Document Number EDCS-540123
Revision Revision 1.3

QSGMII Specification

The Quad Serial Gigabit Media Independent Interface (QSGMII) is designed to satisfy the following requirements:

- Convey 4 ports of network data and port speed between a 10/100/1000 PHY and a MAC with significantly less signal pins than required for GMII & SGMII.
- Operate in both half and full duplex and at all port speeds.
- This implementation can be extended to other port to channel ratios. However, this is outside the scope of this document.

Change History

Revision	Date	Description	Author
1.3	July 20,2009	Add Low Power Idle (LPI) signaling for Energy Efficient Ethernet (IEEE 802.3az)	Hugh Barrass
1.2	September 7,2007	Updated Bit[13] and Bit[0] on Table3 and added a statement of possible removal of the first byte of frame in 10/ 100Mbit/s operation to match SGMII spec.	Akin Koyuncuoglu
1.1	June 20, 2007	Reword Note1, Added a requirement to disable running disparity check at receiver.	Akin Koyuncuoglu
1.0	April 17, 2007	Updated Interconnect Loss Template-Figure 11 and Channel Loss Budget-Table 9. Updated Differential and common mode return loss parameters and differential voltage values in electrical section.	Akin Koyuncuoglu, Warren Meggitt
0.5	January 4, 2007	Updated PCS Receive for carrier_detect function. Updated /I/ Idle Code Group Selection for Transmission. Added a note for running disparity support by Framers. Updated Figure2, Figure4, Figure5,Figure6, Table1 and Table2	Akin Koyuncuoglu
0.4	December 4, 2006	Updated legal section in the end of the document	Akin Koyuncuoglu
0.3	September 1, 2006	Updated Electrical Specification	Akin Koyuncuoglu, Warren Meggitt
0.2	Nov. 3, 2005	Initial Release	Jeff Provost

Definitions

MII – Media Independent Interface: A digital interface that provides a 4-bit wide datapath between a 10/100 Mbit/s PHY and a MAC sublayer. Since MII is a subset of GMII, in this document, we will use the term “GMII” to cover all of the specification regarding the MII interface.

GMII – Gigabit Media Independent Interface: A digital interface that provides an 8-bit wide datapath between a 1000 Mbit/s PHY and a MAC sublayer. It also supports the 4-bit wide MII interface as defined in the IEEE 802.3z specification. In this document, the term “GMII” covers all 10/100/1000 Mbit/s interface operations.

SGMII – Serial Gigabit Media Independent Interface: A digital interface that provides a 1.25 Gbps serial dual-data-rate datapath between a 1000 Mbit/s PHY and a MAC sublayer. Refer to ENG-46158 or *ftp://ftp-eng.cisco.com/smii/smii.html* for details.

LPI – Low Power Idle: An alternative form of idle signaling that is used by the MAC to indicate that the PHY may enter a low power state and signal this change of state to the link partner; and is used by the PHY to signal to the MAC that the link partner has entered a low power state. The functions are defined by IEEE 802.3az in IEEE 802.3 clauses 22, 24, 25 (for 100Mb/s); 35, 36, 40, 70 (for 1Gb/s); 46, 48, 49, 55, 71, 72 (for 10Gb/s); and 78 (for overall descriptions).

Overview

QSGMII uses two data signals in each direction to convey frame data and link rate information between a multi-port 10/100/1000 PHY and Ethernet MAC. The data signals operate at 5.0 Gbps using CDR technology to recover the clock at the MAC and PHY interfaces. Due to the high speed of operation, each of these signal pairs are realized as differential pairs thus optimizing signal integrity while minimizing system noise. When EEE is supported the clocking of the signals may be stopped during idle periods according to the status and control settings in the PHY.

Figure 1 compares the IEEE 802.3 PCS reference diagram before and after the QSGMII modification.

The IEEE 802.3 reference diagrams in Figure 2 shows where the modification to the receive PCS function occurs.

The transmit and receive data paths leverage the 1000BASE-X PCS defined in the IEEE 802.3z specification (clause 36). Four ports of traditional GMII data signals (TXD/RXD), data valid signals (TX_EN/RX_DV), and error signals (TX_ER/RX_ER) are muxed, encoded, and serialized. Carrier Sense (CRS) is derived/inferred from RX_DV, and collision (COL) is logically derived in the MAC when RX_DV and TX_EN are simultaneously asserted. There is a small block in the PHY transmit path to suppress TX_ER in full duplex mode when TX_EN is not asserted. Since four 1.25Gbps SGMII ports are interleaved onto a single link, the data-rate becomes 5.0 Gbps.

Figure 3 illustrates the resulting bit times on the QSGMII 5.0 Gbps link.

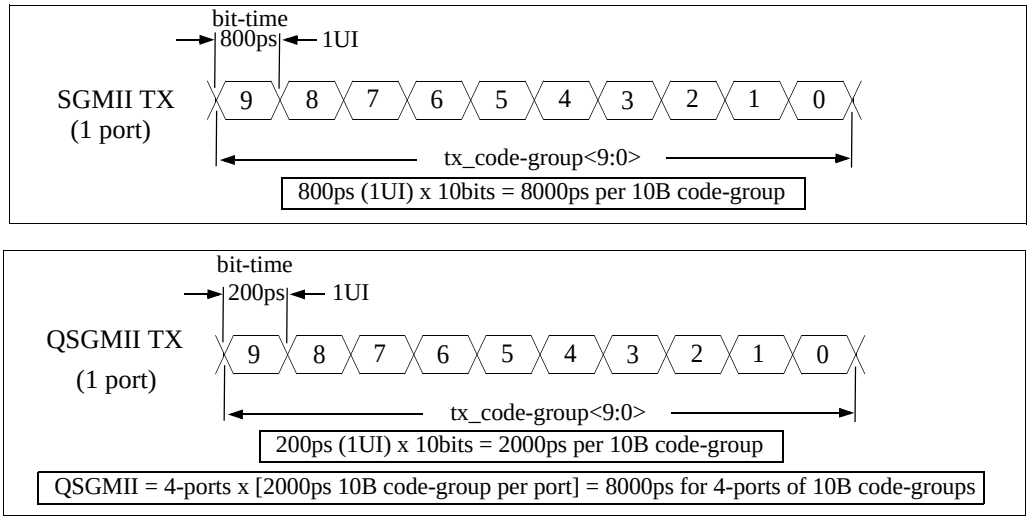

Figure 3 10B Encoded Data on the 5.0 Gbps QSGMII Link vs. SGMII 1.25 Gbps

In order to determine the port number based time slots the port 0 transmit side incorporates a “K28.5” swapper function that modifies the IDLE /I/ and Configuration /C/ ordered_sets by replacing /K28.5/ with /K28.1/ every time /K28.5/ occurs as shown in table 1. Note that the swapper operates on the GMII octets (8 + control), not on the 10B code-group directly. The data will appear on the QSGMII link in the order: port 0 first, then port 1,then port 2, and lastly port 3. Port 0 data then appears on the link again, and so on.

802.3z Section 36.2.4.12 explains the rules for running disparity by sending out one of the two IDLE /I/ ordered_sets whenever the GMII is idle. However, since 8B/10B encoder is detached

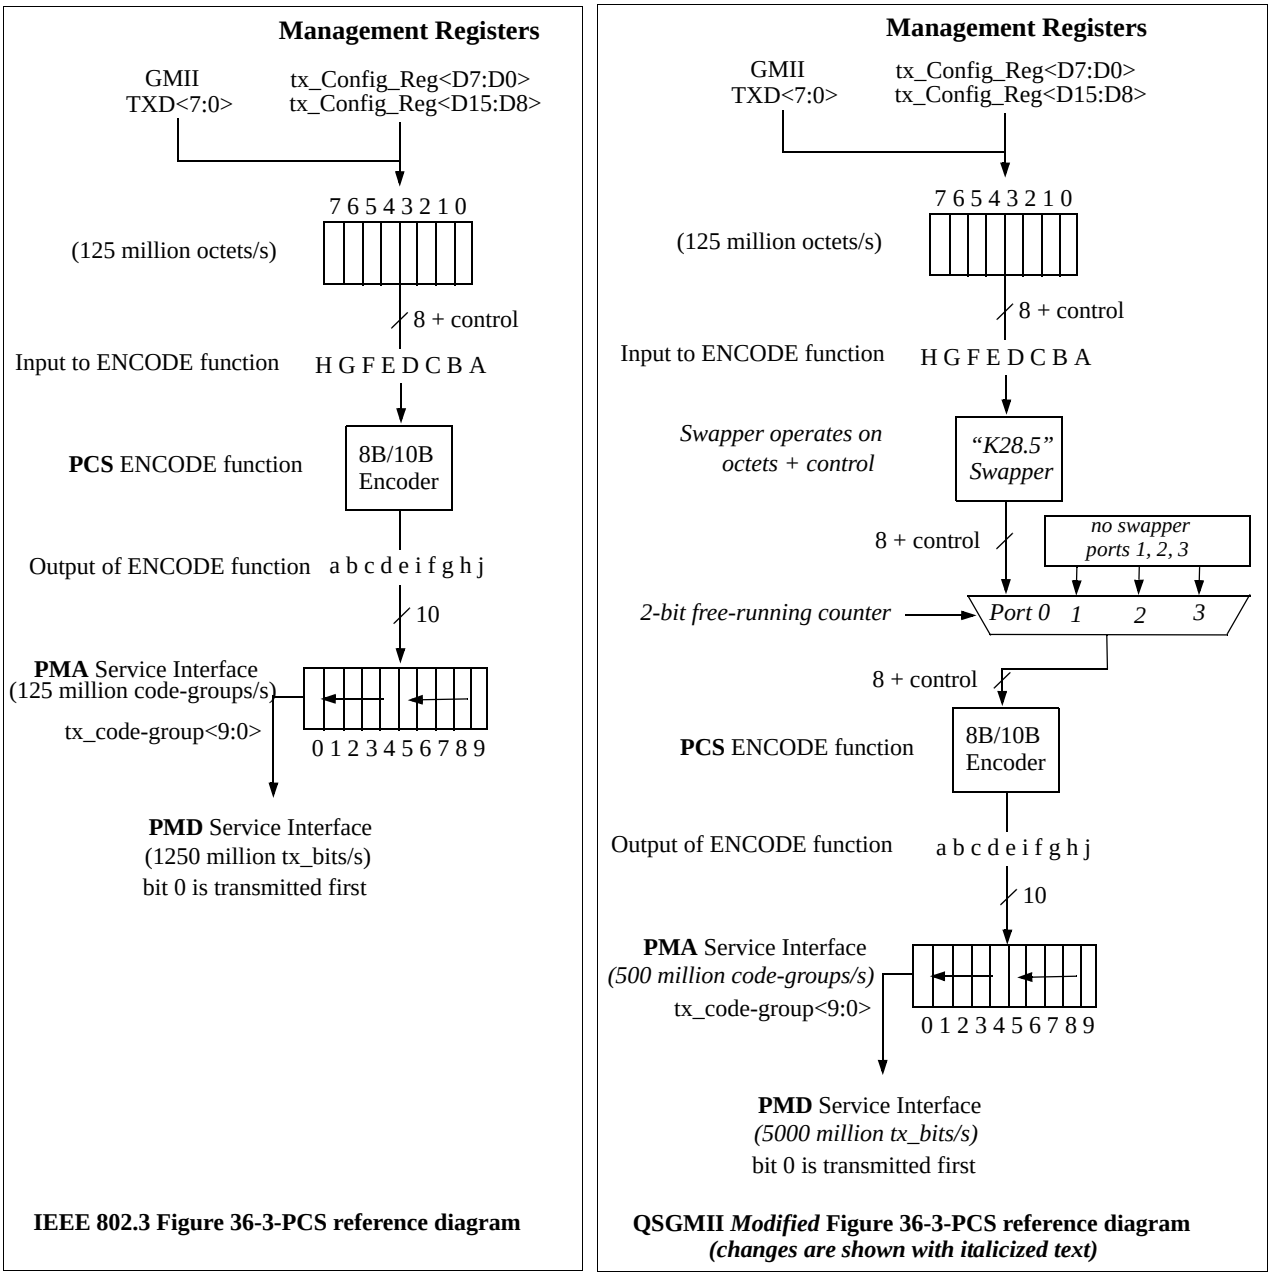

Figure 1 "Standard" and Modified Transmit Path Diagrams per IEEE 802.3 PCS/PMA

from PCS Transmit Function, it is no longer feasible to use /I1/ and /I2/ ordered_sets to force the disparity. Therefore, the transmitter may be simplified to only generate /I1/ ordered_sets. This change requires more functionality from the framer as documented in Note1.

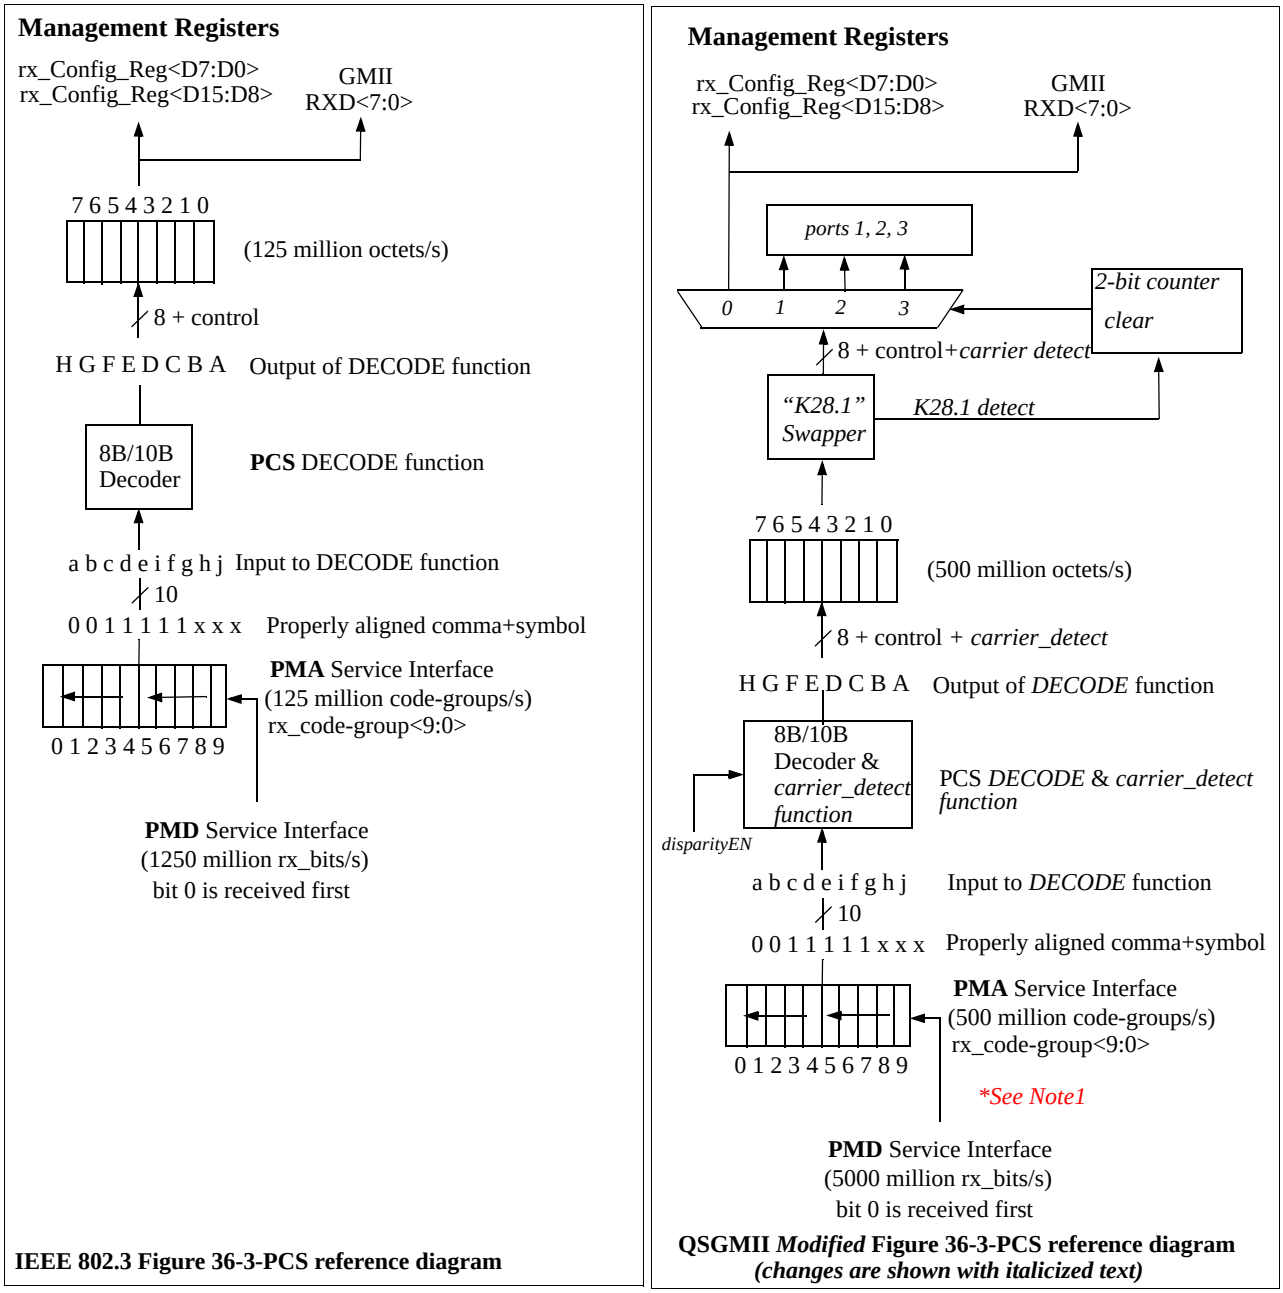

Figure 2 "Standard" and Modified Receive Path Diagram per IEEE 802.3 PCS/PMA

Note1: QSGMII Receivers should not rely upon receipt of /I2/ ordered_sets for proper operation.

Code	Ordered_Set	Number of Code Groups	Port 0 "pre-swapper" Encoding	Port 0 "post-swapper" Encoding
/**C**/	**Configuration**		Alternating /C1/ and /C2/	Alternating /C1/ and /C2/
/C1/	Configuration 1	4	/K28.5/D21.5/Config_Reg	/K28.1/D21.5/Config_Reg
/C2/	Configuration 2	4	/K28.5/D2.2/Config_Reg	/K28.1/D2.2/Config_Reg

table 1 Port 0 "K28.5" Swapper Definition

Code	Ordered_Set	Number of Code Groups	Port 0 ”pre-swapper” Encoding	Port 0 “post-swapper” Encoding
/I/	**IDLE**		Correcting /I1/	Correcting /I1/
/I1/	IDLE 1	2	/K28.5/D5.6/	/K28.1/D5.6/

table 1 Port 0 “K28.5” Swapper Definition

The receive 10B code-groups pass through a “K28.1” swapper that undoes the modification of the IDLE /I/ and Configuration /C/ ordered_sets by replacing /K28.1/ with /K28.5/ for every occurrence. The K28.1 swapper also clears the de-mux (sets the counter to 2’b00) in order to determine the port 0 data according to table 2. Note that the swapper operates on the GMII octets (8 + control + carrier_detect), not on the 10B code-group directly.

On the receive side, carrier_detect function is done on 10B code-groups and this is shown in Figure 2. Please note that, there is a new carrier_detect function that needs to operate on K28.1 for Port0. K28.1 can be locally converted to K28.5 to generate the carrier_detect function. Please refer to 802.3z Section 36.2.5.1.4 for the carrier_detect function that operates on K28.5.

Due to the nature of QSGMII, bit errors on the link may cause a running disparity error to propogate across ports. A software register bit that would enable/disable running disparity checking at the receiver is required. Disabling running disparity checking at the receiver prevents error propogation to other ports. It is not necessary to disable ALL disparity checking in the decoder to prevent error propogation to other ports. It is only necessary to disable the disparity checks that rely on the running disparity value from the previous symbol. Note that code violations due to invalid code-words (and current symbol running disparity errors) should continue to be detected regardless of the state of running disparity checking. Please also note that, 802.3z Section 36.2.4.6 and DECODE([/x/]) function in Section 36.2.5.1.4 will be affected by this requirement.

Code	Ordered_Set	Number of Code Groups	Port 0 ”pre-swapper” Encoding	Port 0 “post-swapper” Encoding
/**C**/	**Configuration**		Alternating /C1/ and /C2/	Alternating /C1/ and /C2/
/C1/	Configuration 1	4	/K28.1/D21.5/Config_Reg	/K28.5/D21.5/Config_Reg
/C2/	Configuration 2	4	/K28.1/D2.2/Config_Reg	/K28.5/D2.2/Config_Reg
/**I**/	**IDLE**		Correcting /I1/	Correcting /I1/
/I1/	IDLE 1	2	/K28.1/D5.6/	/K28.5/D5.6/

table 2 Port 0 “K28.1” Swapper Definition

Ports 1-3 are unchanged from SGMII specification.

Control information, as specified in table 3, is transferred from the PHY to the MAC to signal the change of the link status. This is achieved by using the Auto-Negotiation functionality defined in Clause 37 of the IEEE Specification 802.3z. Instead of the ability advertisement, the PHY sends the control information via its tx_config_Reg[15:0] as specified in table 3 whenever the link status changes. Upon receiving control information, the MAC acknowledges the update of link status by asserting bit 14 of its tx_config_reg{15:0] as specified in table 3.

The link_timer inside the Auto-Negotiation block has been changed from 10 msec to 1.6 msec to ensure a prompt update of the link status.

Bit Number	tx_config_Reg[15:0] sent from the PHY to the MAC	tx_config_Reg[15:0] sent from the MAC to the PHY
15	Link: 1 = link up, 0 = link down	0: Reserved for future use
14	Reserved for Auto-Negotiation acknowledge as specified in 802.3z	1
13	0: Reserved for future use	0: Reserved for future use
12	Duplex mode: 1 = full duplex, 0 = half duplex	0: Reserved for future use

table 3 Definition of Control Information passed between links via tx_config_Reg[15:0]

Bit Number	tx_config_Reg[15:0] sent from the PHY to the MAC	tx_config_Reg[15:0] sent from the MAC to the PHY
11:10	Speed: Bit 11, 10: 1 1 = Reserved 1 0 = 1000 Mbps: 1000BASE-TX, 1000BASE-X 0 1 = 100 Mbps: 100BASE-TX, 100BASE-FX 0 0 = 10 Mbps: 10BASET, 10BASE2, 10BASE5	0: Reserved for future use
9	EEE capability:1= supported, 0 = not supported	0: Reserved for future use
9:1	EEE clock stop capability: 1= supported, 0 = not supported	EEE clock stop enable: 1= enabled, 0 = not enabled
7:1	0: Reserved for future use	0: Reserved for future use
0	1	1

table 3 Definition of Control Information passed between links via tx_config_Reg[15:0]

QSGMII's 5.0 Gbps transfer rate is excessive for PHYs operating at 10 or 100 Mbit/s. When these situations occur, the interface "elongates" the frame by replicating each frame byte 10 times for 100 Mbit/s and 100 times for 10 Mbit/s. This frame elongation takes place "above" the 802.3z PCS layer, thus the start frame delimiter only appears once per frame. The 802.3z PCS layer may remove the first byte of the "elongated" frame.

Implementation Specification

This section discusses how this QSGMII interface shall be implemented by incorporating and modifying the PCS layer of the IEEE Specification 802.3z. Figure 4 illustrates the connections between the MAC and the PHY in a system utilizing QSGMII, as well as the overall scheme of muxing, demuxing 4 ports into a single QSGMII channel.

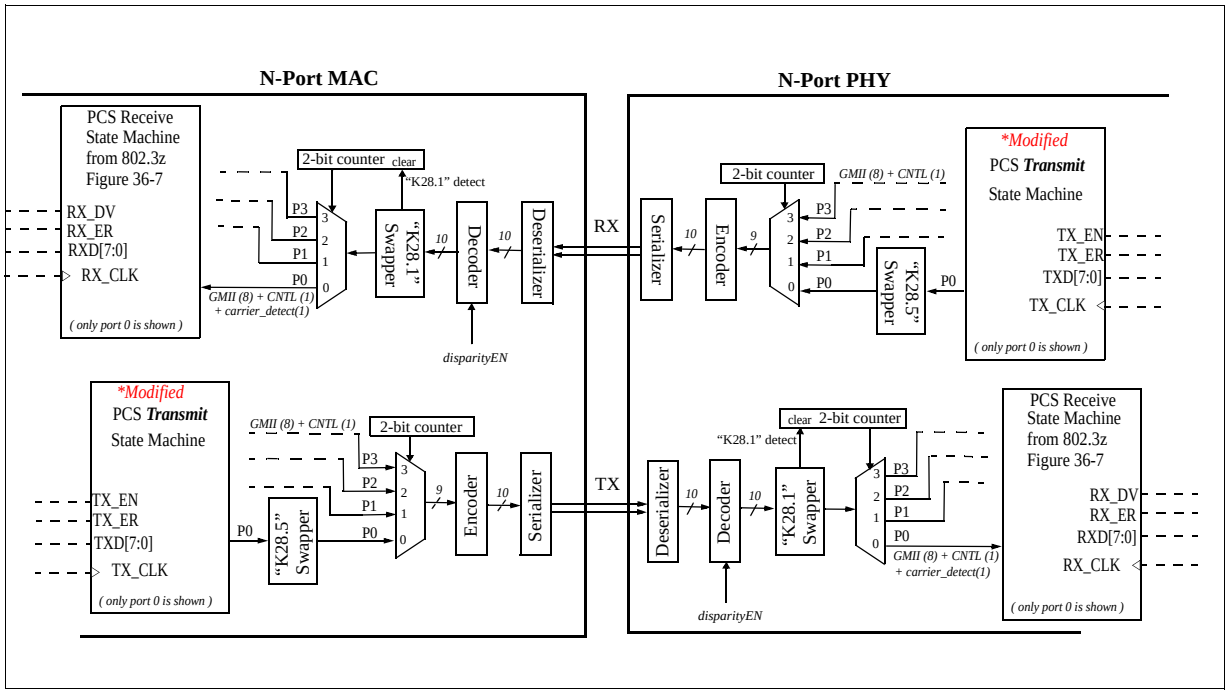

Figure 4 QSGMII Connectivity

Signal Mapping at the PHY side

Figure 5 shows the PHY functional block diagram. It illustrates that the PCS layer shall be modified and incorporated at the PHY side in the QSGMII interface.

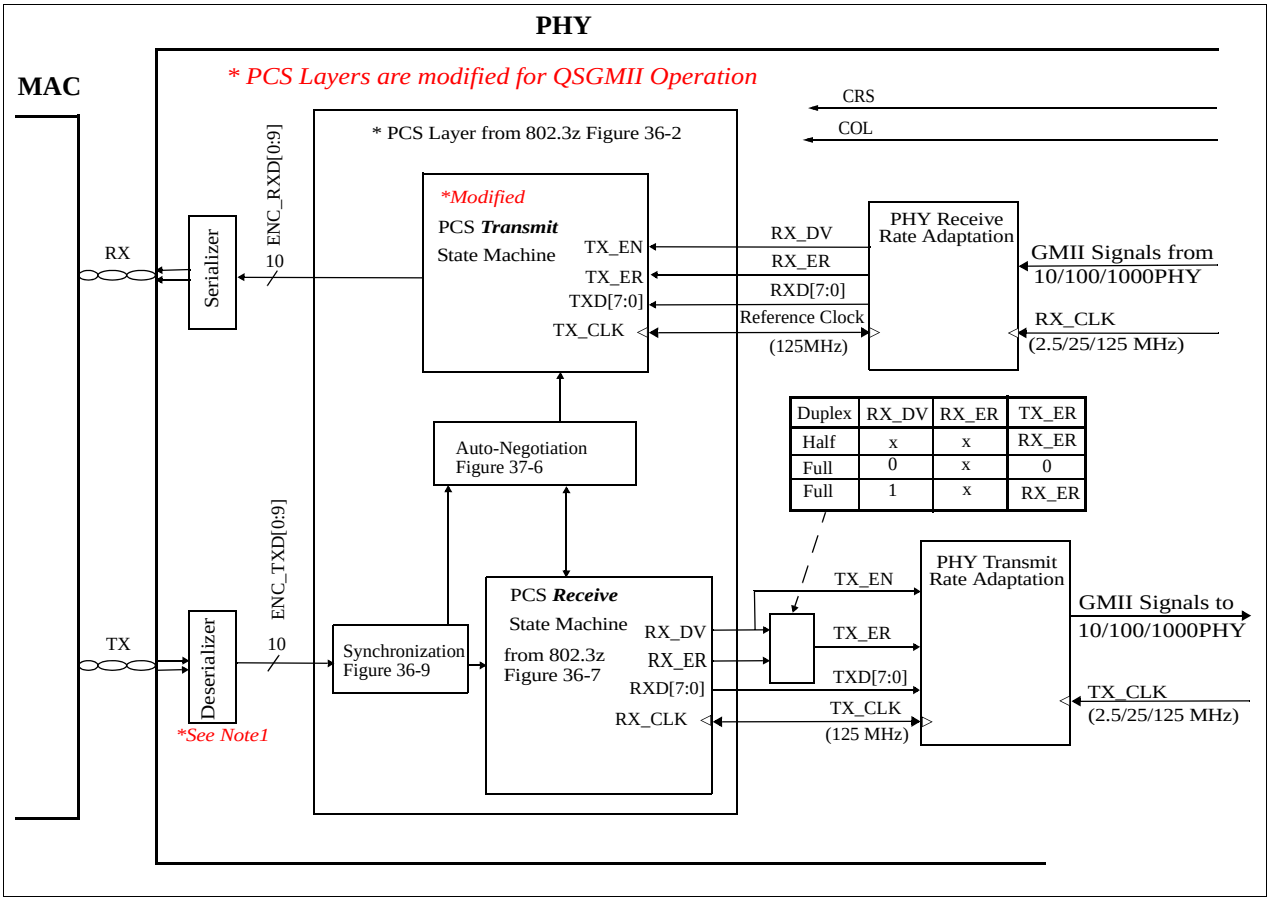

Figure 5 PHY Functional Block For One Port

At the receive side, GMII signals come in at 10/100/1000 Mbit/s clocked at 2.5/25/125 MHz. The PHY passes these signals through the PHY Receive Rate Adaptation to output the 8-bit data RXD[7:0] in Reference Clock domain. Please note that since 4 ports are multiplexed, recovered RX_CLK can not be used to output the 8-bit data RXD[7:0]. RXD is sent to the PCS Transmit State Machine to generate an encoded 10-bit segment ENC_RXD[0:9]. The PHY serializes ENC_RXD[0:9] to create RX and sends it to the MAC at 5.0 Gbps.

802.3z Section 36.2.4.12 explains the rules for running disparity by sending out one of the two IDLE /I/ ordered_sets whenever the GMII is idle. However, since 8B/10B encoder is detached from PCS Transmit Function, it is no longer feasible to use /I1/ and /I2/ ordered_sets to force the disparity. Therefore, the transmitter may be simplified to only generate /I1/ ordered_sets. This change requires more functionality from the framer as documented in Note1.

Due to the nature of QSGMII, bit errors on the link may cause a running disparity error to propogate across ports. A software register bit that would enable/disable running disparity checking at the receiver is required. Disabling running disparity checking at the receiver prevents error propogation to other ports. It is not necessary to disable ALL disparity checking in the decoder to prevent error propogation to other ports. It is only necessary to disable the disparity checks that rely on the running disparity value from the previous symbol. Note that

code violations due to invalid code-words (and current symbol running disparity errors) should continue to be detected regardless of the state of running disparity checking. Please also note that, 802.3z Section 36.2.4.6 and DECODE([/x/]) function in Section 36.2.5.1.4 will be affected by this requirement.

At the transmit side, the PHY deserializes TX to recover encoded ENC_TXD[0:9]. The PHY passes ENC_TXD[0:9] through the PCS Receive State Machine to recover the GMII signals. In the mean time, Synchronization block checks ENC_TXD[0:9] to determine the synchronization status between links, and to realign if it detects the loss of synchronization. The decoded GMII signals have to pass the PHY Transmit Rate Adaptation block to output data segments according to the PHY port speed.

To make the PCS layer from 802.3z work properly, the PHY must provide a frame beginning with at least two preamble symbols followed by a SFD symbol. To be more specific, at the beginning of a frame, RXD[7:0] in Figure 5 shall be {8'h55, 8'h55, (8'h55...), 8'h55} followed by valid frame data.

Some legacy end points will drop frames when RX_ER asserts during the first clock after a frame ends. The receive PCS state machine generates this signalling at the end of certain frames. To avoid this problem, there is a small block in the PHY transmit path to suppress TX_ER in full duplex mode when RX_DV (from the PHY receive PCS state machine) is not asserted.

Low Power Idle operation requires that the LPI symbols are passed across the QSGMII in the manner defined in 802.3 Clause 36. If the EEE clock stop capability is asserted (by the PHY) then the MAC may stop the signals on the QSGMII may be stopped in the same manner as 1000BASE-KX (defined in 802.3 clauses 36 and 70), if and only if all 4 transmit channels are in transmitting LPI. If the EEE clock stop enable is asserted (by the MAC) then the PHY may stop the signals on the QSGMII in the same manner. In cases where the MAC or PHY requires constant signaling on the QSGMII (i.e. EEE clock stop enable or capable = 0), it is assumed that the LPI signal latency is the same as the datapath signal latency and therefore the PHY timing constraints remain unchanged. In cases where the MAC allows the QSGMII clocking to be stopped during Low Power Idle, the system must negotiate an extra 20 uS of wake time via LLDP (see 802.3 clause 78) in the receive direction on all 4 channels to allow time for the QSGMII to wake and resynchronize after the PHY has woken. Similarly, if the MAC stops the QSGMII clocking in the transmit direction, then the system must wait for an additional 20uS between sending normal IDLE on the first channel to wake and sending data on any channel to allow extra time for the QSGMII to wake and resynchronize in the transmit direction.

Signal Mapping at the MAC Side

Figure 6 shows the MAC functional block diagram. It illustrates that the PCS layer shall be modified and incorporated at the MAC side in the QSGMII interface.

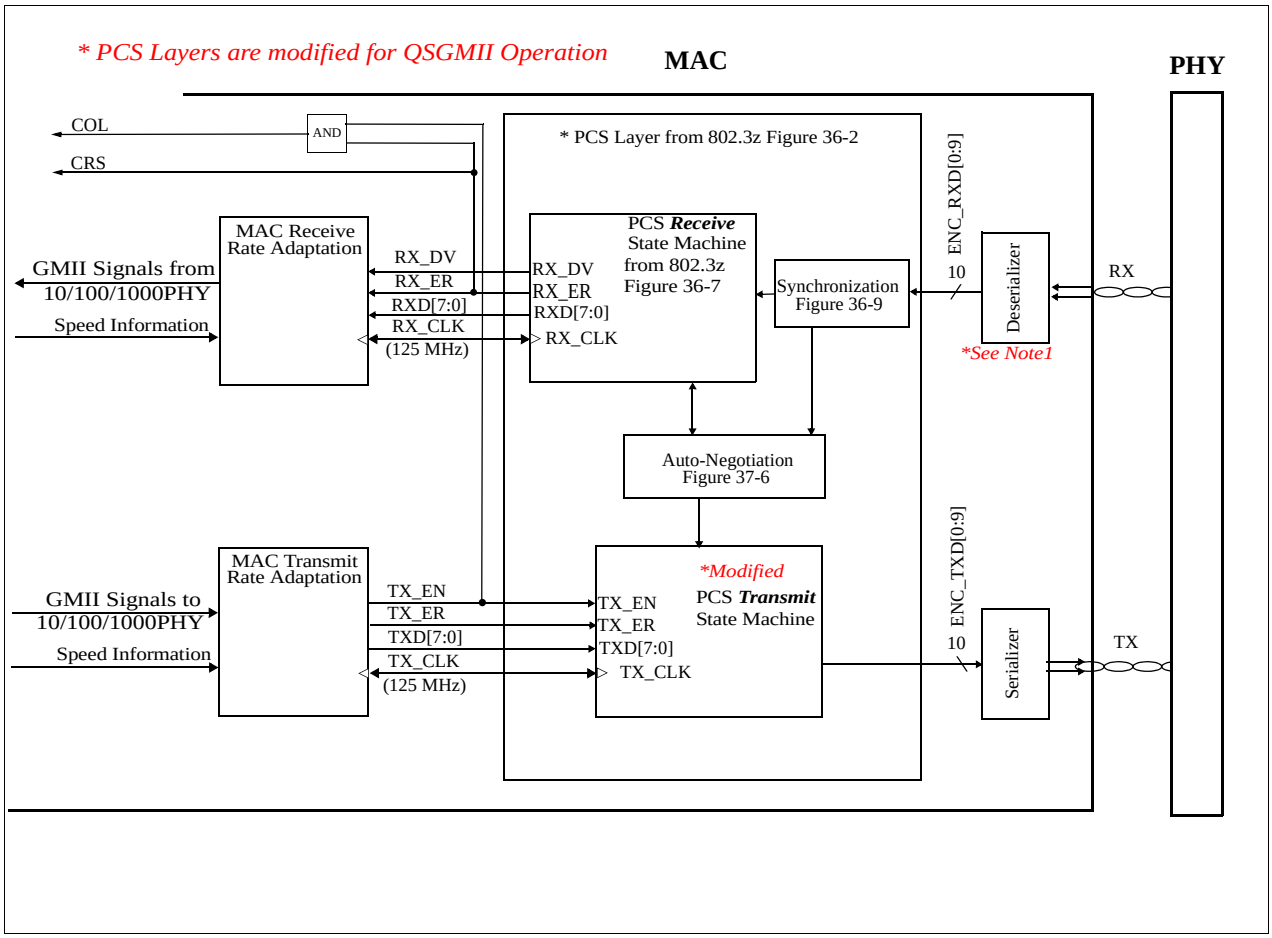

Figure 6 MAC Functional Block for One Port

At the receive side, the MAC deserializes RX to recover encoded ENC_RXD[0:9]. The MAC passes ENC_RXD[0:9] through the PCS Receive State Machine to recover the GMII signals. In the mean time, Synchronization block checks ENC_RXD[0:9] to determine the synchronization status between links, and to realign once it detects the loss of synchronization. The decoded GMII signals have to pass the MAC Receive Rate Adaptation block to output data segments according to the PHY port speed, passed from the PHY to MAC via Auto-Negotiation process.

At the transmit side, GMII signals come in at 10/100/1000 Mbit/s data clocked at 2.5/25/125 MHz. The MAC passes these signals through the MAC Transmit Rate Adaptation to output the 8-bit data TXD[7:0] in 125MHz clock domain. TXD is sent to the PCS Transmit State Machine to generate an encoded 10-bit segment ENC_TXD[0:9]. The MAC serializes ENC_TXD[0:9] to create TX and sends it to the PHY at 5.0 Gbps

Control Information Exchanged Between Links

As described in Overview, it is necessary for the PHY to pass control information to the MAC to notify the change of the link status. QSGMII interface uses Auto-Negotiation block to pass the control information via tx_config_Reg[15:0].

If the PHY detects the link change, it starts its Auto-Negotiation process, switching its Transmit block from “data” to “configuration” state and sending out the updated control information via tx_config_Reg[15:0]. The Receive block in the MAC receives and decodes control information, and starts the MAC’s Auto-Negotiation process. The Transmit block in the MAC acknowledges the update of link status via tx_config_Reg[15:0] with bit 14 asserted, as specified in table 3. Upon receiving the acknowledgement from the MAC, the PHY completes the auto-negotiation process and returns to the normal data process.

As specified in Overview, inside the QSGMII interface the Auto-Negotiation link_timer has been changed from 10 msec to 1.6 msec, ensuring a prompt update of the link status. The expected latency for the update of link is 3.4 msec (two link_timer time + an acknowledgement process).

Data Information Transferred Between Links

Below we briefly describe at receive side how GMII signals get transferred across from the PHY and recovered at the MAC by using the 8B/10B transmission code. The same method applies to the transmit side.

According to the assertion and deassertion of RX_DV, the PHY encodes the Start_of_Packet delimiter (SPD /S/) and the End_Of_Packet delimiter (EPD) to signal the beginning and end of each packet. The MAC recovers RX_DV signal by detecting these two delimiters.

The PHY encodes the Error_Propagation(/V/) ordered_set to indicate a data transmission error. The MAC asserts RX_ER signal whenever it detects this ordered_set.

CRS is not directly encoded and passed to the MAC. To regenerate CRS, the MAC shall uses signal RX_DV before it is being passed to the MAC Receive Rate Adaptation block as shown in Figure 6.

The MAC decodes ENC_RXD[0:9] to recover RXD[7:0].

Figure 7 illustrates how the MAC samples data in 100 Mbit/s mode. The GMII data in 100 Mbit/s mode get replicated ten times after passing through the PHY Receive Rate Adaptation to generate RXD[7:0]. The modified PCS Transmit State Machine encodes RXD[7:0] to create ENC_RXD[0:9]. As noted in the Overview, the SPD (/S/) only appears once per frame. SAMPLE_EN is a MAC *internal* signal to enable the MAC sampling of data starting at the first data segment (/S/) once every ten data segments in 100 Mbit/s mode.

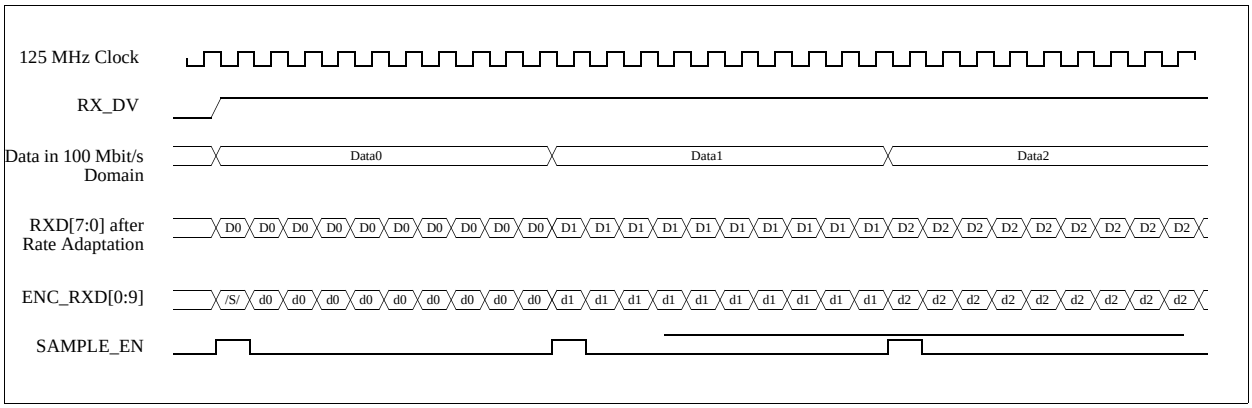

Figure 7 Data Sampling in 100 Mbit/s mode

Electrical Specification

CML, (Current Mode Logic) is by far the most common serdes IO standard in use today. The signal swing provided by the CML output is small, resulting in low power consumption. The driver and receiver are often self-terminated, eliminating external components and minimizing transmission line impedance discontinuity effects on timing and signal integrity.

The following section details the requirements for the high speed electrical interface that will operate at 5Gsym/s using NRZ coding (hence 1 bit per symbol at electrical level). Connections are point to point balanced differential pair with 100 Ohm nominal differential impedance and signalling is unidirectional. Clock and data are embedded hence CDR is required in the receiver. The link should operate with a BER of 10^{-15}. It supports both AC and DC coupled operation. However, DC coupling of PHY to MAC is required since it optimizes system cost, complexity, and signal integrity.

This section is based on the Optical Internetworking Forum's(OIF) Common Electrical I/O CEI-6G-SR Short Reach Standard IA#OIF-CEI-02.0 with some modifications listed in below sections.

Characteristic	Load Type 0(AC Coupling)	Load Type 1(DC Coupling)	Units
R_Zvtt	>1K	<30	Ohms
Nominal Vtt	Undefined	1.2	V

table 4 Definition of Load Types

Parameter	Symbol	Min	Typ	Max	Units	Notes
Baud Rate	T_Baud		5.000		Gsym/s	*a*
Output Differential Voltage (into floating Load Rload=100Ohm)	T_Vdiff	400		*900*	mVppd	*b*
Differential Resistance	T_Rd	80	100	120	Ohms	
Recommended Output Rise and Fall Times (20% to 80%)	T_tr,T_tf	30			ps	
Differential Output Return Loss (100Mhz to *2.5 GHz*)	T_SDD22			-8	dB	c

table 5 Transmitter Output Electrical Specification

Parameter	Symbol	Min	Typ	Max	Units	Notes
Differential Output Return Loss (*2.5 GHz* to 5 GHz)	T_SDD22				dB	[d]
Common Mode Return Loss (100Mhz to *2.5 GHz*)	T_SCC22			-6	dB	[e]
Transmitter Common Mode Noise	T_Ncm			5% of T-Vdiff	mVppd	
Output current into or out of driver pins when either SHORT to GND or each other				100	mA	[f]
Output Common Mode Voltage See Note [g], See Note [h], See Note [i]	T_Vcm	0.0 735		1.8 1135	V mV	[j] Load Type0 Load Type 1

table 5 Transmitter Output Electrical Specification

a. CEI-6G-SR is defined to operate between baud rates of 4.976 and 6.375 Gsym/s, However QSGMII will operate at 5 Gsym/s with a tolerance of +/-100ppm.
b. Absolute driver output voltage shall be between -0.1V and 1.9V with respect to local ground.See Figure 10 for details.
c. See Figure 9
d. See Figure 9
e. See Figure 9
f. +/- 100 mA
g. For both Load Types: R_Rdin=100 Ohms+/- 20 Ohms.For Vcm definition, see Figure 10.
h. For Load Type 1: R_ZVtt<30Ohms; Vtt is defined follows: Load Type 1: R_Vtt =1.2V +5%/-8%
i. DC Coupling compliance is Type 1. It is acceptable for a Transmitter to restrict the range of T_Vdiff in order to comply with the specified T_Vcm range. For a transmitter which supports multiple T_Vdiff levels, it is acceptable for a Transmitter to claim DC Compliance if it meets the T_Vcm ranges for at least one of its T_Vdiff setting as long as those settings that are compliant are indicated.
j. Load Type 0 with min T_Vdiff, AC Coupling or floating load.

Parameter	Symbol	Min	Typ	Max	Units	Notes
Uncorrelated High Probability Jitter	T_UHPJ			0.15	UIpp	[a]
Duty Cycle Distortion	T_DCD			0.05	UIpp	
Total Jitter	T_TJ			0.30	UIpp	[b]
Eye Mask	T_X1			0.15	UI	[c]
Eye Mask	T_X2			0.40	UI	[d]
Eye Mask	T_Y1	200			mV	[e]
Eye Mask	T_Y2			*450*	mV	[f]

table 6 Transmitter Output Jitter Specifications

a. This parameter is defined as: Jitter distribution where the value of the jitter show no correlation to any signal level being transmitted. Formally defined as deterministic jitter, T_DJ
b. The link will operate with a BER of 10^{-15}
c. See Figure 8
d. See Figure 8
e. See Figure 8

f. See Figure 8

Parameter	Symbol	Min	Typ	Max	Units	Notes
RX Baud rate	R_Baud		5.000		GSym/s	[a]
Input Differential Voltage	R_Vdiff	*100*		*900*	mVppd	[b]
Differential Resistance	R_Rdin	80	100	120	Ohms	
Bias Voltage Source Impedance (Load Type 1)	R_Zvtt			30	Ohms	[c]
Differential Input Return Loss (100MHz to *2.5 GHz*)	R_SDD1			-8	dB	[d]
Differential Input Return Loss (*2.5 GHz* to 5 GHz)	R_SDD1				dB	[e]
Common Mode Input Return Loss (100MHz to *2.5 GHz*)	R_SCC1			-6	dB	[f]
Termination Voltage	R_Vtt	Not specified 1.2 -8%	Not specified	Not specified 1.2 +5%	V V	[g] R_Vtt floating, Load Type0 R_Vtt=1.2V Nominal, Load Type1
Input Common Mode Voltage See Note [h]	R_Vrcm	-0.05 720		1.85 R_Vtt-10	V mV	[i] R_Vtt floating, Load Type0 R_Vtt=1.2V Nominal, Load Type1
Wander divider	n			10		[j]

table 7 Receiver Electrical Input Specifications

a. CEI-6G-SR is defined to operate between baud rates of 4.976 and 6.375 Gsym/s, However QSGMII will operate at 5 Gsym/s with a tolerance of +/-100ppm.
b. Min Value is changed from the standard and reduced to 100 mV.
c. Load Type 1 is with DC Coupling.
d. See Figure 9
e. See Figure 9
f. See Figure 9
g. For floating load, input resistance must be > 1K Ohms.
h. For Vcm definition, see Figure 10.
i. Input Common Mode voltage for AC-Coupled or floating load input with min T_Vdiff
j. See Figure 2-27 & Figure 2-28 in CEI-6G-SR document for details

Parameter	Symbol	Min	Typ	Max	Units	Notes
Bounded High Probability Jitter	R_BHPJ			0.45	UIpp	[a]
Sinusoidal Jitter,maximum	R_SJ-max			5	UIpp	

table 8 Receiver Input Jitter Tolerance Specifications

Parameter	Symbol	Min	Typ	Max	Units	Notes
Sinusodial Jitter, High Frequency	R_SJ-hf			0.05	UIpp	
Total Jitter(Does not include Sinu-sodial Jitter)	R_TJ			0.60	UIpp	b
Eye Mask	R_X1			0.30	UI	c
Eye Mask	R_Y1			*50*	mV	d
Eye Mask	R_Y2			*450*	mV	e

table 8 Receiver Input Jitter Tolerance Specifications

a. This is the sum of Uncorrelated Bounded High Probability Jitter (0.15 UI) and Correlated Bounded High Probability Jitter (0.30 UI).

Uncorrelated Bounded High Probability Jitter: Jitter distribution where the value of the jitter show no correlation to any signal level being transmitted. Formally defined as deterministic jitter, T_DJ

Correlated Bounded High Probability Jitter: Jitter distribution where the value of the jitter shows a strong correlation to the signal level being transmitted. This jitter may considered as being equalisable due to its correlation to the signal level

b. The link will operate with a BER of 10^{-15}

c. See Figure 8

d. See Figure 8.

e. See Figure 8.

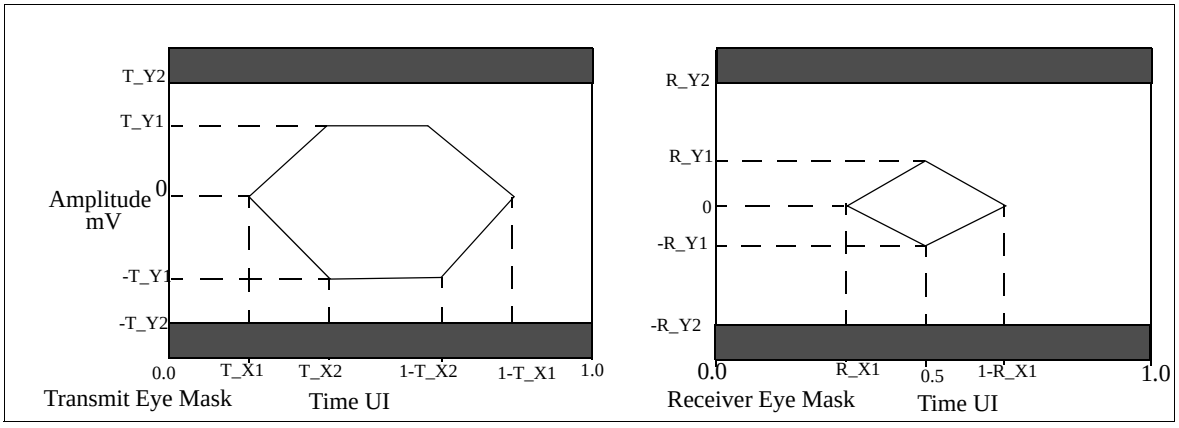

Figure 8 Driver and Receiver Eye Mask

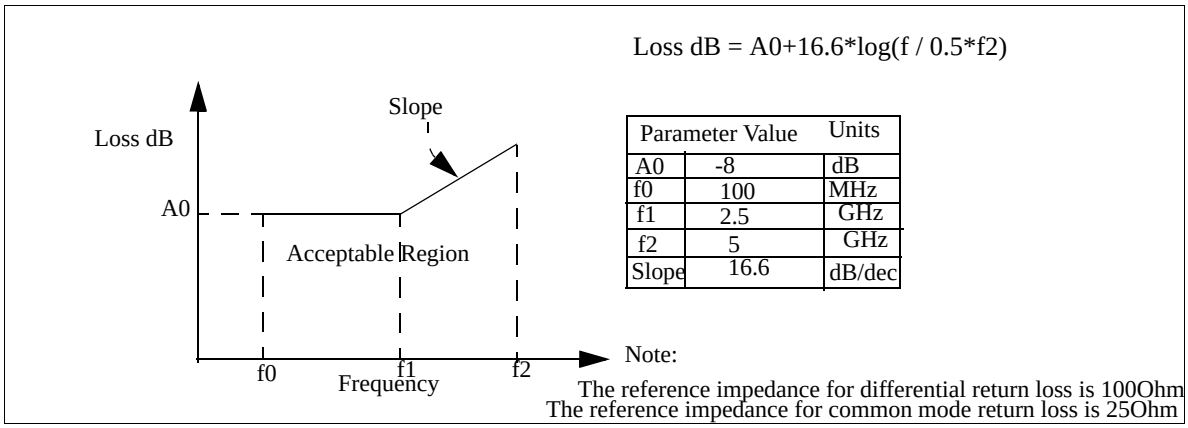

Parameter	Value	Units
A0	-8	dB
f0	100	MHz
f1	2.5	GHz
f2	5	GHz
Slope	16.6	dB/dec

Figure 9 Driver and Receiver Differential Return Loss

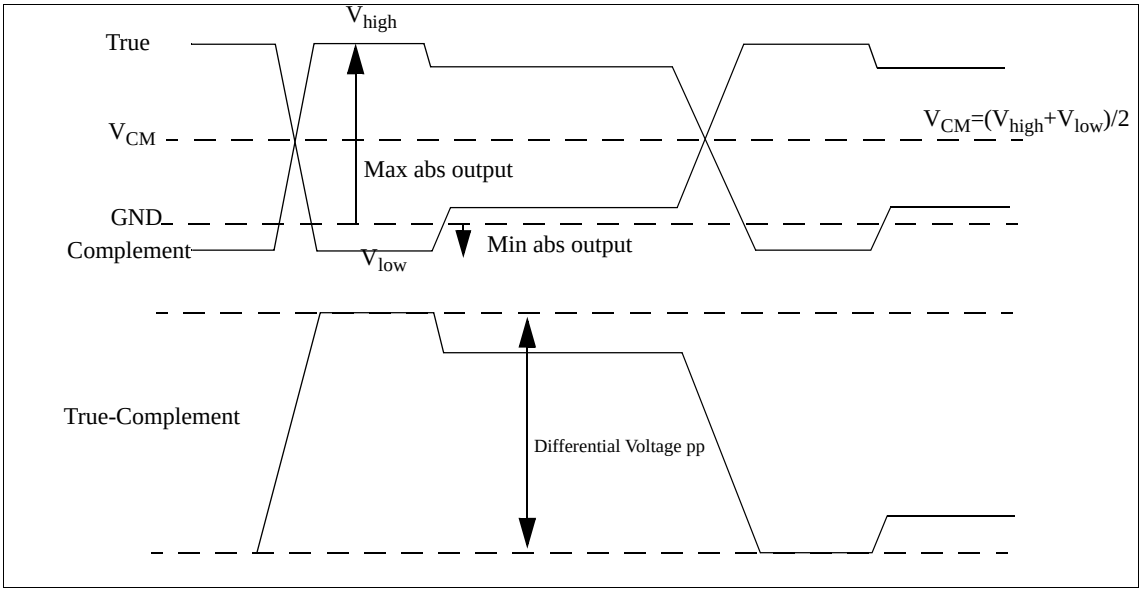

Figure 10 Definition of Driver Amplitude and Swing

QSGMII channel has a total loss budget of 12 dB as shown in table 9.

Parameter	**Condition**	**Loss (dB)**
Transmitter package loss margin	at 2.5 Ghz	0.5
Receiver Package loss margin	at 2.5 Ghz	0.5
Channel Interconnect loss margin	at 2.5 Ghz	10.0
Channel Crosstalk, Reflection and other loss margin	at 2.5 Ghz	1.0
Total Channel loss	at 2.5 Ghz	**12.0**

table 9 QSGMII Channel Loss Budget

Figure 11 shows the interconnect loss template for the channel. At 2.5 GHz, maximum allowed interconnect loss is -10.0 dB which represents a typical 20 inch trace on FR4 PCB.

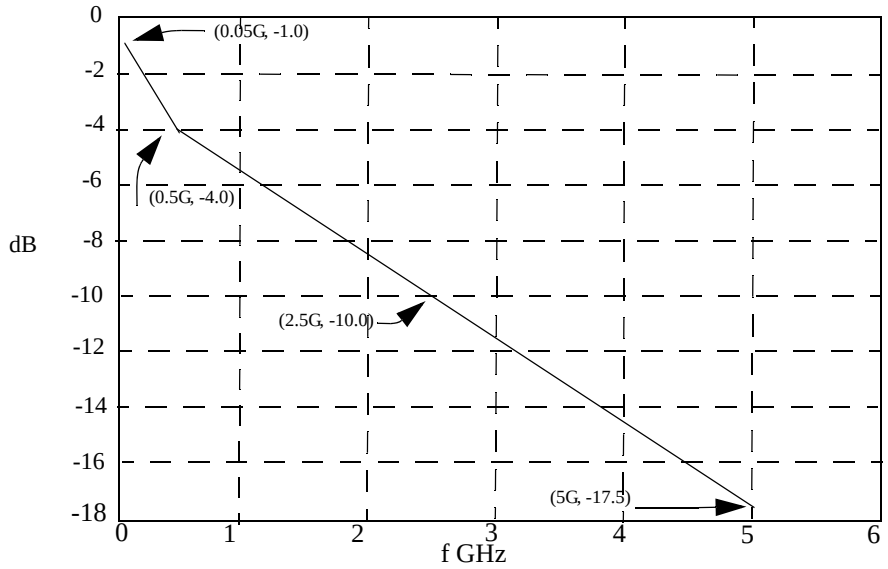

Figure 11 Interconnect Loss Template for QSGMII Channel

Cisco Systems Intellectual Property

Cisco encourages others to adopt this specification. To the extent that Cisco has proprietary rights to information contained in this specification, any company wishing to use this specification may do so under reasonable, non-discriminatory terms, with reciprocity, to implement and fully comply with the specification.

The reasonable non-discriminatory terms are:

Cisco will not assert any claims of any patents essential for implementing this specification that are owned or controlled by Cisco against any party for making, using, selling, importing or offering for sale a product that implements the specification, provided, however, that Cisco retains the right to assert its patents (including the right to claim past royalties) against any party that asserts a patent it owns or controls (either directly or indirectly) against any products that comply with this specification. Cisco retains the right to assert its patents against any product or portion thereof that is not necessary for compliance with the specification.

Royalty-bearing licenses are available to anyone who prefers that option.

For information contact:

Dan Lang
Senior IP Counsel
Cisco Systems
408-526-6672
standards-ipr@cisco.com