

Features

November 2003

- 8 10/100 Mbps auto-negotiating RMII ports
- 1 10/100 Mbps auto-negotiating MII/serial port (port 8) that can be used as a WAN uplink or as a 9th port
- Operates stand-alone or can be cascaded with a second MDS108 to reach 16 ports
 - XLink expansion MII port (port 8)
 - Operates at 100/200/300/400 Mbps
- External I²C EEPROM for power-up configuration
 - Default mode allows operation without external EEPROM
- Up to 8 port-based VLANs
- Full wirespeed layer 2 switching on all ports (up to 2.679 M packets per second)
- Internal 1 K MAC address table
 - Auto address learning
 - Auto address aging
- Leading-edge Quality of Service (QoS) capabilities provided based on 802.1 p and IP TOS/DS field
 - 2 queues per output port
 - Packet scheduling based on Weighted Round Robin (WRR) and Weighted Random Early Detection/Drop (WRED)
 - With flow control disabled, can drop packets during congestion using WRED
 - 2 levels of packet drop provided
- Provides port-based prioritization of packets on up to 4 ports

Ordering Information

MDS108AL 208 Pin PQFP

-40°C to +85°C

- Input ports are defined to be high or low priority
- Allows explicit identification of IP phone ports
- Supports both full and half duplex ports
- Ports 0 & 1 can be trunked to provide a 200 Mbps link to another switch or server
- Port 7 can be used to mirror traffic from the other 7 ports (0-6)
- Utilizes a single low-cost external pipelined, SyncBurst SRAM (SBRAM) for buffer memory
 - 256 KB or 512 KB (1 chip)
- Flow control capabilities
 - Provides back-pressure for half duplex
 - 802.3x flow control for full duplex
- Supports external parallel port for configuration updates
- Special power-saving mode for inactive ports
- Ability to support WinSock 2.0 and Windows2000 smart applications
- Transmit delay control capabilities
 - Assures maximum delay (< 1 ms)
 - Supports mixed voice/data networks
- Optimized pin-out for easy board layout

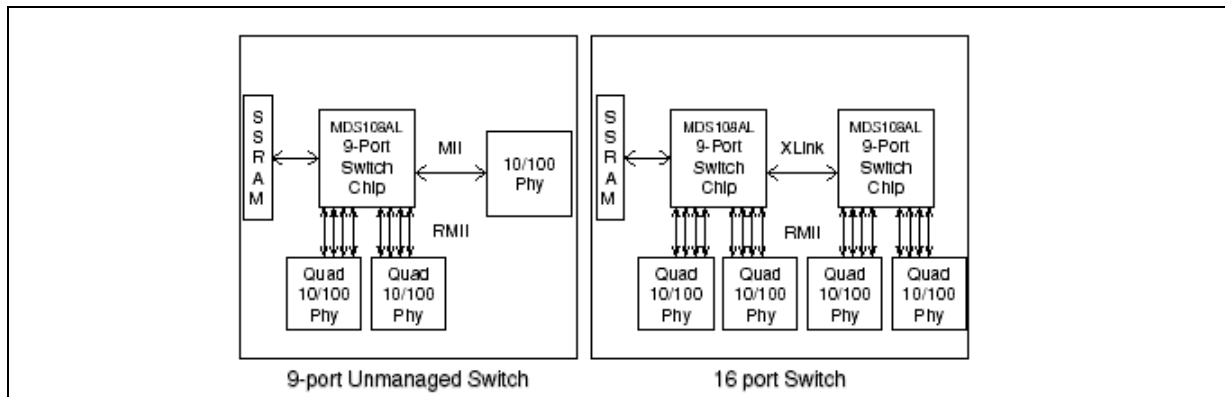


Figure 1 - System Block Diagram

Description

The MDS108 is a fully integrated 9-port Ethernet switch designed to support the low-cost requirements of unmanaged switch applications. The MDS108 provides features that are normally not associated with plug-and-play technology, while not requiring an external processor to facilitate their utilization.

The MDS108 begins operating immediately at power-up, learning addresses automatically, and forwarding packets at full wire-speed to any of its eight output ports or the XLink expansion port. The default configuration allows operation without using an external EEPROM.

With an EEPROM to configure the device at power-up, the MDS108 provides flexible features: port trunking, port mirroring, port-based VLANs and (QoS) capabilities that are usually associated only with managed switches.

The built-in intelligence capabilities of the MDS108 allows it to recognize and offer packet prioritization using Zarlink Semiconductor's QoS. Packets are prioritized based upon their layer 2 VLAN priority tag or the layer 3 Type-Of-Service/ Differentiated Services (TOS/DS) field. This priority can be defined as transmit and/or drop priority.


The MDS108 can be used to create an 8-port unmanaged switch with one WAN router port by connecting a CPU (ARM or MPC 850) to the additional MII port (port 8). The only external components needed are the physical layer transceivers and a single SDRAM, resulting in a low, total system cost.

Designed to support the requirements of converging networks, the MDS108 utilizes a power conserving architecture. To further enhance this power management, the chip automatically detects when a switch port is not being utilized, and turns off the logic associated with that port, thereby saving power and reducing the current load on the switch power supply.

Operating at 66 MHz internally, and with a 66 MHz interface to the external SDRAM, the MDS108 sustains full wire-speed switching on all 9 ports.

The chip is packaged in a small 208 pin Plastic Quad Flat-Pak (PQFP) package.

MDS108 Physical Pinout



Pin Reference Table

Pin #	Pin Name
1	L_A[7]
2	L_A[8]
3	VDD (CORE)
4	L_A[9]
5	L_A[10]
6	L_A[11]
7	VSS
8	L_A[12]
9	L_A[13]
10	L_A[14]
11	VDD
12	NC
13	NC
14	NC
15	NC
16	NC
17	NC
18	VSS (CORE)
19	M0_TXEN
20	M0_TXD[0]
21	M0_TXD[1]
22	M0_CRSDV
23	M0_RXD[0]
24	M0_RXD[1]
25	VDD
26	M1_TXEN
27	M1_TXD[0]
28	M1_TXD[1]
29	M1_CRSDV
30	M1_RXD[0]
31	M1_RXD[1]
32	VSS
33	M2_TXEN
34	M2_TXD[0]

35	M2_TXD[1]
36	M2_CRSDV
37	M2_RXD[0]
38	M2_RXD[1]
39	VDD (CORE)
40	M3_TXEN
41	M3_TXD[0]
42	M3_TXD[1]
43	M3_CRSDV
44	M3_RXD[0]
45	M3_RXD[1]
46	VSS (CORE)
47	NC
48	NC
49	NC
50	NC
51	NC
52	NC
53	NC
54	NC
55	NC
56	NC
57	NC
58	NC
59	VDD
60	M4_TXEN
61	M4_TXD[0]
62	M4_TXD[1]
63	M4_CRSDV
64	M4_RXD[0]
65	M4_RXD[1]
66	VSS
67	M5_TXEN
68	M5_TXD[0]
69	M5_TXD[1]
70	M5_CRSDV

71	M5_RXD[0]
72	M5_RXD[1]
73	VDD (CORE)
74	M6_TXEN
75	M6_TXD[0]
76	M6_TXD[1]
77	M6_CRSDV
78	M6_RXD[0]
79	M6_RXD[1]
80	VSS (CORE)
81	M7_TXEN
82	M7_TXD[0]
83	M7_TXD[1]
84	M7_CRSDV
85	M7_RXD[0]
86	M7_RXD[1]
87	VDD
88	NC
89	NC
90	NC
91	NC
92	NC
93	NC
94	VSS
95	M_CLK
96	VDD (CORE)
97	M8_RXDV/S8_CRSDV
98	M8_COL/S8_COL
99	VSS
100	M8_RXCLK/S8_RXCLK
101	VDD
102	M8_RXD[0]/S8_RXD
103	M8_RXD[1]
104	M8_RXD[2]
105	M8_RXD[3]
106	VSS (CORE)

107	M8_TXCLK/S8_TXCLK	142	RSTIN#	178	L_D[13]
108	VDD	143	MIRROR_CONTROL[0]	179	L_D[14]
109	M8_TXEN[0]/S8_TXEN	144	MIRROR_CONTROL[1]	180	VSS
110	M8_TXD[0]/S8_TXD	145	MIRROR_CONTROL[2]	181	L_D[15]
111	M8_TXD[1]	146	MIRROR_CONTROL[3]	182	L_D[16]
112	M8_TXD[2]	147	VDD	183	L_D[17]
113	M8_TXD[3]	148	SCLK	184	VDD (CORE)
114	M8_LINK/S8_LINK	149	VSS	185	L_D[18]
115	M8_DUPLEX/S8_DUPLEX	150	L_A[2]	186	L_D[19]
116	M8_SPEED	151	L_A[17]	187	L_D[20]
117	VSS	152	VDD	188	L_D[21]
118	M8_REFCLK	153	L_CLK	189	VSS
119	VDD	154	VSS	190	L_D[22]
120	M_MDC	155	L_WE#	191	L_D[23]
121	VSS	156	L_OE#	192	L_D[24]
122	M_MDIO	157	L_ADSC#	193	L_D[25]
123	SCL	158	L_A[16]	194	VDD
124	SDA	159	VDD (CORE)	195	L_D[26]
125	TEST#	160	L_A[15]	196	L_D[27]
126	TRUNK_ENABLE	161	L_D[0]	197	L_D[28]
127	STROBE	162	VSS	198	VSS (CORE)
128	DATA0	163	L_D[1]	199	L_D[29]
129	ACK	164	L_D[2]	200	L_D[30]
130	VDD (CORE)	165	L_D[3]	201	L_D[31]
131	TSTOUT[0]	166	VDD	202	VDD
132	TSTOUT[1]	167	L_D[4]	203	L_A[18]
133	TSTOUT[2]	168	L_D[5]	204	L_A[3]
134	TSTOUT[3]	169	L_D[6]	205	L_A[4]
135	TSTOUT[4]	170	L_D[7]	206	L_A[5]
136	TSTOUT[5]	171	VSS (CORE)	207	VSS
137	TSTOUT[6]	172	L_D[8]	208	L_A[6]
138	TSTOUT[7]	173	L_D[9]		
139	T_MODE	174	L_D[10]		
140	VSS (CORE)	175	VDD		
141	RSTOUT#	176	L_D[11]		
		177	L_D[12]		

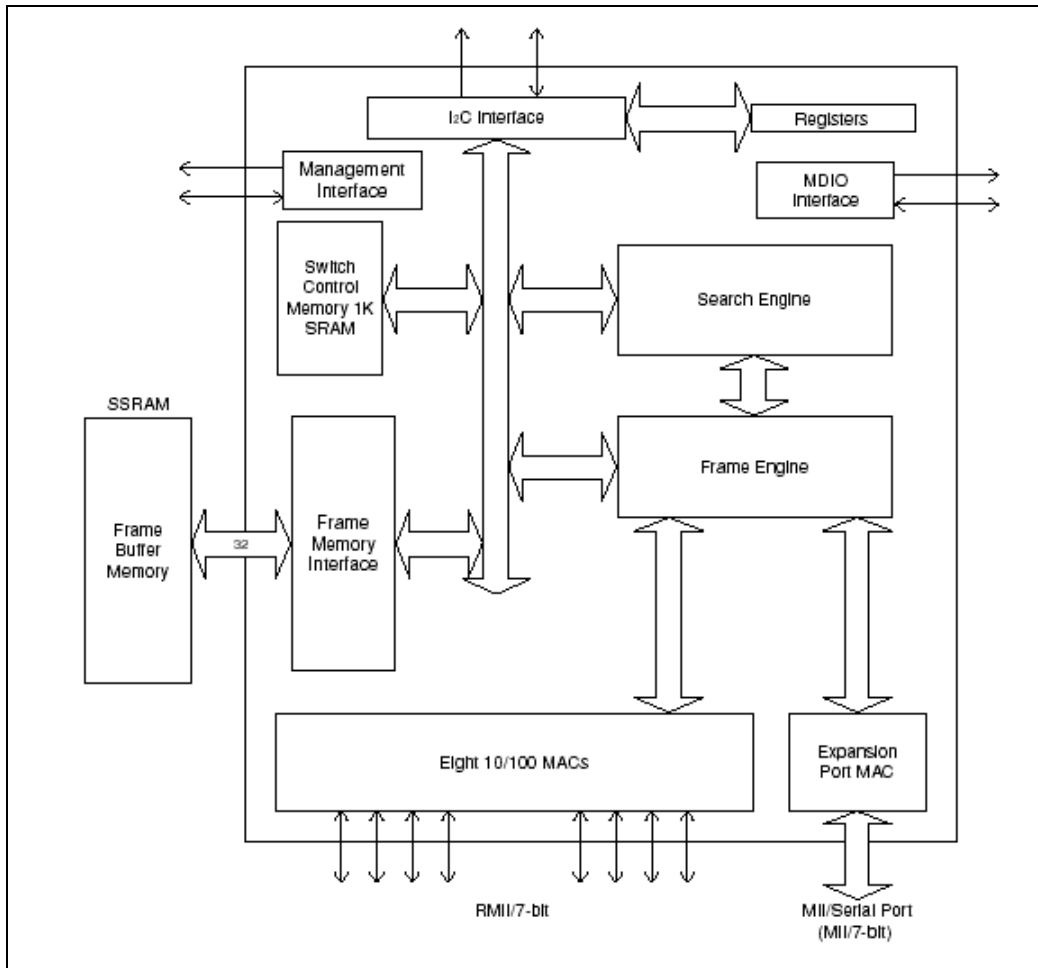


Figure 2 - MDS108 Block Diagram

1.0 Functional Operation

The MDS108 is designed to provide a cost-effective layer 2 switching solution, using technology from the Zarlink family to offer a highly integrated product for the unmanaged, Differentiated Services (DS) ready, Ethernet switch market.

Nine 10/100 Media Access Controllers (MAC) provide the protocol interface. These MACs perform the required packet checks to ensure that each packet that is provided to the Frame Engine has met all the IEEE 802.1 standards. Each MAC supports half duplex “back pressure,” and full duplex 802.3x “PAUSE” Flow Control.

The PHY addresses for the 8 RMII MACs are from 08h to 0Fh. These 8 ports are denoted as ports 0 to 7. The PHY address for the uplink MAC is 10h. This port is denoted as port 8.

Data packets longer than 1518 (1522 with VLAN tag) bytes and shorter than 64 bytes are dropped, and the MDS108 is designed to support minimum interframe gaps between incoming packets.

The Frame Engine (FE) is the primary packet buffering and forwarding engine within the MDS108. As such, the FE controls the storage of packets into and out of the external frame buffer memory, keeps track of frame buffer availability and schedules packet transmissions. While packet data is being buffered, the FE extracts the necessary information from each packet header and sends it to the Search Engine for processing. Search results returned to the FE initiate the scheduling of packet transmission. When a packet is chosen for transmission, the FE reads the packet from external buffer memory and places it in the output FIFO of the output port.

2.0 Address Learning and Aging

The MDS108 is able to begin address learning and packet forwarding shortly after power-up is completed. The Search Engine examines the contents of its internal Switch Database Memory for each valid packet that is received on a MDS108 input port.

Unknown source and destination MAC addresses are detected when the Search Engine does not find a match within its database. These unknown source MAC addresses are learned by creating a new entry in the switch database memory, and storing the necessary resulting information in that location. Subsequent searches to a learned destination MAC address will return the new contents of that MAC Control Table (MCT) entry.

After each source address search the MCT entry aging flag is updated. MCT entries that have not been accessed during a user configurable time period (1 to 67,108 seconds) will be removed.

This aging time period can be configured using the 16-bit value stored in the register’s MAC Address Aging Timer Low and High (AGETIME_LOW, AGETIME_HIGH). The aging period is defined as the bit concatenation of AGETIME_HIGH with AGETIME_LOW, multiplied by 1024 ms. For example, if AGETIME_LOW = 25, and AGETIME_HIGH = 01 (in hexadecimal), then the concatenated value 125 is equal to decimal 293. Multiplying 293 by 1024 ms, we determine that the corresponding aging time is 300 ms. In fact, 300 ms is the default aging time for the MDS108.

The aging of all MCT entries is checked once during each time period. If the MCT entry has not been exercised before the end of the next time period, it will be deleted.

3.0 Quality of Service

The MDS108 applies Zarlink Semiconductor’s architecture to provide new QoS capabilities for unmanaged switch applications. Similar to the QoS capabilities of the other Zarlink chipset members, the MDS108 offers two transmit queues per output port.

The FE manages the output transmission queues for all MDS108 ports. Once the destination address search is complete, and the switch decision is sent back to the FE, the packet is inserted into the appropriate output queue. Whether the packet is inserted into a high or low priority queue is determined by either the VLAN tag information or the Type of Service/Differentiated Services (TOS/DS) field in the IP header. Either of these priority fields can be

used to select the transmission priority (as per the USE_TOS bit in the FCR register). The mapping of the priority field values into either the high or low priority queue can be configured using the MDS108 configuration registers AVPM and TOSPM.

If the system uses the TOS/DS field to prioritize packets, there are two choices regarding which bits of the TOS/DS field are used. Bits [0:2] of the TOS byte (known as the IP precedence field) or bits [3:5] of the TOS byte (known as the Delay/Throughput Reliability or (DTR) field) can be used to resolve the transmission queue priority. Either bit group, [0:2] or [3:5], can also be used to resolve packet drop precedence, as per bits 6 and 7 of the register FCBST.

The MDS108 utilizes Weighted Round Robin (WRR) to schedule packets for transmission. To enable MDS108's intelligent QoS scheduling capabilities, the use of an external EEPROM to change the default register configurations is required.

Weighted Round Robin is an efficient method to ensure that each of the transmission queues receives at least a minimum service level. With two output transmission queues, the MDS108 will transmit X packets from the high priority queue before transmitting a single packet from the low priority queue. The MDS108 allows the designer to set the high priority weight X to a value between 1 and 15. If both queues contain packets, and the high priority weight is set to the value 4, then the MDS108 will transmit 4 high priority packets before transmitting each low priority packet.

The MDS108 also employs a proprietary mechanism to ensure the timely delivery of high priority packets. When the latency of high priority packets reaches a threshold, the MDS108 will override the WRR weights and transmit only high priority packets until the high priority packet delays are below the threshold. This threshold limit is 1 ms (last-in-first-out). The MDS108's proprietary scheduling algorithm is also designed to push low priority traffic through the device faster, if necessary, to unplug congested queues.

Loading the appropriate values into the configuration registers enables the QoS scheduling capabilities of the MDS108. QoS for packet transmission is enabled by performing the following four steps:

1. Select the TOS/DS or VLAN Priority Tag field as the decision-maker for IP packet scheduling. The selection is made using bit 7 of the Flooding Control Register (FCR).
 - FCR[7] = 0: Use VLAN Priority Tag field to determine the transmission priority, if this Tag field exists
 - FCR[7] = 1: Use TOS/DS field for IP packet priority resolution
2. Select which TOS/DS subfield to use as the decision-maker for packet transmission priority if the TOS/DS field was selected in step 1. The selection is made using bit 6 of the FCB Buffer Low Threshold Register (FCBST).
 - FCBST[6] = 0: Use DTR subfield to resolve the transmission priority
 - FCBST[6] = 1: Use IP precedence subfield¹ to resolve the transmission priority
3. Enable QoS using bit 5 of the Transmission Scheduling Control Register (AXSC). Set the transmission queue weight for the high priority queue using bits 0 to 3.
4. Create the mapping from the value in the TOS/DS or VLAN Priority Tag field to the corresponding high or low priority output queue. The mapping is created using the VLAN Priority Map (AVPM) and TOS Priority Map (TOSPM) registers.

Note that for half duplex operation, the priority queues² must be enabled using bit 7 in the Transmission Scheduling Control (AXSC) register to use QoS scheduling.

When QoS and flow control are enabled, the MDS108 will utilize enhanced WRR to schedule packet transmission, and will use either back pressure or 802.3X flow control to handle buffer congestion. When QoS is enabled and flow control is disabled, the MDS108 will utilize enhanced WRR to schedule packet transmission, and will use Weighted Random Early Detection/Drop (WRED) to drop random packets in order to handle buffer congestion. Because of

1. IP precedence and DTR subfields are referred to as TOS/DS[0:2] and TOS/DS[3:5] in the IP TOS/DS byte.

2. In Half Duplex mode, QoS scheduling functions are disabled by default.

WRED, only a few packet flows are slowed down while the remaining see no impact from the network traffic congestion.

WRED is a method of handling traffic congestion in the absence of flow control mechanisms¹. When flow control is enabled, all devices that are connected to a switch node that is exercising flow control are effectively unable to transmit, including nodes that are not directly responsible for the congestion problem. This inability to transmit during flow control periods would wreak havoc with voice packets, or other high priority packet flows, and therefore flow control is not recommended for networks that mix voice and data traffic.

WRED allows traffic to continue flowing into ports on a switch, and randomly drops packets with different probabilities based upon each packet's priority markings. As the switch congestion increases, the probability of dropping an incoming packet increases, and as congestion decreases, the probability of dropping an incoming packet decreases. Not surprisingly, packets designated high-drop are sacrificed with higher odds during congestion than packets designated low-drop.

The following table summarizes the WRED operation of the MDS108. It lists the buffer thresholds at which each drop probability takes effect.

	WRED Threshold		Drop Percentage	
	Condition for High Priority Queue	Condition for Low Priority Queue	Drop Percentage for High-Drop Packet	Drop Percentage for Low-Drop Packet
Level 0	Total buffer space available in device is \leq LPBT		50%	0%
Level 1	24 buffers occupied	72 buffers occupied	75%	25%
Level 2		84 buffers occupied	100%	50%

Table 1 - WRED Operation of the MDS108

The WRED packet drop capabilities of the MDS108 are enabled by performing the following four steps:

- Select the TOS/DS or VLAN Tag field as the decision-maker for dropping packets. The selection is made using bit 7 of the Flooding Control Register (FCR).
 - FCR[7] = 0: Use VLAN Priority Tag field to resolve the drop level, if this field exists.
 - FCR[7] = 1: Use TOS/DS field for IP packet drop level resolution.
- Select which TOS/DS Tag subfield to use for dropping packets, provided that the TOS/DS field was selected in step 1. The selection is made using bit 7 of the FCB Buffer Low Threshold Register (FCBST).
 - FCBST[7] = 0: Use DTR subfield to resolve the drop precedence.
 - FCBST[7] = 1: Use IP precedence subfield to resolve the drop precedence.
- Create the mapping from the values in the TOS/DS or VLAN Tag field to the packet flags representing high or low-drop precedence. The mapping is created using the VLAN Discard Map (AVDM) and TOS Discard Map (TOSDM) registers.
- Make sure that the desired ports are flow control disabled, using the ECR1Px registers.

Note that to apply the WRED QoS function of the MDS108, flow control must be disabled.

1. Flow control, of course, provides the advantage of not dropping packets. However, its primary disadvantage is that a flow controlled port may experience head-of-line blocking. This means that if even 1 packet is destined to a congested output port, then all other packets originating from the same source may, in the worst case, be delayed – even if these other packets have uncongested destinations. On the other hand, WRED may cause some packet loss, but with no such head-of-line blocking problem. Which method of handling traffic congestion should be chosen will depend on the application.

3.1 A Few Examples

- **No QoS Scheduling Desired At All.** The default setting for the MDS108 is no QoS scheduling at all. Packets are transmitted using a simple first-in-first-out (FIFO) approach, without the reordering that would result from prioritization. All destinations use 1 queue only. QoS scheduling can be disabled for the entire chip using AXSC[5].
- **No QoS Scheduling Desired for Half-Duplex Ports.** It is possible to disable QoS for half duplex ports in the MDS108. Indeed, this is the default setting, because it is difficult to assure quality of service for half-duplex ports, which tend to experience unpredictable delay. All destinations configured as half-duplex use 1 transmission queue only in this setting. QoS scheduling can be disabled for all half-duplex ports using AXSC[7].
- **QoS Scheduling for Some Destinations, But Not Others.** The MDS108 does not support this feature. The three options offered are: (a) all ports are QoS-enabled, (b) no ports are QoS-enabled, and (c) only full-duplex ports are QoS enabled. Of course, the MDS108 will still exhibit single-queue scheduling at a port if all packets destined for it are marked with a single transmission priority.
- **No Flow Control Desired.** It is possible to disable flow control for the entire chip, regardless of the individual port settings. During congestion, some packets will be lost. This capability is located in AXSC[6].
- **Flow Control Desired for Some Sources, But Not Others.** By configuring each port separately using bit 0 in the ECR1Px registers, one may enable flow control for some ports, but not others. Flow control cannot be globally disabled in AXSC[6] if this function is to be achieved. At a congested destination, an incoming packet from a flow control enabled source will trigger a flow control message sent back to that source. On the other hand, an incoming packet from a flow control disabled source may or may not be dropped, as per WRED, but will never trigger flow control.
- **Scheduling vs. Dropping.** Using the configuration registers in the MDS108, as in earlier examples, a port may or may not have flow control enabled, and a port may or may not have QoS scheduling enabled. All four combinations are permissible parameter settings, and which one is chosen depends on application.

In one common application, suppose voice, critical data and web traffic packets originate from the same set of input ports, and are destined for the same set of output ports. The MDS108's enhanced WRR scheduling can be used to ensure a delay bound for the voice packets. Furthermore, because we want to guarantee that web traffic congestion does not block critical data or voice, we must disable flow control and use WRED to intelligently drop packets.

On the other hand, if the goal were file transfer without any packet dropping, one would enable the MDS108's flow control function, which halts incoming traffic when the system is congested. QoS scheduling can be disabled, both because flow control may make quality of service unpredictable, and because, in any case, delay is not critical in this application.

3.2 Port-Based Prioritization

Some applications may require an explicit prioritization of packets based upon the port the packet originates from. Defining specific ports of a switch to be IP Phone ports is a specific example that makes use of MDS108's ability to assign default priorities to ports.

The MDS108 can be configured to provide specific priority definitions on up to four ports (ports 0 – 3). These user defined port priorities override the packet priority markings (VLAN tag or TOS/DS), and the new priority is applied to all packets that enter the switch from that port. These port priority definitions are configured in the Port Priority (PTPRI) register. There are two bits for each of the four ports that can support port-based priorities. The EN bit allows the designer to turn on port priorities for each port, and the P bit allows the designer to select either high (1) or low (0) transmission priority for all packets that enter the switch through that port.

When port priorities are enabled, the remaining ports will provide QoS based upon the VLAN Tag or TOS DS field mappings in the configuration registers. Only those ports that have port priorities enabled will override the priority mappings.

4.0 Buffer Management

The MDS108 stores each input packet into the external frame buffer memory while determining the destination the packet is to be forwarded to. The total number of packets that can be stored in the frame buffer memory depends upon the size of the external SDRAM that is utilized. For a 256 KB SDRAM, the MDS108 can buffer 170 packets. For a 512 KB SDRAM, the MDS108 can buffer 340 packets.

In order to provide good quality of service characteristics, the MDS108 must carefully allocate the available buffer space. Such careful allocation can be accomplished using the external EEPROM to load the appropriate values into the MDS108 configuration registers. The Low-Drop Precedence Buffer Threshold (LPBT) register assures that traffic designated as low-drop actually receives reserved buffer space. The designer can set the minimum number of buffers reserved for low-drop unicast traffic, by setting this register with a value between 0 and 255. Unreserved buffers are treated as shared, and are accessible to all types of incoming traffic.

To set the maximum number of buffers permitted for all multicast packets, use the Multicast Buffer Control Register (MBCR). Unlike the LPBT register, the MBCR register does not define a reserved area of buffer memory, but instead provides a bound on the number of multicast packets that can be buffered at any one time.

During operation, the MDS108 will continuously monitor the amount of frame buffer memory that is available, and when the unused buffer space falls below a designer configurable threshold, the MDS108 will initiate flow control if enabled or WRED if not. This threshold is set using the FCB Buffer Low Threshold (FCBST) register.

5.0 Virtual LANs

The MDS108 provides the designer the ability to define a single port-based Virtual LAN (VLAN) for each of the nine ports. This VLAN is individually defined for each port using the Port Control Registers (ECR1Px[6:4]). Bits [6:4] allow the designer to define a VLAN ID (value between 0 – 7) for each port.

When packets arrive at an input of the MDS108, the search engine will determine the VLAN ID for that port, and then determine which of the other ports are also members of that VLAN by matching their assigned VLAN ID values. The packet will then be transmitted to each port with the same VLAN ID as the source port.

6.0 Port Trunking

Port trunking allows the designer to configure the MDS108, such that ports 0 and 1 are defined as a single logical port. This provides a 200 Mbps link to a switch or server utilizing two 100 Mbps ports in parallel.

Ports 0 and 1 can be trunked by pulling the TRUNK_EN pin to the high state. In this mode, the source MAC addresses of all packets received from the trunk are checked against the MCT database to ensure that they have a port ID of 0 or 1. Packets that have a port ID other than 0 and 1 will cause the MDS108 to learn the new MAC address for this port change.

On transmission, the trunk port is determined by hashing the source and destination MAC addresses. This provides a mapping between each MAC address and an associated trunk port. Subsequent packets with the same MAC address will always utilize the same trunk port.

The MDS108 also provides a safe fail-over mode for port trunking. If one of the two ports goes down, as identified by the port's link status signal, then the MDS108 will switch all traffic over to the remaining port in the trunk. Thus, the trunk link is maintained, albeit at a lower effective bandwidth.

7.0 Port Mirroring

Utilizing the 4 port mirroring control pins provides the ability to enable or disable port mirroring, select which of the remaining 7 ports is to be mirrored, and choose whether the receive or transmit data is being mirrored. The control for this function is shown in the following table.

When enabled, port mirroring will allow the user to monitor traffic going through the switch on output Port 7. If the port mirroring control pins MIR_CTL[3:0] are left floating, the MDS108 will operate with the port mirroring function disabled. When port mirroring is enabled, the user must configure Port 7 to operate in the same mode as the port it is mirroring (autonegotiation, duplex, speed, flow control).

Mirrored Port	Mirror_Control [3]	Mirror_Control [2]	Mirror_Control [1]	Mirror_Control [0]
Port 0 RX	1	0	0	0
Port 0 TX	0	0	0	0
Port 1 RX	1	0	0	1
Port 1 TX	0	0	0	1
Port 2 RX	1	0	1	0
Port 2 TX	0	0	1	0
Port 3 RX	1	0	1	1
Port 3 TX	0	0	1	1
Port 4 RX	1	1	0	0
Port 4 TX	0	1	0	0
Port 5 RX	1	1	0	1
Port 5 TX	0	1	0	1
Port 6 RX	1	1	1	0
Port 6 TX	0	1	1	0
Disabled	X	1	1	1

Table 2 - Port Mirroring Configuration

8.0 Power Saving Mode in MAC

The MDS108 was designed to be power efficient. When the internal MAC sections detect that the external port is not receiving or transmitting packets, it will shut off and conserve power. When new packet data is loaded into the output transmit FIFO of a MAC in power saving mode, the MAC will return to life and begin operating immediately.

When the MAC is in power saving mode and new packet data is received on the RMII, the MAC will return to life and receive data normally into the receive FIFO. This wakeup occurs when the MAC sees the Carrier Sense Data Valid (CRS_DV) signal asserted.

Using this method, the switch will turn off all MAC sections during periods when there is no network activity (at night, for example), and save power. For large networks this power savings could be large. To achieve the maximum power efficiency, the designer should use a physical layer transceiver that utilizes "Wake-On-LAN" technology.

9.0 XLink Cascade Port

Port 8 on the MDS108 is the XLink port. The XLink port provides a means of packet communication between two MDS108 chips. When cascaded via the XLink port, two MDS108 devices are capable of providing up to sixteen 10/100 Mbps ports.

The XLink port utilizes an industry standard MII interface. This expansion port is capable of operating at several data rates, depending upon configuration. Multiple data rates are achieved by multiplying the MII reference clock for the XLink interface (M8_REFCLK). When operating at greater than 200 Mbps, the MDS108 must utilize a higher frequency system clock in order to be able to support the greater switching bandwidth requirements. With a 66 MHz system clock the MDS108 can support full wire-speed forwarding on all 8 ports, and full wire speed on the XLink port for expansion port data rates of 200 Mbps and 100 Mbps. With a 80 MHz system clock, the MDS108 can support full wire-speed operation on all 8 ports and the XLink at data rates up to 400 Mbps.

The XLink frequency of operation is determined by strap options during power-up. The frame buffer memory address pins L_A[10:9] provide the frequency of operation selection as shown in the table.

Xlink Data Rate	L_A[10:9] Strap Level
100 Mbps	11
200 Mbps	10
300 Mbps	01
400 Mbps	00

Table 3 - XLink Port Strap Options

When two MDS108 devices are cascaded to form a 16 port switch, the XLink port on each MDS108 is treated as if it were a standard switch output. Standard Ethernet packet protocols are utilized, traditional packet integrity checks are performed at the receiving MDS108 XLink port, and standard length Ethernet packets are transmitted.

Refer to the XLink Application Note (MSAN-216).

10.0 EEPROM I²C Interface

A simple 2 wire serial interface is provided to allow the configuration of the MDS108 from an external EEPROM. The MDS108 utilizes a 1 K bit EEPROM with an I²C interface.

11.0 Management Interface

The MDS108 uses a standard parallel port interface to provide external CPU access to the internal registers. This parallel interface is composed of 3 pins: DATA0, STROBE, and ACK. The DATA0 pin provides the address and data content input to the MDS108, while the ACK pin provides the corresponding output to the external CPU. The STROBE pin is provided as the clock for both serial data streams. Any of the MDS108 internal registers can be modified through this parallel port interface¹.




Figure 3 - Write Command

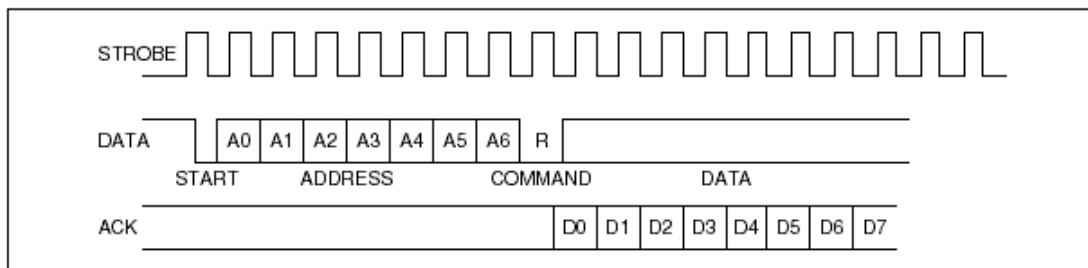


Figure 4 - Read Command

Each management interface transfer consists of four parts:

1. A START pulse – occurs when DATA is sampled high when STROBE is rising followed by DATA being sampled low when STROBE falls.
2. Register address strobed into DATA0 pin, using the high level of the STROBE pin.
3. Either a read or write command (see waveforms above).
4. Data to be written provided on DATA0, or data to be read provided on ACK.

Any command can be aborted in the middle by sending an ABORT pulse to the MDS108. An ABORT pulse occurs when DATA is sampled low and STROBE is rising, followed by DATA being sampled high when STROBE falls.

1. The 3-bit parallel interface is not "parallel" in the usual sense of the word; it is actually a synchronous serial architecture. However, the MDS108 management interface adheres to IEEE 1284 parallel port standards.

12.0 Configuration Register Definitions

The MDS108 registers can be accessed via the parallel interface and/or the I²C interface. Some registers are only accessible through the parallel interface. The access method for each register is listed in the individual register definitions. Each register is 8 bits wide.

12.1 GCR - Global Control Register

- Access: parallel interface, Write Only
- Address: h30

Bit 0	Save configuration into EEPROM Write '1' followed by a '0'	(Default = 0)
Bit 1	Save configuration into EEPROM and reset system Write '1' (self-clearing due to reset)	(Default = 0)
Bit 2	Start Built-In Self-Test (BIST) Write '1' followed by a '0'	(Default = 0)
Bit 3	Reset system Write '1' (self-clearing due to reset)	(Default = 0)
Bit [7:4]	Reserved	

12.2 DCR - Device Status and Signature Register

- Access: parallel interface, Read Only
- Address: h31

Bit 0	Busy writing configuration from I ² C 1: Activity 0: No Activity
Bit 1	Busy reading configuration from I ² C 1: Activity 0: No Activity
Bit 2	Built-In Self-Test (BIST) in progress 1: BIST In-Progress 0: Normal Mode
Bit 3	RAM error during BIST 1: RAM Error 0: No Error
Bits [5:4]	Reserved
Bits [7:6]	Revision number 00: Initial Silicon 01: Second Silicon

12.3 DA – DA Register

- Access: parallel interface, Read Only
- Address: h36

Returns 8-bit value DA (hexadecimal) if the parallel port connection is good. Otherwise, returns some other value indicating failure.

12.4 MBCR – Multicast Buffer Control Register

- Access: parallel interface and I²C, Read/Write
- Address: h00

Bit [7:0]	MAX_CNT_LMT	Maximum number of multicast frames allowed to be buffered inside the device at any one time	(Default = 80)
-----------	-------------	---	----------------

12.5 FCBST – FCB BUFFER LOW THRESHOLD

- Access: parallel interface and I²C, Read/Write
- Address: h01

Bits [5:0]	BUF_LOW_TH	Buffer Low Threshold – the number of free buffers below which flow control or WRED is triggered	(Default = 3F)
Bit 6		Use IP precedence subfield (TOS[0:2]) for Transmission Priority	(Default = 0)
Bit 7		Use IP precedence subfield (TOS[0:2]) for Drop Level	(Default = 0)

Note that, for Bits 6 and 7, Default = 0 means to use TOS[3:5].

12.6 LPBT – Low Drop Precedence Buffer Threshold

- Access: parallel interface and I²C, Read/Write
- Address: h02

Bits [7:0]	LOW_DROP_CNT	Number of frame buffers reserved for low-drop traffic	(Default 3F)
------------	--------------	---	--------------

12.7 FCR – Flooding Control Register

- Access: parallel interface and I²C, Read/Write
- Address: h03

Bits [3:0]	U2MR	Maximum number of flooded frames allowed within any time interval indicated by TimeBase bits (violations are discarded)	(Default = 8)
Bits [6:4]	TimeBase	000 = 100 μs 001 = 200 μs 010 = 400 μs 011 = 800 μs 100 = 1.6 μs 101 = 3.2 μs 110 = 6.4 μs 111 = 100 μs	(Default = 000)
Bit 7	USE_TOS	Use TOS instead of VLAN priority for IP packet	(Default = 0)

12.8 AVTCL – VLAN Type Code Register Low

- Access: parallel interface and I²C, Read/Write
- Address: h04

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

12.9 AVTCH – VLAN Type Code Register High

- Access: parallel interface and I²C, Read/Write
- Address: h05

Bit [7:0]	VLANType_HIGH	Upper 8 bits of VLAN type code	(Default 81)
-----------	---------------	--------------------------------	--------------

12.10 AVPM – VLAN Priority Map

- Access: parallel interface and I²C, Read/Write
- Address: h06

Map VLAN tag into 2 transmit queues (0 = low priority, 1 = high priority)

Bit 0	Mapped priority of tag value 0	(Default 0)
Bit 1	Mapped priority of tag value 1	(Default 0)
Bit 2	Mapped priority of tag value 2	(Default 0)
Bit 3	Mapped priority of tag value 3	(Default 0)
Bit 4	Mapped priority of tag value 4	(Default 0)
Bit 5	Mapped priority of tag value 5	(Default 0)
Bit 6	Mapped priority of tag value 6	(Default 0)
Bit 7	Mapped priority of tag value 7	(Default 0)

12.11 AVDM – VLAN Discard Map

- Access: parallel interface and I²C, Read/Write
- Address: h07

Map VLAN tag into frame discard levels (0 = low drop, 1 = high drop).

Bit 0	Frame discard for tag value 0	(Default 0)
Bit 1	Frame discard for tag value 1	(Default 0)
Bit 2	Frame discard for tag value 2	(Default 0)
Bit 3	Frame discard for tag value 3	(Default 0)
Bit 4	Frame discard for tag value 4	(Default 0)
Bit 5	Frame discard for tag value 5	(Default 0)
Bit 6	Frame discard for tag value 6	(Default 0)
Bit 7	Frame discard for tag value 7	(Default 0)

12.12 TOSPM – TOS Priority Map

- Access: parallel interface and I²C, Read/Write
- Address: h08

Map TOS field in IP packet into 2 transmit queues (0 = low priority, 1 = high priority).

Bit 0	Mapped priority when TOS is 0	(Default 0)
Bit 1	Mapped priority when TOS is 1 ¹	(Default 0)
Bit 2	Mapped priority when TOS is 2	(Default 0)
Bit 3	Mapped priority when TOS is 3	(Default 0)
Bit 4	Mapped priority when TOS is 4	(Default 0)
Bit 5	Mapped priority when TOS is 5	(Default 0)
Bit 6	Mapped priority when TOS is 6	(Default 0)
Bit 7	Mapped priority when TOS is 7	(Default 0)

1. TOS = 1 means the appropriate 3-bit TOS subfield is "001".

12.13 PTPRI – Port Priority

- Access: parallel interface and I²C, Read/Write
- Address: h09

Enable and configure port-based priorities for ports 0, 1, 2, and 3

Bit 0	EN0	Port 0: Enable; 1 = enabled	(Default 0)
Bit 1	P0	Port 0: Priority; 1 = high, 0 = low	(Default 0)
Bit 2	EN1	Port 1: Enable; 1 = enabled	(Default 0)
Bit 3	P0	Port 1: Priority; 1 = high, 0 = low	(Default 0)

Bit 4	EN2	Port 2: Enable; 1 = enabled	(Default 0)
Bit 5	P2	Port 2: Priority; 1 = high, 0 = low	(Default 0)
Bit 6	EN3	Port 3: Enable; 1 = enabled	(Default 0)
Bit 7	P3	Port 3: Priority; 1 = high, 0 = low	(Default 0)

12.14 TOSDM – TOS Discard Map

- Access: parallel interface and I²C, Read/Write
- Address: h0A

Map TOS into frame discard levels (0 = low drop, 1 = high drop).

Bit 0	Frame discard when TOS is 0	(Default 0)
Bit 1	Frame discard when TOS is 1	(Default 0)
Bit 2	Frame discard when TOS is 2	(Default 0)
Bit 3	Frame discard when TOS is 3	(Default 0)
Bit 4	Frame discard when TOS is 4	(Default 0)
Bit 5	Frame discard when TOS is 5	(Default 0)
Bit 6	Frame discard when TOS is 6	(Default 0)
Bit 7	Frame discard when TOS is 7	(Default 0)

12.15 AXSC – Transmission Scheduling Control Register

- Access: parallel interface and I²C, Read/Write
- Address: h0B

Bits [3:0]	Transmission Queue Service Weight for high priority queue	(Default F)
Bit [5]	Global Quality of Service Enable	(Default 0)
Bit [6]	Global Flow Control Disable – if 1, flow control is disabled globally; if 0, each port's flow control settings are separately configurable	(Default 0)
Bit [7]	Half Duplex Priority Enable – if 0, priority is disabled for all half duplex ports; if 1, priority is enabled unless AXSC[5] = 0	(Default 0)

12.16 MII_OP0 – MII Register Option 0

- Access by parallel interface and I²C, Read/Write
- Address: h0C

Permits a non-standard address for the Phy Status Register. When low and high Address bytes are 0, the MDS108 will use the standard address.

Bit [7:0]	Low order address byte	(Default 00)
-----------	------------------------	--------------

12.17 MII_OP1 – MII Register Option 1

- Access: parallel interface and I²C, Read/Write
- Address: h0D

Bit [7:0] High order address byte (Default 00)

12.18 AGETIME_LOW – MAC Address Aging Timer Low

- Access: parallel interface and I²C, Read/Write
- Address: h0E

Bit [7:0] Low byte of the MAC address aging timer (Default 25)

12.19 AGETIME_HIGH – MAC Address Aging Timer High

- Access: parallel interface and I²C, Read/Write
- Address: h0F

Bit [7:0] High byte of the MAC address aging timer. (Default 01)

The aging time is based on the following formula:
 {AGETIME_HIGH, AGETIME_LOW} x 1024 ms

The default setting provides a 300 second aging time.

12.20 ECR1P0 – Port 0 Control Register

- Access: parallel interface and I²C, Read/Write
- Address: h10

Bits [3:0] Port Mode (Default 0000)

Bit [3] 1 – Force configuration based on Bits [2:0]
 0 – Autonegotiate and advertise based on Bits [2:0]

Bit [2] 1 – 10 Mbps
 0 – 100 Mbps

Bit [1] 1 – Half Duplex
 0 – Full Duplex

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

Bits [6:4] PVID Port-based VLAN ID (Default 000)

Bit [7] Reserved

12.21 ECR1P1 – Port 1 Control Register

- Access: parallel interface and I²C, Read/Write
- Address: h11

Bits [3:0]	Port Mode		(Default 0000)
Bit [3]		1 – Force configuration based on Bits [2:0] 0 – Autonegotiate and advertise based on Bits [2:0]	
Bit [2]		1 – 10 Mbps 0 – 100 Mbps	
Bit [1]		1 – Half Duplex 0 – Full Duplex	
Bit [0]		1 – Flow Control Off 0 – Flow Control On	
Bits [6:4]	PVID	Port-based VLAN ID	(Default 000)
Bit [7]	Reserved		

12.22 ECR1P2 – Port 2 Control Register

- Access: parallel interface and I²C, Read/Write
- Address: h12

Bits [3:0]	Port Mode		(Default 0000)
Bit [3]		1 – Force configuration based on Bits [2:0] 0 – Autonegotiate and advertise based on Bits [2:0]	
Bit [2]		1 – 10 Mbps 0 – 100 Mbps	
Bit [1]		1 – Half Duplex 0 – Full Duplex	
Bit [0]		1 – Flow Control Off 0 – Flow Control On	
Bits [6:4]	PVID	Port-based VLAN ID	(Default 000)
Bit [7]	Reserved		

12.23 ECR1P3 – Port 3 Control Register

- Access: parallel interface and I²C, Read/Write
- Address: h13

Bits [3:0]	Port Mode		(Default 0000)
------------	-----------	--	----------------

Bit [3]		1 – Force configuration based on Bits [2:0] 0 – Autonegotiate and advertise based on Bits [2:0]	
Bit [2]		1 – 10 Mbps 0 – 100 Mbps	
Bit [1]		1 – Half Duplex 0 – Full Duplex	
Bit [0]		1 – Flow Control Off 0 – Flow Control On	
Bits [6:4]	PVID	Port-based VLAN ID	(Default 000)
Bit [7]	Reserved		

12.24 ECR1P4 – Port 4 Control Register

- Access: parallel interface and I²C, Read/Write
- Address: h14

Bits [3:0]	Port Mode		(Default 0000)
Bit [3]		1 – Force configuration based on Bits [2:0] 0 – Autonegotiate and advertise based on Bits [2:0]	
Bit [2]		1 – 10 Mbps 0 – 100 Mbps	
Bit [1]		1 – Half Duplex 0 – Full Duplex	
Bit [0]		1 – Flow Control Off 0 – Flow Control On	
Bits [6:4]	PVID	Port-based VLAN ID	(Default 000)
Bit [7]	Reserved		

12.25 ECR1P5 – Port 5 Control Register

- Access: parallel interface and I²C, Read/Write
- Address: h15

Bits [3:0]	Port Mode		(Default 0000)
Bit [3]		1 – Force configuration based on Bits [2:0] 0 – Autonegotiate and advertise based on Bits [2:0]	
Bit [2]		1 – 10 Mbps 0 – 100 Mbps	

Bit [1]		1 – Half Duplex 0 – Full Duplex	
Bit [0]		1 – Flow Control Off 0 – Flow Control On	
Bits [6:4]	PVID	Port-based VLAN ID	(Default 000)
Bit [7]	Reserved		

12.26 ECR1P6 – Port 6 Control Register

- Access: parallel interface and I²C, Read/Write
- Address: h16

Bits [3:0]	Port Mode		(Default 0000)
Bit [3]		1 – Force configuration based on Bits [2:0] 0 – Autonegotiate and advertise based on Bits [2:0]	
Bit [2]		1 – 10 Mbps 0 – 100 Mbps	
Bit [1]		1 – Half Duplex 0 – Full Duplex	
Bit [0]		1 – Flow Control Off 0 – Flow Control On	
Bits [6:4]	PVID	Port-based VLAN ID	(Default 000)
Bit [7]	Reserved		

12.27 ECR1P7 – PORT 7 CONTROL REGISTER

- Access: parallel interface and I²C, Read/Write
- Address: h17

Bits [3:0]	Port Mode		(Default 0000)
Bit [3]		1 – Force configuration based on Bits [2:0] 0 – Autonegotiate and advertise based on Bits [2:0]	
Bit [2]		1 – 10 Mbps 0 – 100 Mbps	
Bit [1]		1 – Half Duplex 0 – Full Duplex	
Bit [0]		1 – Flow Control Off 0 – Flow Control On	
Bits [6:4]	PVID	Port-based VLAN ID	(Default 000)
Bit [7]	Reserved		

12.28 ECR1P8 – Port 8 Control Register

- Access: parallel interface and I²C, Read/Write
- Address: h18

Bits [3:0]	Port Mode		(Default 0000)
Bit [3]		1 – Force configuration based on Bits [2:0] 0 – Autonegotiate and advertise based on Bits [2:0]	
Bit [2]		1 – 10 Mbps 0 – 100 Mbps	
Bit [1]		1 – Half Duplex 0 – Full Duplex	
Bit [0]		1 – Flow Control Off 0 – Flow Control On	
Bits [6:4]	PVID	Port-based VLAN ID	(Default 000)
Bit [7]	Reserved		

12.29 FC_0 – Flow Control Byte 0

- Access: parallel interface and I²C, Read/Write
- Address: h19

The flow control hold time parameter is the length of time a flow control message is effectual (i.e. halts incoming traffic) after being received. The hold time is measured in units of “slots,” the time it takes to transmit 64 bytes at wire-speed. The default setting is 32 slots, or for a 100 Mbps port, approximately 164 s.

Bits [7:0] Flow control hold time byte 0 (Default FF)

12.30 FC_1 – Flow Control Byte 1

- Access: parallel interface and I²C, Read/Write
- Address: h1A

Bits [7:0] Flow control hold time byte 1 (Default 00)

12.31 FC_2 – Flow Control CRC Byte 0

- Access: parallel interface and I²C, Read/Write
- Address: h1B

Bits [7:0] Flow control frame CRC byte 0 (Default 96)

12.32 FC_3 – Flow Control CRC Byte 1

- Access: parallel interface and I²C, Read/Write
- Address: h1C

Bits [7:0] Flow control frame CRC byte 1 (Default 8E)

13.0 MDS108 Pin Descriptions

Note:

- # Active low signal
- I Input signal
- S Input signal with Schmitt-Trigger
- O Output signal
- OD Open-Drain driver
- I/O Input & Output signal
- SL Slew Rate Controlled
- D Pulldown
- U Pullup
- 5 5V Tolerance

Pin No(s).	Symbol	Type	Name & Functions
Frame Buffer Memory Interface			
201, 200, 199, 197, 196, 195, 193, 192, 191, 190, 188, 187, 186, 185, 183, 182, 181, 179, 178, 177, 176, 174, 173, 172, 170, 169, 168, 167, 165, 164, 163, 161	L_D[31:0]	I/O, U, SL	Databus to Frame Buffer Memory
203, 151, 158, 160, 10, 9, 8, 6, 5, 4, 2, 1, 208, 206, 205, 204, 150	L_A[18:2]	I/O, U, SL	Address pins for buffer memory
153	L_CLK	O	Frame Buffer Memory Clock
155	L_WE#	O, SL	Frame Buffer Memory Write Enable
156	L_OE#	O	Frame Buffer Memory Output Enable
157	L_ADSC#	O, SL	Frame Buffer Memory Address Status Control
MII Management Interface			
120	M_MDC	O	MII Management Data Clock
122	M_MDIO	I/O, U	MII Management Data I/O
I²C Interface (Serial EEPROM Interface)			
123	SCL	O, U, 5	I ² C Data Clock
124	SDA	I/O, U, OD, 5	I ² C Data I/O
Parallel Port Management Interface			
127	STROBE	I, U, S, 5	Strobe Pin
128	DATA0	I, U, 5	Data Pin
129	ACK	O, U, OD, 5	Acknowledge Pin
Port 0 RMII Interface			
24, 23	M0_RXD[1:0]	I, U	Port 0 Receive Data
22	M0_CRSDV	I, D	Port 0 Carrier Sense and Data Valid
21, 20	M0_TXD[1:0]	O	Port 0 Transmit Data
19	M0_TXEN	O	Port 0 Transmit Enable
Port 1 RMII Interface			
31, 30	M1_RXD[1:0]	I, U	Port 1 Receive Data
29	M1_CRSDV	I, D	Port 1 Carrier Sense and Data Valid

Pin No(s).	Symbol	Type	Name & Functions
28, 27	M1_TXD[1:0]	O	Port 1 Transmit Data
26	M1_TXEN	O	Port 1 Transmit Enable
Port 2 RMII Interface			
38, 37	M2_RXD[1:0]	I, U	Port 2 Receive Data
36	M2_CRSDV	I, D	Port 2 Carrier Sense and Data Valid
35, 34	M2_TXD[1:0]	O	Port 2 Transmit Data
33	M2_TXEN	O	Port 2 Transmit Enable
Port 3 RMII Interface			
45, 44	M3_RXD[1:0]	I, U	Port 3 Receive Data
43	M3_CRSDV	I, D	Port 3 Carrier Sense and Data Valid
42, 41	M3_TXD[1:0]	O	Port 3 Transmit Data
40	M3_TXEN	O	Port 3 Transmit Enable
Port 4 RMII Interface			
65, 64	M4_RXD[1:0]	I, U	Port 4 Receive Data
63	M4_CRSDV	I, D	Port 4 Carrier Sense and Data Valid
62, 61	M4_TXD[1:0]	O	Port 4 Transmit Data
60	M4_TXEN	O	Port 4 Transmit Enable
Port 5 RMII Interface			
72, 71	M5_RXD[1:0]	I, U	Port 5 Receive Data
70	M5_CRSDV	I, D	Port 5 Carrier Sense and Data Valid
69, 68	M5_TXD[1:0]	O	Port 5 Transmit Data
67	M5_TXEN	O	Port 5 Transmit Enable
Port 6 RMII Interface			
79, 78	M6_RXD[1:0]	I, U	Port 6 Receive Data
77	M6_CRSDV	I, D	Port 6 Carrier Sense and Data Valid
76, 75	M6_TXD[1:0]	O	Port 6 Transmit Data
74	M6_TXEN	O	Port 6 Transmit Enable
Port 7 RMII Interface			
86, 85	M7_RXD[1:0]	I, U	Port 7 Receive Data
84	M7_CRSDV	I, D	Port 7 Carrier Sense and Data Valid
83, 82	M7_TXD[1:0]	O	Port 7 Transmit Data
81	M7_TXEN	O	Port 7 Transmit Enable
Port 8 MII Interface			
105, 104, 103, 102	M8_RXD[3:0]	I, U	Port 8 Receive Data
113, 112, 111, 110	M8_TXD[3:0]	O	Port 8 Transmit Data
109	M8_TXEN	O	Port 8 Transmit Enable
97	M8_RXDV	I, D	Port 8 Receive Data Valid
100	M8_RXCLK	I, U	Port 8 Receive Clock
107	M8_TXCLK	I/O, U	Port 8 Transmit Clock

Pin No(s).	Symbol	Type	Name & Functions
114	M8_LINK	I, U	Port 8 Link Status
116	M8_SPEED	I/O, U	Port 8 Speed Select (100 Mb = 1)
115	M8_DUPLEX	I, U	Port 8 Full Duplex Select (half duplex = 0)
98	M8_COL	I, U	Port 8 Collision Detect
118	M8_REFCLK	I/O, U	Port 8 Reference Clock
Port 8 Serial Interface			
102	S8_RXD	I, U	Port 8 Serial Receive Data
100	S8_RXCLK	I, U	Port 8 Serial Receive Clock
97	S8_CRS_DV	I, D	Port 8 Serial Carrier Sense and Data Valid
110	S8_TXD	O	Port 8 Serial Transmit Data
107	S8_TXCLK	I	Port 8 Serial Transmit Clock
109	S8_TXEN	O	Port 8 Serial Transmit Enable
98	S8_COL	I, U	Port 8 Serial Collision Detect
114	S8_LINK	I, U	Port 8 Link Status
115	S8_DUPLEX	I, U	Port 8 Full-Duplex Select (half duplex = 0)
Miscellaneous Control Pins			
95	M_CLK	I	Reference RMII Clock
148	SCLK	I	System Clock (50 - 80 MHz)
125	TEST#	I, U	Manufacturing Pin. Leave as No Connect (NC)
126	TRUNK_EN	I, D	Port Trunking Enable
146, 145, 144, 143	MIR_CTL[3:0]	I/O, U	Port Mirroring Control
142	RESIN#	I, S	Reset Pin
141	RESETOUT#	O	PHY Reset Pin
Test Pins			
139	TMODE#	I/O, U	Manufacturing Pin. Puts device into test mode for ATE test. Leave as No Connect (NC)
138, 137, 136, 135	TSTOUT[7:4]	O	Test Outputs
134, 133, 132, 131	TSTOUT[3:0]	I/O, U	Test Outputs
NC Pins			
12, 13, 14, 15, 16, 17, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 88, 89, 90, 91, 92, 93	N/C Reserved		No Connect
Power Pins			
3, 39, 73, 96, 130, 159, 184	VDD (Core)	Input	+3.3 Volt DC Supply for Core Logic (7 pins)
11, 25, 59, 87, 101, 108, 119, 147, 152, 166, 175, 194, 202	VDD	Input	+3.3 Volt DC Supply for I/O Pads (13 pins)
18, 46, 80, 106, 140, 171, 198	VSS (Core)	Input	Ground for Core Logic (7 pins)
7, 32, 66, 94, 99, 117, 121, 149, 154, 162, 180, 189, 207	VSS	Input	Ground for I/O Pads (13 pins)

13.1 M8_REFCLK Speed Requirement

The M8_REFCLK pin can be either input or output pin depending on the Port 8 configuration.

- As an output pin, its output frequency is always equal to half of the M_CLK
- As an input, it provides a reference clock source to the Port 8 MAC when it is configured in speed-up mode
- M8_REFCLK input frequency is equal to 25% of the Port 8 speed (data rate)

Port 8 Configuration		M8_REFCLK	
Mode	Speed	Input/Output	Freq.
Reg.	10/100 M	Output	M_CLK/2
2x	200 M	Input	50 MHz
3x	300 M	Input	75 MHz
4x	400 M	Input	100 MHz

13.2 Strap Options

The Strap options are relevant during the initial power-on period, when reset is asserted. During reset, the MDS108 will examine the boot strap address pin to determine its value and modify the internal configuration of the chip accordingly.

"1" mean Pull-up

"0" means Pull-down with an external 1 K Ohm.

Default value is 1, (all boot strap pins have internal pull up resistor).

Pin No(s).	Symbol	Name & Functions
Boot Strap Pins		
206 (L_A[5])	Memory Size	1 - Memory size = 256 KB 0 - Memory size = 512 KB
208 (L_A [6])	EEPROM	1 - NO EEPROM Installed 0 - EEPROM Installed ¹
5, 4 (L_A [10:9])	XLINK Speed	11 - 100 Mbps 10 - 200 Mbps 01 - 300 Mbps 00 - 400 Mbps (0 - Pull down, 1 - Pull up)
151 (L_A[17])	Port 8 MII/Serial	1 - MII Mode for port 8 0 - Serial Mode for port 8
150 (L_A[2])	Link Polarity	Link Polarity for serial interface 1 - Active Low 0 - Active High
204 (L_A[3])	FDX Polarity	Full/Half Duplex Polarity for serial interface 1 - Active Low 0 - Active High
205 (L_A[4])	SPD100 Polarity	Speed polarity for serial interface 1 - Active Low 0 - Active High
2 (L_A[8])	Device ID	Use in cascade mode only
133 (TST[2])	SBRAM Self Test	For Board/System Manufacturing Test ² 1 - Disable 0 - Enable

Note 1: If the MDS108AL is configured from EEPROM preset (L_A[6] pulled down at reset), it will try to load its configuration from the EEPROM. If the EEPROM is blank or not preset, it will not boot up. The parallel port can be used to program the EEPROM at any time.

Note 2: During normal power-up, the MDS108 will run through an external SBRAM memory test to ensure that there are no memory interface problems. If a problem is detected, the chip will stop functioning. To facilitate board debug in the event that a system stops functioning, the MDS108AL can be put into a continuous SBRAM self test mode to allow an operator to determine if there are stuck pins in the memory interface (using network analyzer).

14.0 Characteristics and Timing

14.1 Absolute Maximum Rating

Storage Temperature	-65°C to +150C
Operating Temperature	-40C to +85C
Maximum Junction Temperature	+125C
Supply Voltage VDD with Respect to Vss	+3.0 V to +3.6 V
Voltage on 5V Tolerant Input Pins	-0.5 V to (VDD +3.3 V)
Voltage on Other Pins	-0.5V to (VDD +0.3 V)

Caution: Stresses above those listed may cause permanent device failure. Functionality at or above these limits is not implied. Exposure to the absolute Maximum Ratings for extended Ratings for extended periods may affect device reliability.

14.2 DC Electrical Characteristics

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

14.3 Recommended Operating Conditions

Symbol	Parameter Description	Min.	Typ.	Max.	Unit
f _{OSC}	Frequency of Operation	50	66	80	MHz
I _{DD}	Supply Current - @ 55 MHz, 8x100M, 100% Full Duplex Traffic (VDD = 3.3V)		580		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		VDD + 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} < VDD) (all pins except those with internal pull-up/ pull-down resistors)			10	μA
I _{OL}	Output Leakage Current (0.1 V < V _{OUT} < VDD)			10	μA
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			29.7	C/W
θ _{ja}	Thermal resistance with 1 m/s air flow			28.8	C/W
θ _{ja}	Thermal resistance with 2 m/s air flow			26.8	C/W
θ _{jc}	Thermal resistance between junction and case			12.6	C/W

14.4 Clock Frequency Specifications

Symbol	Parameter	(Hz)	Note:
C1	SCLK - Core System Clock Input	50 M	
C2	M_CLK - RMIID Port Clock	50 M	
C3	M8_REFCLK - MII Reference Clock	25 M	
C4	L_CLK - Frame Buffer Memory Clock	50 M	L_CLK = SCLK
C5	M_MDC - MII Management Data Clock	1.56 M	M_MDC = SCLK/32
C6	SCL - I ² C Data Clock	50 K	SCL = M_CLK/1000

Suggestion Clock rate for various configurations:

Configuration		Input			Output		
Port 0-7	Port 8	SCLK	M_CLK (RMII)	M8_REF	L_CLK	M_MDC	SCL
100 M RMII	Not Used	55 M	50 M	--	=SCLK	=SCLK/32	50 K
100 M RMII	10/100 M MII	60 M	50 M	--	=SCLK	=SCLK/32	50 K
100 M RMII	200 M MII	66.66 M	50 M	50 M	=SCLK	=SCLK/32	50 K
100 M RMII	300 M MII	75 M	50 M	75 M	=SCLK	=SCLK/32	50 K
100 M RMII	400 M MII	80 M	50 M	100 M	=SCLK	=SCLK/32	50 K

14.5 AC Timing Characteristics

14.5.1 Frame Buffer Memory Interface:




Figure 5 - Frame Buffer Memory Interface

Symbol	Parameter	50 MHz		Note
		Min. (ns)	Max. (ns)	
L1	L_D[31:0] input setup time	5		
L2	L_D[31:0] input hold time	0		
L3	L_D[31:0] output valid delay	1	8	$C_L = 30 \text{ pF}$
L4	L_A[18:2] output valid delay	1	8	$C_L = 50 \text{ pF}$
L6	L_ADSC# output valid delay	1	8	$C_L = 50 \text{ pF}$
L8	L_WE# output valid delay	1	8	$C_L = 30 \text{ pF}$
L9	L_OE# output valid delay	1	8	$C_L = 30 \text{ pF}$

Table 4 - Frame Buffer Memory Interface Timing

14.6 RMII Timing Requirements

Symbol	Parameter	50 MHz		Note:
		Min. (ns)	Max. (ns)	
M1	M_CLK			Reference Input Clock
M2	M[7:0]_RXD[1:0] input setup time	4		
M3	M[7:0]_RXD[1:0] input hold time	1		
M4	M[7:0]_CRS_DV input setup time	4		
M5	M[7:0]_TXEN output delay time	1	11	$C_L = 30$ pF
M6	M[7:0]_TXD[1:0] output delay time	1	11	$C_L = 30$ pF
M7	M[7:0]_LINK input setup time	4		

Table 5 - RMII Timing Requirements

14.7 MII Timing Requirements




Figure 6 - Transmit Timing

Symbol	Parameter	Time		Unit
		Min.	Max.	
1	M8_TXCLK rise to M8_TXD[3:0] inactive delay	5	20	ns
2	M8_TXCLK rise to M8_TXD[3:0] active delay	5	20	ns
3	M8_TXCLK rise to M8_TXEN active delay	5	20	ns
4	M8_TXCLK rise of last M8_TXD bit to M8_TXEN inactive delay	5	20	ns
5	M8_TXCLK High wide	25	Inf.	ns
6	M8_TXCLK Low wide	25	Inf.	ns
	M8_TXCLK input rise time require		5	ns
	M8_TXCLK input fall time require		5	ns

*Inf. = infinite

Table 6 - Transmit Timing Requirements






Figure 7 - Receive Timing

Symbol	Parameter	Time		Unit
		Min.	Max.	
1	M8_RXD[3:0] Low input setup time	10		ns
2	M8_RXD[3:0] Low input hold time		5	ns
3	M8_RXD[3:0] High input setup time	10		ns
4	M8_RXD[3:0] High input hold time		5	ns
5	M8_RXDV Low input setup time	10		ns
6	M8_RXDV Low input hold time		5	ns
7	M8_RXDV High input setup time	10		ns
8	M8_RXDV High input hold time		5	ns
9	M8_RXCLK High wide	25	Inf.	ns
10	M8_RXCLK Low wide	25	Inf.	ns
	M8_RXCLK input rise time require		5	ns
	M8_RXCLK input fall time require		5	ns

Table 7 - Receive Timing Requirements



Symbol	Control Dimensions in millimetres		
	MIN	Nominal	MAX
A	—	—	4.10
A1	0.25	—	—
A2	3.20	3.32	3.60
D	30.60 BSC		
D1	28.00 BSC		
E	30.60 BSC		
E1	28.00 BSC		
L	0.45	0.60	0.75
e	0.50 BSC		
b	0.17	0.20	0.27
c	0.09	0.15	0.20
θ	0°	—	7°
ccc	—	0.08	—
N	208		
ND	52		
NE	52		

Conforms to JEDEC MO-143

Notes:

1. Pin 1 indicator may be a corner chamfer, dot or both.
2. Controlling dimensions are in millimeters.
3. The top package body size may be smaller than the bottom package body size by a max. of 0.15 mm.
4. Dimension D1 and E1 do not include mould protusion.
5. Dimension b does not include dambar protusion.

© Zarlink Semiconductor 2003 All rights reserved.

ISSUE	1			
ACN	213984			
DATE	3Feb03			
APPRD.				



Package Code	QB
Previous package codes:	GP / L
	Package Outline for 208 Lead MQFP (28x28x3.32mm) + 2.6 mm (footprint)
	GPD00828



**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