ML4002-28-8W

Technical Reference

QSFP Electrical Passive Loopback Module

MSA Compliant

Innovation for the next generation

Table of Contents

1 Overview

The **ML4002-28-8W** is a QSFP28 passive electrical loopback module which is a hot pluggable form factor designed for high speed testing application for QSFP host ports. The **ML4002-28-8W** is designed for 100 Gigabit Ethernet applications and provides 4x28G RX and TX lanes, I2C module management interface and all the QSFP SFF hardware signals.
The **ML4002-28-8W** loops back 4-lane 28 Gb/s transmit data from the Host back to 4-lane 28 Gb/s receive data port to the Host.
The **ML4002-28-8W** provides programmable power dissipation up to 7.75 W allowing the module to emulate all the QSFP28 power classes. It also provides a voltage sense, insertion counter, power staging, LED blinking rate, upper temperature cut off and temperature sensor.

1.1 ML4002-28-8W QSFP Passive Module | Key Features

- Power Consumption up to 7.75 W
- Operation up to 28G per lane
- Dual LED indicator
- Custom Memory Maps
- 100% at rate AC testing, on each unit
- Temperature range from 0 ° to 80 °C
- MSA Compliant Memory Map
- High performance signal integrity traces
- Temperature Monitoring
- Voltage Monitoring
- Insertion Counter
- Power Staging
- Hot pluggable module
- Micro controller based

1.2 LED Indicator

Green (Solid) – Module is in high power mode.
Amber (Solid) – Module is in low power mode.
Green/ Amber (Blinking) – Module is in high power mode and Voltage or Temperature Alarm is triggered.

1.3 Recommended Operating Conditions

Parameter	Symbol	Notes/Conditions	Min	Typ	Max	Units
Operating Temperature	T_A		0		80	°C
Supply Voltage	VCC	Main Supply Voltage	3.00	3.3	3.5	V

Data Rate	Rb	Guaranteed to work at 28 Gbps per lane	0		100	Gbps
Input/output Load Resistance	RL	AC-Coupled, Differential	90	100	110	Ω
Power Class		Programmable to Emulate all power classes	0		7.75	W

2 QSFP Family

Function	ML4002-28	ML4012-28-C5	ML4012-28	ML4002-28-8W
Insertion Counter	YES	YES	NO	YES
Programmable Power Dissipation	YES	YES	NO (Set At Factory)	YES
Maximum Power Dissipation	3.5 W	5.0 W	4.5 W	7.75 W
Power Ramp-up / Staging	NO	YES	NO	YES
Voltage Sense	NO	YES	NO	YES
Temperature Monitor	YES	YES	NO	YES
Micro-Controller Based	YES	YES	NO	YES
EEPROM Based	NO	NO	YES	YES
LED	YES	YES	YES	YES
Thermal Cut-Off	YES	YES	NO	YES
Temperature Range	0-80˚C	0-80˚C	0-70˚C	0-80˚C

3 Functional Description

3.1 Management Data Interface – I2C

The **ML4002-28-8W** supports the I2C interface. This QSFP+ specification is based on the SFF8436 specification. Address 128 Page00 is used to indicate the use of the QSFP+ memory map rather than the QSFP memory map.

3.2 I2C Signals, Addressing and Frame Structure

3.2.1 I2C Frame

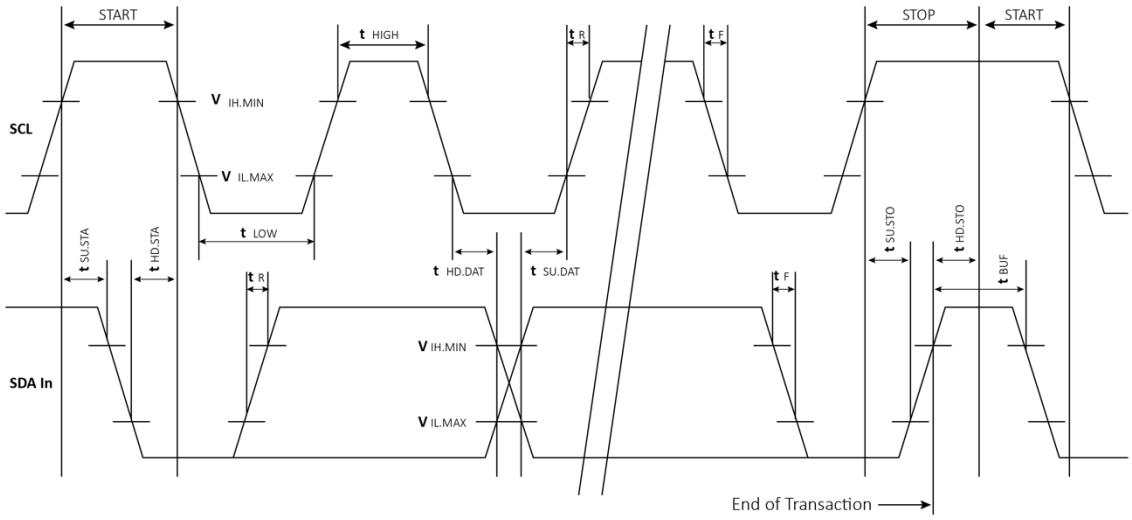

Figure 1: QSFP Timing Diagram

Before initiating a 2-wire serial bus communication, the host should provide setup time on the ModSelL line of all modules on the 2-wire bus. The host should not change the ModSelL line of any module until the 2-wire serial bus communication is complete and the hold time requirement is satisfied. The 2-wire serial interface address of the QSFP+ module is 1010000X (A0h). In order to allow access to multiple QSFP+ modules on the same 2-wire serial bus, the QSFP+ pinout includes a ModSelL or module select pin. This pin (which is pulled high or deselected in the module) must be held low by the host to select the module of interest and allow communication over the 2-wire serial interface. The module will not respond to or accept 2-wire serial bus instructions unless it is selected.

3.2.2 Management Timing Parameters

The timing parameters for the 2-Wire interface to the QSFP module are shown in the table below:

Parameter	Symbol	Min	Max	Unit
Clock Frequency	f_{SCL}	30	400	kHz
Clock Pulse Width Low	t_{LOW}	1.2		us
Clock Pulse Width High	t_{High}	1.1		us
Time bus free before new transmission can start	t_{BUF}	20.8		us
Input Rise Time (400KHz)	$T_{R,400}$	300		ns
Input Fall Time (400KHz)	$T_{F,400}$	300		ns
ModSelL Setup Time	Host_select_setup	2		ms
ModSelL Hold Time	Host_select_hold	10		us
Serial Interface Clock Holdoff (clock stretching)	T_clock_hold		500	us

3.2.3 Device Addressing and Operation

Serial Clock (SCL): The host supplied SCL input to QSFP transceivers is used to positive-edge clock data into each QSFP device and negative-edge clock data out of each device.

Serial Data (SDA): The SDA pin is bi-directional for serial data transfer. This pin is open-drain or open-collector driven.

Master/Slave: QSFP transceivers operate only as slave devices. The host must provide a bus master for SCL and initiate all read/write communication.

Device Address: Each QSFP is hard wired at the device address A0h.

Multiple Devices per SCL/SDA: While QSFP transceivers are compatible with point-to-point SCL/SDA, they can share a single SCL/SDA bus by using the QSFP ModSelL line.

Clock and Data Transitions: The SDA pin is normally pulled high with an external device. Data on the SDA pin may change only during SCL low time periods. Data changes during SCL high periods indicate a START or STOP condition. All addresses and data words are serially transmitted to and from the QSFP in 8-bit words. Every Byte on the SDA line must be 8-bits long. Data is transferred with the most significant bit (MSB) first.

START Condition: A high-to-low transition of SDA with SCL high is a START condition, which must precede any other command.

STOP Condition: A low-to-high transition of SDA with SCL high is a STOP condition.

Acknowledge: After sending each 8-bit word, the transmitter releases the SDA line for one-bit time, during which the receiver is allowed to pull SDA low (zero) to acknowledge (ACK) that it has received each word.

Memory (Management Interface) Reset: After an interruption in protocol, power loss, or system reset, the QSFP Module management interface can be reset. Memory reset is intended *only* to reset the QSFP transceiver management interface (to correct a hung bus). No other module functionality is implied.

1. Clock up to 9 cycles.
2. Look for SDA high in each cycle while SCL is high.
3. Create a Start condition as SDA is high.

Device Addressing: QSFP devices require an 8-bit device address word following a start condition to enable a read or write operation. The device address word consists of a mandatory sequence for the first seven most significant bits in Figure 2. This is common to all QSFP devices.

1	0	1	0	0	0	0	R/W
MSB							LSB

Figure 2: QSFP Device Address

The eighth bit of the device address is the read/write operating select bit. A read operation is initiated if this bit is set high and a write operation is initiated if this bit is set low. Upon compare of the device address (with ModSelL in the low state) the QSFP Module will output a zero (ACK) on the SDA line to acknowledge the address.

3.3 I2C Read/Write Functionality

3.3.1 QSFP+ Memory Address Counter (Read and Write Operations)

QSFP+ devices maintain an internal data word address counter containing the last address accessed during the latest read or write operation, incremented by one. The address counter is incremented whenever a data word is received or sent by the module. This address stays valid between operations as long as QSFP+ power is maintained. The address “roll over” during read and writes operations is from the last byte of the 128-byte memory page to the first byte of the same page.

3.3.2 Read Operations

3.3.2.1 Current Address Read

A current address read operation requires only the device address read word (10100001) be sent, see Figure 3 below.

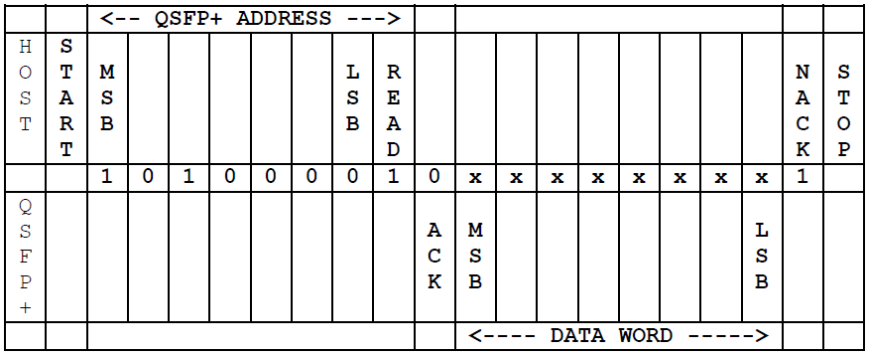

Figure 3: QSFP+ Current Address Read Operation

Once acknowledged by the QSFP+, the current address data word is serially clocked out. The host does not respond with an acknowledgement, but does generate a STOP condition once the data word is read.

3.3.2.2 Random Read

A random read operation requires a “dummy” write operation to load in the target byte address as shown in Figure 3 below. This is accomplished by the following sequence.

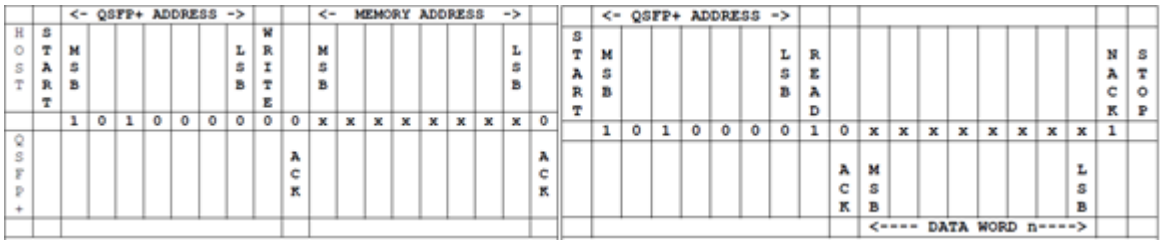

Figure 4: QSFP+ Random Read

The target 8-bit data word address is sent following the device address write word (10100000) and acknowledged by the QSFP+. The host then generates another START condition (aborting the dummy write without incrementing the counter) and a current address read by sending a device read address (10100001). The QSFP+ acknowledges the device address and serially clocks out the requested data word. The host does not respond with an acknowledgement, but does generate a STOP condition once the data word is read.

3.3.2.3 Sequential Read

Sequential reads are initiated by a current address read (Figure 4). To specify a sequential read, the host responds with an acknowledgement (instead of a STOP) after each data word. As long

as the QSFP+ receives an acknowledgement, it will serially clock out sequential data words. The sequence is terminated when the host responds with a NACK and a STOP instead of an acknowledgement.

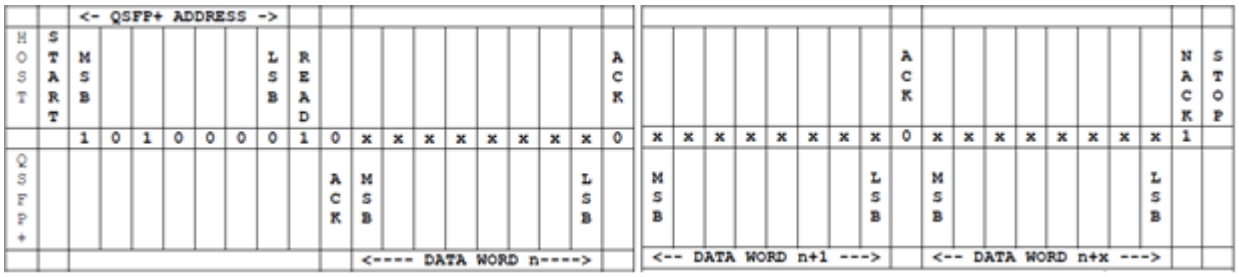

Figure 5: Sequential Address Read Starting at QSFP+ Current Address

3.4 Low Speed Signals

In addition to the 2-wire serial interface the module has the following low speed pins for control and status:

3.4.1 ModSelL

The ModSelL is an input pin. When held low by the host, the module responds to 2-wire serial communication commands and LED will be solid Green. The ModSelL allows the use of multiple QSFP+ modules on a single 2-wire interface bus. When the ModSelL is “High”, the module will not respond to or acknowledge any 2-wire interface communication from the host and the LED will be toggling. ModSelL signal input node must be biased to the “High” state in the module.

3.4.2 ResetL

The ResetL pin must be pulled to Vcc in the QSFP+ module. A low level on the ResetL pin for longer than the minimum pulse length (t_Reset_init) initiates a complete module reset, returning all user module settings to their default state. Module Reset Assert Time (t_init) starts on the rising edge after the low level on the ResetL pin is released. During the execution of a reset (t_init) the host should disregard all status bits until the module indicates a completion of the reset interrupt. The module indicates this by asserting “low” an IntL signal with the Data_Not_Ready bit negated. Note that on power up (including hot insertion) the module should post this completion of reset interrupt without requiring a reset.

For testing purpose ResetL pin state can be read via register address 145 of memory page 2. The data read from address 145 will reflect the hardware signal level (1 if ResetL is high and 0xFF if low).
Note that when ResetL is asserted the module resets the I2C slave and will not respond to I2C master communication.

3.4.3 LPMode

The Module should be in low power mode if the LPMode pin is in the high state, or if the Power_override bit is in the high state and the Power_set bit is also high. The module should be in high power mode if the LPMode pin is in the low state, or the Power_override bit is high and the Power_set bit is low. Note that the default state for the Power_override bit is low.
In low power mode module will stop all power dissipation and LED changes color to orange, in high power mode the power dissipation will be set to value inside register 98 and LED will be green.
For testing purpose LPMode pin state can be read via register address 146 of memory page 2. The data read from address 146 will reflect the hardware signal level (1 if LPMode is high and 0 if low).

3.4.4 ModPrsL

ModPrsL is pulled up to Vcc_Host on the host board and grounded in the module. The ModPrsL is asserted “Low” when inserted and deasserted “High” when the module is physically absent from the host connector.

3.4.5 IntL

IntL is an output pin, when “Low”, it indicates a possible module operational fault or a status critical to the host system. The host identifies the source of the interrupt using the 2-wire serial interface. The IntL pin is an open collector output and must be pulled to host supply voltage on the host board.
For testing purpose IntL can be set via register write to address 147 of memory page 2, by writing 1 to this address the IntL hardware signal will be set high, writing 0 to address 147 will reset IntL to low.

3.5 QSFP Memory Map

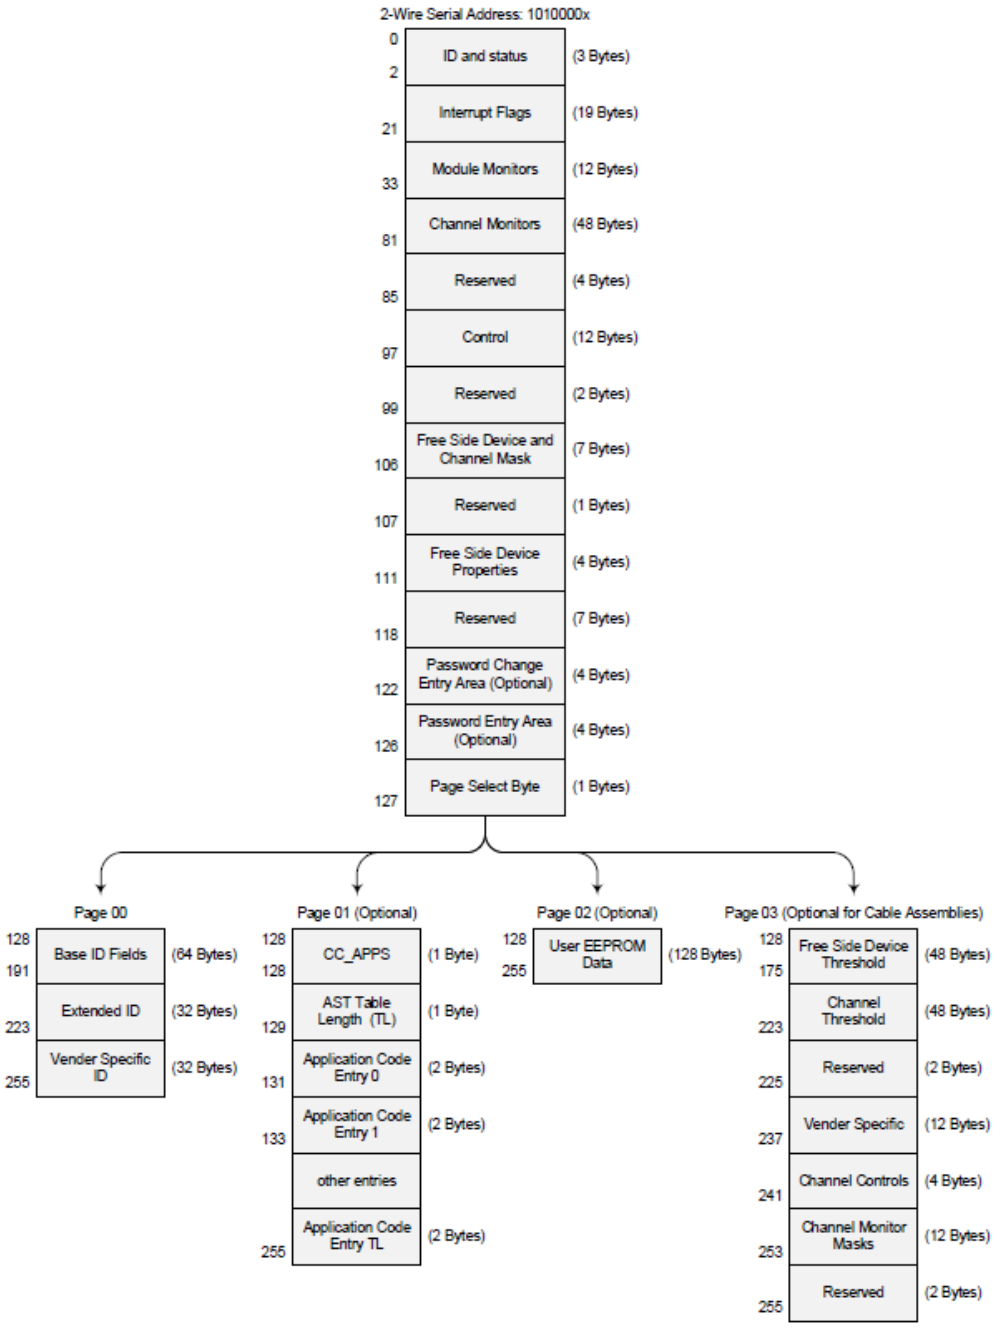

Figure 6: QSFP+ Memory Map

This section defines the Memory Map for QSFP transceiver used for serial ID, digital monitoring and certain control functions. The interface is mandatory for all QSFP devices.
The structure of the memory is shown in Figure 5. The memory space is arranged into a lower, single page, address space of 128 bytes and multiple upper address space pages. This structure permits timely access to addresses in the lower page, e.g. Interrupt Flags and Monitors. Less time critical entries, e.g. serial ID information and threshold settings are available with the Page Select function.

The structure also provides address expansion by adding additional upper pages as needed. For example, in Figure 5 upper pages 01 and 02 are optional. Upper page 01 allows implementation of Application Select Table, and upper page 02 provides user read/write space. The lower page and upper pages 00 and 03 are always implemented.

3.5.1 Low Memory Map

The lower 128 bytes of the 2-wire serial bus address space, see Table below, is used to access a variety of measurements and diagnostic functions, a set of control functions, and a means to select which of the various upper memory map pages are accessed on subsequent reads. This portion of the address space is always directly addressable and thus is chosen for monitoring and control functions that may need to be repeatedly accessed. The definition of Identifier field is the same as page 00h Byte 128.

Byte Address	Description	Type
0	Identifier (1 Byte)	Read-Only
1-2	Status (2 Bytes)	Read-Only
3-21	Interrupt Flags (19 Bytes)	Read-Only
22-33	Module Monitors (12 Bytes)	Read-Only
34-81	Channel Monitors (48 Bytes)	Read-Only
82-85	Reserved (4 Bytes)	Read-Only
86-97	Control (12 Bytes)	Read/Write
98-99	Reserved (2 Bytes)	Read/Write
100-106	Module and Channel Masks (7 Bytes)	Read/Write
107-118	Reserved (12 Bytes)	Read/Write
119-122	Password Change Entry Area (optional) (4 Bytes)	Read/Write
123-126	Password Entry Area (optional) (4 Bytes)	Read/Write
127	Page Select Byte	Read/Write

3.6 ML4002-28-8W Specific Functions

3.6.1 Temperature Monitor

The **ML4002-28-8W** has an internal temperature sensor in order to continuously monitor the module temperature. The temperature sensor readings are present in low-memory registers 22-23 as specified by QSFP SFF. Internally measured Module temperature are represented as a 16-bit signed two's complement value in increments of 1/256 degrees Celsius, yielding a total range of –128 to +127 °C that is considered valid between –40 and +125 °C. Temperature accuracy is vendor specific but must be better than ±3 degrees Celsius over specified operating temperature and voltage. Please see vendor specification for details on location of temperature sensor.

Address	Bit	Name	Description	Type
22	All	Temperature MSB	Internally measured module temperature	RO
23	All	Temperature LSB	Internally measured module temperature	

3.6.2 Insertion Counter

The Insertion counter contains the number of times the module was plugged in a host. The insertion counter is incremented every time the module goes in initializing sequence, as it is nonvolatile it is always saved, and can be read anytime from registers 141 and 142 on memory page 2.

Address	Bit	Name	Description	Type
141	All	Insertion Counter MSB		RO
142	All	Insertion Counter LSB	LSB unit = 1 insertion	

3.6.3 Programmable Power Dissipation and Thermal Emulation

Register 98 is used for power spots control over I2C. It is an 8-bit data wide register.
The consumed power changes accordingly when the value in this register is changed (only when in high power mode). In Low power mode the module automatically turns off the power spots.
The values written in this register are permanently stored. The PWM can also be used for module thermal emulation.
The module contains 3 thermal spots positioned where the optical transceivers usually are in an optical module that are heated relative to the power control register value. The first spot is 3.37 W and the second is 2.32 W. These can be turned either on or off allowing a power consumption of 0 or 100% of the overall power. A third spot of 1.97 W can be controlled using PWM, thus allowing a power consumption that covers all the range from 0 to 1.97 W with around 31 mW precision.

Address	Bit	Name	Description	Type
98	0:5	1.97 W PWM Controller	Power consumption varies linearly with the value set Bits [5:0] val = 0 to 63 corresponds to 0 to 1.97 W power consumption	RW
	6	2.32 W Static Power Spot Controller	Bit 6 = 0 : Adds 0 W to the power consumption Bit 6 = 1 : Adds 2.32 W to the power consumption	
	7	3.37 W Static Power Spot Controller	Bit 7 = 0 : Adds 0 W to the power consumption Bit 7 = 1 : Adds 3.37 W to the power consumption	

3.6.4 Cut-Off Temperature

To avoid overheating the module, a Cut-Off Temperature is pre-defined.
The module is continuously monitoring the temperature and checking its value against the Cut-Off temperature. Once the module temperature reaches the cut-off temperature, the power spots will automatically turn off in order to prevent overheating. Once the temperature is 5 degrees below cut-off value, the power goes back to its previous value.
The Cut-Off temperature for the **ML4002-28-8W** is 80 ˚C.

3.6.5 Power Staging

In order to limit the module inrush current, the firmware drives PWM transitionally through multiple stages; each stage enables 5% of the overall power, a programmable delay separates each stage from the next.
The user can set the programmable delay from register 143 and 144 from page 02.
After the delay value is changed, a module reset is required so that the new delay value becomes effective.

Address	Bit	Name	Description	Type
143	All	Programmable delay MSB		RW
144	All	Programmable delay LSB	LSB unit = 1 µs	

3.6.6 Voltage Sense

A voltage sense circuit is available in the **ML4002-28-8W** that allows to measure the internal module supplied voltage Vcc. Measured values range from 0 to 6528 mV with a 0.1 mV precision.

Address	Bit	Name	Description	Type
26	All	Supply Voltage MSB		RO
27	All	Supply Voltage LSB	LSB unit = 0.1 mV	

3.6.7 Low Speed Control Signals Pin State

Name	Address	Bit Number	Description	Type
ResetL	145 (Page 02)	0	1 if ResetL is High, 0xFF if ResetL is Low	RO
LPMode	146 (Page 02)	0	1 if LPMode is High, 0 if LPMode is Low	
IntL	147 (Page 02)	0	1 to set IntL to High, 0 to set IntL to Low	RW

4 QSFP28 Pin Allocation

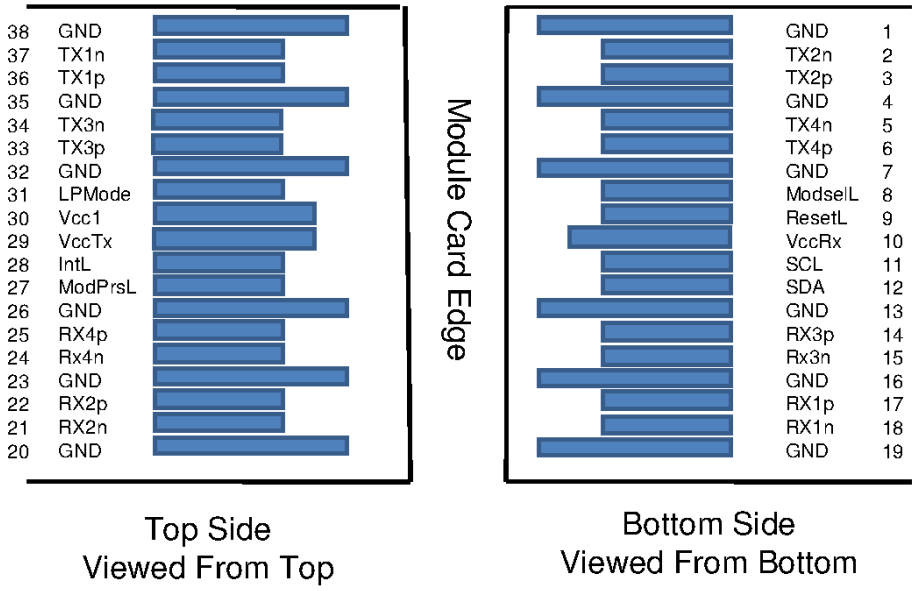

Figure 7: QSFP+ Module Pin Map

Pin#	Pin name	Logic	Description
1	GND		Power Ground
2	Tx2n	CML-I	Transmitter Inverted Data Input
3	Tx2p	CML-I	Transmitter Non-Inverted Data Input
4	GND		Power Ground
5	Tx4n	CML-I	Transmitter Inverted Data Input
6	Tx4p	CML-I	Transmitter Non-Inverted Data Input
7	GND		Power Ground
8	ModSelL	LVTTL I	Module Select
9	ResetL	LVTTL I	Module Reset
10	Vcc Rx		+3.3 V Power supply receiver
11	SCL	LVCMOS – I/O	2-wire serial interface clock
12	SDA	LVCMOS – I/O	2-wire serial interface data
13	GND		Power Ground
14	Rx3p	CML-O	Receiver Non-Inverted Data Output
15	Rx3n	CML-O	Receiver Inverted Data Output
16	GND		Power Ground
17	Rx1p	CML-O	Receiver Non-Inverted Data Output
18	Rx1n	CML-O	Receiver Inverted Data Output
19	GND		Power Ground
20	GND		Power Ground

21	Rx2n	CML-O	Receiver Inverted Data Output
22	Rx2p	CML-O	Receiver Non-Inverted Data Output
23	GND		Power Ground
24	Rx4n	CML-O	Receiver Inverted Data Output
25	Rx4p	CML-O	Receiver Non-Inverted Data Output
26	GND		Power Ground
27	ModPrsl	LVTTL O	Module Present
28	IntL	LVTTL O	Interrupt
29	Vcc Tx		+3.3 V Power supply transmitter
30	Vcc1		+3.3 V Power Supply
31	LPMode	LVTTL I	Low Power Mode
32	GND		Power Ground
33	Tx3p	CML-I	Transmitter Non-Inverted Data Input
34	Tx3n	CML-I	Transmitter Inverted Data Input
35	GND		Power Ground
36	Tx1p	CML-I	Transmitter Non-Inverted Data Input
37	Tx1n	CML-I	Transmitter Inverted Data Input
38	GND		Power Ground

5 High Speed Signals

High speed signals are electrically lopped back from TX side to RX side of the module, all differential TX pairs are connected to the corresponding RX pairs, and the signals are AC coupled as specified by QSFP MSA HW specs.
The Passive traces connecting TX to RX pairs are designed to support a data rate up to 28 Gbps.

5.1 Trace Length

Net NAme	Etch Lengths (mil)
RX1N	821.84
RX1P	864.26
RX2N	819.72
RX2P	862.14
RX3N	508.64
RX3P	551.07
RX4N	503.34
RX4P	545.77
TX1N	920.94
TX1P	878.52
TX2N	964.92
TX2P	922.5
TX3N	613.52
TX3P	571.09
TX4N	578.28
TX4P	535.86

5.2 IL and RL Graphs

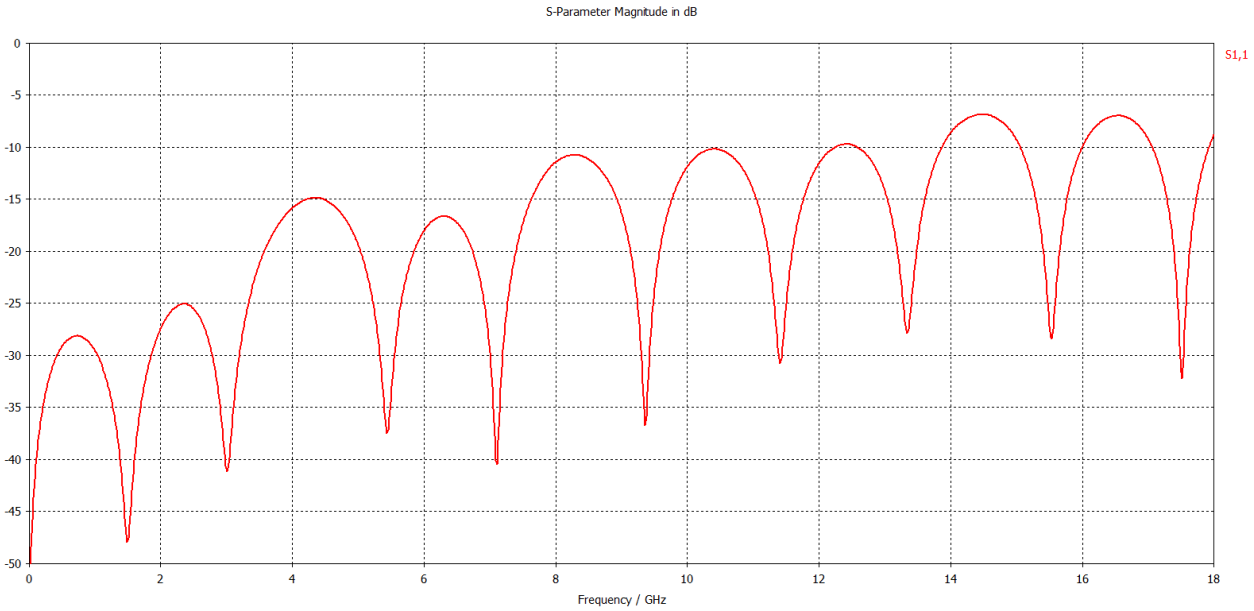

Figure 8: S11

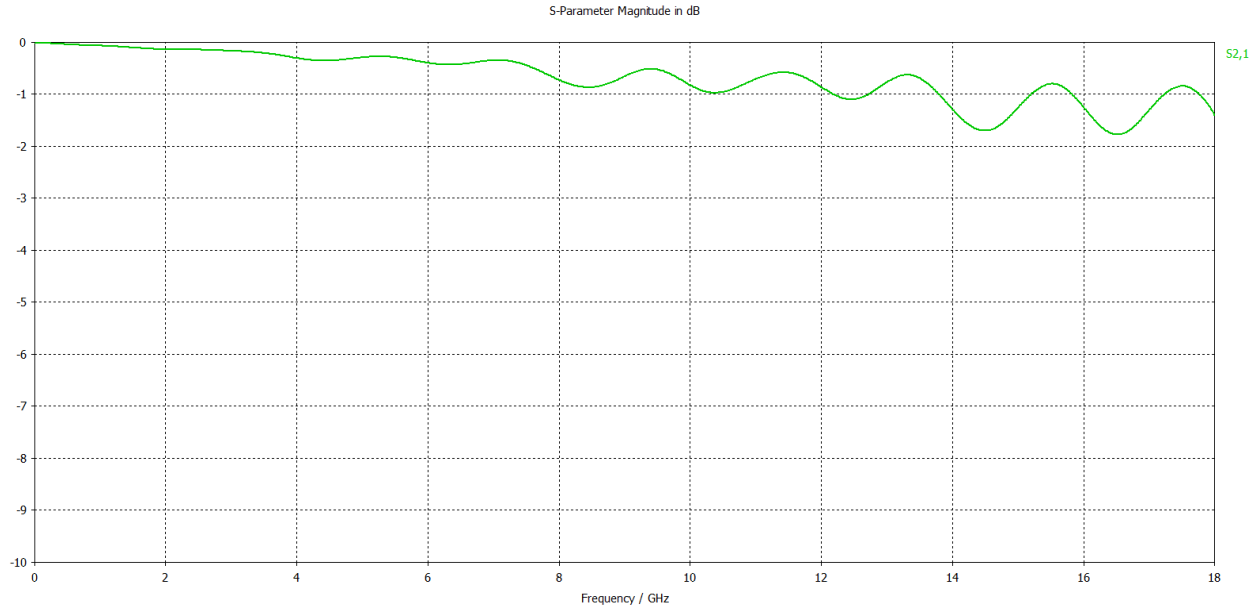

Figure 9: S21

Revision History

Revision number	Date	Description
0.1	02/06/2016	▪ Preliminary
0.2	13/06/2016	▪ Fixed operating conditions max power
0.3	15/07/2019	▪ Fixed max power dissipation value
0.4	9/1/2021	▪ Update Format ▪ Update power values in section 3.6.3 ▪ Update max power value