

Features

- Integrated Single-Chip 10/100 Mbps Ethernet Switch
- 24 10/100 Mbps Autosensing, Fast Ethernet Ports with RMIi or Serial Interface (7WS). Each port can independently use one of the two interfaces
- Serial interface
- Supports one Frame Buffer Memory domain with SRAM at 100 MHz
- Supports SRAM domain memory size 1 MB or 2 MB
- Applies centralized shared memory architecture
- Up to 64 K MAC addresses
- Maximum throughput is 2.4 Gbps non-blocking
- High performance packet forwarding (7.143 M packets per second) at full wire speed
- Full Duplex Ethernet IEEE 802.3x Flow Control
- Backpressure flow control for Half Duplex ports
- Supports Ethernet multicasting and broadcasting and flooding control
- Supports per-system option to enable flow control for best effort frames even on QoS-enabled ports
- Load sharing among trunked ports can be based on source MAC and/or destination MAC

Ordering Information

MVTX2601AG 553 Pin HSBGA

-40°C to 85°C

- Port Mirroring to a dedicated mirroring port or port 23 in unmanaged mode
- Full set of LED signals provided by a serial interface
- 2 port trunking groups with up to 4 10/100 ports per group
- Built-In Self Test for internal and external SRAM
- Traffic Classification
 - 4 transmission priorities for Fast Ethernet ports with 2 dropping levels
 - Classification based on:
 - Port based priority
 - VLAN Priority field in VLAN tagged frame
 - DS/TOS field in IP packet
 - UDP/TCP logical ports: 8 hard-wired and 8 programmable ports, including one programmable range
 - The precedence of the above classifications is programmable

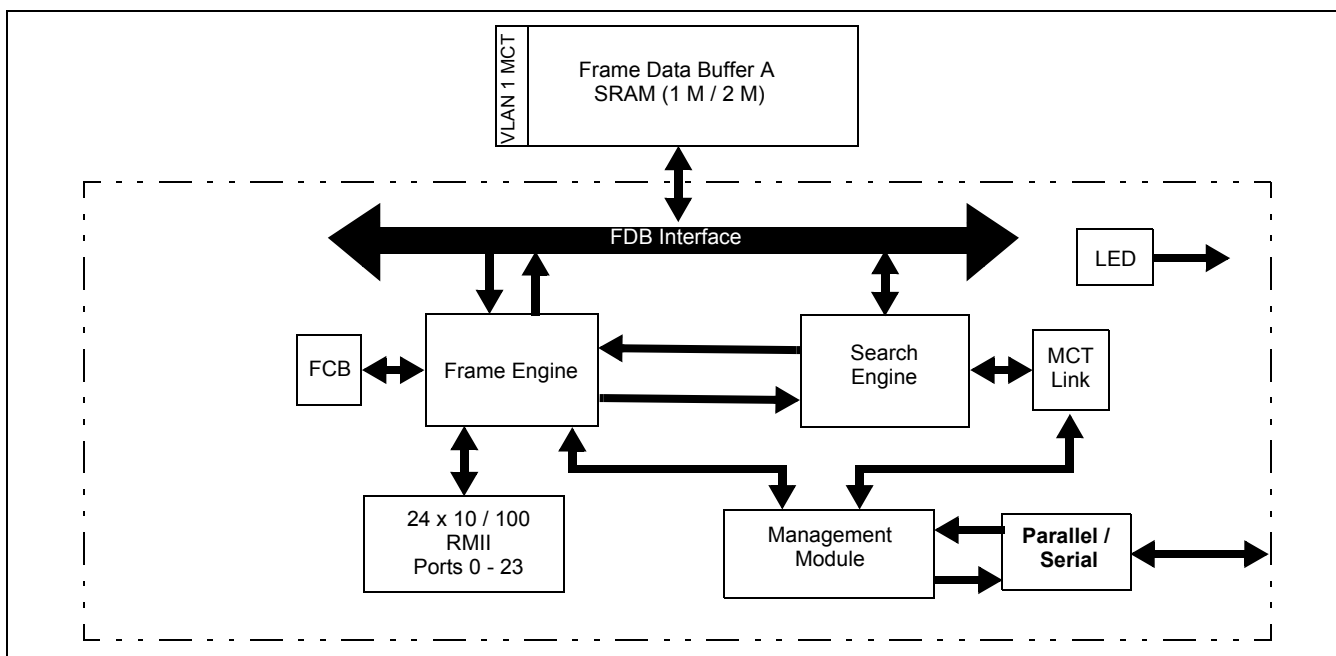


Figure 1 - MVTX2601 System Block Diagram

- QoS Support
- Supports IEEE 802.1p/Q Quality of Service with 4 transmission priority queues with delay bounded, strict priority, and WFQ service disciplines
 - Provides 2 levels of dropping precedence with WRED mechanism
 - User controls the WRED thresholds
 - Buffer management: per class and per port buffer reservations
 - Port-based priority: VLAN priority in a tagged frame can be overwritten by the priority of Port VLAN ID
- Hardware auto-negotiation through serial management interface (MDIO) for Ethernet ports
- Built-in reset logic triggered by system malfunction
- I²C EEPROM for configuration

Description

The MVTX2601 is a high density, low cost, high performance, non-blocking Ethernet switch chip. A single chip provides 24 ports at 10/100 Mbps, and a CPU interface for managed and unmanaged switch applications.

The chip supports up to 64 K MAC addresses. The centralized shared memory architecture permits a very high performance packet forwarding rate at up to 3.57 M packets per second at full wire speed. The chip is optimized to provide low-cost, high-performance workgroup switching.

The Frame Buffer Memory domains utilize cost-effective, high-performance synchronous SRAM with aggregate bandwidth of 6.4 Gbps to support full wire speed on all ports simultaneously.

With delay bounded, strict priority, and/or WFQ transmission scheduling, and WRED dropping schemes, the MVTX2601 provides powerful QoS functions for various multimedia and mission-critical applications. The chip provides 4 transmission priorities and two levels of dropping precedence. Each packet is assigned a transmission priority and dropping precedence based on the VLAN priority field in a VLAN tagged frame, or the DS/TOS field, or the UDP/TCP logical port fields in IP packets. The MVTX2601 recognizes a total of 16 UDP/TCP logical ports, 8 hard-wired and 8 programmable (including one programmable range).

The MVTX2601 supports two groups of port trunking/load sharing. Each 10/100 group can contain up to 4 ports. Port trunking/load sharing can be used to group ports between interlinked switches to increase the effective network bandwidth.

In half-duplex mode, all ports support backpressure flow control to minimize the risk of losing data during long activity bursts. In full-duplex mode IEEE 802.3x flow control is provided. The MVTX2601 also supports a per-system option to enable flow control for best effort frames, even on QoS-enabled ports.

The MVTX2601 is fabricated using 0.25 micron technology. Inputs, however, are 3.3 V tolerant, and the outputs are capable of directly interfacing to LVTTTL levels. The MVTX2601 is packaged in a 553-pin Ball Grid Array package.

Table of Contents

1.0 Block Functionality	9
1.1 Frame Data Buffer (FDB) Interfaces	9
1.2 10/100 MAC Module (RMAC)	9
1.3 Configuration Interface Module	9
1.4 Frame Engine	9
1.5 Search Engine	9
1.6 LED Interface	9
1.7 Internal Memory	9
2.0 System Configuration	10
2.1 Configuration Mode	10
2.2 I ² C Interface	10
2.2.1 Start Condition	10
2.2.2 Address	10
2.2.3 Data Direction	10
2.2.4 Acknowledgment	10
2.2.5 Data	10
2.2.6 Stop Condition	11
2.3 Synchronous Serial Interface	11
2.3.1 Write Command	11
2.3.2 Read Command	12
3.0 MVTX2601 Data Forwarding Protocol	12
3.1 Unicast Data Frame Forwarding	12
3.2 Multicast Data Frame Forwarding	13
4.0 Memory Interface	13
4.1 Overview	13
4.2 Detailed Memory Information	13
4.3 Memory Requirements	13
5.0 Search Engine	14
5.1 Search Engine Overview	14
5.2 Basic Flow	14
5.3 Search, Learning and Aging	15
5.3.1 MAC Search	15
5.3.2 Learning	15
5.3.3 Aging	15
5.4 Quality of Service	15
5.5 Priority Classification Rule	16
5.6 Port Based VLAN	17
5.7 Memory Configurations	17
6.0 Frame Engine	20
6.1 Data Forwarding Summary	20
6.2 Frame Engine Details	21
6.2.1 FCB Manager	21
6.2.2 Rx Interface	21
6.2.3 RxDMA	21
6.2.4 TxQ Manager	21
6.3 Port Control	21
6.4 TxDMA	21
7.0 Quality of Service and Flow Control	21
7.1 Model	21
7.2 Four QoS Configurations	23
7.3 Delay Bound	23

Table of Contents

7.4 Strict Priority and Best Effort	24
7.5 Weighted Fair Queuing	24
7.6 WRED Drop Threshold Management Support	24
7.7 Buffer Management	25
7.7.1 Dropping When Buffers Are Scarce	26
7.8 MVTX2601 Flow Control Basics	26
7.8.1 Unicast Flow Control	27
7.8.2 Multicast Flow Control	27
7.9 Mapping to IETF Diffserv Classes	27
8.0 Port Trunking	28
8.1 Features and Restrictions	28
8.2 Unicast Packet Forwarding	28
8.3 Multicast Packet Forwarding	28
8.4 Trunking	29
9.0 Port Mirroring	29
9.1 Port Mirroring Features	29
9.2 Setting Registers for Port Mirroring	29
10.0 GPSI (7WS) Interface	30
10.1 GPSI connection	30
10.2 SCAN LINK and SCAN COL interface	31
11.0 LED Interface	31
11.1 LED Interface Introduction	31
11.2 Port Status	31
11.3 LED Interface Timing Diagram	32
12.0 Register Definition	33
12.1 MVTX2601 Register Description	33
12.2 Group 0 Address MAC Ports Group	36
12.2.1 ECR1Pn: Port N Control Register	36
12.2.2 ECR2Pn: Port N Control Register	37
12.3 Group 1 Address VLAN Group	38
12.3.1 AVTCL – VLAN Type Code Register Low	38
12.3.2 AVTCH – VLAN Type Code Register High	38
12.3.3 PVMAP00_0 – Port 00 Configuration Register 0	38
12.3.4 PVMAP00_1 – Port 00 Configuration Register 1	38
12.3.5 PVMAP00_2 – Port 00 Configuration Register 2	39
12.3.6 PVMAP00_3 – Port 00 Configuration Register 3	39
12.4 Port Configuration Register	39
12.4.1 PVMODE	40
12.5 Group 2 Address Port Trunking Group	40
12.5.1 TRUNK0_MODE– Trunk group 0 mode	40
12.5.2 TRUNK1_MODE – Trunk group 1 mode	41
12.6 Group 4 Address Search Engine Group	41
12.6.1 TX_AGE – Tx Queue Aging timer	41
12.6.2 AGETIME_LOW – MAC address aging time Low	41
12.6.3 AGETIME_HIGH –MAC address aging time High	42
12.6.4 SE_OPMODE – Search Engine Operation Mode	42
12.7 Group 5 Address Buffer Control/QOS Group	43
12.7.1 FCBAT – FCB Aging Timer	43
12.7.2 QOSC – QOS Control	43
12.7.3 FCR – Flooding Control Register	43
12.7.4 AVPML – VLAN Priority Map	44

Table of Contents

12.7.5 AVPMM – VLAN Priority Map	45
12.7.6 AVPMH – VLAN Priority Map	45
12.7.7 TOSPML – TOS Priority Map	45
12.7.8 TOSPMM – TOS Priority Map	46
12.7.9 TOSPMH – TOS Priority Map	46
12.7.10 AVDM – VLAN Discard Map	46
12.7.11 TOSDML – TOS Discard Map	47
12.7.12 BMRC - Broadcast/Multicast Rate Control	47
12.7.13 UCC – Unicast Congestion Control	48
12.7.14 MCC – Multicast Congestion Control	48
12.7.15 PR100 – Port Reservation for 10/100 ports	48
12.7.16 SFCB – Share FCB Size	49
12.7.17 C2RS – Class 2 Reserve Size	49
12.7.18 C3RS – Class 3 Reserve Size	49
12.7.19 C4RS – Class 4 Reserve Size	49
12.7.20 C5RS – Class 5 Reserve Size	50
12.7.21 C6RS – Class 6 Reserve Size	50
12.7.22 C7RS – Class 7 Reserve Size	50
12.7.23 Classes Byte Limit Set 0	50
12.7.24 Classes Byte Limit Set 1	51
12.7.25 Classes Byte Limit Set 2	51
12.7.26 Classes Byte Limit Set 3	51
12.7.27 Classes WFQ Credit Set 0	51
12.7.28 Classes WFQ Credit Set 1	52
12.7.29 Classes WFQ Credit Set 2	52
12.7.30 Classes WFQ Credit Set 3	52
12.7.31 RDRC0 – WRED Rate Control 0	53
12.7.32 RDRC1 – WRED Rate Control 1	53
12.7.33 User Defined Logical Ports and Well Known Ports	53
12.7.33.1 USER_PORT0_(0~7) – User Define Logical Port (0~7)	54
12.7.33.2 USER_PORT_[1:0]_PRIORITY - User Define Logic Port 1 and 0 Priority	54
12.7.33.3 USER_PORT_[3:2]_PRIORITY - User Define Logic Port 3 and 2 Priority	55
12.7.33.4 USER_PORT_[5:4]_PRIORITY - User Define Logic Port 5 and 4 Priority	55
12.7.33.5 USER_PORT_[7:6]_PRIORITY - User Define Logic Port 7 and 6 Priority	55
12.7.33.6 USER_PORT_ENABLE [7:0] – User Define Logic 7 to 0 Port Enables	55
12.7.33.7 WELL_KNOWN_PORT [1:0] PRIORITY- Well Known Logic Port 1 and 0 Priority	55
12.7.33.8 WELL_KNOWN_PORT [3:2] PRIORITY- Well Known Logic Port 3 and 2 Priority	56
12.7.33.9 WELL_KNOWN_PORT [5:4] PRIORITY- Well Known Logic Port 5 and 4 Priority	56
12.7.33.10 WELL_KNOWN_PORT [7:6] PRIORITY- Well Known Logic Port 7 and 6 Priority	56
12.7.33.11 WELL_KNOWN_PORT_ENABLE [7:0] – Well Known Logic 7 to 0 Port Enables	56
12.7.33.12 RLOWL – User Define Range Low Bit 7:0	57
12.7.33.13 RLOWH – User Define Range Low Bit 15:8	57
12.7.33.14 RHIGHL – User Define Range High Bit 7:0	57
12.7.33.15 RHIGHH – User Define Range High Bit 15:8	57
12.7.33.16 RPRIORITY – User Define Range Priority	57
12.8 Group 6 Address MISC Group	57
12.8.1 MII_OP0 – MII Register Option 0	57
12.8.2 MII_OP1 – MII Register Option 1	58
12.8.3 FEN – Feature Register	58
12.8.4 MIIC0 – MII Command Register 0	59
12.8.5 MIIC1 – MII Command Register 1	59
12.8.6 MIIC2 – MII Command Register 2	59

Table of Contents

12.8.7 MIIC3 – MII Command Register 3	59
12.8.8 MIID0 – MII Data Register 0	60
12.8.9 MIID1 – MII Data Register 1	60
12.8.10 LED Mode – LED Control	60
12.8.11 CHECKSUM - EEPROM Checksum	60
12.9 Group 7 Address Port Mirroring Group	61
12.9.1 MIRROR1_SRC – Port Mirror source port	61
12.9.2 MIRROR1_DEST – Port Mirror destination	61
12.9.3 MIRROR2_SRC – Port Mirror source port	61
12.9.4 MIRROR2_DEST – Port Mirror destination	62
12.10 Group F Address CPU Access Group	62
12.10.1 GCR-Global Control Register	62
12.10.2 DCR-Device Status and Signature Register	63
12.10.3 DCR1-Chip status	63
12.10.4 DPST – Device Port Status Register	64
12.10.5 DTST – Data read back register	65
12.10.6 PLLCR - PLL Control Register	65
12.10.7 LCLK - LA_CLK delay from internal OE_CLK	65
12.10.8 OECLK - Internal OE_CLK delay from SCLK	66
12.10.9 DA – DA Register	66
13.0 BGA and Ball Signal Descriptions	67
13.1 BGA Views (Top-View)	67
13.1.1 Encapsulated View	67
13.2 Ball – Signal Descriptions	68
13.2.1 Ball Signal Descriptions	68
13.3 Ball – Signal Name	76
13.4 AC/DC Timing	81
13.4.1 Absolute Maximum Ratings	81
13.4.2 DC Electrical Characteristics	81
13.4.3 Recommended Operating Conditions	82
13.4.4 Typical Reset & Bootstrap Timing Diagram	82
13.5 Local Frame Buffer SDRAM Memory Interface	83
13.5.1 Local SDRAM Memory Interface	83
13.6 AC Characteristics	84
13.6.1 Reduced Media Independent Interface	84
13.6.2 LED Interface	85
13.6.3 SCANLINK SCANCOL Output Delay Timing	85
13.6.4 MDIO Input Setup and Hold Timing	87
13.6.5 I ² C Input Setup Timing	88
13.6.6 Serial Interface Setup Timing	89

List of Figures

Figure 1 - MVTX260 1 System Block Diagram	1
Figure 2 - Data Transfer Format for I2C Interface	10
Figure 3 - MVTX2601 SRAM Interface Block Diagram (DMAs for 10/1000 Ports Only)	13
Figure 4 - Memory Map	14
Figure 5 - Priority Classification Rule	16
Figure 6 - Memory Configuration for 1 Bank, 1 Layer, 1 MB Total	18
Figure 7 - Memory Configuration for 1 Bank, 2 Layers, 2 MB Total	19
Figure 8 - Memory Configuration for 1 Bank, 1 Layer, 2 MB	20
Figure 9 - Buffer Partition Scheme Used to Implement MVTX2601 Buffer Management	26
Figure 10 - GPSI (7WS) Mode Connection Diagram	30
Figure 11 - SCAN LINK and SCAN COLLISION Status Diagram	31
Figure 12 - Timing Diagram of LED Interface	32
Figure 13 - Typical Reset & Bootstrap Timing Diagram	82
Figure 14 - Local Memory Interface – Input Setup and Hold Timing	83
Figure 15 - Local Memory Interface - Output Valid Delay Timing	83
Figure 16 - AC Characteristics – Reduced Media Independent Interface	84
Figure 17 - AC Characteristics – Reduced Media Independent Interface	85
Figure 18 - AC Characteristics – LED Interface	85
Figure 19 - SCANLINK SCANCOL Output Delay Timing	85
Figure 20 - SCANLINK, SCANCOL Setup Timing	86
Figure 21 - MDIO Input Setup and Hold Timing	87
Figure 22 - MDIO Output Delay Timing	87
Figure 23 - I ² C Input Setup Timing	88
Figure 24 - I ² C Output Delay Timing	88
Figure 25 - Serial Interface Setup Timing	89
Figure 26 - Serial Interface Output Delay Timing	89

List of Tables

Table 1 - Memory Configuration	14
Table 2 - PVMAP Register	17
Table 3 - Supported Memory Configurations (SBRAM Mode)	17
Table 4 - Options for Memory Configuration	18
Table 5 - Two Dimensional World Traffic	22
Table 6 - Four QoS Configurations for a 10/100Mbps Port	23
Table 7 - WRED Drop Thresholds	24
Table 8 - Mapping between MVTX2601 and IETF Diffserv Classes for 10/100 Ports	27
Table 9 - MVTX2601 Features Enabling IETF Diffserv Standards	28
Table 10 - Reset & Bootstrap Timing	83
Table 11 - AC Characteristics – Reduced Media Independent Interface	84
Table 12 - AC Characteristics – LED Interface	85
Table 13 - SCANLINK, SCANCOL Timing	86
Table 14 - MDIO Timing	87
Table 15 - I2C Timing	88
Table 16 - Serial Interface Timing	89

1.0 Block Functionality

1.1 Frame Data Buffer (FDB) Interfaces

The FDB interface supports pipelined synchronous burst SRAM (SBRAM) memory at 100 MHz. To ensure a non-blocking switch, one memory domain with a 64 bit wide memory bus is required. At 100 MHz, the aggregate memory bandwidth is 6.4 Gbps, which is enough to support 24 10/100 Mbps.

The Switching Database is also located in the external SBRAM; it is used for storing MAC addresses and their physical port number.

1.2 10/100 MAC Module (RMAC)

The 10/100 Media Access Control module provides the necessary buffers and control interface between the Frame Engine (FE) and the external physical device (PHY). The MVTX2601 has two interfaces, RMI or Serial (only for 10 M). The 10/100 MAC of the MVTX2601 device meets the IEEE 802.3 specification. It is able to operate in either Half or Full Duplex mode with a back pressure/flow control mechanism. In addition, it will automatically retransmit upon collision for up to 16 total transmissions. The PHY addresses for 24 10/100 MAC are from 08h to 1Fh.

1.3 Configuration Interface Module

The MVTX2601 supports a serial and an I²C interface, which provides an easy way to configure the system. Once configured, the resulting configuration can be stored in an I²C EEPROM.

1.4 Frame Engine

The main function of the frame engine is to forward a frame to its proper destination port or ports. When a frame arrives, the frame engine parses the frame header (64 bytes) and formulates a switching request, sent to the search engine to resolve the destination port. The arriving frame is moved to the FDB. After receiving a switch response from the search engine, the frame engine performs transmission scheduling based on the frame's priority. The frame engine forwards the frame to the MAC module when the frame is ready to be sent.

1.5 Search Engine

The Search Engine resolves the frame's destination port or ports according to the destination MAC address (L2). It also performs MAC learning, priority assignment and trunking functions.

1.6 LED Interface

The LED interface provides a serial interface for carrying 24 port status signals.

1.7 Internal Memory

Several internal tables are required and are described as follows:

- Frame Control Block (FCB) - Each FCB entry contains the control information of the associated frame stored in the FDB, e.g., frame size, read/write pointer, transmission priority, etc.
- MCT Link Table - The MCT Link Table stores the linked list of MCT entries that have collisions in the external MAC Table. The external MAC table is located in the FDB Memory.

Note: the external MAC table is located in the external SBRAM Memory.

2.0 System Configuration

2.1 Configuration Mode

The MVTX2601 can be configured by EEPROM (24C02 or compatible) via an I²C interface at boot time or via a synchronous serial interface during operation.

2.2 I²C Interface

The I²C interface uses two bus lines, a serial data line (SDA) and a serial clock line (SCL). The SCL line carries the control signals that facilitate the transfer of information from EEPROM to the switch. Data transfer is 8-bit serial and bidirectional, at 50 Kbps. Data transfer is performed between master and slave IC using a request / acknowledgment style of protocol. The master IC generates the timing signals and terminates data transfer. Figure 2 depicts the data transfer format.

START	SLAVE ADDRESS	R/W	ACK	DATA 1 (8 bits)	ACK	DATA 2	ACK	DATA M	ACK	STOP
-------	---------------	-----	-----	-----------------	-----	--------	-----	--------	-----	------

Figure 2 - Data Transfer Format for I²C Interface

2.2.1 Start Condition

Generated by the master (in our case, the MVTX2601). The bus is considered to be busy after the Start condition is generated. The Start condition occurs if while the SCL line is High, there is a High-to-Low transition of the SDA line.

Other than in the Start condition (and Stop condition), the data on the SDA line must be stable during the High period of SCL. The High or Low state of SDA can only change when SCL is Low. In addition, when the I²C bus is free, both lines are High.

2.2.2 Address

The first byte after the Start condition determines which slave the master will select. The slave in our case is the EEPROM. The first seven bits of the first data byte make up the slave address.

2.2.3 Data Direction

The eighth bit in the first byte after the Start condition determines the direction (R/W) of the message. A master transmitter sets this bit to W; a master receiver sets this bit to R.

2.2.4 Acknowledgment

Like all clock pulses, the acknowledgment-related clock pulse is generated by the master. However, the transmitter releases the SDA line (High) during the acknowledgment clock pulse. Furthermore, the receiver must pull down the SDA line during the acknowledge pulse so that it remains stable Low during the High period of this clock pulse. An acknowledgment pulse follows every byte transfer.

If a slave receiver does not acknowledge after any byte, then the master generates a Stop condition and aborts the transfer.

If a master receiver does not acknowledge after any byte, then the slave transmitter must release the SDA line to let the master generate the Stop condition.

2.2.5 Data

After the first byte containing the address, all bytes that follow are data bytes. Each byte must be followed by an acknowledge bit. Data is transferred MSB first.

2.2.6 Stop Condition

Generated by the master. The bus is considered to be free after the Stop condition is generated. The Stop condition occurs if while the SCL line is High, there is a Low-to-High transition of the SDA line.

The I²C interface serves the function of configuring the MVTX2601 at boot time. The master is the MVTX2601 and the slave is the EEPROM memory.

2.3 Synchronous Serial Interface

The synchronous serial interface serves the function of configuring the MVTX2601 *not* at boot time but via a PC. The PC serves as master and the MVTX2601 serves as slave. The protocol for the synchronous serial interface is nearly identical to the I²C protocol. The main difference is that there is no acknowledgment bit after each byte of data transferred.

The unmanaged MVTX2601 uses a synchronous serial interface to program the internal registers. To reduce the number of signals required, the register address, command and data are shifted in serially through the D0 pin. STROBE- pin is used as the shift clock. AUTOFD- pin is used as data return path.

Each command consists of four parts.

START pulse

Register Address

Read or Write command

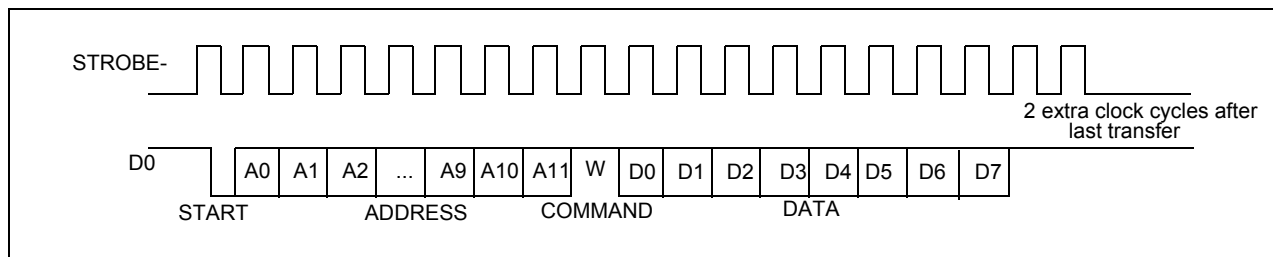
Data to be written or read back

Any command can be aborted in the middle by sending a ABORT pulse to the MVTX2601.


A START command is detected when D0 is sampled high when STROBE- rise and D0 is sampled low when STROBE- fall.

An ABORT command is detected when D0 is sampled low when STROBE- rise and D0 is sampled high when STROBE- fall.

2.3.1 Write Command



2.3.2 Read Command



All registers in MVTX2601 can be modified through this synchronous serial interface.

3.0 MVTX2601 Data Forwarding Protocol

3.1 Unicast Data Frame Forwarding

When a frame arrives, it is assigned a handle in memory by the Frame Control Buffer Manager (FCB Manager). An FCB handle will always be available because of advance buffer reservations.

The memory (SRAM) interface consists of a 64-bit bus connected to SRAM bank. The Receive DMA (RxDMA) is responsible for multiplexing the data and the address. On a port's "turn," the RxDMA will move 8 bytes (or up to the end-of-frame) from the port's associated Rx FIFO into memory (Frame Data Buffer, or FDB).

Once an entire frame has been moved to the FDB and a good end-of-frame (EOF) has been received, the Rx interface makes a switch request. The RxDMA arbitrates among multiple switch requests.

The switch request consists of the first 64 bytes of a frame, containing among other things, the source and destination MAC addresses of the frame. The search engine places a switch response in the switch response queue of the frame engine when done. Among other information, the search engine will have resolved the destination port of the frame and will have determined that the frame is unicast.

After processing the switch response, the Transmission Queue Manager (TxQ manager) of the frame engine is responsible for notifying the destination port that it has a frame to forward to it. But first, the TxQ manager has to decide whether or not to drop the frame, based on global FDB reservations and usage, as well as TxQ occupancy at the destination. If the frame is not dropped, then the TxQ manager links the frame's FCB to the correct per-port-per-class TxQ. Unicast TxQ's are linked lists of transmission jobs, represented by their associated frames' FCB's. There is one linked list for each transmission class for each port. There are 4 transmission classes for each of the 24 10/ 100 ports

The TxQ manager is responsible for scheduling transmission among the queues representing different classes for a port. When the port control module determines that there is room in the MAC Transmission FIFO (Tx FIFO) for another frame, it requests the handle of a new frame from the TxQ manager. The TxQ manager chooses among the head-of-line (HOL) frames from the per-class queues for that port, using a Zarlink Semiconductor scheduling algorithm.

The Transmission DMA (Tx DMA) is responsible for multiplexing the data and the address. On a port's turn, the Tx DMA will move 8 bytes (or up to the EOF) from memory into the port's associated Tx FIFO. After reading the EOF, the port control requests a FCB release for that frame. The Tx DMA arbitrates among multiple buffer release requests.

The frame is transmitted from the Tx FIFO to the line.

3.2 Multicast Data Frame Forwarding

After receiving the switch response, the TxQ manager has to make the dropping decision. A global decision to drop can be made, based on global FDB utilization and reservations. If so, then the FCB is released and the frame is dropped. In addition, a selective decision to drop can be made, based on the TxQ occupancy at some subset of the multicast packet's destinations. If so, then the frame is dropped at some destinations but not others and the FCB is not released.

If the frame is not dropped at a particular destination port, then the TxQ manager formats an entry in the multicast queue for that port and class. Multicast queues are physical queues (unlike the linked lists for unicast frames). There are 2 multicast queues for each of the 24 10/100 ports. The queue with higher priority has room for 32 entries and the queue with lower priority has room for 64 entries. There is one multicast queue for every two priority classes. For the 10/100 ports to map the 8 transmit priorities into 2 multicast queues, the 2 LSB are discarded.

During scheduling, the TxQ manager treats the unicast queue and the multicast queue of the same class as one logical queue. The older head of line of the two queues is forwarded first.

The port control requests a FCB release only after the EOF for the multicast frame has been read by all ports to which the frame is destined.

4.0 Memory Interface

4.1 Overview

The MVTX2601 provides a 64-bit wide SRAM bank. Each DMA can read and write from the SRAM bank. The following figure provides an overview of the MVTX2601 SRAM bank.

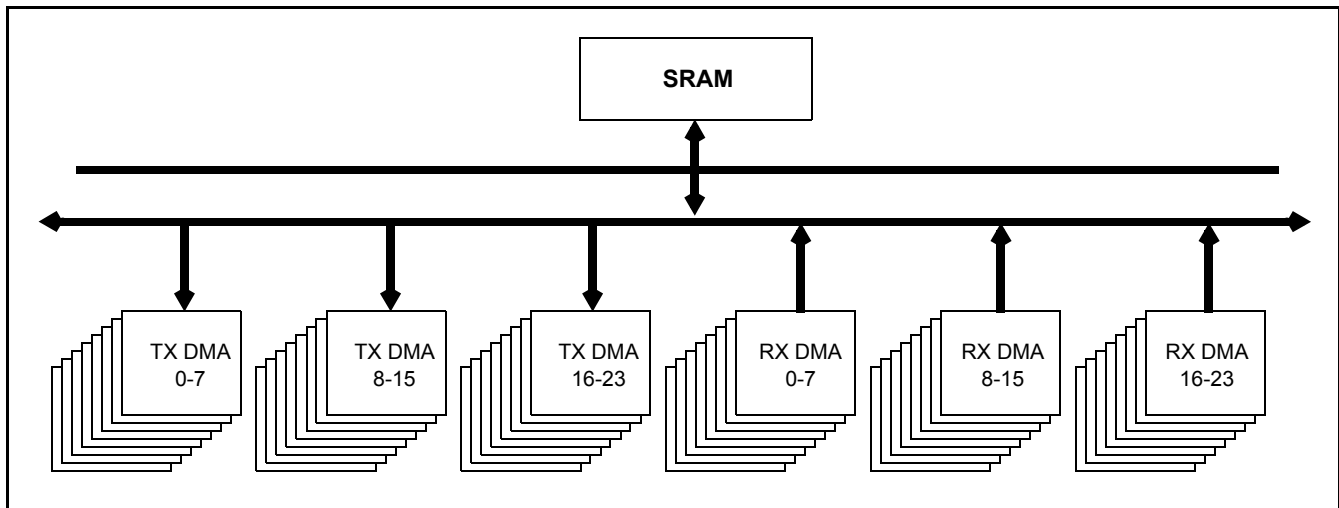


Figure 3 - MVTX2601 SRAM Interface Block Diagram (DMAs for 10/1000 Ports Only)

4.2 Detailed Memory Information

Because the bus for each bank is 64 bits wide, frames are broken into 8-byte granules, written to and read from memory.

4.3 Memory Requirements

To support 64 K MAC address, 2 MB memory is required. When VLAN support is enabled, 512 entries of the MAC address table are used for storing the VLAN ID at VLAN Index Mapping Table.

Up to 1 K Ethernet frame buffers are supported and they will use 1.5 MB of memory. Each frame uses 1536 bytes. The maximum system memory requirement is 2 MB. If less memory is desired, the configuration can scale down.

Memory Bank	Frame Buffer	Max MAC Address
1 M	1 K	32 K
2 M	2 K	64 K

Table 1 - Memory Configuration




Figure 4 - Memory Map

5.0 Search Engine

5.1 Search Engine Overview

The MVTX2601 search engine is optimized for high throughput searching, with enhanced features to support:

- Up to 64 K MAC addresses
- 2 groups of port trunking
- Traffic classification into 4 transmission priorities, and 2 drop precedence levels
- Flooding, Broadcast, Multicast Storm Control
- MAC address learning and aging
- Port based VLAN

5.2 Basic Flow

Shortly after a frame enters the MVTX2601 and is written to the Frame Data Buffer (FDB), the frame engine generates a Switch Request, which is sent to the search engine. The switch request consists of the first 64 bytes of the frame, which contain all the necessary information for the search engine to perform its task. When the search engine is done, it writes to the Switch Response Queue and the frame engine uses the information provided in that queue for scheduling and forwarding.

In performing its task, the search engine extracts and compresses the useful information from the 64-byte switch request. Among the information extracted are the source and destination MAC addresses, the transmission and discard priorities, whether the frame is unicast or multicast. Requests are sent to the external SRAM to locate the associated entries in the external hash table.

When all the information has been collected from external SRAM, the search engine has to compare the MAC address on the current entry with the MAC address for which it is searching. If it is not a match, the process is repeated on the internal MCT Table. All MCT entries other than the first of each linked list are maintained internal to the chip. If the desired MAC address is still not found, then the result is either learning (source MAC address unknown) or flooding (destination MAC address unknown).

In addition, port based VLAN information is used to select the correct set of destination ports for the frame (for multicast) or to verify that the frame's destination port is associated with the VLAN (for unicast).

If the destination MAC address belongs to a port trunk, then the trunk number is retrieved instead of the port number. But on which port of the trunk will the frame be transmitted? This is easily computed using a hash of the source and destination MAC addresses.

As stated earlier, when all the information is compiled the switch response is generated. The search engine also interacts with the CPU with regard to learning and aging.

5.3 Search, Learning and Aging

5.3.1 MAC Search

The search block performs source MAC address and destination MAC address searching. As we indicated earlier, if a match is not found, then the next entry in the linked list must be examined and so on until a match is found or the end of the list is reached.

The port based VLAN bitmap is used to determine whether the frame should be forwarded to the outgoing port. When the egress port is not included in the ingress port VLAN bitmap, the packet is discarded.

The MAC search block is also responsible for updating the source MAC address timestamp and the VLAN port association timestamp, used for aging.

5.3.2 Learning

The learning module learns new MAC addresses and performs port change operations on the MCT database. The goal of learning is to update this database as the networking environment changes over time. Learning and port change will be performed based on memory slot availability only.

5.3.3 Aging

Aging time is controlled by register 400h and 401h.

The aging module scans and ages MCT entries based on a programmable "age out" time interval. As we indicated earlier, the search module updates the source MAC address timestamps for each frame it processes. When an entry is ready to be aged, the entry is removed from the table.

5.4 Quality of Service

Quality of Service (QoS) refers to the ability of a network to provide better service to selected network traffic over various technologies. Primary goals of QoS include dedicated bandwidth, controlled jitter and latency (required by some real-time and interactive traffic) and improved loss characteristics.

Traditional Ethernet networks have had no prioritization of traffic. Without a protocol to prioritize or differentiate traffic, a service level known as "best effort" attempts to get all the packets to their intended destinations with minimum delay; however, there are no guarantees. In a congested network or when a low-performance switch/router is overloaded, "best effort" becomes unsuitable for delay-sensitive traffic and mission-critical data transmission.

The advent of QoS for packet-based systems accommodates the integration of delay-sensitive video and multimedia traffic onto any existing Ethernet network. It also alleviates the congestion issues that have previously plagued such “best effort” networking systems. QoS provides Ethernet networks with the breakthrough technology to prioritize traffic and ensure that a certain transmission will have a guaranteed minimum amount of bandwidth.

Extensive core QoS mechanisms are built into the MVTX2601 architecture to ensure policy enforcement and buffering of the ingress port, as well as weighted fair-queue (WFQ) scheduling at the egress port.

In the MVTX2601, QoS-based policies sort traffic into a small number of classes and mark the packets accordingly. The QoS identifier provides specific treatment to traffic in different classes, so that different quality of service is provided to each class. Frame and packet scheduling and discarding policies are determined by the class to which the frames and packets belong. For example, the overall service given to frames and packets in the premium class will be better than that given to the standard class; the premium class is expected to experience lower loss rate or delay.

The MVTX2601 supports the following QoS techniques:

- In a port-based setup, any station connected to the same physical port of the switch will have the same transmit priority
- In a tag-based setup, a 3-bit field in the VLAN tag provides the priority of the packet. This priority can be mapped to different queues in the switch to provide QoS.
- In a TOS/DS-based set up, TOS stands for “Type of Service” that may include “minimize delay,” “maximize throughput” or “maximize reliability.” Network nodes may select routing paths or forwarding behaviours that are suitably engineered to satisfy the service request.
- In a logical port-based set up, a logical port provides the application information of the packet. Certain applications are more sensitive to delays than others; using logical ports to classify packets can help speed up delay sensitive applications such as VoIP.

5.5 Priority Classification Rule

Figure 5 shows the MVTX2601 priority classification rule.

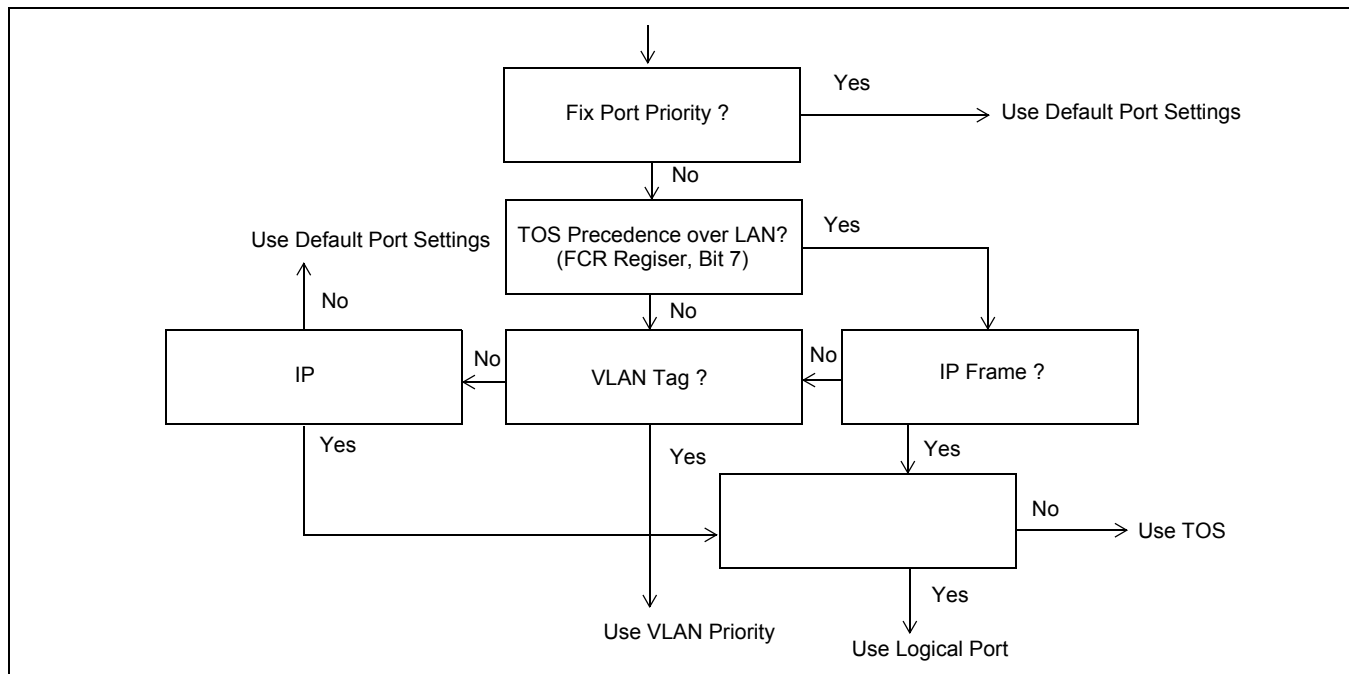


Figure 5 - Priority Classification Rule

5.6 Port Based VLAN

An administrator can use the PVMAP Registers to configure the MVTX2601 for port-based VLAN. For example, ports 1-3 might be assigned to the Marketing VLAN, ports 4-6 to the Engineering VLAN, and ports 7-9 to the Administrative VLAN. The MVTX2601 determines the VLAN membership of each packet by noting the port on which it arrives. From there, the MVTX2601 determines which outgoing port(s) is/are eligible to transmit each packet, or whether the packet should be discarded.

	Destination Port Numbers Bit Map				
Port Registers	23	...	2	1	0
Register for Port #0 PVMAP00_0[7:0] to PVMAP00_2[7:0]	0		1	1	0
Register for Port #1 PVMAP01_0[7:0] to PVMAP01_2[7:0]	0		1	0	1
Register for Port #2 PVMAP02_0[7:0] to PVMAP02_2[7:0]	0		0	0	0
...					
Register for Port #23 PVMAP23_0[7:0] to PVMAP23_2[7:0]	0		0	0	0

Table 2 - PVMAP Register

For example, in the above table, a "1" denotes that an outgoing port is eligible to receive a packet from an incoming port. A 0 (zero) denotes that an outgoing port is not eligible to receive a packet from an incoming port.

In this example:

- Data packets received at port #0 are eligible to be sent to outgoing ports 1 and 2
- Data packets received at port #1 are eligible to be sent to outgoing ports 0, and 2
- Data packets received at port #2 are not eligible to be sent to ports 0 and 1

5.7 Memory Configurations

The MVTX2601 supports the following memory configurations. It supports 1 M and 2 M configurations.

Configuration	1 M (Bootstrap pin TSTOUT7 = open)	2 M (Bootstrap pin TSTOUT7 = pull down)	Connections
Single Layer (Bootstrap pin TSTOUT13 = open)	Two 128 K x 32 SRAM/bank or One 128 K x 64 SRAM/bank	Two 256 K x 32 SRAM/bank	Connect 0E# and WE#
Double Layer (Bootstrap pin TSTOUT13 = pull down)	NA	Four 128 K x 32 SRAM/bank or Two 128 K x 64 SRAM/bank	Connect 0E0# and WE0# Connect 0E1# and WE1#

Table 3 - Supported Memory Configurations (SBRAM Mode)

	Frame data Buffer					
	Only Bank A		Bank A and Bank B		Bank A and Bank B	
	1 M (SRAM)	2 M (SRAM)	1 M/bank (SRAM)	2 M/bank (SRAM)	1 M/bank (ZBT SRAM)	2 M/bank (ZBT SRAM)
MVTX2601	X	X				
MVTX2602	X	X				
MVTX2603			X	X		
MVTX2603 (Gigabit ports in 2 giga mode)					X (125 Mhz)	X (125 Mhz)
MVTX2604			X	X		
MVTX2604 (Gigabit ports in 2 giga mode)					X (125 Mhz)	X (125 Mhz)

Table 4 - Options for Memory Configuration

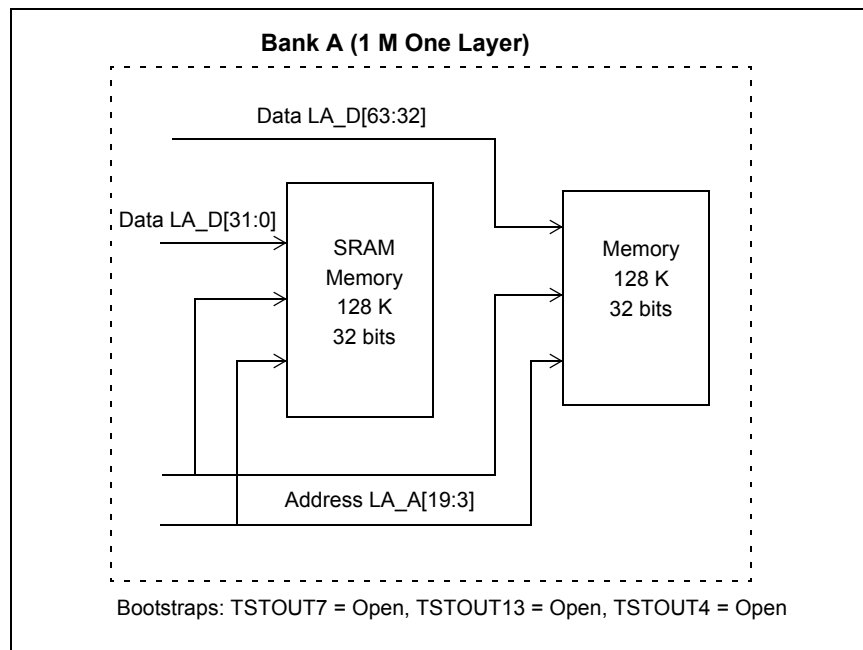


Figure 6 - Memory Configuration for 1 Bank, 1 Layer, 1 MB Total

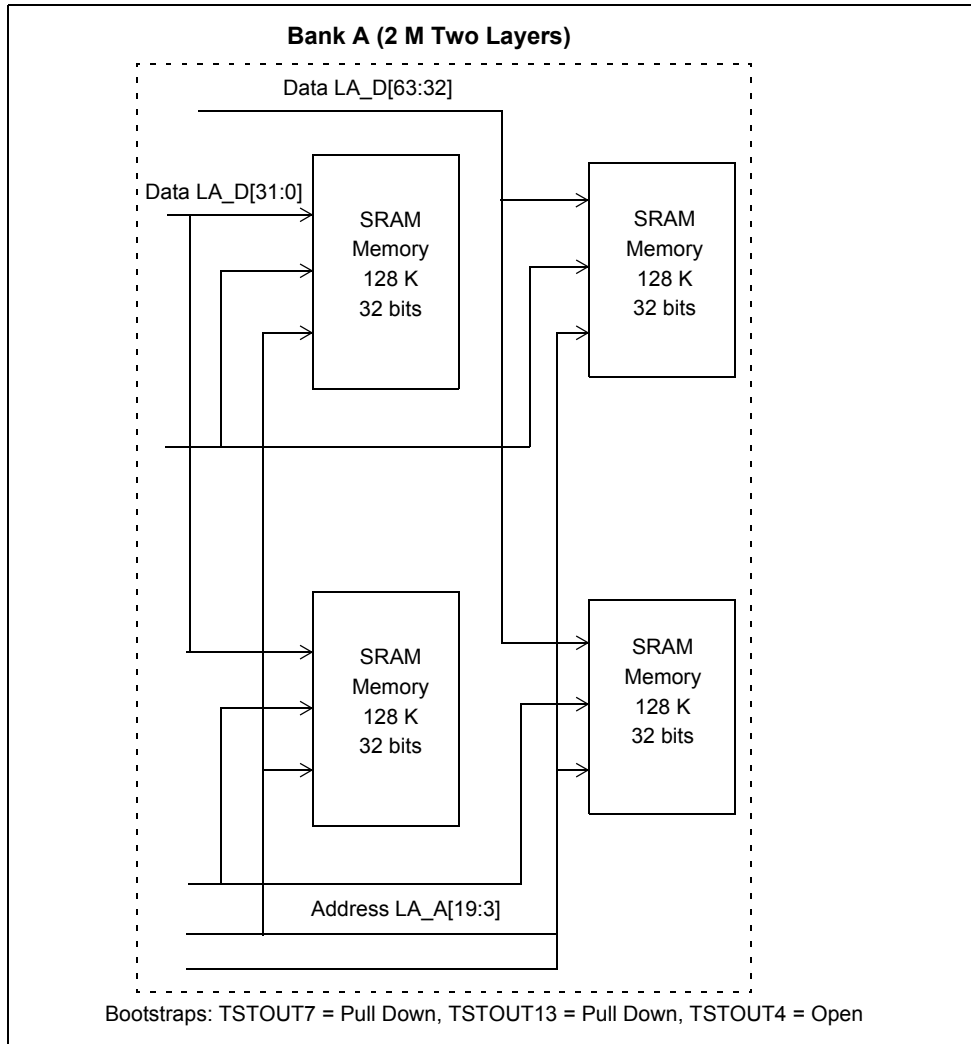


Figure 7 - Memory Configuration for 1 Bank, 2 Layers, 2 MB Total

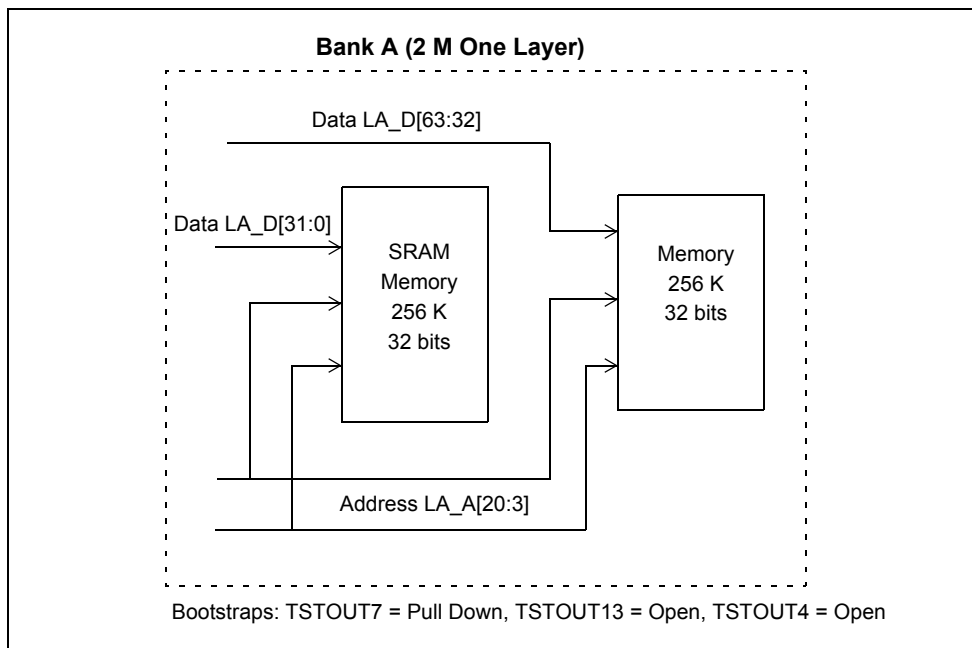


Figure 8 - Memory Configuration for 1 Bank, 1 Layer, 2 MB

6.0 Frame Engine

6.1 Data Forwarding Summary

When a frame enters the device at the RxMAC, the RxDMA will move the data from the MAC RxFIFO to the FDB. Data is moved in 8-byte granules in conjunction with the scheme for the SRAM interface.

A switch request is sent to the Search Engine. The Search Engine processes the switch request.

A switch response is sent back to the Frame Engine and indicates whether the frame is unicast or multicast and its destination port or ports.

A Transmission Scheduling Request is sent in the form of a signal notifying the TxQ manager. Upon receiving a Transmission Scheduling Request, the device will format an entry in the appropriate Transmission Scheduling Queue (TxSch Q) or Queues. There are four TxSch Q for each 10/100, one for each priority. Creation of a queue entry either involves linking a new job to the appropriate linked list if unicast or adding an entry to a physical queue if multicast.

When the port is ready to accept the next frame, the TxQ manager will get the head-of-line (HOL) entry of one of the TxSch Qs, according to the transmission scheduling algorithm (so as to ensure per-class quality of service). The unicast linked list and the multicast queue for the same port-class pair are treated as one logical queue. The older HOL between the two queues goes first. For 10/100 ports multicast queue 0 is associated with unicast queue 0 and multicast queue 1 is associated with unicast queue 2.

The TxDMA will pull frame data from the memory and forward it granule-by-granule to the MAC TxFIFO of the destination port.

6.2 Frame Engine Details

This section briefly describes the functions of each of the modules of the MVTX2601 frame engine.

6.2.1 FCB Manager

The FCB manager allocates FCB handles to incoming frames and releases FCB handles upon frame departure. The FCB manager is also responsible for enforcing buffer reservations and limits. The default values can be determined by referring to Chapter 8. In addition, the FCB manager is responsible for buffer aging and for linking unicast forwarding jobs to their correct TxSch Q. The buffer aging can be enabled or disabled by the bootstrap pin and the aging time is defined in register FCBAT.

6.2.2 Rx Interface

The Rx interface is mainly responsible for communicating with the RxMAC. It keeps track of the start and end of frame and frame status (good or bad). Upon receiving an end of frame that is good, the Rx interface makes a switch request.

6.2.3 RxDMA

The RxDMA arbitrates among switch requests from each Rx interface. It also buffers the first 64 bytes of each frame for use by the search engine when the switch request has been made.

6.2.4 TxQ Manager

First, the TxQ manager checks the per-class queue status and global reserved resource situation and using this information, makes the frame dropping decision after receiving a switch response. If the decision is not to drop, the TxQ manager requests that the FCB manager link the unicast frame's FCB to the correct per-port-per-class TxQ. If multicast, the TxQ manager writes to the multicast queue for that port and class. The TxQ manager can also trigger source port flow control for the incoming frame's source if that port is flow control enabled. Second, the TxQ manager handles transmission scheduling; it schedules transmission among the queues representing different classes for a port. Once a frame has been scheduled, the TxQ manager reads the FCB information and writes to the correct port control module.

6.3 Port Control

The port control module calculates the SRAM read address for the frame currently being transmitted. It also writes start of frame information and an end of frame flag to the MAC TxFIFO. When transmission is done, the port control module requests that the buffer be released.

6.4 TxDMA

The TxDMA multiplexes data and address from port control and arbitrates among buffer release requests from the port control modules.

7.0 Quality of Service and Flow Control

7.1 Model

Quality of service is an all-encompassing term for which different people have different interpretations. In general, the approach to quality of service described here assumes that we do not know the offered traffic pattern. We also assume that the incoming traffic is not policed or shaped. Furthermore, we assume that the network manager knows his applications, such as voice, file transfer or web browsing and their relative importance. The manager can then subdivide the applications into classes and set up a service contract with each. The contract may consist of

bandwidth or latency assurances per class. Sometimes it may even reflect an estimate of the traffic mix offered to the switch. As an added bonus, although we do not assume anything about the arrival pattern, if the incoming traffic is policed or shaped, we may be able to provide additional assurances about our switch's performance.

Table 5 shows examples of QoS applications with three transmission priorities, but best effort (P0) traffic may form a fourth class with no bandwidth or latency assurances.

Goals	Total Assured Bandwidth (user defined)	Low Drop Probability (low-drop)	High Drop Probability (high-drop)
Highest transmission priority, P3	50 Mbps	Apps: phone calls, circuit emulation. Latency: < 1 ms. Drop: No drop if P3 not oversubscribed.	Apps: training video. Latency: < 1 ms. Drop: No drop if P3 not oversubscribed; first P3 to drop otherwise.
Middle transmission priority, P2	37.5 Mbps	Apps: interactive apps, Web business. Latency: < 4-5 ms. Drop: No drop if P2 not oversubscribed.	Apps: non-critical interactive apps. Latency: < 4-5 ms. Drop: No drop if P2 not oversubscribed; first P2 to drop otherwise.
Low transmission priority, P1	12.5 Mbps	Apps: emails, file backups. Latency: < 16 ms desired, but not critical. Drop: No drop if P1 not oversubscribed.	Apps: casual web browsing. Latency: < 16 ms desired, but not critical. Drop: No drop if P1 not oversubscribed; first to drop otherwise.
Total	100 Mbps		

Table 5 - Two Dimensional World Traffic

A class is capable of offering traffic that exceeds the contracted bandwidth. A well-behaved class offers traffic at a rate no greater than the agreed-upon rate. By contrast, a misbehaving class offers traffic that exceeds the agreed-upon rate. A misbehaving class is formed from an aggregation of misbehaving microflows. To achieve high link utilization, a misbehaving class is allowed to use any idle bandwidth. However, such leniency must not degrade the quality of service (QoS) received by well-behaved classes.

As Table 6 illustrates, the six traffic types may each have their own distinct properties and applications. As shown, classes may receive bandwidth assurances or latency bounds. In the table, P3, the highest transmission class, requires that all frames be transmitted within 1 ms, and receives 50% of the 100 Mbps of bandwidth at that port.

Best-effort (P0) traffic forms a fourth class that only receives bandwidth when none of the other classes have any traffic to offer. It is also possible to add a fourth class that has strict priority over the other three; if this class has even one frame to transmit, then it goes first. In the MVTX2601, each 10/100 Mbps port will support four total classes, and each 1000 Mbps port will support eight classes. We will discuss the various modes of scheduling these classes in the next section.

In addition, each transmission class has two subclasses, high-drop and low-drop. Well-behaved users should rarely lose packets. But poorly behaved users – users who send frames at too high a rate – will encounter frame loss and the first to be discarded will be high-drop. Of course, if this is insufficient to resolve the congestion, eventually some low-drop frames are dropped and then all frames in the worst case.

Table 6 shows that different types of applications may be placed in different boxes in the traffic table. For example, casual web browsing fits into the category of high-loss, high-latency-tolerant traffic, whereas VoIP fits into the category of low-loss, low-latency traffic.

7.2 Four QoS Configurations

There are four basic pieces to QoS scheduling in the MVTX2601: strict priority (SP), delay bound, weighted fair queuing (WFQ), and best effort (BE). Using these four pieces, there are four different modes of operation, as shown in Tables 4 and 5. For 10/100 Mbps ports, these modes are selected by the following registers:

QOSC24 [7:6]	CREDIT_C00
QOSC28 [7:6]	CREDIT_C10
QOSC32 [7:6]	CREDIT_C20
QOSC36 [7:6]	CREDIT_C30

	P3	P2	P1	P0
Op1 (default)	Delay Bound			BE
Op2	SP	Delay Bound		BE
Op3	SP	WFQ		
Op4	WFQ			

Table 6 - Four QoS Configurations for a 10/100Mbps Port

The default configuration for a 10/100 Mbps port is three delay-bounded queues and one best-effort queue. The delay bounds per class are 0.8 ms for P3, 2 ms for P2, and 12.8 ms for P1. Best effort traffic is only served when there is no delay-bounded traffic to be served.

We have a second configuration for a 10/100 Mbps port in which there is one strict priority queue, two delay bounded queues and one best effort queue. The delay bounds per class are 3.2 ms for P2 and 12.8 ms for P1. If the user is to choose this configuration, it is important that P3 (SP) traffic be either policed or implicitly bounded (e.g., if the incoming P3 traffic is very light and predictably patterned). Strict priority traffic, if not admission-controlled at a prior stage to the MVTX2601 can have an adverse effect on all other classes' performance.

The third configuration for a 10/100 Mbps port contains one strict priority queue and three queues receiving a bandwidth partition via WFQ. As in the second configuration, strict priority traffic needs to be carefully controlled. In the fourth configuration, all queues are served using a WFQ service discipline.

7.3 Delay Bound

In the absence of a sophisticated QoS server and signaling protocol, the MVTX2601 may not know the mix of incoming traffic ahead of time. To cope with this uncertainty, our delay assurance algorithm dynamically adjusts its scheduling and dropping criteria, guided by the queue occupancies and the due dates of their head-of-line (HOL) frames. As a result, we assure latency bounds for all admitted frames with high confidence, even in the presence of system-wide congestion. Our algorithm identifies misbehaving classes and intelligently discards frames at no detriment to well-behaved classes. Our algorithm also differentiates between high-drop and low-drop traffic with a weighted random early drop (WRED) approach. Random early dropping prevents congestion by randomly dropping a percentage of high-drop frames even before the chip's buffers are completely full, while still largely sparing low-drop frames. This allows high-drop frames to be discarded early, as a sacrifice for future low-drop frames. Finally, the delay bound algorithm also achieves bandwidth partitioning among classes.

7.4 Strict Priority and Best Effort

When strict priority is part of the scheduling algorithm, if a queue has even one frame to transmit, it goes first. Two of our four QoS configurations include strict priority queues. The goal is for strict priority classes to be used for IETF expedited forwarding (EF), where performance guarantees are required. As we have indicated, it is important that strict priority traffic be either policed or implicitly bounded, so as to keep from harming other traffic classes.

When best effort is part of the scheduling algorithm, a queue only receives bandwidth when none of the other classes have any traffic to offer. Two of our four QoS configurations include best effort queues. The goal is for best effort classes to be used for non-essential traffic, because we provide no assurances about best effort performance. However, in a typical network setting, much best effort traffic will indeed be transmitted and with an adequate degree of expediency.

Because we do not provide any delay assurances for best effort traffic, we do not enforce latency by dropping best effort traffic. Furthermore, because we assume that strict priority traffic is carefully controlled before entering the MVTX2601, we do not enforce a fair bandwidth partition by dropping strict priority traffic. To summarize, dropping to enforce bandwidth or delay does not apply to strict priority or best effort queues. We only drop frames from best effort and strict priority queues when global buffer resources become scarce.

7.5 Weighted Fair Queuing

In some environments – for example, in an environment in which delay assurances are not required, but precise bandwidth partitioning on small time scales is essential, WFQ may be preferable to a delay-bounded scheduling discipline. The MVTX2601 provides the user with a WFQ option with the understanding that delay assurances can not be provided if the incoming traffic pattern is uncontrolled. The user sets four WFQ “weights” such that all weights are whole numbers and sum to 64. This provides per-class bandwidth partitioning with error within 2%.

In WFQ mode, though we do not assure frame latency, the MVTX2601 still retains a set of dropping rules that helps to prevent congestion and trigger higher level protocol end-to-end flow control.

As before, when strict priority is combined with WFQ, we do not have special dropping rules for the strict priority queues, because the input traffic pattern is assumed to be carefully controlled at a prior stage. However, we do indeed drop frames from SP queues for global buffer management purposes. In addition, queue P0 for a 10/100 port are treated as best effort from a dropping perspective, though they still are assured a percentage of bandwidth from a WFQ scheduling perspective. What this means is that these particular queues are only affected by dropping when the global buffer count becomes low.

7.6 WRED Drop Threshold Management Support

To avoid congestion, the Weighted Random Early Detection (WRED) logic drops packets according to specified parameters. The following table summarizes the behavior of the WRED logic.

In KB (kilobytes)	P3	P2	P1	High Drop	Low Drop
Level 1 $N \geq 120$	$P3 \geq AKB$	$P2 \geq AKB$	$P1 \geq AKB$	X%	0%
Level 2 $N \geq 140$				Y%	Z%
Level 3 $N \geq 160$				100%	100%

Table 7 - WRED Drop Thresholds

P_x is the total byte count, in the priority queue x . The WRED logic has three drop levels, depending on the value of N , which is based on the number of bytes in the priority queues. If delay bound scheduling is used, N equals $P3*16+P2*4+P1$. If using WFQ scheduling, N equals $P3+P2+P1$. Each drop level from one to three has defined

high-drop and low-drop percentages, which indicate the minimum and maximum percentages of the data that can be discarded. The X, Y Z percent can be programmed by the register RDRC0, RDRC1. In Level 3, all packets are dropped if the bytes in each priority queue exceed the threshold. Parameters A, B, C are the byte count thresholds for each priority queue. They can be programmed by the QOS control register (refer to the register group 5). See Programming Qos Registers Application Note for more information.

7.7 Buffer Management

Because the number of FDB slots is a scarce resource and because we want to ensure that one misbehaving source port or class cannot harm the performance of a well-behaved source port or class, we introduce the concept of buffer management into the MVTX2601. Our buffer management scheme is designed to divide the total buffer space into numerous reserved regions and one shared pool as shown in Figure 9 on page 26.

As shown in the figure, the FDB pool is divided into several parts. A reserved region for temporary frames stores frames prior to receiving a switch response. Such a temporary region is necessary, because when the frame first enters the MVTX2601, its destination port and class are as yet unknown and so the decision to drop or not needs to be temporarily postponed. This ensures that every frame can be received first before subjecting them to the frame drop discipline after classifying.

Six reserved sections, one for each of the first six priority classes, ensure a programmable number of FDB slots per class. The lowest two classes do not receive any buffer reservation. Furthermore, even for 10/100 Mbps ports, a frame is stored in the region of the FDB corresponding to its class. As we have indicated, the eight classes use only four transmission scheduling queues for 10/100 Mbps ports, but as far as buffer usage is concerned, there are still eight distinguishable classes.

Another segment of the FDB reserves space for each of the 24 ports. One parameter can be set for the source port reservation for 10/100 Mbps. These 24 reserved regions make sure that no well-behaved source port can be blocked by another misbehaving source port.

In addition, there is a shared pool which can store any type of frame. The frame engine allocates the frames first in the six priority sections. When the priority section is full or the packet has priority 1 or 0, the frame is allocated in the shared pool. Once the shared pool is full the frames are allocated in the section reserved for the source port.

The following registers define the size of each section of the frame data buffer:

PR100 - Port Reservation for 10/100 Ports

SFCB - Share FCB Size

C2RS - Class 2 Reserve Size

C3RS - Class 3 Reserve Size

C4RS - Class 4 Reserve Size

C5RS - Class 5 Reserve Size

C6RS - Class 6 Reserve Size

C7RS- Class 7 Reserve Size




Figure 9 - Buffer Partition Scheme Used to Implement MVTX2601 Buffer Management

7.7.1 Dropping When Buffers Are Scarce

Summarizing the two examples of local dropping discussed earlier in this chapter:

- If a queue is a delay-bounded queue, we have a multi level WRED drop scheme, designed to control delay and partition bandwidth in case of congestion
- If a queue is a WFQ-scheduled queue, we have a multi level WRED drop scheme, designed to prevent congestion

In addition to these reasons for dropping, we also drop frames when global buffer space becomes scarce. The function of buffer management is to make sure that such dropping causes as little blocking as possible.

7.8 MVTX2601 Flow Control Basics

Because frame loss is unacceptable for some applications, the MVTX2601 provides a flow control option. When flow control is enabled, scarcity of buffer space in the switch may trigger a flow control signal; this signal tells a source port that is sending a packet to this switch to temporarily hold off.

While flow control offers the clear benefit of no packet loss, it also introduces a problem for quality of service. When a source port receives an Ethernet flow control signal, all microflows originating at that port, well-behaved or not are halted. A single packet destined for a congested output can block other packets destined for uncongested outputs. The resulting head-of-line blocking phenomenon means that quality of service cannot be assured with high confidence when flow control is enabled.

In the MVTX2601, each source port can independently have flow control enabled or disabled. For flow control enabled ports, by default all frames are treated as lowest priority during transmission scheduling. This is done so that those frames are not exposed to the WRED Dropping scheme. Frames from flow control enabled ports feed to only one queue at the destination, the queue of lowest priority. What this means is that if flow control is enabled for a given source port, then we can guarantee that no packets originating from that port will be lost, but at the possible

expense of minimum bandwidth or maximum delay assurances. In addition, these “downgraded” frames may only use the shared pool or the per-source reserved pool in the FDB; frames from flow control enabled sources may not use reserved FDB slots for the highest six classes (P2-P7).

The MVTX2601 does provide a system-wide option of permitting normal QoS scheduling (and buffer use) for frames originating from flow control enabled ports. When this programmable option is active, it is possible that some packets may be dropped even though flow control is on. The reason is that intelligent packet dropping is a major component of the MVTX2601’s approach to ensuring bounded delay and minimum bandwidth for high priority flows.

7.8.1 Unicast Flow Control

For unicast frames, flow control is triggered by source port resource availability. Recall that the MVTX2601’s buffer management scheme allocates a reserved number of FDB slots for each source port. If a programmed number of a source port’s reserved FDB slots have been used then flow control Xoff is triggered.

Xon is triggered when a port is currently being flow controlled and all of that port’s reserved FDB slots have been released.

Note that the MVTX2601’s per-source-port FDB reservations assure that a source port that sends a single frame to a congested destination will not be flow controlled.

7.8.2 Multicast Flow Control

In unmanaged mode, flow control for multicast frames is triggered by a global buffer counter. When the system exceeds a programmable threshold of multicast packets, Xoff is triggered. Xon is triggered when the system returns below this threshold.

In addition, each source port has a 23-bit port map recording which port or ports of the multicast frame’s fanout were congested at the time Xoff was triggered. All ports are continuously monitored for congestion and a port is identified as uncongested when its queue occupancy falls below a fixed threshold. When all those ports that were originally marked as congested in the port map have become uncongested, then Xon is triggered and the 23-bit vector is reset to zero.

7.9 Mapping to IETF Diffserv Classes

For 10/100 Mbps ports, the classes of Table 6 are merged in pairs—one class corresponding to NM+EF, two AF classes and a single BE class.

VTX	P3	P2	P1	P0
IETF	NM+EF	AF0	AF1	BE0

Table 8 - Mapping between MVTX2601 and IETF Diffserv Classes for 10/100 Ports

Features of the MVTX2601 that correspond to the requirements of their associated IETF classes are summarized in the table below.

Network management (NM) and Expedited forwarding (EF)	<ul style="list-style-type: none"> • Global buffer reservation for NM and EF • Option of strict priority scheduling • No dropping if admission controlled
Assured forwarding (AF)	<ul style="list-style-type: none"> • Programmable bandwidth partition, with option of WFQ service • Option of delay-bounded service keeps delay under fixed levels even if not admission-controlled • Random early discard, with programmable levels • Global buffer reservation for each AF class
Best effort (BE)	<ul style="list-style-type: none"> • Service only when other queues are idle means that QoS not adversely affected • Random early discard, with programmable levels • Traffic from flow control enabled ports automatically classified as BE

Table 9 - MVTX2601 Features Enabling IETF Diffserv Standards

8.0 Port Trunking

8.1 Features and Restrictions

A port group (i.e., trunk) can include up to 4 physical ports but all of the ports in a group must be in the same MVTX2601.

Load distribution among the ports in a trunk for unicast is performed using hashing based on source MAC address and destination MAC address. Three other options include source MAC address only, destination MAC address only and source port (in bidirectional ring mode only). Load distribution for multicast is performed similarly.

The MVTX2601 also provides a safe fail-over mode for port trunking automatically. If one of the ports in the trunking group goes down, the MVTX2601 will automatically redistribute the traffic over to the remaining ports in the trunk.

8.2 Unicast Packet Forwarding

The search engine finds the destination MCT entry and if the status field says that the destination port found belongs to a trunk then the group number is retrieved instead of the port number. In addition, if the source address belongs to a trunk then the source port's trunk membership register is checked.

A hash key, based on some combination of the source and destination MAC addresses for the current packet, selects the appropriate forwarding port.

8.3 Multicast Packet Forwarding

For multicast packet forwarding, the device must determine the proper set of ports from which to transmit the packet based on the VLAN index and hash key.

Two functions are required in order to distribute multicast packets to the appropriate destination ports in a port trunking environment.

Determining one forwarding port per group. For multicast packets, all but one port per group, the forwarding port, must be excluded.

Preventing the multicast packet from looping back to the source trunk.

The search engine needs to prevent a multicast packet from sending to a port that is in the same trunk group with the source port. This is because when we select the primary forwarding port for each group, we do not take the source port into account. To prevent this, we simply apply one additional filter, so as to block that forwarding port for this multicast packet.

8.4 Trunking

2 trunk groups are supported. Groups 0 and 1 can trunk up to 4 10/100 ports. The supported combinations are shown in the following table.

Group 0	Port 0	Port 1	Port 2	Port 3
	✓	✓		
	✓	✓	✓	
	✓	✓	✓	✓

Select via trunk0_mode register

Group 1	Port 4	Port 5	Port 6	Port 7
	✓	✓		
	✓	✓	✓	✓

Select via trunk1_mode register

The trunks are individually enabled/disabled by controlling pin trunk 0,1.

9.0 Port Mirroring

9.1 Port Mirroring Features

The received or transmitted data of any 10/100 port in the MVTX2601 chip can be “mirrored” to any other port. We support two such mirrored source-destination pairs. A mirror port cannot also serve as a data port. Please refer to the Port Mirroring Application Note for further details.

9.2 Setting Registers for Port Mirroring

- **MIRROR1_SRC**: Sets the source port for the first port mirroring pair. Bits [4:0] select the source port to be mirrored. An illegal port number is used to disable mirroring (which is the default setting). Bit [5] is used to select between ingress (Rx) or egress (Tx) data.
- **MIRROR1_DEST**: Sets the destination port for the first port mirroring pair. Bits [4:0] select the destination port to be mirrored. The default is port 23.
- **MIRROR2_SRC**: Sets the source port for the second port mirroring pair. Bits [4:0] select the source port to be mirrored. An illegal port number is used to disable mirroring (which is the default setting). Bit [5] is used to select between ingress (Rx) or egress (Tx) data.
- **MIRROR2_DEST**: Sets the destination port for the second port mirroring pair. Bits [4:0] select the destination port to be mirrored. The default is port 0.

10.0 GPSI (7WS) Interface

10.1 GPSI connection

The 10/100 RMII ethernet port can function in GPSI (7WS) mode when the corresponding TXEN pin is strapped low with a 1 K pull down resistor. In this mode, the TXD[0], TXD[1], RXD[0] and RXD[1] serve as TX data, TX clock, RX data and RX clock respectively. The link status and collision from the PHY are multiplexed and shifted into the switch device through external glue logic. The duplex of the port can be controlled by programming the ECR register.

The GPSI interface can be operated in port based VLAN mode only.

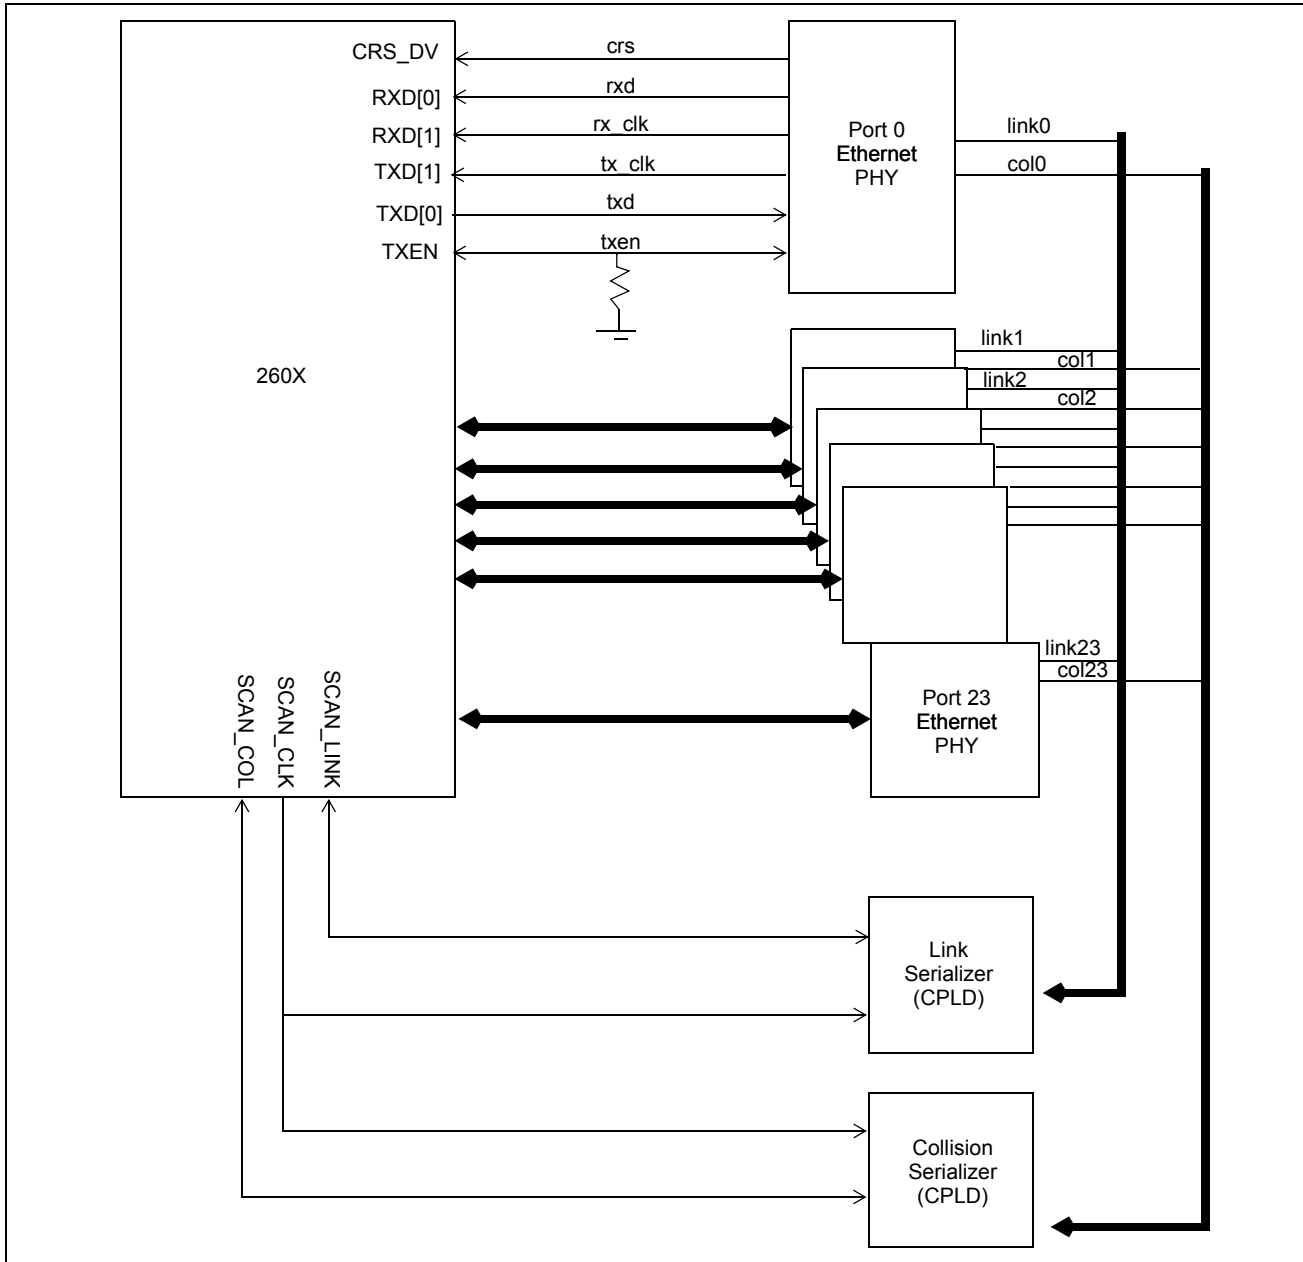


Figure 10 - GPSI (7WS) Mode Connection Diagram

10.2 SCAN LINK and SCAN COL interface

An external CPLD logic is required to take the link signals and collision signals from the GPSI PHYs and shift them into the switch device. The switch device will drive out a signature to indicate the start of the sequence. After that, the CPLD should shift in the link and collision status of the PHYs as shown in the figure. The extra link status indicates the polarity of the link signal. One indicates the polarity of the link signal is active high.

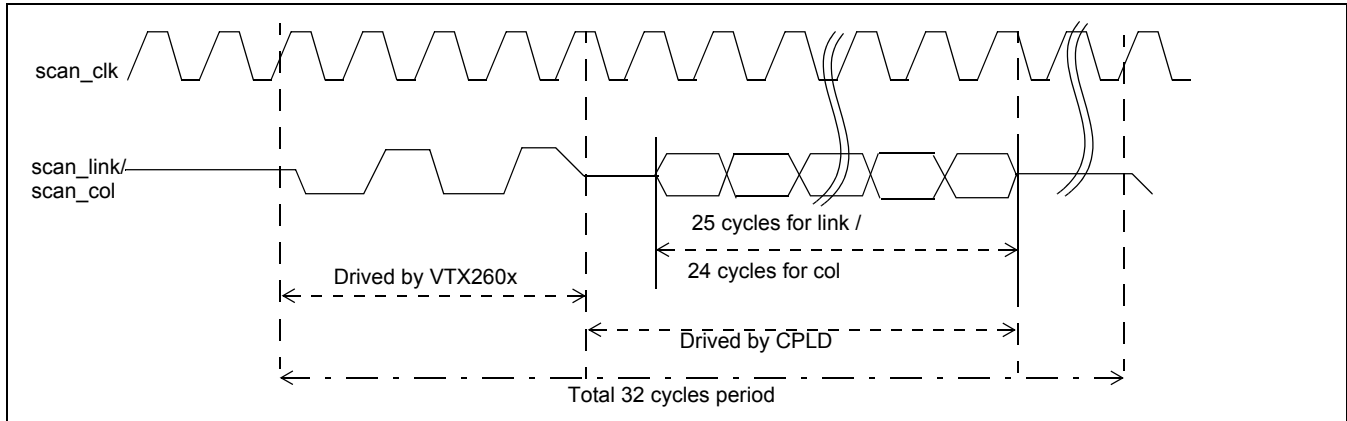


Figure 11 - SCAN LINK and SCAN COLLISION Status Diagram

11.0 LED Interface

11.1 LED Interface Introduction

A serial output channel provides port status information from the MVTX2601 chips. It requires three additional pins.

- LED_CLK at 12.5 MHz
- LED_SYN a sync pulse that defines the boundary between status frames
- LED_DATA a continuous serial stream of data for all status LEDs that repeats once every frame time

A low cost external device (44 pin PAL) is used to decode the serial data and to drive an LED array for display. This device can be customized for different needs.

11.2 Port Status

In the MVTX2601, each port has 8 status indicators, each represented by a single bit. The 8 LED status indicators are:

- Bit 0: Flow control
- Bit 1: Transmit data
- Bit 2: Receive data
- Bit 3: Activity (where activity includes either transmission or reception of data)
- Bit 4: Link up
- Bit 5: Speed (1= 100 Mb/s; 0= 10 Mb/s)
- Bit 6: Full-duplex
- Bit 7: Collision

Eight clocks are required to cycle through the eight status bits for each port.

When the LED_SYN pulse is asserted, the LED interface will present 256 LED clock cycles with the clock cycles providing information for the following ports.

Port 0 (10/100): cycles #0 to cycle #7

Port 1 (10/100): cycles#8 to cycle #15

Port 2 (10/100): cycle #16 to cycle #23

...

Port 22 (10/100): cycle #176 to cycle #183

Port 23 (10/100): cycle #184 to cycle #191

Reserved: cycle #192 to cycle #199

Reserved: cycle #200 to cycle #207

Byte 26 (additional status): cycle #208 to cycle #215

Byte 27 (additional status): cycle #216 to cycle #223

Cycles #224 to 256 present data with a value of zero.

Byte 26 and byte 27 provides bist status

- 26[0]: Reserved
- 26[1]: Reserved
- 26[2]: initialization done
- 26[3]: initialization start
- 26[4]: checksum ok
- 26[5]: link_init_complete
- 26[6]: bist_fail
- 26[7]: ram_error
- 27[0]: bist_in_process
- 27[1]: bist_done

11.3 LED Interface Timing Diagram

The signal from the MVTX2601 to the LED decoder is shown in Figure 12.




Figure 12 - Timing Diagram of LED Interface

12.0 Register Definition

12.1 MVTX2601 Register Description

Register	Description	CPU Addr (Hex)	R/W	I ² C Addr (Hex)	Default	Notes
ETHERNET Port Control Registers Substitute [N] with Port number (0..17h)						
ECR1P"N"	Port Control Register 1 for Port N	000 + 2 x N	R/W	000-018	020	
ECR2P"N"	Port Control Register 2 for Port N	001 + 2 x N	R/W	01B-033	000	
VLAN Control Registers Substitute [N] with Port number (0..17h)						
AVTCL	VLAN Type Code Register Low	100	R/W	036	000	
AVTCH	VLAN Type Code Register High	101	R/W	037	081	
PVMAP"N"_0	Port "N" Configuration Register 0	102 + 4N	R/W	038-050	0FF	
PVMAP"N"_1	Port "N" Configuration Register 1	103 + 4N	R/W	053-06B	0FF	
PVMAP"N"_2	Port "N" Configuration Register 2	104 + 4N	R/W	06E-086	0FF	
PVMAP"N"_3	Port "N" Configuration Register 3	105 + 4N	R/W	089-0A1	007	
PVMODE	VLAN Operating Mode	170	R/W	0A4	000	
TRUNK Control Registers						
TRUNK0_MODE	Trunk Group 0 Mode	203	R/W	0A5	003	
TRUNK1_MODE	Trunk Group 1 Mode	20B	R/W	0A6	003	
Search Engine Configurations						
TX_AGE	Transmission Queue Aging Time	325	R/W	0A7	008	
AGETIME_LOW	MAC Address Aging Time Low	400	R/W	0A8	2M:05C / 4M:02E	
AGETIME_HIGH	MAC Address Aging Time High	401	R/W	0A9	000	
SE_OPMODE	Search Engine Operating Mode	403	R/W	NA	000	
Buffer Control and QOS Control						
FCBAT	FCB Aging Timer	500	R/W	0AA	0FF	
QOSC	QOS Control	501	R/W	0AB	000	
FCR	Flooding Control Register	502	R/W	0AC	008	
AVPML	VLAN Priority Map Low	503	R/W	0AD	000	
AVPMM	VLAN Priority Map Middle	504	R/W	0AE	000	
AVPMH	VLAN Priority Map High	505	R/W	0AF	000	
TOSPML	TOS Priority Map Low	506	R/W	0B0	000	

Register	Description	CPU Addr (Hex)	R/W	I ² C Addr (Hex)	Default	Notes
TOSPMML	TOS Priority Map Middle	507	R/W	0B1	000	
TOSPMH	TOS Priority Map High	508	R/W	0B2	000	
AVDM	VLAN Discard Map	509	R/W	0B3	000	
TOSDML	TOS Discard Map	50A	R/W	0B4	000	
BMRC	Broadcast/Multicast Rate Control	50B	R/W	0B5	000	
UCC	Unicast Congestion Control	50C	R/W	0B6	1M:008 / 2M:010	
MCC	Multicast Congestion Control	50D	R/W	0B7	050	
PR100	Port Reservation for 10/100 Ports	50E	R/W	0B8	1M:035 / 2M:058	
SFCB	Share FCB Size	510	R/W	0BA	1M:046 / 2M:0E6	
C2RS	Class 2 Reserve Size	511	R/W	0BB	000	
C3RS	Class 3 Reserve Size	512	R/W	0BC	000	
C4RS	Class 4 Reserve Size	513	R/W	0BD	000	
C5RS	Class 5 Reserve Size	514	R/W	0BE	000	
C6RS	Class 6 Reserve Size	515	R/W	0BF	000	
C7RS	Class 7 Reserve Size	516	R/W	0C0	000	
QOSC"N"	QOS Control (N=0-39)	517-53E	R/W	0C1-0D2	000	
RDRC0	WRED Drop Rate Control 0	553	R/W	0FB	08F	
RDRC1	WRED Drop Rate Control 1	554	R/W	0FC	088	
USER_PORT"N"_LOW	User Define Logical Port "N" Low (N=0-7)	580 + 2N	R/W	0D6-0DD	000	
USER_PORT"N"_HIGH	User Define Logical Port "N" High	581 + 2N	R/W	0DE-0E5	000	
USER_PORT1:0_PRIORITY	User Define Logic Port 1 and 0 Priority	590	R/W	0E6	000	
USER_PORT3:2_PRIORITY	User Define Logic Port 3 and 2 Priority	591	R/W	0E7	000	
USER_PORT5:4_PRIORITY	User Define Logic Port 5 and 4 Priority	592	R/W	0E8	000	

Register	Description	CPU Addr (Hex)	R/W	I ² C Addr (Hex)	Default	Notes
USER_PORT7:6_PRIORITY	User Define Logic Port 7 and 6 Priority	593	R/W	0E9	000	
USER_PORT_ENABLE	User Define Logic Port Enable	594	R/W	0EA	000	
WLPP10	Well known Logic Port Priority for 1 and 0	595	R/W	0EB	000	
WLPP32	Well known Logic Port Priority for 3 and 2	596	R/W	0EC	000	
WLPP54	Well known Logic Port Priority for 5 and 4	597	R/W	0ED	000	
WLPP76	Well-known Logic Port Priority for 7 & 6	598	R/W	0EE	000	
WLPE	Well known Logic Port Enable	599	R/W	0EF	000	
RLOWL	User Define Range Low Bit7:0	59A	R/W	0F4	000	
RLOWH	User Define Range Low Bit 15:8	59B	R/W	0F5	000	
RHIGHL	User Define Range High Bit 7:0	59C	R/W	0D3	000	
RHIGHH	User Define Range High Bit 15:8	59D	R/W	0D4	000	
RPRIORITY	User Define Range Priority	59E	R/W	0D5	000	
MISC Configuration Registers						
MII_OP0	MII Register Option 0	600	R/W	0F0	000	
MII_OP1	MII Register Option 1	601	R/W	0F1	000	
FEN	Feature Registers	602	R/W	0F2	010	
MIIC0	MII Command Register 0	603	R/W	N/A	000	
MIIC1	MII Command Register 1	604	R/W	N/A	000	
MIIC2	MII Command Register 2	605	R/W	N/A	000	
MIIC3	MII Command Register 3	606	R/W	N/A	000	
MIID0	MII Data Register 0	607	RO	N/A	N/A	
MIID1	MII Data Register 1	608	RO	N/A	N/A	
LED	LED Control Register	609	R/W	0F3	000	
SUM	EEPROM Checksum Register	60B	R/W	0FF	000	
Port Mirroring Controls						
MIRROR1_SRC	Port Mirror 1 Source Port	700	R/W	N/A	07F	

Register	Description	CPU Addr (Hex)	R/W	I ² C Addr (Hex)	Default	Notes
MIRROR1_DEST	Port Mirror 1 Destination Port	701	R/W	N/A	017	
MIRROR2_SRC	Port Mirror 2 Source Port	702	R/W	N/A	OFF	
MIRROR2_DEST	Port Mirror 2 Destination Port	703	R/W	N/A	000	
Device Configuration Register						
GCR	Global Control Register	F00	R/W	N/A	000	
DCR	Device Status and Signature Register	F01	RO	N/A	N/A	
DCR1	Chip status	F02	RO	N/A	N/A	
DPST	Device Port Status Register	F03	R/W	N/A	000	
DTST	Data read back register	F04	RO	N/A	N/A	
DA	DA Register	FFF	RO	N/A	DA	

12.2 Group 0 Address MAC Ports Group

12.2.1 ECR1Pn: Port N Control Register

- I²C Address 000-018; CPU Address:0000+2xN (N = port number)
- Accessed by serial interface and I²C (R/W)

7	6	5	4	3	2	1	0
Sp State		A-FC	Port Mode				

Bit [0] 1 - Flow Control Off
 0 - Flow Control On

- When Flow Control On:
 - In half duplex mode, the MAC transmitter applies back pressure for flow control

In full duplex mode, the MAC transmitter sends Flow Control frames when necessary. The MAC receiver interprets and processes incoming flow control frames. The Flow Control Frame Received counter is incremented whenever a flow control is received.

- When Flow Control off:
 - In half duplex mode, the MAC Transmitter does not assert flow control by sending flow control frames or jamming collision.
 - In full duplex mode, the Mac transmitter does not send flow control frames. The MAC receiver does not interpret or process the flow control frames. The Flow Control Frame Received counter is not incremented.

Bit [1] 1 - Half Duplex - Only 10/100 mode
 0 - Full Duplex

- Bit [2] 1 - 10 Mbps
 0 - 100 Mbps
- Bit [4:3] • 00 - Automatic Enable Auto Neg. This enables hardware state machine for auto-negotiation.
 • 01 - Limited Disable auto Neg. This disables hardware for speed auto-negotiation. Poll MII for link status.
 • 10 - Link Down. Disable auto Neg. state machine and force link down (disable the port)
 • 11 - Link Up. User ERC1 [2:0] for config.
- Bit [5] • Asymmetric Flow Control Enable
 • 0 - Disable asymmetric flow control
 • 1 - Enable asymmetric flow control
 • Asymmetric Flow Control Enable. When this bit is set and flow control is on (bit [0] = 0, don't send out a flow control frame. But MAC receiver interprets and process flow control frames. Default is 0
- Bit [7:6] SS - Spanning tree state Default is 11
 00 – Blocking: Frame is dropped
 01 - Listening: Frame is dropped
 10 - Learning: Frame is dropped. Source MAC address is learned.
 11 - Forwarding: Frame is forwarded. Source MAC address is learned.

12.2.2 ECR2Pn: Port N Control Register

- I²C Address: 01B-035; CPU Address:0001+2xN
- Accessed by and serial interface and I²C (R/W)

7	6	5	4	3	2	1	0
	QoS Sel	Reserve	DisL	Ftf	Futf		

- Bit [0]: • Filter untagged frame (**Default 0**)
 • 0: Disable
 • 1: All untagged frames from this port are discarded
- Bit [1]: • Filter Tag frame (**Default 0**)
 • 0: Disable
 • 1: All tagged frames from this port are discarded
- Bit [2]: • Learning Disable (**Default 0**)
 • 1 Learning is disabled on this port
 • 0 Learning is enabled on this port
- Bit [3]: • Must be set to '1'

- Bit [5:4:]
- QoS mode selection (Default 00)
 - Determines which of the 4 sets of QoS settings is used for 10/100 ports.
 - Note that there are 4 sets of per-queue byte thresholds, and 4 sets of WFQ ratios programmed. These bits select among the 4 choices for each 10/100 port. Refer to QoS Application Note.
 - 00: select class byte limit set 0 and classes WFQ credit set 0
 - 01: select class byte limit set 1 and classes WFQ credit set 1
 - 10: select class byte limit set 2 and classes WFQ credit set 2
 - 11: select class byte limit set 3 and classes WFQ credit set 3
- Bit [7:6] Reserved

12.3 Group 1 Address VLAN Group

12.3.1 AVTCL – VLAN Type Code Register Low

- I²C Address 036; CPU Address:h100
- Accessed by serial interface and I²C (R/W)

Bit [7:0]:

- LANType_LOW: Lower 8 bits of the VLAN type code (Default 00)

12.3.2 AVTCH – VLAN Type Code Register High

- I²C Address 037; CPU Address:h101
- Accessed by serial interface and I²C (R/W)

Bit [7:0]:

- VLANType_HIGH: Upper 8 bits of the VLAN type code (Default is 81)

12.3.3 PVMAP00_0 – Port 00 Configuration Register 0

- I²C Address 038, CPU Address:h102)
- Accessed by serial interface and I²C (R/W)

Bit [7:0]:

- VLAN Mask for ports 7 to 0 (Default FF)

This register indicates the legal egress ports. A “1” on bit 7 means that the packet can be sent to port 7. A “0” on bit 7 means that any packet destined to port 7 will be discarded. This register works with registers 1 and 2 to form a 24-bit mask to all egress ports.

12.3.4 PVMAP00_1 – Port 00 Configuration Register 1

- I²C Address h53, CPU Address:h103
- Accessed by serial interface and I²C (R/W)

Bit [7:0]:

- VLAN Mask for ports 15 to 8 (Default is FF)

12.3.5 PVMAP00_2 – Port 00 Configuration Register 2

- I²C Address h6E, CPU Address:h104
- Accessed by serial interface and I²C (R/W)

Bit [7:0]: • VLAN Mask for ports 23 to 16 (Default FF)

12.3.6 PVMAP00_3 – Port 00 Configuration Register 3

- I²C Address h89, CPU Address:h105)
- Accessed by serial interface and I²C (R/W)

7	6	5	3	2	0
FP en	Drop	Default tx priority			

Bit [2:0]: Reserved (Default 7)

Bit [5:3]: Default Transmit priority. Used when Bit [7] = 1 (**Default 0**)

- 000 Transmit Priority Level 0 (Lowest)
- 001 Transmit Priority Level 1
- 010 Transmit Priority Level 2
- 011 Transmit Priority Level 3
- 100 Transmit Priority Level 4
- 101 Transmit Priority Level 5
- 110 Transmit Priority Level 6
- 111 Transmit Priority Level 7 (Highest)

Bit [6]: Default Discard priority (**Default 0**)

- 0 – Discard Priority Level 0 (Lowest)
- 1 – Discard Priority Level 7(Highest)

Bit [7]: Enable Fix Priority (**Default 0**)

- 0 Disable fix priority. All frames are analyzed. Transmit Priority and Discard Priority are based on VLAN Tag, TOS field or Logical Port.
- 1 Transmit Priority and Discard Priority are based on values programmed in bit [6:3]

12.4 Port Configuration Register

- **PVMAP01_0,1,2,3** I²C Address h39,54,6F,8A; CPU Address:h106,107,108,109)
- **PVMAP02_0,1,2,3** I²C Address h3A,55,70,8B; CPU Address:h10A, 10B, 10C, 10D)
- **PVMAP03_0,1,2,3** I²C Address h3B,56,71,8C; CPU Address:h10E, 10F, 110, 111)
- **PVMAP04_0,1,2,3** I²C Address h3C,57,72,8D; CPU Address:h112, 113, 114, 115)
- **PVMAP05_0,1,2,3** I²C Address h3D,58,73,8E; CPU Address:h116, 117, 118, 119)
- **PVMAP06_0,1,2,3** I²C Address h3E,59,74,8F; CPU Address:h11A, 11B, 11C, 11D)
- **PVMAP07_0,1,2,3** I²C Address h3F,5A,75,90; CPU Address:h11E, 11F, 120, 121)
- **PVMAP08_0,1,2,3** I²C Address h40,5B,76,91; CPU Address:h122, 123, 124, 125)
- **PVMAP09_0,1,2,3** I²C Address h41,5C,77,92; CPU Address:h126, 127, 128, 129)
- **PVMAP10_0,1,2,3** I²C Address h42,5D,78,93; CPU Address:h12A, 12B, 12C, 12D)
- **PVMAP11_0,1,2,3** I²C Address h43,5E,79,94; CPU Address:h12E, 12F, 130, 131)
- **PVMAP12_0,1,2,3** I²C Address h44,5F,7A,95; CPU Address:h132, 133, 134, 135)

- **PVMAP13_0,1,2,3** I²C Address h45,60,7B,96; CPU Address:h136, 137, 138, 139)
- **PVMAP14_0,1,2,3** I²C Address h46,61,7C,97; CPU Address:h13A, h13B, 13C, 13D)
- **PVMAP15_0,1,2,3** I²C Address h47,62,7D,98; CPU Address:h13E, 13F, 140, 141)
- **PVMAP16_0,1,2,3** I²C Address h48,63,7E,99; CPU Address:h142, 143, 144, 145)
- **PVMAP17_0,1,2,3** I²C Address h49,64,7F,9A; CPU Address:h146, 147, 148, 149)
- **PVMAP18_0,1,2,3** I²C Address h4A,65,80,9B; CPU Address:h14A, 14B, 14C, 14D)
- **PVMAP19_0,1,2,3** I²C Address h4B,66,81,9C; CPU Address:h14E, 14F, 150, 151)
- **PVMAP20_0,1,2,3** I²C Address h4C,67,82,9D; CPU Address:h152, 153, 154, 155)
- **PVMAP21_0,1,2,3** I²C Address h4D,68,83,9E; CPU Address:h156, 157, 158, 159)
- **PVMAP22_0,1,2,3** I²C Address h4E,69,84,9F; CPU Address:h15A, 15B, 15C, 15D)
- **PVMAP23_0,1,2,3** I²C Address h4F,6A,85,A0; CPU Address:h15E, 15F, 160, 161)

12.4.1 PVMODE

- I²C Address: h0A4, CPU Address:h170
- Accessed by serial interface, and I²C (R/W)

7	5	4	3	2	1	0
	SM0		DF	SL		

- Bit [0]:
- Reserved
 - Must be '0'
- Bit [1]:
- Slow learning
 - Same function as SE_OP MODE bit 7. Either bit can enable the function; both need to be turned off to disable the feature.
- Bit [2]:
- Disable dropping of frames with destination MAC addresses 0180C2000001 to 0180C200000F (Default = 0)
 - 0: Drop all frames in this range
 - 1: Disable dropping of frames in this range
- Bit [3]:
- Reserved
- Bit [4]:
- Support MAC address 0
 - 0: MAC address 0 is not learned.
 - 1: MAC address 0 is learned.
- Bit [7:5]:
- Reserved

12.5 Group 2 Address Port Trunking Group

12.5.1 TRUNK0_MODE– Trunk group 0 mode


- I²C Address h0A5; CPU Address:203
- Accessed by serial interface and I²C (R/W)

7	4	3	2	1	0
		Hash Select		Port Select	

- Bit [1:0]:
- Port selection in unmanaged mode. Input pin TRUNK0 enable/disable trunk group 0.
 - 00 Reserved
 - 01 Port 0 and 1 are used for trunk0
 - 10 Port 0,1 and 2 are used for trunk0
 - 11 Port 0,1,2 and 3 are used for trunk0
- Bit [3:2]
- Hash Select. The Hash selected is valid for Trunk 0, 1 and 2. (Default 00)
 - 00 Use Source and Destination Mac Address for hashing
 - 01 Use Source Mac Address for hashing
 - 10 Use Destination Mac Address for hashing
 - 11 Use source destination MAC address and ingress physical port number for hashing

12.5.2 TRUNK1_MODE – Trunk group 1 mode

- I²C Address h0A6; CPU Address:20B
- Accessed by serial interface and I²C (R/W)




- Bit [1:0]:
- Port selection in unmanaged mode. Input pin TRUNK1 enable/disable trunk group 1.
 - 00 Reserved
 - 01 Port 4 and 5 are used for trunk1
 - 10 Reserved
 - 11 Port 4, 5, 6 and 7 are used for trunk1

12.6 Group 4 Address Search Engine Group

12.6.1 TX_AGE – Tx Queue Aging timer

- I²C Address: h07;CPU Address:h325
- Accessed by serial interface (RW)



- Bit [5:0]: Unit of 100 ms (Default 8)
- Disable transmission queue aging if value is zero. Aging timer for all ports and queues.
- For no packet loss flow control, this register must be set to 0

12.6.2 AGETIME_LOW – MAC address aging time Low

- I²C Address h0A8; CPU Address:h400
- Accessed by serial interface and I²C (R/W)
- Bit [7:0] Low byte of the MAC address aging timer
- MAC address aging is enable/disable by boot strap TSTOUT9

12.6.3 AGETIME_HIGH –MAC address aging time High

- I²C Address h0A9; CPU Address h401
- Accessed by serial interface and I²C (R/W)
- Bit [7:0]: High byte of the MAC address aging timer
- The default setting provide 300 seconds aging time. Aging time is based on the following equation:
- {AGETIME_TIME,AGETIME_LOW} X (# of MAC address entries in the memory x 100 μsec). Number of MAC entries = 32 K when 1 MB is used. Number of MAC entries = 64 K when 2 MB is used.

12.6.4 SE_OPMODE – Search Engine Operation Mode

- CPU Address:h403
- Accessed by serial interface (R/W)
- {SE_OPMODE} X(# of entries 100 usec)


7	6	5	0
SL	DMS		

- Bit [5:0]:
- Reserved
- Bit [6]:
- Disable MCT speedup aging
 - 1 – Disable speedup aging when MCT resource is low
 - 0 – Enable speedup aging when MCT resource is low
- Bit [7]:
- Slow Learning
 - 1– Enable slow learning. Learning is temporary disabled when search demand is high
 - 0 – Learning is performed independent of search demand

12.7 Group 5 Address Buffer Control/QOS Group

12.7.1 FCBAT – FCB Aging Timer


- I²C Address h0AA; CPU Address:h500



- Bit [7:0]:
- FCB Aging time. Unit of 1ms. (**Default FF**)
 - This function is for buffer aging control. It is used to configure the aging time, and can be enabled/ disabled through bootstrap pin. It is not recommended to use this function for normal operation.

12.7.2 QOSC – QOS Control


- I²C Address h0AB; CPU Address:h501
- Accessed by serial interface and I²C (R/W)



- Bit [0]:
- QoS frame lost is OK. Priority will be available for flow control enabled source only when this bit is set (**Default 0**)
- Bit [4]:
- Per VLAN Multicast Flow Control (**Default 0**)
 - 0 – Disable
 - 1 – Enable
- Bit [5]:
- Reserved
- Bit [6]:
- Select TOS bits for Priority (**Default 0**)
 - 0 - Use TOS [4:2] bits to map the transmit priority
 - 1 - Use TOS [7:5] bits to map the transmit priority
- Bit [7]:
- Select TOS bits for Drop Priority (**Default 0**)
 - 0 - Use TOS[4:2] bits to map the drop priority
 - 1 - Use TOS[7:5] bits to map the drop priority

12.7.3 FCR – Flooding Control Register


- I²C Address h0AC; CPU Address:h502
- Accessed by serial interface and I²C (R/W)



- Bit [3:0]:
- U2MR: Unicast to Multicast Rate. Units in terms of time base defined in bits [6:4]. This is used to limit the amount of flooding traffic. The value in U2MR specifies how many packets are allowed to flood within the time specified by bit [6:4]. To disable this function, program U2MR to 0. (**Default = 8**)
- Bit [6:4]:
- TimeBase:
 - 000 = 100 us**
 - 001 = 200 us**
 - 010 = 400 us**
 - 011 = 800 us**
 - 100 = 1.6 ms**
 - 101 = 3.2 ms**
 - 110 = 6.4 ms**
 - 111 = 100 us (same as 000)**
 - (Default = 000)
- Bit [7]:
- Select VLAN tag or TOS (IP packets) to be preferentially picked to map transmit priority and drop priority (**Default = 0**).
 - 0 – Select VLAN Tag priority field over TOS
 - 1 – Select TOS over VLAN tag priority field

12.7.4 AVPML – VLAN Priority Map

- I²C Address h0AD; CPU Address:h503
- Accessed by serial interface and I²C (R/W)




Registers AVPML, AVPMM, and AVPMH allow the eight VLAN priorities to map into eight internal level transmit priorities. Under the internal transmit priority, seven is highest priority where as zero is the lowest. This feature allows the user the flexibility of redefining the VLAN priority field. For example, programming a value of 7 into bit 2:0 of the AVPML register would map VLAN priority 0 into internal transmit priority 7. The new priority is used inside the 2601. When the packet goes out it carries the original priority.

- Bit [2:0]:
- Priority when the VLAN tag priority field is 0 (**Default 0**)
- Bit [5:3]:
- Priority when the VLAN tag priority field is 1 (**Default 0**)
- Bit [7:6]:
- Priority when the VLAN tag priority field is 2 (**Default 0**)

12.7.5 AVPMM – VLAN Priority Map

- I²C Address h0AE, CPU Address:h504
- Accessed by serial interface and I²C (R/W)




Map VLAN priority into eight level transmit priorities:

- Bit [0]: • Priority when the VLAN tag priority field is 2 (**Default 0**)
- Bit [3:1]: • Priority when the VLAN tag priority field is 3 (**Default 0**)
- Bit [6:4]: • Priority when the VLAN tag priority field is 4 (**Default 0**)
- Bit [7]: • Priority when the VLAN tag priority field is 5 (**Default 0**)

12.7.6 AVPMH – VLAN Priority Map

- I²C Address h0AF, CPU Address:h505
- Accessed by serial interface and I²C (R/W)




Map VLAN priority into eight level transmit priorities:

- Bit [1:0]: • Priority when the VLAN tag priority field is 5 (**Default 0**)
- Bit [4:2]: • Priority when the VLAN tag priority field is 6 (**Default 0**)
- Bit [7:5]: • Priority when the VLAN tag priority field is 7 (**Default 0**)

12.7.7 TOSPML – TOS Priority Map

- I²C Address h0B0, CPU Address:h506
- Accessed by serial interface and I²C (R/W)




Map TOS field in IP packet into eight level transmit priorities

- Bit [2:0]: • Priority when the TOS field is 0 (**Default 0**)
- Bit [5:3]: • Priority when the TOS field is 1 (**Default 0**)
- Bit [7:6]: • Priority when the TOS field is 2 (**Default 0**)

12.7.8 TOSPM – TOS Priority Map

- I²C Address h0B1, CPU Address:h507
- Accessed by serial interface and I²C (R/W)




Map TOS field in IP packet into four level transmit priorities

- Bit [0]: • Priority when the TOS field is 2 (**Default 0**)
- Bit [3:1]: • Priority when the TOS field is 3 (**Default 0**)
- Bit [6:4]: • Priority when the TOS field is 4 (**Default 0**)
- Bit [7]: • Priority when the TOS field is 5 (**Default 0**)

12.7.9 TOSPMH – TOS Priority Map

- I²C Address h0B2, CPU Address:h508
- Accessed by serial interface and I²C (R/W)




Map TOS field in IP packet into four level transmit priorities:

- Bit [1:0]: • Priority when the TOS field is 5 (**Default 0**)
- Bit [4:2]: • Priority when the TOS field is 6 (**Default 0**)
- Bit [7:5]: • Priority when the TOS field is 7 (**Default 0**)

12.7.10 AVDM – VLAN Discard Map

- I²C Address h0B3, CPU Address:h509
- Accessed by serial interface and I²C (R/W)



Map VLAN priority into frame discard when low priority buffer usage is above threshold

- Bit [0]: • Frame drop priority when VLAN tag priority field is 0 (**Default 0**)
- Bit [1]: • Frame drop priority when VLAN tag priority field is 1 (**Default 0**)
- Bit [2]: • Frame drop priority when VLAN tag priority field is 2 (**Default 0**)
- Bit [3]: • Frame drop priority when VLAN tag priority field is 3 (**Default 0**)
- Bit [4]: • Frame drop priority when VLAN tag priority field is 4 (**Default 0**)
- Bit [5]: • Frame drop priority when VLAN tag priority field is 5 (**Default 0**)

- Bit [6]: • Frame drop priority when VLAN tag priority field is 6 (**Default 0**)
- Bit [7]: • Frame drop priority when VLAN tag priority field is 7 (**Default 0**)

12.7.11 TOSDML – TOS Discard Map

- I²C Address h0B4, CPU Address:h50A
- Accessed by serial interface and I²C (R/W)

7	6	5	4	3	2	1	0
FDT7	FDT6	FDT5	FDT4	FDT3	FDT2	FDT1	FDT0

Map TOS into frame discard when low priority buffer usage is above threshold

- Bit [0]: • Frame drop priority when TOS field is 0 (**Default 0**)
- Bit [1]: • Frame drop priority when TOS field is 1 (**Default 0**)
- Bit [2]: • Frame drop priority when TOS field is 2 (**Default 0**)
- Bit [3]: • Frame drop priority when TOS field is 3 (**Default 0**)
- Bit [4]: • Frame drop priority when TOS field is 4 (**Default 0**)
- Bit [5]: • Frame drop priority when TOS field is 5 (**Default 0**)
- Bit [6]: • Frame drop priority when TOS field is 6 (**Default 0**)
- Bit [7]: • Frame drop priority when TOS field is 7 (**Default 0**)

12.7.12 BMRC - Broadcast/Multicast Rate Control

- I²C Address h0B5, CPU Address:h50B)
- Accessed by serial interface and I²C (R/W)


7	4	3	0
Broadcast Rate		Multicast Rate	

- This broadcast and multicast rate defines for each port the number of packet allowed to be forwarded within a specified time. Once the packet rate is reached, packets will be dropped. To turn off the rate limit, program the field to 0. Timebase is based on register 502 [6:4].

- Bit [3:0] : • Multicast Rate Control Number of multicast packets allowed within the time defined in bits 6 to 4 of the Flooding Control Register (FCR). (**Default 0**).
- Bit [7:4] : • Broadcast Rate Control Number of broadcast packets allowed within the time defined in bits 6 to 4 of the Flooding Control Register (FCR). (**Default 0**)

12.7.13 UCC – Unicast Congestion Control


- I²C Address h0B6, CPU Address: 50C
- Accessed by serial interface and I²C (R/W)



- Bit [7:0] :
- Number of frame count. Used for best effort dropping at B% when destination port's best effort queue reaches UCC threshold and shared pool is all in use. Granularity 1 frame. (Default: h10 for 2 MB or h08 for 1 MB)

12.7.14 MCC – Multicast Congestion Control


- I²C Address h0B7, CPU Address: 50D
- Accessed by serial interface and I²C (R/W)



- Bit [4:0]:
- In multiples of two. Used for triggering MC flow control when destination multicast port's best effort queue reaches MCC threshold. (Default 0x10)
- Bit [7:5]:
- Flow control reaction period (Default 2) Granularity 4 uSec.

12.7.15 PR100 – Port Reservation for 10/100 ports


- I²C Address h0B8, CPU Address 50E
- Accessed by serial interface and I²C (R/W)



- Bit [3:0]:
- Per port buffer reservation.
 - Define the space in the FDB reserved for each 10/100 port. Expressed in multiples of 4 packets. For each packet 1536 bytes are reserved in the memory.
- Bits [7:4]:
- Expressed in multiples of 4 packets. Threshold for dropping all best effort frames when destination port best efforts queues reach UCC threshold and shared pool all used and source port reservation is at or below the PR100[7:4] level. Also the threshold for initiating UC flow control.
 - Default:
 - h58 for configuration with 2 MB;
 - h35 for configuration with 1 MB;

12.7.16 SFCB – Share FCB Size


- I²C Address h0BA), CPU Address 510
- Accessed by serial interface and I²C (R/W)



- Bits [7:0]:
- Expressed in multiples of 4 packets. Buffer reservation for shared pool.
 - Default:
 - hE6 for configuration with memory of 2 MB;
 - h46 for configuration with memory of 1 MB;

12.7.17 C2RS – Class 2 Reserve Size


- I²C Address h0BB, CPU Address 511
- Accessed by serial interface and I²C (R/W)



- Buffer reservation for class 2 (third lowest priority). Granularity 1. **(Default 0)**

12.7.18 C3RS – Class 3 Reserve Size


- I²C Address h0BC, CPU Address 512
- Accessed by serial interface and I²C (R/W)



- Buffer reservation for class 3. Granularity 1. **(Default 0)**

12.7.19 C4RS – Class 4 Reserve Size


- I²C Address h0BD, CPU Address 513
- Accessed by serial interface and I²C (R/W)



- Buffer reservation for class 4. Granularity 1. **(Default 0)**

12.7.20 C5RS – Class 5 Reserve Size


- I²C Address h0BE; CPU Address 514
- Accessed by serial interface and I²C (R/W)



- Buffer reservation for class 5. Granularity 1. **(Default 0)**

12.7.21 C6RS – Class 6 Reserve Size


- I²C Address h0BF; CPU Address 515
- Accessed by serial interface and I²C (R/W)



- Buffer reservation for class 6 (second highest priority). Granularity 1. **(Default 0)**

12.7.22 C7RS – Class 7 Reserve Size

- I²C Address h0C0; CPU Address 516
- Accessed by serial interface and I²C (R/W)



- Buffer reservation for class 7 (highest priority). Granularity 1. **(Default 0)**

12.7.23 Classes Byte Limit Set 0

- Accessed by serial interface and I²C (R/W):
 - C — QOSC00 – BYTE_C01 (I²C Address h0C1, CPU Address 517)
 - B — QOSC01 – BYTE_C02 (I²C Address h0C2, CPU Address 518)
 - A — QOSC02 – BYTE_C03 (I²C Address h0C3, CPU Address 519)

QOSC00 through QOSC02 represents one set of values A-C for a 10/100 port when using the Weighted Random Early Drop (WRED) Scheme described in Chapter 7.7. There are four such sets of values A-C specified in Classes Byte Limit Set 0, 1, 2 and 3.

Each 10/ 100 port can choose one of the four Byte Limit Sets as specified by the QoS Select field located in bits 5 to 4 of the ECR2n register. The values A-C are per-queue byte thresholds for random early drop. QOSC02 represents A and QOSC00 represents C.

Granularity when Delay bound is used: QOSC02: 128 bytes, QOSC01: 256 bytes. QOSC00: 512 bytes.
Granularity when WFQ is used: QOSC02: 512 bytes, QOSC01: 512 bytes, QOSC00: 512 bytes.

12.7.24 Classes Byte Limit Set 1

- Accessed by serial interface and I²C (R/W):
 - C - QOSC03 – BYTE_C11 (I²C Address h0C4, CPU Address 51a)
 - B - QOSC04 – BYTE_C12 (I²C Address h0C5, CPU Address 51b)
 - A - QOSC05 – BYTE_C13 (I²C Address h0C6, CPU Address 51c)

QOSC03 through QOSC05 represents one set of values A-C for a 10/100 port when using the Weighted Random Early Detect (WRED) Scheme.

Granularity when Delay bound is used: QOSC05: 128 bytes, QOSC04: 256 bytes. QOSC03: 512 bytes.
Granularity when WFQ is used: QOSC05: 512 bytes, QOSC04: 512 bytes, QOSC03: 512 bytes.

12.7.25 Classes Byte Limit Set 2

- Accessed by serial interface and I²C (R/W):
 - C - QOSC06 – BYTE_C21 (CPU Address 51d)
 - B - QOSC07 – BYTE_C22 (CPU Address 51e)
 - A - QOSC08 – BYTE_C23 (CPU Address 51f)

QOSC06 through QOSC08 represents one set of values A-C for a 10/100 port when using the Weighted Random Early Detect (WRED) Scheme.

Granularity when Delay bound is used: QOSC08: 128 bytes, QOSC07: 256 bytes. QOSC06: 512 bytes.
Granularity when WFQ is used: QOSC08: 512 bytes, QOSC07: 512 bytes, QOSC06: 512 bytes.

12.7.26 Classes Byte Limit Set 3

- Accessed by serial interface and I²C (R/W):
 - C - QOSC09 – BYTE_C31 (CPU Address 520)
 - B - QOSC10 – BYTE_C32 (CPU Address 521)
 - A - QOSC11 – BYTE_C33 (CPU Address 522)

QOSC09 through QOSC011 represents one set of values A-C for a 10/100 port when using the Weighted Random Early Detect (WRED) Scheme.

Granularity when Delay bound is used: QOSC11: 128 bytes, QOSC10: 256 bytes. QOSC09: 512 bytes.
Granularity when WFQ is used: QOSC11: 512 bytes, QOSC10: 512 bytes, QOSC09: 512 bytes.

12.7.27 Classes WFQ Credit Set 0

- Accessed by serial interface (R/W)
 - W0 - QOSC24[5:0] – CREDIT_C00 (CPU Address 52f)
 - W1 - QOSC25[5:0] – CREDIT_C01 (CPU Address 530)
 - W2 - QOSC26[5:0] – CREDIT_C02 (CPU Address 531)
 - W3 - QOSC27[5:0] – CREDIT_C03 (CPU Address 532)

QOSC24 through QOSC27 represents one set of WFQ parameters for a 10/100 port. There are four such sets of values. The granularity of the numbers is 1 and their sum must be 64. QOSC27 corresponds to W3 and QOSC24 corresponds to W0.

QOSC24[7:6]: Priority service type for the ports select this parameter set. Option 1 to 4.

QOSC25[7]: Priority service allow flow control for the ports select this parameter set.

QOSC25[6]: Flow control pause best effort traffic only

Both flow control allow and flow control best effort only can take effect only the priority type is WFQ.

12.7.28 Classes WFQ Credit Set 1

- Accessed by serial interface (R/W)
 - W0 - QOSC28[5:0] – CREDIT_C10 (CPU Address 533)
 - W1 - QOSC29[5:0] – CREDIT_C11 (CPU Address 534)
 - W2 - QOSC30[5:0] – CREDIT_C12 (CPU Address 535)
 - W3 - QOSC31[5:0] – CREDIT_C13 (CPU Address 536)

QOSC28 through QOSC31 represents one set of WFQ parameters for a 10/100 port. There are four such sets of values. The granularity of the numbers is 1 and their sum must be 64. QOSC31 corresponds to W3 and QOSC28 corresponds to W0.

QOSC28[7:6]: Priority service type for the ports select this parameter set. Option 1 to 4.

QOSC29[7]: Priority service allow flow control for the ports select this parameter set.

QOSC29[6]: Flow control pause best effort traffic only

12.7.29 Classes WFQ Credit Set 2

- Accessed by serial interface (R/W)
 - W0 - QOSC32[5:0] – CREDIT_C20 (CPU Address 537)
 - W1 - QOSC33[5:0] – CREDIT_C21 (CPU Address 538)
 - W2 - QOSC34[5:0] – CREDIT_C22 (CPU Address 539)
 - W3 - QOSC35[5:0] – CREDIT_C23 (CPU Address 53a)

QOSC35 through QOSC32 represents one set of WFQ parameters for a 10/100 port. There are four such sets of values. The granularity of the numbers is 1 and their sum must be 64. QOSC35 corresponds to W3 and QOSC32 corresponds to W0.

QOSC32[7:6]: Priority service type for the ports select this parameter set. Option 1 to option 4.

QOSC33[7]: Priority service allow flow control for the ports select this parameter set.

QOSC33[6]: Flow Control pause best effort traffic only

12.7.30 Classes WFQ Credit Set 3

- Accessed by serial interface (R/W)
 - W0 - QOSC36[5:0] – CREDIT_C30 (CPU Address 53b)
 - W1 - QOSC37[5:0] – CREDIT_C31 (CPU Address 53c)
 - W2 - QOSC38[5:0] – CREDIT_C32 (CPU Address 53d)
 - W3 - QOSC39[5:0] – CREDIT_C33 (CPU Address 53e)

QOSC39 through QOSC36 represents one set of WFQ parameters for a 10/100 port. There are four such sets of values. The granularity of the numbers is 1 and their sum must be 64. QOSC39 corresponds to W3 and QOSC36 corresponds to W0.


QOSC36[7:6]: Priority service type for the ports select this parameter set. Option 1 to option 4.

QOSC37[7]: Priority service allow flow control for the ports select this parameter set.

QOSC37[6]: Flow Control pause best effort traffic only

12.7.31 RDRC0 – WRED Rate Control 0

- I²C Address 0FB, CPU Address 553
- Accessed by serial Interface and I²C (R/W)




Bits [7:4]: • Corresponds to the frame drop percentage X% for WRED. Granularity 6.25%.

Bits [3:0]: • Corresponds to the frame drop percentage Y% for WRED. Granularity 6.25%.

See Programming QoS Registers application note for more information.

12.7.32 RDRC1 – WRED Rate Control 1

- I²C Address 0FC, CPU Address 554
- Accessed by serial Interface and I²C (R/W)



Bits [7:4]: • Corresponds to the frame drop percentage Z% for WRED. Granularity 6.25%.

Bits [3:0]: • Corresponds to the best effort frame drop percentage B%, when shared pool is all in use and destination port best effort queue reaches UCC. Granularity 6.25%.

See Programming QoS Register Application Note for more information.

12.7.33 User Defined Logical Ports and Well Known Ports

The MVTX2601 supports classifying packet priority through layer 4 logical port information. It can be setup by 8 Well Known Ports, 8 User Defined Logical Ports and 1 User Defined Range. The 8 Well Known Ports supported are:

- 0:23
- 1:512
- 2:6000
- 3:443
- 4:111
- 5:22555
- 6:22
- 7:554



Their respective priority can be programmed via Well_Known_Port [7:0] priority register. Well_Known_Port_Enable can individually turn on/off each Well Known Port if desired.

Similarly, the User Defined Logical Port provides the user programmability to the priority, plus the flexibility to select specific logical ports to fit the applications. The 8 User Logical Ports can be programmed via User_Port 0-7 registers. Two registers are required to be programmed for the logical port number. The respective priority can be programmed to the User_Port [7:0] priority register. The port priority can be individually enabled/disabled via User_Port_Enable register.

The User Defined Range provides a range of logical port numbers with the same priority level. Programming is similar to the User Defined Logical Port. Instead of programming a fixed port number, an upper and lower limit need to be programmed, they are: {RHIGHH, RHIGHL} and {RLOWH, RLOWL} respectively. If the value in the upper limit is smaller or equal to the lower limit, the function is disabled. Any IP packet with a logical port that is less than the upper limit and more than the lower limit will use the priority specified in RRIORITY.

12.7.33.1 USER_PORT0_(0~7) – User Define Logical Port (0~7)


- USER_PORT_0 - I²C Address h0D6 + 0DE; CPU Address 580(Low) + 581(High)
- USER_PORT_1 - I²C Address h0D7 + 0DF; CPU Address 582 + 583
- USER_PORT_2 - I²C Address h0D8 + 0E0; CPU Address 584 + 585
- USER_PORT_3 - I²C Address h0D9 + 0E1; CPU Address 586 + 587
- USER_PORT_4 - I²C Address h0DA + 0E2; CPU Address 588 + 589
- USER_PORT_5 - I²C Address h0DB + 0E3; CPU Address 58a + 58b
- USER_PORT_6 - I²C Address h0DC + 0E4; CPU Address 58c + 58d
- USER_PORT_7 - I²C Address h0DD + 0E5; CPU Address 58e + 58f
- Accessed by serial interface and I²C (R/W)



- (Default 00) This register is duplicated eight times from PORT 0 through PORT 7 and allows the definition of eight separate ports.

12.7.33.2 USER_PORT_[1:0]_PRIORITY - User Define Logic Port 1 and 0 Priority

- I²C Address h0E6, CPU Address 590
- Accessed by serial interface and I²C (R/W)



- The chip allows the definition of the priority

- Bits [3:0]: • Priority setting, transmission + dropping, for logic port 0
- Bits [7:4]: • Priority setting, transmission + dropping, for logic port 1 (Default 00)

12.7.33.3 USER_PORT_[3:2]_PRIORITY - User Define Logic Port 3 and 2 Priority

- I²C Address h0E7, CPU Address 591
- Accessed by serial interface and I²C (R/W)

7	5	4	3	1	0
Priority 3	Drop	Priority 2	Drop		

12.7.33.4 USER_PORT_[5:4]_PRIORITY - User Define Logic Port 5 and 4 Priority

- I²C Address h0E8, CPU Address 592
- Accessed by serial interface and I²C (R/W)

7	5	4	3	1	0
Priority 5	Drop	Priority 4	Drop		

- (Default 00)

12.7.33.5 USER_PORT_[7:6]_PRIORITY - USER DEFINE LOGIC PORT 7 AND 6 PRIORITY

- I²C Address h0E9, CPU Address 593
- Accessed by serial interface and I²C (R/W)

7		5	4	3		1	0
Priority 7	Drop	Priority 6	Drop				

- (Default 00)

12.7.33.6 USER_PORT_ENABLE [7:0] – User Define Logic 7 to 0 Port Enables

- I²C Address h0EA, CPU Address 594
- Accessed by serial interface and I²C (R/W)

7	6	5	4	3	2	1	0
P7	P6	P5	P4	P3	P2	P1	P0

- (Default 00)

12.7.33.7 WELL_KNOWN_PORT [1:0] PRIORITY- Well Known Logic Port 1 and 0 Priority

- I²C Address h0EB, CPU Address 595
- Accessed by serial interface and I²C (R/W)

7	5	4	3	1	0
Priority 1	Drop	Priority 0	Drop		

- Priority 0 - Well known port 23 for telnet applications.
- Priority 1 - Well known port 512 for TCP/UDP
- (Default 00)

12.7.33.8 WELL_KNOWN_PORT [3:2] PRIORITY- Well Known Logic Port 3 and 2 Priority

- I²C Address h0EC, CPU Address 596
- Accessed by serial interface and I²C (R/W)

	7	5 4	3	1	0
Priority 3	Drop	Priority 2	Drop		

- Priority 2 - Well known port 6000 for XWIN.
- Priority 3 - Well known port 443 for http. sec
- (Default 00)

12.7.33.9 WELL_KNOWN_PORT [5:4] PRIORITY- Well Known Logic Port 5 and 4 Priority

- I²C Address h0ED, CPU Address 597
- Accessed by serial interface and I²C (R/W)

	7	5 4	3	1	0
Priority 5	Drop	Priority 4	Drop		

- Priority 4 - Well known port 111 for sun rpe.
- Priority 5 - Well known port 22555 for IP Phone call setup.
- (Default 00)

12.7.33.10 WELL_KNOWN_PORT [7:6] PRIORITY- Well Known Logic Port 7 and 6 Priority

- I²C Address h0EE, CPU Address 598
- Accessed by serial interface and I²C (R/W)

	7	5 4	3	1	0
Priority 7	Drop	Priority 6	Drop		

- Priority 6 - Well known port 22 for ssh.
- Priority 7 - Well known port 554 for rtsp.
- (Default 00)

12.7.33.11 WELL_KNOWN_PORT_ENABLE [7:0] – Well Known Logic 7 to 0 Port Enables

- I²C Address h0EF, CPU Address 599
- Accessed by serial interface and I²C (R/W)

7	6	5	4	3	2	1	0
P7	P6	P5	P4	P3	P2	P1	P0

- 1 - Enable
- 0 - Disable
- Default 00)

12.7.33.12 RLOWL – User Define Range Low Bit 7:0

- I²C Address h0F4, CPU Address: 59a
- Accessed by serial interface and I²C (R/W)
- (Default 00)

12.7.33.13 RLOWH – User Define Range Low Bit 15:8

- I²C Address h0F5, CPU Address: 59b
- Accessed by serial interface and I²C (R/W)
- (Default 00)

12.7.33.14 RHIGHL – User Define Range High Bit 7:0


- I²C Address h0D3, CPU Address: 59c
- Accessed by serial interface and I²C (R/W)
- (Default 00)

12.7.33.15 RHIGHH – User Define Range High Bit 15:8

- I²C Address h0D4, CPU Address: 59d
- Accessed by serial interface and I²C (R/W)
- (Default 00)

12.7.33.16 RRIORITY – User Define Range Priority

- I²C Address h0D5, CPU Address: 59e
- Accessed by serial interface and I²C (R/W)




- RLOW and RHIGH form a range for logical ports to be classified with priority specified in RRIORITY.

Bit [3:1] • Transmit Priority

Bits [0]: • Drop Priority

12.8 Group 6 Address MISC Group**12.8.1 MII_OP0 – MII Register Option 0**


- I²C Address F0, CPU Address:h600
- Accessed by serial interface and I²C (R/W)



- Bits [7]:
 - Half duplex flow control feature
 - 0 = Half duplex flow control always enable
 - 1 = Half duplex flow control by negotiation
- Bits [6]:
 - Link partner reset auto-negotiate disable
- Bits [5]:
 - Disable jabber detection. This is for HomePNA application or any serial operation slower than 10 Mbps.
 - 1 = disable
 - 0 = enable
- Bit [4:0]:
 - Vendor specified link status register address (null value means don't use it) (Default 00); used when the Linkup bit position in the PHY is non-standard.

12.8.2 MII_OP1 – MII Register Option 1


- I²C Address F1, CPU Address:h601
- Accessed by serial interface and I²C (R/W)



- Bits [3:0]:
 - Duplex bit location in vendor specified register
- Bits [7:4]:
 - Speed bit location in vendor specified register (Default 00)

12.8.3 FEN – Feature Register

- I²C Address F2, CPU Address:h602)
- Accessed by serial interface and I²C (R/W)



- Bits [1:0]:
 - Reserved **(Default 0)**
- Bit [2]:
 - Support DS EF Code. **(Default 0)**
 - When 101110 is detected in DS field (TOS [7:2]), the frame priority is set for 110 and drop is set for 0.
- Bit [5:3]:
 - Reserved **(Default 010)**
- Bit [6]:
 - Disable MII Management State Machine
 - 0: Enable MII Management State Machine **(Default 0)**
 - 1: Disable MII Management State Machine
- Bit [7]:
 - Disable using MCT link list structure
 - 0: Enable using MCT Link List structure **(Default 0)**
 - 1: Disable using MCT Link List structure

12.8.4 MIIC0 – MII Command Register 0

- CPU Address:h603
- Accessed by serial interface only (R/W)
- Bit [7:0] MII Data [7:0]

Note: Before programming MII command: set FEN[6], check MIIC3, making sure no RDY, and no VALID; then program MII command.


12.8.5 MIIC1 – MII Command Register 1

- CPU Address:h604
- Accessed by serial interface only (R/W)
- Bit [7:0] MII Data [15:8]

Note: Before programming MII command: set FEN[6], check MIIC3, making sure no RDY and no VALID; then program MII command.

12.8.6 MIIC2 – MII Command Register 2

- CPU Address:h605
- Accessed by serial interface only (R/W)




Bits [4:0]: • REG_AD – Register PHY Address

Bit [6:5] • OP – Operation code “10” for read command and “01” for write command

Note: Before programming MII command: set FEN[6], check MIIC3, making sure no RDY and no VALID; then program MII command. Writing to this register will initiate a serial management cycle to the MII management interface. For detail information, please refer to the PHY Control Application Note.

12.8.7 MIIC3 – MII Command Register 3

- CPU Address:h606
- Accessed by serial interface only (R/W)



Bits [4:0]: • PHY_AD – 5 Bit PHY Address

Bit [6] • VALID – Data Valid from PHY (Read Only)

Bit [7] • RDY – Data is returned from PHY (Ready Only)

Note: Before programming MII command: set FEN[6], check MIIC3, making sure no RDY and no VALID; then program MII command.

12.8.8 MIID0 – MII Data Register 0

- CPU Address:h607
- Accessed by serial interface only (RO)
- Bit [7:0] MII Data [7:0]

12.8.9 MIID1 – MII Data Register 1

- CPU Address:h608
- Accessed by serial interface only (RO)
- Bit [7:0] MII Data [15:8]

12.8.10 LED Mode – LED Control

- CPU Address:h609
- Accessed by serial interface and I²C (R/W)

7	5	4	3	2	1	0
			Clock rate	Hold Time		

- Bit [0]: • Reserved (Default 0)
- Bit [2:1]: • Hold time for LED signal (**Default= 00**)
 00 = 8 msec 01 = 16 msec
 10 = 32 msec 11 = 64 msec
- Bit [4:3]: • LED clock frequency (**Default 0**)
 For 100 MHz SCLK,
 00 = 100 M/8 = 12.5 MHz 01 = 100 M/16 = 6.25 MHz
 10 = 100 M/32 = 3.125 MHz 11 = 100 M/64 = 1.5625 MHz
 For 125 MHz SCLK
 00 = 125 M/64 = 1953 KHz 01 = 125 M/128 = 977 KHz
 10 = 125 M/512 = 244 KHz 11 = 125 M/1024 = 122 KHz
- Bit [7:6]: • Reserved. Must be 0. (**Default 0**)

12.8.11 CHECKSUM - EEPROM Checksum

- I²C Address FF, CPU Address:h60b
- Accessed by serial interface and I²C (R/W)

- Bit [7:0]: • (Default 0)

Before requesting that the MVTX2601 updates the EEPROM device, the correct checksum needs to be calculated and written into this checksum register. When the MVTX2601 boots from the EEPROM the checksum is calculated and the value must be zero. If the checksum is not zeroed the MVTX2601 does not start and pin CHECKSUM_OK is set to zero.

The checksum formula is: FF


$$\sum \text{I2C register} = 0$$

$$I=0$$

12.9 Group 7 Address Port Mirroring Group

12.9.1 MIRROR1_SRC – Port Mirror source port


- CPU Address 700
- Accessed by serial interface (R/W) (Default 7F)



- Bit [4:0]:
- Source port to be mirrored. Use illegal port number to disable mirroring.
- Bit [5]:
- 1 – select ingress data
 - 0 – select egress data
- Bit [7]:
- Must be '1'

12.9.2 MIRROR1_DEST – Port Mirror destination


- CPU Address 701
- Accessed by serial interface (R/W) (Default 17)



- Bit [4:0]:
- Port Mirror Destination

12.9.3 MIRROR2_SRC – Port Mirror source port


- CPU Address 702
- Accessed by serial interface (R/W) (Default FF)



- Bit [4:0]:
- Source port to be mirrored. Use illegal port number to disable mirroring.
- Bit [5]:
- 1 – select ingress data
 - 0 – select egress data
- Bit [7]
- Must be 1

12.9.4 MIRROR2_DEST – Port Mirror destination

- CPU Address 703
- Accessed by serial interface (R/W) (Default 00)




Bit [4:0]: • Port Mirror Destination

12.10 Group F Address CPU Access Group

12.10.1 GCR-Global Control Register

- CPU Address: hF00
- Accessed by serial interface. (R/W)



- Bit [0]: • Store configuration (**Default = 0**)
- Write '1' followed by '0' to store configuration into external EEPROM
- Bit [1]: • Store configuration and reset (**Default = 0**)
- Write '1' to store configuration into external EEPROM and reset chip
- Bit [2]: • Start BIST (**Default = 0**)
- Write '1' followed by '0' to start the device's built-in self-test. The result is found in the DCR register.
- Bit [3]: • Soft Reset (Default = 0)
- Write '1' to reset chip
- Bit [7:4]: • Reserved

12.10.2 DCR-Device Status and Signature Register

- CPU Address: hF01
- Accessed by serial interface. (RO)

7	6	5	4	3	2	1	0
Revision	Signature	RE	BinP	BR	BW		

- Bit [0]:
- 1: Busy writing configuration to I²C
 - 0: Not busy writing configuration to I²C
- Bit [1]:
- 1: Busy reading configuration from I²C
 - 0: Not busy reading configuration from I²C
- Bit [2]:
- 1: BIST in progress
 - 0: BIST not running
- Bit [3]:
- 1: RAM Error
 - 0: RAM OK
- Bit [5:4]:
- Device Signature
 - 01: MVTX2601 device
- Bit [7:6]:
- Revision
 - 00: Initial Silicon
 - 01: XA1 Silicon

12.10.3 DCR1-Chip status

- CPU Address: hF02
- Accessed by serial interface (RO)

7	6	0
CIC		


- Bit [7]
- Chip initialization completed

12.10.4 DPST – Device Port Status Register

- CPU Address:hF03
 - Accessed by serial interface (R/W)
- Bit [4:0]:
- Read back index register. This is used for selecting what to read back from DTST. **(Default 00)**
 - 5'b00000 - Port 0 Operating mode and Negotiation status
 - 5'b00001 - Port 1 Operating mode/Neg status
 - 5'b00010 - Port 2 Operating mode/Neg status
 - 5'b00011 - Port 3 Operating mode/Neg status
 - 5'b00100 - Port 4 Operating mode/Neg status
 - 5'b00101 - Port 5 Operating mode/Neg status
 - 5'b00110 - Port 6 Operating mode/Neg status
 - 5'b00111 - Port 7 Operating mode/Neg status
 - 5'b01000 - Port 8 Operating mode/Neg status
 - 5'b01001 - Port 9 Operating mode/Neg status
 - 5'b01010 - Port 10 Operating mode/Neg status
 - 5'b01011 - Port 11 Operating mode/Neg status
 - 5'b01100 - Port 12 Operating mode/Neg status
 - 5'b01101 - Port 13 Operating mode/Neg status
 - 5'b01110 - Port 14 Operating mode/Neg status
 - 5'b01111 - Port 15 Operating mode/Neg status
 - 5'b10000 - Port 16 Operating mode/Neg status
 - 5'b10001 - Port 17 Operating mode/Neg status
 - 5'b10010 - Port 18 Operating mode/Neg status
 - 5'b00011 - Port 19 Operating mode/Neg status
 - 5'b10100 - Port 20 Operating mode/Neg status
 - 5'b10101 - Port 21 Operating mode/Neg status
 - 5'b10110 - Port 22 Operating mode/Neg status
 - 5'b10111 - Port 23 Operating mode/Neg status

12.10.5 DTST – Data read back register

- CPU Address: hF04
- Accessed by serial interface (RO)
- This register provides various internal information as selected in DPST bit [4:0]. Refer to the PHY Control Application Note.



When bit is 1:

- Bit [0] – Flow control enable
- Bit [1] – Full duplex port
- Bit [2] – Fast Ethernet port
- Bit [3] – Link is down
- Bit [7:4] – Reserved

12.10.6 PLLCR - PLL Control Register

- CPU Address: hF05
- Accessed by serial interface (RW)

Bit [3] Must be '1'

Bit [7] Selects strap option or LCLK/OECLK registers
 0 - Strap option (default)
 1 - LCLK/OECLK registers

12.10.7 LCLK - LA_CLK delay from internal OE_CLK

- CPU Address: hF06
- Accessed by serial interface (RW)

PD[12:10]	LCLK	Delay
000b	80h	8 Buffers Delay
001b	40h	7 Buffers Delay
010b	20h	6 Buffers Delay
011b	10h	5 Buffers Delay (Recommend)
100b	08h	4 Buffers Delay
101b	04h	3 Buffers Delay
110b	02h	2 Buffers Delay
111b	01h	1 Buffers Delay

The LCLK delay from SCLK is the sum of the delay programmed in here and the delay in OECLK register.

12.10.8 OECLK - Internal OE_CLK delay from SCLK

- CPU Address: hF07
- Accessed by serial interface (RW)

The OE_CLK is used for generating the OE0 and OE1 signals.

PD[15:13]	OECLK	Delay
000b	80h	8 Buffers Delay
001b	40h	7 Buffers Delay (Recommend)
010b	20h	6 Buffers Delay
011b	10h	5 Buffers Delay
100b	08h	4 Buffers Delay
101b	04h	3 Buffers Delay
110b	02h	2 Buffers Delay
111b	01h	1 Buffers Delay

12.10.9 DA – DA Register

- CPU Address: hFFF
- Accessed by CPU and serial interface (RO)
- Always return 8'h **DA**. Indicate the serial port connection is good.

13.0 BGA and Ball Signal Descriptions

13.1 BGA Views (Top-View)

13.1.1 Encapsulated View

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29						
A			LA_D4	LA_D7	LA_D10	LA_D13	LA_D15	LA_D18	LA_OE0	LA_A8	LA_A13	LA_A16	LA_A19	LA_D32	LA_D35	LA_D38	LA_D41	LA_D44	OE_CLK0	LA_CLK1	TRUNK1	MIRROR3	MIRROR1	SCL	SDA	STROBE	TSTOUT7								
B		LA_D1	LA_D3	LA_D6	LA_D9	LA_D12	LA_D14	LA_DSC	LA_OE1	LA_A7	LA_A12	LA_A15	LA_A18	LA_D33	LA_D36	LA_D39	LA_D42	LA_D45	OE_CLK1	LA_CLK2	LA_D62	MIRROR5	MIRROR2	RESE	RESE	D0	TSTOUT8	TSTOUT3							
C	LA_CLK	LA_D0	LA_D2	LA_D5	LA_D8	LA_D11	LA_D3	LA_OE2	LA_WDE	TMO	LA_A14	LA_A17	LA_A20	LA_D34	LA_D37	LA_D40	LA_D43	OE_CLK2	LA_CLK3	P_D	TRUNK0	MIRROR3	MIRROR0	AUTO	TSTOUT11	TSTOUT9	TSTOUT4	TSTOUT0							
D	AGND	LA_D17	LA_D19	LA_D21	LA_D23	LA_D25	LA_D27	LA_D29	LA_D31	LA_A6	LA_A10	LA_A14	LA_A17	LA_D49	LA_D51	LA_D53	LA_D55	LA_D57	LA_D59	LA_D61	LA_D63	LA_D47	SCANCOL	SCANCLK	TSTOUT14	TSTOUT13	TSTOUT12	TSTOUT10	TSTOUT5						
E	SCLK	LA_D16	LA_D18	LA_D20	LA_D22	LA_D24	LA_D26	LA_D28	LA_D30	LA_A5	LA_A9	LA_A11	LA_WDE1	LA_D48	LA_D50	LA_D52	LA_D54	LA_D56	LA_D58	LA_D60	RESE	LA_D46	SCANLINK	TSTOUT15	RESE	RESE	SCANMOD	TSTOUT6	TSTOUT2						
F	AVC	RESIN	SCANEN	RESE	RESE	VCC										VCC										RESE	RESE	RESE	RESE	RESE					
G	RESE	RESE	RESE	RESE	RESE	VCC										VCC										RESE	RESE	RESE	RESE	RESE					
H	RESE	RESE	RESE	RESE	RESE	VCC										VCC										RESE	RESE	RESE	RESE	RESE					
J	RESE	RESE	RESE	RESE	RESE	VCC										VCC										RESE	RESE	RESE	RESE	RESE					
K	RESE	RESE	RESE	RESE	RESE	VCC										VCC										RESE	RESE	RESE	RESE	RESE					
L	RESE	RESE	RESE	RESE	RESE	VCC										VCC										RESE	RESE	RESE	RESE	RESE					
M	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	RESE	RESE	RESE					
N	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	RESE	RESE	RESE					
P	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	RESE	RESE	RESE					
R	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	RESE	RESE	RESE					
T	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	RESE	RESE	RESE					
U	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	RESE	RESE	RESE					
V	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	RESE	RESE	RESE					
W	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	RESE	RESE	RESE					
Y	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	RESE	RESE	RESE					
AA	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	RESE	RESE	RESE					
AB	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	RESE	RESE	RESE					
AC	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	M23	M23	M23					
AD	RESE	RESE	RESE	RESE	RESE	VDD	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	VSS	RESE	RESE	M23	M23	M23					
AE	M0_T	M0_T	M0_T	M3_T	M3_T	M3_R	M5_T	M5_T	M5_R	M8_T	M8_T	M8_R	M10_T	M10_T	M10_R	M13_T	M13_T	M13_R	M15_T	M15_T	M15_R	M16_T	M16_T	M16_R	M18_T	M18_T	M18_R	M20_T	M20_T	M20_R	M22_T				
AF	M0_R	M0_R	M0_C	M3_T	M3_C	M3_R	M5_T	M5_C	M5_R	M8_T	M8_C	M8_R	M10_T	M10_C	M10_R	M13_T	M13_C	M13_R	M14_C	M14_C	M14_R	M16_C	M16_C	M16_R	M17_C	M17_C	M17_R	M18_C	M18_C	M18_R	M20_C	M20_C	M20_R	M22_C	M22_C
AG	M1_T	M1_T	M1_T	M2_T	M2_C	M4_T	M4_C	M6_T	M6_C	M7_T	M7_C	M9_T	M9_C	M11_T	M11_C	M12_T	M12_C	M14_T	M14_C	M15_C	M15_C	M16_C	M16_C	M18_C	M18_C	M19_C	M19_C	M21_C	M21_C	M21_R	M22_T	M22_T			
AH		M1_R	M1_C	M2_T	M2_R	M4_T	M4_R	M6_T	M6_R	M7_T	M7_R	M9_T	M9_R	M11_T	M11_R	M12_T	M12_R	M14_T	M14_R	M13_C	M13_C	M15_C	M15_C	M17_C	M17_C	M19_C	M19_C	M21_T	M21_T	M21_R	M22_T	M22_T			
AJ			M1_R	M2_T	M2_R	M4_T	M4_R	M6_T	M6_R	M7_T	M7_R	M9_T	M9_R	M11_T	M11_R	M12_T	M12_R	M14_T	M14_R	M16_C	M16_C	M17_C	M17_C	M19_C	M19_C	M21_T	M21_T	M21_R	M21_R	M21_R					
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29						

13.2 Ball – Signal Descriptions

All pins are CMOS type; all Input Pins are 5 Volt tolerance; and all Output Pins are 3.3 CMOS drive.

13.2.1 Ball Signal Descriptions

Ball No(s)	Symbol	I/O	Description
I²C Interface Note: Use I²C and Serial control interface to configure the system			
A24	SCL	Output	I ² C Data Clock
A25	SDA	I/O-TS with pull up	I ² C Data I/O
Serial Control Interface			
A26	STROBE	Input with weak internal pull up	Serial Strobe Pin
B26	D0	Input	Serial Data Input
C25	AUTOFD	Output with pull up	Serial Data Output (AutoFD)
Frame Buffer Interface			
D20, B21, D19, E19, D18, E18, D17, E17, D16, E16, D15, E15, D14, E14, D13, E13, D21, E21, A18, B18, C18, A17, B17, C17, A16, B16, C16, A15, B15, C15, A14, B14, D9, E9, D8, E8, D7, E7, D6, E6, D5, E5, D4, E4, D3, E3, D2, E2, A7, B7, A6, B6, C6, A5, B5, C5, A4, B4, C4, A3, B3, C3, B2, C2	LA_D[63:0]	I/O-TS with pull up	Frame Bank A– Data Bit [63:0]
C14, A13, B13, C13, A12, B12, C12, A11, B11, C11, D11, E11, A10, B10, D10, E10, A8, C7	LA_A[20:3]	Output	Frame Bank A – Address Bit [20:3]
B8	LA_ADSC#	Output with pull up	Frame Bank A Address Status Control
C1	LA_CLK	Output	Frame Bank A Clock Input
C9	LA_WE#	Output with pull up	Frame Bank A Write Chip Select for one layer SRAM application
D12	LA_WE0#	Output with pull up	Frame Bank A Write Chip Select for lower layer of two layers SRAM application
E12	LA_WE1#	Output with pull up	Frame Bank A Write Chip Select for upper layer of two layers SRAM application

Ball No(s)	Symbol	I/O	Description
C8	LA_OE#	Output with pull up	Frame Bank A Read Chip Select for one layer SRAM application
A9	LA_OE0#	Output with pull up	Frame Bank A Read Chip Select for lower layer of two layers SRAM application
B9	LA_OE1#	Output with pull up	Frame Bank A Read Chip Select for upper layer of two layers SRAM application
Fast Ethernet Access Ports [23:0] RMII			
R28	M_MDC	Output	MII Management Data Clock – (Common for all MII Ports [23:0])
P28	M_MDIO	I/O-TS with pull up	MII Management Data I/O – (Common for all MII Ports – [23:0])
R29	M_CLKI	Input	Reference Input Clock
AC29, AE28, AJ27, AF27, AJ25, AF24, AH23, AE19, AF21, AJ19, AF18, AJ17, AJ15, AF15, AJ13, AF12, AJ11, AJ9, AF9, AJ7, AF6, AJ5, AJ3, AF1	M[23:0]_RXD[1]	Input with weak internal pull up resistors.	Ports [23:0] – Receive Data Bit [1]
AC28, AF28, AH27, AE27, AH25, AE24, AF22, AF20, AE21, AH19, AH20, AH17, AH15, AE15, AH13, AE12, AH11, AH9, AE9, AH7, AE6, AH5, AH2, AF2	M[23:0]_RXD[0]	Input with weak internal pull up resistors	Ports [23:0] – Receive Data Bit [0]
AC27, AF29, AG27, AF26, AG25, AG23, AF23, AG21, AH21, AF19, AF17, AG17, AG15, AF14, AG13, AF11, AG11, AG9, AF8, AG7, AF5, AG5, AH3, AF3	M[23:0]_CRS_DV	Input with weak internal pull down resistors.	Ports [23:0] – Carrier Sense and Receive Data Valid
AD29, AG28, AJ26, AE26, AJ24, AE23, AJ22, AJ20, AE20, AJ18, AJ21, AJ16, AJ14, AE14, AJ12, AE11, AJ10, AJ8, AE8, AJ6, AE5, AJ4, AG1, AE1	M[23:0]_TXEN	I/O- TS with pull up, slew	Ports [23:0] – Transmit Enable Strap option for RMII/GPSI

Ball No(s)	Symbol	I/O	Description
AD27, AH28, AG26, AE25, AG24, AE22, AJ23, AG20, AE18, AG18, AE16, AG16, AG14, AE13, AG12, AE10, AG10, AG8, AE7, AG6, AE4, AG4, AG3, AE3	M[23:0]_TXD[1]	Output, slew	Ports [23:0] – Transmit Data Bit [1]
AD28, AG29, AH26, AF25, AH24, AG22, AH22, AE17, AG19, AH18, AF16, AH16, AH14, AF13, AH12, AF10, AH10, AH8, AF7, AH6, AF4, AH4, AG2, AE2	M[23:0]_TXD[0]	Output, slew	Ports [23:0] – Transmit Data Bit [0]
LED Interface			
C29	LED_CLK/TSTOUT0	I/O- TS with pull up	LED Serial Interface Output Clock
D29	LED_SYN/TSTOUT1	I/O- TS with pull up	LED Output Data Stream Envelope
E29	LED_BIT/TSTOUT2	I/O- TS with pull up	LED Serial Data Output Stream
B27, A27, E28, D28, C28, B28	TSTOUT[8:3]	I/O- TS with pull up	(Reserved)
C27	INIT_DONE/TSTOUT 9	I/O- TS with pull up	System start operation
D27	INIT_START/TSTOU T10	I/O- TS with pull up	Start initialization
C26	CHECKSUM_OK/TS TOUT11	I/O- TS with pull up	EEPROM read OK
D26	FCB_ERR/TSTOUT1 2	I/O- TS with pull up	FCB memory self test fail
D25	MCT_ERR/TSTOUT1 3	I/O- TS with pull up	MCT memory self test fail
D24	BIST_IN_PRC/TSTO UT14	I/O- TS with pull up	Processing memory self test
E24	BIST_DONE/TSTOU T15	I/O- TS with pull up	Memory self test done
Trunk Enable			
C22	TRUNK0	Input w/ weak internal pull down resistors	Trunk Port Enable

Ball No(s)	Symbol	I/O	Description															
A21	TRUNK1	Input w/ weak internal pull down resistors	Trunk Port Enable															
Test Facility																		
U3	T_MODE0	I/O-TS	Test Pin – Set Mode upon Reset, and provides NAND Tree test output during test mode (Pull Up)															
C10	T_MODE1	I/O-TS	Test Pin – Set Mode upon Reset, and provides NAND Tree test output during test mode (Pull Up) <table border="0"> <tr> <td>T_MODE1</td> <td>T_MODE0</td> <td></td> </tr> <tr> <td>0</td> <td>0</td> <td>NandTree</td> </tr> <tr> <td>0</td> <td>1</td> <td>Reserved</td> </tr> <tr> <td>1</td> <td>0</td> <td>reserved</td> </tr> <tr> <td>1</td> <td>1</td> <td>Regular operation</td> </tr> </table> <p>T_MODE0 and T_MODE1 are used for manufacturing tests. The signals should both be set to 1 for regular operation.</p>	T_MODE1	T_MODE0		0	0	NandTree	0	1	Reserved	1	0	reserved	1	1	Regular operation
T_MODE1	T_MODE0																	
0	0	NandTree																
0	1	Reserved																
1	0	reserved																
1	1	Regular operation																
F3	SCAN_EN	Input with pull down	Scan Enable 0 - Normal mode (unconnected)															
E27	SCANMODE	Input with pull down	1 - Enables Test mode 0 - Normal mode (unconnected)															
System Clock, Power and Ground Pins																		
E1	SCLK	Input	System Clock at 100 MHz															
K12, K13, K17, K18 M10, N10, M20, N20, U10, V10, U20, V20, Y12, Y13, Y17, Y18	VDD	Power	+2.5 Volt DC Supply															
F13, F14, F15, F16, F17, N6, P6, R6, T6, U6, N24, P24, R24, T24, U24, AD13, AD14, AD15, AD16, AD17	VCC	Power	+3.3 Volt DC Supply															

Ball No(s)	Symbol	I/O	Description
M12, M13, M14, M15, M16, M17, M18, N12, N13, N14, N15, N16, N17, N18, P12, P13, P14, P15, P16, P17, P18, R12, R13, R14, R15, R16, R17, R18, T12, T13, T14, T15, T16, T17, T18, U12, U13, U14, U15, U16, U17, U18, V12, V13, V14, V15, V16, V17, V18,	VSS	Power Ground	Ground
F1	AVCC	Analog Power	Analog +2.5 Volt DC Supply
D1	AGND	Analog Ground	Analog Ground
MISC			
D22	SCANCOL	Input	Scans the Collision signal of Home PHY
D23	SCANCLK	Input/ output	Clock for scanning Home PHY collision and link
E23	SCANLINK	Input	Link up signal from Home PHY
F2	RESIN#	Input	Reset Input
G2	RESETOUT#	Output	Reset PHY

Ball No(s)	Symbol	I/O	Description
B22, F4, F5, G4, G5, H4, H5, J4, J5, K4, K5, L4, L5, M4, M5, N4, N5, G3, H1, H2, H3, J1, J2, J3, K1, K2, K3, L1, L2, L3, M1, M2, M3, U4, U5, V4, V5, W4, W5, Y4, Y5, AA4, AA5, AB4, AB5, AC4, AC5, AD4, AD5, W1, Y1, Y2, Y3, AA1, AA2, AA3, AB1, AB2, AB3, AC1, AC2, AC3, AD1, AD2, AD3, N3, N2, N1, P3, P2, P1, R5, R4, R3, R2, R1, T5, T4, T3, T2, T1, W3, W2, V1, G1, V3, P4, P5, V2, U1, U2, U26, U25, V26, V25, W26, W25, Y27, Y26, AA26, AA25, AB26, AB25, AC26, AC25, AD26, AD25, T28, U28, R25, U29, T29, U27, V29, V28, V27, W29, W28, W27, Y29, Y28, Y25, AA29, AA28, AA27, AB29, AB28, AB27, T26, R26, T27, T25, P29, G26, G25, H26, H25, J26, J25, K25, K26, M25, L26, M26, L25, N26, N25, P26, P25, F28, G28, E25, G29, F29, G27, H29, H28, H27, J29, J28, J27, K29, K28, K27, L29, L28, L27, M29, M28, M27, F26, E26, F27, F25, N29, B24, E20, B25	Reserved	I/O-TS	Reserved Pins. Leave unconnected.
Bootstrap Pins (Default= pull up, 1= pull up 0= pull down)			
After reset TSTOUT0 to TSTOUT15 are used by the LED interface.			
C29	TSTOUT0		Reserved
D29	TSTOUT1	Default: Enable (1)	RMII MAC Power Saving Enable 0 - No power saving 1 - Power saving
C28, B28, E29	TSTOUT[4:2]		Reserved
D28	TSTOUT5	Default: SCLK (1)	Scan Speed 0 - ¼ SCLK(HPNA) 1 - SCLK

Ball No(s)	Symbol	I/O	Description
E28	TSTOUT6		Reserved
A27	TSTOUT7	Default: 128 K x 32 or 128 K x 64 (1)	Memory Size 0 - 256 K x 32 or 256 K x 64 (4 M total) 1 - 128 K x 32 or 128 K x 64 (2 M total)
B27	TSTOUT8	Default: Not Installed (1)	EEPROM Installed 0 - EEPROM installed 1 - EEPROM not installed
C27	TSTOUT9	Default: MCT aging enable (1)	MCT Aging 0 - MCT aging disable 1 - MCT aging enable
D27	TSTOUT10	Default: FCB aging enable (1)	FCB Aging 0 - FCB aging disable 1 - FCB aging enable
C26	TSTOUT11	Default: Timeout reset enable (1)	Timeout Reset 0 - Time out reset disable 1 - Time out reset enable. Issue reset if any state machine did not go back to idle for 5 Sec.
D26	TSTOUT12		Reserved
D25	TSTOUT13	Default: Single depth (1)	FDB RAM depth (1 or 2 layers) 0 - Two layers 1 - One layer
D24	TSTOUT14		Reserved
E24	TSTOUT15	Default: Normal operation	SRAM Test Mode 0 - Enable test mode 1 - Normal operation
AD29, AG28, AJ26, AE26, AJ24, AE23, AJ22, AJ20, AE20, AJ18, AJ21, AJ16, AJ14, AE14, AJ12, AE11, AJ10, AJ8, AE8, AJ6, AE5, AJ4, AG1, AE1,	M[23:0]_TXEN	Default: RMII	0 – GPSI 1 - RMII
C21	P_D	Must be pulled-down	Reserved - Must be pulled-down.
C19, B19, A19	OE_CLK[2:0]	Default: 111	Programmable delay for internal OE_CLK from SCLK input. The OE_CLK is used for generating the OE0 and OE1 signals Suggested value is 001.

Ball No(s)	Symbol	I/O	Description
C20, B20, A20	LA_CLK[2:0]	Default: 111	Programmable delay for LA_CLK from internal OE_CLK. The LA_CLK delay from SCLK is the sum of the delay programmed in here and the delay in P_D[15:13]. Suggested value is 011.
B22, A22, C23, B23, A23, C24	MIRROR[5:0]	Default: 111111	Dedicated Port Mirror Mode. The first 5 bits select the port to be mirrored. The last bit selects either ingress or egress data.

Note:

= Active low signal

Input = Input signal

In-ST = Input signal with Schmitt-Trigger

Output = Output signal (Tri-State driver)

Out-OD = Output signal with Open-Drain driver

I/O-TS = Input & Output signal with Tri-State driver

I/O-OD = Input & Output signal with Open-Drain driver

13.3 Ball – Signal Name

Ball No.	Signal Name	Ball No.	Signal Name	Ball No.	Signal Name
D20	LA_D[63]	D3	LA_D[19]	A9	LA_OE0#
B21	LA_D[62]	E3	LA_D[18]	B9	LA_OE1#
D19	LA_D[61]	D2	LA_D[17]	F4	RESERVED
E19	LA_D[60]	E2	LA_D[16]	F5	RESERVED
D18	LA_D[59]	A7	LA_D[15]	G4	RESERVED
E18	LA_D[58]	B7	LA_D[14]	G5	RESERVED
D17	LA_D[57]	A6	LA_D[13]	H4	RESERVED
E17	LA_D[56]	B6	LA_D[12]	H5	RESERVED
D16	LA_D[55]	C6	LA_D[11]	J4	RESERVED
E16	LA_D[54]	A5	LA_D[10]	J5	RESERVED
D15	LA_D[53]	B5	LA_D[9]	K4	RESERVED
E15	LA_D[52]	C5	LA_D[8]	K5	RESERVED
D14	LA_D[51]	A4	LA_D[7]	L4	RESERVED
E14	LA_D[50]	B4	LA_D[6]	L5	RESERVED
D13	LA_D[49]	C4	LA_D[5]	M4	RESERVED
E13	LA_D[48]	A3	LA_D[4]	M5	RESERVED
D21	LA_D[47]	B3	LA_D[3]	N4	RESERVED
E21	LA_D[46]	C3	LA_D[2]	N5	RESERVED
A18	LA_D[45]	B2	LA_D[1]	G3	RESERVED
B18	LA_D[44]	C2	LA_D[0]	H1	RESERVED
C18	LA_D[43]	C14	LA_A[20]	H2	RESERVED
A17	LA_D[42]	A13	LA_A[19]	H3	RESERVED
B17	LA_D[41]	B13	LA_A[18]	J1	RESERVED
C17	LA_D[40]	C13	LA_A[17]	J2	RESERVED
A16	LA_D[39]	A12	LA_A[16]	J3	RESERVED
B16	LA_D[38]	B12	LA_A[15]	K1	RESERVED
C16	LA_D[37]	C12	LA_A[14]	K2	RESERVED
A15	LA_D[36]	A11	LA_A[13]	K3	RESERVED
B15	LA_D[35]	B11	LA_A[12]	L1	RESERVED
C15	LA_D[34]	C11	LA_A[11]	L2	RESERVED
A14	LA_D[33]	D11	LA_A[10]	L3	RESERVED
B14	LA_D[32]	E11	LA_A[9]	M1	RESERVED
D9	LA_D[31]	A10	LA_A[8]	M2	RESERVED

Ball No.	Signal Name	Ball No.	Signal Name	Ball No.	Signal Name
E9	LA_D[30]	B10	LA_A[7]	M3	RESERVED
D8	LA_D[29]	D10	LA_A[6]	U4	RESERVED
E8	LA_D[28]	E10	LA_A[5]	U5	RESERVED
D7	LA_D[27]	A8	LA_A[4]	V4	RESERVED
E7	LA_D[26]	C7	LA_A[3]	V5	RESERVED
D6	LA_D[25]	B8	LA_DSC#	W4	RESERVED
E6	LA_D[24]	C1	LA_CLK	W5	RESERVED
D5	LA_D[23]	C9	LA_WE#	Y4	RESERVED
E5	LA_D[22]	D12	LA_WE0#	Y5	RESERVED
D4	LA_D[21]	E12	LA_WE1#	AA4	RESERVED
E4	LA_D[20]	C8	LA_OE#	AA5	RESERVED
AB4	RESERVED	U2	RESERVED	AH7	M[4]_RXD[0]
AB5	RESERVED	R28	MDC	AE6	M[3]_RXD[0]
AC4	RESERVED	P28	MDIO	AH5	M[2]_RXD[0]
AC5	RESERVED	R29	M_CLK	AH2	M[1]_RXD[0]
AD4	RESERVED	AC29	M[23]_RXD[1]	AF2	M[0]_RXD[0]
AD5	RESERVED	AE28	M[22]_RXD[1]	AC27	M[23]_CRS_DV
W1	RESERVED	AJ27	M[21]_RXD[1]	AF29	M[22]_CRS_DV
Y1	RESERVED	AF27	M[20]_RXD[1]	AG27	M[21]_CRS_DV
Y2	RESERVED	AJ25	M[19]_RXD[1]	AF26	M[20]_CRS_DV
Y3	RESERVED	AF24	M[18]_RXD[1]	AG25	M[19]_CRS_DV
AA1	RESERVED	AH23	M[17]_RXD[1]	AG23	M[18]_CRS_DV
AA2	RESERVED	AE19	M[16]_RXD[1]	AF23	M[17]_CRS_DV
AA3	RESERVED	AF21	M[15]_RXD[1]	AG21	M[16]_CRS_DV
AB1	RESERVED	AJ19	M[14]_RXD[1]	AH21	M[15]_CRS_DV
AB2	RESERVED	AF18	M[13]_RXD[1]	AF19	M[14]_CRS_DV
AB3	RESERVED	AJ17	M[12]_RXD[1]	AF17	M[13]_CRS_DV
AC1	RESERVED	AJ15	M[11]_RXD[1]	AG17	M[12]_CRS_DV
AC2	RESERVED	AF15	M[10]_RXD[1]	AG15	M[11]_CRS_DV
AC3	RESERVED	AJ13	M[9]_RXD[1]	AF14	M[10]_CRS_DV
AD1	RESERVED	AF12	M[8]_RXD[1]	AG13	M[9]_CRS_DV
AD2	RESERVED	AJ11	M[7]_RXD[1]	AF11	M[8]_CRS_DV
AD3	RESERVED	AJ9	M[6]_RXD[1]	AG11	M[7]_CRS_DV
N3	RESERVED	AF9	M[5]_RXD[1]	AG9	M[6]_CRS_DV
N2	RESERVED	AJ7	M[4]_RXD[1]	AF8	M[5]_CRS_DV

Ball No.	Signal Name	Ball No.	Signal Name	Ball No.	Signal Name
N1	RESERVED	AF6	M[3]_RXD[1]	AG7	M[4]_CRS_DV
P3	RESERVED	AJ5	M[2]_RXD[1]	AF5	M[3]_CRS_DV
P2	RESERVED	AJ3	M[1]_RXD[1]	AG5	M[2]_CRS_DV
P1	RESERVED	AF1	M[0]_RXD[1]	AH3	M[1]_CRS_DV
R5	RESERVED	AC28	M[23]_RXD[0]	AF3	M[0]_CRS_DV
R4	RESERVED	AF28	M[22]_RXD[0]	AD29	M[23]_TXEN
R3	RESERVED	AH27	M[21]_RXD[0]	AG28	M[22]_TXEN
R2	RESERVED	AE27	M[20]_RXD[0]	AJ26	M[21]_TXEN
R1	RESERVED	AH25	M[19]_RXD[0]	AE26	M[20]_TXEN
T5	RESERVED	AE24	M[18]_RXD[0]	AJ24	M[19]_TXEN
T4	RESERVED	AF22	M[17]_RXD[0]	AE23	M[18]_TXEN
T3	RESERVED	AF20	M[16]_RXD[0]	AJ22	M[17]_TXEN
T2	RESERVED	AE21	M[15]_RXD[0]	AJ20	M[16]_TXEN
T1	RESERVED	AH19	M[14]_RXD[0]	AE20	M[15]_TXEN
W3	RESERVED	AH20	M[13]_RXD[0]	AJ18	M[14]_TXEN
W2	RESERVED	AH17	M[12]_RXD[0]	AJ21	M[13]_TXEN
V1	RESERVED	AH15	M[11]_RXD[0]	AJ16	M[12]_TXEN
G1	RESERVED	AE15	M[10]_RXD[0]	AJ14	M[11]_TXEN
V3	RESERVED	AH13	M[9]_RXD[0]	AE14	M[10]_TXEN
P4	RESERVED	AE12	M[8]_RXD[0]	AJ12	M[9]_TXEN
P5	RESERVED	AH11	M[7]_RXD[0]	AE11	M[8]_TXEN
V2	RESERVED	AH9	M[6]_RXD[0]	AJ10	M[7]_TXEN
U1	RESERVED	AE9	M[5]_RXD[0]	AJ8	M[6]_TXEN
AE8	M[5]_TXEN	AH8	M[6]_TXD[0]	G27	RESERVED
AJ6	M[4]_TXEN	AF7	M[5]_TXD[0]	H29	RESERVED
AE5	M[3]_TXEN	AH6	M[4]_TXD[0]	H28	RESERVED
AJ4	M[2]_TXEN	AF4	M[3]_TXD[0]	H27	RESERVED
AG1	M[1]_TXEN	AH4	M[2]_TXD[0]	J29	RESERVED
AE1	M[0]_TXEN	AG2	M[1]_TXD[0]	J28	RESERVED
AD27	M[23]_TXD[1]	AE2	M[0]_TXD[0]	J27	RESERVED
AH28	M[22]_TXD[1]	U26	RESERVED	K29	RESERVED
AG26	M[21]_TXD[1]	U25	RESERVED	K28	RESERVED
AE25	M[20]_TXD[1]	V26	RESERVED	K27	RESERVED
AG24	M[19]_TXD[1]	V25	RESERVED	L29	RESERVED
AE22	M[18]_TXD[1]	W26	RESERVED	L28	RESERVED

Ball No.	Signal Name	Ball No.	Signal Name	Ball No.	Signal Name
AJ23	M[17]_TXD[1]	W25	RESERVED	L27	RESERVED
AG20	M[16]_TXD[1]	Y27	RESERVED	M29	RESERVED
AE18	M[15]_TXD[1]	Y26	RESERVED	M28	RESERVED
AG18	M[14]_TXD[1]	AA26	RESERVED	M27	RESERVED
AE16	M[13]_TXD[1]	AA25	RESERVED	G26	RESERVED
AG16	M[12]_TXD[1]	AB26	RESERVED	G25	RESERVED
AG14	M[11]_TXD[1]	AB25	RESERVED	H26	RESERVED
AE13	M[10]_TXD[1]	AC26	RESERVED	H25	RESERVED
AG12	M[9]_TXD[1]	AC25	RESERVED	J26	RESERVED
AE10	M[8]_TXD[1]	AD26	RESERVED	J25	RESERVED
AG10	M[7]_TXD[1]	AD25	RESERVED	K25	RESERVED
AG8	M[6]_TXD[1]	U27	RESERVED	K26	RESERVED
AE7	M[5]_TXD[1]	V29	RESERVED	M25	RESERVED
AG6	M[4]_TXD[1]	V28	RESERVED	L26	RESERVED
AE4	M[3]_TXD[1]	V27	RESERVED	M26	RESERVED
AG4	M[2]_TXD[1]	W29	RESERVED	L25	RESERVED
AG3	M[1]_TXD[1]	W28	RESERVED	N26	RESERVED
AE3	M[0]_TXD[1]	W27	RESERVED	N25	RESERVED
AD28	M[23]_TXD[0]	Y29	RESERVED	P26	RESERVED
AG29	M[22]_TXD[0]	Y28	RESERVED	P25	RESERVED
AH26	M[21]_TXD[0]	Y25	RESERVED	F28	RESERVED
AF25	M[20]_TXD[0]	AA29	RESERVED	G28	RESERVED
AH24	M[19]_TXD[0]	AA28	RESERVED	E25	RESERVED
AG22	M[18]_TXD[0]	AA27	RESERVED	G29	RESERVED
AH22	M[17]_TXD[0]	AB29	RESERVED	F29	RESERVED
AE17	M[16]_TXD[0]	AB28	RESERVED	F26	RESERVED
AG19	M[15]_TXD[0]	AB27	RESERVED	E26	RESERVED
AH18	M[14]_TXD[0]	R26	RESERVED	F25	RESERVED
AF16	M[13]_TXD[0]	T25	RESERVED	E24	BIST_DONE/TSTOUT[15]
AH16	M[12]_TXD[0]	T26	RESERVED	D24	BIST_IN_PRC/TSTOUT[14]
AH14	M[11]_TXD[0]	T28	RESERVED	D25	MCT_ERR/TSTOUT[13]
AF13	M[10]_TXD[0]	U28	RESERVED	D26	FCB_ERR/TSTOUT[12]
AH12	M[9]_TXD[0]	R25	RESERVED	C26	CHECKSUM_OK/TSTOUT[11]
AF10	M[8]_TXD[0]	U29	RESERVED	D27	INIT_START/TSTOUT[10]

Ball No.	Signal Name	Ball No.	Signal Name	Ball No.	Signal Name
AH10	M[7]_TXD[0]	T29	RESERVED	C27	INIT_DONE/TSTOUT[9]
B27	TSTOUT[8]	U18	VSS	N12	VSS
A27	TSTOUT[7]	V12	VSS	N13	VSS
E28	TSTOUT[6]	V13	VSS	K17	VDD
D28	TSTOUT[5]	V14	VSS	K18	VDD
C28	TSTOUT[4]	V15	VSS	M10	VDD
B28	TSTOUT[3]	V16	VSS	N10	VDD
E29	LED_BIT/TSTOUT[2]	V17	VSS	M20	VDD
D29	LED_SYN/TSTOUT[1]	V18	VSS	N20	VDD
C29	LED_CLK/TSTOUT[0]	N14	VSS	U10	VDD
N29	RESERVED	N15	VSS	V10	VDD
P29	RESERVED	N16	VSS	U20	VDD
F3	SCAN_EN	N17	VSS	V20	VDD
E1	SCLK	N18	VSS	Y12	VDD
U3	T_MODE0	P12	VSS	Y13	VDD
C10	T_MODE1	P13	VSS	Y17	VDD
B24	RESERVED	P14	VSS	Y18	VDD
A21	TRUNK1	P15	VSS	K12	VDD
C22	TRUNK0	P16	VSS	K13	VDD
A26	STROBE	C19	OE_CLK2	M16	VSS
B26	D0	B19	OE_CLK1	M17	VSS
C25	AUTOFD	A19	OE_CLK0	M18	VSS
A24	SCL	R13	VSS	F16	VCC
A25	SDA	R14	VSS	F17	VCC
F1	AVCC	R15	VSS	N6	VCC
D1	AGND	R16	VSS	P6	VCC
D22	SCANCOL	R17	VSS	R6	VCC
E23	SCANLINK	R18	VSS	T6	VCC
E27	SCANMODE	T12	VSS	U6	VCC
N28		T13	VSS	N24	VCC
N27		T14	VSS	P24	VCC
F2	RESIN#	T15	VSS	R24	VCC
G2	RESETOUT#	T16	VSS	T24	VCC

Ball No.	Signal Name	Ball No.	Signal Name	Ball No.	Signal Name
B22	MIRROR5	T17	VSS	U24	VCC
A22	MIRROR4	T18	VSS	AD13	VCC
C23	MIRROR3	U12	VSS	AD14	VCC
B23	MIRROR2	U13	VSS	AD15	VCC
A23	MIRROR1	U14	VSS	AD16	VCC
C24	MIRROR0	U15	VSS	AD17	VCC
D23	SCANCLK	U16	VSS	F13	VCC
T27	RESERVED	U17	VSS	F14	VCC
F27	RESERVED	M12	VSS	F15	VCC
C20	LA_CLK2	M13	VSS		
B20	LA_CLK1	M14	VSS		
A20	LA_CLK0	M15	VSS		
C21	P_D	P17	VSS		
E20	RESERVED	P18	VSS		
B25	RESERVED	R12	VSS		

13.4 AC/DC Timing

13.4.1 Absolute Maximum Ratings

Storage Temperature	-65°C to +150°C
Operating Temperature	-40°C to +85°C
Maximum Junction Temperature	+125°C
Supply Voltage VCC with Respect to VSS	+3.0 V to +3.6 V
Supply Voltage VDD with Respect to VSS	+2.38 V to +2.75 V
Voltage on Input Pins	-0.5 V to (VCC + 3.3 V)

Caution: Stress above those listed may damage the device. Exposure to the Absolute Maximum Ratings for extended periods may affect device reliability. Functionality at or above these limits is not implied.

13.4.2 DC Electrical Characteristics

VCC = 3.0 V to 3.6 V (3.3v +/- 10%) T_{AMBIENT} = -40°C to +85°C

VDD = 2.5V +10% - 5%

13.4.3 Recommended Operating Conditions

Symbol	Parameter Description	Min.	Typ.	Max.	Unit
f_{osc}	Frequency of Operation		100		MHz
I_{CC}	Supply Current – @ 100 MHz ($V_{CC}=3.3$ V)			450	mA
I_{DD}	Supply Current – @ 100 MHz ($V_{DD}=2.5$ V)			1500	mA
V_{OH}	Output High Voltage (CMOS)	2.4			V
V_{OL}	Output Low Voltage (CMOS)			0.4	V
V_{IH-TTL}	Input High Voltage (TTL 5 V tolerant)	2.0		$V_{CC} + 2.0$	V
V_{IL-TTL}	Input Low Voltage (TTL 5 V tolerant)			0.8	V
I_{IL}	Input Leakage Current (0.1 V < V_{IN} < V_{CC})(all pins except those with internal pull-up/pull-down resistors)			10	μ A
I_{OL}	Output Leakage Current (0.1 V < V_{OUT} < V_{CC})			10	μ A
C_{IN}	Input Capacitance			5	pF
C_{OUT}	Output Capacitance			5	pF
$C_{I/O}$	I/O Capacitance			7	pF
θ_{ja}	Thermal resistance with 0 air flow			11.2	C/W
θ_{ja}	Thermal resistance with 1 m/s air flow			10.2	C/W
θ_{ja}	Thermal resistance with 2 m/s air flow			8.9	C/W
θ_{jc}	Thermal resistance between junction and case			3.1	C/W
θ_{jb}	Thermal resistance between junction and board			6.6	C/W

13.4.4 Typical Reset & Bootstrap Timing Diagram

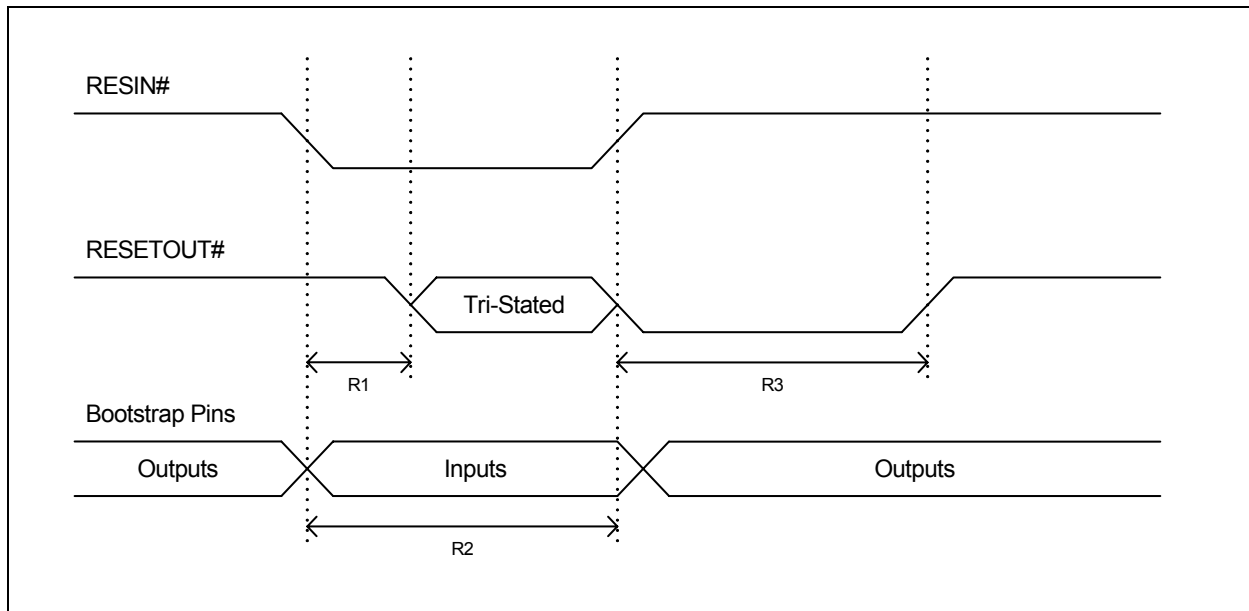


Figure 13 - Typical Reset & Bootstrap Timing Diagram

Symbol	Parameter	Min.	Typ.	Note
R1	Delay until RESETOUT# is tri-stated		10ns	RESETOUT# state is then determined by the external pull-up/down resistor
R2	Bootstrap stabilization	1 μ s	10 μ s	Bootstrap pins sampled on rising edge of RESIN# ^a
R3	RESETOUT# assertion		2 ms	

Table 10 - Reset & Bootstrap Timing

a. The TSTOUT[8:0] pins will switch over to the LED interface functionality in 3 SCLK cycles after RESIN# goes high

13.5 Local Frame Buffer SBRAM Memory Interface

13.5.1 Local SBRAM Memory Interface




Figure 14 - Local Memory Interface – Input Setup and Hold Timing

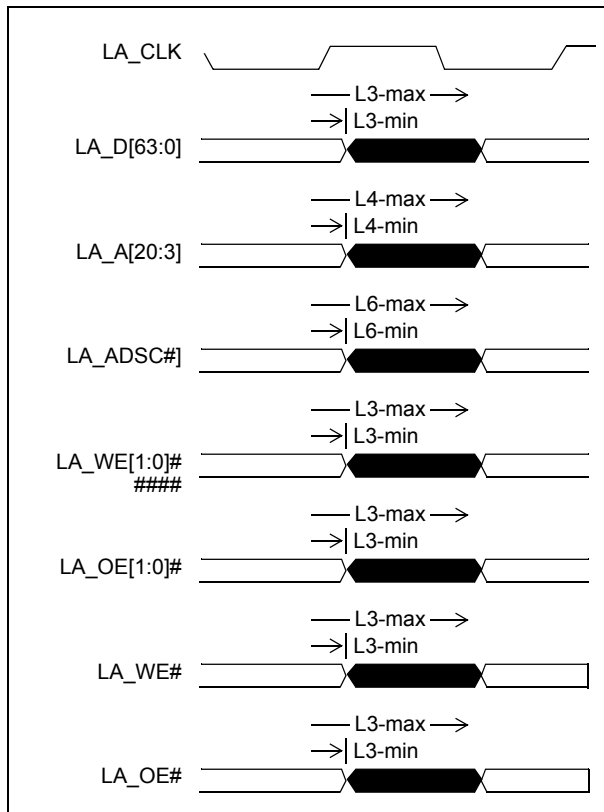



Figure 15 - Local Memory Interface - Output Valid Delay Timing

AC Characteristics – Local frame buffer SBRAM Memory Interface

Symbol	Parameter	-100 MHz		Note
		Min. (ns)	Max. (ns)	
L1	LA_D[63:0] input set-up time	4		
L2	LA_D[63:0] input hold time	1.5		
L3	LA_D[63:0] output valid delay	1.5	7	$C_L = 25$ pf
L4	LA_A[20:3] output valid delay	2	7	$C_L = 30$ pf
L6	LA_ADSC# output valid delay	1	7	$C_L = 30$ pf
L7	LA_WE[1:0]#output valid delay	1	7	$C_L = 25$ pf
L8	LA_OE[1:0]# output valid delay	-1	1	$C_L = 25$ pf
L9	LA_WE# output valid delay	1	7	$C_L = 25$ pf
L10	LA_OE# output valid delay	1	5	$C_L = 25$ pf

13.6 AC Characteristics**13.6.1 Reduced Media Independent Interface**

Symbol	Parameter	-50 MHz		Note
		Min. (ns)	Max. (ns)	
M2	M[23:0]_RXD[1:0] Input Setup Time	4		
M3	M[23:0]_RXD[1:0] Input Hold Time	1		
M4	M[23:0]_CRS_DV Input Setup Time	4		
M5	M[23:0]_CRS_DV Input Hold Time	1		
M6	M[23:0]_TXEN Output Delay Time	2	11	$C_L = 20$ pF
M7	M[23:0]_TXD[1:0] Output Delay Time	2	11	$C_L = 20$ pF

Table 11 - AC Characteristics – Reduced Media Independent Interface**Figure 16 - AC Characteristics – Reduced Media Independent Interface**




Figure 17 - AC Characteristics – Reduced Media Independent Interface

13.6.2 LED Interface




Figure 18 - AC Characteristics – LED Interface

Symbol	Parameter	Variable FREQ.		Note
		Min. (ns)	Max. (ns)	
LE5	LED_SYN Output Valid Delay	-1	7	C _L = 30 pf
LE6	LED_BIT Output Valid Delay	-1	7	C _L = 30 pf

Table 12 - AC Characteristics – LED Interface

13.6.3 SCANLINK SCANCOL Output Delay Timing




Figure 19 - SCANLINK SCANCOL Output Delay Timing




Figure 20 - SCANLINK, SCANCOL Setup Timing

Symbol	Parameter	-25 MHz		Note
		Min. (ns)	Max (ns)	
C1	SCANLINK input set-up time	20		
C2	SCANLINK input hold time	2		
C3	SCANCOL input setup time	20		
C4	SCANCOL input hold time	1		
C5	SCANLINK output valid delay	0	10	$C_L = 30$ pf
C7	SCANCOL output valid delay	0	10	$C_L = 30$ pf

Table 13 - SCANLINK, SCANCOL Timing

13.6.4 MDIO Input Setup and Hold Timing




Figure 21 - MDIO Input Setup and Hold Timing




Figure 22 - MDIO Output Delay Timing

Symbol	Parameter	1 MHz		Note
		Min. (ns)	Max. (ns)	
D1	MDIO input setup time	10		
D2	MDIO input hold time	2		
D3	MDIO output delay time	1	20	$C_L = 50$ pf

Table 14 - MDIO Timing

13.6.5 I²C Input Setup Timing




Figure 23 - I²C Input Setup Timing




Figure 24 - I²C Output Delay Timing

Symbol	Parameter	50 KHz		Note
		Min. (ns)	Max. (ns)	
S1	SDA input setup time	20		
S2	SDA input hold time	1		
S3*	SDA output delay time	4 usec	6 usec	C _L = 30 pf
* Open Drain Output. Low to High transistor is controlled by external pullup resistor.				

Table 15 - I²C Timing

13.6.6 Serial Interface Setup Timing




Figure 25 - Serial Interface Setup Timing







Figure 26 - Serial Interface Output Delay Timing

Symbol	Parameter	Min. (ns)	Max. (ns)	Note
D1	D0 setup time	20		
D2	D0 hold time	3 μ s		
D3	AutoFd output delay time	1	50	$C_L = 100$ pf
D4	Strobe low time	5 μ s		
D5	Strobe high time	5 μ s		

Table 16 - Serial Interface Timing



DIMENSION	MIN	MAX
A	2.20	2.46
A1	0.50	0.70
A2	1.17 REF	
D	37.30	37.70
D1	34.50 REF	
E	37.30	37.70
E1	34.50 REF	
b	0.60	0.90
e	1.27	
N	553	
Conforms to JEDEC MS - 034		



NOTE:

1. CONTROLLING DIMENSIONS ARE IN MM
2. DIMENSION "b" IS MEASURED AT THE MAXIMUM SOLDER BALL DIAMETER
3. SEATING PLANE IS DEFINED BY THE SPHERICAL CROWNS OF THE SOLDER BALLS.
4. N IS THE NUMBER OF SOLDER BALLS
5. NOT TO SCALE.
6. SUBSTRATE THICKNESS IS 0.56 MM

© Zarlink Semiconductor 2002 All rights reserved.

ISSUE	1			
ACN	213932			
DATE	20Jan03			
APPRD.				



Previous package codes:

BH / G

Package Code GK

Package Outline for 553 Ball HSBGA (37.5x37.5x2.33mm)

GPD00818



**For more information about all Zarlink products
visit our Web Site at
www.zarlink.com**

Information relating to products and services furnished herein by Zarlink Semiconductor Inc. or its subsidiaries (collectively "Zarlink") is believed to be reliable. However, Zarlink assumes no liability for errors that may appear in this publication, or for liability otherwise arising from the application or use of any such information, product or service or for any infringement of patents or other intellectual property rights owned by third parties which may result from such application or use. Neither the supply of such information or purchase of product or service conveys any license, either express or implied, under patents or other intellectual property rights owned by Zarlink or licensed from third parties by Zarlink, whatsoever. Purchasers of products are also hereby notified that the use of product in certain ways or in combination with Zarlink, or non-Zarlink furnished goods or services may infringe patents or other intellectual property rights owned by Zarlink.

This publication is issued to provide information only and (unless agreed by Zarlink in writing) may not be used, applied or reproduced for any purpose nor form part of any order or contract nor to be regarded as a representation relating to the products or services concerned. The products, their specifications, services and other information appearing in this publication are subject to change by Zarlink without notice. No warranty or guarantee express or implied is made regarding the capability, performance or suitability of any product or service. Information concerning possible methods of use is provided as a guide only and does not constitute any guarantee that such methods of use will be satisfactory in a specific piece of equipment. It is the user's responsibility to fully determine the performance and suitability of any equipment using such information and to ensure that any publication or data used is up to date and has not been superseded. Manufacturing does not necessarily include testing of all functions or parameters. These products are not suitable for use in any medical products whose failure to perform may result in significant injury or death to the user. All products and materials are sold and services provided subject to Zarlink's conditions of sale which are available on request.

Purchase of Zarlink's I²C components conveys a licence under the Philips I²C Patent rights to use these components in and I²C System, provided that the system conforms to the I²C Standard Specification as defined by Philips.

Zarlink, ZL and the Zarlink Semiconductor logo are trademarks of Zarlink Semiconductor Inc.

Copyright Zarlink Semiconductor Inc. All Rights Reserved.

TECHNICAL DOCUMENTATION - NOT FOR RESALE
