

Features

- Integrated Single-Chip 10/100/1000M Ethernet Switch
 - Twenty-four 10/100 Mbps auto-negotiating Fast Ethernet (FE) ports with RMII or GPSI (7WS) interface options per port
 - Two 10/100/1000M auto-negotiating Gigabit Ethernet (GE) ports with GMII, TBI, and MII interface options per port
- Operates stand-alone or can be cascaded
 - Stacking port supports hot swap in managed configuration
 - Stacking port supports 2G mode for inter-chip communications
- Supports two Frame Data Buffer (FDB) memory domains (2 MB or 4 MB) with pipelined, sync-burst SRAM at 100 MHz
 - Applies centralized shared memory architecture
- L2 Switching
 - MAC address self learning, up to 64K MAC addresses
 - Supports port-based VLAN
- High performance packet classification and switching at full-wire speed
- CPU access supports the following interface options:
 - Serial interface in unmanaged mode, with optional

Ordering Information

MVTX2603AG 553 Pin HSBGA
MVTX2603AG2 553 Pin HSBGA**

**Pb Free Tin/Silver/Copper

-40°C to 85°C

- I²C EEPROM support
- 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
- QoS Support
 - 4 transmission priorities for Fast Ethernet ports and 8 transmission classes for Gigabit ports
 - Per-queue weighted random early discard (WRED) with 2 drop precedence levels
 - Scheduling using delay bounded (DB), strict priority (SP), and Weighted Fair Queuing (WFQ) disciplines
 - User controlled WRED thresholds
 - Buffer management: per-class, shared, and per-port buffer reservations
- Classification based on:
 - Port-based priority: priority in a frame can be overwritten by the priority of port

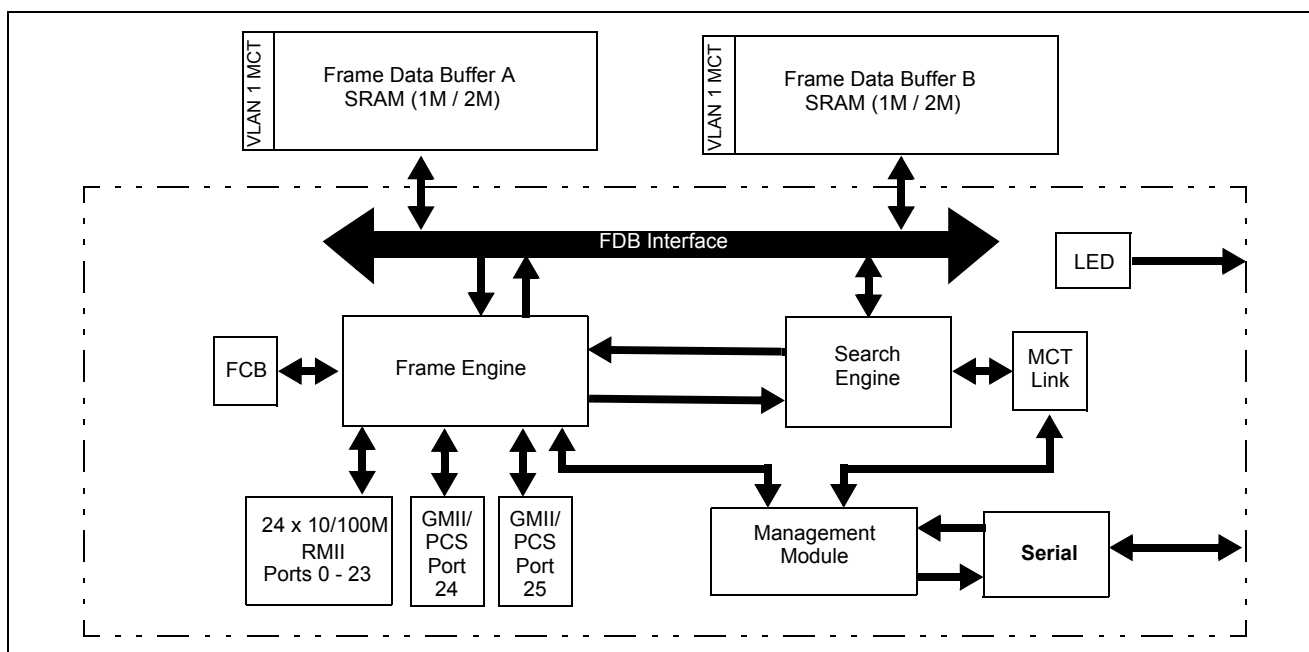


Figure 1 - System Block Diagram

- VLAN Priority field in VLAN tagged frame (IEEE 802.1p)
- DS/TOS field in IP packet
- UDP/TCP logical ports: 8 hard-wired and 8 programmable ports, including one programmable range
- The drop precedence of the above classifications is programmable
- Supports IEEE 802.3ad link aggregation
 - 3 port trunking groups
 - one group for the 2 Gigabit ports
 - two groups for 10/100 ports, with up to 4 ports per group
 - Load sharing among trunked ports can be based on:
 - Source and/or destination MAC address
 - Source port (Gigabit ports only)
- Port Mirroring
 - supports either a dedicated mirroring port or port 23 in unmanaged mode
- Full Duplex Ethernet IEEE 802.3x Flow Control
- Backpressure flow control for Half Duplex ports
- Full set of LED signals provided by a serial interface, or 6 LED signals dedicated to Gigabit port status only (without serial interface)
- Built-in reset logic triggered by system malfunction
- Built-in self test (BIST) for internal and external SRAM

Description

The MVTX2603 is a high density, low cost, high performance, non-blocking Ethernet switch chip. A single chip provides 24 ports at 10/100 Mbps, 2 ports at 1000 Mbps and a CPU interface for unmanaged switch applications. The Gigabit ports can also support 10/100 M and 2 G modes.

The chip supports up to 64K MAC addresses and port-based Virtual LANs (VLANs). The centralized shared memory architecture permits a very high performance packet forwarding rate at full wire speed. The chip is optimized to provide low-cost, high-performance workgroup switching.

Two 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. In the 2G-mode configuration, two ZBT-SRAM domains are needed.

With delay bounded, strict priority, and/or WFQ transmission scheduling and WRED dropping schemes, the MVTX2603 provides powerful QoS functions for various multimedia and mission-critical applications. The chip provides 4 transmission priorities (8 priorities per Gigabit port) and 2 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 MVTX2603 recognizes a total of 16 UDP/TCP logical ports, 8 hard-wired and 8 programmable (including one programmable range).

The MVTX2603 supports 3 groups of port trunking/load sharing. One group is dedicated to the two Gigabit ports, and the other two groups to 10/100 ports, where 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 MVTX2603 also supports a per-system option to enable flow control for best effort frames, even on QoS-enabled ports.

The Physical Coding Sublayer (PCS) is integrated on-chip to provide a direct 10-bit interface (TBI) for connection to SERDES chips. The PCS can be bypassed to provide a GMII interface.

The MVTX2603 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 MVTX2603 is packaged in a 553-pin Ball Grid Array package.

Changes Summary

The following table captures the changes from the April 2006 issue.

Revision Date	Summary of Changes
July 2006	Updated G3 "Gn_RXDV Input Setup Times" min. from 2 ns to 1.2 ns.

The April 2006 issue is the starting point for the change summary section.

Revision Date	Summary of Changes
April 2006	<ul style="list-style-type: none">- Added Pb-free order code (MVTX2603AG2)- Corrected ECR1Pn default value (should be 0xC0)- Corrected PR100 default value (should be 0x35)- Corrected SFCB default value (should be 0x46)

Table of Contents

1.0 BGA and Ball Signal Descriptions	10
1.1 BGA Views (Top-View)	10
1.1.1 Encapsulated view in unmanaged mode	10
1.2 Ball – Signal Descriptions	11
1.3 Ball – Signal Name	21
1.4 Signal Mapping and Internal Pull Up/Down Configuration	27
2.0 Block Functionality	29
2.1 Frame Data Buffer (FDB) Interfaces	29
2.2 MAC Modules	29
2.2.1 RMII MAC Module (RMAC)	29
2.2.1.1 GPSI Interface	29
2.2.1.2 SCANLINK and SCANCOL interface	30
2.2.2 GMII MAC Module (GMAC)	31
2.2.2.1 Physical Coding Sublayer (PCS) Module	31
2.2.2.2 TBI Interface	32
2.2.3 PHY Addresses	32
2.3 Frame Engine	33
2.4 Search Engine	33
2.5 LED Interface	33
2.5.1 Port Status	33
2.5.2 LED Interface Timing Diagram	34
2.6 Internal Memory	34
2.7 Timeout Reset Monitor	35
3.0 System Configuration (Stand-alone and Stacking)	35
3.1 Management and Configuration	35
3.1.1 I2C Interface	35
3.1.1.1 Start Condition	35
3.1.1.2 Address	35
3.1.1.3 Data Direction	35
3.1.1.4 Acknowledgment	35
3.1.1.5 Data	36
3.1.1.6 Stop Condition	36
3.1.2 Synchronous Serial Interface	36
3.1.2.1 Write Command	36
3.1.2.2 Read Command	37
3.2 Stacking	37
4.0 Data Forwarding Protocol	38
4.1 Unicast Data Frame Forwarding	38
4.2 Multicast Data Frame Forwarding	38
5.0 Memory Interface	39
5.1 Overview	39
5.2 Memory Requirements	40
5.2.1 ZBT Support	40
5.3 Memory Configurations	40
6.0 Search Engine	44
6.1 Search Engine Overview	44
6.2 Basic Flow	44
6.3 Search, Learning, and Aging	45
6.3.1 MAC Search	45
6.3.2 Learning	45
6.3.3 Aging	45

Table of Contents

6.4 Port-Based VLAN	45
6.5 Quality of Service	46
6.5.1 Priority Classification Rule	47
7.0 Frame Engine	47
7.1 Data Forwarding Summary	47
7.2 Frame Engine Details	48
7.2.1 FCB Manager	48
7.2.2 Rx Interface	48
7.2.3 RxDMA	48
7.2.4 TxQ Manager	48
7.2.5 Port Control	48
7.2.6 TxDMA	48
8.0 Quality of Service and Flow Control	49
8.1 Model	49
8.2 Four QoS Configurations	50
8.3 Delay Bound	51
8.4 Strict Priority and Best Effort	51
8.5 Weighted Fair Queuing	51
8.6 Shaper	52
8.7 Rate Control	52
8.8 WRED Drop Threshold Management Support	53
8.9 Buffer Management	53
8.9.1 Dropping When Buffers Are Scarce	55
8.10 Flow Control Basics	55
8.10.1 Unicast Flow Control	55
8.10.2 Multicast Flow Control	56
8.11 Mapping to IETF DiffServ Classes	56
9.0 Port Trunking	57
9.1 Features and Restrictions	57
9.2 Unicast Packet Forwarding	57
9.3 Multicast Packet Forwarding	58
9.4 Unmanaged Trunking	58
10.0 Port Mirroring	59
10.1 Port Mirroring Features	59
10.2 Setting Registers for Port Mirroring	59
11.0 Register Definition	60
11.1 Register Description	60
11.2 Indirectly Accessed Registers	64
11.2.1 (Group 0 Address) MAC Ports Group	64
11.2.1.1 ECR1Pn: Port n Control Register 1	64
11.2.1.2 ECR2Pn: Port n Control Register 2	64
11.2.1.3 GGControl – Extra GIGA Port Control	66
11.2.2 (Group 1 Address) VLAN Group	67
11.2.2.1 AVTCL – VLAN Type Code Register Low	67
11.2.2.2 AVTCH – VLAN Type Code Register High	67
11.2.2.3 PVMAP00_0 – Port 00 Configuration Register 0	67
11.2.2.4 PVMAP00_1 – Port 00 Configuration Register 1	68
11.2.2.5 PVMAP00_2 – Port 00 Configuration Register 2	68
11.2.2.6 PVMAP00_3 – Port 00 Configuration Register 3	68
11.2.2.7 PVMAPnn_0,1,2,3 – Port nn Configuration Registers	69
11.2.2.8 PVMODE	70

Table of Contents

11.2.3 (Group 2 Address) Port Trunking Groups	71
11.2.3.1 TRUNK0_MODE– Trunk group 0 mode	71
11.2.3.2 TRUNK1_MODE – Trunk group 1 mode	71
11.2.3.3 (Group 3 Address) CPU Port Configuration GroupTX_AGE – Tx Queue Aging timer	72
11.2.4 (Group 4 Address) Search Engine Group	72
11.2.4.1 AGETIME_LOW – MAC address aging time Low	72
11.2.4.2 AGETIME_HIGH –MAC address aging time High	72
11.2.5 (Group 5 Address) Buffer Control/QOS Group	72
11.2.5.1 FCBAT – FCB Aging Timer	72
11.2.5.2 QOSC – QOS Control	73
11.2.5.3 FCR – Flooding Control Register	73
11.2.5.4 AVPML – VLAN Tag Priority Map	74
11.2.5.5 AVPMM – VLAN Priority Map	74
11.2.5.6 AVPMH – VLAN Priority Map	75
11.2.5.7 TOSPML – TOS Priority Map	75
11.2.5.8 TOSPMM – TOS Priority Map	75
11.2.5.9 TOSPMH – TOS Priority Map	76
11.2.5.10 AVDM – VLAN Discard Map	76
11.2.5.11 TOSDML – TOS Discard Map	76
11.2.5.12 BMRC - Broadcast/Multicast Rate Control	77
11.2.5.13 UCC – Unicast Congestion Control	77
11.2.5.14 MCC – Multicast Congestion Control	77
11.2.5.15 PR100 – Port Reservation for 10/100 ports	78
11.2.5.16 PRG – Port Reservation for Giga ports	78
11.2.5.17 SFCB – Share FCB Size	79
11.2.5.18 C2RS – Class 2 Reserve Size	79
11.2.5.19 C3RS – Class 3 Reserve Size	79
11.2.5.20 C4RS – Class 4 Reserve Size	80
11.2.5.21 C5RS – Class 5 Reserve Size	80
11.2.5.22 C6RS – Class 6 Reserve Size	80
11.2.5.23 C7RS – Class 7 Reserve Size	80
11.2.5.24 QOSC00~02 - Classes Byte Limit Set 0	80
11.2.5.25 QOSC03~05 - Classes Byte Limit Set 1	81
11.2.5.26 QOSC12~17 - Classes Byte Limit Giga Port 1	81
11.2.5.27 QOSC18~23 - Classes Byte Limit Giga Port 2	81
11.2.5.28 QOSC40~47 - Classes WFQ Credit Port G1	82
11.2.5.29 QOSC48~55 - Classes WFQ Credit Port G2	82
11.2.5.30 QOSC56~57 - Class 6 Shaper Control Port G1	82
11.2.5.31 QOSC58~59 - Class 6 Shaper Control Port G2	83
11.2.5.32 RDRC0 – WRED Rate Control 0	83
11.2.5.33 RDRC1 – WRED Rate Control 1	83
11.2.5.34 USER_PORT0~7)_L/H – User Define Logical Port (0~7)	84
11.2.5.35 USER_PORT_[1:0]_PRIORITY - User Define Logic Port 1 and 0 Priority	85
11.2.5.36 USER_PORT_[3:2]_PRIORITY - User Define Logic Port 3 and 2 Priority	85
11.2.5.37 USER_PORT_[5:4]_PRIORITY - User Define Logic Port 5 and 4 Priority	85
11.2.5.38 USER_PORT_[7:6]_PRIORITY - User Define Logic Port 7 and 6 Priority	85
11.2.5.39 USER_PORT_ENABLE[7:0] – User Define Logic 7 to 0 Port Enables	85
11.2.5.40 WELL_KNOWN_PORT[1:0]_PRIORITY- Well Known Logic Port 1 and 0 Priority	86
11.2.5.41 WELL_KNOWN_PORT[3:2]_PRIORITY- Well Known Logic Port 3 and 2 Priority	86
11.2.5.42 WELL_KNOWN_PORT [5:4]_PRIORITY- Well Known Logic Port 5 and 4 Priority	86
11.2.5.43 WELL_KNOWN_PORT [7:6]_PRIORITY- Well Known Logic Port 7 and 6 Priority	86
11.2.5.44 WELL_KNOWN_PORT_ENABLE [7:0] – Well Known Logic 7 to 0 Port Enables	87

Table of Contents

11.2.5.45	RLOWL – User Define Range Low Bit 7:0	87
11.2.5.46	RLOWH – User Define Range Low Bit 15:8	87
11.2.5.47	RHIGHL – User Define Range High Bit 7:0	87
11.2.5.48	RHIGHH – User Define Range High Bit 15:8	87
11.2.5.49	RRIORITY – User Define Range Priority	87
11.2.6	(Group 6 Address) MISC Group	88
11.2.6.1	MII_OP0 – MII Register Option 0	88
11.2.6.2	MII_OP1 – MII Register Option 1	88
11.2.6.3	FEN – Feature Register	88
11.2.6.4	MIIC0 – MII Command Register 0	90
11.2.6.5	MIIC1 – MII Command Register 1	90
11.2.6.6	MIIC2 – MII Command Register 2	90
11.2.6.7	MIIC3 – MII Command Register 3	90
11.2.6.8	MIID0 – MII Data Register 0	91
11.2.6.9	MIID1 – MII Data Register 1	91
11.2.6.10	LED Mode – LED Control	91
11.2.6.11	DEVICE Mode	91
11.2.6.12	CHECKSUM - EEPROM Checksum	92
11.2.7	(Group F Address) CPU Access Group	92
11.2.7.1	GCR-Global Control Register	92
11.2.7.2	DCR - Device Status and Signature Register	93
11.2.7.3	DCR1 - Chip Status	94
11.2.7.4	DPST – Device Port Status Register	95
11.2.7.5	DTST – Data read back register	96
11.2.7.6	DA – Dead or Alive Register	96
11.3	TBI Registers	96
11.3.1	Control Register	96
11.3.2	Status Register	97
11.3.3	Advertisement Register	98
11.3.4	Link Partner Ability Register	98
11.3.5	Expansion Register	99
11.3.6	Extended Status Register	99
11.4	Characteristics and Timing	100
11.4.1	Absolute Maximum Ratings	100
11.4.2	DC Electrical Characteristics	100
11.4.3	Recommended Operating Conditions	101
11.5	AC Characteristics and Timing	102
11.5.1	Typical Reset & Bootstrap Timing Diagram	102
11.5.2	Local Frame Buffer SDRAM Memory Interface	103
11.5.2.1	Local SDRAM Memory Interface A	103
11.5.2.2	Local SDRAM Memory Interface B	105
11.5.3	Reduced Media Independent Interface	107
11.5.4	Gigabit Media Independent Interface	108
11.5.5	Ten Bit Interface	109
11.5.6	LED Interface	110
11.5.7	SCANLINK, SCANCOL Interface	111
11.6	MDIO Interface	112
11.6.1	I2C Interface	113
11.6.2	Synchronous Serial Interface	114

List of Figures

Figure 1 - System Block Diagram	1
Figure 2 - GPSI (7WS) Mode Connection Diagram	30
Figure 3 - SCANLINK and SCANCOL Status Diagram	31
Figure 4 - TBI Connection	32
Figure 5 - Timing Diagram of LED Interface	34
Figure 6 - Data Transfer Format for I2C Interface	35
Figure 7 - SRAM Interface Block Diagram (DMAs for 10/100 Ports Only)	39
Figure 8 - Memory Configuration	40
Figure 9 - Memory Configuration For 1 M/bank, 1 Layer	42
Figure 10 - Memory Configuration For 2 M/bank, 2 Layers	42
Figure 11 - Memory Configuration For 2 M/bank, 1 Layer	43
Figure 12 - ZBT Memory Configuration For 2 M/bank, 2 Layers	43
Figure 13 - ZBT Memory Configuration For 2 M/bank, 1 Layer	44
Figure 14 - Priority Classification Rule	47
Figure 15 - Buffer Partition Scheme Used to Implement Buffer Management	54
Figure 16 - Typical Reset & Bootstrap Timing Diagram	102
Figure 17 - Local Memory Interface – Input Setup and Hold Timing	103
Figure 18 - Local Memory Interface - Output Valid Delay Timing	103
Figure 19 - Local Memory Interface – Input Setup and Hold Timing	105
Figure 20 - Local Memory Interface - Output Valid Delay Timing	105
Figure 21 - AC Characteristics – Reduced Media Independent Interface	107
Figure 22 - AC Characteristics – Reduced Media Independent Interface	107
Figure 23 - AC Characteristics- Gigabit Media Independent Interface	108
Figure 24 - AC Characteristics – Gigabit Media Independent Interface	108
Figure 25 - AC Characteristics – Ten Bit Interface (RX)	109
Figure 26 - AC Characteristics –Ten Bit Interface (TX)	109
Figure 27 - AC Characteristics – LED Interface	110
Figure 28 - SCANLINK, SCANCOL Output Delay Timing	111
Figure 29 - SCANLINK, SCANCOL Setup Timing	111
Figure 30 - MDIO Input Setup and Hold Timing	112
Figure 31 - MDIO Output Delay Timing	112
Figure 32 - I2C Input Setup Timing	113
Figure 33 - I2C Output Delay Timing	113
Figure 34 - Serial Interface Setup Timing	114
Figure 35 - Serial Interface Output Delay Timing	114

1.0 BGA and Ball Signal Descriptions

1.1 BGA Views (Top-View)

1.1.1 Encapsulated view in unmanaged mode

	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_D 4	LA_D 7	LA_D 10	LA_D 13	LA_D 15	LA_A 4	LA_O E0#	LA_A 8	LA_A 13	LA_A 16	LA_A 19	LA_D 33	LA_D 36	LA_D 39	LA_D 42	LA_D 45	OE_C LK0	LA_C LK0	TRUN K1	MIRR OR4	MIRR OR1	SCL	SDA	STRO BE	TSTO UT7	A														
B		LA_D 1	LA_D 3	LA_D 6	LA_D 9	LA_D 12	LA_D 14	LA_D DSC#	LA_O E1#	LA_A 7	LA_A 12	LA_A 15	LA_A 18	LA_D 32	LA_D 35	LA_D 38	LA_D 41	LA_D 44	OE_C LK1	LA_C LK1	LA_D 62	MIRR OR5	MIRR OR2	TRUN K2	RSVD	D0	TSTO UT8	TSTO UT3	B													
C	LA_C LK	LA_D 2	LA_D 5	LA_D 8	LA_D 11	LA_A 3	LA_O E#	LA_W E#	T_MO DE1	LA_A 11	LA_A 14	LA_A 17	LA_A 20	LA_D 34	LA_D 37	LA_D 40	LA_D 43	OE_C LK2	LA_C LK2	P_D	TRUN K0	MIRR OR3	MIRR OR0	AUTO FD	TSTO UT11	TSTO UT9	TSTO UT4	TSTO UT0	C													
D	VSSA	LA_D 17	LA_D 19	LA_D 21	LA_D 23	LA_D 25	LA_D 27	LA_D 29	LA_D 31	LA_A 6	LA_A 10	LA_W E0#	LD_D 49	LA_D 51	LA_D 53	LA_D 55	LA_D 57	LA_D 59	LA_D 61	LA_D 63	LA_D 47	SCAN COL	SCAN COL	TSTO UT14	TSTO UT13	TSTO UT12	TSTO UT10	TSTO UT5	TSTO UT1	D												
E	SCLK	LA_D 16	LA_D 18	LA_D 20	LA_D 22	LA_D 24	LA_D 26	LA_D 28	LA_D 30	LA_A 5	LA_A 9	LA_W E1#	LA_D 48	LA_D 50	LA_D 52	LA_D 54	LA_D 56	LA_D 58	LA_D 60	RSVD	LA_D 46	NC	SCAN LNK	TSTO UT15	G1_C RS	G1_T XER	SCAN MD	TSTO UT6	TSTO UT2	E												
F	VDDA	RESI N#	SCAN EN	LB_D 63	LB_D 62														G1_T XCK	G1_T XEN	G1_M TXC	G1_R XDV	G1_R XCK	F																		
G	LB_C LK	RESO UT#	LB_D 47	LB_D 61	LB_D 60														G1_T XD14	G1_T XD15	G1_R XD15	G1_R XER	G1_C OL	G																		
H	LB_D 46	LB_D 45	LB_D 44	LB_D 59	LB_D 58														G1_T XD12	G1_T XD13	G1_R XD12	G1_R XD13	G1_R XD14	H																		
J	LB_D 43	LB_D 42	LB_D 41	LB_D 57	LB_D 56														G1_T XD10	G1_T XD11	G1_R XD9	G1_R XD10	G1_R XD11	J																		
K	LB_D 40	LB_D 39	LB_D 38	LB_D 55	LB_D 54	VDD		VDD															G1_T XD9	G1_T XD8	G1_R XD6	G1_R XD7	G1_R XD8	K														
L	LB_D 37	LB_D 36	LB_D 35	LB_D 53	LB_D 52														G1_T XD4	G1_T XD6	G1_R XD3	G1_R XD4	G1_R XD5	L																		
M	LB_D 34	LB_D 33	LB_D 32	LB_D 51	LB_D 50	VDD		VSS						VDD															G1_T XD7	G1_T XD5	G1_R XD0	G1_R XD1	G1_R XD2	M								
N	LB_A 18	LB_A 19	LB_A 20	LB_D 49	LB_D 48	VCC	VDD		VSS						VDD															VCC	G1_T XD2	G1_T XD3	NC	NC	GREF CLK1	N						
P	LB_A 15	LB_A 16	LB_A 17	LB_W E0#	LB_W E1#	VCC	VSS		VSS						VSS															VCC	G1_T XD0	G1_T XD1	NC	MDIO	GREF CLK0	P						
R	LB_A 10	LB_A 11	LB_A 12	LB_A 13	LB_A 14	VCC	VSS		VSS						VSS															VCC	G0_C RS	G0_T XER	NC	MDC	M_CLK	R						
T	LB_A 5	LB_A 6	LB_A 7	LB_A 8	LB_A 9	VCC	VSS		VSS						VSS															VCC	G0_T XCK	G0_T XEN	G0_M TXC	G0_R XDV	G0_R XCK	T						
U	LB_O E0#	LB_O E1#	T_MO DE0	LB_D 31	LB_D 30	VCC	VDD		VSS						VSS															VDD						VCC	G0_T XD14	G0_T XD15	G0_R XD15	G0_R XER	G0_C OL	U
V	LB_A DSC#	LB_O E#	LB_W E#	LB_D 29	LB_D 28	VDD		VSS						VSS															VDD						G0_T XD12	G0_T XD13	G0_R XD12	G0_R XD13	G0_R XD14	V		
W	LB_D 15	LB_A 3	LB_A 4	LB_D 27	LB_D 26																								G0_T XD10	G0_T XD11	G0_R XD9	G0_R XD10	G0_R XD11	W								
Y	LB_D 14	LB_D 13	LB_D 12	LB_D 25	LB_D 24	VDD		VDD																									G0_R XD6	G0_T XD8	G0_R XD9	G0_R XD7	G0_R XD8	Y				
AA	LB_D 11	LB_D 10	LB_D 9	LB_D 23	LB_D 22																								G0_T XD6	G0_T XD7	G0_R XD3	G0_R XD4	G0_R XD5	AA								
AB	LB_D 8	LB_D 7	LB_D 6	LB_D 21	LB_D 20																								G0_T XD4	G0_T XD5	G0_R XD0	G0_R XD1	G0_R XD2	AB								
AC	LB_D 5	LB_D 4	LB_D 3	LB_D 19	LB_D 18																								G0_T XD2	G0_T XD3	M23_CRS	M23_RXD0	M23_RXD1	AC								
AD	LB_D 2	LB_D 1	LB_D 0	LB_D 17	LB_D 16	VCC					VCC																		G0_T XD0	G0_T XD1	M23_TXD1	M23_TXD0	M23_TXEN	AD								
AE	M0_T XEN	M0_T XD0	M0_T XD1	M3_T XD1	M3_T XEN	M3_R XD0	M5_T XD1	M5_T XEN	M5_R XD0	M8_T XD1	M8_T XEN	M8_R XD0	M10_TXD1	M10_TXEN	M10_RXD0	M13_TXD1	M16_TXD0	M15_TXD1	M16_RXD1	M15_TXEN	M15_RXD0	M18_TXD1	M18_TXEN	M18_RXD0	M20_TXD1	M20_TXEN	M20_RXD0	M22_RXD1	NC	AE												
AF	M0_R XD1	M0_R XD0	M0_C RS	M3_T XD0	M3_C RS	M3_R XD1	M5_T XD0	M5_C RS	M5_R XD1	M8_T XD0	M8_C RS	M8_R XD1	M10_TXD0	M10_CRS	M10_RXD1	M13_TXD0	M13_CRS	M13_RXD1	M14_CRS	M16_RXD0	M15_RXD1	M17_RXD0	M17_CRS	M18_RXD1	M18_CRS	M19_RXD1	M19_CRS	M20_TXD0	M20_CRS	M20_RXD1	M22_RXD0	M22_CRS	AF									
AG	M1_T XEN	M1_T XD0	M1_T XD1	M2_T XD1	M2_C RS	M4_T XD1	M4_C RS	M6_T XD1	M6_C RS	M7_T XD1	M7_C RS	M9_T XD1	M9_C RS	M11_TXD1	M11_CRS	M12_TXD1	M12_CRS	M14_TXD1	M15_TXD0	M16_CRS	M16_TXD1	M18_CRS	M18_CRS	M19_TXD1	M19_CRS	M21_TXD1	M21_CRS	M22_TXEN	M22_TXD0	AG												
AH		M1_R XD0	M1_C RS	M2_T XD0	M2_R XD0	M4_T XD0	M4_R XD0	M6_T XD0	M7_R XD0	M7_T XD0	M7_R XD0	M9_T XD0	M9_R XD0	M11_TXD0	M11_RXD0	M12_TXD0	M12_RXD0	M14_TXD0	M14_RXD0	M13_RXD0	M15_CRS	M17_TXD0	M17_RXD1	M19_TXD0	M19_RXD0	M21_TXD0	M21_RXD0	M22_TXD1	AH													
AJ		M1_R XD1	M2_T XEN	M2_R XD1	M4_T XEN	M4_R XD1	M6_T XEN	M7_R XD1	M7_T XEN	M7_R XD1	M9_T XEN	M9_R XD1	M11_TXEN	M11_RXD1	M12_TXEN	M12_RXD1	M14_TXEN	M14_RXD1	M16_TXEN	M13_TXEN	M17_TXEN	M17_TXD1	M19_TXEN	M19_RXD1	M21_TXEN	M21_RXD1	AJ															
	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													

1.2 Ball – Signal Descriptions

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

Notes:

= Active low signal

Weak internal pull-up/down resistors are nominal 100k ohm

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

Ball No(s)	Symbol	I/O	Description
CPU BUS Interface in Unmanaged Mode - Use I2C and Serial control interface to configure the system			
A24	SCL	Output	I2C Data Clock
A25	SDA	I/O-TS with weak internal pull-up	I2C Data I/O
A26	STROBE	Input with weak internal pull-up	Serial Strobe Pin
B26	DATAIN (D0)	Input with weak internal pull-up	Serial Data Input (D0)
C25	DATAOUT (AUTOFD)	Output with weak internal 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 weak internal 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]

Ball No(s)	Symbol	I/O	Description
B8	LA_ADSC#	Output with weak internal pull-up	Frame Bank A Address Status Control
C1	LA_CLK	Output	Frame Bank A Clock Input
C9	LA_WE#	Output with weak internal pull-up	Frame Bank A Write Chip Select for one layer SRAM configuration
D12	LA_WE0#	Output with weak internal pull-up	Frame Bank A Write Chip Select for lower layer of two layers SRAM configuration
E12	LA_WE1#	Output with weak internal pull-up	Frame Bank A Write Chip Select for upper layer of two layers SRAM configuration
C8	LA_OE#	Output with weak internal pull-up	Frame Bank A Read Chip Select for one bank SRAM configuration
A9	LA_OE0#	Output with weak internal pull-up	Frame Bank A Read Chip Select for lower layer of two layers SRAM configuration
B9	LA_OE1#	Output with weak internal pull-up	Frame Bank A Read Chip Select for upper layer of two layers SRAM configuration
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	LB_D[63:0]	I/O-TS with weak internal pull-up	Frame Bank B– Data Bit [63:0]
N3, N2, N1, P3, P2, P1, R5, R4, R3, R2, R1, T5, T4, T3, T2, T1, W3, W2	LB_A[20:3]	Output	Frame Bank B – Address Bit [20:3]
V1	LB_ADSC#	Output with weak internal pull-up	Frame Bank B Address Status Control
G1	LB_CLK	Output with weak internal pull-up	Frame Bank B Clock Input
V3	LB_WE#	Output with weak internal pull-up	Frame Bank B Write Chip Select for one layer SRAM configuration

Ball No(s)	Symbol	I/O	Description
P4	LB_WE0#	Output with weak internal pull-up	Frame Bank B Write Chip Select for lower layer of two layer SRAM configuration
P5	LB_WE1#	Output with weak internal pull-up	Frame Bank B Write Chip Select for upper layer of two layers SRAM configuration
V2	LB_OE#	Output with weak internal pull-up	Frame Bank B Read Chip Select for one layer SRAM configuration
U1	LB_OE0#	Output with weak internal pull-up	Frame Bank B Read Chip Select for lower layer of two layers SRAM configuration
U2	LB_OE1#	Output with weak internal pull-up	Frame Bank B Read Chip Select for upper layer of two layers SRAM configuration
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 weak internal pull-up	MII Management Data I/O – (Common for all MII Ports –[23:0])
R29	M_CLK	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	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	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	Ports [23:0] – Carrier Sense and Receive Data Valid

Ball No(s)	Symbol	I/O	Description
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, slew with weak internal pull-up	Ports [23:0] – Transmit Enable Bootstrap option for RMII/GPSI
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]
GMII/TBI Gigabit Ethernet Access Ports 0 & 1 (also referred to ports 25 & 26)			
U26, U25, V26, V25, W26, W25, Y27, Y26, AA26, AA25, AB26, AB25, AC26, AC25, AD26, AD25	G0_TXD[15:0]	Output	Transmit Data Bit [15:0] [7:0] - GMII [9:0] - TBI [15:0] - 2G
T28	G0_RX_DV	Input with weak internal pull-down	Receive Data Valid - GMII/MII
U28	G0_RX_ER	Input with weak internal pull-up	Receive Error - GMII/MII
R25	G0_CRSD/ G0_LINK	Input with weak internal pull-down	Carrier Sense - GMII/MII Link Status - TBI
U29	G0_COL/ G0_RBC1	Input with weak internal pull-up	Collision Detected - GMII/MII Receive Byte Clock 1 - TBI
T29	G0_RXCLK G0_RBC0	Input with weak internal pull-up	Receive Clock - GMII/MII Receive Byte Clock 0 - TBI
U27, V29, V28, V27, W29, W28, W27, Y29, Y28, Y25, AA29, AA28, AA27, AB29, AB28, AB27	G0_RXD[15:0]	Input with weak internal pull-up	Receive Data Bit [15:0] [7:0] - GMII [9:0] - TBI [15:0] - 2G

Ball No(s)	Symbol	I/O	Description
T26	G0_TX_EN	Output with weak internal pull-up	Transmit Data Enable - GMII/MII
R26	G0_TX_ER	Output with weak internal pull-up	Transmit Error - GMII/MII
T27	G0_MTXCLK	Input with weak internal pull-down	MII Mode Transmit Clock
T25	G0_TXCLK	Output	Gigabit Transmit Clock
P29	GREFCLK0	Input with weak internal pull-up	Gigabit Reference Clock
G26, G25, H26, H25, J26, J25, K25, K26, M25, L26, M26, L25, N26, N25, P26, P25	G1_TXD[15:0]	Output	Transmit Data Bit [15:0] [7:0] - GMII [9:0] - TBI [15:0] - 2G
F28	G1_RX_DV	Input with weak internal pull-down	Receive Data Valid - GMII/MII
G28	G1_RX_ER	Input with weak internal pull-up	Receive Error - GMII/MII
E25	G1_CRSD/ G1_LINK	Input with weak internal pull-down	Carrier Sense - GMII/MII Link Status - TBI
G29	G1_COL/ G1_RBC1	Input with weak internal pull-up	Collision Detected - GMII/MII Receive Byte Clock 1 - TBI
F29	G1_RXCLK G1_RBC0	Input with weak internal pull-up	Receive Clock - GMII/MII Receive Byte Clock 0 - TBI
G27, H29, H28, H27, J29, J28, J27, K29, K28, K27, L29, L28, L27, M29, M28, M27	G1_RXD[15:0]	Input with weak internal pull-up	Receive Data Bit [15:0] [7:0] - GMII [9:0] - TBI [15:0] - 2G
F26	G1_TX_EN	Output with weak internal pull-up	Transmit Data Enable - GMII/MII
E26	G1_TX_ER	Output with weak internal pull-up	Transmit Error - GMII/MII
F27	G1_MTXCLK	Input with weak internal pull-down	MII Mode Transmit Clock
F25	G1_TXCLK	Output	Gigabit Transmit Clock
N29	GREFCLK1	Input with weak internal pull-up	Gigabit Reference Clock
LED Interface			
C29	LED_CLK/ TSTOUT0	I/O-TS with weak internal pull-up	LED Serial Interface Output Clock

Ball No(s)	Symbol	I/O	Description
D29	LED_SYN/ TSTOUT1	I/O-TS with weak internal pull-up	LED Output Data Stream Envelope
E29	LED_BIT/ TSTOUT2	I/O-TS with weak internal pull-up	LED Serial Data Output Stream
B28	LED_G1_RXTX#/ TSTOUT3	I/O-TS with weak internal pull-up	LED for Gigabit port 1 (receive + transmit)
C28	LED_G1_DPCOL#/ TSTOUT4	I/O-TS with weak internal pull-up	LED for Gigabit port 1 (full duplex + collision)
D28	LED_G1_LINK#/ TSTOUT5	I/O-TS with weak internal pull-up	LED for Gigabit port 1
E28	LED_G2_RXTX#/ TSTOUT6	I/O-TS with weak internal pull-up	LED for Gigabit port 2 (receive + transmit)
A27	LED_G2_DPCOL#/ TSTOUT7	I/O-TS with weak internal pull-up	LED for Gigabit port 2 (full duplex + collision)
B27	LED_G2_LINK#/ TSTOUT8	I/O-TS with weak internal pull-up	LED for Gigabit port 2
C27	INIT_DONE/ TSTOUT9	I/O-TS with weak internal pull-up	System start operation
D27	INIT_START/ TSTOUT10	I/O-TS with weak internal pull-up	Start initialization
C26	CHECKSUM_OK/ TSTOUT11	I/O-TS with weak internal pull-up	EEPROM read OK
D26	FCB_ERR/ TSTOUT12	I/O-TS with weak internal pull-up	FCB memory self test fail
D25	MCT_ERR/ TSTOUT13	I/O-TS with weak internal pull-up	MCT memory self test fail
D24	BIST_IN_PRC/ TSTOUT14	I/O-TS with weak internal pull-up	Processing memory self test
E24	BIST_DONE/ TSTOUT15	I/O-TS with weak internal pull-up	Memory self test done
Test Facility			
U3, C10	T_MODE0, T_MODE1	I/O-TS Must be externally pulled-up	Test Pins. Manufacturing test option. 00 – Test mode – Set test mode upon reset, and provides NANDTree test output during test mode 01 - Reserved - Do not use 10 - Reserved - Do not use 11 – Normal mode Use external pull-ups for normal mode

Ball No(s)	Symbol	I/O	Description
F3	SCAN_EN	Input with weak internal pull-down	Scan Enable. Manufacturing test option. Should not be connected for proper operation.
E27	SCANMODE	Input with weak internal pull-down	Scan Mode Enable. Manufacturing test option. 1 – Enable Test mode 0 - Normal mode (open) Should not be connected for proper operation.
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
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	VDDA	Analog Power	Analog +2.5 Volt DC Supply
D1	VSSA	Analog Ground	Analog Ground
MISC			
D22	SCANCOL	I/O	Scans the Collision signal of Home PHY
D23	SCANCLK	Output	Clock for scanning Home PHY collision and link

Ball No(s)	Symbol	I/O	Description
E23	SCANLINK	I/O	Link up signal from Home PHY
F2	RESIN#	Input	Reset Input
G2	RESETOUT#	Output	Reset PHY
E22, N27, N28, P27, R27, AE29	NC	NC	No Internal Connect
B25, E20	RSVD	N/A	Reserved. Leave unconnected.
Bootstrap Pins (1= pull-up 0= pull-down) (Default = 1 due to weak internal pull-ups)			
After reset TSTOUT0 to TSTOU15 are used by the LED interface.			
C29	TSTOUT0	Input (Reset Only) with weak internal pull-up	Polarity for Gn_LINK in TBI mode 1 – active high 0 – active low
D29	TSTOUT1	Input (Reset Only) with weak internal pull-up	RMII MAC Power Saving Enable 1 – power saving 0 – No power saving
E29	TSTOUT2	Input (Reset Only) with weak internal pull-up Must be externally pulled-down	Manufacturing Option. Must be '0'.
B28	TSTOUT3	Input (Reset Only) with weak internal pull-up	Giga Module Detect Enable 1 – Hot swap disable 0 – Hot swap enable
C28	TSTOUT4	Input (Reset Only) with weak internal pull-up	Memory is SDRAM/ZBT 1 – Pipeline SDRAM 0 – ZBT
D28	TSTOUT5	Input (Reset Only) with weak internal pull-up	Scan Speed: ¼ SCLK or SCLK 1 - SCLK 0 – ¼ SCLK (HPNA)
E28	TSTOUT6	Input (Reset Only) with weak internal pull-up	Reserved
A27	TSTOUT7	Input (Reset Only) with weak internal pull-up	Memory Size 1 - 128 K x 32 or 128 K x 64 (1 M/bank, 2 M total) 0 - -256 K x 32 or 256 K x 64 (2 M/bank, 4 M total)
B27	TSTOUT8	Input (Reset Only) with weak internal pull-up	EEPROM Installed 1 – EEPROM not installed 0 – EEPROM installed

Ball No(s)	Symbol	I/O	Description
C27	TSTOUT9	Input (Reset Only) with weak internal pull-up	MCT Aging 1 – MCT aging enable 0 – MCT aging disable
D27	TSTOUT10	Input (Reset Only) with weak internal pull-up Must be externally pulled-down	Manufacturing Option. Must be '0'.
C26	TSTOUT11	Input (Reset Only) with weak internal pull-up	Timeout Reset 1 – Time out reset enable 0 – Time out reset disable If enabled, issue reset if any state machine did not go back to idle for 5sec.
D26	TSTOUT12	Input (Reset Only) with weak internal pull-up	Manufacturing Option. Must be '1'.
D25	TSTOUT13	Input (Reset Only) with weak internal pull-up	FDB RAM depth (1 or 2 layers) 1 – 1 layer 0 – 2 layer
D24	TSTOUT14	Input (Reset Only) with weak internal pull-up	Reserved
E24	TSTOUT15	Input (Reset Only) with weak internal pull-up	SRAM Test Mode 1 – Normal operation 0 – Enable test mode
T26, R26	G0_TXEN, G0_TXER	Input (Reset Only) with weak internal pull-up	Gigabit Port 0 G0_TXEN G0_TXER 1 1 TBI 1 0 GMII 0 1 2G 0 0 MII
F26, E26	G1_TXEN, G1_TXER	Input (Reset Only) with weak internal pull-up	Gigabit Port 1 G1_TXEN G1_TXER 1 1 TBI 1 0 GMII 0 1 2G 0 0 MII

Ball No(s)	Symbol	I/O	Description
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	Input (Reset Only) with weak internal pull-up	1 – RMII 0 – GPSI
C21	P_D	Input (Reset Only) with weak internal pull-up Must be externally pulled-down	Manufacturing Option. Must be '0'.
C19, B19, A19	OE_CLK[2:0]	Input (Reset Only) with weak internal pull-up Recommend 001 with external pull-downs on P_D[15:14]OE_CLK[2:1].	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.
C20, B20, A20	L_CLK[2:0]	Input (Reset Only) with weak internal pull-up Recommend 011 with external pull-down on P_D[12]L_CLK[2].	Programmable delay for LA_CLK and LB_CLK from internal OE_CLK. The LA_CLK and LB_CLK delay from SCLK is the sum of the delay programmed in here and the delay in OE_CLK[2:0]. Suggested value is 011.
B22, A22, C23, B23, A23, C24	MIRROR[5:0]	Input (Reset Only) with weak internal pull-up	Dedicated Port Mirror Mode. The first 5 bits ([4:0]) select the port to be mirrored. The last bit ([5]) selects either ingress or egress data.
C22	TRUNK0	Input (Reset Only) with weak internal pull-down	Trunk Group 0 Enable 0 – Disable 1 – Enable
A21	TRUNK1	Input (Reset Only) with weak internal pull-down	Trunk Group 1 Enable 0 – Disable 1 – Enable
B24	TRUNK2	Input (Reset Only) with weak internal pull-down	Trunk Group 2 Enable 0 – Disable 1 – Enable

1.3 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	LB_D[63]
E19	LA_D[60]	E2	LA_D[16]	F5	LB_D[62]
D18	LA_D[59]	A7	LA_D[15]	G4	LB_D[61]
E18	LA_D[58]	B7	LA_D[14]	G5	LB_D[60]
D17	LA_D[57]	A6	LA_D[13]	H4	LB_D[59]
E17	LA_D[56]	B6	LA_D[12]	H5	LB_D[58]
D16	LA_D[55]	C6	LA_D[11]	J4	LB_D[57]
E16	LA_D[54]	A5	LA_D[10]	J5	LB_D[56]
D15	LA_D[53]	B5	LA_D[9]	K4	LB_D[55]
E15	LA_D[52]	C5	LA_D[8]	K5	LB_D[54]
D14	LA_D[51]	A4	LA_D[7]	L4	LB_D[53]
E14	LA_D[50]	B4	LA_D[6]	L5	LB_D[52]
D13	LA_D[49]	C4	LA_D[5]	M4	LB_D[51]
E13	LA_D[48]	A3	LA_D[4]	M5	LB_D[50]
D21	LA_D[47]	B3	LA_D[3]	N4	LB_D[49]
E21	LA_D[46]	C3	LA_D[2]	N5	LB_D[48]
A18	LA_D[45]	B2	LA_D[1]	G3	LB_D[47]
B18	LA_D[44]	C2	LA_D[0]	H1	LB_D[46]
C18	LA_D[43]	C14	LA_A[20]	H2	LB_D[45]
A17	LA_D[42]	A13	LA_A[19]	H3	LB_D[44]
B17	LA_D[41]	B13	LA_A[18]	J1	LB_D[43]
C17	LA_D[40]	C13	LA_A[17]	J2	LB_D[42]
A16	LA_D[39]	A12	LA_A[16]	J3	LB_D[41]
B16	LA_D[38]	B12	LA_A[15]	K1	LB_D[40]
C16	LA_D[37]	C12	LA_A[14]	K2	LB_D[39]
A15	LA_D[36]	A11	LA_A[13]	K3	LB_D[38]
B15	LA_D[35]	B11	LA_A[12]	L1	LB_D[37]
C15	LA_D[34]	C11	LA_A[11]	L2	LB_D[36]
A14	LA_D[33]	D11	LA_A[10]	L3	LB_D[35]
B14	LA_D[32]	E11	LA_A[9]	M1	LB_D[34]

Ball No.	Signal Name	Ball No.	Signal Name	Ball No.	Signal Name
D9	LA_D[31]	A10	LA_A[8]	M2	LB_D[33]
E9	LA_D[30]	B10	LA_A[7]	M3	LB_D[32]
D8	LA_D[29]	D10	LA_A[6]	U4	LB_D[31]
E8	LA_D[28]	E10	LA_A[5]	U5	LB_D[30]
D7	LA_D[27]	A8	LA_A[4]	V4	LB_D[29]
E7	LA_D[26]	C7	LA_A[3]	V5	LB_D[28]
D6	LA_D[25]	B8	LA_DSC#	W4	LB_D[27]
E6	LA_D[24]	C1	LA_CLK	W5	LB_D[26]
D5	LA_D[23]	C9	LA_WE#	Y4	LB_D[25]
E5	LA_D[22]	D12	LA_WE0#	Y5	LB_D[24]
D4	LA_D[21]	E12	LA_WE1#	AA4	LB_D[23]
E4	LA_D[20]	C8	LA_OE#	AA5	LB_D[22]
AB4	LB_D[21]	U2	LB_OE1#	AH7	M[4]_RXD[0]
AB5	LB_D[20]	R28	MDC	AE6	M[3]_RXD[0]
AC4	LB_D[19]	P28	MDIO	AH5	M[2]_RXD[0]
AC5	LB_D[18]	R29	M_CLK	AH2	M[1]_RXD[0]
AD4	LB_D[17]	AC29	M[23]_RXD[1]	AF2	M[0]_RXD[0]
AD5	LB_D[16]	AE28	M[22]_RXD[1]	AC27	M[23]_CRS_DV
W1	LB_D[15]	AJ27	M[21]_RXD[1]	AF29	M[22]_CRS_DV
Y1	LB_D[14]	AF27	M[20]_RXD[1]	AG27	M[21]_CRS_DV
Y2	LB_D[13]	AJ25	M[19]_RXD[1]	AF26	M[20]_CRS_DV
Y3	LB_D[12]	AF24	M[18]_RXD[1]	AG25	M[19]_CRS_DV
AA1	LB_D[11]	AH23	M[17]_RXD[1]	AG23	M[18]_CRS_DV
AA2	LB_D[10]	AE19	M[16]_RXD[1]	AF23	M[17]_CRS_DV
AA3	LB_D[9]	AF21	M[15]_RXD[1]	AG21	M[16]_CRS_DV
AB1	LB_D[8]	AJ19	M[14]_RXD[1]	AH21	M[15]_CRS_DV
AB2	LB_D[7]	AF18	M[13]_RXD[1]	AF19	M[14]_CRS_DV
AB3	LB_D[6]	AJ17	M[12]_RXD[1]	AF17	M[13]_CRS_DV
AC1	LB_D[5]	AJ15	M[11]_RXD[1]	AG17	M[12]_CRS_DV
AC2	LB_D[4]	AF15	M[10]_RXD[1]	AG15	M[11]_CRS_DV
AC3	LB_D[3]	AJ13	M[9]_RXD[1]	AF14	M[10]_CRS_DV
AD1	LB_D[2]	AF12	M[8]_RXD[1]	AG13	M[9]_CRS_DV
AD2	LB_D[1]	AJ11	M[7]_RXD[1]	AF11	M[8]_CRS_DV

Ball No.	Signal Name	Ball No.	Signal Name	Ball No.	Signal Name
AD3	LB_D[0]	AJ9	M[6]_RXD[1]	AG11	M[7]_CRS_DV
N3	LB_A[20]	AF9	M[5]_RXD[1]	AG9	M[6]_CRS_DV
N2	LB_A[19]	AJ7	M[4]_RXD[1]	AF8	M[5]_CRS_DV
N1	LB_A[18]	AF6	M[3]_RXD[1]	AG7	M[4]_CRS_DV
P3	LB_A[17]	AJ5	M[2]_RXD[1]	AF5	M[3]_CRS_DV
P2	LB_A[16]	AJ3	M[1]_RXD[1]	AG5	M[2]_CRS_DV
P1	LB_A[15]	AF1	M[0]_RXD[1]	AH3	M[1]_CRS_DV
R5	LB_A[14]	AC28	M[23]_RXD[0]	AF3	M[0]_CRS_DV
R4	LB_A[13]	AF28	M[22]_RXD[0]	AD29	M[23]_TXEN
R3	LB_A[12]	AH27	M[21]_RXD[0]	AG28	M[22]_TXEN
R2	LB_A[11]	AE27	M[20]_RXD[0]	AJ26	M[21]_TXEN
R1	LB_A[10]	AH25	M[19]_RXD[0]	AE26	M[20]_TXEN
T5	LB_A[9]	AE24	M[18]_RXD[0]	AJ24	M[19]_TXEN
T4	LB_A[8]	AF22	M[17]_RXD[0]	AE23	M[18]_TXEN
T3	LB_A[7]	AF20	M[16]_RXD[0]	AJ22	M[17]_TXEN
T2	LB_A[6]	AE21	M[15]_RXD[0]	AJ20	M[16]_TXEN
T1	LB_A[5]	AH19	M[14]_RXD[0]	AE20	M[15]_TXEN
W3	LB_A[4]	AH20	M[13]_RXD[0]	AJ18	M[14]_TXEN
W2	LB_A[3]	AH17	M[12]_RXD[0]	AJ21	M[13]_TXEN
V1	LB_ADSC#	AH15	M[11]_RXD[0]	AJ16	M[12]_TXEN
G1	LB_CLK	AE15	M[10]_RXD[0]	AJ14	M[11]_TXEN
V3	LB_WE#	AH13	M[9]_RXD[0]	AE14	M[10]_TXEN
P4	LB_WE0#	AE12	M[8]_RXD[0]	AJ12	M[9]_TXEN
P5	LB_WE1#	AH11	M[7]_RXD[0]	AE11	M[8]_TXEN
V2	LB_OE#	AH9	M[6]_RXD[0]	AJ10	M[7]_TXEN
U1	LB_OE0#	AE9	M[5]_RXD[0]	AJ8	M[6]_TXEN
AE8	M[5]_TXEN	AH8	M[6]_TXD[0]	G27	G1_RXD[15]
AJ6	M[4]_TXEN	AF7	M[5]_TXD[0]	H29	G1_RXD[14]
AE5	M[3]_TXEN	AH6	M[4]_TXD[0]	H28	G1_RXD[13]
AJ4	M[2]_TXEN	AF4	M[3]_TXD[0]	H27	G1_RXD[12]
AG1	M[1]_TXEN	AH4	M[2]_TXD[0]	J29	G1_RXD[11]
AE1	M[0]_TXEN	AG2	M[1]_TXD[0]	J28	G1_RXD[10]
AD27	M[23]_TXD[1]	AE2	M[0]_TXD[0]	J27	G1_RXD[9]

Ball No.	Signal Name	Ball No.	Signal Name	Ball No.	Signal Name
AH28	M[22]_TXD[1]	U26	G0_TXD[15]	K29	G1_RXD[8]
AG26	M[21]_TXD[1]	U25	G0_TXD[14]	K28	G1_RXD[7]
AE25	M[20]_TXD[1]	V26	G0_TXD[13]	K27	G1_RXD[6]
AG24	M[19]_TXD[1]	V25	G0_TXD[12]	L29	G1_RXD[5]
AE22	M[18]_TXD[1]	W26	G0_TXD[11]	L28	G1_RXD[4]
AJ23	M[17]_TXD[1]	W25	G0_TXD[10]	L27	G1_RXD[3]
AG20	M[16]_TXD[1]	Y27	G0_TXD[9]	M29	G1_RXD[2]
AE18	M[15]_TXD[1]	Y26	G0_TXD[8]	M28	G1_RXD[1]
AG18	M[14]_TXD[1]	AA26	G0_TXD[7]	M27	G1_RXD[0]
AE16	M[13]_TXD[1]	AA25	G0_TXD[6]	G26	G1_TXD[15]
AG16	M[12]_TXD[1]	AB26	G0_TXD[5]	G25	G1_TXD[14]
AG14	M[11]_TXD[1]	AB25	G0_TXD[4]	H26	G1_TXD[13]
AE13	M[10]_TXD[1]	AC26	G0_TXD[3]	H25	G1_TXD[12]
AG12	M[9]_TXD[1]	AC25	G0_TXD[2]	J26	G1_TXD[11]
AE10	M[8]_TXD[1]	AD26	G0_TXD[1]	J25	G1_TXD[10]
AG10	M[7]_TXD[1]	AD25	G0_TXD[0]	K25	G1_TXD[9]
AG8	M[6]_TXD[1]	U27	G0_RXD[15]	K26	G1_TXD[8]
AE7	M[5]_TXD[1]	V29	G0_RXD[14]	M25	G1_TXD[7]
AG6	M[4]_TXD[1]	V28	G0_RXD[13]	L26	G1_TXD[6]
AE4	M[3]_TXD[1]	V27	G0_RXD[12]	M26	G1_TXD[5]
AG4	M[2]_TXD[1]	W29	G0_RXD[11]	L25	G1_TXD[4]
AG3	M[1]_TXD[1]	W28	G0_RXD[10]	N26	G1_TXD[3]
AE3	M[0]_TXD[1]	W27	G0_RXD[9]	N25	G1_TXD[2]
AD28	M[23]_TXD[0]	Y29	G0_RXD[8]	P26	G1_TXD[1]
AG29	M[22]_TXD[0]	Y28	G0_RXD[7]	P25	G1_TXD[0]
AH26	M[21]_TXD[0]	Y25	G0_RXD[6]	F28	G1_RX_DV
AF25	M[20]_TXD[0]	AA29	G0_RXD[5]	G28	G1_RX_ER
AH24	M[19]_TXD[0]	AA28	G0_RXD[4]	E25	G1_CRS
AG22	M[18]_TXD[0]	AA27	G0_RXD[3]	G29	G1_COL
AH22	M[17]_TXD[0]	AB29	G0_RXD[2]	F29	G1_RXCLK
AE17	M[16]_TXD[0]	AB28	G0_RXD[1]	F26	G1_TX_EN
AG19	M[15]_TXD[0]	AB27	G0_RXD[0]	E26	G1_TX_ER
AH18	M[14]_TXD[0]	R26	G0_TX_ER	F25	G1_TXCLK

Ball No.	Signal Name	Ball No.	Signal Name	Ball No.	Signal Name
AF16	M[13]_TXD[0]	T25	G0_TXCLK	E24	BIST_DONE/TSTOUT[15]
AH16	M[12]_TXD[0]	T26	G0_TX_EN	D24	BIST_IN_PRC/TSTOUT[14]
AH14	M[11]_TXD[0]	T28	G0_RX_DV	D25	MCT_ERR/TSTOUT[13]
AF13	M[10]_TXD[0]	U28	G0_RX_ER	D26	FCB_ERR/TSTOUT[12]
AH12	M[9]_TXD[0]	R25	G0_CRS	C26	CHECKSUM_OK/TSTOUT[11]
AF10	M[8]_TXD[0]	U29	G0_COL	D27	INIT_START/TSTOUT[10]
AH10	M[7]_TXD[0]	T29	G0_RXCLK	C27	INIT_DONE/TSTOUT[9]
B27	G2_LINK#/TSTOUT[8]	U18	VSS	N12	VSS
A27	G2_DPCOL#/TSTOUT[7]	V12	VSS	N13	VSS
E28	G2_RXTX#/TSTOUT[6]	V13	VSS	K17	VDD
D28	G1_LINK#/TSTOUT[5]	V14	VSS	K18	VDD
C28	G1_DPCOL#/TSTOUT[4]	V15	VSS	M10	VDD
B28	G1_RXTX#/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	REF_CLK1	N15	VSS	V10	VDD
P29	REF_CLK0	C19	OE_CLK2	U20	VDD
F3	SCAN_EN	B19	OE_CLK1	V20	VDD
E1	SCLK	A19	OE_CLK0	Y12	VDD
U3	T_MODE0	P12	VSS	Y13	VDD
C10	T_MODE1	P13	VSS	Y17	VDD
B24	TRUNK2	P14	VSS	Y18	VDD
A21	TRUNK1	P15	VSS	K12	VDD
C22	TRUNK0	P16	VSS	K13	VDD
A26	STROBE	N16	VSS	M16	VSS
B26	D0	N17	VSS	M17	VSS

Ball No.	Signal Name	Ball No.	Signal Name	Ball No.	Signal Name
C25	AUTOFD	N18	VSS	M18	VSS
A24	SCL	R13	VSS	F16	VCC
A25	SDA	R14	VSS	F17	VCC
F1	VDDA	R15	VSS	N6	VCC
D1	VSSA	R16	VSS	P6	VCC
D22	SCANCOL	R17	VSS	R6	VCC
E23	SCANLINK	R18	VSS	T6	VCC
E27	SCANMODE	T12	VSS	U6	VCC
N28	NC	T13	VSS	N24	VCC
N27	NC	T14	VSS	P24	VCC
F2	RESIN#	T15	VSS	R24	VCC
G2	RESETOUT#	T16	VSS	T24	VCC
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	G0_MTXCLK	U17	VSS	F14	VCC
F27	G1_MTXCLK	M12	VSS	F15	VCC
C20	L_CLK2	M13	VSS		
B20	L_CLK1	M14	VSS		
A20	L_CLK0	M15	VSS		
C21	P_D	P17	VSS		
E20	RSVD	P18	VSS		
B25	RSVD	R12	VSS		

1.4 Signal Mapping and Internal Pull Up/Down Configuration

The MVTX2603 Fast Ethernet ports (0-23) support 2 interface options: RMII & GPSI. The table below summarizes the interface signals required for each interface and how they relate back to the Pin Symbol name shown in “Ball – Signal Descriptions” on page 11.

Notes:

I – Input

O – Output

NC - No Connect

Fast Ethernet Ports Pin Symbol	RMII Mode (Bootstrap Mn_TXEN='1')	GPSI Mode (Bootstrap Mn_TXEN='0')
Mn_RXD0	Mn_RXD0 (I)	Mn_RXD (I)
Mn_RXD1	Mn_RXD1 (I)	Mn_RXCLK (I)
Mn_CRS_DV	Mn_CRS_DV (I)	Mn_CRS (I)
Mn_TXD0	Mn_TXD0 (O)	Mn_TXD (O)
Mn_TXD1	Mn_TXD1 (O)	Mn_TXCLK (I)
Mn_TXEN	Mn_TXEN (O)	Mn_TXEN (O)
M_CLK	M_CLK (I)	M_CLK (I)
SCANCLK	NC	SCANCLK (O)
SCANLINK	NC	SCANLINK (IO)
SCANCOL	NC	SCANCOL (IO)

Table 1 - Fast Ethernet Ports Signal Mapping In Different Operation Mode

The MVTX2603 Gigabit Ethernet ports supports 4 interface options: 2G, GMII, TBI & MII. The table below summarizes the interface signals required for each interface and how they relate back to the Pin Symbol name shown in “Ball – Signal Descriptions” on page 11.

Notes:

I – Input
 O – Output
 U – Pull-up
 D – Pull-down
 NC – No Connect

Gigabit Ports Pin Symbol	No Module (Bootstrap TSTOUT3='0')	TBI Mode (Bootstrap Gn_TXEN='1' and Gn_TXER='1')	GMII Mode (Bootstrap Gn_TXEN='1' and Gn_TXER='0')	MI Mode (Bootstrap Gn_TXEN='0' and Gn_TXER='0')
Gn_RXD[3:0]	(U)	Gn_RXD[3:0] (I)	Gn_RXD[3:0] (I)	Gn_RXD[3:0] (I)
Gn_RXD[7:4]	(U)	Gn_RXD[7:4] (I)	Gn_RXD[7:4] (I)	NC (U)
Gn_RXD[9:8]	(U)	Gn_RXD[9:8] (I)	NC (U)	NC (U)
Gn_RXDV	(D)	NC (D)	Gn_RXDV (I)	Gn_RXDV (I)
Gn_RXER	(U)	NC (U)	Gn_RXER (I)	Gn_RXER (I)
Gn_CRS	(D)	Gn_LINK (I)	Gn_CRS (I)	Gn_CRS (I)
Gn_COL	(U)	Gn_RBC1 (I)	Gn_COL (I)	Gn_COL (I)
Gn_RXCLK	(U)	Gn_RBC0 (I)	Gn_RXCLK (I)	Gn_RXCLK (I)
Gn_TXD[3:0]	(O)	Gn_TXD[3:0] (O)	Gn_TXD[3:0] (O)	Gn_TXD[3:0] (O)
Gn_TXD[7:4]	(O)	Gn_TXD[7:4] (O)	Gn_TXD[7:4] (O)	NC (O)
Gn_TXD[9:8]	(O)	Gn_TXD[9:8] (I)	NC (O)	NC (O)
Gn_TXEN	(U)	NC (U)	Gn_TXEN (O)	Gn_TXEN (O)
Gn_TXER	(U)	NC (U)	Gn_TXER (O)	Gn_TXER (O)
Gn_TXCLK	(O)	Gn_TXCLK (O)	Gn_TXCLK (O)	NC (O)
GREFCLKn	(U)	GREFCLKn (I)	GREFCLKn (I)	REFCLKn (I)
Gn_MTXCLK	(D)	NC (D)	Gn_MTXCLK (I)	Gn_MTXCLK (I)

Table 2 - Gigabit Ethernet Ports Signal Mapping in Different Operation Mode

2.0 Block Functionality

2.1 Frame Data Buffer (FDB) Interfaces

The FDB interface supports pipelined synchronous burst SRAM (SBRAM) memory at 100 MHz. To ensure a non-blocking switch, two memory domains with a 64-bit wide memory bus are required. At 100 MHz, the aggregate memory bandwidth is 12.8 Gbps which is enough to support 24 10/100 M and 2 10/100/1000 M ports at full wire speed switching. For the 2G-mode stacking application, pipelined ZBT-SRAM memory running at 125 MHz is required.

The Switching Database is also located in the external SRAM; it is used for storing MAC addresses and their physical port number. It is duplicated and stored in both memory domains. Therefore, when the system updates the contents of the switching database it has to write the entry to both domains at the same time.

2.2 MAC Modules

2.2.1 RMI MAC Module (RMAC)

The 10/100 M Media Access Control (RMAC) module provides the necessary buffers and control interface between the Frame Engine (FE) and the external physical device (PHY).

The MVTX2603 RMAC implements two interfaces, RMI or GPSI (7WS) (only for 10 M), and fully 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 RMACs are from 08h to 1Fh. These twenty-four ports are denoted as ports 0 to 23.

2.2.1.1 GPSI Interface

The 10/100 M RMI 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 2 - GPSI (7WS) Mode Connection Diagram

2.2.1.2 SCANLINK and SCANCOL 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.

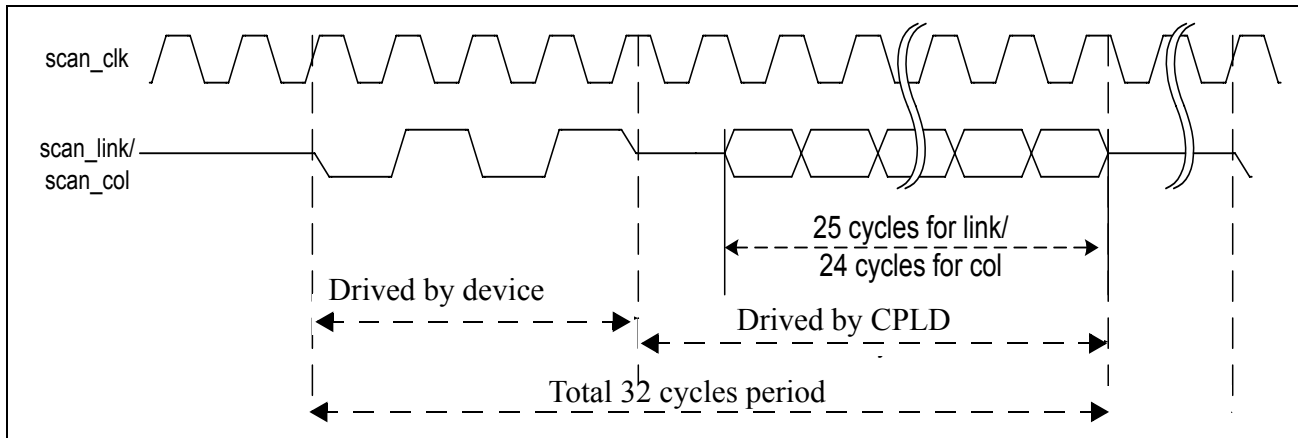


Figure 3 - SCANLINK and SCANCOL Status Diagram

2.2.2 GMII MAC Module (GMAC)

The 10/100/1000 M Media Access Control (MAC) module provides the necessary buffers and control interface between the Frame Engine (FE) and the external physical device (PHY). The MVTX2603 GMAC implements both GMII and MII interface, which offers a simple migration from 10/100 M to 1000 M.

The GMAC of the MVTX2603 meets the IEEE 802.3Z specification. It is able to operate in 10/100M either Half or Full Duplex mode with a back pressure/flow control mechanism or in 1G Full duplex mode with flow control mechanism. Furthermore, it will automatically retransmit upon collision for up to 16 total transmissions.

The PHY addresses for the two GMACs are 01h and 02h. These two ports are denoted as ports 25 (G0) and 26 (G1).

For fiber optics media, the MVTX2603 implements the Physical Code Sublayer (PCS) interface. The PCS includes an 8B10B encoder and decoder, auto-negotiation and Ten Bit Interface (TBI)

2.2.2.1 Physical Coding Sublayer (PCS) Module

For the MVTX2603, the 1000BASE-X PCS module is designed internally and may be utilized in the absence of GMII. The PCS incorporates all the functions required by the GMII to include encoding (decoding) 8B GMII data to (from) 8B/10B TBI format for PHY communication and generating Collision Detect (COL) signals for half-duplex mode. It also manages the auto-negotiation process by informing the management entity that the PHY is ready for communications. The on-chip PCS may be disabled if a PCS block exists within the Gigabit PHY. The TBI interface provides a uniform interface for all 1000 Mbps PHY implementations.

The PCS comprises the PCS Transmit, Synchronization, PCS Receive and auto-negotiation processes for 1000BASE-X.

The PCS Transmit process sends the TBI signals TXD[9:0] to the physical medium and generates the GMII Collision Detect (COL) signal based on whether a reception is occurring simultaneously with transmission. Additionally, the Transmit process generates an internal "transmitting" flag and monitors auto-negotiation to determine whether to transmit data or to reconfigure the link.

The PCS Synchronization process determines whether or not the receive channel is operational.

The PCS Receive process receives the TBI signals RXD[9:0] from the physical medium, and generates the GMII RXD[7:0] signals and the internal "receiving" flag for use by the Transmit processes.

The PCS auto-negotiation process allows the MVTX2603 to exchange configuration information between two devices that share a link segment and to automatically configure the link for the appropriate speed of operation for both devices.

2.2.2.2 TBI Interface

The TBI interface can be used for 1000M fiber operation. In this mode, the MVTX2603 is connected to the SERDES as shown in Figure 4. To enable the PCS module and TBI interface, the corresponding Gn_TXEN and Gn_TXER pins need to be boot strapped.

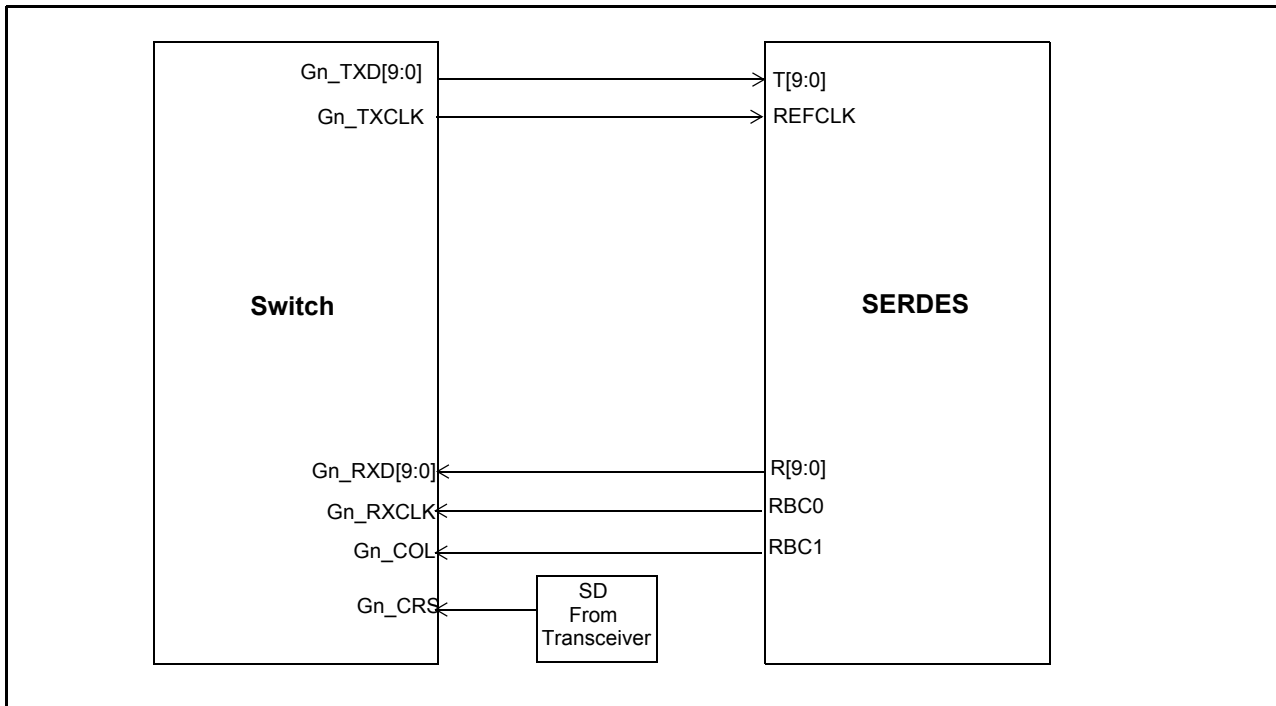


Figure 4 - TBI Connection

2.2.3 PHY Addresses

The table below provides an overview of the PHY addresses required for each port in order for the MDIO auto-negotiation to work between the MVTX2603 MAC and the PHY device. If a different PHY address is used, then the port must be manually brought up and the PHY will need to be polled for link status via the MIIC/D registers.

MAC Port	PHY Address
GMAC Port 0	0x01
GMAC Port 1	0x02
RMAC Port 0	0x08
RMAC Port 1	0x09
...	...
RMAC Port 23	0x1F
CMAC Port	N/A

Table 3 - PHY Addresses

2.3 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 which is 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.

2.4 Search Engine

The search engine resolves the frame's destination port or ports by searching the appropriate MVTX2603 databases. To achieve its objective, the search engine may use the destination MAC address, IP multicast address (IP multicast packet), and VLAN fields in the packet header. The search engine is also responsible for MAC and VLAN learning, assignment of transmission priority based on IEEE 802.1p or IP TOS/DS fields, and port trunking functions.

2.5 LED Interface

The LED interface provides a serial interface for carrying 24 + 2 port status signals. It can also provide direct status pins (6) for the two Gigabit ports.

A serial output channel provides port status information from the MVTX2603 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 non-serial interface is also allowed, but in this case, only the Gigabit ports will have status LEDs.

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.

2.5.1 Port Status

In the MVTX2603, 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 (10/100M): cycles #0 to cycle #7

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

...

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

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

Gigabit Port 0: cycle #192 to cycle #199

Gigabit Port 1: 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.

The first two bits of byte 26 provides the speed information for the Gigabit ports while the remainder of byte 26 and byte 27 provides bist status.

26[0]: G0 port (1= port G0 is operating at Gigabit speed; 0= speed is either 10 or 100 Mb/s depending on speed bit of Port 24)

26[1]: G1 port (1= port G1 is operating at Gigabit speed; 0= speed is either 10 or 100 Mb/s depending on speed bit of Port 25)

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

2.5.2 LED Interface Timing Diagram

The signal from the MVTX2603 to the LED decoder is shown in Figure 5.




Figure 5 - Timing Diagram of LED Interface

2.6 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.
- Network Management (NM) Database - The NM database contains the information in the statistics counters and MIB.
- MAC address Control Table (MCT) Link Table - The MCT Link Table stores the linked list of MCT entries that have collisions in the external MAC Table.

Note that the external MAC table is located in the external SRAM Memory.

2.7 Timeout Reset Monitor

The MVTX2603 supports a state machine monitoring block which can trigger a reset if any state machine is determined to be stuck in a non-idle state for more than 5 seconds. This feature is enabled via a bootstrap pin (TSTOUT11).

3.0 System Configuration (Stand-alone and Stacking)

3.1 Management and Configuration

Only one mode is supported in the MVTX2603: unmanaged. In unmanaged mode, the MVTX2603 has no CPU but can be configured by EEPROM using an I²C interface at bootup, or via a synchronous serial interface otherwise.

In unmanaged mode, the MVTX2603 can be configured by EEPROM (24C02 or compatible) via an I²C interface at boot time, or via a synchronous serial interface during operation.

3.1.1 I²C Interface

The I²C interface serves the function of configuring the MVTX2603 at boot time. The master is the MVTX2603, and the slave is the EEPROM memory.

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 6 depicts the data transfer format. The slave address is the memory address of the EEPROM. Refer to “Register Definition” on page 60 for I²C address for each register.

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

Figure 6 - Data Transfer Format for I²C Interface

3.1.1.1 Start Condition

Generated by the master (in our case, the MVTX2603). 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.

3.1.1.2 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.

3.1.1.3 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.

3.1.1.4 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.

3.1.1.5 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.

3.1.1.6 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.

3.1.2 Synchronous Serial Interface

The synchronous serial interface (SSI) serves the function of configuring the MVTX2603, not at boot time, but via a PC. The PC serves as master and the MVTX2603 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 MVTX2603 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 MVTX2603.

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

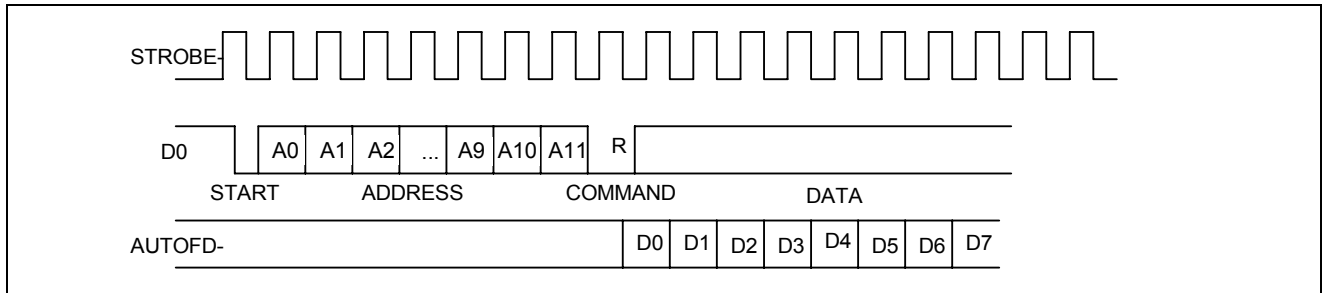
An ABORT command is detected when D0 is sampled low when STROBE rise and D0 is sampled high when STROBE fall.

All registers in MVTX2603 can be modified through this synchronous serial interface.

3.1.2.1 Write Command



3.1.2.2 Read Command




3.2 Stacking


The MVTX2603 supports expanded port count by providing stacking capabilities. The Gigabit port is used as the link between boxes, and each Gigabit port can be accelerated to 2 Gbps, if desired (in conjunction with ZBT memory domains at 125 MHz). If both Gigabit ports are used in 2G-mode for this purpose, this provides a total of 4 Gbps of bandwidth between devices.

In addition to a standard back-to-back configuration of devices, the MVTX2603 also provides more powerful stacking alternatives:

- Unidirectional ring configuration. Up to 32 devices. Devices are connected by one Gigabit link, which can be accelerated to 2 Gbps, if desired. Flow control cannot be enabled in this configuration, because of the inherent inefficiency in sending flow control messages upstream in a unidirectional ring.



- Bidirectional ring configuration. Up to 32 devices. Devices are connected by two Gigabit links, forming two rings, one clockwise and one counter clockwise. The total outgoing bandwidth can be as much as 4 Gbps. Flow control may be enabled in this configuration. The outgoing direction of a packet (clockwise or counter clockwise) is selected using a hash key for load distribution. The hash key can be a function of source MAC address, destination MAC address, both MAC addresses, or source port. This configuration provides fault-tolerance when one of the stacking links fail.



- Cascade Stacking configuration. Up to 32 devices. Devices are connected to form a list configuration. Devices are connected by two Gigabit links, except the two devices at both ends, where one Gigabit port is used as an uplink port. Flow control may be enabled in this configuration.

4.0 Data Forwarding Protocol

4.1 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 two 64-bit buses, connected to two SRAM banks, A and B. 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 M ports and 8 classes for each of the two Gigabit ports – a total of 112 unicast queues.

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.

4.2 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 M ports. The queue with higher priority has room for 32 entries and the queue with lower priority has room for 64 entries. There are 4 multicast queues for each of the two Gigabit ports. The size of the queues are: 32 entries (higher priority queue), 32 entries, 32 entries and 64 entries (lower priority queue). There is one multicast queue for every two priority classes. For the 10/100 M ports to map the 8 transmit priorities into 2 multicast queues, the 2 LSB are discarded. For the Gigabit ports to map the 8 transmit priorities into 4 multicast queues, the 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.

5.0 Memory Interface

5.1 Overview

The MVTX2603 provides two 64-bit wide SRAM banks, SRAM Bank A and SRAM Bank B. Each DMA can read and write from both bank A and bank B. The following figure provides an overview of the MVTX2603 SRAM banks.

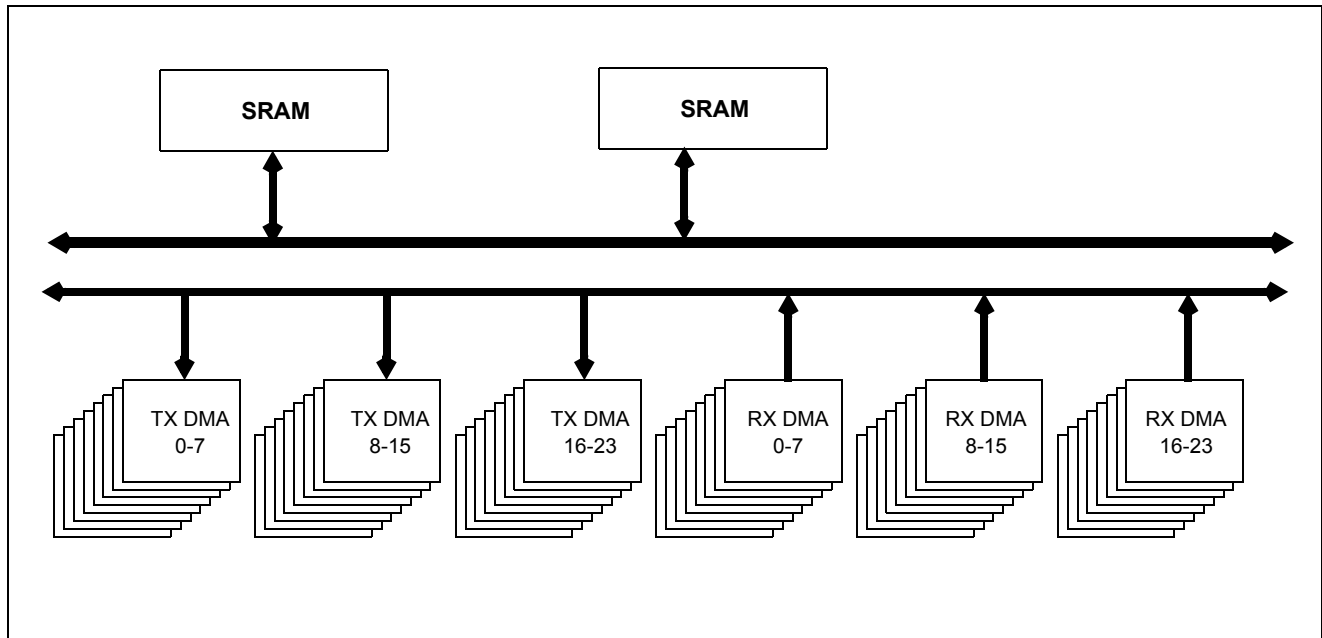


Figure 7 - SRAM Interface Block Diagram (DMAs for 10/100 Ports Only)

Because the bus for each bank is 64 bits wide, frames are broken into 8-byte granules, written to and read from memory. The first 8-byte granule gets written to Bank A, the second 8-byte granule gets written to Bank B and so on in alternating fashion. When reading frames from memory, the same procedure is followed, first from A, then from B and so on.

The reading and writing from alternating memory banks can be performed with minimal waste of memory bandwidth. What's the worst case? For any speed port, in the worst case, a 1-byte-long EOF granule gets written to Bank A. This means that a 7-byte segment of Bank A bandwidth is idle, and furthermore, the next 8-byte segment of Bank B bandwidth is idle, because the first 8 bytes of the next frame will be written to Bank A, not B. This scenario results in a maximum 15 bytes of waste per frame, which is always acceptable because the interframe gap is 20 bytes.

The CPU management port gets treated like any other port, reading and writing to alternating memory banks starting with Bank A. The VLAN Index Mapping Table and Mac Address Table are duplicated in Bank A and B. When the CPU writes an entry to the VLAN Index Mapping Table it has to write the same data in bank A and bank B. Search engine data is written to both banks in parallel. In this way, a search engine read operation can be performed by either bank at any time without a problem.

5.2 Memory Requirements

To speed up searching and decrease memory latency, the external MAC address database is duplicated in both memory banks. To support 64 K MAC address, 2 MB/bank memory is required.

Up to 2 K Ethernet frame buffers are supported and they will use 3 MB of memory. Each frame uses 1536 bytes. The maximum system memory requirement is 4 MB. If less memory is desired, the configuration can scale down.

Bank A	Bank B	Tagged-based VLAN	Max. Frame Buffers	Max MAC Address
1 M	1 M	Disable	1 K	32 K
2 M	2 M	Disable	2 K	64

Figure 8 - Memory Configuration

5.2.1 ZBT Support

The MVTX2603 supports Zero Bus Turnaround (ZBT) SRAM. ZBT is a synchronous SRAM architecture that is optimized for networking and telecommunications applications. It can significantly increase the switch's internal bandwidth when compared to standard Pipeline SyncBurst SRAM.

The ZBT architecture is optimized for switching and other applications with highly random READs and WRITES. ZBT SRAMs eliminate all idle cycles when turning the data bus around from a WRITE operation to a READ operation (or vice versa). This feature results in dramatic performance improvements in systems that have such traffic patterns (that is, frequent and random read and write access to the SRAM).

ZBT memory is only needed when the MVTX2603 is configured in 2G-mode operation. Some limitations of the 2G-mode and/or ZBT memory usage is highlighted below:

- tag-based VLAN is not support
- IP multicast switching is not supported
- only port-based priority mapping is supported

Please refer to the 2G-Mode and ZBT-SRAM Memory Application Note, MSAN-212, for further details.

5.3 Memory Configurations

The MVTX2603 supports pipelined SBRAM with 1 M and 2 M per bank configurations. For detail connection information, please reference the Memory Interface Application Note, MSAN-211.

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

Table 4 - Supported Memory Configurations (SBRAM Mode)

In the 2G-mode operation, the MVTX2603 supports a 4 M ZBT-SRAM configuration, with 2 M per domain (bank).

ZBT-SRAM Configurations	2 M per bank	Connections
Single Layer (Bootstrap pin TSTOUT13 = open)	Two 256 K x 32 ZBT/bank or One 256 K x 64 ZBT/bank	Connect ADS# to Layer 0 CS# pin
Double Layer (Bootstrap pin TSTOUT13 = pulled down)	Four 128 K x 32 ZBT/bank or Two 128 K x 64 ZBT/bank	Connect ADS# to Layer 0 CS# pin Connect 0E# to Layer 1 CS# pin

Table 5 - Supported Memory Configurations (ZBT Mode)

	Only Bank A		Bank A and Bank B		Bank A and Bank B
	1 M (SBRAM)	2 M (SBRAM)	1 M/bank (SBRAM)	2 M/bank (SBRAM)	2 M/bank (ZBT-SRAM)
ZL50415	X	X			
ZL50416	X	X			
ZL50417			X	X	
ZL50418			X	X	
MVTX2601	X	X			
MVTX2602	X	X			
MVTX2603			X	X	
MVTX2603 (Gigabit ports in 2G-mode)					X
MVTX2604			X	X	
MVTX2604 (Gigabit ports in 2G-mode)					X

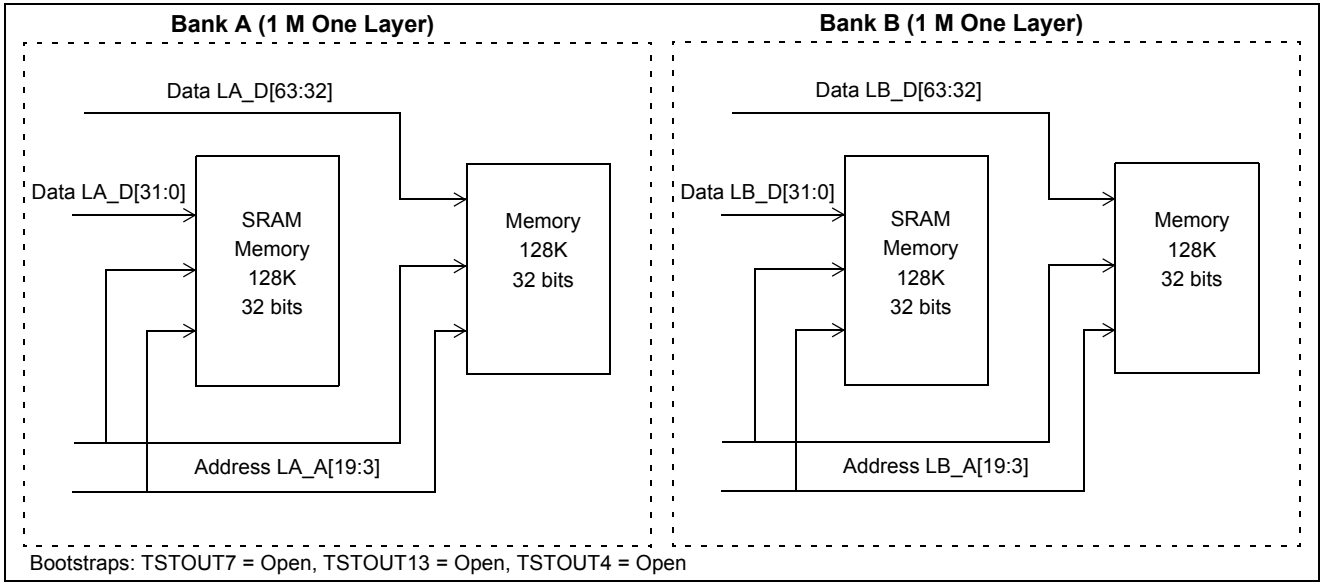


Figure 9 - Memory Configuration For 1 M/bank, 1 Layer

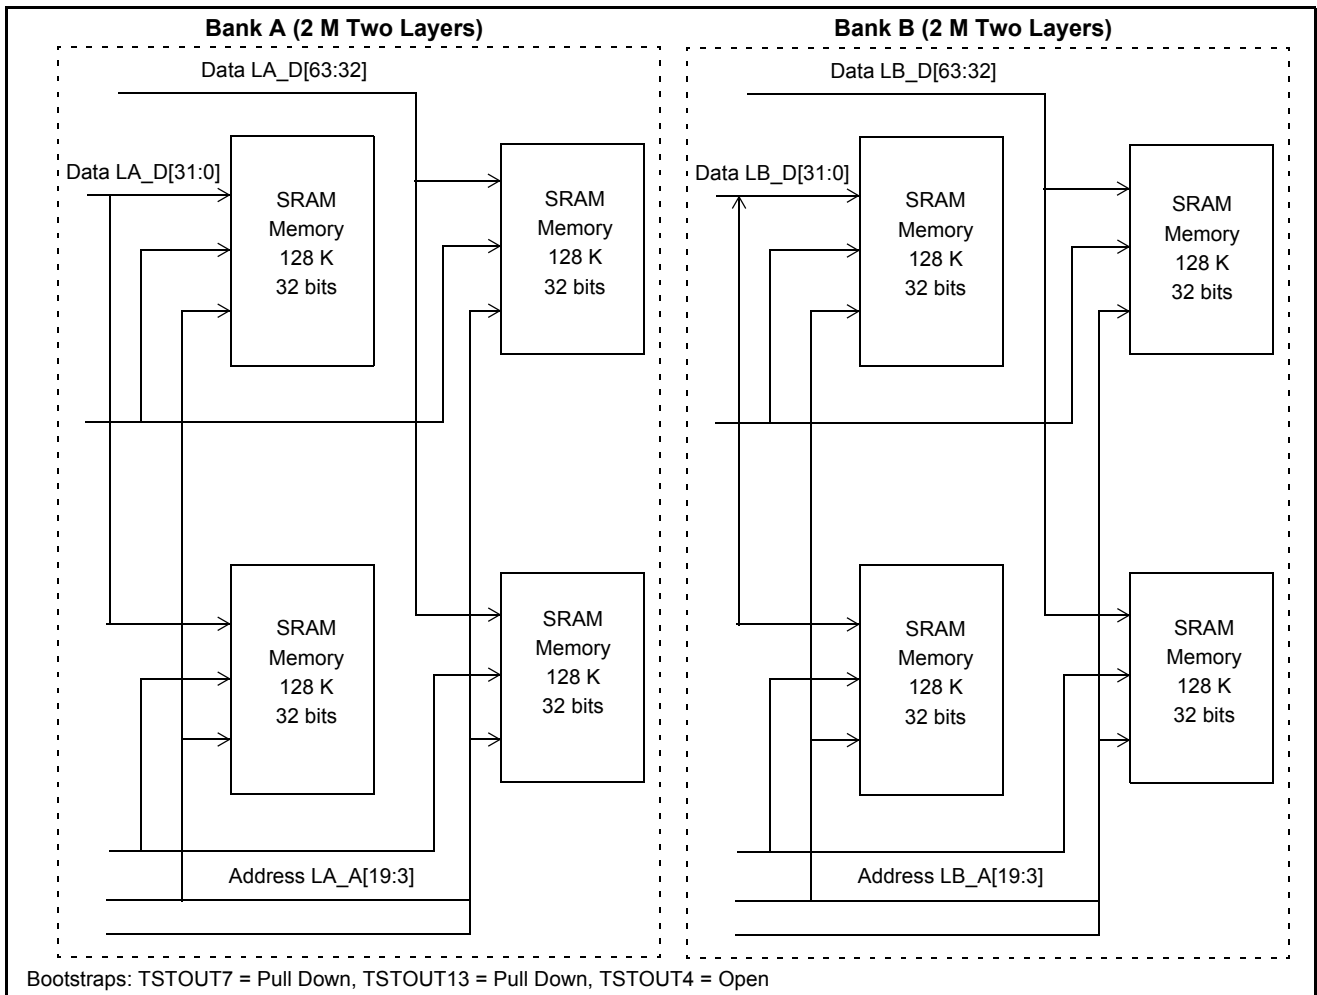


Figure 10 - Memory Configuration For 2 M/bank, 2 Layers

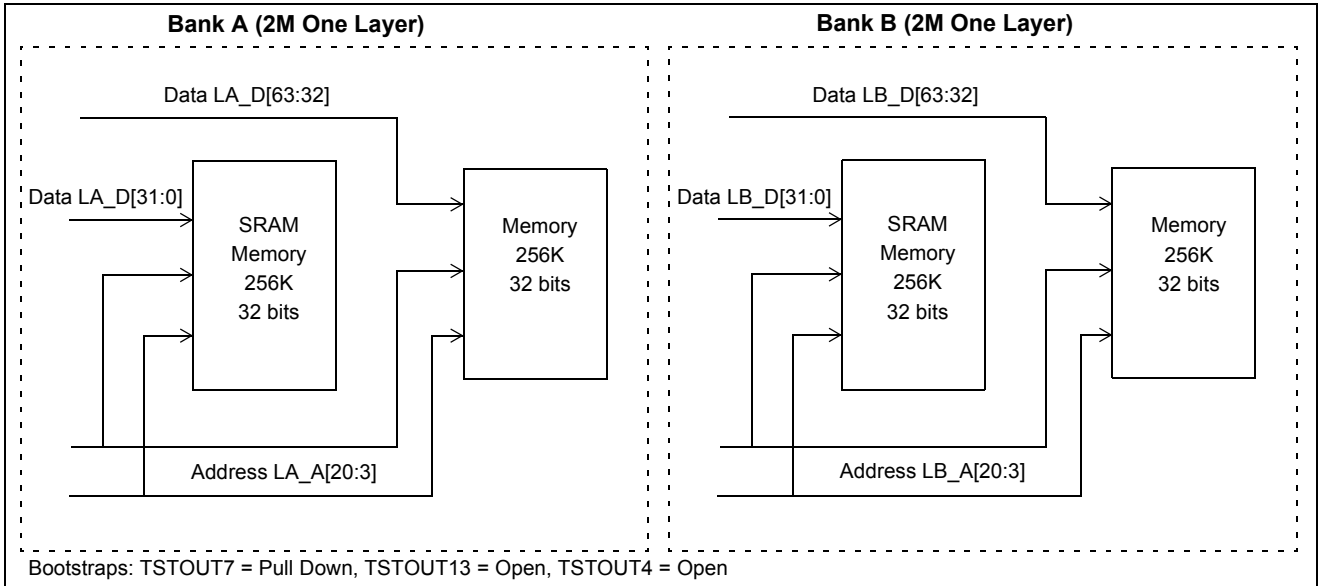


Figure 11 - Memory Configuration For 2 M/bank, 1 Layer

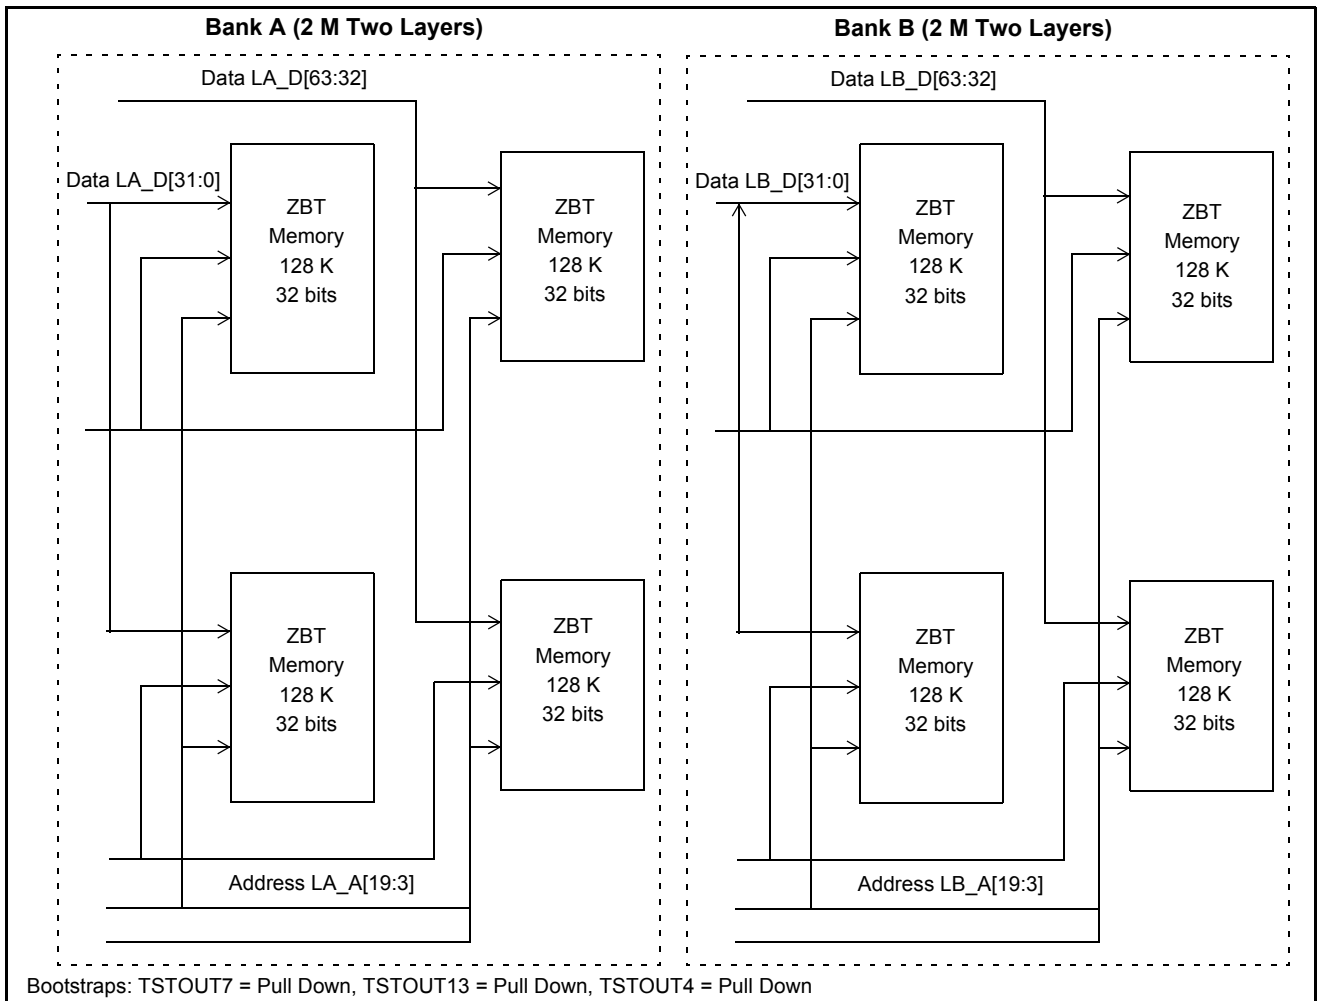


Figure 12 - ZBT Memory Configuration For 2 M/bank, 2 Layers

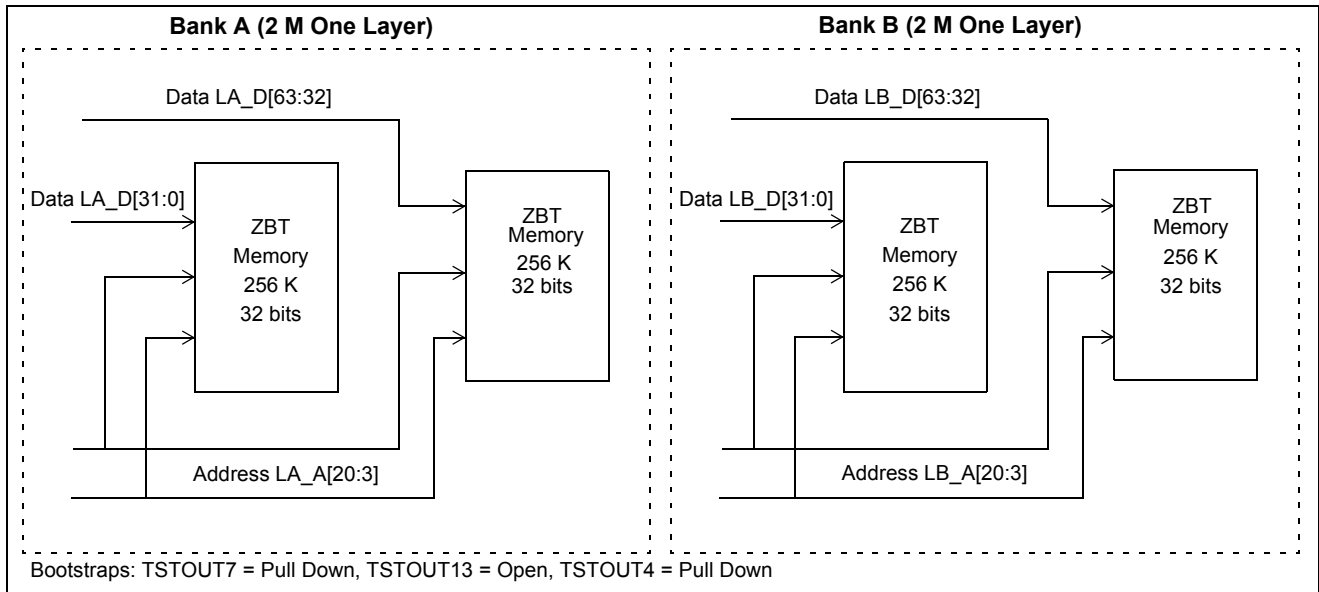


Figure 13 - ZBT Memory Configuration For 2 M/bank, 1 Layer

6.0 Search Engine

6.1 Search Engine Overview

The MVTX2603 search engine is optimized for high throughput searching, with enhanced features to support:

- Up to 64 K MAC addresses
- Port-based VLAN
- 3 groups of port trunking (1 for the two Gigabit ports and 2 others)
- Traffic classification into 4 (or 8 for Gigabit) transmission priorities and 2 drop precedence levels
- Flooding, Broadcast, Multicast Storm Control
- MAC address learning and aging

6.2 Basic Flow

Shortly after a frame enters the MVTX2603 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, and VLAN ID. 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, 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.

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

6.3 Search, Learning, and Aging

6.3.1 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.

In port-based VLAN mode, a bitmap is used to determine whether the frame should be forwarded to the outgoing port. The bitmap is not dynamic. Ports cannot enter and exit groups because of real-time learning made by a CPU.

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

6.3.2 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.

6.3.3 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.

6.4 Port--Based VLAN

An administrator can use the PVMAP registers to configure the MVTX2603 for port-based VLAN (See "Register Definition" on page 60.). 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 MVTX2603 determines the VLAN membership of each packet by noting the port on which it arrives. From there, the MVTX2603 determines which outgoing port(s) is/are eligible to transmit each packet or whether the packet should be discarded.

Port Registers	Destination Port Numbers Bit Map				
	26	...	2	1	0
Register for Port #0 PVMAP00_0[7:0] to PVMAP00_3[2:0]	0		1	1	0
Register for Port #1 PVMAP01_0[7:0] to PVMAP01_3[2:0]	0		1	0	1
Register for Port #2 PVMAP02_0[7:0] to PVMAP02_3[2:0]	0		0	0	0
...					
Register for Port #26 PVMAP26_0[7:0] to PVMAP26_3[2:0]	0		0	0	0

Table 6 - 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.

6.5 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 MVTX2603 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 MVTX2603, 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 MVTX2603 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.

6.5.1 Priority Classification Rule

Figure 14 shows the MVTX2603 priority classification rule.

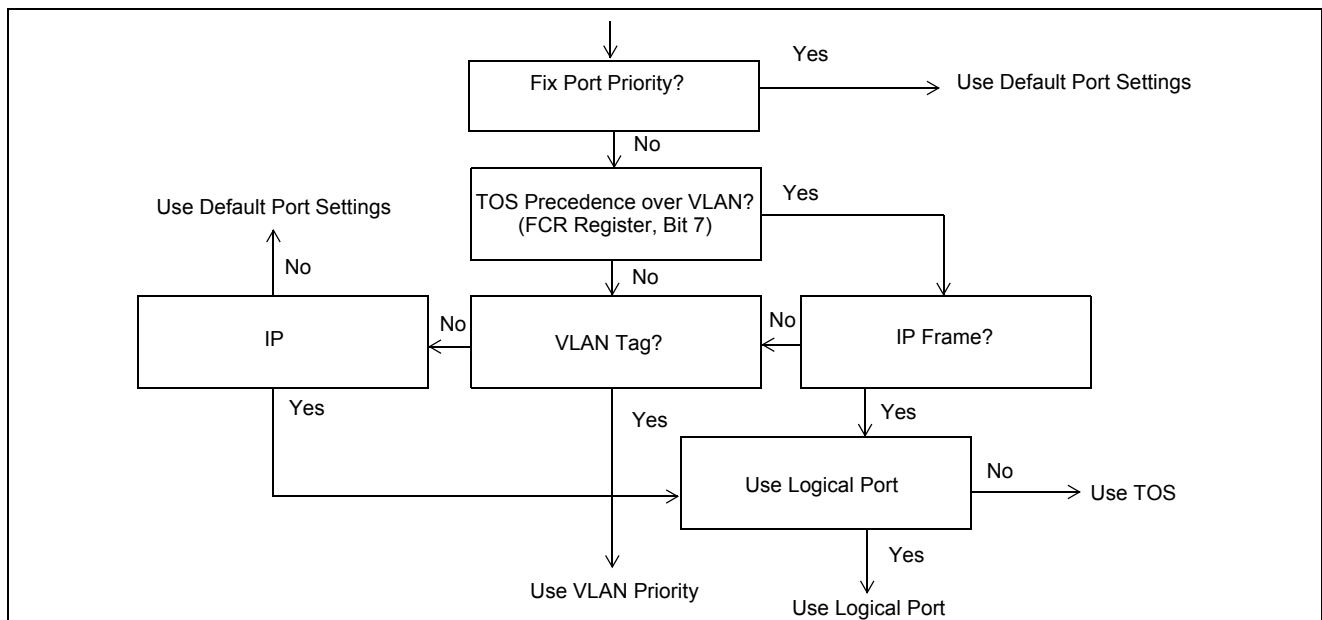


Figure 14 - Priority Classification Rule

7.0 Frame Engine

7.1 Data Forwarding Summary

When a frame enters the device at the RxMAC, the RxDMA will move the data from the MAC Rx FIFO 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 and a switch response is sent back to the Frame Engine. This response 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 4 TxSch Q for each 10/100 M port (and 8 per Gigabit port), 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 (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 M ports multicast queue 0 is associated with unicast queue 0 and multicast queue 1 is associated with unicast queue 2. For Gigabit ports multicast queue 0 is associated with unicast queue 0, multicast queue 1 with unicast queue 2, multicast queue 2 with unicast queue 4 and multicast queue 3 with unicast queue 6.

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

7.2 Frame Engine Details

This section briefly describes the functions of each of the modules of the MVTX2603 frame engine.

7.2.1 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. 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.

7.2.2 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.

7.2.3 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.

7.2.4 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.

7.2.5 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.

7.2.6 TxDMA

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

8.0 Quality of Service and Flow Control

8.1 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 7 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. Gigabit ports actually have eight total transmission priorities.

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 7 - 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 7 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 MVTX2603, each 10/100 M port will support four total classes and each Gigabit 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 7 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.

8.2 Four QoS Configurations

There are four basic pieces to QoS scheduling in the MVTX2603: 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 the tables below. For 10/100 M ports, the following registers select these modes:

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 8 - Four QoS Configurations for a 10/100 M Port

QOSC40 [7:6] and QOSC48 [7:6] select these modes for the first and second Gigabit ports, respectively.

	P7	P6	P5	P4	P3	P2	P1	P0
Op1 (default)	Delay Bound						BE	
Op2	SP		Delay Bound				BE	
Op3	SP		WFQ					
Op4	WFQ							

Table 9 - Four QoS Configurations for a Gigabit Port

The default configuration for a 10/100 M port is three delay-bounded queues and one best-effort queue. The delay bounds per class are 0.8 ms for P3, 3.2 ms for P2, and 12.8 ms for P1. For a Gigabit port, we have a default of six delay-bounded queues and two best-effort queues. The delay bounds for a Gigabit port are 0.16 ms for P7 and P6, 0.32 ms for P5, 0.64 ms for P4, 1.28 ms for P3, and 2.56 ms for P2. For a Gigabit port, where there are two best-effort queues, P1 has strict priority over P0. 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 M 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 MVTX2603, can have an adverse effect on all other classes' performance. For a Gigabit port, P7 and P6 are both SP classes and P7 has strict priority over P6. In this case, the delay bounds per class are 0.32 ms for P5, 0.64 ms for P4, 1.28 ms for P3, and 2.56 ms for P2.

The third configuration for a 10/100 M 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.

8.3 Delay Bound

In the absence of a sophisticated QoS server and signaling protocol, the MVTX2603 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.

8.4 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 MVTX2603, 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.

8.5 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 MVTX2603 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” (eight for Gigabit ports) 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 MVTX2603 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 M port (and queues P0 and P1 for a Gigabit 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.

8.6 Shaper

Although traffic shaping is not a primary function of the MVTX2603, the chip does implement a shaper for expedited forwarding (EF). Our goal in shaping is to control the peak and average rate of traffic exiting the MVTX2603. Shaping is limited to the two Gigabit ports only, and only to class P6 (the second highest priority). This means that class P6 will be the class used for EF traffic. If shaping is enabled for P6, then P6 traffic must be scheduled using strict priority. With reference to Table 7, only the middle two QoS configurations may be used.

Peak rate is set using a programmable whole number, no greater than 64. For example, if the setting is 32, then the peak rate for shaped traffic is $32/64 * 1000 \text{ Mbps} = 500 \text{ Mbps}$. Average rate is also a programmable whole number, no greater than 64 and no greater than the peak rate. For example, if the setting is 16, then the average rate for shaped traffic is $16/64 * 1000 \text{ Mbps} = 250 \text{ Mbps}$. As a consequence of the above settings in our example, shaped traffic will exit the MVTX2603 at a rate always less than 500 Mbps and averaging no greater than 250 Mbps. See Programming QoS Register application note for more information.

Also, when shaping is enabled, it is possible for a P6 queue to explode in length if fed by a greedy source. The reason is that a shaper is by definition not work-conserving; that is, it may hold back from sending a packet even if the line is idle. Though we do have global resource management, we do nothing to prevent this situation locally. We assume SP traffic is policed at a prior stage to the MVTX2603.

8.7 Rate Control

The MVTX2603 provides a rate control function on its 10/100 M ports. This rate control function applies to the outgoing traffic aggregate on each 10/100 M port. It provides a way of reducing the outgoing average rate below full wire speed. Note that the rate control function does not shape or manipulate any particular traffic class. Furthermore, though the average rate of the port can be controlled with this function, the peak rate will still be full line rate.

Two principal parameters are used to control the average rate for a 10/100 M port. A port's rate is controlled by allowing, on average, M bytes to be transmitted every N microseconds. Both of these values are programmable. The user can program the number of bytes in 8-byte increments and the time may be set in units of 10 ms.

The value of M/N will, of course, equal the average data rate of the outgoing traffic aggregate on the given 10/100 M port. Although there are many (M,N) pairs that will provide the same average data rate performance, the smaller the time interval N, the "smoother" the output pattern will appear.

In addition to controlling the average data rate on a 10/100 M port, the rate control function also manages the maximum burst size at wire speed. The maximum burst size can be considered the memory of the rate control mechanism; if the line has been idle for a long time, to what extent can the port "make up for lost time" by transmitting a large burst? This value is also programmable, measured in 8-byte increments.

Example: Suppose that the user wants to restrict Fast Ethernet port P's average departure rate to 32 Mbps – 32% of line rate – when the average is taken over a period of 10 ms. In an interval of 10 ms, exactly 40000 bytes can be transmitted at an average rate of 32 Mbps.

So how do we set the parameters? The rate control parameters are contained in an internal RAM block accessible through the CPU port (See Programming QoS Registers application note and Processor interface application note). The data format is shown below.

63:40	39:32	31:16	15:0
0	Time interval	Maximum burst size	Number of bytes

As we indicated earlier, the number of bytes is measured in 8-byte increments, so the 16-bit field “Number of bytes” should be set to $40000/8$, or 5000. In addition, the time interval has to be indicated in units of 10 ms. Though we want the average data rate on port P to be 32 Mbps when measured over an interval of 10 ms, we can also adjust the maximum number of bytes that can be transmitted at full line rate in any single burst. Suppose we wish this limit to be 12 kilobytes. The number of bytes is measured in 8-byte increments, so the 16-bit field “Maximum burst size” is set to $12000/8$, or 1500.

8.8 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 $N \geq 120$	$P3 \geq AKB$	$P2 \geq BKB$	$P1 \geq CKB$	X%	0%
Level 2 $N \geq 140$				Y%	Z%
Level 3 $N \geq 160$				100%	100%

Table 10 - 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, ZLAN-05, for more information.

8.9 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 MVTX2603. Our buffer management scheme is designed to divide the total buffer space into numerous reserved regions and one shared pool as shown in Figure 15 on page 54.

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 MVTX2603, 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 10/100M 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 M 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 ports — 26 ports for Ethernet. Two parameters can be set, one for the source port reservation for 10/100 M ports, and one for the source port reservation for Gigabit ports. These 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 M Ports

PRG- Port Reservation for Gigabit 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 15 - Buffer Partition Scheme Used to Implement Buffer Management

8.9.1 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.

8.10 Flow Control Basics

Because frame loss is unacceptable for some applications, the MVTX2603 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 MVTX2603, 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. This means 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 MVTX2603 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 MVTX2603’s approach to ensuring bounded delay and minimum bandwidth for high priority flows.

8.10.1 Unicast Flow Control

For unicast frames, flow control is triggered by source port resource availability. Recall that the MVTX2603’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 MVTX2603’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.

8.10.2 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 managed mode, per-VLAN flow control is used for multicast frames. In this case, flow control is triggered by congestion at the destination. How so? The MVTX2603 checks each destination to which a multicast packet is headed. For each destination port, the occupancy of the lowest-priority transmission multicast queue (measured in number of frames) is compared against a programmable congestion threshold. If congestion is detected at even one of the packet's destinations then Xoff is triggered.

In addition, each source port has a 26-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 26-bit vector is reset to zero.

The MVTX2603 also provides the option of disabling VLAN multicast flow control.

Note: If per-Port flow control is on, QoS performance will be affected.

8.11 Mapping to IETF DiffServ Classes

The mapping between priority classes discussed in this chapter and elsewhere is shown below.

MVTX2603	P7	P6	P5	P4	P3	P2	P1	P0
IETF	NM	EF	AF0	AF1	AF2	AF3	BE0	BE1

Table 11 - Mapping between MVTX2603 and IETF DiffServ Classes for Gigabit Ports

MVTX2603	P3	P2	P1	P0
IETF	NM+EF	AF0	AF1	BE0

Table 12 - Mapping between MVTX2603 and IETF DiffServ Classes for 10/100 M Ports

As Table 11 illustrates, P7 is used solely for network management (NM) frames. P6 is used for expedited forwarding service (EF). Classes P2 through P5 correspond to an assured forwarding (AF) group of size 4. Finally, P0 and P1 are two best effort (BE) classes.

For 10/100 M ports, the classes of Table 11 are merged in pairs, as shown in Table 12 — one class corresponding to NM+EF, two AF classes, and a single BE class.

Features of the MVTX2603 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 • Shaper for EF traffic on Gigabit ports • Option of strict priority scheduling • No dropping if admission controlled
Assured forwarding (AF)	<ul style="list-style-type: none"> • Four AF classes for Gigabit ports • 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"> • Two BE classes for Gigabit ports • 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 13 - MVTX2603 Features Enabling IETF DiffServ Standards

9.0 Port Trunking

9.1 Features and Restrictions

A port group (i.e., trunk) can include up to 4 physical ports but when using stack all of the ports in a group must be in the same MVTX2603.

The two Gigabit ports may also be trunked together. There are three trunk groups total including the option to trunk Gigabit ports.

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 MVTX2603 also provides a safe fail-over mode for port trunking automatically. If one of the ports in the trunking group goes down, the MVTX2603 will automatically redistribute the traffic over to the remaining ports in the trunk in unmanaged mode.

9.2 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 as specified in the Trunk_Hash registers.

9.3 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 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.

9.4 Unmanaged Trunking

In unmanaged mode, 3 trunk groups are supported. Groups 0 and 1 can trunk up to 4 10/100 M ports. Group 2 can trunk 2 Gigabit 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

Group 2	Port 25 (Gigabit 0)	Port 26 (Gigabit 1)
	✓	✓

In unmanaged mode, the trunks are individually enabled/disabled by controlling pin TRUNK0,1,2.

10.0 Port Mirroring

10.1 Port Mirroring Features

The received or transmitted data of any 10/100 M port in the MVTX2603 chip can be “mirrored” to any other port. We support two such mirrored source-destination pairs. A mirror port can not also serve as a data port.

Please refer to the Port Mirroring Application Note, MSAN-210, for further details.

10.2 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.

11.0 Register Definition

11.1 Register Description

Register	Description	CPU Addr (Hex)	R/W	I ² C Addr (Hex)	Default	Notes
0. Ethernet Port Control Registers (substitute n with port number (0..1Ah))						
ECR1Pn	Port Control Register 1 for Port n	000+2n	R/W	000+n	0C0	
ECR2Pn	Port Control Register 2 for Port n	001+2n	R/W	01B+n	000	
GGC	Extra GIGA bit control register	036	R/W	NA	000	
1. VLAN Control Registers (substitute n with port number (0..1Ah))						
AVTCL	VLAN Type Code Register Low	100	R/W	036	000	
AVTCH	VLAN Type Code Register High	101	R/W	037	081	
PVMAPn_0	Port n Configuration Register 0	102+4n	R/W	038+n	0FF	
PVMAPn_1	Port n Configuration Register 1	103+4n	R/W	053+n	0FF	
PVMAPn_2	Port n Configuration Register 2	104+4n	R/W	06E+n	0FF	
PVMAPn_3	Port n Configuration Register 3	105+4n	R/W	089+n	007	
PVMODE	VLAN Operating Mode	170	R/W	0A4	000	
2. 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	
3. CPU Port Configuration						
TX_AGE	Transmission Queue Aging Time	325	R/W	0A7	008	
4. Search Engine Configurations						
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	
5. 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	

Register	Description	CPU Addr (Hex)	R/W	I ² C Addr (Hex)	Default	Notes
TOSPML	TOS Priority Map Low	506	R/W	0B0	000	
TOSPMM	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	2M:008/ 4M:010	
MCC	Multicast Congestion Control	50D	R/W	0B7	050	
PR100	Port Reservation for 10/100 Ports	50E	R/W	0B8	2M:035/ 4M:036	
PRG	Port Reservation for Giga Ports	50F	R/W	0B9	2M:035/ 4M:058	
SFCB	Share FCB Size	510	R/W	0BA	2M:046/ 4M:064	
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	
QOSCn	QOS Control (n=0 - 5)	517-51C	R/W	0C1-0C6	000	
	QOS Control (n=12 - 23)	523-52E	R/W	0C7-0D2	000	
RDRC0	WRED Drop Rate Control 0	553	R/W	0FB	08F	
RDRC1	WRED Drop Rate Control 1	554	R/W	0FC	088	
USER_PORTn_LOW	User Define Logical Port n Low	580+2n	R/W	0D6+n	000	(n=0-7)
USER_PORTn_HIGH	User Define Logical Port n High	581+2n	R/W	0DE+n	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	

Register	Description	CPU Addr (Hex)	R/W	I ² C Addr (Hex)	Default	Notes
USER_PORT5:4_PRIORITY	User Define Logic Port 5 and 4 Priority	592	R/W	0E8	000	
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	
6. 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	NA	000	
MIIC1	MII Command Register 1	604	R/W	NA	000	
MIIC2	MII Command Register 2	605	R/W	NA	000	
MIIC3	MII Command Register 3	606	R/W	NA	000	
MIID0	MII Data Register 0	607	RO	NA	NA	
MIID1	MII Data Register 1	608	RO	NA	NA	
LED	LED Control Register	609	R/W	0F3	000	
DEVICE	Device id and test	60A	R/W	NA	000	
SUM	EEPROM Checksum Register	60B	R/W	0FF	000	

Register	Description	CPU Addr (Hex)	R/W	I ² C Addr (Hex)	Default	Notes
F. Device Configuration Register						
GCR	Global Control Register	F00	R/W	NA	000	
DCR	Device Status and Signature Register	F01	RO	NA	NA	
DCR1	Giga Port status	F02	RO	NA	NA	
DPST	Device Port Status Register	F03	R/W	NA	000	
DTST	Data read back register	F04	RO	NA	NA	
DA	Dead or Alive Register	FFF	RO	NA	DA	

11.2 Indirectly Accessed Registers

11.2.1 (Group 0 Address) MAC Ports Group

11.2.1.1 ECR1Pn: Port n Control Register 1

I²C Address 000+n; CPU Address:0000+2n (n = port number)

Accessed by CPU, serial interface and I²C (R/W)

7	6	5	4	0
SS	A-FC	Port Mode		

- Bit [0] 1 - Flow Control Disabled
0 - Flow Control Enabled (Default)
- Bit [1] 1 - Half Duplex - Only in 10/100 mode
0 - Full Duplex (Default)
- Bit [2] 1 - 10 Mbps
0 - 100 Mbps (Default)
- Bit [4:3] 00 - Enable Auto-Negotiation
This enables hardware state machine for auto-negotiation. (Default)
01 - Limited Disable Auto-Negotiation
This disables hardware state machine for speed auto-negotiation (use ECR1Pn[2:0] for configuration). Hardware will still poll PHY for link status.
10 - Force Link Down
Disable the port. Hardware does not talk to PHY.
11 - Force Link Up
The configuration in ECR1Pn[2:0] is used for (speed/duplex/flow control) setup. Hardware does not talk to PHY.
- Bit [5] Asymmetric Flow Control Enable.
0 - Disable asymmetric flow control (Default)
1 - Enable Asymmetric flow control
- When this bit is set and flow control is on (bit [0] = 0), the device does not send out flow control frames, but it's receiver interprets and processes flow control frames.
- Bit [7:6] SS - Spanning tree state (IEEE 802.1D spanning tree protocol)
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. (Default)

11.2.1.2 ECR2Pn: Port n Control Register 2

I²C Address: 01B+n; CPU Address:0001+2n (n = port number)

Accessed by CPU and serial interface (R/W)

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

- Bit [0]: Filter untagged frame
 0: Disable (Default)
 1: All untagged frames from this port are discarded or follow security option when security is enable
- Bit [1]: Filter Tag frame
 0: Disable (Default)
 1: All tagged frames from this port are discarded or follow security option when security is enable
- Bit [2]: Learning Disable
 0: Learning is enabled on this port (Default)
 1: Learning is disabled on this port
- Bit [3]: Must be '1'
- Bit [5:4] QOS mode selection. Determines which of the 4 sets of QoS settings is used for 10/100 ports.
- 00: select class byte limit set 0 and classes WFQ credit set 0 (Default)
 - 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

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 Programming QoS Registers Application Note, ZLAN-05.

- Bit[7:6] Security Enable. The MVTX2603 checks the incoming data for one of the following conditions:
- If the source MAC address of the incoming packet is in the MAC table and is defined as secure address but the ingress port is not the same as the port associated with the MAC address in the MAC table.
 - A MAC address is defined as secure when its entry at MAC table has static status and bit 0 is set to 1. MAC address bit 0 (the first bit transmitted) indicates whether the address is unicast or multicast. As source addresses are always unicast bit 0 is not used (always 0). MVTX2603 uses this bit to define secure MAC addresses.
 - If the port is set as learning disable and the source MAC address of the incoming packet is not defined in the MAC address table or the MAC address is not associated to the ingress port.
 - If the port is configured to filter untagged frames and an untagged frame arrives
 - If the port is configured to filter tagged frames and a tagged frame arrives
- If any one of the conditions is met, the packet is forwarded based on these setting.
- 00 – Disable port security, forward packets as usual. (Default)
 - 01 – Discard violating packets
 - 10 – Forward violating packets as usual and also to the CPU for inspection
 - 11 – Forward violating packets to the CPU for inspection

11.2.1.3 GGControl – Extra GIGA Port Control

CPU Address:h036

Accessed by CPU and serial interface (R/W)

7	6	5	4	3	2	1	0
DF	DI	MiiB	RstA	DF	DI	MiiA	RstA

- Bit [0]: Reset GIGA port 0
- 0: Normal operation (**default**)
 - 1: Reset Gigabit port 0. Normally used when a new Phy is connected (Hot swap).
- Bit [1]: GIGA port 0 use MII interface (10/100M)
- 0: Gigabit port operations at 1000 mode (**default**)
 - 1: Gigabit port operations at 10/100 mode
- Bit [2]: Device information insertion enable for Gigabit port 0
- 0: Disable preamble stack device ID insertion (**default**).
 - 1: Insert stack device ID into the preamble (must be enabled for ring mode).
- Bit [3]: GIGA port 0 direct flow control (MAC to MAC connection). The MVTX2603 supports direct flow control mechanism; the flow control frame is therefore not sent through the Gigabit port data path.
- 0: Direct flow control disabled (**default**)
 - 1: Direct flow control enabled
- Bit [4]: Reset GIGA port 1
- 0: Normal operation (**default**)
 - 1: Reset Gigabit port 1
- Bit [5]: GIGA port 1 use MII interface (10/100M)
- 0: Gigabit port operates at 1000 mode (**default**)
 - 1: Gigabit port operates at 10/100 mode
- Bit [6]: Device information attach enable for Gigabit port 1
- 0: Disable preamble stack device ID insertion (**default**)
 - 1: Insert stack device ID into the preamble (must be enabled for ring mode).
- Bit [7]: GIGA port 1 direct flow control (MAC to MAC connection). MVTX2603 supports direct flow control mechanism; the flow control frame is therefore not sent through the Gigabit port data path.
- 0: Direct flow control disabled (**default**)
 - 1: Direct flow control enabled

11.2.2 (Group 1 Address) VLAN Group

11.2.2.1 AVTCL – VLAN Type Code Register Low

I²C Address 036; CPU Address:h100

Accessed by CPU, serial interface and I²C (R/W)

Bit [7:0]: VLANType_LOW: Lower 8 bits of the VLAN type code (**Default 00**)

11.2.2.2 AVTCH – VLAN Type Code Register High

I²C Address 037; CPU Address:h101

Accessed by CPU, serial interface and I²C (R/W)

Bit [7:0]: VLANType_HIGH: Upper 8 bits of the VLAN type code (**Default 0x81**)

11.2.2.3 PVMAP00_0 – Port 00 Configuration Register 0

I²C Address 038, CPU Address:h102

Accessed by CPU, serial interface and I²C (R/W)

In Port-based VLAN Mode

Bit [7:0]: VLAN Mask for ports 7 to 0 (Default 0xFF)

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, 2 and 3 to form a 27 bit mask to all egress ports.

11.2.2.4 PVMAP00_1 – Port 00 Configuration Register 1

I²C Address h53, CPU Address:h103

Accessed by CPU, serial interface and I²C (R/W)

In Port-based VLAN Mode

Bit [7:0]: VLAN Mask for ports 15 to 8 (Default 0xFF)

11.2.2.5 PVMAP00_2 – Port 00 Configuration Register 2

I²C Address h6E, CPU Address:h104

Accessed by CPU, serial interface and I²C (R/W)

In Port-based VLAN Mode

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

11.2.2.6 PVMAP00_3 – Port 00 Configuration Register 3

I²C Address h89, CPU Address:h105

Accessed by CPU, serial interface and I²C (R/W)

In Port-based VLAN Mode

Bit [0]: Reserved (**Default 1**)

Bit [2:1]: VLAN Mask for ports 26 to 25 (Gigabit ports) (Default 3)

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. Used when Bit[7]=1 (**Default 0**)

- 0 - Discard Priority Level 0 (Lowest)
- 1 - Discard Priority Level 1(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 or Logical Port.
- 1 Transmit Priority and Discard Priority are based on values programmed in bit [6:3]

11.2.2.7 PVMAPnn_0,1,2,3 – Port nn Configuration Registers

PVMAP01_0,1,2,3 I²C Address h39,54,6F,8A; CPU Address:h106,107,108,109 (Port 1)

PVMAP02_0,1,2,3 I²C Address h3A,55,70,8B; CPU Address:h10A, 10B, 10C, 10D (Port 2)

PVMAP03_0,1,2,3 I²C Address h3B,56,71,8C; CPU Address:h10E, 10F, 110, 111 (Port 3)

PVMAP04_0,1,2,3 I²C Address h3C,57,72,8D; CPU Address:h112, 113, 114, 115 (Port 4)

PVMAP05_0,1,2,3 I²C Address h3D,58,73,8E; CPU Address:h116, 117, 118, 119 (Port 5)

PVMAP06_0,1,2,3 I²C Address h3E,59,74,8F; CPU Address:h11A, 11B, 11C, 11D (Port 6)

PVMAP07_0,1,2,3 I²C Address h3F,5A,75,90; CPU Address:h11E, 11F, 120, 121 (Port 7)

PVMAP08_0,1,2,3 I²C Address h40,5B,76,91; CPU Address:h122, 123, 124, 125 (Port 8)

PVMAP09_0,1,2,3 I²C Address h41,5C,77,92; CPU Address:h126, 127, 128, 129 (Port 9)

PVMAP10_0,1,2,3 I²C Address h42,5D,78,93; CPU Address:h12A, 12B, 12C, 12D (Port 10)

PVMAP11_0,1,2,3 I²C Address h43,5E,79,94; CPU Address:h12E, 12F, 130, 131 (Port 11)

PVMAP12_0,1,2,3 I²C Address h44,5F,7A,95; CPU Address:h132, 133, 134, 135 (Port 12)

PVMAP13_0,1,2,3 I²C Address h45,60,7B,96; CPU Address:h136, 137, 138, 139 (Port 13)

PVMAP14_0,1,2,3 I²C Address h46,61,7C,97; CPU Address:h13A, h13B, 13C, 13D (Port 14)

PVMAP15_0,1,2,3 I²C Address h47,62,7D,98; CPU Address:h13E, 13F, 140, 141 (Port 15)

PVMAP16_0,1,2,3 I²C Address h48,63,7E,99; CPU Address:h142, 143, 144, 145 (Port 16)

PVMAP17_0,1,2,3 I²C Address h49,64,7F,9A; CPU Address:h146, 147, 148, 149 (Port 17)

PVMAP18_0,1,2,3 I²C Address h4A,65,80,9B; CPU Address:h14A, 14B, 14C, 14D (Port 18)

PVMAP19_0,1,2,3 I²C Address h4B,66,81,9C; CPU Address:h14E, 14F, 150, 151 (Port 19)

PVMAP20_0,1,2,3 I²C Address h4C,67,82,9D; CPU Address:h152, 153, 154, 155 (Port 20)

PVMAP21_0,1,2,3 I²C Address h4D,68,83,9E; CPU Address:h156, 157, 158, 159 (Port 21)

PVMAP22_0,1,2,3 I²C Address h4E,69,84,9F; CPU Address:h15A, 15B, 15C, 15D (Port 22)

PVMAP23_0,1,2,3 I²C Address h4F,6A,85,A0; CPU Address:h15E, 15F, 160, 161 (Port 23)

PVMAP25_0,1,2,3 I²C Address h51,6C,87,A2; CPU Address:h166, 167, 168, 169 (Port 25 - Gigabit port 0)

PVMAP26_0,1,2,3 I²C Address h52,6D,88,A3; CPU Address:h16A, 16B, 16C, 16D (Port 26 - Gigabit port 1)

11.2.2.8 PVMODE

I²C Address: h0A4, CPU Address:h170

Accessed by CPU, serial interface (R/W)

7	6	5	4	3	2	1	0
MAC05	MMA	STP	SM0	rPCS	DF	SL	Vmod


- Bit [0]:
- VLAN Mode (Default = 0)
 - 1 Tagged-based VLAN Mode
 - 0 Port-based VLAN Mode
- Bit [1]:
- Slow learning (Default = 0)
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]:
- Disable Reset PCS (Default = 0)
 - 0: Enable reset PCS. PCS FIFO will be reset when received a PCS symbol error.
 - 1: Disable reset PCS
- Bit [4]:
- Support MAC address 0 (Default = 0)
 - 0: MAC address 0 is not learned.
 - 1: MAC address 0 is learned.
- Bit [5]:
- Disable IEEE multicast control frame (0180C2000000 to 0180C200000F) to CPU in managed mode (Default = 0)
 - 0: Packet is forwarded to CPU
 - 1: Packet is forwarded as multicast
- Bit [6]:
- Multiple MAC addresses (Default = 0)
 - 0: Single MAC address is assigned to CPU. Registers MAC0 to MAC5 are used to program the CPU MAC address.
 - 1: One block of 32 MAC addresses are assigned to CPU. The block is defined in an increase way from the MAC address programmed in registers MAC0 to MAC5.
- Bit [7]:
- Disable registers MAC 5 – 0 (CPU MAC address) in comparison with Ethernet frame destination MAC address. When disable, unicast frames are not forward to CPU. (Default = 0)
 - 1: Disable
 - 0: Enable

11.2.3 (Group 2 Address) Port Trunking Groups

11.2.3.1 TRUNK0_MODE– Trunk group 0 mode

I²C Address h0A5; CPU Address:203

Accessed by CPU, serial interface and I²C (R/W)




- Bit [1:0]:
- Port selection in unmanaged mode. Input pin TRUNK0 enable/disable trunk group 0 in unmanaged mode.
 - 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

11.2.3.2 TRUNK1_MODE – Trunk group 1 mode

I²C Address h0A6; CPU Address:20B

Accessed by CPU, serial interface and I²C (R/W)




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

11.2.3.3 (Group 3 Address) CPU Port Configuration Group TX_AGE – Tx Queue Aging timer

I²C Address: h07; CPU Address: h324

Accessed by CPU, serial interface (RW)



Bit [5:0]: Unit of 100ms (**Default 8**)

Disable transmission queue aging if value is zero. Aging timer for all ports and queues.

This register must be set to 0 for 'No Packet Loss Flow Control Test'.

11.2.4 (Group 4 Address) Search Engine Group

11.2.4.1 AGETIME_LOW – MAC address aging time Low

I²C Address h0A8; CPU Address: h400

Accessed by CPU, serial interface and I²C (R/W)

The MVTX2603 removes the MAC address from the data base and sends a Delete MAC Address Control Command to the CPU. MAC address aging is enable/disable by boot strap TSTOUT9.

Bit [7:0] Low byte of the MAC address aging timer.

11.2.4.2 AGETIME_HIGH –MAC address aging time High

I²C Address h0A9; CPU Address h401

Accessed by CPU, 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 entries in the memory X 100µsec). Number of MAC entries = 32K when 1 MB is used per Bank. Number of entries = 64K when 2 MB is used per Bank.

11.2.5 (Group 5 Address) Buffer Control/QOS Group

11.2.5.1 FCBAT – FCB Aging Timer

I²C Address h0AA; CPU Address: h500

Accessed by CPU, serial interface and I²C (R/W)



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

11.2.5.2 QOSC – QOS Control

I²C Address h0AB; CPU Address:h501

Accessed by CPU, serial interface and I²C (R/W)

7	6	5	4	3	1	0
Tos-d	Tos-p	PMcq	VF1c			L

- 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]:
- Select processor multicast queue size
 - 0 = 16 entries
 - 1 = 64 entries
- 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

11.2.5.3 FCR – Flooding Control Register

I²C Address h0AC; CPU Address:h502

Accessed by CPU, serial interface and I²C (R/W)

7	6	4	3	0
Tos	TimeBase	U2MR		

- 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]: Time Base: (Default = 000)

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.

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

11.2.5.4 AVPML – VLAN Tag Priority Map

I²C Address h0AD; CPU Address:h503

Accessed by CPU, serial interface and I²C (R/W)



Registers AVPML, AVPMM and AVPMH allow the eight VLAN Tag priorities to map into eight Internal level transmit priorities. Under the Internal transmit priority, seven is the 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 packet VLAN priority 0 into Internal transmit priority 7. The new priority is used inside the MVTX2603. 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**)

11.2.5.5 AVPMM – VLAN Priority Map

I²C Address h0AE, CPU Address:h504

Accessed by CPU, 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)**

11.2.5.6 AVPMH – VLAN Priority Map

I²C Address h0AF, CPU Address:h505

Accessed by CPU, 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)**

11.2.5.7 TOSPML – TOS Priority Map

I²C Address h0B0, CPU Address:h506

Accessed by CPU, 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)**

11.2.5.8 TOSPM – TOS Priority Map

I²C Address h0B1, CPU Address:h507

Accessed by CPU, serial interface and I²C (R/W)




Map TOS field in IP packet into eight 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)**

11.2.5.9 TOSPMH – TOS Priority Map

I²C Address h0B2, CPU Address:h508

Accessed by CPU, serial interface and I²C (R/W)




Map TOS field in IP packet into eight 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)**

11.2.5.10 AVDM – VLAN Discard Map

I²C Address h0B3, CPU Address:h509

Accessed by CPU, 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)**

11.2.5.11 TOSDML – TOS Discard Map

I²C Address h0B4, CPU Address:h50A

Accessed by CPU, serial interface and I²C (R/W)




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)

11.2.5.12 BMRC - Broadcast/Multicast Rate Control

I²C Address h0B5, CPU Address:h50B)

Accessed by CPU, serial interface and I²C (R/W)




This broadcast and multicast rate defines for each port, the number of packets 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. Time base is based on register FCR [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)

11.2.5.13 UCC – Unicast Congestion Control

I²C Address h0B6, CPU Address: 50C

Accessed by CPU, 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/bank or h08 for 1 MB/bank)
-------------	--

11.2.5.14 MCC – Multicast Congestion Control

I²C Address h0B7, CPU Address: 50D

Accessed by CPU, serial interface and I²C (R/W)

7	5	4	0
FC reaction period		Multicast congest threshold	

Bit [4:0]: In multiples of two frames (granularity). Used for triggering MC flow control when destination port's multicast best effort queue reaches MCC threshold.(Default 0x10)

Bit [7:5]: Flow control reaction period (Default 2) Granularity 4uSec.

11.2.5.15 PR100 – Port Reservation for 10/100 ports

I²C Address h0B8, CPU Address 50E

Accessed by CPU, serial interface and I²C (R/W)

7	4	3	0
Buffer low threshold		SP Buffer reservation	

Bit [3:0]: Per source port buffer reservation.

Define the space in the FDB reserved for each 10/100 port and CPU. 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 reaches UCC threshold, shared pool is all used and source port reservation is at or below the PR100[7:4] level. Also the threshold for initiating UC flow control.

- Default:
 - h36 for 24+2 configuration with memory 2 MB/bank;
 - h24 for 24+2 configuration with 1MB/bank;

11.2.5.16 PRG – Port Reservation for Giga ports

I²C Address h0B9, CPU Address 50F

Accessed by CPU, serial interface and I²C (R/W)

7	4	3	0
Buffer low threshold		SP buffer reservation	

Bit [3:0]: Per source port buffer reservation.

Define the space in the FDB reserved for each Gigabit port. Expressed in multiples of 16 packets. For each packet 1536 bytes are reserved in the memory.


Bits [7:4]: Expressed in multiples of 16 packets. Threshold for dropping all best effort frames when destination port best effort queues reach UCC threshold, shared pool is all used and source port reservation is at or below the PRG[7:4] level. Also the threshold for initiating UC flow control.

- Default:
 - h58 for memory 2 MB/bank;
 - h35 for 1 MB/bank;

11.2.5.17 SFCB – Share FCB Size

I²C Address h0BA), CPU Address 510

Accessed by CPU, serial interface and I²C (R/W)




Bits [7:0]: Expressed in multiples of 4 packets. Buffer reservation for shared pool.

- Default:
 - h64 for 24+2 configuration with memory of 2 MB/bank;
 - h14 for 24+2 configuration with memory of 1 MB/bank;

11.2.5.18 C2RS – Class 2 Reserve Size

I²C Address h0BB, CPU Address 511

Accessed by CPU, serial interface and I²C (R/W)




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





11.2.5.19 C3RS – Class 3 Reserve Size

I²C Address h0BC, CPU Address 512

Accessed by CPU, serial interface and I²C (R/W)



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

11.2.5.20 C4RS – Class 4 Reserve SizeI²C Address h0BD, CPU Address 513Accessed by CPU, serial interface and I²C (R/W)Buffer reservation for class 4. Granularity 1. **(Default 0)****11.2.5.21 C5RS – Class 5 Reserve Size**I²C Address h0BE; CPU Address 514Accessed by CPU, serial interface and I²C (R/W)Buffer reservation for class 5. Granularity 1. **(Default 0)****11.2.5.22 C6RS – Class 6 Reserve Size**I²C Address h0BF; CPU Address 515Accessed by CPU, serial interface and I²C (R/W)Buffer reservation for class 6 (second highest priority). Granularity 1. **(Default 0)****11.2.5.23 C7RS – Class 7 Reserve Size**I²C Address h0C0; CPU Address 516Accessed by CPU, serial interface and I²C (R/W)Buffer reservation for class 7 (highest priority). Granularity 1. **(Default 0)****11.2.5.24 QOSC00~02 - Classes Byte Limit Set 0**Accessed by CPU; 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. There are four such sets of values A-C specified in Classes Byte Limit Set 0, 1, 2, and 3. For CPU port A-C values are defined using register CPUQOSC1, 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.

11.2.5.25 QOSC03~05 - Classes Byte Limit Set 1

Accessed by CPU, 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 Drop (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.

11.2.5.26 QOSC12~17 - Classes Byte Limit Giga Port 1

Accessed by CPU, serial interface and I²C (R/W):

F - QOSC12 – BYTE_C2_G1 (I²C Address h0C7, CPU Address 523)

E - QOSC13 – BYTE_C3_G1 (I²C Address h0C8, CPU Address 524)

D - QOSC14 – BYTE_C4_G1 (I²C Address h0C9, CPU Address 525)

C - QOSC15 – BYTE_C5_G1 (I²C Address h0CA, CPU Address 526)

B - QOSC16 – BYTE_C6_G1 (I²C Address h0CB, CPU Address 527)

A - QOSC17 – BYTE_C7_G1 (I²C Address h0CC, CPU Address 528)

QOSC12 through QOSC17 represent the values A-F for Gigabit port 1. They are per-queue byte thresholds for random early drop. QOSC17 represents A, and QOSC12 represents F.

Granularity when Delay bound is used: QOSC17 and QOSC16: 256 bytes, QOSC15 and QOSC14: 512 bytes, QOSC13 and QOSC12: 1024 bytes.

Granularity when WFQ is used: QOSC17 to QOSC12: 1024 bytes

11.2.5.27 QOSC18~23 - Classes Byte Limit Giga Port 2

Accessed by CPU, serial interface and I²C (R/W)

F - QOSC18 – BYTE_C2_G2 (I²C Address h0CD, CPU Address 529)

E - QOSC19 – BYTE_C3_G2 (I²C Address h0CE, CPU Address 52a)

D - QOSC20 – BYTE_C4_G2 (I²C Address h0CF, CPU Address 52b)

C - QOSC21 – BYTE_C5_G2 (I²C Address h0D0, CPU Address 52c)

B - QOSC22 – BYTE_C6_G2 (I²C Address h0D1, CPU Address 52d)

A - QOSC23 – BYTE_C7_G2 (I²C Address h0D2, CPU Address 52e)

QOSC18 through QOSC23 represent the values A-F for Gigabit port 2. They are per-queue byte thresholds for random early drop. QOSC23 represents A, and QOSC18 represents F.

Granularity when Delay bound is used: QOSC23 and QOSC22: 256 bytes, QOSC21 and QOSC20: 512 bytes, QOSC19 and QOSC18: 1024 bytes.

Granularity when WFQ is used: QOSC23 to QOSC18: 1024 bytes

11.2.5.28 QOSC40~47 - Classes WFQ Credit Port G1

Accessed by CPU and serial interface

W0 - QOSC40[5:0] - CREDIT_C0_G1(CPU Address 53f)

[7:6]: Priority service type. Option 1 to 4.

W1 - QOSC41[5:0] - CREDIT_C1_G1 (CPU Address 540)

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

[6]: Flow control pause best effort traffic only

W2 - QOSC42[5:0] - CREDIT_C2_G1 (CPU Address 541)

W3 - QOSC43[5:0] - CREDIT_C3_G1 (CPU Address 542)

W4 - QOSC44[5:0] - CREDIT_C4_G1 (CPU Address 543)

W5 - QOSC45[5:0] - CREDIT_C5_G1 (CPU Address 544)

W6 - QOSC46[5:0] - CREDIT_C6_G1 (CPU Address 545)

W7 - QOSC47[5:0] - CREDIT_C7_G1 (CPU Address 546)

QOSC40 through QOSC47 represents the set of WFQ parameters for Gigabit port 24. The granularity of the numbers is 1 and their sum must be 64. QOSC47 corresponds to W7 and QOSC40 corresponds to W0. In the 2G trunk configuration, the sum of all values QOSC40 through QOSC47 must be equal to 128.

11.2.5.29 QOSC48~55 - Classes WFQ Credit Port G2

Accessed by CPU and serial interface

W0 - QOSC48[5:0] - CREDIT_C0_G2(CPU Address 547)

[7:6]: Priority service type. Option 1 to 4

W1 - QOSC49[5:0] - CREDIT_C1_G2(CPU Address 548)

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

[6]: Flow control pause best effort traffic only

W2 - QOSC50[5:0] - CREDIT_C2_G2(CPU Address 549)

W3 - QOSC51[5:0] - CREDIT_C3_G2(CPU Address 54a)

W4 - QOSC52[5:0] - CREDIT_C4_G2(CPU Address 54b)

W5 - QOSC53[5:0] - CREDIT_C5_G2(CPU Address 54c)

W6 - QOSC54[5:0] - CREDIT_C6_G2(CPU Address 54d)

W7 - QOSC55[5:0] - CREDIT_C7_G2(CPU Address 54e)

QOSC48 through QOSC55 represents the set of WFQ parameters for Gigabit port 2. The granularity of the numbers is 1 and their sum must be 64. QOSC55 corresponds to W7 and QOSC48 corresponds to W0. In the 2G trunk configuration, the sum of all values QOSC48 through QOSC55 must be equal to 128.

11.2.5.30 QOSC56~57 - Class 6 Shaper Control Port G1

Accessed by CPU and serial interface

QOSC56[5:0] - TOKEN_RATE_G1 (CPU Address 54f).

Programs the average rate for gigabit port 1. When equal to 0, shaper is disable. Granularity is 1.

QOSC57[7:0] - TOKEN_LIMIT_G1 (CPU Address 550).

Programs the maximum counter for gigabit port 1. Granularity is 16 bytes.

Shaper is implemented to control the peak and average rate for outgoing traffic with priority 6 (queue 6). Shaper is limited to gigabit ports and queue P6 when it is in strict priority. QOSC41 programs the peak rate for gigabit port 1. See Programming QoS Registers Application Note for more information.

11.2.5.31 QOSC58~59 - Class 6 Shaper Control Port G2

Accessed by CPU and serial interface

QOSC58[5:0] – TOKEN_RATE_G2 (CPU Address 551).

Programs the average rate for gigabit port 2. When equal to 0, shaper is disabled. Granularity is 1.

QOSC59[7:0] – TOKEN_LIMIT_G2 (CPU Address 552).

Programs the maximum counter for gigabit port 2. Granularity is 16 bytes.

Shaper is implemented to control the peak and average rate for outgoing traffic with priority 6 (queue 6). Shaper is limited to gigabit ports and queue P6 when it is in strict priority. QOSC49 programs the peak rate for gigabit port 2. See Programming QoS Register application note for more information.

11.2.5.32 RDRC0 – WRED Rate Control 0

I²C Address 0FB, CPU Address 553

Accessed by CPU, Serial Interface and I²C (R/W)

7	4	3	0
X Rate		Y Rate	

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

11.2.5.33 RDRC1 – WRED Rate Control 1

I²C Address 0FC, CPU Address 554

Accessed by CPU, Serial Interface and I²C (R/W)

7	4	3	0
Z Rate		B Rate	

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 Registers application note for more information

User Defined Logical Ports and Well Known Ports

The MVTX2603 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:

- 23
- 512
- 6000
- 443
- 111
- 22555
- 22
- 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.

11.2.5.34 USER_PORT0~7)_L/H – USER DEFINE LOGICAL PORT (0~7)

USER_PORT0_L/H - I²C Address h0D6 + 0DE; CPU Address 580(Low) + 581(high)

USER_PORT1_L/H - I²C Address h0D7 + 0DF; CPU Address 582 + 583

USER_PORT2_L/H - I²C Address h0D8 + 0E0; CPU Address 584 + 585

USER_PORT3_L/H - I²C Address h0D9 + 0E1; CPU Address 586 + 587


USER_PORT4_L/H - I²C Address h0DA + 0E2; CPU Address 588 + 589

USER_PORT5_L/H - I²C Address h0DB + 0E3; CPU Address 58A + 58B

USER_PORT6_L/H - I²C Address h0DC + 0E4; CPU Address 58C + 58D

USER_PORT7_L/H - I²C Address h0DD + 0E5; CPU Address 58E + 58F

Accessed by CPU, serial interface and I²C (R/W)



(Default 00) This register is duplicated eight times from PORT 0 through PORT 7 and allows the CPU to define eight separate ports.

11.2.5.35 USER_PORT_[1:0]_PRIORITY - User Define Logic Port 1 and 0 PriorityI²C Address h0E6, CPU Address 590Accessed by CPU, serial interface and I²C (R/W)

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

The chip allows the CPU to define 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)

11.2.5.36 USER_PORT_[3:2]_PRIORITY - User Define Logic Port 3 and 2 PriorityI²C Address h0E7, CPU Address 591Accessed by CPU, serial interface and I²C (R/W)

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

11.2.5.37 USER_PORT_[5:4]_PRIORITY - User Define Logic Port 5 and 4 PriorityI²C Address h0E8, CPU Address 592Accessed by CPU, serial interface and I²C (R/W)

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

(Default 00)

11.2.5.38 USER_PORT_[7:6]_PRIORITY - USER DEFINE LOGIC PORT 7 AND 6 PRIORITYI²C Address h0E9, CPU Address 593Accessed by CPU, serial interface and I²C (R/W)

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

(Default 00)

11.2.5.39 USER_PORT_ENABLE[7:0] – User Define Logic 7 to 0 Port EnablesI²C Address h0EA, CPU Address 594Accessed by CPU, serial interface and I²C (R/W)

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

(Default 00)

11.2.5.40 WELL_KNOWN_PORT[1:0]_PRIORITY- Well Known Logic Port 1 and 0 PriorityI²C Address h0EB, CPU Address 595Accessed by CPU, 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)

11.2.5.41 WELL_KNOWN_PORT[3:2]_PRIORITY- Well Known Logic Port 3 and 2 PriorityI²C Address h0EC, CPU Address 596Accessed by CPU, 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)

11.2.5.42 WELL_KNOWN_PORT [5:4]_PRIORITY- Well Known Logic Port 5 and 4 PriorityI²C Address h0ED, CPU Address 597Accessed by CPU, 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 remote procedure call.

Priority 5 - Well Known port 22555 for IP Phone call setup.

(Default 00)

11.2.5.43 WELL_KNOWN_PORT [7:6]_PRIORITY- WELL KNOWN LOGIC PORT 7 AND 6 PRIORITYI²C Address h0EE, CPU Address 598Accessed by CPU, serial interface and I²C (R/W)

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

Priority 6 - well know port 22 for ssh.

Priority 7 – well Known port 554 for rtsp.

(Default 00)

11.2.5.44 WELL KNOWN_PORT_ENABLE [7:0] – Well Known Logic 7 to 0 Port EnablesI²C Address h0EF, CPU Address 599Accessed by CPU, 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)

11.2.5.45 RLOWL – USER DEFINE RANGE LOW BIT 7:0I²C Address h0F4, CPU Address: 59aAccessed by CPU, serial interface and I²C (R/W)

[7:0] Lower 8 bit of the User Define Logical Port Low Range (Default 00)

11.2.5.46 RLOWH – User Define Range Low Bit 15:8I²C Address h0F5, CPU Address: 59bAccessed by CPU, serial interface and I²C (R/W)

[7:0] Upper 8 bit of the User Define Logical Port Low Range (Default 00)

11.2.5.47 RHIGHL – User Define Range High Bit 7:0I²C Address h0D3, CPU Address: 59cAccessed by CPU, serial interface and I²C (R/W)

[7:0] Lower 8 bit of the User Define Logical Port High Range (Default 00)

11.2.5.48 RHIGHH – User Define Range High Bit 15:8I²C Address h0D4, CPU Address: 59dAccessed by CPU, serial interface and I²C (R/W)

[7:0] Upper 8 bit of the User Define Logical Port High Range (Default 00)

11.2.5.49 RRIORITY – User Define Range PriorityI²C Address h0D5, CPU Address: 59eAccessed by CPU, serial interface and I²C (R/W)

7	4	3	0
		Range Transmit Priority	Drop

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

11.2.6 (Group 6 Address) MISC Group

11.2.6.1 MII_OP0 – MII Register Option 0

I²C Address F0, CPU Address:h600

Accessed by CPU, serial interface and I²C (R/W)

7	6	5	4	0
hfc	1prst	DisJ	Vendor Spc. Reg Addr	

- 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 applications or any serial operation slower than 10 Mbps.
 0 = Enable
 1 = Disable
- Bit [4:0]: Vendor specified link status register address (null value means don't use it) (Default 00). This is used if the Linkup bit position in the PHY is non-standard.

11.2.6.2 MII_OP1 – MII Register Option 1

I²C Address F1, CPU Address:h601

Accessed by CPU, serial interface and I²C (R/W)

7	4	3	0
Speed bit location		Duplex bit location	

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

11.2.6.3 FEN – Feature Register

I²C Address F2, CPU Address:h602

Accessed by CPU, serial interface and I²C (R/W)

7	6	5	4	3	2	1	0
DML	Mii	Rp	IP Mul	V-Sp	DS	RC	SC

-
- Bits [0]: Statistic Counter Enable **(Default 0)**
- 0 – Disable
 - 1 – Enable (all ports)
- When statistic counter is enable, an interrupt control frame is generated to the CPU, every time a counter wraps around. This feature requires an external CPU.
- Bits [1]: Rate Control Enable **(Default 0)**
- 0 – Disable
 - 1 – Enable; Must also set ECR2Pn[3] = 1
- This bit enables/disables the rate control for all 10/100 ports. To start rate control in a 10/100 port the rate control memory must be programmed. This feature requires an external CPU. See Programming QoS Registers Application Note and Processor Interface Application Note for more information.
- Bit [2]: Support DS EF Code. **(Default 0)**
- 0 – Disable
 - 1 – Enable (all ports)
- When 101110 is detected in DS field (TOS[7:2]), the frame priority is set for 110 and drop is set for 0.
- Bit [3]: Enable VLAN spanning tree support **(Default 0)**
- 0 – Disable
 - 1 – Enable
- When VLAN spanning tree is enable the registers ECR1Pn are NOT used to program the port spanning tree status. The port status is programmed using the Control Command Frame.
- Bit [4]: Disable IP Multicast Support **(Default 1)**
- 0 – Enable IP Multicast Support
 - 1 – Disable IP Multicast Support
- When enable, IGMP packets are identified by search engine and are passed to the CPU for processing. IP multicast packets are forwarded to the IP multicast group members according to the VLAN port mapping table.
- Bit [5]: Enable report to CPU**(Default 0)**
- 0 – Disable report to CPU
 - 1 – Enable report to CPU
- When disable new VLAN port association report, new MAC address report or aging reports are disable for all ports. When enable, register SE_OPEMODE is used to enable/disable selectively each function.
- Bit [6]: Disable MII Management State Machine **(Default 0)**
- 0: Enable MII Management State Machine
 - 1: Disable MII Management State Machine
-

Bit [7]: Disable using MCT Link List structure (**Default 0**)

0 – Enable using MCT Link structure

1 - Disable using MCT Link List structure

11.2.6.4 MIIC0 – MII Command Register 0

CPU Address:h603

Accessed by CPU and 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.

11.2.6.5 MIIC1 – MII Command Register 1

CPU Address:h604

Accessed by CPU and 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.

11.2.6.6 MIIC2 – MII Command Register 2

CPU Address:h605

Accessed by CPU and serial interface only (R/W)



Bit [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.

11.2.6.7 MIIC3 – MII Command Register 3

CPU Address:h606

Accessed by CPU and 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. Writing this register will initiate a serial management cycle to the MII management interface.

11.2.6.8 MIID0 – MII Data Register 0

CPU Address:h607

Accessed by CPU and serial interface only (RO)

Bit [7:0] - MII Data [7:0]

11.2.6.9 MIID1 – MII Data Register 1

CPU Address:h608

Accessed by CPU and serial interface only (RO)

Bit [7:0] - MII Data [15:8]

11.2.6.10 LED Mode – LED Control

CPU Address:h609

Accessed by CPU, 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 100MHz SCLK

00 = 100MHz/8 = 12.5 MHz	01 = 100MHz/16 = 6.25 MHz
10 = 100MHz/32 = 3.125 MHz	11 = 100MHz/64 = 1.5625 MHz

For 125 MHz SCLK

00 = 125MHz/64 = 1953 KHz	01 = 125MHz/128 = 977 KHz
10 = 125MHz/512 = 244 KHz	11 = 125MHz/1024 = 122 KHz

Bit [7:5]: Reserved. Must be set to '0' (**Default 0**)

11.2.6.11 DEVICE Mode

CPU Address:h60a

Accessed by CPU and serial interface (R/W)

7	4	3	2	1	0
Device ID			LgFrm		

- Bit [1:0]: Reserved. Must be set to '0' (Default 0)
- Bit [2]: Support < = 1536 frames
0: < = 1518 bytes (< = 1522 bytes with VLAN tag) (Default)
1: < = 1536 bytes
- Bit [3]: Reserved. Must be set to '0' (**Default 0**)
- Bit [7:4]: DEVICE ID (**Default 0**). This is for stacking operation. This is the stack ID for loop topology.

11.2.6.12 CHECKSUM - EEPROM Checksum

I²C Address FF, CPU Address:h60b

Accessed by CPU, serial interface and I²C (R/W)

Bit [7:0]: (**Default 0**)

This register is used in unmanaged mode only. Before requesting that the MVTX2603 updates the EEPROM device, the correct checksum needs to be calculated and written into this checksum register. The checksum formula is:

$$\sum_{i=0}^{FF} i^2C \text{ register} = 0$$

When the MVTX2603 boots from the EEPROM the checksum is calculated and the value must be zero. If the checksum is not zeroed the MVTX2603 does not start and pin CHECKSUM_OK is set to zero.

11.2.7 (Group F Address) CPU Access Group

11.2.7.1 GCR-Global Control Register

CPU Address: hF00

Accessed by CPU and serial interface. (R/W)

7	5	4	3	2	1	0
		Reset	Bist	SR	SC	

- 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

11.2.7.2 DCR - Device Status and Signature Register

CPU Address: hF01

Accessed by CPU and 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 (not writing configuration to I²C)

Bit [1]: 1: Busy reading configuration from I²C
0: Not busy (not 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: MVTX2603 device

Bit [7:6]: Revision
00: Initial Silicon
01: XA1 Silicon
10: Production Silicon

11.2.7.3 DCR1 - Chip Status

CPU Address: hF02

Accessed by CPU and serial interface. (RO)

7	6	4	3	2	1	0
CIC		GIGA1	GIGA0			

- Bit [1:0]: Giga port 0 strap option
- 00 – 100 Mb MII mode
 - 01 – 2G mode
 - 10 – GMII
 - 11 – TBI
- Bit [3:2] Giga port 1 strap option
- 00 – 100 Mb MII mode
 - 01 – 2G mode
 - 10 – GMII
 - 11 – TBI
- Bit [7] Chip initialization completed

11.2.7.4 DPST – Device Port Status Register

CPU Address:hF03

Accessed by CPU and 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 and Negotiation status
 - 5'b00010 - Port 2 Operating mode and Negotiation status
 - 5'b00011 - Port 3 Operating mode and Negotiation status
 - 5'b00100 - Port 4 Operating mode and Negotiation status
 - 5'b00101 - Port 5 Operating mode and Negotiation status
 - 5'b00110 - Port 6 Operating mode and Negotiation status
 - 5'b00111 - Port 7 Operating mode and Negotiation status
 - 5'b01000 - Port 8 Operating mode and Negotiation status
 - 5'b01001 - Port 9 Operating mode and Negotiation status
 - 5'b01010 - Port 10 Operating mode and Negotiation status
 - 5'b01011 - Port 11 Operating mode and Negotiation status
 - 5'b01100 - Port 12 Operating mode and Negotiation status
 - 5'b01101 - Port 13 Operating mode and Negotiation status
 - 5'b01110 - Port 14 Operating mode and Negotiation status
 - 5'b01111 - Port 15 Operating mode and Negotiation status
 - 5'b10000 - Port 16 Operating mode and Negotiation status
 - 5'b10001 - Port 17 Operating mode and Negotiation status
 - 5'b10010 - Port 18 Operating mode and Negotiation status
 - 5'b00011 - Port 19 Operating mode and Negotiation status
 - 5'b10100 - Port 20 Operating mode and Negotiation status
 - 5'b10101 - Port 21 Operating mode and Negotiation status
 - 5'b10110 - Port 22 Operating mode and Negotiation status
 - 5'b10111 - Port 23 Operating mode and Negotiation status
 - 5'b11001 - Port 25 Operating mode/Neg status (Gigabit 1)
 - 5'b11010 - Port 26 Operating mode/Neg status (Gigabit 2)

11.2.7.5 DTST – Data read back register

CPU Address: hF04

Accessed by CPU and serial interface (RO)

This register provides various internal information as selected in DPST bit[4:0]. Refer to the PHY Control Application Note.

7	6	5	4	3	2	1	0
MD	Info	Sig	Giga	Inkdn	FE	Fdpx	FcEn

When bit is 1:

Bit [0] – Flow control enable

Bit [1] – Full duplex port

Bit [2] – Fast Ethernet port

Bit [3] – Link is down

Bits [7:4] for GE ports only:

Bit [4] – Giga port

Bit [5] – Signal detect (PCS mode only)

Bit [6] - 2G signal detect (2G mode only)

Bit [7] – Module detected (for hot swap purpose)

11.2.7.6 DA – Dead or Alive Register

CPU Address: hFFF

Accessed by CPU and serial interface (RO)

Always return 8'h **DA**. Indicate the CPU interface or serial port connection is good.

11.3 TBI Registers

Two sets of TBI registers are used for configure the two Gigabit ports if they are operating in TBI mode. These TBI registers are located inside the switching chip and they are accessed through the MII command and MII data registers.

11.3.1 Control Register

MII Address: h00

Read/Write

Bit [15]	Reset PCS logic and all TBI registers 1 = Reset. 0 = Normal operation.
Bit [14]	Reserved. Must be programmed with “0”.
Bit [13]	Speed selection (See bit 6 for complete details)

Bit [12]	Auto Negotiation Enable 1 = Enable auto-negotiation process. 0 = Disable auto-negotiation process (Default).
Bit [11:10]	Reserved. Must be programmed with "0"
Bit [9]	Restart Auto Negotiation. 1 = Restart auto-negotiation process. 0 = Normal operation (Default).
Bit [8:7]	Reserved.
Bit [6]	Speed Selection Bit[6][13] 1 1 = Reserved 1 0 =1000 Mb/s (Default) 0 1 =100 Mb/s 0 0 =10 Mb/s
Bit [5:0]	Reserved. Must be programmed with "0".

11.3.2 Status Register

MII Address: h01

Read Only

Bit [15:9]	Reserved. Always read back as "0".
Bit [8]	Reserved. Always read back as "1".
Bit [7:6]	Reserved. Always read back as "0".
Bit [5]	Auto-Negotiation Complete 1 = Auto-negotiation process completed. 0 = Auto-negotiation process not completed.
Bit [4]	Reserved. Always read back as "0"
Bit [3]	Reserved. Always read back as "1"
Bit [2]	Link Status 1 = Link is up. 0 = Link is down.
Bit [1]	Reserved. Always read back as "0".
Bit [0]	Reserved. Always read back as "1".

11.3.3 Advertisement Register

MII Address: h04

Read/Write

Bit [15]	Next Page 1 = Has next page capabilities. 0 = Do not has next page capabilities (Default).
Bit [14]	Reserved. Always read back as “0”. Read Only.
Bit [13:12]	Remote Fault. Default is “0”.
Bit [11:9]	Reserved. Always read back as “0”. Read Only.
Bit [8:7]	Pause. Default is “00”
Bit [6]	Half Duplex 1 = Support half duplex (Default). 0 = Do not support half duplex.
Bit [5]	Full duplex 1 = Support full duplex (Default). 0 = Do not support full duplex.
Bit [4:0]	Reserved. Always read back as “0”. Read Only.

11.3.4 Link Partner Ability Register

MII Address: h05

Read Only

Bit [15]	Next Page 1 = Has next page capabilities. 0 = Do not has next page capabilities.
Bit [14]	Acknowledge
Bit [13:12]	Remote Fault.
Bit [11:9]	Reserved. Always read back as “0”.
Bit [8:7]	Pause.
Bit [6]	Half Duplex 1 = Support half duplex. 0 = Do not support half duplex.
Bit [5]	Full duplex 1 = Support full duplex. 0 = Do not support full duplex.
Bit [4:0]	Reserved. Always read back as “0”.

11.3.5 Expansion Register

MII Address: h06

Read Only

Bit [15:2]	Reserved. Always read back as “0”.
Bit [1]	Page Received. 1 = A new page has been received. 0 = A new page has not been received.
Bit [0]	Reserved. Always read back as “0”.

11.3.6 Extended Status Register

MII Address: h15

Read Only

Bit [15]	1000 Full Duplex 1 = Support 1000 full duplex operation (Default). 0 = Do not support 1000 full duplex operation.
Bit [14]	1000 Half Duplex 1 = Support 1000 half duplex operation (Default). 0 = Do not support 1000 half duplex operation.
Bit [13:0]	Reserved. Always read back as “0”.

11.4 Characteristics and Timing

11.4.1 Absolute Maximum Ratings

Storage Temperature	-65°C to +150°C
Operating Temperature	-40°C to +85°C
Maximum Junction Temperature	+125°C
Supply Voltage V_{CC} with Respect to V_{SS}	+3.0V to +3.6V
Supply Voltage V_{DD} with Respect to V_{SS}	+2.38V to +2.75V
Voltage on Input Pins	+0.5V to ($V_{CC} + 3.3V$)

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.

11.4.2 DC Electrical Characteristics

$$V_{CC} = 3.3V \pm 10\%$$

$$T_{AMBIENT} = -40^{\circ}C \text{ to } +85^{\circ}C$$

$$V_{DD} = 2.5V +10\% / -5\%$$

11.4.3 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

11.5 AC Characteristics and Timing

11.5.1 Typical Reset & Bootstrap Timing Diagram

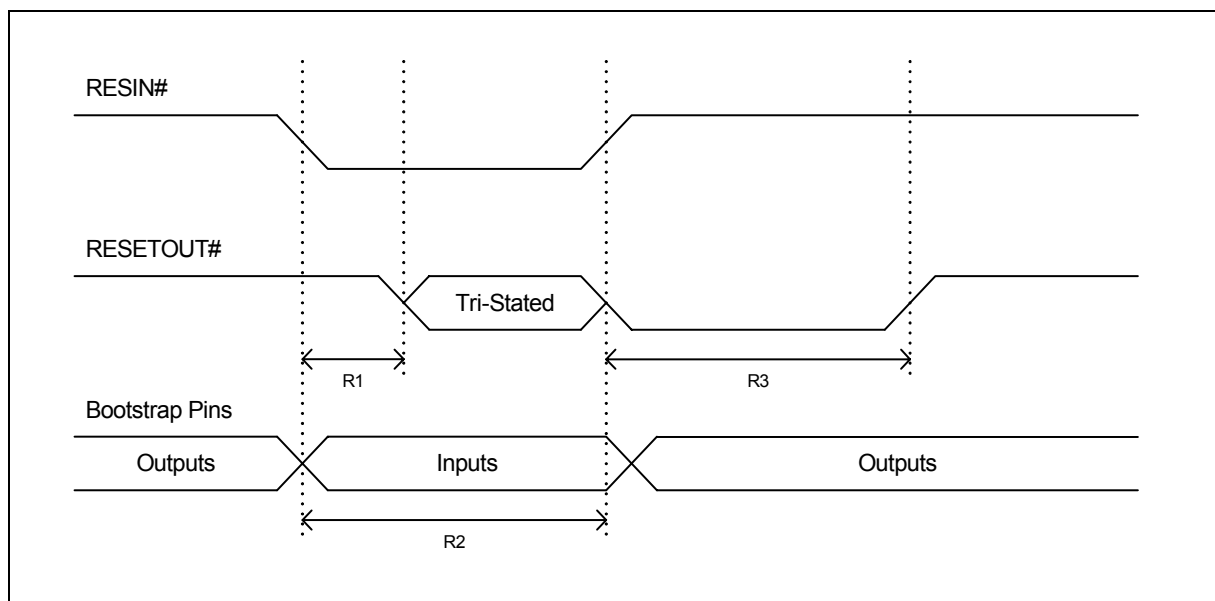


Figure 16 - Typical Reset & Bootstrap Timing Diagram

Symbol	Parameter	Min.	Typ.	Note:
R1	Delay until RESETOUT# is tri-stated		10 ns	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 14 - 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

11.5.2 Local Frame Buffer SBRAM Memory Interface

11.5.2.1 Local SBRAM Memory Interface A




Figure 17 - Local Memory Interface – Input Setup and Hold Timing

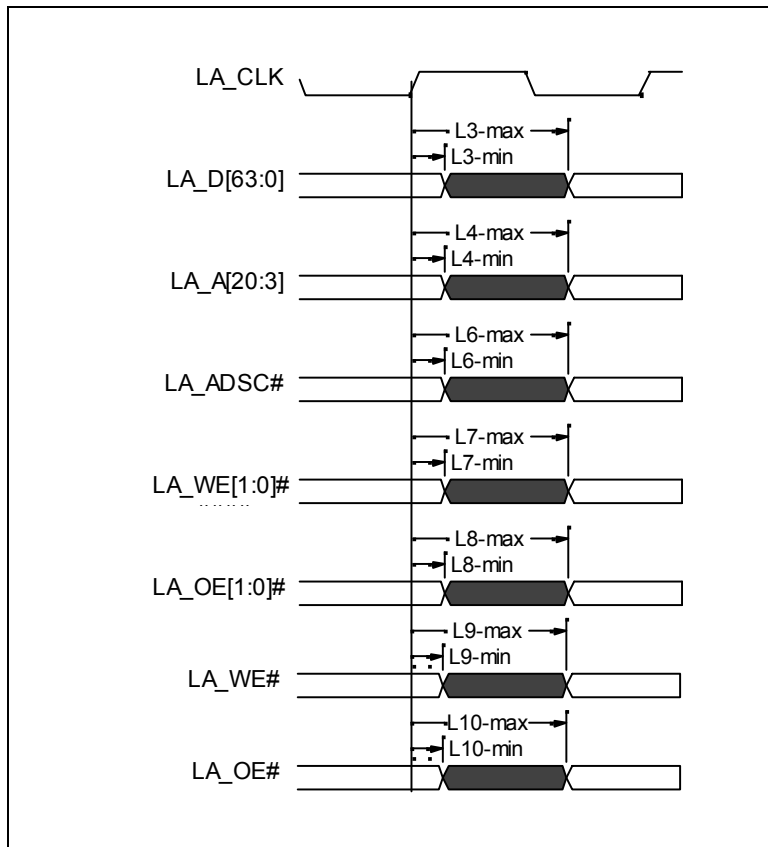


Figure 18 - Local Memory Interface - Output Valid Delay Timing

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

Table 15 - AC Characteristics – Local Frame Buffer SBRAM Memory Interface

11.5.2.2 Local SBRAM Memory Interface B




Figure 19 - Local Memory Interface – Input Setup and Hold Timing

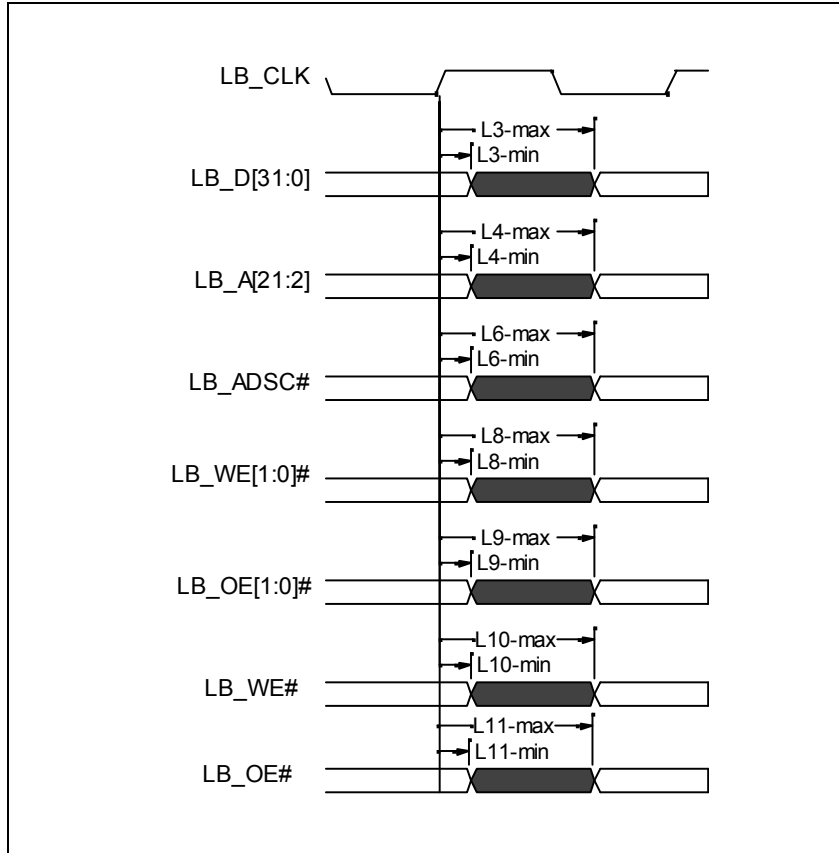


Figure 20 - Local Memory Interface - Output Valid Delay Timing

Symbol	Parameter	-100 MHz		Note:
		Min. (ns)	Max. (ns)	
L1	LB_D[63:0] input set-up time	4		
L2	LB_D[63:0] input hold time	1.5		
L3	LB_D[63:0] output valid delay	1.5	7	C _L = 25 pf
L4	LB_A[20:3] output valid delay	2	7	C _L = 30 pf
L6	LB_ADSC# output valid delay	1	7	C _L = 30 pf
L8	LB_WE[1:0]#output valid delay	1	7	C _L = 25 pf
L9	LB_OE[1:0]# output valid delay	-1	1	C _L = 25 pf
L10	LB_WE# output valid delay	1	7	C _L = 25 pf
L11	LB_OE# output valid delay	1	5	C _L = 25 pf

Table 16 - AC Characteristics – Local Switch Database SBRAM Memory Interface

11.5.3 Reduced Media Independent Interface

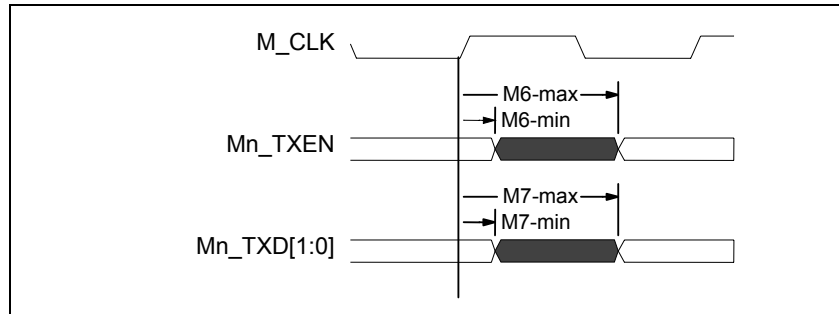


Figure 21 - AC Characteristics – Reduced Media Independent Interface

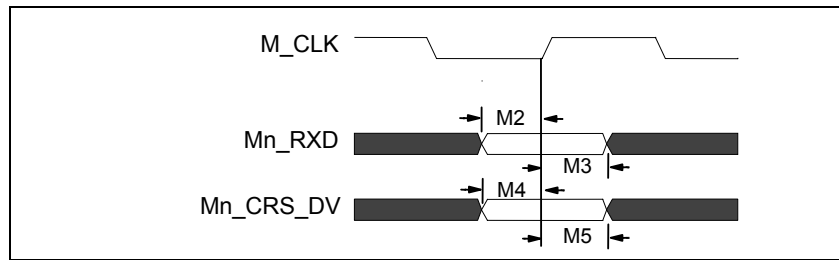


Figure 22 - AC Characteristics – Reduced Media Independent Interface

Symbol	Parameter	M_CLK=50 MHz		Note
		Min. (ns)	Max. (ns)	
M2	Mn_RXD[1:0] Input Setup Time	4		
M3	Mn_RXD[1:0] Input Hold Time	1		
M4	Mn_CRS_DV Input Setup Time	4		
M5	Mn_CRS_DV Input Hold Time	1		
M6	Mn_TXEN Output Delay Time	2	11	C _L = 20 pF
M7	Mn_TXD[1:0] Output Delay Time	2	11	C _L = 20 pF

Table 17 - AC Characteristics – Reduced Media Independent Interface

11.5.4 Gigabit Media Independent Interface




Figure 23 - AC Characteristics- Gigabit Media Independent Interface




Figure 24 - AC Characteristics – Gigabit Media Independent Interface

Symbol	Parameter	-125 Mhz		Note
		Min. (ns)	Max. (ns)	
G1	Gn_RXD[7:0] Input Setup Times	2		
G2	Gn_RXD[7:0] Input Hold Times	0.5		
G3	Gn_RXDV Input Setup Times	1.2		
G4	Gn_RXDV Input Hold Times	0.5		
G5	Gn_TXD[7:0] Output Delay Times	1	6	$C_L = 20$ pf
G6	Gn_TXEN Output Delay Times	1	6.5	$C_L = 20$ pf

Table 18 - AC Characteristics – Gigabit Media Independent Interface

11.5.5 Ten Bit Interface




Figure 25 - AC Characteristics – Ten Bit Interface (RX)




Figure 26 - AC Characteristics –Ten Bit Interface (TX)

Symbol	Parameter	(TXCLK=125 MHz RXCLK0/1=62.5 MHz)		Note
		Min. (ns)	Max .(ns)	
T1	Gn_RXD[9:0] Input Setup Times in reference to G_RXCLK0	2		
T2	Gn_RXD[9:0] Input Hold Times in reference to G_RXCLK0	1		
T3	Gn_RXD[9:0] Input Setup Times in reference to G_RXCLK1	2		
T4	Gn_RXD[9:0] Input Hold Times in reference to G_RXCLK1	1		
T5	Gn_TXD[9:0] Output Delay Times	1	6	C _L = 20 pf

Table 19 - AC Characteristics – Ten Bit Interface

11.5.6 LED Interface




Figure 27 - AC Characteristics – LED Interface

Symbol	Parameter	Variable FREQ.		Note
		Min. (ns)	Max. (ns)	
LE5	LED_SYN Output Valid Delay	-1	7	$C_L = 30 \text{ pf}$
LE6	LED_BIT Output Valid Delay	-1	7	$C_L = 30 \text{ pf}$

Table 20 - AC Characteristics – LED Interface

11.5.7 SCANLINK, SCANCOL Interface

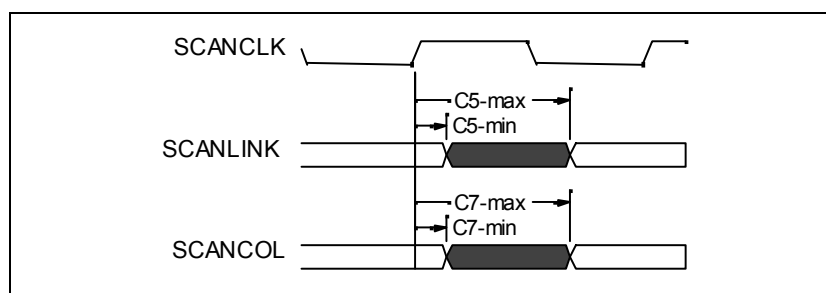


Figure 28 - SCANLINK, SCANCOL Output Delay Timing

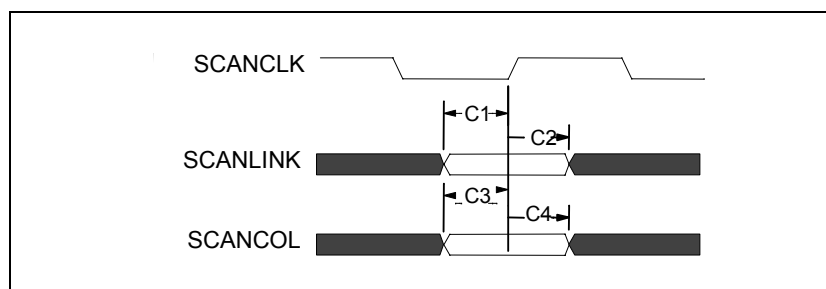


Figure 29 - SCANLINK, SCANCOL Setup Timing

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	$C_L = 30\text{pf}$
C7	SCANCOL output valid delay	0	10	$C_L = 30\text{pf}$

Table 21 - SCANLINK, SCANCOL Timing

11.6 MDIO Interface




Figure 30 - MDIO Input Setup and Hold Timing




Figure 31 - 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 22 - MDIO Timing

11.6.1 I²C Interface




Figure 32 - I²C Input Setup Timing




Figure 33 - 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 23 - I²C Timing

11.6.2 Synchronous Serial Interface




Figure 34 - Serial Interface Setup Timing







Figure 35 - 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 24 - 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
