Positive responses are recorded, and will eventually result in a data transfer. Polling continues independent of the data transfer state.

Data transfers are initiated with the assertion of LTENB & LTSX, and complete without pausing.

Figure 52 - Loop Any-PHY Transmit Master



18 ABSOLUTE MAXIMUM RATINGS

Maximum rating are the worst case limits that the device can withstand without sustaining permanent damage. They are not indicative of normal mode operation conditions.

Table 48 - Absolute Maximum Ratings

Parameter	Symbol	Value
Storage Temperature	T _{ST}	-40°C to +125°C
Supply Voltage	V _{DD}	-0.3V to +4.6V
Voltage on Any Pin	V _{IN}	0V to V _{DDO} +0.5V
Static Discharge Voltage		±1000 V
Latch-Up Current		±100 mA
DC Input Current	I _{IN}	±10 mA
Lead Temperature		+230°C
Junction Temperature	T _J	+150°C

19 D.C. CHARACTERISTICS

$T_A = -40^{\circ}\text{C}$ to $+85^{\circ}\text{C}$, $V_{DD} = 3.3\text{ V} \pm 0.3\text{V}$, $V_{PCH} = 2.5\text{ V} \pm 8\%$

(Typical Conditions: $T_A = 25^{\circ}\text{C}$, $V_{DD} = 3.3\text{ V}$, $V_{PCH} = 2.5\text{ V}$)

Table 49 - D.C. Characteristics

Symbol	Parameter	Min	Typ	Max	Units	Conditions
VDD	Pin Power Supply	3.0	3.3	3.6	Volts	
PCH	Core Power Supply	2.3	2.5	2.7	Volts	
VIL	Input Low Voltage	-0.3		0.8	Volts	
VIH	Input High Voltage	2.0		5.5V	Volts	
VOL	Output or Bidirectional Low Voltage			0.4	Volts	VDD = min, IOL = 2 mA, VI = VIL
VOH	Output or Bidirectional High Voltage	2.4			Volts	VDD = min, IOH = -2 mA, VI = VIH
VT+	Schmitt Input High Threshold Voltage	1.39	1.82	2.06	Volts	For pins RSTB & TRSTB
VT-	Schmitt Input Low Threshold Voltage	0.8	1.24	1.46	Volts	For pins RSTB & TRSTB
VTH	Schmitt Input Hysteresis Voltage		0.51		Volts	For pins RSTB & TRSTB
IILPU	Input Low Leak Current	+10		+100	μA	VIL = GND. Notes 1, 3
IIHPU	Input High Leak Current	-10		+10	μA	VIH = VDDO. Notes 1, 3
IIL	Input Low Leak Current	-10		+10	μA	VIL = GND. Notes 2, 3

Symbol	Parameter	Min	Typ	Max	Units	Conditions
I _{IH}	Input High Current	-10		+10	μA	V _{IH} = V _{DDO} . Notes 2, 3
C _{IN}	Input Capacitance		6		pF	Excluding Package, Package Typically 1 pF
C _{OUT}	Output Capacitance		6		pF	Excluding Package, Package Typically 1 pF
C _{IO}	Bidirectional Capacitance		6		pF	Excluding Package, Package Typically 1 pF
I _{DDOP}	Operating Current			513 (core) 99 (I/O unloaded) 410 (I/O loaded)	mA	T _A = 85 degC V _{DD} = V _{DD} (max),, P _{CH} = P _{CH} (max) SYSCLK = 80MHz, BCLK = 66MHz WRCLK, WTCLK, LRCLK, LTCLK = 52MHz, all Any-PHY I/F in master modes 1.66Mcells/s Aggregate throughput I/O loaded: 80pF on Loop Any-PHY I/F 40pF on WAN Any-PHY I/F, SSRAM I/F 20pF on SDRAM I/F 50pF on uP I/F

Notes on D.C. Characteristics:

1. Input pin or bi-directional pin with internal pull-up resistor.
2. Input pin or bi-directional pin without internal pull-up resistor
3. Negative currents flow into the device (sinking), positive currents flow out of the device (sourcing).

20 A.C. TIMING CHARACTERISTICS

($T_A = -40^{\circ}\text{C}$ to $+85^{\circ}\text{C}$, $V_{DD} = 3.3\text{ V} \pm 0.3\text{V}$, $V_{PCH} = 2.5\text{ V} \pm 8\%$)

Notes on Input Timing:

1. When a set-up time is specified between an input and a clock, the set-up time is measured from the 50% point of the input to the 50% point of the clock.
2. When a hold time is specified between a clock and an input, the hold time is measured from the 50% point of the clock to the 50% point of the input.

Notes on Output Timing:

1. Output time is measured between the 50% point of the clock to the 50% point of the output.

Figure 53 - RSTB Timing




Table 50 - RTSB Timing

Symbol	Description	Min	Max	Units
tV_{RSTB}	RSTB Pulse Width	100		ns

Figure 54 - Synchronous I/O Timing




Table 51 - SYSCLK Timing

Symbol	Description	Min	Max	Units
fCLK	Frequency, SYSCLK	40	80	MHz
DCLK	Duty Cycle, SYSCLK	40	60	%

Table 52 - Cell Buffer SDRAM Interface

Symbol	Description	Min	Max	Units
T_s	Input Set-up time to SYSCLK	2.4		ns
T_h	Input Hold time to SYSCLK	0.7		ns
T_p	SYSCLK High to Output Valid	1.25	7.0	ns
T_z	SYSCLK High to Output High-Impedance	1.25	8.0	ns
T_{zb}	SYSCLK High to Output Driven	1.25		ns

Maximum output propagation delays are measured with a 20pF load on the outputs.

Minimum output propagation delays are measured with a 10 pF load on the outputs.

Table 53 - Context Memory ZBT & Late Write SSRAM Interface

Symbol	Description	Min	Max	Units
T_s	Input Set-up time to SYSCLK	2.4		ns

Symbol	Description	Min	Max	Units
Th	Input Hold time to SYSCLK	0.7		ns
Tp	SYSCLK High to Output Valid	1.5	8.25	ns
Tz	SYSCLK High to Output High-Impedance	1.25	8.25	ns
Tzb	SYSCLK High to Output Driven	1.5		ns

Maximum output propagation delays are measured with a 40pF load on the outputs.

Minimum output propagation delays are measured with a 0 pF load on the outputs.

Table 54 - Microprocessor Interface

Symbol	Description	Min	Max	Units
fCLK	Frequency, BCLK	0	66	MHz
DCLK	Duty Cycle, BCLK	40	60	%
Ts	Input Set-up time to BCLK	3.0		ns
Th	Input Hold time to BCLK	0.1		ns
Tp	BCLK High to Output Valid	2.0	8.0	ns
Tz	BCLK High to Output High-Impedance	2.0	11.0	ns
Tzb	BCLK High to Output Driven	2.0		ns

Maximum output propagation delays are measured with a 50pF load on the outputs.

Minimum output propagation delays are measured with a 0 pF load on the outputs.

Table 55 - Loop Any-PHY Transmit Interface

Symbol	Description	Min	Max	Units
fCLK	LTCLK Frequency	5	52	
DCLK	LTCLK Duty Cycle	40	60	%
Ts	Input Set-up time to LTCLK	4.0		ns
Th	Input Hold time to LTCLK	0.0		ns

Symbol	Description	Min	Max	Units
Tp	LTCLK High to Output Valid	1.75	10.0	ns
Tz	LTCLK High to Output High-Impedance	1.75	12.0	ns
Tzb	LTCLK High to Output Driven	1.75		ns

Maximum output propagation delays are measured with an 80pF load on the loop outputs.

Minimum output propagation delays are measured with a 0 pF load on the outputs.

Table 56 - WAN Any-PHY Transmit Interface

Symbol	Description	Min	Max	Units
fCLK	WTCLK Frequency	0	52	
DCLK	WTCLK Duty Cycle	40	60	%
Ts	Input Set-up time to WTCLK	3.5		ns
Th	Input Hold time to WTCLK	0.0		ns
Tp	WTCLK High to Output Valid	1.75	8.0	ns
Tz	WTCLK High to Output High-Impedance	1.75	12.0	ns
Tzb	WTCLK High to Output Driven	1.75		ns

Maximum output propagation delays are measured with a 40pF load on the WAN outputs.

Minimum output propagation delays are measured with a 0 pF load on the outputs.

Table 57 - Loop Any-PHY Receive Interface

Symbol	Description	Min	Max	Units
fCLK	LRCLK Frequency	0	52	MHz
DCLK	LRCLK Duty Cycle	40	60	%
Ts	Input Set-up time to LRCLK	3.0		ns
Th	Input Hold time to LRCLK	0.25		ns
Tp	LRCLK High to Output Valid	1.75	9.0	ns

Symbol	Description	Min	Max	Units
Tz	LRCLK High to Output High-Impedance	1.75	12.0	ns
Tzb	LRCLK High to Output Driven	1.75		ns

Maximum output propagation delays are measured with an 50pF load on the loop outputs.

Minimum output propagation delays are measured with a 0 pF load on the outputs.

Table 58 - WAN Any-PHY Receive Interface

Symbol	Description	Min	Max	Units
fCLK	WRCLK Frequency	0	52	MHz
DCLK	WRCLK Duty Cycle	40	60	%
Ts	Input Set-up time to WRCLK	2.0		ns
Th	Input Hold time to WRCLK	0.35		ns
Tp	WRCLK High to Output Valid	1.75	9.0	ns
Tz	WRCLK High to Output High-Impedance	1.75	12.0	ns
Tzb	WRCLK High to Output Driven	1.75		ns

Maximum output propagation delays are measured with a 40pF load on the WAN outputs.

Minimum output propagation delays are measured with a 0 pF load on the outputs.

20.1 JTAG INTERFACE

Figure 55 - JTAG Port Interface Timing

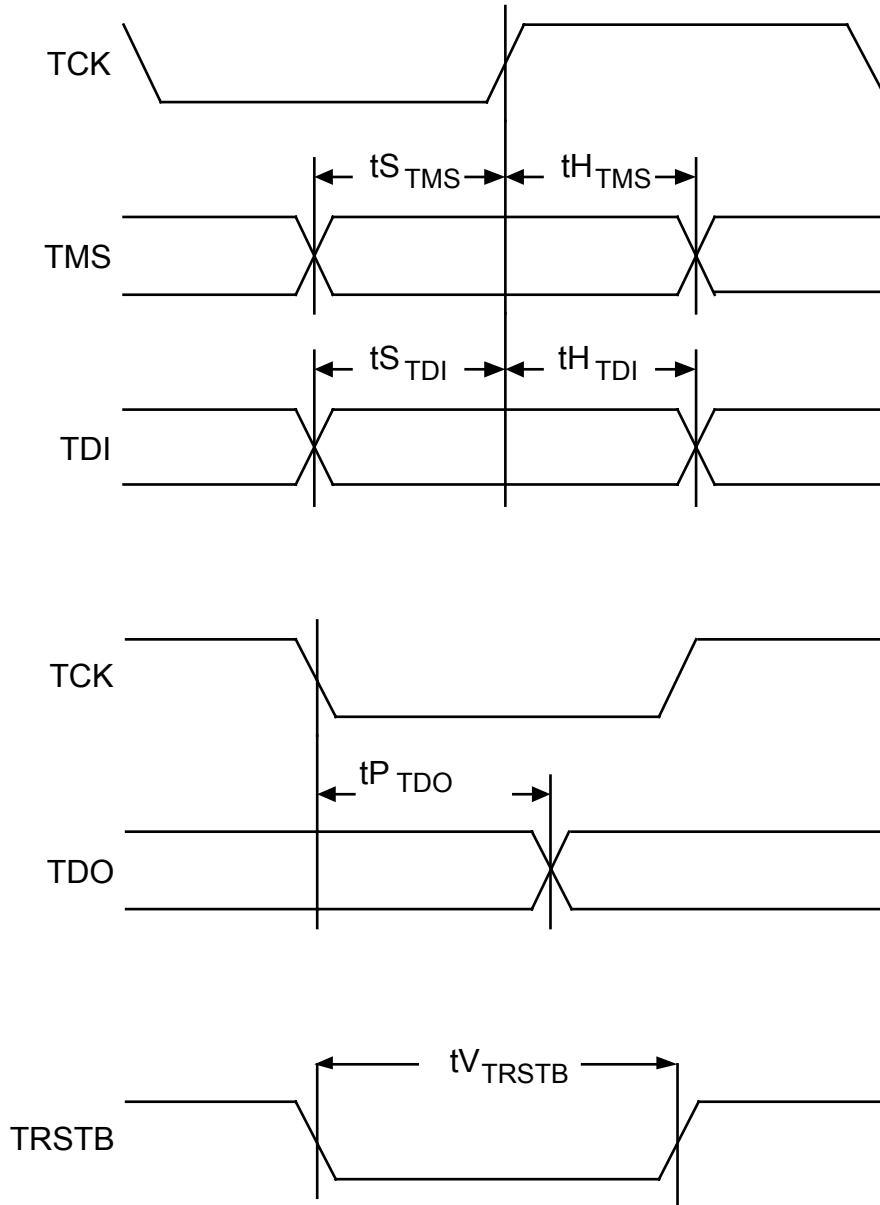


Table 59 - JTAG Port Interface

Symbol	Description	Min	Max	Units
	TCK Frequency		5	MHz

Symbol	Description	Min	Max	Units
	TCK Duty Cycle	40	60	%
tSTMS	TMS Setup time to TCK	50		ns
tHTMS	TMS Hold time to TCK	50		ns
tSTDI	TDI Setup time to TCK	50		ns
tHTDI	TDI Hold time to TCK	50		ns
tPTDO	TCK Low to TDO Valid	2	50	ns
tVTRSTB	TRSTB Pulse Width	100		ns

21 ORDERING AND THERMAL INFORMATION

Part No.	Description
PM7326-BI	352 Ball Grid Array Package (SBGA)

Theta JC < 1 degC/W


Theta JA degC/W	Forced Air (Linear Feet per Minute)					
	Conv	100	200	300	400	500
@ 2.86W						
Dense Board	18.5	16.4	15.1	14.2	13.8	13.7
JEDEC Board	13.0	12.0	11.3	10.8	10.3	9.7

Notes:

- DENSE Board is defined as a 3S3P board and consists of a 3x3 array of device PM7326 located as close to each other as board design rules allow. All PM7326 devices are assumed to be dissipating 2.86 Watts. Θ_{JA} listed is for the device in the middle of the array.
- JEDEC Board Θ_{JA} is the measured value for a single thermal device in the same package on a 2S2P board following EIA/JESD 51-3

22 MECHANICAL INFORMATION

Figure 56 - Mechanical Drawing 352 Pin Ball Grid Array (SBGA)



- NOTES: 1) ALL DIMENSIONS IN MILLIMETER.
 2) DIMENSION AAA DENOTES COPLANARITY
 3) DIMENSION BBB DENOTES PARALLEL
 4) DIMENSION CCC DENOTES FLATNESS

PACKAGE TYPE: 352 PIN THERMAL BALL GRID ARRAY																
BODY SIZE: 35 x 35 x 1.45 MM																
DIM.	A	A1	A2	D	DI	E	EI	M,N	e	b	AAA	BBB	CCC	DDD	P	S
MIN.	1.41	0.56	0.85	34.90	31.65	34.90	31.65			0.60				0.15	0.20	
NOM.	1.54	0.63	0.91	35.00	31.75	35.00	31.75	26x26	1.27	0.75				0.33	0.30	
MAX.	1.67	0.70	0.97	35.10	31.85	35.10	31.85			0.90	0.15	0.15	0.20	0.50	0.35	0.635

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

(604) 415-4533

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

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.

© 2000 PMC-Sierra, Inc.

PMC-981224 (P6) ref PMC-970790 (R5)

Issue date: April 2000