



Universidad de
Oviedo



ESCUELA POLITÉCNICA DE INGENIERÍA DE GIJÓN

MÁSTER UNIVERSITARIO EN INGENIERÍA DE TELECOMUNICACIÓN

TRABAJO FIN DE MÁSTER N° 202002

FPGA-SIDE INVESTIGATIONS FOR 56 Gbps 4-PAM SERIAL DATA TRANSMISSION

AUTOR: AZA VILLAMOR, ANDREA

TUTOR: D. ARIAS PEREZ DE AZPEITIA, MANUEL

**COTUTORES: PLETL, JOSEF
LASKA, VALENTIN**

FECHA: 07 - 2020

ACRONYMS

ISI	Inter-Symbol Interference	LFSR	Linear Feedback Shift Register
PAM	Pulse Amplitude Modulation	PAM2	PAM 2-levels
PAM4	PAM 4-levels	NRZ	Non-return-to zero
DTLE	Discrete Time Linear Equalizer	CTLE	Continuous Time Linear Equalizer
DFE	Decision Feedback Equalization	FEC	Forward Error Correction
RS-FEC	Reed-Solomon FEC	PRBS	Pseudo Random Binary Sequence
QPRBS	Quaternary PRBS	SNR	Signal-to-Noise Ratio
BW	Bandwidth	TX/RX	transmitter/receiver
xcvr	transceiver	FPGA	Field-programmable gate array
BER	Bit Error Rate	SER	Symbol Error Rate
VSR	Very Short Range	MR	Medium-Range
LR	Long Range	PCB	Printed Circuit Board
TD	Transition Density	CEI	Common Electrical I/O
OIF	Optical Internetworking Forum	FFE	Forward Feedback Equalizer
PIO	Parallel Input Output	I2C	Inter-Integrated Circuit
PHY	physical layer	CDR	Clock Data Recovery
LSR	Linear Shift Register Internetworking Forum	SMA	SubMiniature version A
QSFP	Quad Small Form-factor Pluggable	QSFP-DD	QSFP double density
TTK	Transceiver Tool-Kit	TSDs	Temperature Sensing Diodes
PMA	Physical Medium Attachment	PCS	Physical Coding Sublayer
S10	Stratix 10	UI	Unit Interval
COM	Channel Operating Margin	DFP	D flip-flop
FIFO	first in, first out	EMIB	Embedded Multi-die Interconnection Bridge
SoC	System-on-Chip	IF	Intermediate frequency
DFT	Design For Test	VGA	Variable Gain Amplifier
MSB	Most Significant Bit	LSB	Less Significant Bit

CONTENTS

0	SUMMARY - ES	2
0.0	ACRÓNIMOS	2
0.1	OBJETIVO	3
0.2	MODULACIÓN PAM4	3
0.2.1	MODULACIONES DE SEÑAL EN COMUNICACIONES SERIE DE ALTA FRECUENCIA	3
0.2.2	DISTORSIÓN DE SEÑAL EN MODULACIÓN PAM4	4
0.2.3	CUANTIFICACIÓN DE PÉRDIDAS DE SEÑAL EN TRANSMISIONES DE ALTA FRECUENCIA	4
0.3	SISTEMA VHDL PARA ANÁLISIS DE RENDIMIENTO DE COMUNICACIONES SERIE PAM4	5
0.3.1	CONTROLADOR DEL TRANSCÉPTOR (XCVR-CONTROLLER)	5
0.3.2	GENERACIÓN VERIFICACIÓN DE LA SECUENCIA BINARIA	6
0.3.2.1	PATRÓN DE TEST PRBS31	6
0.3.2.2	GENERADOR CHECKER PRBS31	6
0.3.3	ANÁLISIS ESTADÍSTICO DE ERROR	7
0.3.4	MUESTREO DE TEMPERATURA	7
0.3.5	SUB-SISTEMA DE CONTROL EMBEBIDO	7
0.4	RESUMEN DE RESULTADOS	8
0.4.1	CÁLCULO DE LA TASA BINARIA DE ERROR (BER)	8
0.4.2	DESCRIPCIÓN Y ANÁLISIS DE RESULTADOS RELEVANTES	8
0.4.3	SÍNTESIS DE RESULTADOS	10
0.5	CONCLUSIONES	11
1	INTRODUCTION	13
1.1	OBJECTIVE	13
1.2	DOCUMENT STRUCTURE	14
2	PAM4 SIGNAL INTEGRITY	15
2.1	HIGH-SPEED SERIAL-LINK SIGNAL MODULATION	15
2.2	PAM2 , PAM4 PERFORMANCE COMPARISON	15
2.3	EYE-DIAGRAM BASED ANALYSIS	16
2.3.1	PAM4 3-EYED DIAGRAM FORMAL PARAMETERS	18
2.3.1.1	EYE ANATOMY	19
2.3.2	ISI IMPACT ON EYE DISTORTION	20
3	EQUALIZATION PARAMETERS	22
3.1	TX EQUALIZATION	22
3.1.1	FIR FILTER BASED TX-EQUALIZER	23
3.2	RX EQUALIZATION	24
3.2.1	CONTINUOUS TIME LINEAR EQUALIZATION , CTLE	24
3.2.2	DIRECT FEEDBACK EQUALIZATION , DFE	24
4	CEI-OIF 4.0	26
4.1	CEI-OIF XCVR REQUIREMENTS	26
4.2	CHANNEL REFERENCE MODELS	27
5	HIGH-DATARATE PAM4 SERIAL LINK ANALYSIS SYSTEM	28
5.1	FUNCTIONAL DIVISION	28
5.1.1	1 ST LEVEL DIVISION	28
5.1.2	2 ND LEVEL DIVISION	29
5.1.2.1	μPROCESSOR EMBEDED SUB-SYSTEM (QSYS-EMBEDDED SYSTEM)	29
5.1.2.2	FPGA DESIGN	30
5.1.2.2.1	XCVR CONTROLLER (NATIVE-PHY-IP E-TILE XCVR)	31
5.1.2.2.2	QPRBS31 PARALLEL GENERATOR/CHECKER	32
5.1.2.2.3	ERROR ANALYZER	32
5.1.2.2.4	DATA ALIGNMET (DESKEW LOGIC)	33
5.1.2.2.5	FREQUENCY MEASUREMET (FREQ-MEASURE)	33
5.1.2.2.6	TIME MEASUREMENT (TST_TIMER)	33
5.1.2.2.7	SYNCHRONIZATION CLOCK GENERATION (CLCK_MGMT_GEN)	33

5.1.2.2.8	XCVR TEMPERATURE MEASUREMENT (TEMP-SENSE)	34
5.1.3	3 RD LEVEL DIVISION (NATIVE-PHY-IP)	35
5.2	DATA INTERFACES	35
5.2.1	STRATIX 10TX BOARD CONNECTORS	37
5.2.2	AVALON MEMORY MAPPED INTERFACES (AVALON-MM)	38
5.2.3	AVALON PARALLEL INPUT/OUTPUT INTERFACES (AVALON-PIO)	39
5.3	CLOCKING	39
5.4	E-TILE NATIVE PHY CONFIGURATION AND CONTROL	40
5.4.1	CONFIGURATION / CONTROL STRUCTURE	40
5.4.2	MINIMUM CONFIGURATION SEQUENCE	42
5.4.3	CONFIGURABLE ATTRIBUTES AND STATUS INDICATORS	42
5.5	HW SYSTEM DEBUGGING AND VERIFICATION	47
5.6	NIOS-II μ PROCESSOR PROGRAM	48
5.6.1	GENERAL FLOW	48
5.6.2	AUTO-SWEEP	50
5.6.3	BAUDRATE-SWEEP	50
5.6.4	TEST CONFIGURATION / RUN	50
6	SERIAL LINK TEST AND RESULT ANALYSIS	52
6.1	MEASUREMENT CALCULATION METHODS	52
6.1.1	BER CALCULATION	53
6.1.2	ON-DIE STATISTICAL EYE DIAGRAM GENERATION	53
6.2	SERIAL LINK TESTS AND RESULTS ANALYSIS	54
6.2.1	TEST DESCRIPTION	54
6.2.1.1	INTERNAL SERIAL LOOPBACK VALIDATION TEST	54
6.2.1.2	QSFP-DD 1.5 M LINK TEST	55
6.2.1.3	QSFP-DD 0.5 M LINK TEST	55
6.2.1.4	QSFP 1.0 M LINK TEST	56
6.2.1.5	SMA 2.4 MM RF CONNECTOR LINK TEST	57
6.2.1.6	QSFP 1.0 M + CONNECTION BOARD + SMA LOOPBACK (WORST CASE LINK TEST)	58
6.2.2	TEST RESULTS ANALYSIS	60
6.2.2.1	INTERNAL SERIAL LOOPBACK VALIDATION TEST	61
6.2.2.2	QSFP-DD 1.5 M LINK TEST	63
6.2.2.3	QSFP-DD 0.5 M LINK TEST	66
6.2.2.4	QSFP 1.0 M LINK TEST	68
6.2.2.5	SMA 2.4 MM RF LINK TEST	71
6.2.2.6	QSFP 1.0 M + CONNECTION BOARD + SMA LOOPBACK (WORST CASE LINK TEST)	72
6.2.2.7	TEST RESULT SUMMARY	75
7	CONCLUSIONS AND FUTURE WORK	77
7.1	CONCLUSIONS	77
7.2	FUTURE WORK	77
	REFERENCES	79
A	PAM4 SIGNAL DEGRADATION AND EQUALIZATION TECHNIQUES	83
A.1	PAM4 SIGNALLING	83
A.2	HIGH FREQUENCY SIGNAL DISTORTION ON ELECTRICAL SERIAL LINKS (INSERTION AND RETURN LOSSES)	85
A.2.1	ISI (INTER-SYMBOL INTERFERENCE)	85
A.2.1.1	ISI REDUCTION USING PAM4	85
A.2.1.2	CHANNEL EQUALIZATION GOALS	86
A.2.1.3	ISI FORMULATION	87
A.2.1.4	TRANSITION DENSITY	87
A.2.2	PAM4,PAM2 ATTENUATION COMPARISON (INSERTION LOSSES)	88
A.3	PAM4 EQUALIZATION	89
A.3.1	DISCRETE TIME LINEAR EQUALIZATION USING FIR FILTERS	89
A.3.2	TX EQUALIZER (DE-EMPHASIS EQUALIZATION)	90
A.3.3	RX EQUALIZER (CTLE + DFE EQUALIZATION STRUCTURE)	93

A.3.3.1	CTLE (CONTINUOUS TIME LINEAR EQUALIZATION)	93
A.3.3.2	DFE (DECISION FEEDBACK EQUALIZATION)	94
A.3.3.3	1/(1+D) SIGNAL PRECODING	95
B	INTEL STRATIX 10 TX ARCHITECTURE	97
B.1	INTEL STRATIX 10TX SIGNAL INTEGRITY DEVELOPMENT KIT OVERVIEW	97
B.2	HYPERFLEX ARCHITECTURE (FPGA ARCHITECTURE)	98
B.2.1	FPGA-XCVR INTERFACE (EMIB, EMBEDDED MULTI-DIE INTERCONNECT BRIDGE)	98
B.2.2	XCVR TILE-STRUCTURE (E-TILES)	99
B.2.2.1	E-TILE STRUCTURE (GXE TRANSCEIVERS)	99
B.2.2.2	E-TILE RESOURCES	100
B.2.2.2.1	S10 E-TILE BANKS	100
B.2.2.2.2	S10 E-TILE CLOCKING	100
B.2.2.3	CDR (CLOCK AND DATA RECOVERY)	101
B.2.3	DOUBLE WIDTH TRANSFER MODE	101
B.2.3.1	DATA ALIGNMENT (DE-SKEW)	103
B.3	SIGNAL PATH	105
B.3.1	GXE TRANSCEIVER DATAPATH	105
B.3.1.1	PMA DATAPATH	106
B.3.2	STRATIX 10 E-TILE XCVRs LOOPBACK MODES	106
B.4	PMA EQUALIZATION	108
B.4.1	RX-EQUALIZER ADAPTATION MODES	108
B.4.2	EQUALIZATION	110
B.4.2.1	SLICER	110
B.5	XCVRs CONTROLLER INTERFACE (NPDM AND NATIVE-PHY IP ARBITRATOR)	110
C	PRBS31 PARALLEL GENERATOR/CHECKER	112
C.1	PRBS31 (QPRBS31) TEST PATTERN	112
C.2	LFSR-BASED SERIAL TEST PATTERN GENERATOR	112
C.3	PRBS PARALLEL GENERATOR/CHECKER	113
C.3.1	PRBS GENERATOR	113
C.3.1.1	TRANSITION MATRIXES SERIAL-PARALLEL TRANSFORMATION	113
C.3.1.2	128-b PARALLEL GENERATOR FEEDBACK EQUATIONS	115
C.3.1.3	PARALLEL GENERATOR STRUCTURE (LFRS + XOR-TREE)	116
C.3.1.4	PAM4 SIGNALING TEST PATTERN (QUATERNARY PRBS31, QPRBS31)	116
C.3.2	PRBS CHECKER	117
C.3.2.1	LOCKED CONDITION	118
C.3.3	VERIFICATION	118
C.3.3.1	SIMULATION BASED VERIFICATION	118
C.3.3.2	OVER FPGA DESIGN VALIDATION	119
D	STATISTICAL ERROR ANALYZER (RS(544,514) FEC SPECIFIC IMPLEMENTATION)	121
D.1	RS(544,514) FORWARD ERROR CORRECTION (FEC)	121
D.1.1	RS-FEC IMPACT ON PAM4 BER	121
D.1.1.1	FEC GAIN	121
D.1.1.2	BURST CORRECTION ENHANCEMENT (INTERLEAVING)	122
D.1.2	FEC ERROR CORRECTION CAPABILITIES	122
D.1.2.1	CORRECTION OF UP TO M RANDOM ERRORS	122
D.1.2.2	CORRECTION OF BURSTS UP TO LENGTH λ	123
D.1.3	RS-FEC(544,514) ERROR CORRECTION CAPABILITIES	123
D.1.3.1	CORRECTION OF UP TO M RANDOM ERRORS	123
D.1.3.2	CORRECTION OF BURSTS UP TO LENGTH λ	124
D.2	STATISTICAL ERROR ANALYZER (RS(544,514) SPECIFIC)	124
D.2.1	BURST DENSITY	124
D.2.1.1	BURST END CRITERIA (MNZ)	124
D.2.2	GAP DEFINITION	125
D.2.3	GAP, BURST LENGTH MEASUREMENT	125
D.2.4	STATISTICAL ERROR ANALYZER STRUCTURE	126

D.2.5	STATISTICS DATA OUTPUT	126
D.3	VALIDATION	127
E	NIOS-2 APP MENU	130

LIST OF FIGURES

0.1 (a) Ejemplo de señales digitales señales digitales con modulaciones PAM4 ; (b) NRZ 3

0.1 Esquema simplificado del sistema de análisis de integridad de señal para transmisiones serie PAM4. 6

0.1 Peor caso de canal evaluado para transmisiones serie PAM4|NRZ de alta frecuencia 9

0.2 Progresión de la tasa binaria de error (BER) para transmisiones NRZ y PAM4 a 22.4 GBs de 2 h de duración, sobre el canal mostrado en la imagen 0.1. 10

0.3 Curvas solapadas de progression de tasa binaria de error (BER) para transmisiones NRZ y PAM4 de 2 h de duración sobre los canales serie evaluados 11

1.1 PAM4 signal integrity analysis system implified schema including main design units involved in signal generation and processing 14

2.1 (a) PAM4 , (b) PAM2 baseband example modulated signaling (this figured is copy of A.1, included here for the sake of clarity). 15

2.2 Possible signal transitions between PAMn modulation symbols for PAM4, PAM2(NRZ) , [1]. 17

2.3 (a) PAM2 ; (b) PAM4 example eye-diagram shapes. 17

2.4 (a) PAM4 eye parameter measuring definition [8]; (b) eye-height measured over BER 1E-6 contour over statistical eye-diagram, EH6 [2]; (c) eye-width measured over BER 1E-6 contour over statistical eye-diagram, EW6 [2]. 18

2.5 (a.1) PAM4 upper inner eye; (a.2) PAM4 mid inner eye; (a.3) PAM4 lower inner eye; (b.1) PAM4 eye transitions +3 -1; (b.2) PAM4 eye transitions +1 -3; (c) PAM4 outer eye (transitions +3 -3), [1] 20

2.6 (a.1) Possible NRZ undistorted signal transitions ; (a.2) eye-diagram for undistorted signal transitions overlapping; (a.3) NRZ distorted signal transitions; (a.4) eye-diagram for distorted signal transitions overlapping, [10]; (b.1) example low distortion eye-diagram; (b.2) mid distortion eye-diagram; (b.3) mid-highdistortion eye-diagram [9]. 21

3.1 Generic xcvr signal path including TX FIR filter-based equalizer , CTLE+DFE RX equalizer 22

3.2 (a) PAM4 eye-diagram at TX-output without pre-distortion; (b) PAM4 eye-diagram at RX-input without pre-distortion; (c) PAM4 eye-diagram at TX-output when pre-distortion applied; (d) PAM4 eye-diagram at RX-input when pre-distortion applied, [10]. 23

3.3 PAM4 signal distortion progression over serial channel when both no TX-equalization is applied, and de-emphasis pre-distortion is applied over transmitted signal. 23

3.4 Common structure for k-coefficients (c_{-1} , c_0 , c_1 , .. c_{k-2}) FIR filter with $k = m+2$ (note that Δ blocks are 1UI delay). 24

3.5 CTLE gain (transfer function) shape for RX-equalization (from CTLE characteristic definition in [12].) 25

3.6 DFE equalizer internal structure usign slicer + feedback FIR filter architecture (note that DFE's input $r'(n)$ is the received signal after CTLE equalization). 25

4.1 (a) Transmit signaling and mapping diagram, minimum functional requirements for PAM4 transmitter; (b) receive signaling and mapping diagram, minimum functional requirements for PAM4 receiver, [8]. 26

4.2 (a) Gray mapping at PAM4 encoded lanes for TX output lanes; (b) and RX-output lanes, where PAM4 symbols are referred to as 0 , 1 , 2 , 3 (notation equivalent to -3 , -1 , +1 , +3, respectively) , [13]. 27

4.3 Informative channel insertion-loos limit-mask for channel compliance at 29.0 Gsym/s , [13]. 27

5.1 PAM4 serial link testing system simplified funtional/conceptual diagram. 29

5.2 PAM4 serial link testing system 1st level functional division. 29

5.3 PAM4 serial link testing system 2nd level functional division Qsys-embedded system. 31

5.4 PAM4 serial link testing system 2nd level functional division FPGA design (note that interfaces to/from QSYS embedded system are only connected if map directly to signals shown , for the sake of simplicity) 34

5.5 Intel temperature-sensor-ipcore block diagram (enables control on internal on-die temperature samples from in-chip TSDs' measurements)¹ , [14]. 34

5.6 Simplified TX (a); RX (b) PMA block diagram for each channel in Stratix 10 E-tile xcvr , [16]. 36

5.7 Data interface widthover simplified TX (a); RX (b) PMA block diagram for each channel in Stratix 10 e-tile xcvr; including path from/to soft external PRBS31 generator/checker. 36

5.8 FPGA tst_timer outputs connection to tst_timer Avalon-PIO registers on one channel_reg_set (one PHY). 37

5.9 'temperature-sensor-ipcore' control/data connection from/to Nios-2 through Avalon-PIO registers ² 37

5.10 Stratix 10TX Transceiver Signal Integrity Development Kit board overview , [17]. 38

5.11 'native-phy-ipcore' reconfig. Avalon-mm interface (Avalon-mm slave,dynamic reconfiguration interface), [11]. 39

5.12 (a) Example input Avalon-PIO register associated latch; (b) input Avalon-pio register controller; (c) example output Avalon-PIO register controller³ 40

5.13 Datapath clocks (signal frequency or lane rate) over simplified TX (a)⁴ ; RX (b) PMA block diagram for each channel in Stratix 10 e-tile xcvr; including path from/to soft external PRBS31 generator/checker. 41

5.14 'native-phy-ipcore' (Stratix 10 E-Tile Transceiver Native Phy) block symbol including shared⁵ dynamic reconfiguration Avalon-mm (reconfig_avmm) interface. 41

5.15 Simplified block diagram for native-phy-ipcore's configuration/control structure used in PAM4 serial link testing system. 42

5.16 Design structure used for preliminary verifications of xcvr configuration sequences observing eye-diagram through TTK eye-viewer⁶ . 47

5.17 Design structure used for verifications of on developed PRBS31 generator/checker integration with native-phy-ipcore by performance comparison to Intel's soft PRBS31 generator and checker ipcores. 48

5.18	(a) Nios-2 console application status info. (stratix 10TX hw info , including configured bitrate for used PHYs; (b) per PHY xcvr status info; (c) e-tile/core temperature , and QSFP-DD modules' status info.)	49
5.19	Ongoing serial link test status info printed on Nios-2 console.	51
6.1	Basic serial link testing set-up reference model (copy of image 6.4 , included here for the sake of clarity)	52
6.2	Example statistical eye-diagram measurement using TTK eye-viewer for on-die statistical eye generation over low distortion PAM2 transmission.	54
6.3	Set-up for PAM2/PAM4 serial link test over QSFP-DD cable of 1.5 m length connecting 2 e-tile xcvr channels on adjacent QSFP-DD modules.	55
6.4	Set-up for PAM2/PAM4 serial link test over QSFP-DD cable of 0.5 m length connecting 2 e-tile xcvr channels on adjacent QSFP-DD modules.	56
6.5	Set-up for PAM2/PAM4 serial link test over QSFP cable of 1 m length connecting 2 e-tile xcvr channels on adjacent QSFP-DD modules.	56
6.6	Set-up for PAM2/PAM4 serial link test over 2.4mm SMA cables of 0.7 m length connecting e-tile xcvr channel in external loopback.	57
6.7	S11 parameter for 0.7m length 2.4mm SMA cables SF101E/11PC24/11PC24/700mm from hubersuhner (parameter curve obtained from provider specific test report).	57
6.8	Set-up for PAM2/PAM4 worst case serial link test over channel composed as QSFP cable of 1 m length + <100 mm connection-board trace + external loopback over connection-board + <100 mm connection-board trace + QSFP cable of 1 m length(a); signal path for worst case serial link tested (b); connection board detail (c).	59
6.9	S11 parameter for 1.0m length SMA cables used for external serial loopback over connection board (parameter curve obtained from Rohde & Schwarz specific test report).	60
6.10	PAM4 and NRZ BER progression over 2 h. test time for reference no-distortion internal serial loopback mode test at 28GBs baudrate.	62
6.11	E-tile xcvr on-die temperature progression over 2 h. test time for reference no-distortion internal serial loopback mode test at 28GBs baudrate.	63
6.12	BER against xcvr on-die temperature progression over 2 h. test time for reference no-distortion internal serial loopback mode test at 28GBs baudrate.	63
6.13	BER measured over 358 s. against baudrate, for baudrate-sweep over QSFP-DD 1.5m-length serial link test.	64
6.14	BER measured over 358 s. against baudrate, for baudrate-sweep over QSFP-DD 1.5m-length serial link test.	64
6.15	PAM4 and NRZ BER progression over 2 h. test time for QSFP-DD 1.5m-length serial link test at 28GBs baudrate.	65
6.16	BER against xcvr on-die temperature progression over 2 h. test for QSFP-DD 1.5m-length serial link test at 28GBs baudrate.	65
6.17	PAM4 eye diagram measured for QSFP-DD 1.5m-length serial link test at 28GBs baudrate (EH6 = 7, EW6 = 30).	65
6.18	NRZ eye diagram measured for QSFP-DD 1.5m-length serial link test at 28GBs baudrate (EH6 = 176, EW6 = 36).	66
6.19	PAM4 and NRZ BER progression over 2 h. test time for QSFP-DD 0.5m-length serial link test at 28GBs baudrate.	67
6.20	E-tile xcvr on-die temperature progression over 2 h. test time for QSFP-DD 0.5m-length serial link test at 28GBs baudrate.	67
6.21	BER against xcvr on-die temperature progression over 2 h. test for QSFP-DD 0.5m-length serial link test at 28GBs baudrate.	67
6.22	PAM4 eye diagram measured for QSFP-DD 0.5m-length serial link test at 28GBs baudrate (EH6 = 9, EW6 = 32).	68
6.23	NRZ eye diagram measured for QSFP-DD 0.5m-length serial link test at 28GBs baudrate (EH6 = 175, EW6 = 36).	68
6.24	PAM4 and NRZ BER progression over 2 h. test time for QSFP 1.0m-length serial link test at 28GBs baudrate.	69
6.25	E-tile xcvr on-die temperature progression over 2 h. test time for QSFP 1.0m-length serial link test at 28GBs baudrate.	69
6.26	BER against xcvr on-die temperature progression over 2 h. test for QSFP 1.0m-length serial link test at 28GBs baudrate.	70
6.27	PAM4 eye diagram measured for QSFP 1.0m-length serial link test at 28GBs baudrate (EH6 = 8, EW6 = 35).	70
6.28	NRZ eye diagram measured for QSFP 1.0m-length serial link test at 28GBs baudrate (EH6 = 181, EW6 = 35).	70
6.29	PAM4 and NRZ BER progression over 2 h. test time for SMA 2.4 mm RF 0.7m-length serial link test at 28GBs baudrate.	71
6.30	E-tile xcvr on-die temperature progression over 2 h. test time for SMA 2.4 mm RF 0.7m-length serial link test at 28GBs baudrate.	72
6.31	BER against xcvr on-die temperature progression over 2 h. test for SMA 2.4 mm RF 0.7m-length serial link test at 28GBs baudrate.	72
6.32	BER measured over 358 s. against baudrate, for baudrate-sweep over 'connection-board loopback set-up' serial link test.	73
6.33	PAM4 and NRZ BER progression over 2 h. test time for 'connection-board loopback set-up' serial link test at 22.4GBs baudrate.	73
6.34	E-tile xcvr on-die temperature progression over 2 h. test time for 'connection-board loopback set-up' serial link test at 22.4GBs baudrate.	74
6.35	BER against xcvr on-die temperature progression over 2 h. test for 'connection-board loopback set-up' serial link test at 22.4GBs baudrate.	74
6.36	PAM4, NRZ overlapped BER progressions over 2 h. test time for serial link tests performed	75
A.1	PAM4 (a) , PAM2 (b) baseband example modulated signaling.	83
A.2	PAM4-eye narrowing over PAM4 simulated 3-eyed diagram with low distortion (A,B indicate where slower PAM4 transition crosses NRZ 'zero-threshold', PAM4 'zero-threshold' respectively).	84
A.3	PAM4-eye narrowing over PAM4 simulated 3-eyed diagram with low distortion (A,B placed where transitions from different levels to zero 'zero-threshold'.)	84
A.4	PAM2(NRZ) zero distortion eye-diagram (a); PAM4 zero distortion 3-eye-diagram(b) -[24]-.	84
A.5	Example typical insertion loss curves for non-ideal channel (a); corresponding delta-like impulse responses (b), [1].	86
A.6	Squared PAM2(NRZ) modulated signal shape (1 V peak-to-peak normalized) used as input signal (a); non-ideal delta-like channel impulse response (causing ISI) (b); PAM2 smoothed signal at channel output (ISI affected)(c) , [1].	86

A.7	Typical channel transfer function (insertion loss curve) for limited BW channel (blue); TX-qualizer + channel composed transfer function (green); RX-equalization transfer function (black); TX-qualizer + channel + RX-qualizer flattened composed transfer function (red) , [1]	87
A.8	Example normalized impulse response $h(t)$, with channel delay $d=2$, sampled at $n \cdot t_b$ where $t_b = 1/f_b$ and f_b =baudrate.	87
A.9	Generic xcvr signal path including TX FIR filter-based equalizer , CTLE+DFE RX equalizer	89
A.10	Common structure for k-coefficients ($c_{-1} , c_0 , c_1 , ..c_{k-2}$) FIR filter with $k = m+2$ (note that Δ blocks are 1UI delay).	90
A.11	3-tap FIR filter structure for equalization on TX-side in Stratix 10 TX signal integrity development kit (1 FIR filter based TX-equalizer is included for each channel on each e-tile xcvr).	90
A.12	PAM4 eye-diagram at TX-output without pre-distortion (a); PAM4 eye-diagram at RX-input without pre-distortion (b);PAM4 eye-diagram at TX-output when pre-distortion applied (c); PAM4 eye-diagram at RX-input when pre-distortion applied (d) , [10].	91
A.13	Example channel impulse response $h(t)$ (causing ISI) (a); example DFE FIR filter impulse response (b); DFE + channel composed impulse response $h'(t)$ ideally cancelling 1st pre- and post- cursors (c); overlapped channel , 1-tap pre-emphasis FIR filter equalizer, composed response example (included for further clarity) [25] (d)	92
A.14	Signal path from TX-side equalizer input ($s(n)$) to RX input ($r(n)$)	92
A.15	Example channel impulse response without equalization(red), and composite response when 3-tap TX FIR equalizer used(blue) (a); example PAM2 signal at channel input when TX-FIR equalizer not used(red), and if used (blue) (b); channel, 3-tap TX FIR equalizer, composed(channel + 3-tap TX FIR equalizer) transfer function , [27].	93
A.16	Example 64-levels PAM4 eye-diagram measured at channel input when 3-tap TX FIR filter based equalizer is used (after pre-distortion) (a); associated PAM4 eye-diagram measured at channel output (pre-distortion compensated) (b), [1].	93
A.17	CTLE gain (transfer function) shape for RX-equalization (from CTLE characteristic definition in [12].)	94
A.18	Channel transfer function (frequency response) transformation when CTLE is used for equalization at RX (a) [28]; example NRZ eye-diagram evolution from channel input to RX-input (after CTLE) (transmitter eye(b.1) , transmitted eye after channel (b.2) , transmitted after after channel and CTLE (b.3)) , [29].	95
A.19	TX/RX equalization structure with FIR filter de-emphasis TX-equalizer; and CTLE + FIR filter DFE RX-equalizer	95
A.20	DFE equalizer internal structure using slicer + feedback FIR filter architecture (note that DFE's input $r'(n)$ is the received signal after CTLE equalization).	96
A.21	Simplified serial signal path schema when 1/1+D pre-codingsweepdecoding is used (Δ is a unit delay used to perform 1/1+D encoding)(a); example demonstrating 1/1+D working principle. [30].	96
B.1	Intel Stratix 10 TX Development Kit Block Diagram (courtesy of Intel [17]).	97
B.2	Intel Stratix 10 TX Transceiver Signal Integrity Development Kit Picture(courtesy of Intel [17]).	97
B.3	H-Tile and E-tile Layout Configuration for Intel Stratix 10TX Device with 5 E-Tiles and 1 H-Tile (144 Transceiver channels) (courtesy of Intel [11]).	98
B.4	Example of die-to-die Heterogeneous Integration using EMIB Technology for Stratix 10 FPGA(courtesy of Intel [34]).	98
B.5	E-tile xcvr internal GXE channels(tx/RX) set (courtesy of Intel [16]).	99
B.6	E-Tile GXE channel xcvr simplified internal structure block diagram(courtesy of Intel [16]).	99
B.7	E-tile transceiver channels block diagram , including E-tiles transceivers and connector placement among S10 E-tile banks (courtesy of Intel [17]).	101
B.8	E-tile transceiver dedicated input reference clocks (courtesy of Intel [17]).	102
B.9	E-tile transceiver dedicated input reference clocks generation schema on Stratix 10TX devices (courtesy of Intel [36]).	103
B.10	Data interfaces width on E-Tile - FPGA core signal path when double width transfer selected (EMIB 64-b parallel interfaces re-allocated) (courtesy of Intel [35]).	103
B.11	S10 E-Tile transceiver GXE used even channels on used E-Tile PHYs (courtesy of Intel [35]).	104
B.12	Parallel data interface bit assignment in 80-b EMIB parallel interface (with 64-b data out of 80-b) (E-tile - S10 FPGA interface) (courtesy of Intel [37]).	104
B.13	Double width transfer data alignment between odd-, even- channels 64-b parallel data lanes (de-skew logic performance) (courtesy of Intel [11]).	105
B.14	E-Tile GXE channel xcvr simplified internal structure block diagram (courtesy of Intel [16]), copy of B.14 included here for the sake of clarity.	105
B.15	E-Tile GXE channel xcvr simplified internal structure block diagram in PMA direct mode, PCS block by-passed (courtesy of Intel [16]).	106
B.16	TX(a) / RX(b) PMA internal signal path simplified block diagram for each GXE channel in S10 E-Tile xcvr (courtesy of Intel [16]).	106
B.17	PMA internal architecture complete block diagram (for each GXE channel in S10 E-Tile xcvr) (courtesy of Intel [11]).	107
B.18	Internal serial loopback path within PMA internal architecture on S10 E-Tile xcvr GXE channel (courtesy of Intel [11]).	107
B.19	Reverse serial loopback path within PMA internal architecture on S10 E-Tile xcvr GXE channel (courtesy of Intel [11]).	108
B.20	Intel S10 TX receiver equalization architecture block diagram including on-die instrumentation (courtesy of Intel [7]).	108
B.21	Example signal level thresholds for PAM4 RX-slicer measuring distortion-,ISI- free eye-diagram (3 symmetric inner-eyes without misalignment).	110
B.22	Arbitration logic structure supporting multiplexation on native-phy-ipcore reconfig-avmm internal reconfiguration interface (courtesy of Intel [11]).	111

- C.1 Generic schema for synchronous n-LFSR(n-DFF LFSR)-based PRBSn generator (c_j defined in GF(2) for $j= 1, .. n$; and $c_n = 1$) 113
- C.2 PRBS31 test pattern serial generator (a); ckecker (b), (generator polynomial $1+x^28+x^31$), [41]. 114
- C.3 PRBS9 serial generator using 9-DFF LFSR structure (generator polynomial $g(x) = 1 + x^5 + x^9$) 114
- C.4 PRBS9 4-b LFSR-based parallel generator (generator polynomial $g(x) = 1 + x^5 + x^9$) 115
- C.5 LFSR + XOR-TREE structure used for PRBS31 parallel generator implementation , [39]. 117
- C.6 Modelsim simulation results for PRBS31 generator/checker verification with own PRBS31 checker fed by altera_avalon_data_patern_generator 119
- D.1 Input to output BER when implementing FEC on serial link transmissions for common FEC codes [8]. 121
- D.2 Typical n-branch interleaving structure for n=4 (Δ represents 1UI delay synchronized latches). 122
- D.3 Developed VHDL statistical error analyzer internal structure. 127
- D.4 Simplified schema for simulation testbench used to validate the statistical error analyzer. 128
- D.5 Modelsim simulation results for statistical error analyzer validation testbench (note that 128-b width signals are shown in hexadecimal format). 129

LIST OF TABLES

5.3	PMA TX-equalization settings (vod pre-tap-(1-3) post-tap) (TX-equalization data in this table is extracted from, [11]	42
5.4	PMA RX-equalization parameters (RX-equalization data in this table is extracted from tables 24 and 41 in, [11].	43
5.5	PMA RX configuration pre-sets (pre-set equalizer parameters for 10G, 28G-LR, 28G-VSR, 56G-LR, 56G-VSR models defined in 4.2, [8])	44
5.6	E-tile channel PMA configuration parameters	45
5.7	Serial link transmission status indicators	46
6.1	Common xcvr and test configuration parameter values used for high-speed serial link testing.	61
6.2	Summary of BER results for serial link tests performed.	76
A.1	PAM2,PAM4 relevant attributes for ISI vulnerability analysis , [24]	88
B.1	Parallel clock (TX_clkout) definitions, [35]	102
B.2	PMA RX-equalization parameters supported ranges (definition for parameters in table can be found in 5.4) [11]	109





0. SUMMARY - ES

Se recoge en este documento un resumen del desarrollo de un sistema para el análisis de integridad de señal en comunicaciones serie de alta velocidad empleando modulación de señal PAM4 (Pulse Amplitude Modulation 4-levels).

Se realiza a modo de introducción un breve análisis de la modulación PAM4 empleada, que justifica su uso para incrementar la tasa de transferencia en comunicaciones serie frente a una modulación binaria.

Se incluye asimismo una versión reducida de los acrónimos empleados, con objeto de dotar este resume de carácter auto-contenido.

0.0. ACRÓNIMOS

PAM	Pulse Amplitude Modulation	PAM2	PAM 2-levels
PAM4	PAM 4-levels	NRZ	Non-return-to zero
QPRBS	Quaternary PRBS	PRBS	Pseudo Random Binary Sequence
CEI	Common Electrical I/O	OIF	Optical Internetworking Forum
BER	Bit Error Rate	FEC	Forward Error Correction
xcvr	transceiver	tx rx	transmitter receiver
S10	Stratix 10	FPGA	Field-programmable gate array
PIO	Parallel Input Output	SMA	SubMiniature version A
QSFP	Quad Small Form-factor Pluggable	QSFP-DD	QSFP double density
E/S	entrada/salida		



0.1. OBJETIVO

Se plantea como objetivo el análisis de rendimiento de la modulación de señal PAM4 (en términos de tasa de bits erróneos, BER), para comunicaciones serie por encima de 30 Gbps.

Para la determinación de la calidad del enlace y la viabilidad de la comunicación PAM4 se emplean como referencia los valores de tasa binaria de error recogidos en la especificación CEI-OIF 4.0 [8].

Al sistema de análisis BER implementado se añade un módulo que realizará un análisis estadístico en tiempo real de la distribución de errores en la comunicación, con objeto de caracterizar el comportamiento del canal a analizar.

Para la implementación del sistema de testeo descrito en este documento se ha empleado una FPGA Stratix 10TX de Intel, que incluye transceptores PAM4 capaces de cursar comunicaciones serie hasta 56 Gbps.

0.2. MODULACIÓN PAM4

0.2.1. MODULACIONES DE SEÑAL EN COMUNICACIONES SERIE DE ALTA FRECUENCIA


Las comunicaciones serie con elevada tasa de transferencia binaria (típicamente hasta 28 Gbps) emplean comúnmente una modulación de señal binaria denotada NRZ o PAM2 en la que cada bit transmitido se representa mediante un nivel de señal, dando lugar a una señal de 2 niveles. Se muestra en la figura 0.1b un ejemplo de señal con modulación NRZ.

Con el empleo de una modulación de señal NRZ, el incremento de la tasa de transferencia en un enlace serie supone un incremento de la frecuencia de señal transmitida.


Dada la naturaleza paso-bajo de los canales de comunicación serie (i.e. el filtrado de las componentes de señal del alta frecuencia), el incremento de la frecuencia de señal enviada (para incrementar la velocidad de transferencia) supone un incremento de degradación de la señal en su paso por el canal. Dicha degradación de señal (típicamente denominada 'pérdida de inserción') se traduce en un incremento del número de bits erróneos en el receptor, que para señales NRZ por encima de los 28 GHz (28 Gbps) impide generalmente la recuperación de la información transmitida a partir de la señal recibida.

Para permitir el incremento de la tasa binaria se ha propuesto el empleo de la modulación PAM4, una modulación digital con 4 niveles de señal para la que cada nivel permite codificar 2 bits.

De este modo, manteniendo constante la frecuencia de señal (y por consiguiente evitando el incremento de degradación en su paso por el canal), se consigue idealmente duplicar la velocidad de transferencia con respecto de la que se obtiene empleando una modulación PAM2 a la misma frecuencia de señal. Se muestra en la figura 0.1a un ejemplo de señal con modulación PAM4.



(a)



(b)

Figure 0.1.- (a) Ejemplo de señales digitales con modulaciones PAM4 ; (b) NRZ .



0.2.2. DISTORSIÓN DE SEÑAL EN MODULACIÓN PAM4

Idealmente, el empleo de una modulación de señal PAM4 permite duplicar la tasa de transferencia binaria con respecto de la tasa de transferencia para una señal NRZ con la misma frecuencia, sin incrementar la distorsión de señal introducida por el canal. Sin embargo, las características de la modulación PAM4 hacen que resulte más vulnerable a la distorsión de señal introducida en el canal (i.e. que la máxima distorsión soportada sin producir errores en recepción sea menor que la soportada por una señal NRZ a la misma frecuencia, luego con igual distorsión).

Para una señal PAM4, los niveles de señal (cada uno representando un símbolo de 2 bits) se encuentran más próximos entre sí que en una señal NRZ (i.e. la distancia entre 2 niveles de señal PAM4 adyacentes es 1/3 de la distancia entre los 2 niveles de una señal NRZ). La distancia reducida entre los niveles de señal hace que la modulación PAM4 sea más vulnerable a la atenuación del nivel de señal en el canal, dado que la distorsión que puede asumir un nivel de señal (símbolo) sin alcanzar el nivel de señal contiguo (i.e. sin provocar un error de símbolo en recepción) es aproximadamente 1/3 de la distorsión soportada por una señal NRZ. De este modo, para la misma frecuencia de señal (luego para la misma atenuación de canal), la tasa binaria de error observada será mayor para una señal PAM4 (a partir de cierto nivel de atenuación de canal). Dicha vulnerabilidad de la modulación PAM4 a la atenuación de canal hace necesario el análisis de la calidad de señal (i.e. de la BER alcanzable) en un enlace serie cuando se emplea una señal PAM4 para incrementar la tasa de transferencia de datos.

0.2.3. CUANTIFICACIÓN DE PÉRDIDAS DE SEÑAL EN TRANSMISIONES DE ALTA FRECUENCIA

Si bien la modulación PAM4 presenta mayor vulnerabilidad a la atenuación de la señal, su empleo para incrementar la tasa de transferencia se justifica mediante la comparación cuantitativa de la degradación de señal sufrida al duplicar la tasa de transferencia binaria si:

- se mantiene el empleo de modulación NRZ, y por consiguiente se duplica la frecuencia de señal para duplicar la tasa binaria;
- o bien se mantiene la frecuencia de señal y se duplica la tasa binaria empleando modulación PAM4.

Para realizar una comparación cuantitativa, la degradación de señal se puede expresar en términos de pérdidas de señal para cada alternativa:

- Cuando se duplica la tasa de transferencia empleando una modulación PAM4, el incremento de vulnerabilidad a la atenuación de señal en el canal se expresa en términos del factor en que se reduce la tolerancia a la distorsión del nivel de señal para cada símbolo enviado. Así, la reducción de 1/3 en tolerancia a la atenuación del nivel de señal (en tensión) soportada por la modulación se puede expresar como una pérdida (o degradación) de potencia de señal de aproximadamente 9.5 dB, con respecto de una señal NRZ a la misma frecuencia:
$$20 \log_{10}(1/3) \approx 9.5 \text{ dB}$$
- Cuando se duplica la tasa de transferencia duplicando la frecuencia de la señal NRZ, el incremento de la pérdida de señal se puede obtener de los límites de pérdida de inserción establecidos en la especificación CEI-OIF 4.0 [8]. Para un incremento de la frecuencia de señal de 14 GHz a 28 GHz (14 Gbps a 28 Gbps para la señal NRZ), i.e., se tiene un incremento de pérdida de señal de aproximadamente 11 dB.

Así, la degradación de señal sufrida al duplicar la tasa binaria resulta inferior cuando dicho incremento de la velocidad de transferencia se realiza manteniendo la frecuencia de señal y empleando una modulación PAM4 (9.5 dB < 11 dB), justificando el empleo de modulación de señal PAM4 en comunicaciones serie de elevada tasa de transferencia.



0.3. SISTEMA VHDL PARA ANÁLISIS DE RENDIMIENTO DE COMUNICACIONES SERIE PAM4

Para evaluar la viabilidad de transferencias de señal PAM4 con tasa binaria en el rango 30 Gbps – 56 Gbps, se diseña un sistema que permite la detección y análisis de errores de bit sobre la señal recibida.

Se plantean como objetivos del sistema de análisis de error:

- La detección de cada bit erróneo para una transmisión PAM4 en curso sobre el canal bajo análisis, que permita el cálculo de la tasa binaria de error (BER) para evaluar la integridad de la señal recibida.
- El análisis de la distribución de errores de bit sobre la señal recibida, incluyendo la caracterización de ráfagas de error en el enlace serie bajo análisis.⁷
- Soporte al análisis de error sobre señales NRZ en el mismo rango de frecuencia de señal (*baudrate*) soportado para transmisiones PAM4, permitiendo analizar la degradación de señal asociada al cambio de modulación.

Se muestra en la figura 0.1 un esquema simplificado del sistema VHDL de análisis de error en comunicaciones serie PAM4, donde el bloque denotado como 'xcvr-control' permite la configuración de los transceptores PAM4 integrados en la placa de la FPGA Stratix 10TX, y la serie de bloques morados en la parte inferior cubre la generación y tratamiento de la secuencia binaria empleada en la transmisión.

Para el sistema mostrado en la figura 0.1, el tratamiento de la secuencia binaria empleada para modular la señal serie se realiza de modo que:

- La secuencia binaria a transmitir se genera en el bloque 'PRBS31_generate' y es recibida por el controlador del transceptor, que modula y transmite la señal serie (tx_serial_data).
- El mismo controlador recibe y demodula la señal serie tras atravesar el canal. La secuencia binaria recibida se procesa en el bloque 'PRBS31_checker', que identifica el patrón de errores que afecta a la señal recibida.
- El patrón de errores extraído de la secuencia recibida sirve de entrada al bloque 'statistical_error_analyzer', que analiza el patrón de error y genera datos para el análisis estadístico del patrón de error.

Se incluye en las sub-secciones 0.3.1-0.3.5 una descripción simplificada de la funcionalidad de cada módulo del sistema mostrado en la figura 0.1.

0.3.1. CONTROLADOR DEL TRANSEPTOR (XCVR-CONTROLLER)

Se trata de un componente proporcionado por Intel para el control de los transceptores PAM4 (sus interfaces principales para el control de la transmisión se muestran en la figura 0.1).

El controlador expone una interfaz de configuración (marcada con línea gris punteada en la figura 0.1) que permite controlar la configuración del transceptor, y por tanto la generación de la señal transmitida: *baudrate* (frecuencia de señal) y modulación de señal (NRZ o PAM4)⁸.

El transceptor recibe la secuencia binaria a transmitir sobre un bus de datos paralelo (tx_parallel_data). Dicha secuencia se serializa y se emplea para modular la señal a transmitir (emitida a través de la interfaz tx_serial_data).

Para el análisis del canal serie, las interfaces serie del controlador del transceptor (tx_serial_data, rx_serial_data) se mapean a conectores de la placa Stratix 10TX y se realiza la conexión tx – rx mediante el canal serie a analizar (de modo que las interfaces tx_serial_data, rx_serial_data quedan conectadas mediante un enlace *loopback* externo).

De forma análoga, la señal recibida en la interfaz rx_serial_data se deserializa, y la secuencia binaria recibida queda accesible sobre el bus paralelo

⁷ nótese que la caracterización de ráfagas de error en el enlace no tiene como objetivo la determinación de la integridad de señal a la entrada del receptor, sino la determinación de estadísticos empleados para aplicar técnicas de corrección de errores en recepción que permitan reducir la tasa de error (BER) observada.

⁸ nótese que cada transceptor PAM4 en la placa Stratix 10TX expone una serie de parámetros configurables que permiten controlar la generación de señal modulada y procesado en recepción, no incluidos en este resumen por simplicidad

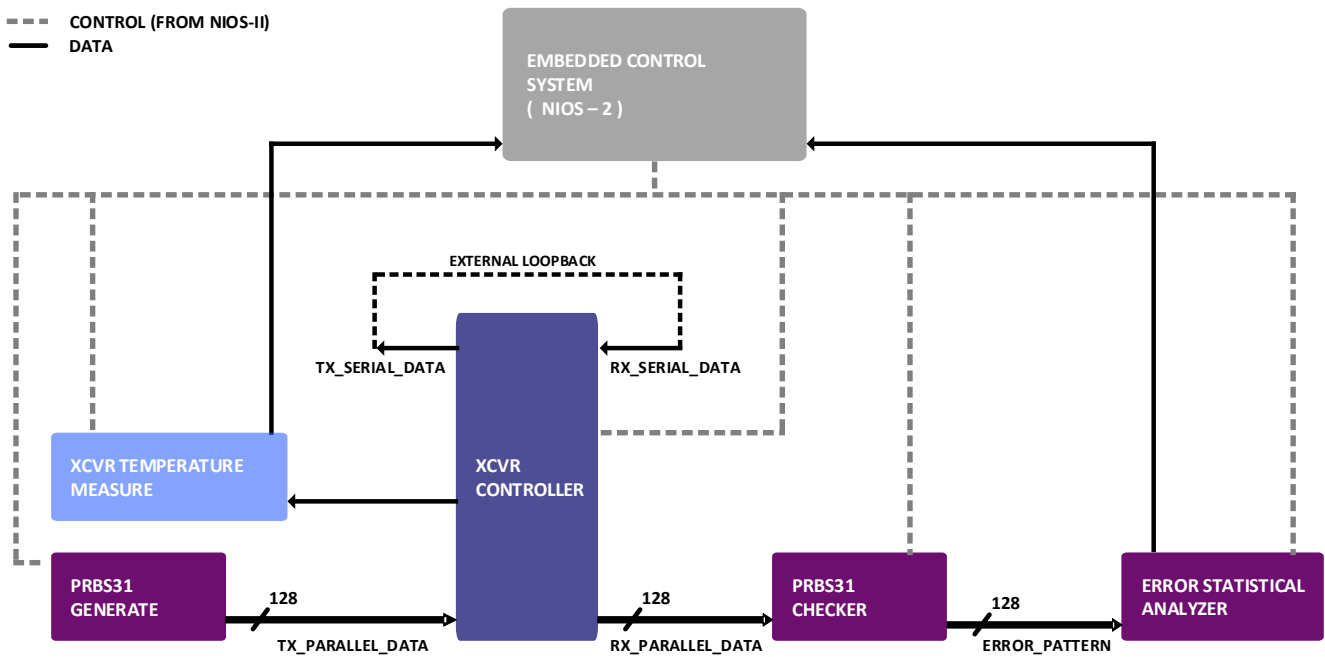


Figure 0.1.- Esquema simplificado del sistema de análisis de integridad de señal para transmisiones serie PAM4.

rx_parallel_data para el análisis de error.

0.3.2. GENERACIÓN|VERIFICACIÓN DE LA SECUENCIA BINARIA

0.3.2.1 PATRÓN DE TEST PRBS31

Para la generación de la señal de modulación PAM4 se emplea el patrón de test PRBS31 (pseudo-random binary sequence de orden 31) definido en las especificaciones [8],[41].

Para una transmisión serie, la probabilidad de error en recepción se maximiza cuando se emplean secuencias binarias aleatorias para modular la señal transmitida.


Una secuencia PRBSn se define como un patrón binario sin periodicidad de longitud 2^n , cuyo comportamiento estadístico coincide con el de una secuencia aleatoria. Si dicha secuencia PRBSn de longitud suficiente, un patrón de test generado por repetición de dicha secuencia (patrón de test PRBSn) se puede asumir como aleatorio, despreciando la periodicidad introducida por la repetición de la secuencia de longitud 2^n .

De este modo, los patrones de test PRBSn permiten testear canales serie en condiciones de máxima probabilidad de error ('aleatoriedad' de la secuencia binaria) empleando una señal de generación determinista, que puede ser recreada en el receptor para identificar el patrón de error sobre la señal recibida.

0.3.2.2 GENERADOR|CHECKER PRBS31

En la figura 0.1, el bloque denotado 'PRBS-generate' genera el patrón de test PRBS31 empleado como secuencia binaria para modular la señal transmitida. El generador se ha diseñado con una estructura síncrona de generación en paralelo, de modo que en cada pulso del reloj de sincronización del diseño se generen simultáneamente n bits consecutivos de la secuencia PRBS31, que se inyectan en el bus paralelo 'tx_parallel_data' para ser transmitidos sobre el canal serie.

La secuencia recibida a la salida del canal sirve como entrada al módulo 'PRBS-checker', que compara la secuencia recibida con el patrón de test PRBS31 para determinar la ocurrencia de errores en la transmisión y obtener el patrón de error que afecta a la secuencia recibida (nótese que el módulo 'PRBS-checker' es capaz de identificar cada bit erróneo, así como su posición en el patrón de test PRBS31, permitiendo la caracterización de las ráfagas de error observadas sobre la secuencia binaria recibida). Dicho patrón de error se expone en el bus 'error_pattern', quedando disponible para el cálculo



de la tasa de error de bit (BER) y el análisis estadístico de la distribución de errores en la secuencia (en términos de error de bit y ráfagas de error).

0.3.3. ANÁLISIS ESTADÍSTICO DE ERROR

'error statistical analyzer' en la figura 0.1 denota al bloque que procesa el patrón de error detectado sobre la secuencia recibida, produciendo datos que permiten realizar un análisis estadístico de la distribución de errores en la comunicación:

- número de bits erróneos desde el inicio de la transmisión;
- longitud de cada ráfaga de error observada (nótese que se denota como ráfaga de error de longitud λ a una sucesión de $(\lambda - 2)$ bits erróneos o no, delimitados por 2 bits erróneos);
- densidad de cada ráfaga de error observada definida como: (no. de bits erróneos)/(longitud de ráfaga).

Sendos datos permiten obtener tanto las distribuciones de probabilidad de longitud de ráfaga y densidad de ráfaga, como la tasa binaria de error (BER).

Si bien en el momento de redactar este documento el bloque 'statistical error analyzer' no ha sido aún integrado en el sistema, su diseño se ha validado mediante simulación.

0.3.4. MUESTREO DE TEMPERATURA

Para analizar la degradación de la calidad de señal (en términos de incremento de la tasa binaria de error) asociada al incremento de temperatura en el transceptor (debido a la transmisión activa) se incorpora un módulo que permite muestrear la temperatura en los chips que soportan los transceptores integrados en la placa Stratix 10TX. Con el análisis de la progresión de temperatura durante la transmisión, se persigue no sólo determinar la degradación del rendimiento de los transceptores con el incremento de temperatura; sino también disociar la degradación de la BER observada debida a la distorsión de canal, de la causada por el calentamiento del transceptor.

En la imagen 0.1, el bloque 'xcvr temperature measure' denota una instancia de dicho módulo, proporcionado por Intel como interfaz de acceso al valor de temperatura en cada uno de los chips que soportan dichos transceptores.

0.3.5. SUB-SISTEMA DE CONTROL EMBEBIDO

Para el control y coordinador del sistema diseñado se incorpora un sub-sistema de control basado en un microprocesador (Intel Nios-II integrado en la placa Stratix 10TX), que permite:

- Simplificar la implementación de las secuencias de configuración del transceptor.
- Coordinar el funcionamiento de los sub-sistemas involucrados en el análisis de la transmisión.
- Recoger información de estado de la transmisión serie ofrecida por dichos sub-sistemas, así como datos relativos al análisis de error sobre la transmisión en curso.
- Interactuar con el sistema para la configuración dinámica del transceptor, así como mostrar información tanto de estado del sistema como de progresión de BER y temperatura durante la transmisión.

Dicho microprocesador se integra en el diseño VHDL en forma de sub-sistema embebido, que incluye los recursos requeridos para la comunicación y sincronización entre el microprocesador y la FPGA sobre la que se despliega el sistema de análisis de transmisiones serie PAM4 descrito. Se incluyen asimismo en el sub-sistema embebido componentes que dan soporte a una conexión de consola Ubuntu con el microprocesador Nios-II, empleada como E/S estándar para el microprocesador. Así, se emplea la conexión de consola de modo que además de mostrar la información de estado de sistema y progresión de la transmisión; soporte la introducción de comandos para realizar cambios de configuración en el transceptor.

Para la obtención de los resultados mostrados la sección 0.4, la información de estado recogida por el microprocesador se almacena en un archivo de log procesado posteriormente para representar las progresiones de las variables de interés (BER y temperatura) durante el test de transmisión.



0.4. RESUMEN DE RESULTADOS

Se recoge en esta sección un resumen de los resultados más significativos obtenidos de los test de transmisión PAM4 de alta velocidad, en términos de tasa binaria de error (BER) medida. Para la interpretación inequívoca de los resultados presentados se incluye asimismo una breve descripción del método empleado para el cálculo de la tasa binaria de error.

0.4.1. CÁLCULO DE LA TASA BINARIA DE ERROR (BER)

Para la representación de la progresión de la tasa binaria de error durante el test de transmisión se emplean las muestras periódicas recogidas del contador de errores de bit detectados, calculando la tasa binaria de error para cada muestra tomada en el curso de la transmisión como:

$$BER = (\text{no. bits erróneos})/(\text{no. de bits enviados})$$

Inicialmente, se ha realizado un test de transmisión libre de distorsión de canal en el que señal serie se transmite entre las interfaces tx_serial_data, rx_serial_data en el interior de la placa Stratix 10TX, evitando así la distorsión introducida por los conectores de la placa y el canal de comunicación. Dicho test se ha realizado con el objetivo de medir la mínima tasa binaria de error alcanzable para transmisiones serie NRZ y PAM4 con *baudrate* = 28 GBs. Sendas tasas binarias de error, medidas en ausencia de distorsión de canal, se emplean como referencia para evaluar la degradación de la tasa de error (BER) introducida por cada uno de los canales analizados.

Para los test iniciales realizados en ausencia de distorsión, el número de errores de bit observados fue 0 para ambas modulaciones de señal, NRZ y PAM4. Cuando el número de errores de bit es 0, la tasa binaria de error (BER) presentada se corresponde con el valor denotado típicamente como BER(CL = 0.95).

BER(CL = 0.95) representa el valor BER tal que la BER observada se encontrará por debajo del valor BER(CL = 0.95) calculado con un nivel de confianza del 95 %, para cualquier repetición del test realizado en las mismas condiciones. BER(CL = 0.95) se obtiene como [21]:

$$BER(CL = 0.95) = 3/(\text{no. bits enviados})$$

0.4.2. DESCRIPCIÓN Y ANÁLISIS DE RESULTADOS RELEVANTES

Para el análisis de rendimiento de la modulación PAM4 en enlaces serie se ha realizado una secuencia de test de transmisión sobre distintos canales, bajo las mismas condiciones de transmisión:

- Cada test se realizó para una transmisión ininterrumpida de 2 h. sobre el canal a evaluar.
- Para cada canal se realizó un test de 2 h. para cada modulación (NRZ y PAM4) empleando la misma frecuencia de señal (*baudrate*), con objeto de observar la degradación de la tasa de error con el cambio de modulación para cada canal.
- Para cada test se tomaron muestras de tasa binaria de error (BER) cada 0.5 s.
Además de la tasa de error, la progresión de temperatura en el transceptor se muestreó para cada test con la misma frecuencia.
- Para cada canal analizado, el test de transmisión se ha realizado a la máxima frecuencia de señal a la que el canal demostró soportar una comunicación PAM4 manteniendo la tasa de error observada por debajo del límite de la especificación CEI-OIF 4.0 [8], BER = 1E-4.

Se presentan en este resumen únicamente los resultados obtenidos para el peor canal serie testeado (siendo el único canal testeado que demostró no soportar la comunicación PAM4 serie a 56 Gbps dentro del límite de BER = 1E-4).

Se muestra en la imagen 0.1 el set-up empleado para el peor caso de enlace serie evaluado, donde las flechas rojas marcan el recorrido de la señal serie transmitida (señal emitida sobre la interfaz 'tx_serial_data' mostrada en la figura 0.1) desde uno de los conectores QSFP-DD en la placa Stratix 10TX.


Siguiendo el recorrido marcado en la imagen 0.1b:

- La señal es conducida a través de un cable QSFP de 1.0 m de longitud hasta una placa de conexionado, que conduce la señal a un conector SMA.
- Desde el conector SMA, se realiza una conexión en *loopback* empleando un cable SMA de 1.0 m de longitud hacia el conector SMA correspondiente en la placa de conexionado.
- Desde el punto anterior, el trayecto de retorno de señal a través del cable QSFP se encuentra señalizado mediante flechas naranjas en la imagen 0.1b.

Se muestran en la figura 0.2 las progresiones de la tasa binaria de error (BER) durante las 2 h. de transmisión para las modulaciones NRZ y PAM4. Sendas curvas BER se corresponden con la tasa de error para transmisiones realizadas a frecuencia de señal 22.4 GHz (*baudrate* = 22.4 GBs), máxima frecuencia de señal para la que el canal descrito soporta una transmisión PAM4 manteniendo una tasa de error BER <1E-4, donde BER = 1E-4 es el límite definido en el estándar CEI-OIF 4.0 [8] (representado en la imagen 0.2 con línea negra punteada).


La tasa de error medida para la transmisión NRZ a 22.4 GBs (curva azul en la imagen 0.2) coincide con la tasa de error medida inicialmente en ausencia de distorsión (i.e. con una transmisión libre de errores de bit), de modo que la curva BER NRZ se corresponde con la progresión del valor BER(CL = 0.95).

La degradación de la tasa binaria de error (al final de la transmisión) desde 1E-14(NRZ) a 1E-5 (aproximadamente) al emplear una modulación PAM4 (con la misma frecuencia de señal, 22.4 GHz) evidencia la vulnerabilidad de la modulación PAM4 a la atenuación de señal introducida en el canal (frente al comportamiento de la transmisión NRZ, capaz de soportar dicha atenuación sin errores de transmisión).



(a)

(b)



(c)

Figure 0.1.- (a) Set-up empleado para el test de transmisión PAM4|NRZ sobre el peor enlace serie analizado, conformado por: cable QSFP de 1.0 m de longitud + placa de conexionado (con longitud de pista <100 mm) + *loopback* sobre cable SMA de 1.0 m de longitud + retorno sobre placa de conexionado (con longitud de pista <100 mm) + retorno sobre cable QSFP de 1.0 m de longitud ; (b) recorrido de señal sobre canal serie: cable QSFP 1.0 m + placa de conexionado + *loopback* sobre SMA + placa de conexionado + cable QSFP 1.0 m (nótese que el recorrido de retorno se encuentra indicado con flechas naranjas); (c) detalle de placa de conexionado.




Figure 0.2.- Progresión de la tasa binaria de error (BER) para transmisiones NRZ y PAM4 a 22.4 GBs de 2 h de duración, sobre el canal mostrado en la imagen 0.1.

0.4.3. SÍNTESIS DE RESULTADOS

Se muestra en la imagen 0.3 la superposición de las curvas BER medidas para cada uno de los test realizados sobre los canales evaluados (nótese que los resultados representados se corresponden con transmisiones serie a 28 GBs excepto para las curvas denotadas como 'BER CONN-BOARD', correspondientes a las tasas de error observadas para transmisiones a 22.4 GBs sobre el canal descrito en la sub-sección anterior). Sendas curvas de progresión de tasa binaria de error se mantienen en valores de $BER < 1E-4$, demostrando la viabilidad de las transmisiones PAM4 a 56 Gbps dentro de los límites de la especificación, para los canales evaluados.

Entre las curvas solapadas, las denotadas como QSFP-DD 0.5 , QSFP 1.0 y QSFP-DD 1.5 se corresponden con test realizados sobre canales de características similares con longitud creciente (0.5 m, 1.0 m , 1.5 m, respectivamente). La degradación de la tasa binaria de error (BER) observada entre las transmisiones PAM4 sobre sendos canales refleja cómo la distorsión sobre la señal PAM4 crece con la longitud de canal (i.e. con el incremento atenuación de potencia de señal al incrementar la longitud del canal), debido a la vulnerabilidad de la modulación PAM4 a la atenuación de señal.




Figure 0.3.- Curvas solapadas de progression de tasa binaria de error (BER) para transmisiones NRZ y PAM4 de 2 h de duración sobre los canales serie evaluados (nótese que las curvas BER mostradas se corresponden con transmisiones a 28GBs, excepto para las curvas donotadas como 'BER CONN-BOARD' que se corresponden con transmisiones la máxima frecuencia de señal soportada por el canal asociado para una BER $1E-4$: 22.4 GBs).

0.5. CONCLUSIONES

Se recoge en esta sub-sección una relación breve de las conclusiones derivadas del sistema desarrollado.

- Se ha completado el diseño e implementación de un generador de secuencias pseudo-aleatorias con arquitectura de generación paralela. Asimismo, se ha desarrollado el verificador análogo para la extracción del patrón de error introducido por la distorsión de canal sobre dicha secuencia pseudo-aleatoria.
- Se ha diseñado e implementado un módulo VHDL para el análisis estadístico del patrón de error en términos de errores de bit aislados, así como de ráfagas de error. Si bien no encuentra integrado en el sistema de análisis de integridad de señal descrito, el diseño de analizador estadístico ha sido verificado mediante simulación.
- Se ha desarrollado un sistema que permite la configuración del transceptor en tiempo real, así como la recogida de datos que permiten analizar las progresiones de tasa binaria de error y temperatura en el interior del transceptor, para transmisiones serie de alta velocidad empleando modulación de señal NRZ o PAM4.
- Se ha demostrado la viabilidad de la comunicación serie empleando modulación PAM4 a 44.8 Gbps para el peor caso de canal evaluado.
- Si bien no se ha incluido en este resumen por simplicidad, se ha observado la independencia de la tasa binaria de error medida (BER) con respecto de la temperatura en el transceptor, para variaciones de la temperatura de hasta 2 °C.





1. INTRODUCTION

High speed serial link transmissions use massively 2 signal coding schemes: NRZ (non-return-to-zero), a.k.a PAM2 (pulse-amplitude-modulation 2-level), and PAM4 (pulse-amplitude-modulation 4-level).

Serial link transmissions over 25 Gbps are moving to PAM4 (4-level pulsed modulation scheme) signaling to overcome channel bandwidth issues affecting high datarate NRZ signaling, that introduce a significant increase in link losses and power consumption for NRZ (PAM2) datarate over 25Gbps (thus digital signal frequency over 25GHz). PAM4 modulation alleviates frequency dependent issues in serial links by cutting down signal frequency to half the frequency of NRZ signaling for a given datarate (as each PAM4 signal level -modulation symbol- represents 2 bits, so that doubling the density of data), thus becoming a feasible alternative to double the achievable datarate over serial links without increasing frequency-dependant link-losses.

2 major aspects must be considered when replacing NRZ by PAM4 modulation in high-speed serial link transmissions:

- Link-loss budget will be maintained/reduced as signal frequency (baudrate) can be decreased.
- FEC strategies may be used to counter the SNR degradation suffered by PAM4 signaling (as the number of modulation symbols increases, the 'space' between the associated signal levels is also reduced causing SNR degradation).

As hw manufacturers demonstrate feasibility of transmissions over 56Gbps [1],[2], common serial-link standards include support for PAM4 for high-speed transmissions in last releases [3] (PCIe 6.0) thus promoting investigations on PAM4 transmission enhancement for NRZ replacement in high-speed transmissions.

PAM4 xcvr on sale are commonly able for up to 57.8Gbps transmissions, but success on 112Gbps PAM4 capable xcvr have already been reported by major manufacturers [4]. Individual research contributions to PAM4 enhancement are also arising as PAM4 adoption increases, focusing mainly on xcvr structure modifications and xcvr equalization alternatives[5], [6].

As PAM4 takes hold, transceivers must be adapted and tested under stress conditions to determine their weaknesses and tolerance to signal impairments. FPGAs (field programmable gate arrays) are commonly used for serial link testing purposes, enabling to generate custom high-speed digital signals patterns to stress the link under test and analyze signal quality/integrity after link distortion. PAM4 capable Intel Stratix 10TX FPGA is used in this work to evaluate PAM4 signal integrity over a series of serial link channels, including QSFP/QSFP-DD cables, SMA 2.4 mm cables and backplane connections(connections over connection board).

1.1. OBJECTIVE

The purpose of this work is to analyze the feasibility of PAM4 transmissions over different links using the E-Tile xcvr of Intel Stratix 10 TX signal development kit FPGA [7]. Evaluation of received signal quality is mainly done in terms of eye-diagram and BER measurements, aiming to determine whether a PAM4 transmission over the channel under test meets pre-FEC BER required by CEI-OIF specification [8].

A testing hw system capable of performing PAM4 signal integrity analysis was designed and implemented in VHDL, adapted to be downloaded to Intel Stratix 10 TX signal development kit FPGA (that includes configurable PAM4 capable xcvr). Intel Stratix 10 TX board includes a microprocessor (Nios-2 microprocessor) that was programmed to configure and control board's xcvr; enabling to:

- optimize equalizer-settings on both TX and RX sides;
- perform parametrized transmission tests over the link;
- collect bitwise transmission error statistics that characterize the signal quality/integrity at receiver side.

In order to test the channel under stress conditions, PRBS31 sequences are used as data source for the signal to be transmitted over the channel. So that, PRBS31 parallel generator and checker were designed and implemented adjusting the design to the needs of the signal evaluation system.

For error statistics collection over the received signal, an error statistical analyzer was also designed and implemented in VHDL. The error statistical analyzer is implemented to be capable of:

- retrieving bitwise error burst length and weight distributions;
- monitoring gap distribution;
- collecting 10-bit symbol based error statistics (by considering the received PRBS31 pattern as sequence of 10-bit symbols with values in $\{0, 1, \dots, 2^{10}-1\}$).

10-bit symbol based error analysis over the received binary sequence is performed to determine the achievable performance after applying RS-FEC(544,5144) Reed-Solomon forward error correction to correct errors occurring in transmissions over the tested link¹.

A preliminary simplified schema of the described PAM4 signal analysis system is shown in figure 5.1 for the sake of clarity.



Figure 1.1.- PAM4 signal integrity analysis system simplified schema including main design units involved in signal generation and processing (note that this schema is a reduced version of the system diagram shown in figure 5.1).

1.2. DOCUMENT STRUCTURE

The remainder of this document is organized as follows: chapter 2 gives a brief introduction to PAM4 signal integrity issues, including a theoretical performance comparison to PAM2 and a description of eye-diagram based signal quality analysis method.

Chapter 3 provides an insight in channel equalization structures.

CEI-OIF specification requirements for PAM4 transmissions are given in chapter 4.

In chapter 5 the developed signal integrity analysis system is depicted, including a functional system division, definition for interfaces involved on HW design and flow description for the link-testing program developed for Nios-2 microprocessor.

Chapter 6 gives detailed description of tests performed, including channel-under-test description; and presents results obtained and transmission feasibility analysis.

Chapters 7.1 and 7.2 conclude this document with a detailed relation on future works and a summary of conclusions extracted from this work.

Further detail on aspects covered in this document can be found in appendixes A, B, C, D.

¹ note that RS-FEC(544,5144) is a non-binary FEC algorithm that works with data blocks of 514 10-bit symbols defined over $GF(2^{10})$ (Galois Field composed of the sub-set of 10-bit binary numbers) So that, the operations required for 10-bit symbol statistics calculation were implemented as operations defined over $GF(2^{10})$ for the error statistical analyzer designed.

2. PAM4 SIGNAL INTEGRITY

This section includes a brief analysis on PAM4 ISI issues and introduces a preliminary comparison with PAM2 in terms of signal distortion suffered on high-speed serial link transmissions (for electrical signals), explaining when PAM4 is expected to outperform PAM2 signaling. For both ISI and signal degradation analysis, that comparison will be based mainly on eye-diagram signal analysis (eye-diagram signal analysis method is described in 2.3), describing the effects of signal jitter and attenuation on eye closure.

Last part of this subsection, 2.3.1.1, describes PAM4 eye anatomy providing exact definition of eye-height and width for the 3-eyed diagram.


2.1. HIGH-SPEED SERIAL-LINK SIGNAL MODULATION

2 modulation schemas are commonly used for high-speed serial-link signals: PAM2(NRZ) and PAM4.² PAMn (pulse-amplitude modulation n-levels) refers to pulsed ('squared') signals that can take n different values (voltage levels), where each level is associated with 1 value in the sub-set $\{0, 1, \dots, n-1\}$. So that, PAMn signals levels represent the possible values for a set of $\log_2(n)$ bits³ referred to as modulation symbols.


PAMn modulation symbols (signal levels) are kept constant for 1UI (see 2.1a,2.1b), being $UI = 1/f_b$ where f_b is the baudrate used in ongoing transmission⁴.

Signal integrity suffers across serial-link paths due to many impairments for both PAM2(a.k.a NRZ) and PAM4 signaling, causing attenuation and signal degradation and thus introducing ISI on both PAMn signals.

ISI impact on PAMn signal integrity depends on both frequency and the number of modulation symbols, 'n'. Therefore, PAM2, PAM4 are not equally affected by signal degradation on serial links.



(a)



(b)

Figure 2.1.- (a) PAM4 , (b) PAM2 baseband example modulated signaling (this figured is copy of A.1, included here for the sake of clarity).

2.2. PAM2 , PAM4 PERFORMANCE COMPARISON

PAM2/PAM4 performance differences for high-frequency serial transmissions lies mainly on signal attenuation (insertion losses) variations with signal frequency, responsible for both frequency-dependent signal level attenuation and high-frequency signal components filtering.

Insertion losses on serial links increase for higher frequencies, meaning that higher frequency signals will suffer from greater channel losses, thus being affected by greater ISI (inter-symbol interference).

For PAMn signals, the attenuation (or filtering) of high-frequencies over serial links introduces smoothness on signal transitions that causes greater ISI on high-frequency PAMn signals. If severe smoothness is introduced, some transitions will not have time enough to reach the 'destination' signal level and transmission errors will occur (i.e., some transmitted symbols will arrive at the receiver with a signal value corresponding to another modulation symbol, thus being detected incorrectly).

PAMn signals' tolerance to frequency-dependent channel attenuation/insertion losses decreases as 'n' increases (i.e. PAMn signals also suffer from

² note that serial-link is used in this document to refer to the data signal path between TX and RX, including both connectors and signal lanes (the later being either backplane channels over connection boards(PCBs) or copper cable channels).

³ for k bits there are 2^k different possible combinations, of which each is associated with 1 signal level; thus if $k = \log_2(n)$, then 2^k results in n different possible combinations and so, in n different signal levels

⁴ as each signal level in PAMn signal represents $\log_2(n)$ bits, bitrate for a given baudrate can be defined as $f = \log_2(n)f_b$, where f_b is the baudrate (bauds/s) and f is bitrate (bit/s)



increasing ISI as 'n' increases), as for higher 'n' values, signal levels (voltage levels associated with modulation symbols) are 'less spaced' (as shown in figure 2.1). As the margin separating modulation levels is reduced, the probability of the channel noise being greater than signal level 'separation' increases, and so does the probability of symbol errors at receiver⁵ (refer to appendix A for thorough analysis on PAM4 integrity).

So that, 2 aspects are considered in PAM2, PAM4 performance comparison :

- at a given baudrate (for the same channel attenuation), PAM4 signal experiences more 'vertical' ISI (ISI introduced by signal voltage level distortion) than a PAM2 signal due to PAM4 'less' distanced signal-levels;
- but channel attenuation (frequency-dependent insertion losses) are lower for PAM4 at a given data rate, as half the PAM2 signal frequency is required to achieve same data rate.

Quantitative analysis demonstrating PAM4 performs better beyond 28GHz is done in A. Results provided quantify the penalty on signal integrity when doubling the transmission data rate by either doubling the frequency of an NRZ signal, or by using PAM4 modulation and keeping signal frequency constant (note that signal integrity penalty is expressed in terms of signal losses, in dB):

- 9.5dB (SNR drop) loss as cost of doubling speed by moving to PAM4 (keeping baudrate)
- 11dB loss as of doubling speed keeping NRZ (thus doubling baudrate = signal frequency for NRZ)

Values given for signal integrity penalties when doubling data rate by either keeping PAM2 signaling or moving to PAM4, reveal that insertion loss for NRZ is greater than SNR-loss penalty associated with PAM4 reduced 'space' between signal levels (9.5 dB < 11 dB), which makes channel kinder to PAM4. So that, though being more vulnerable to signal level attenuation (as separation between signal levels are 1/3 the separation when using NRZ), PAM4's greater resilience to ISI at a given data rate on bandwidth-limited electrical channels is the primary reason for switching from PAM2(NRZ) to PAM4 (see figures A for further info on PAM4 ISI issues).

PAM4 reduces the signal frequency required in half for a given data rate, but ISI associated with insertion losses is not reduced down to 1/2 of the insertion loss affecting the equivalent NRZ signal (considering a PAM2 signal of double frequency in this comparison, for 'constant' data rate). The non-ideal ISI reduction can be explained if considering signal transitions shown in PAM4 eye-diagram (note that in explanation on eye-diagram based signal analysis is given in the following section. PAM2, PAM4 example eye diagrams in figures 2.3a2.3b can help to clarify this brief explanation).

PAM4 symbols would ideally have a 2UI duration (thus eye-diagram would have a 2UI 'horizontal' aperture, being UI the duration of the equivalent NRZ signal at same data rate)⁶, but transitions from/to 'mid'/'inner' signal levels reduces symbol-level duration (thus the eye-aperture) to approximately 1/2UI to 2/3UI (far less than expected 2UI, but still enough to double data density while keeping the eye opened)⁷.

Besides, as ISI introduced by high-frequency components attenuation/filtering (due to channel insertion losses) affects PAM4 signal in transitions (high-frequencies filtering smoothes signal-level transitions), the rate of symbol transitions implying a change in signal voltage level has influence over the probability of PAM4 signal suffering ISI (for consecutive equal PAM4 symbols PAM4 signal level remains constant over consecutive UIs, so the lack of signal-level transition removes the possibility of observing ISI derived from a smoothed transitions). If considering all possible combinations between 2 consecutive PAM4 signal levels, TD (transition density) refers to the rate of combinations for which the consecutive symbols are different, thus implying a change in PAM4 signal level. 2.2 shows all possible symbol transitions for both PAM2, PAM4, from which transition density can be calculated as (num. of transitions between different levels, or symbols)/(total num. of transitions), thus retrieving transition densities of 50 % for PAM2 and 75 % for PAM4:

(for PAM 4, i.e.) 4 possible transitions occurring between same levels are shown in figure 2.2, and 12 possible transitions between different PAM4 levels, so transition density is obtained as:

$$TD = 12/(12 + 4) = 0.75 (75 \%)$$

So that, though suffering less impact on BER when ISI occurs, PAM4 is more prone to suffer from ISI (as greater TD means more probability of ISI occurrence).

2.3. EYE-DIAGRAM BASED ANALYSIS

PAM2 and PAM4 eye-diagram shapes are shown in figure 2.3. An eye diagram is commonly used as an indicator of signal quality for high-speed digital signals and can be used to analyze how channel distortion would affect a high-speed serial transmission.

An eye-diagram is generated from the signal received right at the input of the receiver (thus after equalization structures), by overlapping signal segments of 1UI duration (being 1UI = 1/f_b where f_b is the baudrate, thus lasting the time that 1 symbol signal-level is ideally hold). If a transition between symbol

⁶ double the width of the NRZ eye corresponding to NRZ signaling at same data rate, as PAM4 signal requires 1/2 the NRZ baudrate to keep same data rate

⁷ more detail on this brief analysis is given in A




Figure 2.2.- Possible signal transitions between PAMn modulation symbols for PAM4, PAM2(NRZ) , [1].


levels occurs, 1UI of the received signal is overlapped beginning at the transition (so that transitions will appear at the edges of the 1UI interval, and the center will correspond to the middle of the signal-level ‘step’ representing the modulation symbol being received).

As a result, when enough UI intervals are overlapped, every possible transition between modulation symbols will have been traced over eye-diagram.


If no channel distortion occurred, the measured eye-diagram would have a ‘rectangular’ shape corresponding the squared signal sent, free of distortion. As signal transitions are smoothed through the bandwidth-limited channel, eye-diagrams takes the ‘eyed-shape’ shown in figure 2.3 .

If the received signal is not affected by ISI (i.e. when channel noise does not introduce distortion enough to make signal levels move far from their original symbol-level, getting closer to other symbol-level values), signal levels associated to different modulation symbols can be distinguished at the receiver, generating opened eyes.

If channel distortion cause transmitted signal-levels to get closer to the level corresponding to adjacent symbols (generating errors at the receiver), eye measured will be closed (as signal levels will be displaced from the discrete symbol-levels occupying margins in between).



(a)



(b)

Figure 2.3.- (a) PAM2 eye-diagram (measured over $2UI_{NRZ}$ interval) ; (b) PAM4 eye-diagram (measured over $1UI_{PAM4}$ interval) (note that $UI = 1/f_b$ refers to baudrate of each modulated signal , thus for a given datarate $1UI_{PAM4} = 2UI_{NRZ}$) , [1].

PAM4 signal analysis using eye-diagram is similar to PAM2/NRZ eye-based analysis, by applying the analysis separately to each of the PAM4’s three-eye diagrams (note that PAM4 inner eyes’ width/height definitions differs slightly from PAM2’s , definition and measurement procedure are given in 2.3.1).

PAM4 eye diagram spans one unit interval (UI) and consists of a single ‘outer eye’ (that corresponds to the eye traced by transitions between higher – lower signal-levels/modulation-symbols, thus between PAM2 symbol-levels), with three ‘inner’ eyes. Inner eye diagrams are referred to as low, middle, and upper eyes (sub-section 2.3.1 provides formal definition for PAM4 inner eyes’ parameters).

As each of the three eyes is analyzed the way that PAM2 eyes are analyzed, EW (eye width) and EH (eye height) are measured separately for the lower, middle and upper eye diagrams. Thus, if observing the figure 2.3 it can be seen that inner eyes’ width (marked in red) does not match the expected

$2UI_{NRZ}$ width⁸.

Because PAM4 has 4 voltage levels, there are transitions between non-adjacent signal levels that take longer time than required for transitions between adjacent levels; so that reducing the ideal $2xUI$ eye-width down to between $1/2 UI_{NRZ}$ and $2/3 UI_{NRZ}$ as introduced in 2.2, [1].

2.3.1. PAM4 3-EYED DIAGRAM FORMAL PARAMETERS

Formal definition for PAM4 3-eyed diagram is giving in this section (as defined in CEI-OIF specification [8]), taking eye-diagrams in figure 2.4 as reference for parameter definitions given.

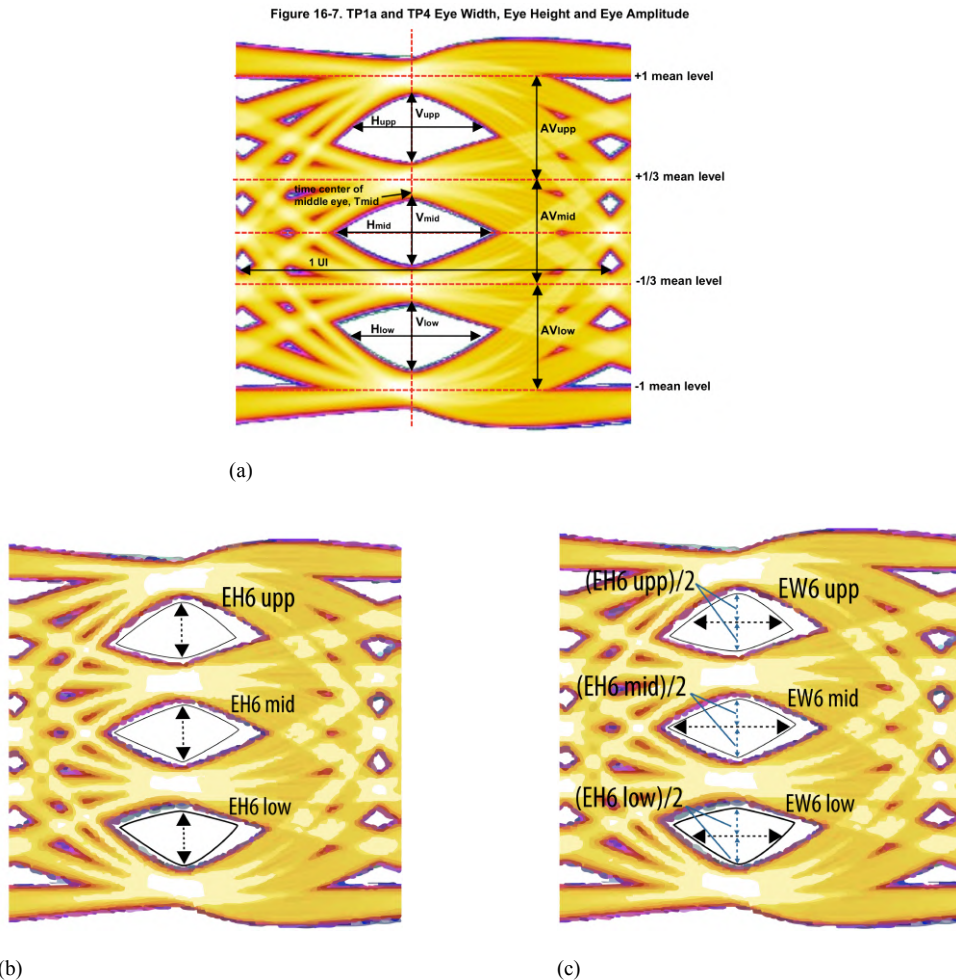


Figure 2.4.- (a) PAM4 eye parameter measuring definition (parameters used in EH_6, EW_6 measurement over statistical eye-diagram defined) [8]; (b) eye-height measured over BER 1E-6 contour over statistical eye-diagram, EH_6 [2]; (c) eye-width measured over BER 1E-6 contour over statistical eye-diagram, EW_6 (from EH_6 determined) [2].

PAM4 eye-diagram relevant parameters (EW_6 and EH_6) are determined from the eye contour related to BER value of 1E-6 (being BER = 1E-6 the value set by CEI-OIF specification as the general maximum allowable BER for PAM4 high-speed serial transmissions when no FEC algorithm is used)⁹.

CEI-OIF specification [8] defines an 'statistical' eye-diagram consisting of a series of eye contours (eye-shaped contours delimiting eye-apertures), where each contour serves as a 'mask' indicating the aperture for each eye that guarantees that BER measures will be below a certain value for PAM4 high-speed serial transmissions showing an eye-aperture greater than the given eye-mask. Therefore, the greater the eye aperture at the 1E-6 contour, the

⁸ as PAM4 requires half the NRZ frequency to achieve same data rate, PAM4 symbols will hold $2xUI$ being $UI = 1/f_b$ where f_b is the NRZ baudrate, thus the time that each NRZ symbol is hold

⁹ note that CEI-OIF contour correspond actually to SER limits; but when a symbol error occurs the sent symbol is most frequently confused with an adjacent one, so that if assuming gray-coding is used (adjacent symbols differs in only 1 bit) $BER \approx SER$



greater channel noise can be assumed 'before' the eye gets closed.

EW6,EH6 values retrieved from Stratix 10TX on-die instrumentation in tests performed are also measured as defined in CEI-OIF specifications[8] . (EW6 and EH6) values for the 3 inner eyes (referred to as $EW6_{low},EH6_{low}$, $EW6_{mid},EH6_{mid}$, $EW6_{upp},EH6_{upp}$ for lower , mid and upper eyes respectively in figure 2.4) are measured as :

1. mid-point of the greatest aperture in mid eye is located, named t_{mid} .
2. vertical line is traced over t_{mid} .
3. for each eye, EH6 is defined as vertical space over that vertical line, between the points where that line intersects eye's 1E-6 contour.
4. an horizontal line is traced over the mid-point of the EH6 trace. EW6 is defined as the 'space' over that horizontal line, between the points where that line intersects eye's 1E-6 contour¹⁰

Since each PAM4 eye is analyzed 'independently' and system's BER is limited by the smallest eye-aperture, CEI-OIF specification defines minimum aperture values for the parameters :

$$EW6 = \min(EW6_{low}, EW6_{mid}, EW6_{upp}) ;$$

$$EH6 = \min(EH6_{low}, EH6_{mid}, EH6_{upp}).$$

(EW6 and EH6) values measured at receiver input (thus after RX-equalization) should be at least $EW6 \geq 0.2$ UI and $EH6 \geq 30$ mV (for 18 GBs –29 GBs signals) for a transmission to be considered successful [8].

2.3.1.1 EYE ANATOMY

A PAM4 eye diagram is shown in figure 2.5 divided in parts to provide insight on how different transitions affect eyes' aperture.

PAM4 has 64 different combinations for 3 consecutive symbols¹¹ , from which NRZ-like combinations (shown in 2.5b.3) that compose the 'outer-eye' are much well behaved.

Series of pieces from PAM4 eye show in figure 2.5, reflect PAM4 sensitivity to noise amplitude caused by the reduced space between multiple symbol-signal-levels (figures 2.5a.1 – 2.5a.3 corresponding to transitions between adjacent levels that form 'inner' eyes, show smallest eye-apertures thus limiting PAM4 ISI tolerance).

¹⁰ that defining eyes' heights with respect to the same point in time, t_{mid} , it is assumed that 3 eyes will be sampled at receiver simultaneously; thus keeping compatibility even with early PAM4 receivers that did not support flexibility on sampling instants to accommodate to eyes misalignment, for the sake of receiver simplicity

¹¹ for each signal segment of 1UI duration traced over the eye diagram , 3 consecutive symbols are involved as previous and post symbols determine the transitions observed at UI window sides

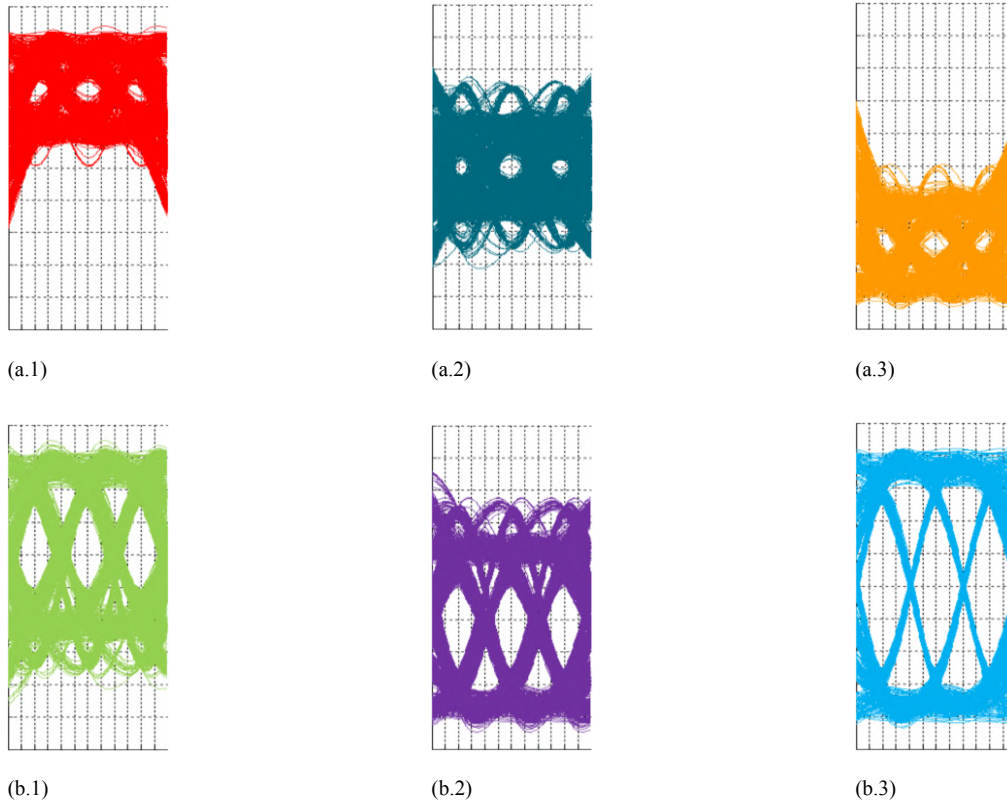


Figure 2.5.- (a.1) PAM4 upper inner eye (transitions $+3 \leftrightarrow +1$); (a.2) PAM4 mid inner eye (transitions $+1 -1$); (a.3) PAM4 lower inner eye (transitions $-1 -3$)¹²; (b.1) PAM4 eye transitions $+3 -1$; (b.2) PAM4 eye transitions $+1 -3$ ¹³; (c) PAM4 outer eye (transitions $+3 -3$)^{14,15} (note that eye pieces are shown for 3UI intervals), [1].

2.3.2. ISI IMPACT ON EYE DISTORTION

2.6a.1 - 2.6a.4 series of figures reflect the effect of insertion losses over eye diagram aperture for PAM2 signal¹⁶, showing how limited bandwidth channels (introducing smoothness in transitions, thus 'horizontal eye-closure') and attenuation/noise (affecting signal level, thus causing 'vertical eye-closure') degrade signal integrity (the former affecting more PAM2: as for a given datarate NRZ requires 2x the PAM4 signal frequency, so symbols are held half the time (PAM4 wider-in-time pulses suffer less ISI); and the latter affecting PAM4: as space between signal levels is 1/3 the space between NRZ's levels, thus being more vulnerable to signal level drops or boosts).

2.6b.1-2.6b.3 shows eye-diagram enclosure for increasing signal frequency (i.e. for increasing channel attenuation, [9]).

¹² (a.1)-(a.3) correspond to transitions between adjacent symbols

¹³ (b.1)-(b.3) correspond to transitions between 1-level separated symbols

¹⁴ (c) corresponds to transitions between 2-level separated symbols

¹⁵ being $-3, -1, +1, +3$ PAM4 modulation symbols

¹⁶ effect of insertion losses over eye-diagram 'horizontal' closure is shown for a PAM2 signal for simplicity, but same concept can be applied individually for each of the 3 eyes in PAM4 eye-diagram


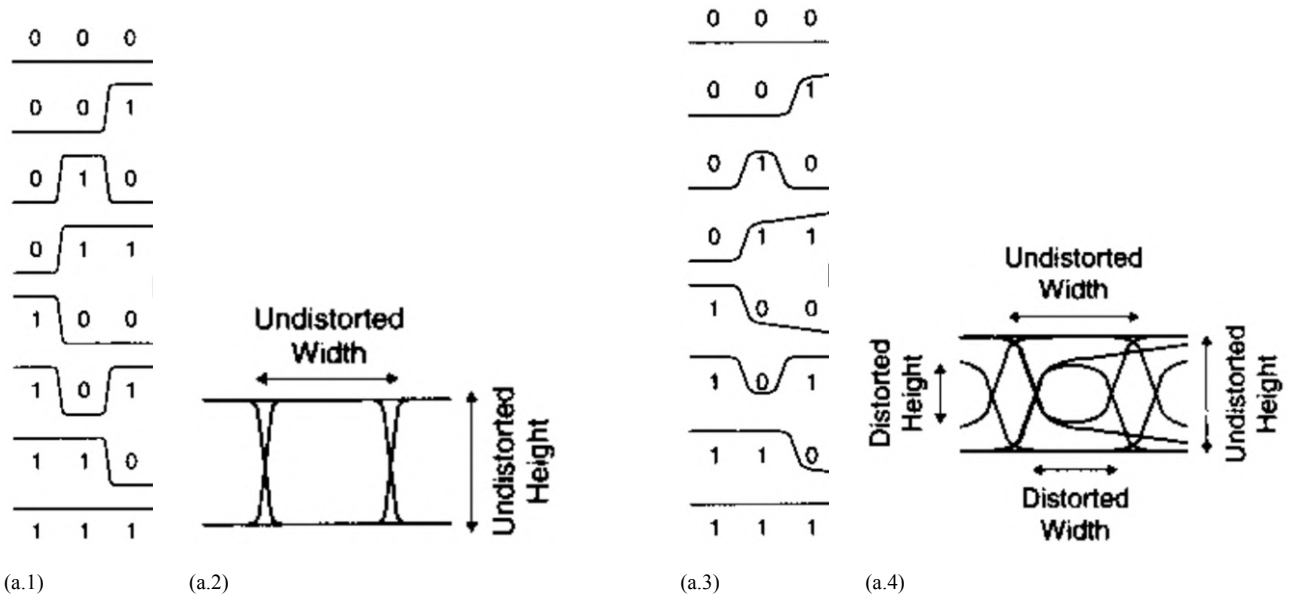


Figure 2.6.- (a.1) Possible NRZ undistorted signal transitions for any combination of 3 consecutive symbols; (a.2) eye-diagram for undistorted signal transitions overlapping; (a.3) NRZ distorted signal transitions (smoothed transitions); (a.4) eye-diagram for distorted signal transitions overlapping, [10]; (b.1) example low distortion eye-diagram ($f_b = f_1$, $f_b = \text{baudrate}$); (b.2) mid distortion eye-diagram ($f_b = f_2$, $2f_1 = f_2$); (b.3) mid-highdistortion eye-diagram ($f_b = f_3$, $4f_1 = f_3$) [9].

3. EQUALIZATION PARAMETERS

PAM4 reduced signal-levels spacing (1/3 of the NRZ's separation) and non-ideal eye-width (reduced down to 1/2 UI – 2/3 UI due to the ISI caused by low-pass nature of the channel), introduce the need for equalization to minimize ISI suffered in serial link transmissions .

Equalization aims to invert signal distortion suffered in serial links so as to recover sent signal shape at the receiver, thus removing ISI introduced in the transmission (see figure 2.6a.3). If observing the impact of channel distortion over eye diagram (detailed in sub-section 2.3.2), smoothness in transitions can be explained as an 'attenuation' affecting the 1st bit after the transition; so that equalization will aim to boost that 1st bit after transitions (alleviating smoothness) to speed-up transitions and so, open the eyes.

This section describes equalization techniques that aim to compensate for both channel response (where ISI compensation lies) and signal degradation (attenuation and jitter, caused by other effects: crosstalk , insertion loss,return loss etc), figure 3.1 (copy of A.9 included here for the sake of clarity) shows the equalization structure used in Intel Stratix 10TX E-Tile xcvs¹⁷ .




Figure 3.1.- Generic xcvr signal path including TX FIR filter-based equalizer , CTLE+DFE RX equalizer (note that in real xcvr structure signal modulator and demodulator would be placed right at the input of TX-equalizer and output of RX-CTLE).

3.1. TX EQUALIZATION

TX-side equalizer is implemented using a FIR filter structure in Stratix 10 E-tile xcvs, and referred to as de-emphasis equalizer.

FIR filter-based de-emphasis equalizer aims to invert channel distortion before it happens by subtracting the ISI that will be introduced in the channel (thus the influence of adjacent symbols in the sequence being transmitted). So that ISI introduced by the channel will be cancelled by de-emphasis pre-distortion, and signal will arrive at the receiver with reduced or null distortion.

In the FIR filter the signal sent is pre-distorted (as shown in image 3.2c) such that when the pre-distorted signal is passes over the channel, 'adding the channel response effect' results into pre-distortion cancelation instead of signal degradation, thus recovering the shape of the sent signal at RX-side. Signal in figure 3.2d reflects the recovery of the desired eye-shape from a pre-distorted signal sent over channel; while not-pre-distorted signal (non.-equalized signal, as shown in figure 3.2a) arrives at RX affected by ISI (reflected in smudgy, blurred eye-diagram shown in figure 3.2b)¹⁸ .

De-emphasis pre-distortion effect is summarized in figure 3.3 where signal progression when no pre-distortion is applied is shown at the top; and bottom signal progression shows the effect of sending a pre-distorted signal through the channel when de-emphasis FIR filter is adjusted to compensate channel distortion.

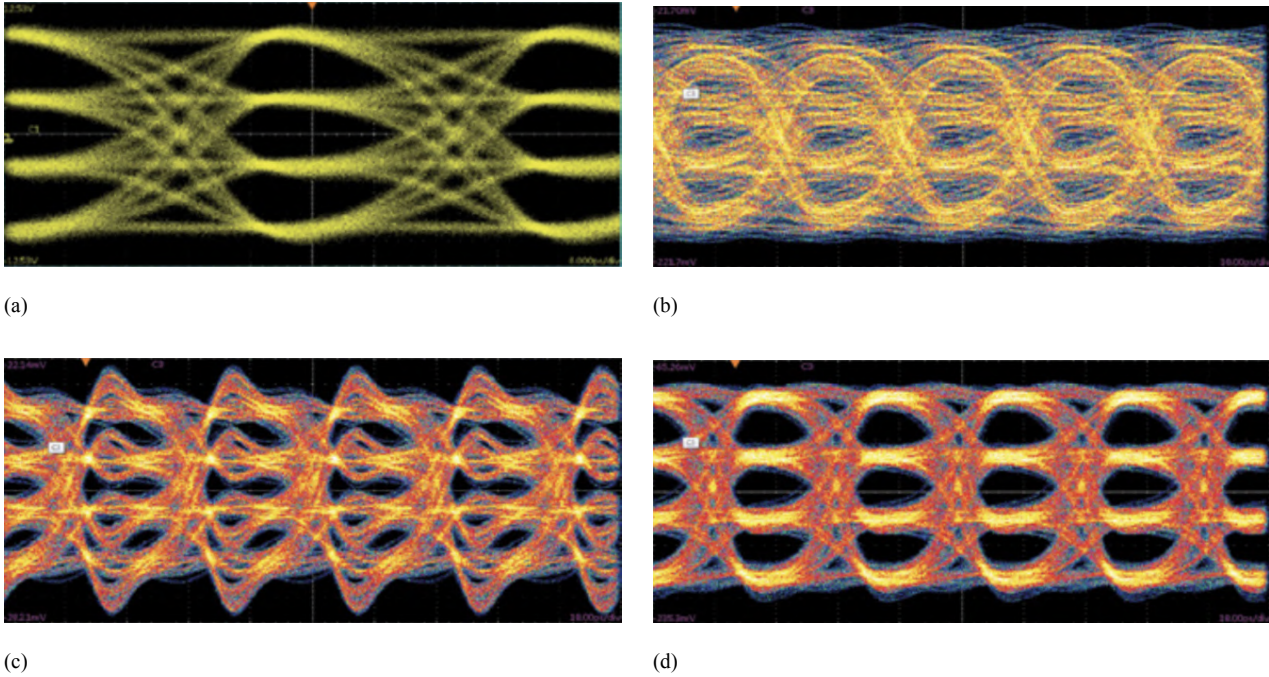


Figure 3.2.- (a) PAM4 eye-diagram measured at TX-output (right before insertion in channel) when no pre-distortion is applied (de-emphasis not used, thus TX signal not pre-distorted); (b) PAM4 eye-diagram measured at RX-input (right after channel) when no pre-distortion is applied (de-emphasis not used, thus RX distortion not compensated); (c) PAM4 eye-diagram measured at TX-output when pre-distortion is applied (de-emphasis pre-distorts signal to invert channel distortion); (d) PAM4 eye-diagram measured at RX-input when pre-distortion is applied (de-emphasis used, thus channel inverts pre-distortion and RX signal eyes re-open), [10].




Figure 3.3.- PAM4 signal distortion progression over serial channel when both no TX-equalization is applied, and de-emphasis pre-distortion is applied over transmitted signal.

3.1.1. FIR FILTER BASED TX-EQUALIZER

Intel Stratix 10TX uses FFE(feed forward equalizer) FIR filter implementation for de-emphasis equalization [11].

FIR filters consist of a series of delay elements used to buffer a limited number of consecutive signal symbols; each signal symbol is weighted by the

¹⁷ note that though many techniques exists this section aims to provide a brief understanding of only the TX/RX equalization structures used in Stratix 10TX E-Tile xcvr; thorough explanation on signal degradation and ISI compensation are given in A

¹⁸ figure 3.2 is a copy of figure A.12 included here for the sake of clarity

corresponding coefficient in the FIR filter structure as shown in figure 3.4¹⁹, where those coefficients $c_{-n} \dots c_0 \dots c_m$ are adjusted to reduce ISI.

In FIR filter schema shown in figure 3.4 c_j coefficient (where $j \neq 0$, assuming $j=0$ is the coefficient applied to the symbol being transmitted) determines the influence of the symbol j over the signal level transmitted when symbol s_0 is being sent (coefficients c_j with $j > 0$ correspond to previous symbols; and c_0 with $j < 0$ to later symbols)²⁰

So that if the serial channel introduces ISI such that symbol s_0 is received as, e.g.: $s_0 + c_1 s_{-1}$ (where $s_{-1} = s(-1)$, see figure 3.4); when TX equalizer 'transmits' $s_0 - c_1 s_{-1}$ instead of s_0 , the receiver gets $(s_0 - c_1 s_{-1}) + c_1 s_{-1} = s_0$, thus recovering the sent symbol without distortion (thorough explanation on FIR filter TX-equalization is provided in A).

Intel Stratix 10 implements FFE de-emphasis equalizer using 3-tapped FIR filter [11], thus $c_{-n} \dots c_0 \dots c_m = c_{-1}, c_0, c_1$




Figure 3.4.- Common structure for k -coefficients ($c_{-1}, c_0, c_1, \dots, c_{k-2}$) FIR filter with $k = m+2$ (note that Δ blocks are 1UI delay).

3.2. RX EQUALIZATION

A '2 stage' CTLE + DFE equalization structure (as shown in figure 3.1) is depicted in this section, as is the structure used for RX-side equalization in Stratix 10TX E-Tile xcvr (note that this section aims to provide a slight insight to ease the understanding of the configurable equalization parameters exposed by PAM4 xcvr integrated in Intel Stratix 10TX board; for further detail on its working principle, see A.3.3).

3.2.1. CONTINUOUS TIME LINEAR EQUALIZATION, CTLE

CTLE (continuous-time linear equalizer) equalizer is placed right at the input of RX and boosts high frequency components of electrical waveforms to invert the low pass filter effects of the channel:

- By boosting high frequency components, CTLE 'recovers' faster signal transitions; thus partially removing the smoothness introduced by BW limited channel, and so reducing the 'horizontal' ISI caused by longer transition times (smoothed transitions).
- To boost high-frequency components CTLE equalizers have a frequency response with the shape shown in figure 3.5 (note that figure 3.5 is a copy of figure A.17 included here for the sake of clarity).

CTLE frequency response's planar segment causes the low frequency signal components to be attenuated in comparison to higher frequencies (weighted by the 'peak' in frequency response as signal passes through CTLE equalizer), thus compensating the higher attenuation suffered by higher-frequencies in serial link. CTLE gain in planar region (low frequency gain) can be optimized as a parameter to achieve maximum eye openings at receiver input

3.2.2. DIRECT FEEDBACK EQUALIZATION, DFE

DFE (direct-feedback equalizer) is composed by a decisor (a.k.a. slicer) whose output feeds a 'm-tapped' FIR filter as shown in figure 3.6 (figure 3.6 is a copy of figure A.20 included here for the sake of clarity).

¹⁹ figure 3.4 is a copy of figure A.10 included here for the sake of clarity

²⁰ note that for digital signals, discrete signal analysis is applied, thus symbol sent over 1UI is represented by sample $s(n)$ where $n=0$ is associated to symbol being transmitted and $n \neq 0$ with adjacent symbols (signal samples take 1UI apart; thus at $1/f_b$, where f_b =baudrate), such that $n > 0$ correspond to later symbols, and $n < 0$ to previous.




Figure 83E-10—Selectable continuous time linear equalizer (CTLE) characteristic

Figure 3.5.- CTLE gain (transfer function) shape for RX-equalization (from CTLE characteristic definition in [12].)

CTLE equalizer's output (corrected by DFE's FIR filter as shown in figure 3.6), $\hat{r}(n)$, is introduced into a decider that determines which modulation symbol corresponds to the signal value at its input (thus, recovering the ideal symbol-voltage from an input signal sample whose voltage level may be distorted). So that, if no error occurred during the serial transmission, $\hat{s}(n) = s(n)$ (as shown in figure 3.6), where $s(n)$ is the symbol transmitted corresponding to the received signal sample at DFE's input, $r'(n)$.

DFE's FIR filter operates on the same basis as the FIR filter in TX de-emphasis equalizer, but subtracting only the influence of symbols received previously, as later symbols cannot be anticipated. So that, if FIR filter's coefficients are well adjusted, the remaining channel distortion affecting samples at DFE's input can be removed, thus eliminating residual ISI.




Figure 3.6.- DFE equalizer internal structure using slicer + feedback FIR filter architecture (note that DFE's input $r'(n)$ is the received signal after CTLE equalization).

DFE uses the values of already demodulated symbols, $\hat{s}(n)$, to correct the incoming symbol (assuming there was no error on previous symbols' decisions). If one decision-error occurs decisor's output $\hat{s}(n)$ does not correspond to the symbol sent by TX, $s(n)$, thus $\hat{s}(n)$ value cannot be used to remove $s(n)$'s influence on later symbols. So that, one symbol error on DFE slicer causes later $r'(n)$ samples to be 'modified' incorrectly (distorted rather than corrected), thus possibly turning out into another decision error, and ultimately generating an error burst at DFE's output²¹.

²¹ note that generally DFE error bursts can be accommodated using RS(544,514) FEC, that can correct bursts up to 30 10-bit symbols, or more when interleaving is applied (D.1.1.2 provides detailed explanation on how to use RS-FEC(544,514) to address DFE's error bursts issue.)

4. CEI-OIF 4.0

CEI-OIF-56G-PAM4 included in CEI-OIF 4.0 release [8] is baseline specification for PAM4 high-speed serial link transmissions.

CEI-OIF 4.0 classifies reference channel models (described in 4.2 section), and provides both high-bounds for BER and occurrence-frequency for bursts of max correctable length (note that FEC use is assumed for PAM4 serial transmissions), and guidance on application of interleaving (refer to D.1.1.2 sub-section for more detail on interleaving techniques) and other techniques to achieve compliant PAM4 transmission.²²

4.1. CEI-OIF XCVR REQUIREMENTS


CEI-OIF specifies both electrical and functional requirements for PAM4 xcvs operating above 14GBs:

- Minimum functional structure for both TX and RX (as shown in figure 4.1 ,extracted from appendixes 16.C.1 and 16.C.2 on [8], respectively).
- PAM4 signalling shall map modulation symbols to a 4-symbol gray-code²³ .
- RS-FEC(544,514) must be supported.
- QPRBS31 patterns shall be used for serial link test (PRBS patterns are used to stress channels under test, C provides further detail on use of PRBS31 pattern and its generation/verification)²⁴ .
- supported baudrates , link length , pre-FEC and post-FEC BER limits are given for each channel reference model

4.1a shows how transmitter process shall map consecutive pairs of bits {A, B} (where A is the bit arriving first) to its corresponding Gray-coded signal level (being modulation signal levels denoted as {0, 1, 2,3 }).


4.1b show how receiver rprocess shall map Gray-coded PAM4 symbols (signal levels) to pairs of bits {A, B} (where A is considered to be the first bit).

Figure 16-17. Transmit signaling and mapping diagram



(a)

Figure 16-18. Receive signaling and mapping diagram



(b)

Figure 4.1.- (a) Transmit signaling and mapping diagram, minimum functional requirements for PAM4 transmitter; (b) receive signaling and mapping diagram, minimum functional requirements for PAM4 receiver, [8].

²² note that this section does not include complete CEI-OIF requirements for PAM4 channels above 14GBs, limiting its contents to the aspect required to ease understanding or support assumptions/statements referring CEI-OIF specifications.

²³ gray codes re-arrange modulation symbols into a set where adjacent symbols differ only in 1 bit; so that if signal-level is distorted getting close enough to an adjacent symbol-level, only 1 bit error occurs. PAM4 signal is gray-coded by mapping each modulation symbol to its 'level-equivalent' in gray-code symbol-set (as shown in figure 4.2). PAM4 symbols will be replaced by their associated symbol in the gray-code before transmission

²⁴ note that besides test data patterns, further electrical requirements and measurement methods are described in CEI-OIF specifications



<p>{0, 0} maps to 0, {0, 1} maps to 1, {1, 1} maps to 2, and {1, 0} maps to 3.</p>	<p>0 maps to {0, 0}, 1 maps to {0, 1}, 2 maps to {1, 1}, and 3 maps to {1, 0}.</p>
--	--

(a)

(b)

Figure 4.2.- (a) Gray mapping at PAM4 encoded lanes for TX output lanes; (b) and RX-output lanes, where PAM4 symbols are referred to as 0, 1, 2, 3 (notation equivalent to -3, -1, +1, +3, respectively), [13].

4.2. CHANNEL REFERENCE MODELS

3 reference channel models are distinguished for PAM4 transmissions up to 56 Gbps (specific requirements are given for each channel reference model):

CEI-56G-PAM4-VSR (very short range)

- chip-to-module interfaces (C2M)
- <10cm , 1 connector

CEI-56G-PAM4-MR (medium range)

- chip-to-chip and midrange backplane C2M
- <50cm , 1 connector

CEI-56G-PAM4-LR (long range)²⁵

- chip-to-chip over backplane or copper cable
- <100cm , 2 connectors
- up to 31dB loss at 14GHz
- supported baudrate fb in 18.0Gsym/s – 29.0Gsym/s
- minimum BER of 1E-4 , assuming FEC to be used to achieve FEC <1E-15

4.3 shows the insertion loss mask indicating the maximum allowed insertion losses for CEI-56G-PAM4-LR compliant channels, considered as a reference for PAM4 analysis.




Figure 4.3.- Informative channel insertion-loss limit-mask for channel compliance at 29.0 Gsym/s , [13].

²⁵ note that only CEI-56G-PAM4-LR channels are completely specified here as only CEI-56G-PAM4-LR channels were tested



5. HIGH-DATARATE PAM4 SERIAL LINK ANALYSIS SYSTEM

This section describes the PAM4 high-datarate serial link analysis system developed.

The hw system designed is capable of both NRZ and PAM4 transmissions at 14 - 28GBs baudrates and supports configuration and control from both design logic and Intel's debugging tools (so that enabling on-die eye-diagram view, refer to section 5.5 for further detail on use of debugging tools).

In section 5.1 system's functional division is included, hierarchically organized in 3 division levels to ease the explanation provided; note that for the sake of simplicity data interfaces between hw entities and clocking domains are described in sub-sequent sections.

Section 5.6 describes a console application developed to enable to configure and run PAM4/NRZ transmissions for link testing over the hw system designed, limiting each transmission test either by the elapsed time or by the number of bits sent, and collecting statistical error data and test status info. during test duration.

Interaction with the described system is available through a Windows embedded Ubuntu console, from which gathered test data can be retrieved for further analysis.

5.1. FUNCTIONAL DIVISION

An overview of the PAM4 serial link analysis system developed is shown in figure 5.1, where wine-coloured blocks correspond to VHDL entities designed from scratch and blue one's are Intel's official ipcores²⁶. System described is controlled from grey-coloured block, which includes a microprocessor on charge of configuring the design logic and collecting data.²⁷

'native-phy-ip e-tile xcvr' is the xcvr controller that provides abstraction over medium access and physical layers, exposing interfaces to input data to be transmitted and output received data. So that, serial link is established between board's connectors (marked as 'external loopback' in figure 5.1) but both transmitted data and received data are handled in VHDL design by an unique 'native-phy-ip e-tile xcvr' entity.²⁸

In figure 5.1, series of wine-colored blocks at the bottom simplifies the datapath within the designed system, flowing left-to-right. PRBS31 test pattern is used as required by CEI-OIF 4.0 specification for PAM4 serial link testing (see appendix C for info on PRBS31 pattern generation). PRBS31 test pattern generator output is connected as input to the xcvr controller that will send the test pattern over a physical lane (serial link) connecting to a controlled receiver. So that, the same xcvr controller receive the test pattern after the serial link and output the received pattern to the PRBS31 checker, on charge of verifying the received sequence detecting mismatches with the known PRBS31 pattern to locate errored bits in case of transmission errors had occurred.

Subsequent sections provide further detail on the function of each entity in the design.

5.1.1. 1ST LEVEL DIVISION

Fig. 5.2 shows 1st level functional division for the conceptual system design reflected in figure 5.1:

1. Qsys-embedded system:
exerts control over FPGA design, controls xcvr configuration sequences, collects transmission error data and transmission status info., and enables to run configurable controlled transmission tests.
2. PAM4 testing FPGA design:
hw design that handles test pattern generation and verification, analyzes the error-pattern detected over received signal, provides access to xcvs on Stratix 10 TX board and manages clocking generation and distribution.

²⁶ ipcores refer to tested VHDL entities for controlling specific purpose hard-coded hw modules included in Stratix 10TX board

²⁷ note that 'grey-colored' control entity is developed as an independent embedded system, integrated in the system designed using Altera's Avalon interfaces (Altera's own interface protocols)

²⁸ note that Stratix 10TX FPGA includes 5 PAM4 capable xcvs, each controlling up to 24 physical lanes (each physical lane includes 1 TX and 1 RX, thus each channel can be understood as 1 complete xcvr itself). System designed uses only 2 out of 5 available xcvs (referred to as PHY: PHY0 - PHY4, more detail on Stratix 10TX structure is provided in appendix B), and one 'native-phy-ip e-tile xcvr' controller entity is required to control each PAM4 xcvr (only 1 'native-phy-ip e-tile xcvr' controller is represented in figure 5.1 for the sake of simplicity).

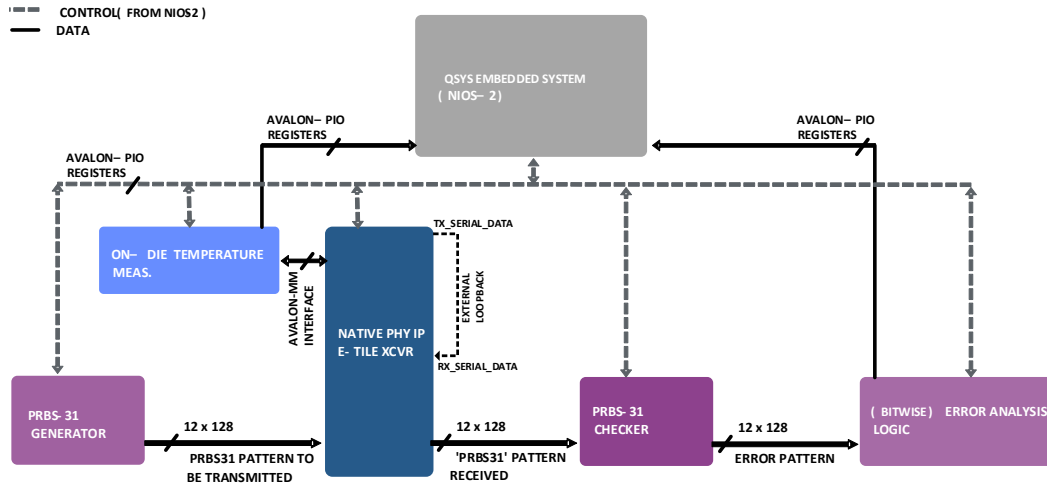


Figure 5.1.- PAM4 serial link testing system simplified funtional/conceptual diagram.

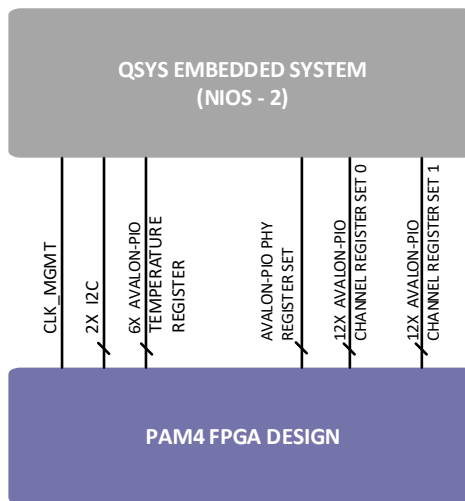



Figure 5.2.- PAM4 serial link testing system 1st level functional division.

5.1.2. 2ND LEVEL DIVISION

5.1.2.1 μ PROCESSOR EMBEDED SUB-SYSTEM (QSYS-EMBEDDED SYSTEM)

Fig. 5.3 shows a 2nd level functional division diagram for the embedded system referred to as 'qsys-embedded system', where Nios-2 is the microprocessor included in Stratix 10TX board and the rest of the elements aim to provide interface to/from Nios-2 microprocessor:

1. Nios-2 microprocessor:
its main function is to control configuration sequences over the configuration interface exposed by each xcvr controller. Nios-2 can be accessed from an ubuntu console and supports interaction through console session while design is running, so that allowing for c/c++ app. execution (more info. on Nios-2 c++ app. developed to enable for dynamic xcvr configuration, test control parameters configuration and transmission status monitoring is given in section 5.6).
2. JTAG-UART bridge:
serves as a bridge between a JTAG connection (used to connect Intel Stratix 10TX board to Nios-2 console host) and an Avalon-mm interface exposed by Nios-2 microprocessor (refer to sub-section 5.2.2 for description of Intel's Avalon-mm interface).



3. I2C master:

serves as a bridge between the Avalon-mm interface exposed by Nios-2 microprocessor and an I2C²⁹ master agent to support communication with I2C modules in the QSFP-DD connectors integrated on Stratix 10TX board (see 5.2.1 for detail on connectors used to perform serial-link external-loopback connections over Stratix 10TX board).

4. 'phy-register-set':

set of Avalon-PIO (parallel input/output) registers³⁰ that can be understood as the interface for sharing data and control info. with FPGA design.

Besides the set of registers, 'phy-register-set' includes an Avalon-mm master bridge that exports an Avalon-mm slave interface for connection to Nios-2 (connection to Nios-2's Avalon-mm master) on one side; and an Avalon-mm master to forward master signaling received from Nios-2's master to the Avalon-mm slave interfaces on each Avalon-PIO (as shown in figure 5.3)³¹

'phy-register-set' entity is defined using tcl script parametrized to instantiate several Avalon-PIO regs per PHY used in design (shown in figure 5.3):

reconfig_mgmt_phy_0_rcnfg	Avalon-PIO registers for each signal involved in Avalon-mm interfaces for native-phy-ip reconfiguration registered signals are connected to the Avalon-mm slave interface exposed by the corresponding PHY controller ('native-phy-ip e-tile xcvr'), through which xcvr configuration and control is enabled
phy_set_bitrate_reg_0_export	(FPGA to Nios-2) holds measured bitrate
phy_set_RXclock_reg_0_export	(FPGA to Nios-2) holds measured frequency for clock recovered form received signal (CDR)
phy_set_counter_1ms_reg_0_export	(FPGA to Nios-2) holds measured milliseconds elapsed from xcvr start
phy_set_control_reg_0_export	(Nios-2 to FPGA) reset and enable control signals
phy_set_control2_reg_0_export	(FPGA to Nios-2) channel selection control (note that each PHY controls up to 12 channels in system designed, so that channel selection mask is used to enable simultaneous operations on different channels)
phy_set_tst_active_reg_0_export	(FPGA to Nios-2) 12-b width test control info register (for each channel indicates whether a test is running over channel or not)
phy_set_tst_rset_cntrl_0_export	(FPGA to Nios-2) 12-b width test control info register
phy_set_tst_en_cntrl2_0_export	(FPGA to Nios-2) 12-b width test control info register

5. 'channel-register-set':

set of Avalon-PIO (parallel input/output) registers for channel related info interfacing (so that, 1 'channel-register-set' is instanced per PHY used in design; and 'channel-register-set' tcl definition is parametrized to place 12 instances of each Avalon-PIO register defined in 'channel-register-set', 1 per used channel in each PHY)

channel_reg_set_0_channel_0	(FPGA to Nios-2) 16-b width holding link status info
channel_reg_set_0_error_count_h_0	(FPGA to Nios-2) 32 upper bits of error counter
channel_reg_set_0_error_count_l_0	(FPGA to Nios-2) 32 lower bits of error counter
channel_reg_set_0_tstmer_reg_0	(FPGA to Nios-2) 32-b width holding msecs elapsed from test start

5.1.2.2 FPGA DESIGN

Fig. 5.4 shows a simplified schema for the functional division of the implemented FPGA design (at 2nd-level of system's functional division) that includes all the hw modules required to support the designed PAM4 serial link testing system (thus providing definition for the hw sub-system on charge of conducting and processing high-speed serial transmissions).

FPA design sub-system includes the resources required for both high-speed serial signal generation at TX-side and error-analysis signal processing at

²⁹ I2C (Inter-Integrated Circuit) master-slave serial bus protocol

³⁰ Avalon-PIO registers are commonly used as interface from/to Nios-2 microprocessor, further detail is given in 5.2.3

³¹ note that Avalon-mm slaves on Avalon-PIO are all connected to the Avalon-mm master in the bridge (thus to Nios-2's Avalon-mm master) using same shared buses

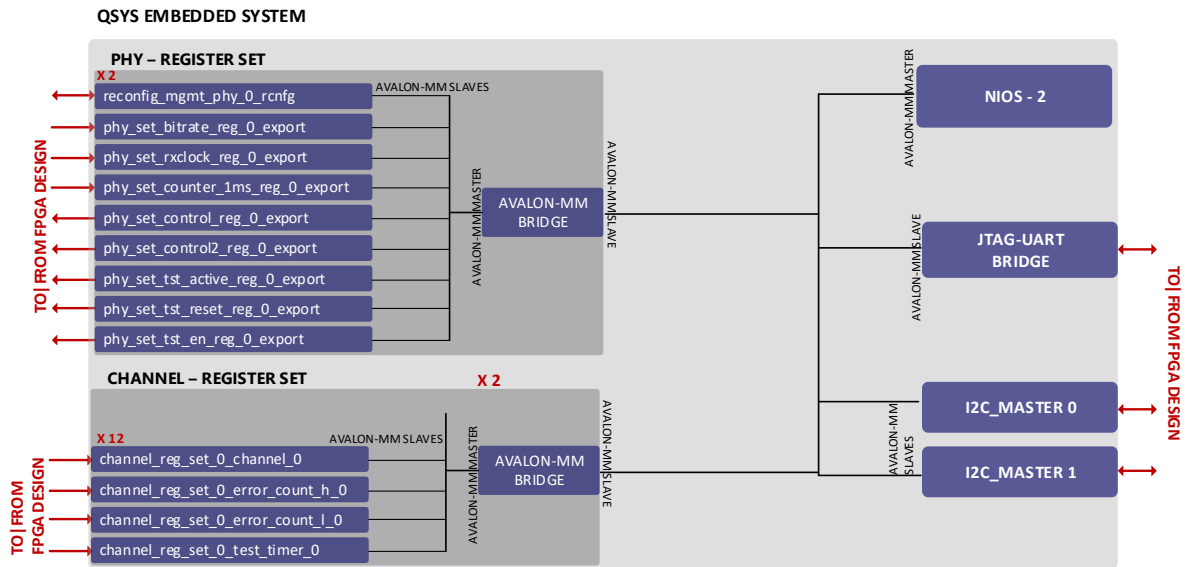


Figure 5.3.- PAM4 serial link testing system 2nd level functional division Qsys-embedded system.

TX-side (signal treatment related modules are wine-colored in figure 5.4, except for the dark-blue block denoted ‘native-phy e-tile xcvr’ representing Intel’s xcvr controller ipcore, that handles signal modulation/demodulation); as well as sub-system control resources.

FPGA sub-system control resources also include hw modules designed/integrated to support:

- sub-system synchronization;
- xcvr status control (including xcvr temperature tracking;
- transmission status monitoring;
- signal frequency control) and configuration, and serial-transmission-test time control.

In addition to the resources that support the serial link testing implementation, FPGA design exposes interfaces used to exchange data with Nios-II microprocessor and receive re-configuration and control requests.

In figure 5.4 the block denoted ‘phy-controller’ includes the resources required to control the 12 xcvr channels used on each used E-tile xcvr (each PHY). So that, blocks marked with red ‘x12’ tag are instantiated 12 times, 1 per channel used. Similarly, 2 instances of the ‘phy-controller’ block are placed in FPGA design (marked with red ‘2x’ tag), 1 for each of the 2 E-tile xcvs used in the serial link testing system designed.

5.1.2.2.1 XCVR CONTROLLER (NATIVE-PHY-IP E-TILE XCVR)

In fig. 5.2 , the module denoted ‘native-phy-ip e-tile xcvr’ represents an instance of Intel’s PHY controller ipcore that provides access to PAM4 capable E-tile xcvs integrated on Statix 10TX board (through a series of interfaces, including serial signal interfaces as shown in figure 5.2), thus enabling to control test signal generation (by configuring xcvr parameters), and obtain xcvr status information.

Intel’s PHY controller ipcore exposes an Avalon-mm interface that enables for xcvr parameters re-configuration (including equalization parameters) and xcvr control.



In addition to the re-configuration interface, the xcvr controller exposes interfaces to handle both transmitted/received binary sequences (over parallel data buses), and the transmitted/received serial signal:

'tx_serial_data' and 'rx_serial_data' interfaces (shown in figure 5.2) are connected to the desired FPGA connector to output/input serial signals to/from the serial link under test³².

'tx_serial_data' and 'rx_serial_data' interfaces are exposed to receive data to be transmitted from design logic; and deliver received data for further analysis (once deserialized).³³

5.2 section provides more detail on 'native-phy-ipcore' interfaces, and supported operations and configuration are included in ipcore's user-guide [11].

5.1.2.2.2 QPRBS31 PARALLEL GENERATOR/CHECKER

QPRBS31 parallel generator/checker blocks are integrated in the serial link testing system designed to generate/check the binary test pattern used to modulate the serial signal, thus being connected directly to de 'tx_parallel_data' and 'rx_parallel_data' interfaces exposed by the xcvr controller (note that both are designed using a parallel generation/checking architecture, to generate/receive the binary sequence over a parallel bus in adaptation to the xcvr controller ipcore).

QPRBS31 parallel generator generates test pattern required as per CEI-OIF specification for serial link PAM4 tests. PRBS generator is designed to generate QPRBS31 pattern at required serial data rate (14-28GB/s) through a 128-b parallel interface as parallel output interface for the generator pattern has its width adjusted to the width of the TX_parallel_data interface in native-phy ipcore-core (more detail on QPRBS31 parallel generator design is given in appendix A).

QPRBS31 generator also generates the deskew pulse required to align serial data at RX-side for PAM4 high-speed serial data transmissions (deskew pulse generation and deskew logic are described in sub-section 5.1.2.2.4)

QPRBS31 parallel checker gets received data from 'RX_parallel_data' interface (after deskewing is applied) in 'native-phy-ipcore' over a 128-b parallel interface. QPRBS checker is able to locate errored bits in each '128-b sub-sequence' received over its parallel interface, and outputs the error pattern over another 128-b parallel interface for further analysis (where a '1' indicates a bit error occurred in the position where non-zero bit is found within the parallel 128-b 'subsequence').

QPRBS31 checker also outputs error-bits and received-bits counters that are registered in Avalon-PIO registers to be used by Nios-2 for BER calculations (note that error and received bits counter values are only valid when 'prbs_locked' signal shown in figure 5.4 is asserted, detail on QPRBS31 checker design is included in appendix C).

QPRBS31 generator/checker are instantiated on per channel basis, thus each of 2 PHYs used in design instantiate 12 QPRBS31 generator/checker pairs, each one connected to the 'TX_parallel_data' and 'RX_parallel_data' interfaces of the corresponding channel.

5.1.2.2.3 ERROR ANALYZER

Performs statistical analysis on error bursts occurring in serial link, using the error pattern output by PRBS31 checker (and received in a 128-b width interface) (refer to appendix D for detail on burst statistical analysis performed and error-statistical-analyzer design).


'error-statistics-analyzer' performs the statistical analysis over received error pattern by measuring:

- measures gap-length;
- burst length;
- and burst density data.

'error-statistics-analyzer' also measures received bits, errored-symbols/block distribution data, number of error-free words (blocks of consecutive bits),

³² for the PHYs used in the system implemented, 8 out of 12 channels can be connected to a QSFP-DD module (for one of the e-tiles used) and 2 out of 12 channels can be connected to 2.4mm SMA TX and RX connectors (for the other PHY); more info on 'outer' signal connections is given in sub-section 5.2.1

³³ note that for PAM4 high-speed links bitrate frequency is not supported by FPGA core, thus transmitted and received data are generated and delivered in 128-bit width parallel buses at a freq = bitrate/128 (details on desig clocking are given in section 5.3). 'tx_serial_data' and 'rx_serial_data' are actually 160-bit width, but only 128 of them hold data, so from here on both will be referred to as 128-bit parallel interfaces



and determines the number of RS-FEC(544,514) uncorrectable words³⁴ (number of error-affected correctable words received can be obtained with further analysis as $((\text{received_bits}/5440) - \text{uncorrectable_words})$).

'error-statistics-analyzer' is shown in figure 5.4 though it has still not been integrated in the system described by the moment of writing this document; nevertheless simulation based verification results are provided in appendix D, demonstrating design's operation.

5.1.2.2.4 DATA ALIGNMENT (DESKEW LOGIC)

'deskew logic' structures are commonly used to align data received over 2 different lanes in same clock cycle.

Stratix 10TX E-tile transceivers can control up to 24 complete xcvr's each (up to 24 duplex physical lanes per PHY); delivering serial data received over a n-bit parallel interface at a freq. $f = \text{bitrate}/n$, where $n \leq 64$.

PAM4 serial transmissions over 30Gbps require 128-b parallel interfaces so that freq. is reduced enough (down to $f = \text{bitrate}/128$) to be supported by FPGA core. Thus Stratix 10TX is configured (for both PAM4 and NRZ in the system designed) to assign the 64-b parallel interface of an of the 12 un-used channels to one of the 12 high-speed channel in use (so that only up to 12 channels can be used per PHY).

The conformed 128-b parallel interface behaves as 2 independent 64-b buffers, as right after deserializing the received data, 64-b are placed in the buffer of the used channel, and the 64-b following bits are placed in the 're-assigned' buffer belonging to the un-used channel (note that de-skew behavior is thoroughly described in [11]). From there on, data on each buffer is treated as if actually coming from 2 different physical channels and misalignments can occur, thus requiring de-skew logic to ensure data alignment.

'de-skew logic' gets data in a 128-b parallel interface from 'rx_serial_data native-phy-ipc' interface and outputs de-skewed data (free of missalignment) over 128-b parallel interface to the PRBS-checker (note that though native-phy-ipc last releases include 'internal' hard-coded de-skew logic, but an Intel's external de-skewer is used in designed system for greater flexibility).

'de-skew pulse' generation is included in QPRBS31 generator.

5.1.2.2.5 FREQUENCY MEASUREMENT (FREQ-MEASURE)

'freq-measure' measures the frequency of an input signal using a 125MHz reference clock.

2 instances of 'freq-measure' are instantiated in system designed to measure the frequency of both 'tx_clkout', 'rx_clkout' (thus bitrate; and clock recovered from received signal)³⁵.

'tx_clkout', 'rx_clkout' measured frequencies drive values of respective Avalon-PIO registers.

5.1.2.2.6 TIME MEASUREMENT (TST_TIMER)

'tst_timer' shown in 5.4 is integrated in design to count milliseconds elapsed from test start using a 125MHz reference clock.

1 'tst_timer' is instantiated per used channel and started when a test over the associated channel is run; and its msec count drives the value of the associated Avalon-PIO register that enables Nios-II to get the time elapsed for each test..

5.1.2.2.7 SYNCHRONIZATION CLOCK GENERATION (CLCK_MGMT_GEN)

'clk_mgmt_gen' is an instance of altera_s10_configuration_clock ipc'core that generates a 250MHz clock from board's resources. 250MHz clock serves as input to a 2x divider in the designed system, so as to obtain a 125MHz clock.

The 125MHz clock is referred to as 'clk_mgmt' in clocking schema shown in figure 5.13 and is the clock used for design synchronization.(details on

³⁴ note that both random-error and burst-correction capabilities of RS-FEC(544,514) are considered 'before' considering a word uncorrectable, thorough detail on 'error-statistics-analyzer' working is given in D

³⁵ note that only 1 instance the 'freq-measure' pair is instantiated per PHY, as same bitrate is used for all channels in same PHY

clocking structure through the system are given in 5.3)

5.1.2.2.8 XCVR TEMPERATURE MEASUREMENT (TEMP-SENSE)

'temp-sense' is an instance of 'altera_s10_temperature_sensor' ipcore that enables for on-die temperature measurement on dies supporting each of the 5 PHYs on Stratix 10TX board and on FPGA core. On-die temperature measurement (based on temperature sensing diodes) supported by Stratix 10TX board is used in the serial link analysis system designed to analyze PAM4 capable xcvr performance degradation as temperature increases.

'altera_s10_temperature_sensor' provided by Intel exposes 2 interfaces; 1 expecting commands requesting temperature measurements, and 1 used to provide response for requested measuring operations (both connected to Avalon-PIO registers to enable temperature measurement from Nios-2).

Fig. 5.5 shows block view for altera_s10_temperature_sensor (commands supported and interaction protocol for temperature requesting is defined in [14])

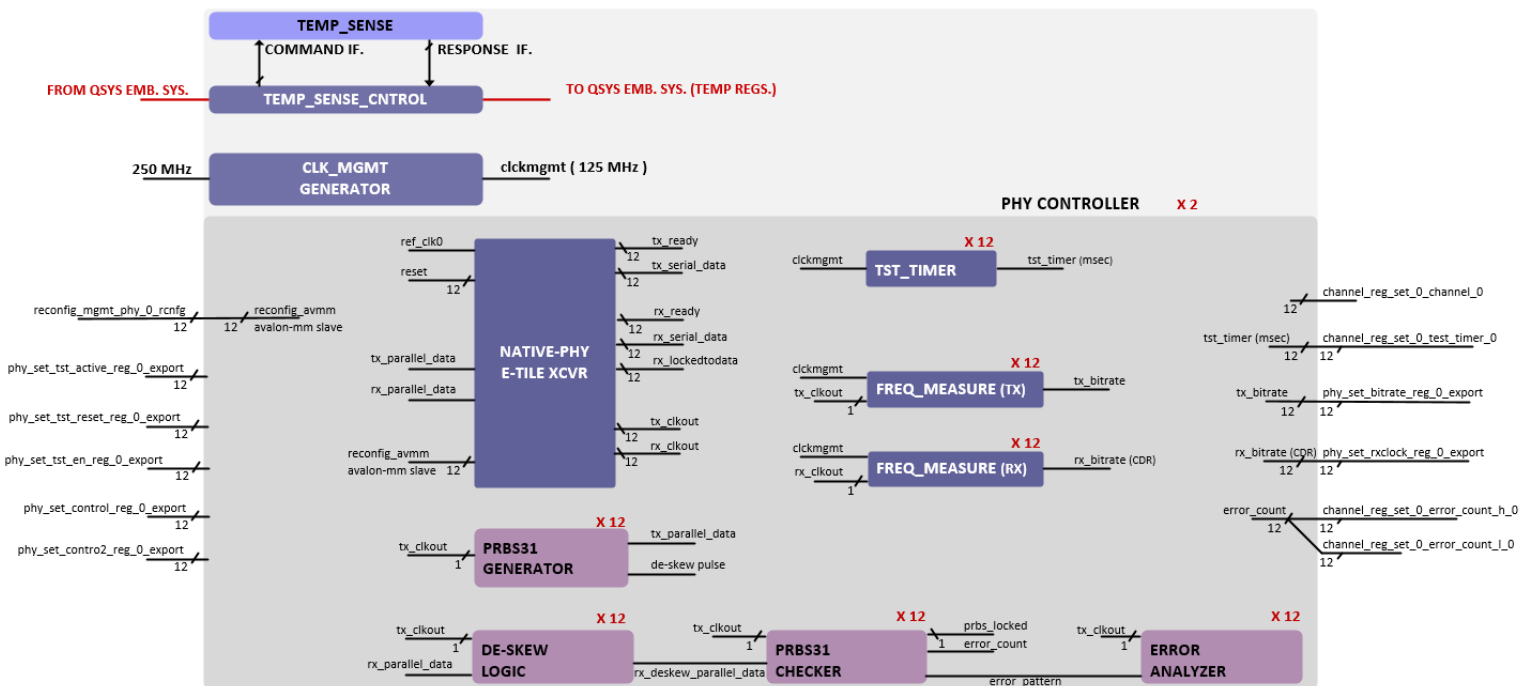


Figure 5.4.- PAM4 serial link testing system 2nd level functional division FPGA design (note that interfaces to/from QSYS embedded system are only connected if map directly to signals shown , for the sake of simplicity)




Figure 5.5.- Intel temperature-sensor-ipcore block diagram (enables control on internal on-die temperature samples from in-chip TSDs' measurements)³⁶ , [14].



5.1.3. 3RD LEVEL DIVISION (NATIVE-PHY-IP)

As aforementioned 'native-phy-ip' denotes Intel's ipcore enabling access to xcvr control.

Each instance of 'native-phy-ip' ipcore allows to control 1 of the 5 PAM4 capable E-tile xcvr on Stratix 10TX board that allows for up to 57.8 Gbps serial transmissions through the TX/RX structures shown in figures 5.6a,5.6b; so that (for a functional division) it can be considered that each 'native-phy-ip' instance stands for the whole PHY (physical E-tile xcvr), thus including both TX,RX structures 'inside' (note that 12 channels are used on each PHY the system designed, thus each of the 2 'native-phy-ip' instances contains 12 instances of the TX,RX structures shown).

Figs. 5.6a,5.6b show simplified schemas of the PHY structure, that do not match directly with underlying FPGA structure, but ease explanation on the system (note that functional diagrams shown may be also adapted for simplicity, but further detail and explanation on full xcvr structure and how the design is placed over are given in appendix B).

gray encoder/decoder:

replaces each 2-bit signal symbol by its corresponding gray-coded in TX, using mapping shown in figure 4.2; and viceversa in RX-side.

serializer/deserializer:

serializer streams data from 128-b parallel TX-parallel-data interface into a single lane (128:1); and viceversa in RX-side to fill 128-b parallel RX-parallel-data interface from bit stream received from serial link.³⁷

NRZ/PAM4 modulator/demod.:

modulator sets serial signal level to the voltage corresponding to the modulation symbol to be sent (thus, 4 levels are used for PAM4) and holds for $T_b = 1/f_b$ (where f_b = baudrate), before changing to the level corresponding to the next modulation symbol received from serializer.

NRZ/PAM4 demodulator takes samples from input signal $1T_b$ apart, turning the analogue received signal into discrete-time sampled signal.

TX/RX equalizer.:

TX equalizer is a 3-tapped FIR filter equalizer structure (as described in section 3.1)

RX equalizer is implemented using CTLE + DFE schema (as described in section 3.2)

CDR :

clock data recovery (CDR) uses the received data signal to recover a clock signal at the received bitrate frequency.

5.2. DATA INTERFACES

Figs. 5.7 - 5.9 aim to provide detail on data interfaces within the system designed (note that as aforementioned some width values used to ease explanations do not match actual parallel buses' width; figures 5.7 - 5.9 indicate actual widths in brackets, but further explanation is given in appendix B).

Fig. 5.7 shows data-interfaces' width within the internal structure(data path) of the E-tile xcvr, for each channel (note that per channel parallel data interface shown does not match actual FPGA structure. B.10 figure in appendix B shows how the 128-b (actually 160-b) interface is actually composed).

Fig. 5.7 shows data interfaces handled in design, as exposed by 'native-phy-ipcore' for the 12 channels used per PHY.

'TX_parallel_data' and 'RX_parallel_data' interfaces are exposed as (12x)(128) (actually 12x160) width bus, where 128 (160) lanes are allocated for each channel in use.

PRBS31 generator/checker data interfaces have actual 128-b width in design thus requiring some mid adaptation logic to be connected to 'TX_deskewed_parallel_data' and 'RX_deskewed_parallel_data' respectively (adaptation required isolates 128 data lanes from the actual 160-b parallel data interface as defined in Intel documentation for S10 board, [11]). PRBS31 generator's adaptation logic also inserts generated de-skew pulse in 'TX_parallel_data' interface.

³⁶ SDM (Secure Device Manager is microprocessor block that manages FPGA configuration and security. TSDs are part of temperature-sensing funtional block included in SDM [15]), thus acces to temperature values must be done through SDM (explanation on SDM exceeds the scope of this document). SDM samples temperature measured on TSDs through an ADC(analog-digital converter).

³⁷ note that 128-b parallel interfaces are actually 160-b width from which 128 are data (see appendix B for further information on actual xcvr structure)

³⁸ actual interface width including control bits is shown in brackets




Figure 5.6.- Simplified TX (a); RX (b) PMA block diagram for each channel in Stratix 10 E-tile xcvr (further explanation on E-tile xcvr structure, including PMA is given in B) , [16].




Figure 5.7.- Data interface width³⁸ over simplified TX (a); RX (b) PMA block diagram for each channel in Stratix 10 e-tile xcvr; including path from/to soft external PRBS31 generator/checker (PRBS31 generator output is marked as 128+1, as de-skew pulse is inserted within the 160-b actual width of native-phy-ipcore TX_parallel_data interface).

Figs. 5.8 - 5.9 serves as an example of how Nios-2 – FPGA connections are done using Avalon-PIO interfaces (further detail on Avalon-PIO interfaces is given in sub-section 5.2.3).

Fig. 5.8 reflects connection schema for test-timer registers associated to each channel on a certain PHY (E-tile PAM4 capable xcvr). ‘test-timer’ registers are connected to the inputs of 12 corresponding d-latches (with same 32-b width than ‘test-timer’ registers). d-latches are controlled by 12 Avalon-PIO register controllers that expose corresponding Avalon-mm slave agents to allow Nios-2 to control read operations over the latched ‘test-timer’ values.

Fig. 5.9 schema serves also to show the controlling structure used for the ‘temperature-sense’ ipcore. ‘command-interface’ shown in figure 5.5 has each bus connected to d-latched registers controlled by Avalon-PIO controllers.

‘altera_s10_temperature_sensor’ receives commands in sequence, and so outputs responses; thus returned temperature values are buffered in an LSR (linear shift register), where 32-b width d-latched registers are controlled and synchronized by Avalon-PIO controllers

(note that design entities also have reset and enable ports, not shown in diagrams for simplicity, that are assigned by applying some control logic to ‘phy_set_control_reg_0_export’ and ‘phy_set_control2_reg_0_export’ registers (described in sub-section 5.1.2.1) controlled by Nios-2 through Avalon-PIO controllers)




Figure 5.8.- FPGA tst_timer outputs connection to tst_timer Avalon-PIO registers on one channel_reg_set (one PHY).




Figure 5.9.- ‘temperature-sensor-ipcore’ control/data connection from/to Nios-2 through Avalon-PIO registers ³⁹.

5.2.1. STRATIX 10TX BOARD CONNECTORS

PHYs used in design are associated with 2 of the 5 E-tile xcvs in Stratix 10TX board.

So that, depending on selected E-tile xcvs, different board connectors are available to output native-phy ipcores’s serial data signals (‘TX_serial_data’ and ‘RX_serial_data’). In the designed system PHY0 and PHY1 are used (B appendix shows location of each xcvr(PHY) over Stratix 10TX board); with 8 PHY0’s channels mapped to the 8 different lanes available in a QSFP-DD module located in same E-tile (referred to as QSFP-DD_1x2 in figure B.2

³⁹ note that Nios-e performs temperature reading in same sequence every time, thus getting subsequent responses from temperature-sensor-ipcore in same order as requested (temp_tile_1 , temp_tile_2 , ... , temp_tile_6 , core_temp) . so that temperature responses are inserted in LSR arranged inversely (core_temp , temp_tile_6 , temp_tile_5 , ... , temp_tile_1) , such that when all readings are completed each temperature value reach its corresponding register.

in B appendix), and 1 PHY1's channel mapped to a TX/RX 2.4mm SMA connector pair ⁴⁰

PHY0,PHY1 channel mapping is done as described in test section 6.2.1 ⁴¹




Figure 5.10.- Stratix 10TX Transceiver Signal Integrity Development Kit board overview , [17].

5.2.2. AVALON MEMORY MAPPED INTERFACES (AVALON-MM)

Fig. 5.11 shows the Avalon-mm interface type exposed by 'native-phy' ipcores for xcvr re-configuration and control.

'Avalon-mm (memory mapped)' are common memory mapped interfaces exposed to read/write a set of addressed registers or values.

'Avalon-mm' interfaces perform a master-slave protocol such that slaves hold the addressed registers and accept read/write requests from Avalon-mm masters; supporting both multiple masters and multiple slaves to communicating over shared Avalon-mm interface's buses without conflict in read/write operations ('Avalon-mm' interface protocol is described in [18], [19]).

'Avalon-mm masters' read/write registers hold by slaves in the shared bus by placing the register address in the address bus, and asserting read/write signal in the shared bus (in case of write operation, new value for the register being written is also placed in data bus).

'Avalon-mm slave' to which the addressed register belongs, answer the request by either placing register's value over data bus (in case of read operation), or taking register's new value from data bus (in case of write operation).

'native-phy' ipcore's re-configuration Avalon-mm interface (shown in figure 5.11) has 19-b width address bus to enable xcvr reconfiguration and control ⁴²; but 'reconfig_address' bus exposed but the ipcore may be wider than 19-b.

'reconfig_avmm' re-configuration interface can be shared for all channels used in PHY, thus 'n' ⁴³ additional bits are added to the Avalon-mm interface's

⁴⁰ note that non-mapped channels form the 12 available on each PHY are still available to perform 'link' tests over in 'internal serial loopback' mode (loopback connection is done between channel's TX/RX within E-tile xcvr, thus keeping serial signal inside Stratix 10TX board, B appendix for further detail on serial loopback modes)

⁴¹ note that native-phy ipcore's clock inputs(signal clocks to drive transmissions over the 12 PHY used channels) must also be connected to clocks available within the corresponding E-tile xcvr

⁴² note that valid address values' meanings for xcvr re-configuration are thoroughly described in Intel documentation,[11]

⁴³ $n \geq \log_2(12)$, being 12 the number of channels in use on each PHY, in the system designed

address bus, to address the channel that is to be configured by the current requested operation (thus, the [n-1:19] segment of 'address' bus is used to identify the channel to be configured. [18:0] bits in address bus provide the register address within the reconfiguration space of the selected channel in PHY).




Figure 5.11.- 'native-phy-ipc' reconfig. Avalon-mm interface (Avalon-mm slave, dynamic reconfiguration interface), [11].

5.2.3. AVALON PARALLEL INPUT/OUTPUT INTERFACES (AVALON-PIO)

'Avalon-PIO registers' can be understood as the common-use interface used to communicate Nios-2 microprocessor with a hw design running over an FPGA.

5.12a and 5.12b figures show the actual structure for both input and output Avalon-PIO registers (note that input/output direction of Avalon-PIO registers is defined from Nios-2 side).

'Avalon-PIO registers' are composed by an 'Avalon-PIO' controller and a d-latched register.

'Avalon-PIO controller' exposes an Avalon-mm slave interface that supports communication with Nios-2 (receives read/write operation requests and outputs response).

'Avalon-PIO controller' has another interface connected to a latched register; either to D input for output Avalon-PIO registers (so that Avalon-PIO controller puts the value indicated by Nios-2 processor, and the value is placed at Q output at the next sync-clock positive edge), or Q output for input registers (working inversely to output Q value set from FPGA desing at D-input, towards Nios-2).

125MHz clock is used for synchronization by all Avalon-PIO registers in the design.

5.3. CLOCKING

Fig. 5.13 shows clocking schema for the complete data-path in the system design, where parallel data interfaces in the 'inner' part of the design (from FPGA core to serializer/desrializer element) work at a frequency $f = \text{bitrate} / 128$ being 128 the de-/serialization factor when serial data is output on 128-b width interfaces.

CDR block in RX-side outputs a clock signal at that frequency $\text{bitrate}/128$, that is assigned to native-py ipc'core 'RX_clkout' interface (show in 5.4 figure), and so used as clock input for further blocks working on 128-b width parallel data interfaces (de-skewer, PRBS31 checker, error analyzer).

Similarly, native-phy ipc'core generates 'TX_clkout' signal from the clock signal used to drive the transmitter (at bitrate frequency)⁴⁵ 'TX_clkout' is used as clock input to generate PRBS31 test pattern to be transmitted.

⁴⁴ note that Q is connected to either input_port in input Avalon-PIO register controller (Nios-2 reads Q value through Avalon-mm readdata port in Avalon-PIO controller); similarly D is connected to output_port in output Avalon-PIO register controller (Nios-2 writes D value through Avalon-mm writedata port in Avalon-PIO controller).

⁴⁵ note that each E-tile xcvr in Stratix 10TX generates the clock at bitrate frequency from a reference clock provided by a 'si5341' clock generator that feeds every E-tile over Stratix 10TX with a configurable reference clock; so that baudrate for all used transmitters within each phy is calculated as $\text{baudrate} = (\text{reference_clock}) \cdot (\text{multiplier})$ where multiplier is a xcvr attribute that can be configured using Avalon-mm re-configuration interface on 'native-phy' ipc'core (when PAM4 is used, bitrate is calculated as $2 \times \text{baudrate}$). PHY0 and PHY1 reference clocks are configured at 280MHz before starting the design




Figure 5.12.- (a) Example input Avalon-PIO register associated latch; (b) input Avalon-pio register controller; (c) example output Avalon-PIO register controller⁴⁴.

'TX_clk_out' and 'RX_clk_out' clocks driving FPGA core design at freq. of bitrate/128 (being 128 the serdes (serialization/des-serialization) factor, thus bitrate/128 the 'speed' for parallel data interfaces) are always slower than 437.5 MHz, for bitrates $\leq 56\text{Gbps}$ (as $56\text{E}9/128 = 437.5\text{ MHz}$).

'clckmgmt' 125MHz clock signal generated by the altera_s10_configuration_clock ipcore is used for design synchronization (including read/write operations over Avalon-PIOs, reset/enable signals, and even PHY re-configuration through the Avalon-mm interface referred to as 'reconfig_avmm').

So that 3 clocking domains can be distinguished in the system designed: 1 clock domain per E-tile at freq. $(280\text{E}6) \cdot (\text{TX_clk_multiplier})$ (where 'TX_clk_multiplier' is configurable on per PHY basis, through 'native-phy' ipcore Avalon-mm interface); and synchronization domain using a 125MHz clock for design synchronization.

5.4. E-TILE NATIVE PHY CONFIGURATION AND CONTROL

5.4.1. CONFIGURATION / CONTROL STRUCTURE

Fig. 5.14 shows all configuration, control, and status ports exposed by each 'native-phy' ipcore instance.

'reconfig_avmm' Avalon-mm interface that allows for xcvr reconfiguration from Nios-2 microprocessor is actually an abstraction of the internal structure of the PHY xcvr in Stratix 10TX board.

Fig. 5.15 shows the complete configuration/control underlying structure for the system designed, where NPDME (native phy debugging master endpoint) exposes an 'exported_reconfig_avmm' interface that is actually the 'reconfig_avmm' interface exposed by 'native-phy' ipcore shown in 5.14.

NPDME[11] interface acts as a multiplexer that manages the access to the internal reconfiguration interface of the E-tile xcvr (thus deciding which of the configuration requests received is applied). NPDME supports the use of Intel's TTK (transceiver toolkit) debugging tool while the system designed is running on the FPGA.

TTK is used during transmission tests to get the eye-diagram view without miss-configuring the xcvr. NPDME assigns higher priority to re-configuration commands received at 'exported_reconfig_avmm' managed by Nios-2, thus only 'read' access is granted to TTK when used for eye-measuring, so that protecting the test results' integrity for the ongoing test (xcvr configurable attributes are collected in 5.4.3).⁴⁷

⁴⁶ where $\text{TX_clkout} = \text{bitrate} / 128 = (\text{reference_clock})(\text{TX_clk_multiplier}) / 128$.

⁴⁷ NPDME multiplexing is done by a hard-coded entity referred to as 'arbitrator', detail on arbitrator operation is given in B.5 in B appendix




Figure 5.13.- Datapath clocks (signal frequency or lane rate) over simplified TX (a)⁴⁶; RX (b) PMA block diagram for each channel in Stratix 10 e-tile xcvs; including path from/to soft external PRBS31 generator/checker.

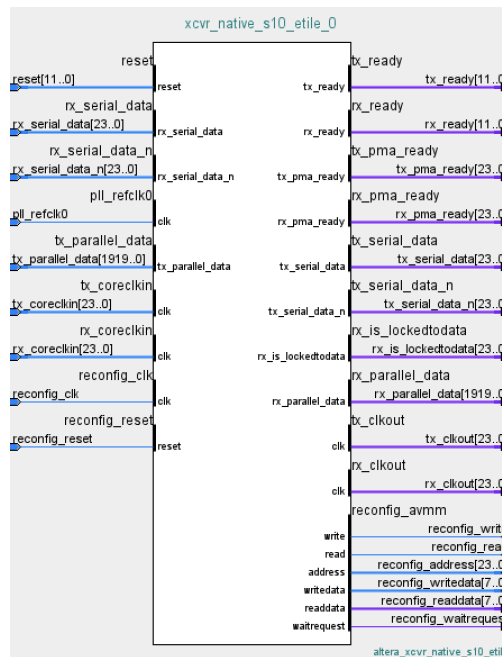


Figure 5.14.- 'native-phy-ipcove' (Stratix 10 E-Tile Transceiver Native Phy) block symbol including shared⁴⁸ dynamic reconfiguration Avalon-mm (reconfig_avmm) interface.

⁴⁸ Avalon-mm dynamic reconfiguration interface shared for all channels in PHY (most significant bits in 'reconfig_address' may be used to address the target channel within PHY to be configured)




Figure 5.15.- Simplified block diagram for native-phy-ipcore's configuration/control structure used in PAM4 serial link testing system.

5.4.2. MINIMUM CONFIGURATION SEQUENCE

Nios-2 microprocessor acts as xcvr reconfiguration controller, this implementing re-configuration sequences for operations that can be requested from Nios-2 console.

If a non-controlled transmission is to be started (when no test is parametrized to test the link under determined control parameters , but link must be started for debugging or diagnosis) there is a minimum configuration requests sequence that must be completed manually from Nios-2 console⁴⁹ :

- 'serial loopback' mode must be configured for the channels to be started
- PMA pre-set RX-equalization configuration must be set
- 'initial adaptation' must be run before starting channels

5.4.3 sub-section includes a series of tables giving explanation on 'serial loopback', PMA pre-set and 'initial adaptation' terms; and all PHY configurable attributes and status indicators available.

5.4.3. CONFIGURABLE ATTRIBUTES AND STATUS INDICATORS

Table 5.3.- PMA TX-equalization settings (vod|pre-tap-(1-3)|post-tap) (TX-equalization data in this table is extracted from, [11])

TX-EQUALIZATION PARAM	RANGE	STEP SIZE	DESCRIPTION
vod (ATTN)	$0 \leq \text{ATTN} \leq 26$ Increment decrement by 1	18.5 mV/step (0.17 dB)	de-emphasis tx-equalization FIR filter's main cursor
pre-tap 1 (PRE1)	$-10 \leq \text{PRE1} \leq 10$ Increment decrement by 2	18.5 mV/step (0.34 dB)	de-emphasis tx-equalization FIR filter's 1st pre-cursor
pre-tap 2 (PRE2)	$-15 \leq \text{PRE2} \leq 15$ Increment decrement by 1	9.25 mV/step	de-emphasis tx-equalization FIR filter's 2nd pre-cursor
pre-tap 3 (PRE3)	$-1 \leq \text{PRE3} \leq 1$ Increment decrement by 1	9.25 mV/step	de-emphasis tx-equalization FIR filter's 3rd pre-cursor
post-tap 1 (POST)	$-18 \leq \text{POST} \leq 18$ Increment decrement by 2	18.5 mV/step (0.34 dB)	de-emphasis tx-equalization FIR filter's 1st post-cursor

⁴⁹ note that steps involved in minimum configuration sequence to start a non-controlled link transmission imply complex configuration sequences partially described in Intel documentation for E-Tile xcvs [11] (xcvr configuration sequences require thorough understanding of Stratix 10TX board structure, thus exceeding the scope of this document)



Table 5.4.- RX PMA equalization parameters (RX-equalization data in this table is extracted from tables 24 and 41 in, [11]).

'initial adaptation' and 'continuous adaptation' columns indicate whether each parameter can be optimized automatically by 'adaptative parametric tuning engine' during an 'initial adaptation' or 'continuous adaptation', respectively (B provides further info on the operation of the 'adaptative parametric tuning engine')

N/A value is used in 'initial adaptation' and 'continuous adaptation' columns to indicate that the corresponding parameter is not a PMA equalization tunable parameter (N/A is applied to parameters that determine how PMA adaptation is done, but are not equalizer's parameters)

'fixed/adaptable' parameters are used to fix the corresponding parameter value so its value will not be changed by the 'adaptative parametric tuning engine' on sub-subsequent PMA adaptations (either initial or continuous)

(footnotes ⁵⁰, ⁵¹, ⁵² in table give clarification on aspects uncovered within this table).

RX-EQUALIZATION PARAM	SUPPORTED VALUES	FIRMWARE DEFAULT	INITIAL ADAPTATION	CONTINUOUS ADAPTATION	POSSIBLE MANUAL OPTIMIZATION CONFIGURATION	DESCRIPTION ¹
GS1	0, 1, 2, 3	0	no	no	yes	CTLE Low Frequency Gain Shaping 1.
GS2	0, 1, 2, 3	0	no	no	yes	CTLE Low Frequency Gain Shaping 2.
RF_B1	0, 1, 2, 3, 4, 5, 6, 7, 8	0	yes	yes	yes	RF_B1 setting
RF_B1 fix adaptable	fixed, adaptable	adaptable ⁵¹	N/A	N/A	yes	RF_B1 setting
RF_B0	0, 1, 2, 3, 4, 5	0	yes	yes	yes	RF_B0 setting
RF_B0 fix adaptable	fixed, adaptable	adaptable ⁵¹	N/A	N/A	yes	RF_B0 setting
RF_A	100, 110, 120, 130, 140, 150, 160	160 (NRZ) 130 (PAM4)	no	no	yes	RF_A setting
GAINLF	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	8	yes	yes	yes	CTLE Low Frequency Gain.
GAINLF fix adaptable	fixed, adaptable	adaptable ⁵¹	N/A	N/A	yes	CTLE Low Frequency Gain.
CTLE LF min	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	0	N/A	N/A	yes	Limits CTLE LF minimum
CTLE LF max	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	15	N/A	N/A	yes	Limits CTLE LF maximum
GAINHF	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	0	yes	yes	yes	Inputs CTLE High Frequency Gain.
GAINHF fix adaptable	fixed, adaptable	adaptable ⁵¹	N/A	N/A	yes	Inputs CTLE High Frequency Gain - fix adaptable options.
CTLE HF min	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	0	N/A	N/A	yes	Limits CTLE HF minimum
CTLE HF max	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	15	N/A	N/A	yes	Limits CTLE HF maximum
RF_P2	-10, -9, -8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10	0	yes	no	no	RF_P2 setting
RF_P2 fix adaptable	fixed, adaptable	adaptable ⁵¹	N/A	N/A	yes	RF_P2 setting -fix adaptable options.
RF_P2 min	-10, -9, -8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10	-10	N/A	N/A	yes	Limits RF_P2 minimum
RF_P2 max	-10, -9, -8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10	10	N/A	N/A	yes	Limits RF_P2 maximum
RF_P1	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	0	yes	yes	no	RF_P1 setting
RF_P1 fix adaptable	fixed, adaptable	adaptable ⁵¹	N/A	N/A	yes	RF_P1 setting -fix adaptable options.
RF_P1 min	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	0	N/A	N/A	yes	Limits RF_P1 minimum
RF_P1 max	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	15	N/A	N/A	yes	Limits RF_P1 maximum
RF_P1 threshold	-15, -14, -13, -12, -11, -10, -9, -8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	- ⁵²	N/A	N/A	yes	Controls the rate RF_P1 adapts
RF_P0	-15, -14, -13, -12, -11, -10, -9, -8, -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	0	yes	yes	no	RF_P0 setting
RF_P0 fix adaptable	fixed, adaptable	adaptable ⁵¹	N/A	N/A	yes	Limits RF_P0 -fix adaptable options.
RF_P0 threshold	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15	- ⁵²	N/A	N/A	yes	Controls the rate RF_P0 adapts
RF_BOT	10, 15, 20, 25, 30, 35, 40, 45, 50	20	no	no	yes	Controls the rate RF_B0 adapts



Table 5.5.- RX PMA configuration pre-sets (pre-set equalizer parameters for 10G, 28G-LR, 28G-VSR, 56G-LR, 56G-VSR models defined in 4.2, [8]) (note that values provided in this table are used for pre-set PMA configuration in PAM4 serial link testing system; and match values used in 'hard-coded' pre-set PMA configurations available when instantiating native-phy-ipcpre, note 3 is used to mark exceptions.)

(footnotes ⁵³, ⁵⁴ in table give clarification on aspects uncovered within this table).

RX-EQUALIZATION PARAM	10G	28G-LR	28G-VSR	56G-LR	56G-VSR
GS1	2	2	fw-default ⁵³	1	0
GS2	1	2	fw-default ⁵³	1	0
RF_B1	5	1	fw-default ⁵³	8	3
RF_B1 fix adaptable	fixed	fixed	adaptable	fixed	fixed
RF_B0	1	1	fw-default ⁵³	2	3
RF_B0 fix adaptable	fixed	fixed	adaptable	adaptable	fixed
RF_A	fw-default ⁵³	130	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
GAINLF	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
GAINLF fix adaptable	adaptable	adaptable	adaptable	adaptable	adaptable
CTLE LF min	fw-default ⁵³	fw-default ⁵³	7 ⁵⁴	fw-default ⁵³	fw-default ⁵³
CTLE LF max	2	3	fw-default ⁵³	2	fw-default ⁵³
GAINHF	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
GAINHF fix adaptable	adaptable	adaptable	adaptable	adaptable	adaptable
CTLE HF min	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
CTLE HF max	fw-default ⁵³	fw-default ⁵³	7 ⁵⁴	fw-default ⁵³	fw-default ⁵³
RF_P2	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
RF_P2 fix adaptable	fixed	fixed	adaptable	adaptable	adaptable
RF_P2 min	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
RF_P2 max	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
RF_P1	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
RF_P1 fix adaptable	adaptable	adaptable	adaptable	adaptable	adaptable
RF_P1 min	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
RF_P1 max	fw-default ⁵³	6	fw-default ⁵³	6 ³	6 ⁵⁴
RF_P1 threshold	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
RF_P0	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
RF_P0 fix adaptable	adaptable	adaptable	adaptable	adaptable	adaptable
RF_P0 threshold	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³	fw-default ⁵³
RF_B0T	10 ³	10	10 ⁵⁴	40	10

⁵⁰ note that Intel does not provide exact definition for RX-equalizer parameters; undefined parameters are included in this table as PAM4 serial link testing developed allows to optimize their values to achieve lower BER values, though how each parameter modifies RX-equalization structure is unknown.

⁵¹ no info. on default value at PMA startup is given Intel's documentation; indicated value is configured as initial value at PAM4 serial link testing system start

⁵² parameter not set in PMA4 serial link testing system; and no info. on default value at PMA startup is given Intel's documentation.

⁵³ fw-default stands for firmware-default (firmware-default values can be found in 5.4)) for firmware default value refer to tab)

⁵⁴ parameter value differs from pre-set available when instantiating native-phy-ipcpre; values used here are provided by Intel testing team based on serial link tests conducted (values used in pre-set configurations available at instantiation can be found in [11], table 43))



Table 5.6.- E-tile|channel PMA configuration params. (note that values marked as default are initial values assigned at PAM4 serial link testing system startup, for all channels used withing both PHYs).

(footnotes ⁵⁵, ⁵⁶, ⁵⁷, ⁵⁸, ⁵⁹ in table give clarification on aspects uncovered within this table).

PMA CONFIG PARAM.	VALUE (DEFAULT)	DESCRIPTION
rx termination mode ⁵⁵	GND , VCC , FLOAT ⁵⁵	set rx termination and tx-driver tri-state behaviour ⁵⁹
serial loopback ⁵⁶	enabled(default), disabled	sets internal serial loopback between channel's tx,rx serial lanes
reverse loopback ⁵⁶	enabled, disabled(default)	sets channel's tx in reverse loopback (tx fed from rx signal)
line encoding	PAM4(default) , NRZ	sets signal modulation
gray encoding	enabled(default), disabled	enables disables gray encoding
1/1+D encoding	enabled, disabled(default)	enables disables 1/1+D pre-coding
swizzle	enabled(default), disabled	allows to invert MSB,LSB on per-symbol basis order before signal modulation
Invert tx polatiry	enabled, disabled(default)	inverts signal polarity over tx differential serial output
Invert rx polatiry	enabled, disabled(default)	inverts signal polarity over rx differential serial output
adaptation mode ⁵⁷	initial-adapt(default),continuous-adapt	sets PMA adaptation mode
tx-clk-divider	$i / i \in \{N / ((50 \leq i \leq 100) \cap (i \% 2 = 0) \cup (i \% 5 = 0))\}$ ⁵⁸	sets tx reference clock (baudrate)
rx-clk-divider	$i / i \in \{N / ((50 \leq i \leq 100) \cap (i \% 2 = 0) \cup (i \% 5 = 0))\}$ ⁵⁸	sets rx reference clock for CDR

⁵⁵ note that RX-termination mode parameter is not configurable,but set initially on PAM4 testing system startup depending on the connection map for each channel used. GND is configured for unused channels; FLOAT sets RX-termination to be undriven for non-connected active channels (in internal serial loopback mode); VCC must be used to set RX-termination driven to VCC_H (in normal operation for channels mapped to external connectors, both TX-termination,RX-termination are driven to VCC_H). [11] table 77 gives more detail on RX-termination modes.

⁵⁶ refer to B.3.2 for further detail on loopback modes

⁵⁷ refer to B.4.1 for further detail on PMA adaptation modes

⁵⁸ note that more values are allowed for both TX-clk-divider, RX-clk-divider ([11] table 79 shows all supported divider values). TX-clk-divider, RX-clk-divider ranges in this table correspond to set of values supported in PAM4 testing system for baudrates limited to 14 – 28GBs

⁵⁹ RX-termination must be set to VVC_H when external AC-coupling is used (thus when mapping RX-termination to external connector); whereas floating RX-termination must be used when AC-coupling is not used



Table 5.7.- Serial link transmission status indicators (note both native-phy-ip xcvr ports; and other signal integrity indicators within PAM4 testing system are included to help understanding parameters involved in signal integrity analysis).

(footnotes ⁶⁰, ⁶¹, ⁶², ⁶³ in table give clarification on aspects uncovered within this table).

STATUS-INDICATOR	VALUE	DESCRIPTION
tx-ready	asserted de-asserted	when asserted indicates tx is locked to PLL driving clock , thus transmitting at the configured bitrate
rx-ready	asserted de-asserted	when asserted indicates both rx-PMA configuration ; and reference clock for CDR operation are correct.
rx-lockedtodata	asserted de-asserted	when asserted indicates CDR is able to recover data and clock from serial received signal ⁶²
prbs-locked	asserted de-asserted	asserted when received signal quality is enough to guarantee PRBS31 pattern checking can be done accurately ⁶³
ppm difference	$(tx-clkout / rx-clkout - 1) \times 1E3$ kHz	measures the difference between tx-bitrate and bitrate detected by CDR (when tx,rx belonging to same channel are connected in any loopback mode , ppm should be 0)
tx-clkout	tx-bitrate / 128	clock signal at frequency at which data to be transmitted is taken from 128-b parallel data interface
recovered clock (CDR) (rx-clkout)	rx-bitrate / 128	clock signal at frequency f/128 where f is the received bitrate detected by CDR
lower even eye	(height measured in last channel adaptation) ⁶⁰	lower eye height for even channel ⁶¹
mid even eye	(height measured in last channel adaptation) ⁶⁰	mid eye height for even channel ⁶¹
upper even eye	(height measured in last channel adaptation) ⁶⁰	upper eye height for even channel ⁶¹
lower odd eye	(height measured in last channel adaptation) ⁶⁰	lower eye height for odd channel ⁶¹
mid odd eye	(height measured in last channel adaptation) ⁶⁰	mid eye height for odd channel ⁶¹
upper odd eye	(height measured in last channel adaptation) ⁶⁰	upper eye height for odd channel ⁶¹

⁶⁰ eye-heights are measured in steps. TTK eye-viewer divides signal-voltage margin in a no. of steps configurable when performing measurement; no. of steps (thus step-size) is indicated for eye-views shown in 6.2.2

⁶¹ in RX-side, serial signal is divided in 64-b sample blocks (note that when using PAM4, signal is divided in blocks of 32 samples corresponding to 64-b blocks); so that de-serializer puts samples corresponding to 64-b over even GXE RX-channel structure , and the following 64-b over odd GXE RX-PMA (as 2 RX-channel parallel interfaces are used for high-speed PAM4 transmissions in PAM4 testing system, 5.1.2.2.4). In RX-channel structure, eye-viewer instrumentation is placed after de-serializer; so that, eyes are measured over 2 RX-channels though signal samples come from 1 single physical lane.

⁶² RX-lockedtodata may be asserted though ISI affecting incoming signal is high enough to degrade signal quality such that prbs-locked signal will be de-asserted.

⁶³ prbs-locked is output by PRBS31 checker developed (C gives explanation on when prbs-locked signal is asserted).

5.5. HW SYSTEM DEBUGGING AND VERIFICATION

PAM4 serial link evaluation system described was verified using Intel's TTK(transceiver toolkit) and signal-tap debugging tools ⁶⁴ (note that unit verification for designed PRBS31 generator and checker , and error-statistical-analyzer is described in appendixes C,D, respectively; along with design description).

TTK debugging tool allows to perform eye-diagram measurements over active channels performing either PAM4 or NRZ transmissions; so that, 'native-phy' ipcore configuration correctness in the designed system was verified by measuring the eye-diagram for internal-serial-loopback transmissions ⁶⁵ , thus expecting a completely opened eye in case the system designed had completed the E-tile xcvr configuration correctly (further information on how TTK performs on-die eye diagram measurements on Stratix 10TX devices is given in sub-section 6.1.2).

Preliminary verification stages described aimed to confirm the implementation of the xcvr configuration sequence; verification tests were performed using Intel's PRBS31 generator and verifier that are hard-coded 'inside' the E-tile xcvr itself; so that isolating xcvr configuration testing from the PRBS31 generator and checker developed. 5.16 figure shows the structure described for verification of xcvr configuration sequences performed by using E-tiles' hard-coded PRBS31 generator and checker to drive transmission over which eye measurements were performed.




Figure 5.16.- Design structure used for preliminary verifications of xcvr configuration sequences observing eye-diagram through TTK eye-viewer⁶⁶ .

Signal-tap was used for verification on developed PRBS31 generator/checker integration in the designed system (Signal-tap adds hw to the design logic that allows to sample on-die signals -without outputting them to board ports- as an internal oscilloscope).

PRBS31 generator and checker integration was verified by duplicating the design shown in 5.16 and replacing hard-coded PRBS31 generator and checker in both 'copies'. One copy feeds 'native-phy' ipcore from PRBS31 generator and checker Intel's official ipcores; and the other integrates the developed PRBS31 generator/checker⁶⁶ .

Signal-tap was configured to take a buffer of 256 signal samples of the 'RX_parallel_data' interface of both 'native-phy' ipcores (signal samples were taken on positive edges of 'RX_clkout' output clock),5.17.

PAM4 transmission was started in 'internal serial loopback' mode, and integration was verified by comparing samples taken from received parallel data interfaces (note that comparison between received signals is possible as both ipcore PRBS31 generator/checker and developed PRBS31 generator/checker start generation aligned at same point of the PRBS31 pattern).

⁶⁴ both TTK and signal-tap are installed with Quartus Prime Pro programmable-logic-device design sw

⁶⁵ in internal serial loopback mode, TX and RX are connected inside the xcvr itself thus removing the serial link connection and so guaranteeing a distortion-free transmission

⁶⁵ JTAG-to- Avalon-mm bridge for TTK connection was omitted here for simplicity

⁶⁶ note that previous functional simulation-based verification of designed PRBS31 generator and checker is described in C

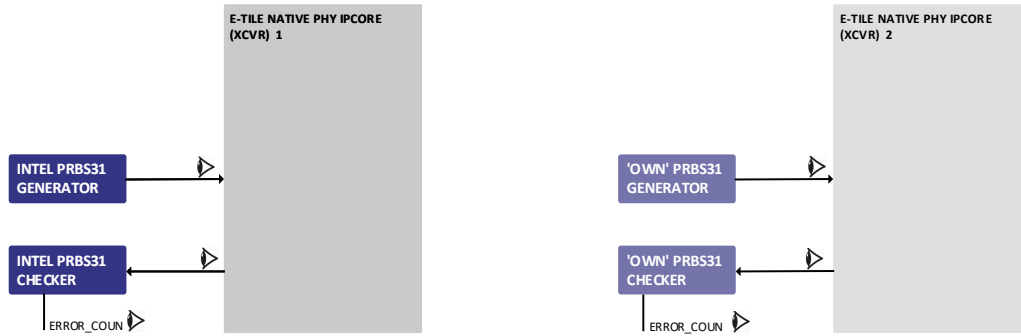


Figure 5.17.- Design structure used for verifications of on developed PRBS31 generator/checker integration with native-phy-ipcore by performance comparison to Intel's soft PRBS31 generator and checker ipcores.

5.6. NIOS-II μ PROCESSOR PROGRAM

This section provides a brief description on the main capabilities of the developed Nios-2 xcvr-control app that supports interaction with FPGA design for xcvr configuration and control.

5.6.1 sub-section gives general insight on the simple program flow, showing available xcvr configuration options and supported operations; and main developed operations are described in sub-sections 5.6.2 - 5.6.4.

5.6.1. GENERAL FLOW

After downloading both hw system design to Stratix 10TX FPGA and sw to Nios-2:

1. Nios-2 resets xcvr equalizers.
2. Performs initial configuration of the xcvr, setting default values for the attributes:


<u>TX RX_clk_divider</u> :	100 (default), thus baudrate = 28E9 GBs	<u>adaptation mode</u> :	initial_adaptataion_half_effort
<u>serial loopback</u> :	enabled (default)	<u>RX-equalization preset</u> :	none (default) (no equalization)
<u>NRZ PAM4</u> :	PAM4 (default)	<u>vod</u> :	0 (default)(no equalization)
<u>gray coding</u> :	enabled (default)	<u>TX_pretap_3</u> :	0 (default)(no equalization)
<u>1/1+D</u> :	enabled (default)	<u>TX_pretap_2</u> :	0 (default)(no equalization)
<u>invert_TX_polarity</u> :	disabled (default)	<u>TX_pretap_1</u> :	0 (default)(no equalization)
<u>invert_RX_polarity</u> :	disabled (default)	<u>TX_posttap_1</u> :	0 (default)(no equalization)

3. Collects and shows status info (as shown in figure 5.18), showing following status indicators:

<u>serial loopback</u>	<u>adaptation mode</u>
<u>reverse loopback</u>	<u>PMA configuration mode</u> : RX-equalizer pre-set
<u>line encoding</u>	<u>TX_ready</u>
<u>gray encoding</u>	<u>RX_freqlocked_1ms</u>
<u>1/1+D encoding</u>	<u>prbs_locked</u>
<u>swizzle</u>	<u>error_count</u>
<u>invert_TX_polarity</u>	<u>bit_count</u>
<u>invert_RX_polarity</u>	

<u>test time</u> :	elapsed from last e-tile xcvr reset , not related to tests run over single channels
<u>measured bitrate</u> :	as (TX_clkout_freq):(128) , being TX_clkout_freq the measured frequency for TX_clkout
<u>measured received bitrate</u> :	as (RX_clkout_freq):(128) , being RX_clkout_freq the measured frequency for RX_clkout recovered by CDR
<u>configured bitrate</u> :	per PHY, as (280E6):(TX_clk_divider)

4. Shows main menu (complete Nios-2 app menu can be found in E), and waits for operation request input (when operation request is received, program flow goes back to 3 after it is completed).



```
stratix 10 tx si board 56gbps PAM4/NRZ soft PRBS31
-----
|board revision      : S10 TX SI kit production (rev B1)
|hardware revision   : 04/27/2020 variant 00
|eTile firmware version : 0x1081.14
|software build date  : Jul  8 2020 16:58:56
-----
|PHY                : 0| 1|
|BANK                : 9C| 9A|
|number of lanes     : 12| 12|
|lane rate (Mbps)    : 56000| 56000|
|selected channel in PHY : 0| 0|
|selected PHY        : 0
-----
```

(a)

```
PHY 0 bank 9C (QSFP-DD 1x2)
-----
lane|chann. : 0| 1| 2| 3| 4| 5| 6| 7| 8| 9| 10| 11|
-----
rx termination : VCC VCC FLOAT FLOAT VCC VCC VCC VCC FLOAT FLOAT VCC VCC
LPM mode       : OFF OFF OFF OFF OFF OFF OFF OFF OFF OFF OFF OFF
serial lpbck   : 1| 1| 1| 1| 1| 1| 1| 1| 1| 1| 1| 1|
reverse lpbck  : 0| 0| 0| 0| 0| 0| 0| 0| 0| 0| 0| 0|
line encoding  : PAM4 PAM4 PAM4 PAM4 PAM4 PAM4 PAM4 PAM4 PAM4 PAM4 PAM4 PAM4
gray encoding  : GRAY GRAY GRAY GRAY GRAY GRAY GRAY GRAY GRAY GRAY GRAY GRAY
1/10 encoding :
swizzle       : SWIZZLE SWIZZLE SWIZZLE SWIZZLE SWIZZLE SWIZZLE SWIZZLE SWIZZLE SWIZZLE SWIZZLE SWIZZLE SWIZZLE
invert tx polarity :
invert rx polarity :
PRBS          : 31 31 31 31 31 31 31 31 31 31 31 31
adaptation mode : 0.5 iADP 0.5 iADP 0.5 iADP 0.5 iADP 0.5 iADP 0.5 iADP 0.5 iADP 0.5 iADP 0.5 iADP
channel type    : DAC DAC C2C C2C DAC DAC DAC DAC C2C C2C DAC DAC
connection type : QSFPDD QSFPDD DEFAULT DEFAULT QSFP28 QSFP28 QSFPDD QSFPDD DEFAULT DEFAULT QSFP28 QSFP28
PMA configuration : DEFAULT DEFAULT DEFAULT DEFAULT DEFAULT DEFAULT DEFAULT DEFAULT DEFAULT DEFAULT DEFAULT DEFAULT
pll lcked(tx_ready) : 1| 1| 1| 1| 1| 1| 1| 1| 1| 1| 1| 1|
rx lcked lms(CDR) : 1| 1| 1| 1| 1| 1| 1| 1| 1| 1| 1| 1|
prbs lcked       : 1| 1| 1| 1| 1| 1| 1| 1| 1| 1| 1| 1|
error count      : 0| 0| 0| 0| 0| 0| 0| 0| 0| 0| 0| 0|
-----
error count (prbschckr) : 0| 0| 0| 0| 0| 0| 0| 0| 0| 0| 0| 0|
bitcount (prbschckr) : 2.58e+10| 2.58e+10| 2.58e+10| 2.58e+10| 2.58e+10| 2.58e+10| 2.58e+10| 2.58e+10| 2.58e+10| 2.58e+10| 2.58e+10| 2.58e+10|
BER (CL=0.95) ch 0 : 4.826255e-10
BER (CL=0.95) ch 1 : 4.826255e-10
BER (CL=0.95) ch 2 : 4.826255e-10
BER (CL=0.95) ch 3 : 4.826255e-10
BER (CL=0.95) ch 4 : 4.826255e-10
BER (CL=0.95) ch 5 : 4.826255e-10
BER (CL=0.95) ch 6 : 4.826255e-10
BER (CL=0.95) ch 7 : 4.826255e-10
BER (CL=0.95) ch 8 : 4.826255e-10
BER (CL=0.95) ch 9 : 4.826255e-10
BER (CL=0.95) ch 10 : 4.826255e-10
BER (CL=0.95) ch 11 : 4.826255e-10
-----
test time PHY 0 : 0h 0m 0s
tx clkout frequency (bitrate) : 437499 kHz
recovered clock freq (CDR)(lane 0) : 437499 kHz
measured ppm difference : 0
```

(b)

```

0 10 20 30 40 50 60 70 80 90
|.....|.....|.....|.....|.....|.....|.....|.....|
core temp : 45°
HSSI_0_0 (1C/D/E/F) temp : no data
HSSI_1_0 (8B)(phy 4) temp : no data (inactive phy)
HSSI_2_0 (8C)(phy 3) temp : no data (inactive phy)
HSSI_0_1 (9A)(phy 1) temp : 49°
HSSI_1_1 (9B)(phy 2) temp : no data (inactive phy)
HSSI_2_1 (9C)(phy 0) temp : 53°
-----
PHY 0 bank 9C DDQ_1x2 L2 plugged in Vendor : Molex Part : 1002971101 QSFP28 Passive Cable with length 1 m
```

(c)

Figure 5.18.- (a) Nios-2 console application status info. (stratix 10TX hw info , including configured bitrate for used PHYs; (b) per PHY xcvr status info; (c) e-tile/core temperature , and QSFP-DD modules' status info.)



5.6.2. AUTO-SWEEP

'auto-sweep PMA parameter' operation finds the TX-equalization selected parameter's⁶⁷ value for which lower BER is achieved. Nios-2 fixes xcvr configuration and sweeps on all parameter's supported values measuring BER for 10 secs, for each supported value. 'auto-sweep' returns the optimum value (or range) found for the selected parameter.

5.6.3. BAUDRATE-SWEEP

'baudrate-sweep' operation allows to find the maximum achievable bitrate over each serial link set-up, to guarantee that the target BER will not be exceeded.

'baudrate -sweep' fixes the xcvr configuration (including TX and RX equalization parameters) and performs BER measurement for a 358 s serial link test, for all supported TX_clk_divider/RX_clk_divider values (thus for all supported baudrates). So that, a BER progression curve can be obtained range of supported baudrates from the BER measurements retrieved; allowing to determine the maximum baudrate before a maximum target BER is reached, for the link tested.

5.6.4. TEST CONFIGURATION / RUN

'config test' option (refer to E for complete Nios-2 app menu) allows to configure a test stop condition to be applied to later started tests and enable/disable temperature monitoring :

1. Stop condition: limited time / limited amount of bits sent.
2. Temperature monitoring: if enabled measures temperature on e-tile to which channel under test belongs, every time test status data is collected , and monitors either absolute e-tile temperature or e-tile temperature variation⁶⁸ :
 - Absolute temperature threshold exceeded (temperature threshold can be configured when 'temperature threshold monitoring' is enabled)
 - Maximum temperature drift exceeded (being temperature drift calculated as max_test_temp - min_test_temp , for all temperature samples measured from test start)

'run test' option allows to start a test over the selected channel (test will run under the stop and temperature monitoring conditions configured when test is started).

1. Requests for channel selection (channel selected acts as RX)
2. Detects matching channel in the E-tile xcvr that will act as TX (depending on serial loopback mode of selected RX channel)⁶⁹ (details on external connection between E-tiles's channels is given in appendix B)
3. Configures TX channel to match the current configuration of the selected RX-channel
4. Creates and starts a thread to control the test (new thread takes test status info every 0.5 ms and stops the xcvr when stop conditions are met.)

'test-control' thread keep updating test status data over Nios-2 console (see figure 5.19):

temperature (if enabled)	test time elapsed
BER	TX_clkout
error count	RX_clkout
sent bits count	

⁶⁷ 'auto-sweep' is supported for vod , pre-tap-3 , pre-tap-2 , pre-tap-1 , and post-tap-1 TX-equalization parameters

⁶⁸ note that e-tile temperature is measured on each test status info collection, only if temperature monitoring is enabled (either for absolute e-tile temperature or e-tile temperature variation)

⁶⁹ when RX selected channel has serial loopback mode disabled (thus assuming an external connection is done), serial loopback mode is disabled for matching TX channel



(collected info is stored in a text log file in console's host, enabling for further processing to analyze test progress)

```
(tst 1)checking temperature sense ready status
(tst 1)ber(CL=0.95)(errcoun = 0) ch 2 : 3.016885e-13
(tst 1)err count ch 2 phy 1 : 0.000000e+00

(tst 1)test time ch 2 phy 1 : 0h 2m 57s
(tst 1)tx clkout frequency (bitrate) : 437499.997000 kHz
(tst 1)recovered clock freq (CDR)(lane 0) : 437499.996000 kHz
(tst 1)measured ppm difference : 0

(tst 1)err count (prbschcker_pam) ch 2 phy 1 : 0.000000e+00
(tst 1)b count (prbschcker_pam) ch 2 phy 1 : 0.000000e+00

(tst 1)tile_4 temp(phy 1) = 42°

tst_1_dat 177572| 9.944032e+12| 437499997.000000| 437499996.000000| 1| 42| 3.016885e-13
```

Figure 5.19.- Ongoing serial link test status info printed on Nios-2 console.



6. SERIAL LINK TEST AND RESULT ANALYSIS

This section presents the serial link test results obtained using the serial link testing system developed in this work, together with an analysis of the results obtained in terms of bit error rate (BER) and eye-diagram aperture. A preliminary description on calculation/measurement methods used to produce the results presented, and test set-up description is given in section 6.1.

Results presented in this section correspond to serial link tests run under same conditions, enabling for fair performance comparison (common test conditions used are defined at the beginning of section 6.2, that includes test results obtained and an analysis on serial link performance for high-datarate PAM4 transmissions conducted). In figure 6.1 (copy of image 6.4 , included here for the sake of clarity) the basic model set-up used to conduct high-datarate transmissions is shown (note that the highlighted cable at the right-side of the image is an example of one of the serial links tested plugged between 2 of the QSFP-DD modules integrated over Stratix 10TX board; similarly, the highlighted cables at the bottom of the image correspond to another serial link channel tested that sets a loopback connection for a TX/RX channel mapped to a SMA 2.4 mm connector).




Figure 6.1.- Basic serial link testing set-up reference model (copy of image 6.4 , included here for the sake of clarity)

6.1. MEASUREMENT | CALCULATION METHODS

'Nios-2' xcvr controller samples error related data every 0.5 ms for each serial link test running over the used xcvr channels (transmission performance over the link under test is mainly analyzed in terms of measured BER), that can be further processed to obtain curves that enable for analysis on channel performance progression over test time.

Besides BER analysis, TTK is used for each test to generate an eye-diagram view from which signal integrity can be analyzed.

6.1.1 and 6.1.2 sub-sections provide detail on BER calculation and on-die statistical eye-diagram generation, respectively.

BER values measured for each test correspond to BER values measured after RX equalizer (note that basic TX/RX equalization settings adjustment was performed before each test in order to measure the minimum achievable BER for each serial link tested):

- TX equalization settings were adjusted by running an 'auto-sweep' over TX-equalizer vod parameter, to configure best supported value for minimum BER.
- RX equalization settings were optimized by running and 'initial-adaptation' for every serial link under test, so that results obtained can be fairly compared.



6.1.1. BER CALCULATION

PAMn SER(symbol error rate) theoretical calculation (where n is the number of symbols in modulation), is given by the expression in equation 1 (assuming n symbols are equiprobable) [20]:

$$SER = \frac{1}{M} \sum_{i=0}^{M-1} \sum_{j=0, j \neq i}^{M-1} P_{ij} \quad (1)$$

where P_{ij} is the probability of receiving symbol j when symbol i was transmitted.

BER can be calculated from SER as:

$$BER = \frac{1}{M} \sum_{i=0}^{M-1} \sum_{j=0, j \neq i}^{M-1} \frac{d_{ij}}{\log_2(M)} P_{ij} \quad (2)$$

where d_{ij} are the number of bit errors occurred when symbol i is received as symbol j, and $\log_2(n)$ the number of bits on each modulation symbol (thus being $d_{ij}/\log_2(n)$ the 'partial' BER when symbol i is received as symbol j; and BER the weighted average calculated adding 'partial' BERs each weighted by its occurrence probability)

BER expression shown in equation 2 cannot be calculated in the design system as P_{ij} probabilities are not known, thus BER calculated in the designed system is estimated as:

$$BER = \text{error_bits}/\text{total_bits} \quad (3)$$

where $\text{total_bits} = (\text{test_timer_ms}) \cdot (\text{bitrate_ms})$, being test_timer_ms the time elapsed from test start in ms; and bitrate_ms the bitrate in bits/ms. ⁷⁰

In case of error-free transmission ($\text{error_bits} = 0$), test BER value given corresponds to the statistical 0.95 confidence level BER (referred to as BER(CL = 0.95)).

BER(CL = 0.95) gives BER value that is assumed to be *qeq* than BER observed if the same test transmission is repeated, with a 0.95 confidence level. BER(CL = 0.95) is calculated as [21]:

$$BER = 3/\text{total_bits} \quad (4)$$

6.1.2. ON-DIE STATISTICAL EYE DIAGRAM GENERATION

TTK eye-viewer allows for both PAM4 and NRZ on-die eye-diagram measurement, using on-die instrumentation in Stratix 10TX devices. TTK does not generate an 'analog' eye diagram view ⁷¹.

TTK generates the statistical eye by shifting both the received signal sampling point both horizontally (in time) and vertically (shifting voltage reference levels); so that, TTK moves the reference point from the center of the eye in horizontal (simulating signal delays) and vertical (simulating signal amplitude distortion). For each point TTK measures BER that would be obtained if eye's reference point were moved that point. So that, TTK associates each point in the eye diagram with a BER value that is colored in color gradient scale, turn out to be eye shaped as shown in 6.2; where the blue eye-shaped area surrounding eye-diagram's center is an area of BER <1E-6, meaning that reference point could be displaced to any position within the blue eye, keeping BER below 1E-6. So that, the greater the allowed movement, greater distortion (in terms of delay or signal level distortion) can be assumed (horizontal movements/shifts of the sampling point (movements in time) determine the tolerance to signal delay; while BER measurements with vertical shifts reveal tolerance to signal level degradation effects (mainly ISI))

⁷⁰ note that for long-term high-speed serial link tests it can be assumed that amount of bits sent is enough for the estimated BER to have stabilized, thus being to the theoretical value given in equation 2

⁷¹ PAM4 e-tile xcvr's are digital xcvr's; so an oscilloscope-like eye diagram cannot be generated, but an statistical eye-diagram based on BER contours as described in CEI-OIF 4.0, 2.C.5.3 Annex – Generation of Statistical Eye [8]




Figure 6.2.- Example statistical eye-diagram measurement using TTK eye-viewer for on-die statistical eye generation over low distortion PAM2 transmission.

6.2. SERIAL LINK TESTS AND RESULTS ANALYSIS

6.2.2 shows the serial link test results obtained from shown sub-sequent link tests (described in 6.2.1), run under common conditions:

- PRBS31 is used as data pattern for all tests.
- For each tested link a baudrate-sweep is run to determine the maximum achievable baudrate for the target BER of 1E-4 established by CEI-OIF specification [8] (note that baudrate-sweep is run after both RX-equalization and TX-equalization parameters are adjusted by performing an 'initial_half_effort' adaptaion and an 'auto-sweep' over vod parameter, respectively (as stated in 6.1)).
- Then a time-limited test will be run for 2h, at the maximum achievable baudrate found, with temperature monitoring enabled; thus allowing to obtain both temperature and BER progression curves over test time. BER against temperature curve will be also presented for each link tested to analyze xcvr performance as on-die temperature increases for the E-tile xcvr in use.
- All tests described are run for both NRZ and PAM4 (at same baudrate⁷²).
- TTK will be used for each test run to generate the eye-diagram view⁷³ for signal integrity analysis.

6.2.1 sub-section describes serial link set-ups tested providing brief detail on how E-tile xcvr's channels map to the board's connector used (note that further detail on channel to connector-lane correspondence is given in B)

6.2.1. TEST DESCRIPTION

6.2.1.1 INTERNAL SERIAL LOOPBACK VALIDATION TEST

Initial test was run over both channel0-(PHY0) and channel0-(PHY1) being both configured with default parameter values, (thus with no TX or RX equalization), in internal serial loopback mode. Serial link tests require 'internal serial loopback' mode to be disabled so that the signal will be output to the connector to be inserted in the channel under tests (thus initial 'internal serial loopback' tests was run only for validation purposes, and the rest of the tests deccribed in this section are run with 'internal serial loopback' mode disables for the PHY channel used).

'internal serial loopback' initial test is run for system validation purposes but also to analyze the minimum measurable BER for 2 h. serial link test; that will serves as a reference for BER results analysis in sub-sequent serial link tests (note that in internal serial loopback the eye-diagram measured

⁷² note that E-tile xcvr's support baudrates up to 28.9Gbps, so NRZ transmissions over 30Gbps cannot be run to compare PAM4/NRZ performances at same bitrate. So that comparison is limited to performance observed at same baudrate

⁷³ statistical eye-diagram is an indicator of both the amplitude and time noise/fluctuations supported while still keeping the BER measured over the serial link under test below certain BER value

corresponds to a distortion-free channel test, apart from the thermal noise introduced by xcvr's active elements that will be common for all tests run over E-tile xcvrs).

6.2.1.2 QSFP-DD 1.5 M LINK TEST

6.3 shows the set-up for a serial link test over a COM compliant QSFP-DD cable of 1.5 m length (refer to [22] for product specification from vendor), under conditions described for QSFP-DD 0.5 m serial link test 6.2.1.2.

(QSFP-DD cable (2015913015 from Mollex) is connected between same 2 QSFP-DD modules)

QSFP-DD 1.5 m test results (for channel0 (PHY0) – channel6 (PHY0) link) and xcvr settings used in test (including RX-equalization and TX-equalization settings ,after the optimization run before test's start) are analyzed in 6.2.2.2.




Figure 6.3.- Set-up for PAM2/PAM4 serial link test over QSFP-DD cable of 1.5 m length (2015913015 from Mollex , [22]) connecting 2 e-tile xcvr channels on adjacent QSFP-DD modules.

6.2.1.3 QSFP-DD 0.5 M LINK TEST

6.4 shows the set-up for a serial link test over a COM compliant QSFP-DD cable of 0.5 m length (refer to [22] for product specification from vendor).

TX_serial_data and RX_serial_data signals of channels 0,1,4,5 in PHY0 are connected to one QSFP-DD module on Stratix 10TX board , and channels 6,7,10,11 in PHY0 to another QSFP-DD module in same PHY.

QSFP-DD cable (2015911005 from Mollex) is connected between those 2 QSFP-DD modules resulting in following connections (refer to B for detail on 'outer' connections):

- channel0 (PHY0) – channel6 (PHY0)
- channel11 (PHY0) – channel7 (PHY0)
- channel4 (PHY0) – channel10 (PHY0)
- channel5 (PHY0) – channel11 (PHY0)

QSFP-DD 0.5 m test results (for channel0 (PHY0) – channel6 (PHY0) link) and xcvr settings used in test (including RX-equalization and TX-equalization settings ,after the optimization run before test's start) are analyzed in 6.2.1.3.




Figure 6.4.- Set-up for PAM2/PAM4 serial link test over QSFP-DD cable of 0.5 m length (2015911005 from Mollex, [22]) connecting 2 e-tile xcvr channels on adjacent QSFP-DD modules.

6.2.1.4 QSFP 1.0 M LINK TEST

6.5 shows the set-up for a serial link test over a COM compliant QSFP cable of 1.0 m length (refer to [23] for product specification from vendor), under conditions described for QSFP-DD serial link tests 6.2.1.3, 6.2.1.2.

QSFP-DD modules are compatible with QSFP connections; and compatibility is kept by using only half of the available channels on QSFP-DD module when a QSFP cable is plugged in. So that, only the following channels 'remain' connected through the external QSFP cable :

- channel4 (PHY0) – channel10 (PHY0)
- channel5 (PHY0) – channel11 (PHY0)

QSFP 1.0 m cable (1002971101 from Mollex) is connected between same 2 QSFP-DD modules)

QSFP 1.0 m test results (for channel4 (PHY0) – channel10 (PHY0) link) and xcvr settings used in test (including RX-equalization and TX-equalization settings ,after the optimization run before test's start) are analyzed in 6.2.2.4.




Figure 6.5.- Set-up for PAM2/PAM4 serial link test over QSFP cable of 1 m length (1002971101 from Mollex) connecting 2 e-tile xcvr channels on adjacent QSFP-DD modules.

6.2.1.5 SMA 2.4 MM RF LINK TEST

6.6 shows the set-up for a serial link test over 2.4mm SMA cables of 0.7 m length⁷⁴ (6.7 shows S11 parameter from provider specific test report).

TX_serial_data and RX_serial_data signals of channel2 in PHY1 are connected to SMAA_TX and SMAA_RX ports shown in 6.6; and external loopback is done using SMA 2.4mm cable (SF101E/11PC24/11PC24/700mm from hubersuhner)

SMA 2.4 mm test results and xcvr settings used in test (including RX-equalization and TX-equalization settings ,after the optimization run before test's start) are analyzed in 6.2.2.5.




Figure 6.6.- Set-up for PAM2/PAM4 serial link test over 2.4mm SMA cables of 0.7 m length (SF101E/11PC24/11PC24/700mm from hubersuhner) connecting e-tile xcvr channel in external loopback.




Figure 6.7.- S11 parameter for 0.7m length 2.4mm SMA cables SF101E/11PC24/11PC24/700mm from hubersuhner (parameter curve obtained from provider specific test report).

⁷⁴ note that LVDS(low voltage differential signaling port) ports are used in PAM4 capable E-tile xcvs, so that 2 SMA cables are required



6.2.1.6 QSFP 1.0 M + CONNECTION BOARD + SMA LOOPBACK (WORST CASE LINK TEST)


6.8a shows the set-up for a serial link test over a link composed by QSDP 1.0 m cable from Stratix 10TX board to a connection board used to perform a loopback connection between TX/RX belonging to the same channel of the QSFP module (note that loopback connection is not done directly over the connection board, but through an external SMA cable link, as shown in 6.8b).

QSFP cable has only one side connected to Stratix 10TX board; and the 2 available channels are in external loopback, thus setting up the following channels:

channel4 (PHY0) – channel4 (PHY0)
channel5 (PHY0) – channel5 (PHY0)


6.9 shows S11 parameter from Rohde & Schwarz specific test report for the 1.0 m SMA cables used to perform external loopback connection between connection board's ports.

6.2.2.6 analyzes test results for channel4 (PHY0) – channel10 (PHY0) link and xcvr settings used in test (including RX-equalization and TX-equalization settings ,after the optimization run before test's start).



(a)

(b)



(c)

Figure 6.8.- Set-up for PAM2/PAM4 worst case serial link test over channel composed as QSFP cable of 1 m length (1002971101 from Mollex,[30]) + <100 mm connection-board trace + external loopback over connection-board + <100 mm connection-board trace (in return to RX-port)+ QSFP cable of 1 m length (1002971101 from Mollex,[23]) (in return to RX-port)(a); signal path for serial link test over QSFP 1.0m + connection-board trace+ connection-board + connection-board trace + QSFP 1.0m (note that return trace after external loopback over SMA cables is marked in orange)(b); connection board detail (c).




Figure 6.9.- S11 parameter for 1.0m length SMA cables used for external serial loopback over connection board (parameter curve obtained from Rohde & Schwarz specific test report).

6.2.2. TEST RESULTS ANALYSIS

In this section, the results corresponding to the serial link tests described in 6.2.1 sub-section are presented and analyzed.

PRBS31 test pattern was used for all transmission tests conducted over the serial links described.

6.1 shows S10 E-Tile transceiver configuration parameters' values used for serial link testing (note that values that are not common to all the tests conducted will be given in the corresponding test-analysis sub-section).

For each link tested, a serial high-speed transmission was conducted for 2 h., at the maximum baudrate supported by the link under test. That maximum baudrate value was determined before starting the 2 h. serial link test, by performing short PAM4 transmissions over the target serial link, and analyzing the BER observed for a 358 s transmission.

BER values measured during 358 s duration transmissions were taken for each baudrate supported by S10 E-Tile xcvr's in the range 14 – 28 GBs (baudrate-sweep). The maximum baudrate used in test is determined as the maximum rate for which the BER value measured over 358 s remained under the $1E-4$ CEI-OIF limit for long range channels⁷⁵

In order to compare PAM4 and NRZ in terms of performance, each test was conducted under same conditions (same parameters' values and same baudrate) using NRZ.

For each test NRZ and PAM4 eye diagrams are obtained (when possible); and BER and temperature progressions over time are represented.

As same baudrate is used for both NRZ and PAM4, eye-diagrams obtained will have equal UI values ($UI = 1/\text{baudrate}$), thus allowing to establish a direct comparison between eyes' heights and widths.

PAM4 eye-diagrams obtained show 2 PAM4 eye-diagrams as TTK represents 2 consecutive UIs, but eye width measured (and used for performance comparison), correspond to 1 single symbol interval (UI).

2 h. test duration was selected as compromise solution for experimental BER measurements.

BER calculation as (no. of error bits)/(no. of sent bits) may require long run tests for the measured BER value to be stabilized, particularly for low distortion serial links (with low amount of bit errors, or even no error occurrence).

⁷⁵ note that baudrate-sweep measures BER for 358 s, so as to account for at least $1E13$ bits sent over the link when PAM4 is used (for the minimum baudrate considered, 14GBs, the amount of bits sent during 358 s using PAM4 is given by $(14E9) \cdot (2) \cdot (358) = 1.0024E13$). So that enabling to measure BER up to $1E-13 < 1E-4$ CEI-OIF target limit.



BER stabilization over the time can be easily understood when considering the case of zero-error transmission, for which BER is given by its 0.95 confidence level value, $BER(CL = 0.95) = 3/(\text{no. of sent bits})$:

assuming that channel remains error-free over the time, $BER(CL = 0.95)$ value depends only on the no. of sent bits. So that if a test aims to determine if BER for the channel under test is below $1E-16$, e.g., $3/1E-16 = 3E16$ bits must be sent at least.

Thus, the lower the BER target, the longer test-time is required.

If low enough BER value is achieved (test runs for enough time), $BER(CL = 0.95)$ seems to stabilize over the time (in logarithmic scale representation), as great amount of sent bits are required to lower the value.

BER curves obtained from serial link tests are represented using logarithmic y-axis graphs.

All tests were performed under fan cooling conditions for security reasons, keeping board's main fan uninterruptedly on.

Following sub-sections include an analysis on the results obtained from tests described in 6.2.1, performed under the given conditions.

6.2.2.7 summarizes the results obtained, including performance comparison for the serial links analyzed.

Table 6.1.- Common xcvr and test configuration parameter values used for high-speed serial link testing.

(footnotes ⁷⁶, ⁷⁷ (corresponding to superscripts 1,2 in table, respectively) give clarification on aspects uncovered within this table).

XCVR CONFIGURATION PARAMETERS	
rx-termination	VCC (LVDS (low voltage differential signalling) ports with 100Ω differential impedance termination)
serial loopback	0 (disabled)
reverse loopback	0 (disabled)
line encoding	NRZ PAM4 (both tested for each serial link under test)
gray encoding	0 for NRZ (disabled) 1 for PAM4 (enabled)
1/(1+D) encoding	0 (disabled)
swizzle	0 for NRZ (disabled) 1 for PAM4 (enabled)
invert tx polarity	0 (disabled)
invert rx polarity	0 (disabled)
adaptation mode	Initial adaptation

XCVR EQUALIZATION PARAMETERS	
rx-PMA equalizer pre-set	56G-VSR ⁷⁶
(tx-equ) VOD	detailed for each test-case
(tx-equ) pre-tap 1	0 (no equalization)
(tx-equ) pre-tap 2	0 (no equalization)
(tx-equ) pre-tap 3	0 (no equalization)
(tx-equ) post-tap 1	0 (no equalization)

TEST CONFIGURATION	
test limit condition	test_time = 7201 s (2 h.)
baudrate	detailed for each test-case
rx-channel (PHY)	detailed for each test-case ⁷⁷

6.2.2.1 INTERNAL SERIAL LOOPBACK VALIDATION TEST

Internal serial loopback test results for 28GBs PAM4 and NRZ transmissions are presented in figures 6.10 - 6.12.

In Internal serial loopback mode the transmitted signal remains inside the xcvr die, thus avoiding the distortion introduced by both cable and connectors. So that, the Internal serial loopback test can be understood as minimum signal distortion test-case, serving as reference for results analysis in sub-sequent serial link tests; and as an indicator of xcvr performance.

For Internal serial loopback mode test, baudrate-sweep to find maximum achievable baudrate was not performed as support for both NRZ and PAM4

⁷⁶ see table 5.5 for 56G-VSR RX-equalizer pre-set values for each equalization parameter.

⁷⁷ note that indicated channel is the channel monitored acting as receive (the corresponding channel at opposite cable edge is used as TX



successful transmissions⁷⁸ up to 28 GBs can be assumed in zero-distortion conditions [11].

6.10 shows both NRZ and PAM4 BER measured over over the 2h. test.

0 bit-errors were reported for both tests (as expected for internal serial loopback mode test) so BER curve shown corresponds to BER (CL = 0.95)=3/(no. of bits sent) calculated values; thus resulting into negative exponential shaped curves (when represented using logarithmic y-axis).

BER curves obtained for zero distortion serial transmission reflect 'BER stabilization' over test time.

BER observed values are lower for PAM4 transmission as the no. of bits sent for PAM4 signaling over test time is double the no. of bits sent when using NRZ at same baudrate; thus resulting into lower BER (CL = 0.95) values even when both NRZ and PAM4 transmissions are error-free.

BER curves shown in 6.10 for both NRZ and PAM4 BER (reaching minimum BER values of 7.4394e-15 and 1.4878e-14 at the end of the 2 h. test, respectively) will be used in sub-sequent section as the minimum reference achievable values, for measured BER analysis.

6.11 shows xcvr's on-die temperature progression over test-time; revealing slight on-die temperature variations due to high-frequency transmissions activity.

On-die temperature tracking for internal serial loopback transmissions would allow to determine the effect of temperature increases on xcvr performance (BER degradation with temperature increase would be caused only by xcvr performance degradation without the influence of channel response variations for internal serial loopback test).

6.12 graphic shows calculated BER for each temperature value sampled, enabling for BER-temperature analysis. BER-temperature results shown in for tests conducted cannot be considered significant as no major on-die temperature variations were observed for any of the tests performed; thus preventing for analysis on temperature effect over measured BER.

Slight temperature variations observed for tests conducted is explained by the cooling fan integrated over the board and kept active over test duration, mitigating xcvr heating.

Further tests for BER-temperature analysis with absence of fan cooling could not be performed by the moment of writing this document as the controlled temperature chamber was not available, thus preventing from performing under secure controlled temperatures to avoid damage on S10 board.

PAM4 and NRZ eye diagrams are not presented for internal serial loopback test, as TTK does not support eye-measurement in internal serial loopback mode for S10 devices.




Figure 6.10.- PAM4 and NRZ BER progression over 2 h. test time for reference no-distortion internal serial loopback mode test at 28GBs baudrate.

⁷⁸ transmissions with BER <1E-4 CEI-OIF limit for LR channels




Figure 6.11.- E-tile xcvr on-die temperature progression over 2 h. test time for reference no-distortion internal serial loopback mode test at 28GBs baudrate.




Figure 6.12.- BER against xcvr on-die temperature progression over 2 h. test time for reference no-distortion internal serial loopback mode test at 28GBs baudrate.

6.2.2.2 QSFP-DD 1.5 M LINK TEST

In this sub-section, results for test transmissions performed over QSFP-DD 1.5 m serial link are described (corresponding to the test set-up described in sub-section 6.2.1.2).

6.13 figure shows the results obtained for baudrate-sweep performed over the target channel, where BER measured is presented for each supported baudrate value (note that absence of BER data for baudrates below 24.08 GBs correspond to error-free BER measurements.)

BER $\approx 1E-8$ was measured for 28GBs transmission over the 358s measuring interval, thus demonstrating support for 56Gbps PAM4 transmissions over the QSFP-DD 1.5 m serial link under test.

So that, the results presented in this section were obtained for 28GBs serial link tests performed over the target serial link.

BER measurements for increasing baudrate shown in 6.13, reflect the xcvr degradation performance for PAM4 transmissions when carrier signal frequency is increased, thus providing an insight on the cost of increasing the transmission bitrate when using PAM4.

Serial link tests were conducted with TX-equalized configured using vod = 3 (optimum vod value for minimum BER over the link under test, obtained by performing vod-sweep before test).

6.14 figure shows the BER progression over test duration, revealing a significantly higher BER value for PAM4 transmission as expected (note that for same baudrate, the signal frequency dependent losses are the same for both NRZ and PAM4 transmissions; but PAM4 vulnerability to signal attenuation



due to its less spaced signal levels, causes greater ISI affecting PAM4 signaling, thus higher BER).

BER shown in figure 6.14 remains almost equal to the ideal distortion-free NRZ BER shown in figure 6.10 for internal serial loopback test, thus revealing NRZ resilience to distortion introduced by the channel under test for 28GBs transmissions.

In contrast to NRZ BER, PAM4 BER shows a great increase (up to $BER \approx 1E-8$) in comparison to the ideal distortion free PAM4 BER $\approx 1E-14$ measured in internal serial loopback test; thus reflecting signal degradation over the channel under test, but still remaining under $BER = 1E-4$ CEI-OIF limit.

6.17 , 6.18 figures show both PAM4 and NRZ measured eye-diagrams , respectively, with measured EH6 , EW6 values:

PAM4: EH6 = 7 , EW6 = 30

NRZ: EH6 = 176 , EW6 = 36

EW6 measured values are similar for NRZ and PAM4 transmissions as expected (note that for same baudrate, the 'horizontal' ISI affecting both signals is almost the same, as PAM4 eyes' width use to be at least 2/3 the width of NRZ eyes at same baudrate).

PAM4 EH6 is on the contrary much shorter than the ideal 1/3 of NRZ EH6 measured, thus revealing the great PAM4 vulnerability to signal attenuation over the channel; explaining PAM4 BER degradation shown in figure 6.14, in comparison to the ideal PAM4 BER shown in 6.10 .

6.15- 6.16 figures show on-die temperature tracked over test duration, revealing again slight temperature variations (as explained in 6.2.2.1 sub-section for internal serial loopback test).

6.16 figure shows different BER values being measured for constant on-die temperature (53°C for NRZ , and 55°C for PAM4), thus confirming BER independency on on-die temperature, for such slight die heating.




Figure 6.13.- BER measured over 358 s. against baudrate, for baudrate-sweep over QSFP-DD 1.5m-length serial link test.




Figure 6.14.- BER measured over 358 s. against baudrate, for baudrate-sweep over QSFP-DD 1.5m-length serial link test.





Figure 6.15.- PAM4 and NRZ BER progression over 2 h. test time for QSFP-DD 1.5m-length serial link test at 28GBs baudrate.




Figure 6.16.- BER against xcvr on-die temperature progression over 2 h. test for QSFP-DD 1.5m-length serial link test at 28GBs baudrate.




Figure 6.17.- PAM4 eye diagram measured for QSFP-DD 1.5m-length serial link test at 28GBs baudrate (EH6 = 7, EW6 = 30).




Figure 6.18.- NRZ eye diagram measured for QSFP-DD 1.5m-length serial link test at 28GBs baudrate (EH6 = 176, EW6 = 36).

6.2.2.3 QSFP-DD 0.5 M LINK TEST

6.19 - 6.23 figures show the results obtained for serial link test performed over QSFP-DD 0.5 m cable as described in 6.2.1.3 sub-section.

QSFP-DD 0.5 m serial link tests for both NRZ and PAM4 were conducted for a TX-equalization vod value: vod = 2 (optimum parameter value for minimum BER, obtained from vod-sweep conducted over the channel under test).

NRZ and PAM4 transmissions were both performed at 28GBs baudrate (note that baudrate-sweep curve is not included here as results are quite similar to the results obtained for QSFP-DD 1.5 m cable (show in 6.10)).

6.19 figure shows BER progression over time, revealing a value of BER $\approx 7E-10$ for PAM4 transmission test, lower than the $1E-8$ value measured for QSFP-DD 1.5 m length; thus reflecting the effect of greater signal attenuation accross longer cables on PAM4 BER, due to its vulnerability to signal strength loss (note that both QSFP-DD 0.5 m and QSFP-DD 1.5 m cables belong to same manufacturer and technologie series, thus allowing for fair attenuation effect comparison).

In contrast, NRZ BER remains almost unaffected by channel distortion, showing BER values close to the ideal distortion-free results shown in 6.10 (as expected if considering the results presented in 6.2.2.2 for 1.0 m longer cable).

6.22 , 6.23 figures show both PAM4 and NRZ measured eye-diagrams , respectively, with measured EH6 , EW6 values:

PAM4: EH6 = 8 , EW6 = 32
NRZ: EH6 = 175 , EW6 = 36

PAM4 EH6 and EW6 values measured, slightly greater than those measured for test conducted over QSFP-DD 1.5 m cable, explain lower PAM4 BER measured over shorter QSFP-DD 0.5 m cable; while reflecting to what extent an slightly greater eye-openness can reduce measured BER in about 2 magnitude orders.

NRZ eye's height and width remain unaffected in comparison to values measured for the longer QSFP-DD 1.5 m (shown in previous sub-section), thus confirming QSFP-DD cables's high-performance for high-speed NRZ transmissions.

6.20 , 6.21 figures show on-die temperature variations plotted against test time and measured BER, respectively.

BER independency on temperature (at least for the slight variations measured) can be assumed from results presented in figures 6.20 , 6.21; analogously to the analyze given in previous sub-sections.




Figure 6.19.- PAM4 and NRZ BER progression over 2 h. test time for QSFP-DD 0.5m-length serial link test at 28GBs baudrate.




Figure 6.20.- E-tile xcvr on-die temperature progression over 2 h. test time for QSFP-DD 0.5m-length serial link test at 28GBs baudrate.




Figure 6.21.- BER against xcvr on-die temperature progression over 2 h. test for QSFP-DD 0.5m-length serial link test at 28GBs baudrate.




Figure 6.22.- PAM4 eye diagram measured for QSFP-DD 0.5m-length serial link test at 28GBs baudrate (EH6 = 9, EW6 = 32).




Figure 6.23.- NRZ eye diagram measured for QSFP-DD 0.5m-length serial link test at 28GBs baudrate (EH6 = 175, EW6 = 36).

6.2.2.4 QSFP 1.0 M LINK TEST

PAM4 and NRZ transmission results for serial link tests over QSFP 1.0 m cable (as described in 6.2.1.4) are shown in figures 6.24 - 6.28.

QSFP 1.0 m serial link tests were conducted with TX vod parameter configured as vod = 8 (obtained from vod-sweep as optimum vod value for minimum BER). PAM4 and NRZ transmissions were both performed at 28GBs baudrate (as explained in previous sub-section, baud-rate curve for QSFP 1.0 m cable is not included here as results obtained for a 358 s BER measurement are quite similar to those presented in sub-section 6.2.2.2 for 1.5 m QSFP-DD cable).

6.24 figure shows PAM4 and NRZ BER progressions over test-time.

PAM4 BER values shown in 6.24 are intermediate values between those obtained for 1.5m-length and 0.5m-length QSFP-DD cables (shown in sub-sections 6.2.2.2, 6.2.2.3, respectively), as expected for a mid-length cable (as signal attenuation affecting PAM4 signaling over the channel increases with cable-length).

NRZ BER, remains also close to the ideal distortion-free curves measured for internal serial loopback test (see 6.2.2.1), as greater space between symbol signal levels, makes NRZ resilient to attenuation through the tested QSFP-DD cables.

6.27, 6.28 figures show both PAM4 and NRZ measured eye-diagrams, holding EH6, EW6 values:

PAM4: EH6 = 8, EW6 = 35

NRZ: EH6 = 181, EW6 = 35

PAM4 EW6, EH6 measured over QSFP 1.0 m are close to values measured for 1.5m-length and 0.5m-length QSFP-DD cables, as expected if considering

the similarity between BER results.

EH6 measured value lies exactly in between values measured for 1.5m-length and 0.5m-length QSFP-DD cables in previous tests, in consistency with the observed BER value, in between BERs measured for both longer and shorter QSFP-DD cables. In the contrary, EW6 shown is greater than EW6 value measured for shorter 0.5m-length QSFP-DD cable (show in previous section); meaning that QSFP 1.0m-length cable tested behaves better in terms of signal jitter ('horizontal' ISI or signal time-delay) over the cable. QSFP 1.0m-length cable belongs to same manufacturer than previously tested QSFP-DD cables, but different manufacturing technology is used⁷⁹, which can explain the unexpectedly greater horizontal eye-openness observed for QSFP 1.0m-length cable.

NRZ eye's height and width, remains similar to those measured for QSFP-DD tested cables, reflecting again NRZ resilience to signal attenuation over such a channel.

6.25, 6.26 figures show on-die temperature progression tracked during the test, reflecting again low temperature variations under the test fan-cooling conditions; and also BER independency on on-die measured temperature, for that low temperature variations.




Figure 6.24.- PAM4 and NRZ BER progression over 2 h. test time for QSFP 1.0m-length serial link test at 28GBs baudrate.




Figure 6.25.- E-tile xcvr on-die temperature progression over 2 h. test time for QSFP 1.0m-length serial link test at 28GBs baudrate.

⁷⁹ QSFP and QSFP-DD cable technologies differ in available lanes per connector, being QSFP-DD a double-density evolution of QSFP technology

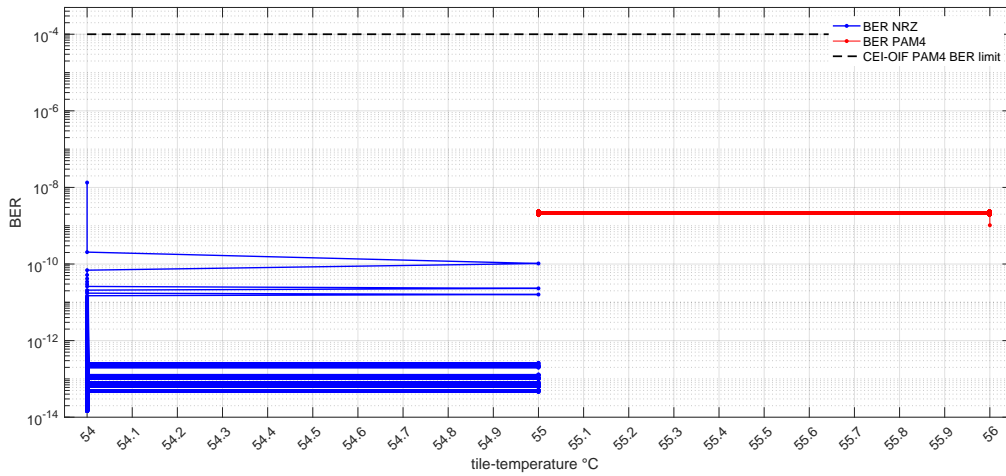


Figure 6.26.- BER against xcvr on-die temperature progression over 2 h. test for QSFP 1.0m-length serial link test at 28GBs baudrate.




Figure 6.27.- PAM4 eye diagram measured for QSFP 1.0m-length serial link test at 28GBs baudrate (EH6 = 8, EW6 = 35).




Figure 6.28.- NRZ eye diagram measured for QSFP 1.0m-length serial link test at 28GBs baudrate (EH6 = 181, EW6 = 35).

6.2.2.5 SMA 2.4 MM RF LINK TEST

6.29 - 6.31 figures show results obtained for serial link test conducted over SMA 2.4 mm RF cables as described in sub-section 6.2.1.5.

Serial transmissions were performed at 28Gbps baudrate for both NRZ and PAM4 signaling, to obtain the results analyzed in this sub-section (note that baudrate-sweep results, from which test baudrate was determined, are not included here as curve obtained for BER measurements over 358s over the SMA 2.4 mm cables are quite similar to those presented for previous tests, see sub-section 6.2.2.2).

Both NRZ and PAM4 transmissions were done with TX configure with vod equalization parameter: vod = 8 (determined as best value for minimum BER, from vod-sweep performed over the channel under test).

6.29 figure shows BER progression over the test time for NRZ and PAM4 test transmissions.

PAM4 BER value reached at the end of the 2 h. test ($BER \approx 1E-13$) is lower than PAM4 BER values obtained for QSFP-DD and QSFP cables analyzed in previous sub-sections; thus revealing greater performance of SMA 2.4 mm cables for PAM4 high-speed serial transmissions.

NRZ BER values also remain close to the ideal zero-distortion values measured for the internal serial loopback test used as reference (as expected if considering SMA 2.4 mm high-performance revealed by low PAM4 BER values measured).

Similarly to the results presented in previous sub-sections, 6.30 and 6.31 figures show short on-die temperature variations over test-time, thus enabling to disregard temperature effect on measured BER, under fan-cooling conditions.

PAM4 and NRZ eye-diagrams could not be measured for this test due to a still unidentified firmware issue.




Figure 6.29.- PAM4 and NRZ BER progression over 2 h. test time for SMA 2.4 mm RF 0.7m-length serial link test at 28Gbps baudrate.




Figure 6.30.- E-tile xcvr on-die temperature progression over 2 h. test time for SMA 2.4 mm RF 0.7m-length serial link test at 28GBs baudrate.




Figure 6.31.- BER against xcvr on-die temperature progression over 2 h. test for SMA 2.4 mm RF 0.7m-length serial link test at 28GBs baudrate.

6.2.2.6 QSFP 1.0 M + CONNECTION BOARD + SMA LOOPBACK (WORST CASE LINK TEST)

PAM4 and NRZ transmission results are analyzed in this sub-section for the worst-case serial link tested (using set-up described in 6.2.1.6)

6.32 figure shows the result obtained for the baudrate-sweep performed over the target channel to determine the maximum supported baudrate (maximum baudrate for which BER measured lies below $1E-4$ CEI-OIF BER limit).


22.4 GBs baudrate can be determined from figure 6.32, as the maximum PAM4 signaling baudrate that guarantees BER remaining below $1E-4$ limit (note that the absence of BER measurements beyond 22.4 GBs in figure 6.32 means that RX was unable to lock to the PRBS31 sequence received due to the great amount of errored bits caused by signal distortion).

So that, test results presented in this section correspond to serial link transmissions performed at 22.4 GBs baudrate for both NRZ and PAM4 signal modulations.

Serial link tests were conducted with TX configured with vod parameter value: vod = 2, determined from vod-sweep as optimum value for minimum BER.

6.33 figure shows BER progression over test time for both NRZ and PAM4 tests conducted. NRZ BER measured BER values, remain close to the ideal no-distortion BER curve measured in internal serial loopback mode; while PAM4 BER values show a significant degradation, increasing up to $BER \approx 1E-5$.

PAM4 BER values observed increase to BER levels close to the $1E-4$ CEI limit, but remain constant over test time, thus enabling to confirm that PAM4



transmissions are viable over the worst-case channel up to 22.4 GBs.

6.34 and 6.35 figures show that temperature variations remain also insignificant, in terms of signal degradation, even for the worst-case link tested.

PAM4 and NRZ eye-diagrams could not be measured for this test due to a still unidentified firmware issue, as for previous test over SMA 2.4 mm cable.




Figure 6.32.- BER measured over 358 s. against baudrate, for baudrate-sweep over 'connection-board loopback set-up' serial link test.




Figure 6.33.- PAM4 and NRZ BER progression over 2 h. test time for 'connection-board loopback set-up' serial link test at 22.4GBs baudrate.




Figure 6.34.- E-tile xcvr on-die temperature progression over 2 h. test time for 'connection-board loopback set-up' serial link test at 22.4GBs baudrate.




Figure 6.35.- BER against xcvr on-die temperature progression over 2 h. test for 'connection-board loopback set-up' serial link test at 22.4GBs baudrate.



6.2.2.7 TEST RESULT SUMMARY

Serial link test results analyzed are summarized in this section, including a brief overall performance comparison between tested links.

6.2 table summarizes BER values observed for serial link tests conducted, where '2 h. BER' columns refer to BER values reached after 2 h. of test transmission; and 'internal serial loopback' row collects results for the ideal no-distortion test performed in internal serial loopback mode as reference for the rest of tests conducted.

6.36 figure shows PAM4,NRZ overlapped BER curves measured for all tested channels , including reference BER traces measured for internal serial loopback mode (note that NRZ BER curves that are not visible over graphic area are stacked below the curve referenced to as 'BER NRZ CONN-BOARD').

The following statements can be concluded from the test result analysis given for each serial link test conducted:

- BER degradation as cable length increases for QSFP-DD/QSFP cables was observed (see overlapped BER curves in figure 6.36).
- SMA 2.4 mm show higher performance for 56Gbps PAM4 transmissions than QSFP-DD cables, even for shorter ones (see overlapped BER curves in figure 6.36).
- viability for PAM4 transmissions up to 22.4 GBs (44.8 Gbps) over worst-case link was demonstrated, while keeping BER under 1E-4 CEI limit.
- NRZ resilience to channel distortion at maximum baudrate tested (28 GBs) was observed for the serial links tested, keeping NRZ BER values close to the ideal no-distortion values measured for internal serial loopback test.
- temperature low variability under fan-cooling conditions was demonstrated; and so BER independency on temperature variations up to 2 °C (note that further tests in temperature controlled environment, with board's fan off, should be done to analyze temperature variations effect on measured BER)




Figure 6.36.- PAM4, NRZ overlapped BER progressions over 2 h. test time for serial link tests performed (note that all BER curves correspond to 28GBs baudrate tests, except 'CONN-BOARD' curves that correspond to 22.4 GBs maximum supported baudrate for the connection-board link set-up tested).



Table 6.2.- Summary of BER results for serial link tests performed.

(footnotes ⁸⁰, ⁸¹, ⁸², ⁸³, ⁸⁴ in table give clarification on aspects uncovered within this table).

	(max) TEST BAUDRATE ⁸⁰	2 h. BER PMA4	2 h. BER NRZ	min BER PMA4	min BER NRZ
QSFP-DD 1.5 m	28 GBs	1.3865e-08	1.4878e-14	2.9597e-10 ⁽⁸¹⁾	1.4878e-14
QSFP-DD 0.5 m	28 GBs	7.4084e-10	7.4390e-15	6.7812e-10	7.4390e-15
QSFP 1.0 m	28 GBs	2.2357e-09	1.4879e-14	1.0302e-09	1.4879e-14
CONNECTION BOARD LOOPBACK	22.4 GBs ⁽⁸⁴⁾	1.8407e-05	1.8597e-14	1.9531e-07	1.8597e-14
SMA 2.4 mm RF	28 GBs	9.1751e-13	1.4879e-14	5.1323e-13	1.4879e-14
INTERNAL SERIAL LOOPBACK	28 GBs	1.4878e-14 ⁽⁸³⁾	7.4394e-15 ⁽⁸³⁾	1.4878e-14 ⁽⁸²⁾	7.4394e-15 ⁽⁸²⁾


⁸⁰ note that for PAM4 tests bitrate is calculated as 2xbaudrate.

⁸¹ PAM4 BER measured over QSP-DD 1.5 m cable shows lower BER values initially, that could correspond to a short period of few symbol transitions in the PRBS31 pattern sent, causing lower ISI.

⁸² note that for NRZ 4 out of 5 tested serial links reach ideal no-distortion BER value measured for internal serial loopback test; but for PAM4, signal vulnerability to channel attenuation affects signal integrity over all tested serial links, preventing measured BER values to reach the no-distortion BER value measured for internal serial loopback test.

⁸³ note that internal serial loopback run without errored bits, so BER values presented correspond to the calculated BER(CL = 0.95) that decreases linearly with the amount of bits sent (thus lineally with test time), explaining the coincidence between 'min. BER' and '2 h. BER' values.

⁸⁴ note that serial link tests over 'connection-board loopback set-up' were performed at lower baudrate = 22.4 GBs.



7. CONCLUSIONS AND FUTURE WORK

7.1. CONCLUSIONS

In this work, several serial links are tested in terms of BER performance for high-speed PAM4 serial transmissions, demonstrating the feasibility for such transmissions even in the worst case channel analyzed.

PAM4 high-speed successful transmissions with BER $<1E-4$ CEI-OIF limit for 56G-LR channels were achieved up to up to 24.53 GBs with partially optimized TX/RX equalization parameters.

PAM4 tests were also performed under same conditions but using NRZ signaling at same baudrate to provide an analysis on signal integrity degradation penalty when moving to PAM4 to double data rate over serial links.

Intel Stratix 10TX board was used to implement a dual serial-link testing design capable for both PAM4 and NRZ signaling 14-28GBs, to support the PAM4 high-data rate signaling analysis; achieving reduced compilation time design achieved by using only 2 out of 5 available PHYs over Stratix 10TX board

PAM4 testing system was designed to allow for real-time xcvr and test parameters configuration through Nios-2 console and support auto-sweep based calibration to automate xcvr equalization settings optimization

PAM4 transmission experiments were carried out using PRBS31 CEI-OIF test pattern. PRBS31 128-b parallel generator/checker were developed to support the data pattern generation and checking. PRBS31 128-b parallel generator/checker design validation tests' results analysis confirms their capability for generating and checking QPRBS31 pattern over 30Gbps. PRBS31 128-b parallel checker was designed to be able for bitwise transmission errors location, outputting the exact error pattern for further statistical analysis.

Furthermore, a 128-b parallel error analyzer was developed to obtain statistical information on burst length and density, and gap length from the received signal. The statistical error analyzer is also capable to generate statistics on 5440-b length block basis, considering limited windows of the received signal. The statistical error analyzer was not integrated in the PAM4 test system, but promising results were obtained for simulation based validation tests.

7.2. FUTURE WORK

From the work presented in this document, further steps aiming to enrich PAM4 serial link analysis are derived.

temperature control

Integration of Stratix 10TX fan control to enable for running long-time tests keeping control over on-die temperature on E-tile xcvs; enabling to analyze to what extent PAM4 transmissions' performance is affected by temperature increases.

Intel's asserts about E-tile xcvs' capability to adjust equalization settings automatically to compensate non-drastic temperature changes, when configured in continuous adaptation mode, can be found in [11]. PAM4 transmission tests may be run in temperature controlled environment with E-tile xcvr configured in continuous adaptation mode to determine to what extent the equalization structure is able to compensate channel distortion caused by temperature changes.

FEC integration

Integration of statistical error analyzer developed in the system designed to enable the analysis on RS-FEC(544,514) capabilities to correct transmission errors.

Investigation on how interleaving (D.1.1.2 sub-section in D) enhances RS-FEC performance to determine the actual BER limit for the custom channel to be used.

Investigation on low effort FEC approach proposed in PCIe 6.0 release to determine the feasibility of replacing RS-FEC(544,514).

PAM4 link testing system enhancement

PAM4 testing system modification to automate thorough equalization settings optimization, performing parameter sweeping for every equalization parameter (considering both TX and RX equalization settings).

PAM4 serial link controller implementation can be derived from the previous step, aiming to develop a serial transmission controller able to adapt RX-equalization settings to channel distortion changes; thus extending xcvr capability to recognize and overcome channel response degradations.






REFERENCES

- [1] “PAM4 Signaling for 56G Serial Link Applications – A Tutorial”; Hongtao Zhang, Brandon Jiao, Yu Liao, and Geoff Zhang - XILIX ; at DesignCon 2016, JANUARY 19-21, 2016 ; <https://www.xilinx.com/publications/events/designcon/2016/slides-pam4signalingfor56gserial-zhang-designcon.pdf> ; last access July 2020.
- [2] “PAM4 signaling fundamentals”; Intel AN-835 (application note-835), 2019.03.12. .
- [3] “PCIe 6.0 Rev 0.5 Electrical Spec Review”; Mohiuddin Mazumder – Intel Corporation ; at PCI Express Base Specification Revision 6.0 Version 0.5 , PCI-SIG Draft Spec. 30 January 2020.
- [4] “112Gbps Serial Transmission over Copper – PAM4 vs PAM8 Signalling”; Min Wu , Kelvin Qiu, Geoff Zhang ; at DesignCon 2017, JANUARY 31-FEB, 2017 ; <https://www.xilinx.com/publications/events/designcon/2017/112gbps-serial-transmission-over-copperpam4-vs-pam8-slides.pdf> ; last access July 2020.
- [5] J. Im et al., “A 40-to-56 Gb/s PAM-4 Receiver With Ten-Tap Direct Decision-Feedback Equalization in 16-nm FinFET”; in IEEE Journal of Solid-State Circuits, vol. 52, no. 12, pp. 3486-3502, Dec. 2017, doi: 0.1109/JSSC.2017.2749432.
- [6] G. Steffan et al., “6.4 A 64Gb/s PAM-4 transmitter with 4-Tap FFE and 2.26pJ/b energy efficiency in 28nm CMOS FDSOI,”; 2017 IEEE International Solid-State Circuits Conference (ISSCC), San Francisco, CA, 2017, pp. 116-117, doi: 10.1109/ISSCC.2017.7870288.
- [7] “Intel Stratix 10TX Device Overview”, S10-TX-OVERVIEW, 2020.03.24 ;<https://www.intel.com/content/www/us/en/programmable/documentation/jzw1474049428757.html#joc1431448703837>; last access July 2020.
- [8] “Common Electrical I/O (CEI) - Electrical and Jitter Interoperability agreements for 6G+ bps, 11G+ bps, 25G+ bps I/O and 56G+ bps ”; OIF-CEI-04.0 ; December 29, 2017.
- [9] “Combating Closed Eye – Design & Measurement of Pre-Emphasis and Equalization for Lossy Channels”; BERTScope , Tektronix
- [10] “Equalization to open eye ”; Ruey-Beei Wu - S. H. Hall & H. L. Heck, High-Speed Digital Designs, ch. 12..
- [11] “E-Tile Transceiver PHY User Guide, UG-20056”; Intel ; 2020.06.02 ; https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug_etile_xcvt_phy.pdf ; last access July 2020.
- [12] “Chip-to-module 100 Gb/s four-lane Attachment Unit Interface (CAUI-4)”; IEEE Draft P802.3bm/D3.1 ; 22nd July 2014. .
- [13] “IEEE Standard for Ethernet - Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters for 200Gb/s and 400Gb/s Operation,”; in IEEE Std 802.3bs-2017 (Amendment to IEEE 802.3-2015 as amended by IEEE’s 802.3bw-2015, 802.3by-2016, 802.3bq-2016, 802.3br-2016, 802.3bn-2016, 802.3bu-2016, 802.3bv-2017, and IEEE 802.3-2015/Cor1-2017), vol., no., pp. 1-372, 12 Dec. 2017, doi: 10.1109/IEEESTD.2017.8207825 .
- [14] “Intel Stratix 10 Analog to Digital Converter User Guide ,Updated for Intel Quartus Prime Design Suite: 19.3 , UG-S10ADC ”; Intel ; 2019.10.09 ;<https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/ug-s10-adc.pdf> ; last access July 2020.
- [15] “Secure Device Manager for Intel Stratix 10 Devices Provides FPGA and SoC Security”, white paper; Ting Lu, Ryan Kenny , Sean Atsatt - Intel Corporation ;<https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/wp/wp-01252-secure-device-manager-for-fpga-soc-security.pdf> ; last access July 2020.
- [16] “Intel FPGA E-Tile Transceiver Basics”; Intel FPGA training course
- [17] “Intel Stratix 10 TX Transceiver Signal Integrity Development Kit User Guide, UG-20150”; Intel ; 2019.07.27 ; <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug-intel-s10-tx-devl-kit.pdf> ; last access July 2020.
- [18] “Avalon Interface Specifications”, Updated for Intel Quartus Prime Design Suite: 20.1 , MNL-AVABUSREF; 2020.05.26 ;https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/manual/mnl_avalon_spec.pdf
- [19] “Avalon Memory-Mapped Interface Specifications ”, version 3.3 ; Altera; May 2017
- [20] G. Haßlinger and O. Hohlfeld, “Analysis of random and burst error codes in 2-state Markov channels,”; 2011 34th International Conference on Telecommunications and Signal Processing (TSP , Budapest, 2011, pp. 178-184, doi: 10.1109/TSP.2011.6043747. .
- [21] “How Do I Measure the Bit Error Rate (BER) to a Given Confidence Level on the J-BERT M8020A and the M8040A High-Performance BERT? ”; J-BERT M8020A and M8040A High-Performance BERT series guides - Keysight Technologies ; <https://www.keysight.com/main/editorial.jsp?ckey=1481106&id=1481106&nid=-11143.0.00&lc=ger&cc=DE>; last access July 2020.



- [22] “Product Specification WSFP-DD Passive Copper Assembly”, , revision A, documento no. 2015910001 PS 000, EC No.:613233, Date 04/11/2017 ; https://www.molex.com/pdm_docs/ps/2015910001-000.pdf ; last access July 2020.
- [23] “100297-1101 zQSFP+ Cable Assembly Test Summary”, , revision 1.0, document no. TS - 100297.-1101 PS 000, Date 2015/10/16 ; https://www.molex.com/pdm_docs/ts/TS-100297-1101-001.pdf ; last access July 2020.
- [24] “PAM4 Signaling in High-Speed Serial Technology: Test, Analysis, and Debug ”; APPLICATION NOTE, Tektronix.
- [25] “Understandig the Pre-emphais and Linear Equalization Features in Stratix IV GX devices”; Altera corporation , application note AN-602-1.0 ; November 2010.
- [26] “Pre-emphasis equalization, High-Speed Circuits and Systems Lab. ”; Yonsei University ; http://tera.yonsei.ac.kr/class/2016_1_2/lecture/Design%208%20Pre-emphasis.pdf ; last access July 2020.
- [27] “Lecture 7: Equalization Introduction & TX FIR Eq ”; San Palermo Analog & Mixed-Signal Center, Texas A&M University ; ECEN720 : High-Speer Links Circuit ans Systems , Spring 2019.
- [28] “Equalization for High-Speed Serial Interfaces in Xilinx 7 Series FPGA Transceivers, white paper : 7 Series FPGAs ”; Harry Fu - XILINX ; March 27, 2012 ; https://www.xilinx.com/support/documentation/white_papers/wp419-7Series-XCVR-Equalization.pdf ; last access July 2020.
- [29] “The Benefits of Using Linear Equalization in Backplane and Cable Applications ”; Michael Peffers , Communications Interfaces – CIF ; Texas Instruments , Application Report SLLA338–June 2013.
- [30] “Precoding proposal for PAM4 modulation ”; Sudeep Bhoja, Will Bliss, Chung Chen, Vasu Parthasarathy, John Wang, Zhongfeng Wang – Broadcom ; at 100 Gb/s Backplane and Cable Task Force IEEE 802.3 Chicago September 2011 ; http://www.ieee802.org/3/bj/public/sep11/parthasarathy_01_0911.pdf ; last access July 2020.
- [31] “Understanding How the New Intel Hyperflex FPGA Architecture Enables Next-Generation High-Performance Systems”, Mike Hutton - Intel Programmable Solutions Group, White paper FPGA ; <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/wp/wp-01231-understanding-how-hyperflex-architecture-enables-high-performance-systems.pdf>; last access July 2020.
- [32] “Intel Custom foundry EMIB (Embedded Multi-die Interconnect Bridge)”, <https://www.intel.com/content/www/us/en/silicon-innovations/6-pillars/emib.html>; last access July 2020.
- [33] R. Mahajan et al., “Embedded Multi-die Interconnect Bridge (EMIB)—A High Density, High Bandwidth Packaging Interconnect,”, i2016 IEEE 66th Electronic Components and Technology Conference (ECTC), Las Vegas, NV, 2016, pp. 557-565, doi : 10.1109/ECTC.2016.201.
- [34] “Enablinnd Next-Generation Platfroms using Intel’s 3D System-in-Package Technology”, Manish Deo Senior Product Marketing Manager - Intel Programmable Solutions Group, White paper FPGA ; <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/wp/wp-01251-enabling-nextgen-with-3d-system-in-package.pdf>; last access July 2020.
- [35] “Intel Stratix 10 Clocking and PLL User Guide”, Updated for Intel Quartus Prime Design Suite : 19.3; UG-S10CLKPLL, 2020.03.06. ; <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/stratix-10/ug-s10-clkpll.pdf>; last access July 2020.
- [36] “Intel Stratix 10TX Signal Integrity Development Kit Board Schematic”, revision B1, document number : 150-0321330-B1 , Wednesday, March 11, 2020 ; https://www.intel.com/content/dam/altera-www/global/en_US/support/boards-kits/stratix10/si_tx/s10tx_si_b1.pdf; last access July 2020.
- [37] “Building an Intel Stratix 10 FPGA Transceiver PHY Layer”; Intel FPGA training course ; https://www.intel.com/content/www/us/en/programmable/customertraining/OLT/S10_XCVR_PHY/presentation_html5.html; last access July 2020.
- [38] “Intel Stratix Forward Error Correction”; Intel AN-846 (application note-846), 2018.07.02. .
- [39] Haiyun Yang, Tianshu Chi, “Parallel pseudo random bit sequence generation with adjustable width”, U.S. Patent 9,747,076 B1, issued August 29,2017.
- [40] “IEEE Standard for Ethernet SECTION SIX, Clause 78 through Clause 95 and Annex 83A through Annc 93C”, IEEE Std 802.3-TM2015 SECTION SIX , IEEE STANDARD FOR ETHERNET.
- [41] “IEEE Standard for Ethernet SECTION FOUR, Clause 44 through Clause 55 and Annex 44A through Annc 55B”, IEEE Std 802.3-2012 SECTION FOUR , IEEE STANDARD FOR ETHERNET.
- [42] K. Zhu and V.Saxena, “From Design to Test: A High-Speed PRBS,”, in IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 26, no. 10, pp. 2099-2107, Oct. 2018, doi : 10.1109/TVLSI.2018.2834373.
- [43] Hirohisa Yamaguchi, “Parallel pseudo random bit sequence generation with adjustable width”, Parallel M-sequence generator circuit” U.S. Patent 6,188,714 B1, issued Feb. 13,2001.



- [44] Anthony R. Bonaccio, Allen P. Haar, “Pseudo-random binary sequence checker with automatic synchronization”, U.S. Patent 7,219,113 B2, issued May 15,2007.
- [45] J. J. O’Reilly, “Series-parallel generation of m-sequences, ”, in Radio and Electronic Engineer, vol. 45, no. 4, pp. 171-176, April 1975, doi: 10.1040/ree.1975.0033.
- [46] Wei-Zen Chen a d Guan-Sheng Huang, “A parallel multi-pattern PRBS generator and BER tester for 40+Gbps Serdes Applications, ”, Proceedings of 2004 IEEE Asia-Pacific Conference on Advanced System Integrated Circuit, Fukuouka, Japan, 2004, pp. 318-321, doi: 10.1109/APA-SIC.2004.1349484.
- [47] “PRBS31Q Example Sequence,”, Tom Palkert, Atul Gupta ; comment #552, MACOM Partners from RF to Light..
- [48] Lyubomirsky, F. I., “PRQS Test Patterns for PAM4”, IEEE 802.3bs 400GbE Task Force, Sep. 14-16, 2015, 16 pages, Coconout Point, Florida, USA.
- [49] Peter Anslow, Ciena, “400GbE AMs and PAM4 test pattern characteristics”, IEEE 802.3bs Task Force, Logic Ad Hoc, 11 December 2015.
- [50] “Embedded Peripherals IP User Guide”, Updated for Intel Quartus Prime Design Suite: 19.4 , UG-01085 ; 2020.01.22 ;https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug_embedded_ip.pdf; last access July 2020.
- [51] “RS(544,514) FEC performance with 4:1 interleaving (updated 2)” ; Peter Anslow, Ciena ; IEEE P802.3ck Task Force, August 2018 ; http://grouper.ieee.org/groups/802/3/ck/public/18_09/anslow_3ck_01_0918.pdf; last access July 2020.
- [52] “Introduction to Reed-Solomon codes,”Bernard Sklar, Digital Communications: Fundamentals and Applications, Second Edition (Prentice-Hall, 2001, ISBN 0-13-084788-7) , Apr 12,2002 ;https://ptgmedia.pearsoncmg.com/images/art_sklar7_reed-solomon/elementLinks/art_sklar7_reed-solomon.pdf; last access July 2020.
- [53] “Forward Error Correction (FEC) techniques for optical communications”, Kamran Azadet, Mark Yu - Lucent Technologies , IEEE 802.3 High-Speed Study Group Plenary meeting, Montreal July 1999 ;<https://web.stanford.edu/class/ee392d/Chap8.pdf>; last access July 2020.
- [54] “chapter-8 Reed-Solomon”,EE392D-Channel Coding: Techniques, Analysis and Design Principles - Winter 2007 (shortened version of course 6.451 Principles of Digital Communication II at MIT), Professor G. David Forney, Standford University ;<https://web.stanford.edu/class/ee392d/Chap8.pdf>; last access July 2020.
- [55] “Intel FPGA E-Tile Clocking” ; Intel FPGA training course ;https://www.intel.com/content/www/us/en/programmable/customertraining/OLT/Etile_clocking/etile_clocking.mp4?; accessible to common public at ;<https://www.youtube.com/watch?v=Enk74lvAff0>; last access July 2020.
- [56] Smith, Wayne & Burns, Karen & Moorhead, Jane. (1993). Using Markov Chains to Model the Error Behavior of Data Communications Channels. 59.



A. PAM4 SIGNAL DEGRADATION AND EQUALIZATION TECHNIQUES

An analysis of distortion caused by insertion and return losses on high-speed serial transmissions is included in this appendix.⁸⁵


PAM4, PAM2(NRZ) modulations are compared in terms of SNR and insertion losses, providing thorough understanding of the advantages of PAM4 signalling over PAM2 in high-datarate links.

A qualitative insight on both TX and RX equalization techniques to reduce the signal degradation suffered by electrical signals is also given. Note that depicted equalization structures are limited to those used in Intel Stratix 10TX E-Tile xcvr.⁸⁶


A.1. PAM4 SIGNALLING

PAM4 is a 4-level pulse amplitude modulation for which each pulse at certain signal level encodes 2 bits (A.1 figure shows a comparison between PAM4 and PAM2 (NRZ) modulated signals, where f_b is baudrate).

PAM4 signals transmits 2 bits per symbol (2 bits per signal level), thus requiring half signal frequency (so half bandwidth) to achieve same datarate than NRZ modulated transmissions⁸⁷



(a)



(b)

Figure A.1.- PAM4 (a), PAM2 (b) baseband example modulated signaling.

PAM4 signalling performance is analyzed in next sections establishing a comparison with NRZ modulation for high-speed serial transmissions. PAM4, PAM2 comparison analysis is mainly done by observing the transformations in the ideal eye-diagram⁸⁸ when considering the signal degradation introduced by a non-ideal transmission channel, and the impact of eye distortions over signal integrity at the receiver.

PAM4, PAM2 eye-diagram shapes are both shown in A.4, revealing how PAM4 signal levels are much closer than in NRZ (separated only 1/3 of PAM2 'distance'). As separation between signal levels is significantly reduced, PAM4 is more vulnerable to channel noise (independently on signal frequency); meaning that lower noise can cause signal level to be increased or reduced such that it reaches an adjacent symbol-level, thus causing an error in transmitted symbol at the receiver.

Besides, PAM4 signal has more possible transitions as each symbol transmitted can take the value of any of the 4 modulation symbols, causing asymmetry in transition times.⁸⁹

That asymmetry makes PAM4 eyes narrower than expected⁹⁰ and introduce potential misalignment(skew) between PAM4 eyes (as shown in A.2,A.3)

⁸⁵ note that a brief comment on PAM4 characteristics is provided in A.1 including only aspects derived from use of 4 signal levels, that have impact on eye opening and thus on PAM4 resilience to ISI (A.1 includes only aspects required to ease the understanding of this appendix; in-depth details on PAM4 eye anatomy are provided in 2 section)

⁸⁶ note that naming used in A.3 sub-section of this appendix do not match the name given to configurable equalization parameters in 3; as the aim of this appendix is only to provide understanding of the meaning of Stratix10 equalization parameters and how they can help improving signal quality and eye opening at the receiver by reverting distortions suffered in transmissions.

⁸⁷ from here on, baudrate and datarate will be referred as f_d and f_b respectively; being baudrate the frequency at which modulation symbols(signal levels) are transmitted, meaning that for same baudrate, PAM4 will transmit twice the bits transmitted by an NRZ signal.

⁸⁸ refer to 2.3 for PAM4 eye-diagram anatomy details

⁸⁹ PAM4 transitions between adjacent symbols will require less transition time than transitions between symbols "00"-"11", for which signal level must change in 2x the level change for transitions between adjacent symbols

⁹⁰ for same datarate, ideal PAM4 eyes will have 2x width than NRZ as PAM4 signal will have 1/2 frequency than NRZ

A.2 marks points A,B over the slowest signal transition, corresponding to the point where transition crosses the horizontal thresholds to the region where signal samples are considered zero (A,B are placed over NRZ,PAM4 thresholds respectively). So that A.2 shows B occurring later, thus causing greater eye-closure (PAM4 inner-eyes suffer from extra-closure corresponding to the horizontal distance between A,B, on each side of the eye).

A.3 indicates the horizontal distance between the points A,B where transitions to zero enter in 'zero region'(occurring even under zero delay conditions), for transitions starting at different signal levels(different previous symbols), causing eye-missalignment.

Thus, when using same baudrate, potential miss-alignment(skew) in PAM4 eye-centers(ideal sampling point associated to each transition) hardens the requirements for flexibility in sampling point within each UI interval⁹¹.




Figure A.2.- PAM4-eye narrowing over PAM4 simulated 3-eyed diagram with low distortion (A,B indicate where slower PAM4 transition crosses NRZ 'zero-threshold', PAM4 'zero-threshold' respectively).






Figure A.3.- PAM4-eye narrowing over PAM4 simulated 3-eyed diagram with low distortion (A,B placed where transitions from different levels to zero 'zero-threshold'.)



(a)



(b)

Figure A.4.- PAM2(NRZ) zero distortion eye-diagram (a); PAM4 zero distortion 3-eye-diagram(b) -[24]-.

⁹¹ refer to (ref to main doc) for detail on eye-center definition/placement



A.2. HIGH FREQUENCY SIGNAL DISTORTION ON ELECTRICAL SERIAL LINKS (INSERTION AND RETURN LOSSES)

Insertion⁹² and return⁹³ losses increase with frequency in electrical lanes(channels) causing 2 major distortions on signal being transmitted :

- frequency dependent attenuation
reducing peak-to-peak signal voltage⁹⁴ that cannot be countered using equalization
- signal degradation due to ISI (inter-symbol interference)⁹⁵

PAM4 relieve both attenuation and ISI for a given datarate (i.e. when considering a PAM4 signal with half frequency than required NRZ to achieve desired datarate) but does not solve neither. Thus next sections analyze the behavior of PAM4 signals in terms of both attenuation and inter-symbol interference to expose to what extent distortion is avoided by using PAM4, and how equalization structures must be to counter remaining distortion.

A.2.1. ISI (INTER-SYMBOL INTERFERENCE)

Insertion loss characteristics for electrical lanes are frequency dependent, causing greater insertion loss for higher frequency signals⁹⁶ .

Insertion loss increase with frequency implies limited channel BW, thus filtering of high-frequency signal components and so smoothing signal transitions⁹⁷ (as shown in A.6c) .

So that, channel limited BW degrades signal quality.

ISI caused by insertion losses can be easily explained from that smoothness, when formulated in time-domain.

Insertion losses curves have typically the shape shown in figure A.5a, which corresponds to a non-ideal delta-like channel response $h(t)$, see figure A.5b .

If a pulsed signal (as shown in A.6a) is sent over such a channel, signal output at RX-side can be obtained as the convolution of the pulsed input signal and channel response $h(t)$, obtaining a distorted version of input signal (as in figure A.6c), where high-frequency components have been filtered, thus smoothing transitions.

So that, wave received at the end of a channel can be considered as if the signal takes time to reach its level.

If no other transition occurs while received signal is changing, it finally reached the transmitted level(symbol); but if another transition occurs before, the previous level will never be reached, causing symbol loss.

So that, if channel distortion is enough to make transitions last more than 1UI ($UI = 1/f_b$ where f_b is baudrate), succeed in transmitting 1 symbol depends on adjacent symbols⁹⁸ (so that insertion loss causes data-dependent losses, as dependency on adjacent symbols values is introduced; thus inter-symbol interference) .

A.2.1.1 ISI REDUCTION USING PAM4

As PAM4 requires half frequency than NRZ for same bitrate, symbol pulses (minimum duration for each signal level) for PAM4 signal is double than for NRZ (as each pulse sends 2 bits instead of 1).

⁹² Insertion loss refers to attenuation of signal strength over a line (note that insertion loss does not include losses due to the reflected signal from the input impedance)

⁹³ return loss refers to the part of the signal being sent over a line that does not go into the channel but returns back towards the transmitter due to impedance mismatches.

⁹⁴ for high-speed serial transmissions differential lanes are commonly used,thus attenuation is defined as compression peak-to-peak 'compression'

⁹⁵ when transmitting a sequence of symbols $s_n s_{(n+1)} \dots s_{l-1} s_l \dots s_m$ (each with 4 possible realizations when using PAM4) ISI refers to the influence of adjacent symbols s_j ($j \neq i$) on the signal value sampled at RX when receiving symbol s_i (in-depth explanation is given in sub-section A.2.1)

⁹⁶ insertion loss does not include losses due to the reflected signal from the input impedance (S11, S22)

⁹⁷ in pulsed(squared) signals, ideal transitions occur in zero time, corresponding to 'infinite-frequency' signal component. So that, when squared signal goes through BW-limited channel, signal high-frequency components will have been filtered at RX side, causing smoothness in signal transitions

⁹⁸ value of adjacent symbols will determine transitions involved,thus if the time required by involved transitions is short enough to reach signal level associated to the symbol that is being transmitted (note that not all transitions take same time)




Figure A.5.- Example typical insertion loss curves for non-ideal channel (a); corresponding delta-like impulse responses (b), [1].




Figure A.6.- Squared PAM2(NRZ) modulated signal shape (1 V peak-to-peak normalized) used as input signal (a); non-ideal delta-like channel impulse response (causing ISI) (b); PAM2 smoothed signal at channel output (ISI affected)(c) , [1].

Thus, when channel distortion is applied over PAM4 'squared' signal, the wider-in-time pulses of PAM4 can assume smoothness observed as longer duration increase probability of reaching symbol-level before next transition occurs, and so prevents signal levels (symbols) to disappear and be merged with next one, causing info losses .

So that, channel equalization become less challenging when using PAM4.

A.2.1.2 CHANNEL EQUALIZATION GOALS

As the typical channel impulse response shape shown in A.5b affects signal transitions introducing smoothness; equalization structures⁹⁹ will aim to boost the amplitude of the 1st data bit after the transition and reduce the influence of adjacent symbols.

So that, TX-equalization structures aim to boost higher-frequencies in order to flatten the channel transfer function¹⁰⁰ as shown in figure A.7¹⁰¹ .

For such a flattened response, signal values in transitions (highest-frequency components in signal) will suffer less attenuation, allowing for faster transitions and thus reducing smoothness. As signal is allowed for more rapid level changes, the probability of not 'having time enough to reach signal level associated with symbol being transmitted' is reduced, and so that the probability of succeeding in transmitting 1 symbol reduces the dependency on the adjacent symbol values (so that, ISI is reduced).

⁹⁹ explained later in this appendix, see A.3

¹⁰⁰ channel transfer function refers to channel characteristic curve in frequency-domain

¹⁰¹ equalization structure generating composed transfer function shown in A.7 is depicted in A.3.3

If considering channel response transformation (when adding equalization) in time-domain, flattening in channel transfer function corresponds to sharpen impulse response ($h(t)$) in time-domain, making it closer to delta-like response (so closer to ideal loss-free channel).




Figure A.7.- Typical channel transfer function (insertion loss curve) for limited BW channel (blue); TX-qualizer + channel composed transfer function (green); RX-equalization transfer function (black); TX-qualizer + channel + RX-qualizer flattened composed transfer function (red) , [1]

A.2.1.3 ISI FORMULATION

In digital transmissions, pulsed signals are sampled at RX side at discrete time instants, so analysis can be done by considering PAM signals to be transmitted as a sequence of discrete samples, $1/f_b$ apart (where f_b is baudrate), $s(n)$ being n the index of the sample corresponding to each symbol in relation with the symbol being transmitted (symbol being transmitted will hold 'index' $n=0$).

For a discrete-time analysis, channel response $h(t)$ is also considered as a set of coefficients $h(n)$, where $n=0, 1 \dots$ (figure A.8 shows how channel impulse response is 'sampled').




Figure A.8.- Example normalized impulse response $h(t)$, with channel delay $d=2$, sampled at $n \cdot t_b$ where $t_b = 1/f_b$ and f_b =baudrate.

In discrete-time time domain, convolution between transmitted signal and channel response ($h(t)*s(n)$), implies that each received symbol $r(n)$ (where 'n' states for the sampling instant and is related to the transmitted symbol $s(n)$) is given by :

$$r(n) = \sum_{i=0}^{inf} h(i) \cdot s(n - i + d) = h(0)s(n + d) + \dots + h(d)s(n) + h(d + 1)s(n - 1) + \dots \quad (5)$$

,where d is the 'position' of the highest-value-sample of $h(t)$ (the channel delay), so that symbol being transmitted $s(n)$ is multiplied by greatest coefficient in $h(n)$ (when $i = d$, $h(i)s(n-i+d) = h(d)s(n)$), see A.8.

As channel discrete impulse response has coefficients $h_i \neq 0$ ($i \neq d$, being d the channel delay), sampled signal value for each received symbol $r(n)$ will have influence of other adjacent symbols.

A.2.1.4 TRANSITION DENSITY

As ISI is data dependent (higher ISI presence for higher frequency of level-changes in data signals), PAM4 transition density should also be considered (in terms of impact on ISI).



Assuming all modulation symbols are equally probable, transition density is an indicator of rate of transitions in signal that imply a change in signal level (i.e. how often a data signal will change voltage level).

TD(transition density) is calculated as :

$$\text{average transition density} = (\text{distinct transitions} = \text{changing level}) / \text{total transitions} = (\text{distinct transitions}) / (\text{distinct transitions} + \text{no.-changing-level-transitions} = \text{num of symbols}) = 12 / (12 + 4) = 12/16 = 0.75$$

PAM2 (NRZ) has lower transition density (0.5), making it more resilient to ISI when channel equalization is not achieved for PAM4 signalling.

A.2.2. PAM4,PAM2 ATTENUATION COMPARISON (INSERTION LOSSES)

While a PAM4 experiences less ISI than a PAM2 signal at a given datarate (as PAM4 signal is half frequency, thus suffering from less insertion losses), it experiences more ISI at a given baudrate. As for a given baudrate both PAM4 and PAM2 signals are same frequency (thus PAM4 is 2x datarate) both suffer same insertion losses; so that, as PAM4 symbols are separated by smaller voltage margin (1/3 of the PAM2 margin), PAM4 is more vulnerable to ISI.¹⁰²

So that, when it comes to attenuation, PAM4 eyes suffer more signal level distortions as its SNR (signal to noise ratio) is lower in comparison to PAM2 SNR.¹⁰³

For same baudrate, noise affecting both PAM4 and PAM2 signals will be same, meaning that PAM4 SNR and PAM2 SNR will only differ in terms of signal level. As PAM4 symbols have 1/3 the voltage margin of PAM2, PAM4 'extra' losses related to lower separation between signal levels can be calculated as :

$$20 \log_{10}(1/3) = -9,542 \text{ dB} \quad (6)$$

, meaning that at same baudrate PAM4 will suffer -9.542 dB of extra attenuation over PAM2.

So that the 'cost' of doubling datarate by moving to PAM4, can be measured as ≈ 9.5 dB of extra attenuation in signal.

Meanwhile, doubling datarate by doubling PAM2 signal frequency will cause an attenuation over PAM2 signal of ≈ 11 dB for a CEI-OIF compliant channel.¹⁰⁴

So that, the 'cost' of doubling datarate (measured in signal attenuation or penalty), will be less if that increase is done by moving to PAM4 modulation rather than increasing PAM2 signal frequency, making PAM4 more attractive for high-datarate transmissions, as lower attenuation implies lower ISI impact.

A.1 sums up PAM2,PAM4 comparison in terms of ISI vulnerability.

	NRZ	PAM4
bits per symbol	1	2
eye diagrams per UI	1	3
SNR loss (at same baudrate)	0 dB	+9.5 dB
different transitions	2	12
rising falling edges	2	6
average transition density	50 %	75 %
skew and compression	absent	present

Table A.1.- PAM2,PAM4 relevant attributes for ISI vulnerability analysis , [24]

¹⁰² when the separation between signal levels is lower, less channel noise is enough to distort the signal such that signal level reaches voltage associated with an adjacent symbol, thus causing errors in transmission.

¹⁰³ SNR is here used to establish a quantitative comparison between the resilience of PAM2,PAM4 to channel noise.

¹⁰⁴ CEI-OIF specification provides channel-dependent max insertion loss 'masks'. If considering CEI-OIF masks, attenuation introduced when doubling frequency from 14-28 GHz will be a minimum 11 dB for lower attenuation case,[8]

A.3. PAM4 EQUALIZATION

Though PAM4 enables a cut down in signal frequency to half the NRZ bandwidth -thus reducing insertion loss-, its reduced SNR makes it more vulnerable to amplitude degradations. So that, PAM4 does not remove the need for channel equalization to minimize ISI.

As stated in A.2.1.2 sub-section, channel equalization goal is to flatten channel frequency response, thus sharpening impulse response for a narrower delta-like $h(t)$.

So that, removing/minimizing $h(-1)$ and $h(1)$ coefficients of $h(n)$ ¹⁰⁵ (named 1st pre- and post- cursors) corresponding to symbols adjacent to the one being transmitted (if assuming zero delay in $h(n)$, $d = 0$), thus removing/minimizing contribution of adjacent symbols to the signal value received and reducing errors in transmissions due to signal distortion.¹⁰⁶

Though many strategies exists, equalization structures described in this appendix are limited to the equalizer structures used in Stratix 10 E-tile xcvs [11]. A.9 figure shows an equalization structure including de-emphasis FIR-filter TX-equalizer and CTLE (continuous time linear equalization)+DFE(decision feedback equalization).

A.7 anticipated how the composite response of a link between RX and TX is transformed as adding equalization structures to compose chain in A.9¹⁰⁷

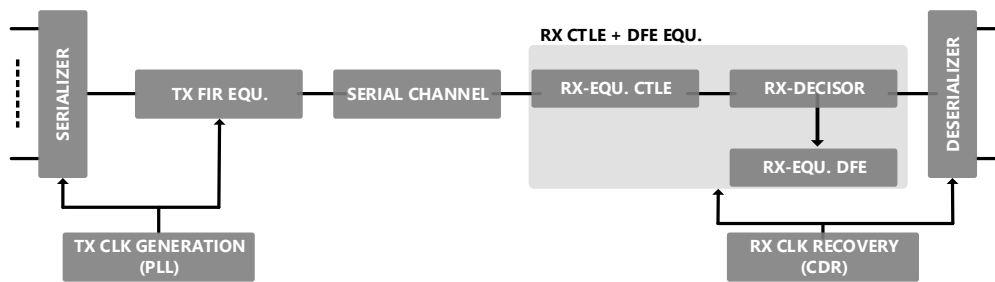


Figure A.9.- Generic xcvr signal path including TX FIR filter-based equalizer, CTLE+DFE RX equalizer (note that in real xcvr structure signal modulator and demodulator would be placed right at the input of TX-equalizer and output of RX-CTLE, respectively).

TX-side equalization is implemented in Stratix 10 e-tile xcvs using a FIR-filter based equalizer; while RX-side equalization relies on a CTLE + DFE structure, [11].

DFE equalizer is composed by an slicer(a decision block on charge of deciding which modulation symbol corresponds to the voltage level sampled¹⁰⁸), and a feedback FIR-filter.

FIR-filters in both TX and RX sides perform discrete-time linear equalization (DTLE); while CTLE performs 'analogue' signal equalization (structure and purpose of each equalization element is explained in subsequent sections).

A.3.1. DISCRETE TIME LINEAR EQUALIZATION USING FIR FILTERS

DTLE uses finite impulse response (FIR) filters (figure A.10 shows the generic structure of a k -coefficients FIR filter, where $k = m+2$).

¹⁰⁵ $h(n)$ refers to sampled version of the channel impulse response

¹⁰⁶ note that ultimate goal of equalization aims to keep the system working within a target BER

¹⁰⁷ note that DFE effect is not included in A.7 as it cannot be directly represented due to its non-linearity

¹⁰⁸ sampled signal values can be affected by distortion, thus having a voltage level different than the 'exact-ideal level' associated to each modulation symbol; so that a decisor is placed before RX in chain (as shown in A.9 figure), to the sampled signal value deducting which modulation symbol was transmitted, and output that 'recovered' symbol to the receiver

FIR filters are built from a shift register(chain of delay elements, each one introducing a 1UI delay¹⁰⁹ , so that a sequence of sent symbols $s_1 s_0 s_{-1} \dots s_{-m}$ (where s_i is sample $s(i)$ of signal passing through FIR filter) are hold at output of registers in chain.

$s_1 s_0 s_{-1} \dots s_{-m}$ are weighted using $s_{-1} s_0 s_1 \dots s_m$ coefficients(referred to as filter taps) and summed up (in linear combination), so that filter output at instant i (corresponding to instant at which input symbol $s(i)$ is processed) is $y(i)$, where $y(i)$ has influence of previous $s_{-1} \dots s_{-m}$ symbols , but can also consider value of later symbols, $s_1 = s(1)$.

- negative tap-coefficient subscripts (c_i where $i < 0$), weight symbols sent after current symbol, $s(0)$, and are named precursor coefficients [10].
- positive tap-coefficient subscripts (c_i where $i > 0$), weight symbols sent before current symbol, $s(0)$, and are named postcursor coefficients [10].




Figure A.10.- Common structure for k -coefficients ($c_{-1}, c_0, c_1, \dots, c_{k-2}$) FIR filter with $k = m+2$ (note that Δ blocks are 1UI delay).

A.3.2. TX EQUALIZER (DE-EMPHASIS EQUALIZATION)

DTLE 3-tap FIR-filter based equalizer is used for equalization on TX side in Stratix 10 e-tile xcvr [11] (figure A.11 shows 3-tap FIR filter structure).

DTLE filter used is referred to as de-emphasis equalizer, meaning that frequency response flattening is done by attenuating low-frequency signal components: de-emphasis filter applies reduction coefficients(attenuation) to adjacent symbols to reduce their influence on symbol being transmitted, thus de-emphasis filter acts over symbols instead of acting over transitions(higher signal frequency components).

So that TX FIR-equalizer performance affect 'low' frequency components (component at $f_b = \text{baudrate}$).




Figure A.11.- 3-tap FIR filter structure for equalization on TX-side in Stratix 10 TX signal integrity development kit (1 FIR filter based TX-equalizer is included for each channel on each e-tile xcvr).

TX de-emphasis filter aims to compensate channel distortion before it happens by inserting in channel a pre-distorted signal version; manipulated such that after channel distortion, received signal will ideally match the original shape of data signal (before passing through de-emphasis FIR equalizer, see set of figures A.12).

¹⁰⁹ 1UI = $1/f_b$, being f_b the baudrate; so that 1UI delay equals 1 symbol delay

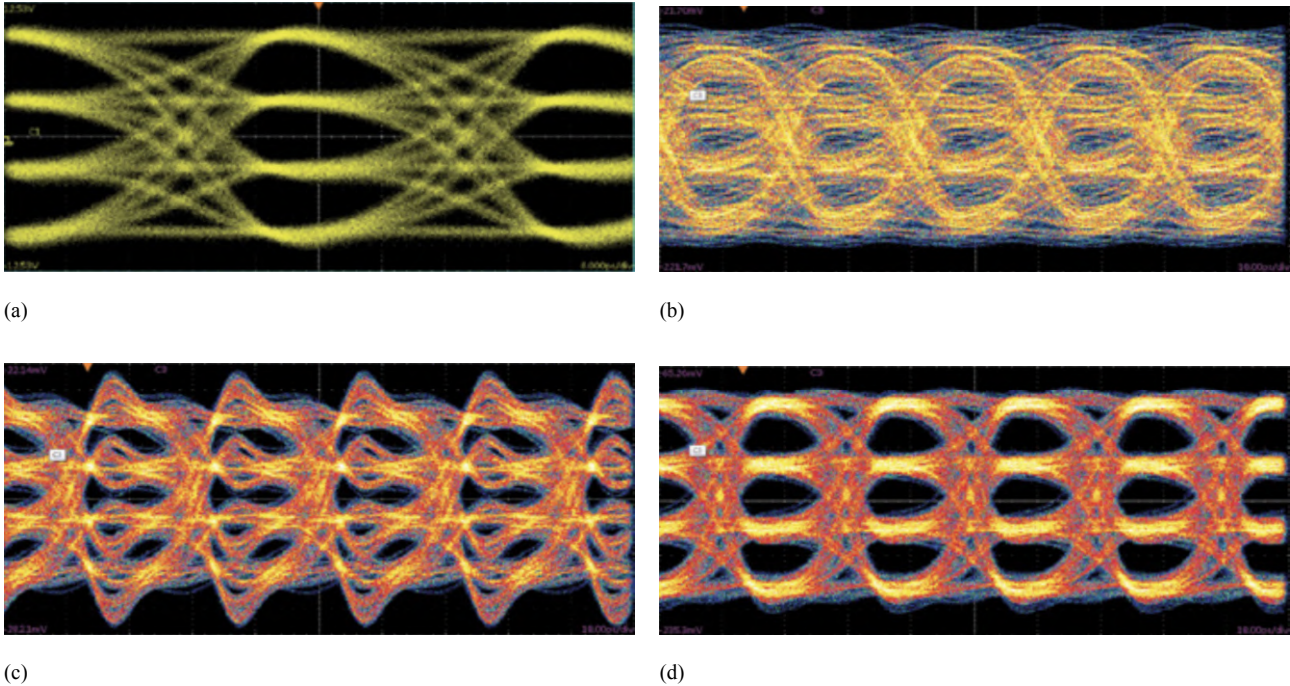


Figure A.12.- PAM4 eye-diagram measured at TX-output (right before insertion in channel), when no pre-distortion is applied (de-emphasis not used, thus TX signal not pre-distorted) (a); PAM4 eye-diagram measured at RX-input (right after channel), when no pre-distortion is applied (de-emphasis not used, thus RX distortion not compensated) (b); PAM4 eye-diagram measured at TX-output when pre-distortion is applied (de-emphasis pre-distorts signal to invert channel distortion) (c); PAM4 eye-diagram measured at RX-input, when pre-distortion is applied (de-emphasis used, thus channel inverts pre-distortion and RX signal eyes re-open) (d) , [10].

If considering a channel impulse response $h(t)$ with shape shown in figure A.8 for which discrete-time version $h(n)$ ¹¹⁰ cursors are marked, received signal after channel is obtained as :

$$r(n) = s(n) * h(n) = \dots + s(1)h(-1) + s(0)h(0) + s(-1)h(1) + \dots + s(-m)h(m)$$

, where $s(n)$ is signal inserted in the channel (the signal to be transmitted)¹¹¹ , ¹¹² .

As there are coefficients $h(i) \neq 0$ such that $h(i) > 0$, the signal level sampled when receiving symbol $r(0)$ will have influence of adjacent symbols (mainly $s(-1)$ and $s(1)$).

If only the influence of $s(-1)$ and $s(1)$ is considered, $r(0)$ can be expressed as $r(0) = s(1)h(-1) + s(0)h(0) + s(-1)h(1)$.

ISI introduced by terms $s(1)h(-1)$, $s(-1)h(1)$ can be reduced by using a 3-tap de-emphasis filter with shape shown in A.11 and coefficients c_{-1}, c_0, c_1 .

Introducing the 3-tapped de-emphasis FIR filter , its coefficients c_{-1}, c_1 will ideally cancel $h(-1)$, $h(1)$ so that the composite response of the set FIR-filter + channel , $h'(t)$ (see figures A.14, A.13) , is such that the received symbol $r(0)$ will be :

$$r(0) = s(1)h'(-1) + s(0)h'(0) + s(-1)h'(1) = s(1)(h(-1) + c_1) + s(0)h'(0) + s(-1)(h(1) + c_1) = s(0)h'(0)$$
¹¹³

, thus removing main ISI component due to adjacent symbols $s(-1)$ and $s(1)$ ¹¹⁴



¹¹⁰ $h(n)$ is set of coefficients (cursors) that is obtained by sampling $h(t)$ at $1/f_b$ intervals ($h(n) = \{ \dots h(-1) h(0) h(1) \dots h(m) \} = \{ h(t=0) h(t=1/f_b) h(t=2/f_b) \dots \}$).

¹¹¹ note that impulse response with zero delay is used for simplicity, thus $h(0)$ is the greatest coefficient of the impulse response, applied to the symbol being transmitted (so that, $h(0)$ is denoted main cursor)

¹¹² signal samples $s(i)$ where ($i < 0$) correspond with symbols that will be sent after symbol being transmitted

¹¹³ as FIR filter taps c_{-1}, c_0, c_1 are adjusted so that c_{-1} and c_1 will cancel $h(1)$ and $h(-1)$, respectively

¹¹⁴ $h'(-1) h'(1)$ will be ≈ 0 and so $s'_0 \approx s_0$ (reduced ISI)



(d)

Figure A.13.- Example channel impulse response $h(t)$ (causing ISI) (a); example DFE FIR filter impulse response (b); DFE + channel composed impulse response $h'(t)$ ideally cancelling 1st pre- and post- cursors (c); overlapped channel , 1-tap pre-emphasis FIR filter equalizer, composed response example (included for further clarity) [25] (d);(note that axis scale and values shown in figures (a),(b),(c) must not be taken into account, this images just aim to show impulse response transformation when DFE FIR filter-based TX-equalizer is used).




Figure A.14.- Signal path from TX-side equalizer input ($s(n)$) to RX input ($r(n)$).

The pre-distortion introduced by the 3-tapped FIR filter means that the signal inserted ($\hat{s}(n)$ in figure A.14) in channel may differ in shape from original signal $s(n)$, so that the shape of the received signal after channel $r(n)$ will match $s(n)$

So that, the purpose of TX pre-emphasis equalizer is to apply delay and inversion to the signal and add to original signal using the proper weight , thereby compensating ISI from nearby symbols [2].

3-tapped FIR-filter TX-equalizer for is able to output 64 different discrete signal levels (thus TX output signal after equalizer $\hat{s}(n)$, is not limited to the 4 signal levels associated to PAM4 symbols), when using PAM4. FIR filter coefficients in TX-equalizer must belong to the same sub-space as modulation symbols (sub-space of naturals represented using 2 bits , $\{0, 1, 2, 3\}$; where +, - operations are defined in modulus 2^2), and each of the 3 terms considered $s(1)h'(-1)$, $s(0)h'(0)$, $s(-1)h'(1)$ will also belong to same sub-space.

So that, each of the 3 terms can hold 4 different values, resulting in up to $4^3 = 64$ possible values that can be distinguished at TX output to send the associated signal level. 64 possible analog-signal-levels at output of transmitter gives TX FIR-filter based equalizer great flexibility to adjust output signal level to the required level.




Figure A.15.- Example channel impulse response without equalization(red), and composite response when 3-tap TX FIR equalizer used(blue) (a); example PAM2 signal at channel input when TX-FIR equalizer not used(red), and if used (blue) (b); channel, 3-tap TX FIR equalizer, composed(channel + 3-tap TX FIR equalizer) transfer function , [27].

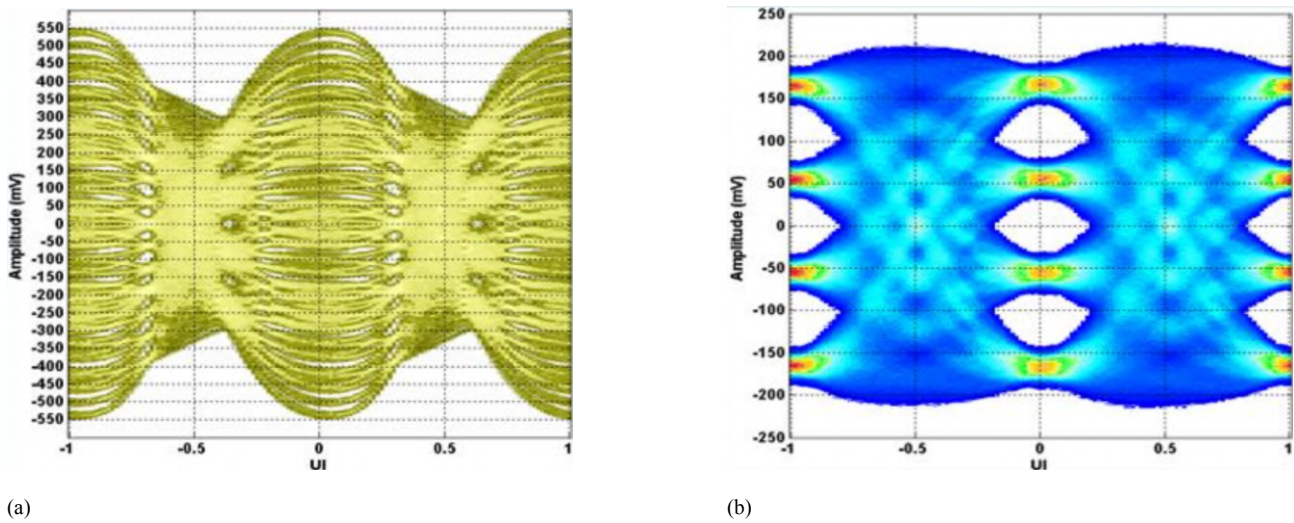


Figure A.16.- Example 64-levels PAM4 eye-diagram measured at channel input when 3-tap TX FIR filter based equalizer is used (after pre-distortion) (a); associated PAM4 eye-diagram measured at channel output (pre-distortion compensated) (b), [1].

If channel response weights symbols s_{-1} , s_1 using ‘coefficients’ $h(1),h(-1)$ respectively at instant of transmitting s_0 , and add it up; received symbol should be $h_{-1}s_1 + s_0 + h_1s_{-1}$, (assuming $h_0 = 1$ for simplicity).

TX can transmit s_0 as $(s_0 - c_{-1}s_1 - c_1s_{-1})$; so that after channel, receiver ideally gets

$$(s_0 - c_{-1}s_1 - c_1s_{-1}) + h_{-1}s_1 + h_1s_{-1} = s_0$$

, if 3-tapped filter coefficients c_i ideally adjusted.

Those 64 levels give reduced step to adjust the output signal level so that $(s_0 - c_{-1}s_1 - c_1s_{-1})$ can be closer to $(s_0 - h_{-1}s_1 - h_1s_{-1})$, thus receiving symbol can be $\approx s_0$.

A.16 figure shows eye-diagram at TX-equalizer output when 64 signal levels can be used to pre-distort sent signal to increase eye-opening after the channel (thus to reduce ISI) .

A.3.3. RX EQUALIZER (CTLE + DFE EQUALIZATION STRUCTURE)

A.3.3.1 CTLE (CONTINUOUS TIME LINEAR EQUALIZATION)

CTLE equalizer has a frequency characteristic curve shape similar to curves shown in figure A.17 where higher frequency components are boosted over lower frequency components (either by boosting high frequencies or attenuating lower ones) .

So that, CTLE reduced the smoothness introduced by BW limited channel in signal transitions, thus 'restoring' transition times (so reducing ISI) and re-opening received eye (A.18b.1-A.18b.3 figures reflects eye-diagram transformation over link including eye re-opening after transitions smoothness is removed by CTLE¹¹⁵)

CTLE equalizer is placed immediately at channel output (as shown in A.19) to perform 'analogue' equalization over received signal before the decisor on charge of determine the transmitted symbol, considering the received signal level.




Figure 83E-10—Selectable continuous time linear equalizer (CTLE) characteristic

Figure A.17.- CTLE gain (transfer function) shape for RX-equalization (from CTLE characteristic definition in [12].)

A.3.3.2 DFE (DECISION FEEDBACK EQUALIZATION)

DFE is a RX-side equalizer structure that, assuming previous symbols were received without error is able to remove their influence on the symbol being received by subtracting a linear combination of m previous symbols, thus removing ISI corresponding caused by contributions due to m previous symbols $\{s_{-1}, s_{-2} \dots s_{-m}\}$ (figure A.20 shows the structure of DFE equalizer).

DFE is nonlinear¹¹⁶ equalizer composed by a decisor(a.k.a slicer) that determines symbol sent from a received signal sampled¹¹⁷ , $\hat{s}(n)$ in figure A.20 and a FIR filter that weights decided/'re-stored' symbols, thus DFE FIR filter input is fed with ideal symbol-voltage-levels $\hat{s}(n)$, to subtract the distortion caused by previous pulses on the current pulse:

- when $r'(n)$ has remaining ISI from m previous symbols $r'(n) \neq r(n)$ (where $r'(n)$ is received signal $r(n)$ after CTLE equalizer , see figure A.20) ,
- If DFE succeeds in eliminating ISI $\hat{s}(n) = r(n)$ (see figure A.20)

So that DFE FIR filter behaves in similar way as TX de-emphasis FIR filter to remove remaining ISI, but operating limited to previous symbols s_{-1} , $s_{-2} \dots s_{-m}$.

DFE operates on the assumption of a low BER¹¹⁸ .

Symbol decisions are fed back through FIR filte, added to the output of the CTLE (see figure A.19), and the resulting signal is sent to the decider (a.k.a. slicer). If DFE feeds back an error, signal is deteriorated at the slicer which can cause another error. So that an error in a symbol decision can 'propagate' causing burst errors when DFE fails.

PAM4 reduced SNR(demonstrated in sub-section A.2.2) makes it prone to errors, increasing the probability for DFE bursts to occur.

¹¹⁵ note that transmitted eye shown corresponds to the eye after a TX-equalizer, where marked lobes represent 'transition-bits' being boosted by de-emphasis filter

¹¹⁶ DFE non-linearity is introduced by decisor/slicer, as decision criteria is usually a non-linear operation

¹¹⁷ note that slicer 'decides' modulation symbol associated to signal level sampled(that will match the symbol sent when no error occur in transmission), and outputs a 'recovered' symbol to the RX, meaning that signal level sampled is 'restored' to the voltage associated with the decided symbol

¹¹⁸ assuming decided symbols $\hat{s}(n)$ are correct, thus 'distortion' subtracted to current is removes ISI (instead of introducing more interference because of subtracting a 'wrong' value)







Figure A.18.- Channel transfer function (frequency response) transformation when CTLE is used for equalization at RX (a) [28]; example NRZ eye-diagram evolution from channel input to RX-input (after CTLE) (transmitter eye(b.1), transmitted eye after channel (b.2), transmitted after after channel and CTLE (b.3)), [29].

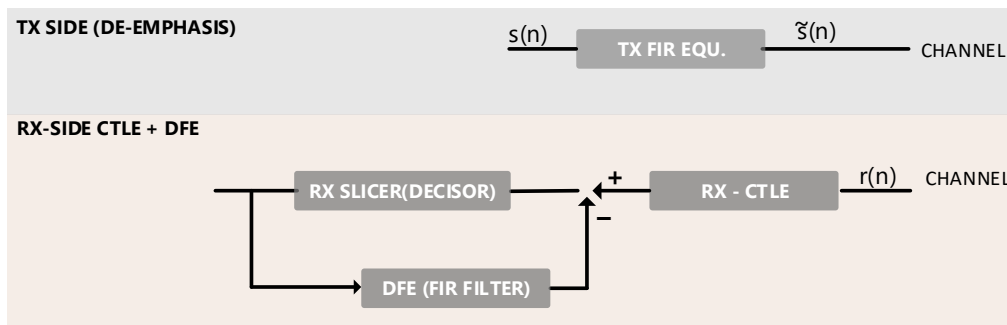


Figure A.19.- TX/RX equalization structure with FIR filter de-emphasis TX-equalizer; and CTLE + FIR filter DFE RX-equalizer

$1/(1+D)$ encoding can be used to address this problem introduced when using PAM4 ($1/(1+D)$ operation is explained in A.3.3.3).¹¹⁹

CTLE + DFE equalization scheme runs into problems when combined with error correction. DFE can introduce error bursts that may exceed the max correctable-burst-length of FEC algorithm being used.

Interleaving can stir error bursts into multiple FEC blocks, spreading an error bursts of certain length in several shorter bursts affecting subsequent blocks; so that, increasing possibility for bursts those reduced bursts to be short enough to become correctable (data interleaving is explained in D.1.1.2 in appendix D).

A.3.3.3 $1/(1+D)$ SIGNAL PRECODING

$1/(1+D)$ precoding reduces error bursts generated by 1 tap DFE into 2 errors per error event (thus, each symbol decision error turns out into 2 errors after $1/(1+D)$ decoding, with no dependence on actual length of burst caused by decision error)¹²⁰.

$1/(1+D)$ operates by pre-coding signal at transmitter side using a $1/(1+D) \bmod 4$ coder; while RX implements a $(1+D) \bmod 4$ decoder after the slicer.

¹¹⁹ note that many xcvr implementations give up on DFE in RX-equalization structure.

¹²⁰ note that $1/(1+D)$ precoding reduces bursts errors generated by decision errors in DFE, but does not reduce error bursts caused by channel distortion

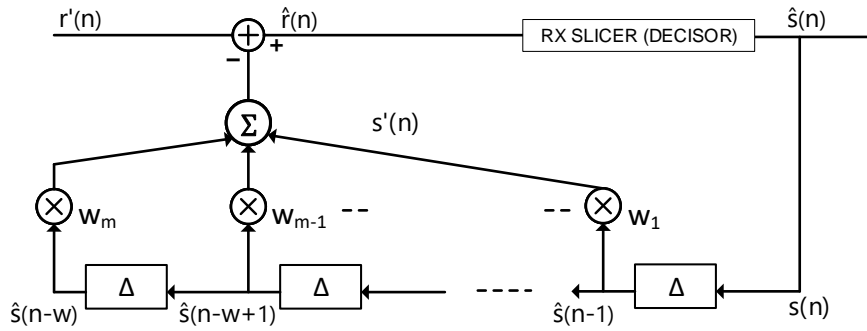



Figure A.20.- DFE equalizer internal structure using slicer + feedback FIR filter architecture (note that DFE's input $r'(n)$ is the received signal after CTLE equalization).

(figure A.21a shows TX/RX structures when $1/(1+D)$ precoding is used. A.21b shows an example that demonstrates DFE operation over TX/RX structures in A.21a).



(a)

Precoder Input : $tx(n)$	-	2	2	2	0	3	2	0	1	3	3	0	0	0	0	2	3	0	3	
Precoder Output : $p(n)$	-	0	2	0	2	2	1	1	3	2	1	2	2	2	2	0	3	1	2	
DFE, Slicer Output : $d(n)$	-	0	1	1	1	3	0	2	2	3	0	3	1	3	1	3	0	3	1	2
Error Event : $p(n) - d(n)$	-	0	1	-1	1	-1	1	-1	1	-1	1	-1	1	-1	1	-1	0	0	0	0
Decoder Output after $1+D$ decoding at rx : $rn(n)$	-	2	1	2	2	0	3	2	0	1	3	3	0	0	0	0	3	3	0	3
			↙																	↘
			Entry Error																	Exit Error

(b)

Figure A.21.- Simplified serial signal path schema when $1/(1+D)$ pre-coding is used (Δ is a unit delay used to perform $1/(1+D)$ encoding)(a); example demonstrating $1/(1+D)$ working principle. [30].

B. INTEL STRATIX 10 TX ARCHITECTURE

Intel Stratix 10TX signal integrity development kit board's structure is partially covered in this appendix.

This appendix includes thorough explanation on the underlying structure supporting the PAM4 serial link testing system (note that the contents of this appendix are limited to the structure segments involved in the PAM4 serial link testing system developed, further detail on complete Stratix 10TX signal integrity development kit are given in [7],[17],[11]).

B.1. INTEL STRATIX 10TX SIGNAL INTEGRITY DEVELOPMENT KIT OVERVIEW

Intel Stratix 10 TX FPGAs feature power-efficient, dual-mode(PAM4/PAM2) transceivers, capable of both 57.8 Gbps PAM4 (Pulse Amplitude Modulation) and 28.9 Gbps NRZ (Non Return to Zero) operation (supporting baudrates up to 28GBs for chip-to-chip, chip-to-module and backplane transmissions [8]).

Intel Stratix 10 TX devices are implemented using Intel 14 nm tri-gate (FinFET) technology¹²¹ to support high-speed FPGA core performance. Intel's hyperflex architecture is used in Stratix 10 TX devices to provide greater design flexibility without performance degradation (refer to [31] more detail on hyperflex architecture for performance improvement)

Intel Stratix 10TX signal integrity development kit allows to:

- evaluate transceiver performance up to 58 Gbps per E- Tile (PAM4 capable xcvr).
- Generate and check pseudo-random binary sequence (PRBS) patterns
- Dynamically change differential output voltage (VOD) pre-emphasis and equalization settings to optimize transceiver performance for each channel

so that enabling for thorough PAM4 signal quality analysis over serial link transmissions.

B.1 shows a block diagram of Intel Stratix 10TX on-board FPGA connection structure, including its interfaces with external connector places over the board as shown in B.2.

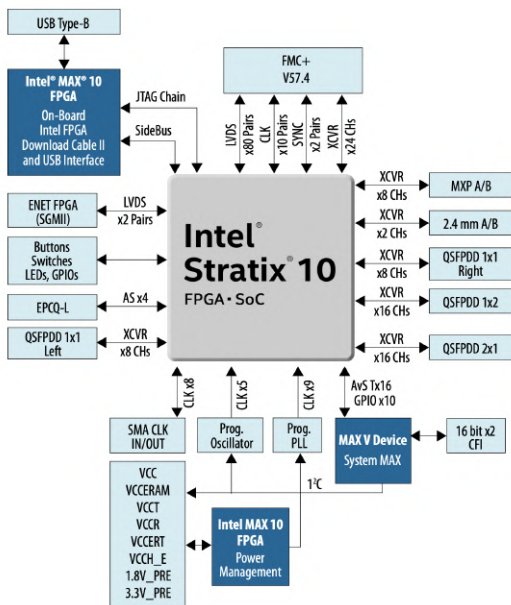


Figure B.1.- Intel Stratix 10 TX Development Kit Block Diagram (courtesy of Intel [17]).




Figure B.2.- Intel Stratix 10 TX Transceiver Signal Integrity Development Kit Picture(courtesy of Intel [17]).

¹²¹ 14 nm tri-gate (FinFET) technology refers to on-die transistor's technology

B.2. HYPERFLEX ARCHITECTURE (FPGA ARCHITECTURE)

Intel Stratix 10TX uses Intel’s hyperflex architecture , which boosts FPGA core performance but also enables for ‘in-chip’ integration with multiple devices on different dies, over a high-data-density, high-speed multi-die interface (EMIB).

So that, Intel Stratix 10TX implements advanced packaging technology based on Intel’s Embedded Multi-die Interconnect Bridge (EMIB), which support the connection of transceiver tiles (solid separately over different dies) to the FPGA fabric [16].

B.3 shows Intel Stratix 10TX Signal Integrity development kit configuration including 5 E-Tile PAM4 capable xcvs (each over separate die, integrated over EMIB interfaces. B.2.1 sub-section provides more detail on EMIB bridge.)

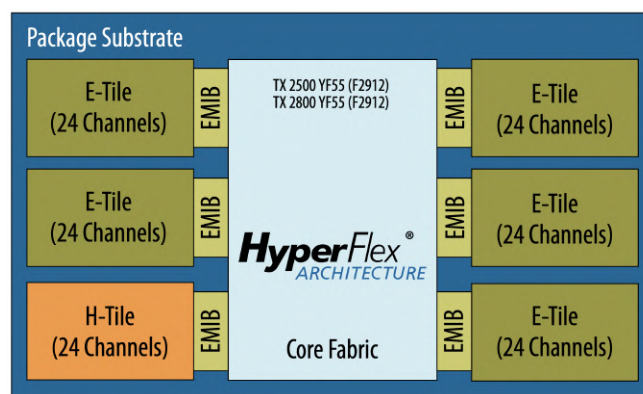


Figure B.3.- H-Tile and E-tile Layout Configuration for Intel Stratix 10TX Device with 5 E-Tiles and 1 H-Tile (144 Transceiver channels) (courtesy of Intel [11]).

B.2.1. FPGA-XCVR INTERFACE (EMIB, EMBEDDED MULTI-DIE INTERCONNECT BRIDGE)

B.4 figure shows FPGA core – E-Tile xcvr connection over EMIB interface.

Intel’s EMIB (Embedded Multi-Die Interconnect Bridge) technology supports ‘in-chip’ integration of multiple devices over different dies, providing a high-bandwidth interface to handle data and control signals for FPGA – e-tile xcvr communication.

So that, e-tile xcvs are implemented on separate dies and built up only of xcvs and resources to connect to and control the xcvs; thus isolating xcvr resources from FPGA SoC without degrading performance¹²² (more details on EMIB technology are provided in [32],[33]).




Figure B.4.- Example of die-to-die Heterogeneous Integration using EMIB Technology for Stratix 10 FPGA(courtesy of Intel [34]).

¹²² EMIB’s high-bandwidth avoids performance degradation due to inter-die communication



B.2.2. XCVR TILE-STRUCTURE (E-TILES)

B.2.2.1 E-TILE STRUCTURE (GXE TRANSCEIVERS)

Inside each E-tile xcvr die, there are 24 independent xcvr channels (as shown in figure B.5), referred to as GXE channels. GXE channels can run both PAM4 and NRZ transmissions up to 28.9GBs (thus up to 57.8Gbps when using PAM4)

B.6 shows a simplified diagram of a GXE channel, where EMIB handles the connection to FPGA core for both data and xcvr-control/status signals (PCS and PMA blocks shown in B.6 are detailed in B.3.1 sub-section).

EMIB resources are allocated to manage data/control signals for each GXE channel (note that all 24 channels share a single EMIB interface, associated with the e-tile xcvr die) and so occurs on FPGA (note that GXE channel structure shown in B.6 includes resources on both FPGA die (left-side of EMIB) and e-tile die (right-side of EMIB)

24 GXE channels within the same E-Tile share same EMIB interface.

In PAM4 transmissions over 30Gbps, the GXE channel performing the transmission needs to double its data-density through the EMIB interface (data density is doubled so that transmitted/received data can be exchanged over a parallel interface at a frequency supported by FPGA core).

As EMIB interface is shared by all 24 GXE channels, EMIB resources can be re-allocated to accommodate doubled data density for high-datarate PAM4 channels.

2 adjacent EMIB interfaces are used to double data-density for each PAM4 GXE channel over 30Gbps (so that, the adjacent GXE channel is disconnected as its EMIB interface is re-assigned to the PAM4 high-speed channel; thus only up to 12 GXE channels - out of 24– are available when PAM4 high-datarate mode is used within the E-tile xcvr. B.2.3 provides more detail on double data-density when PAM4 high-datarate mode is used).

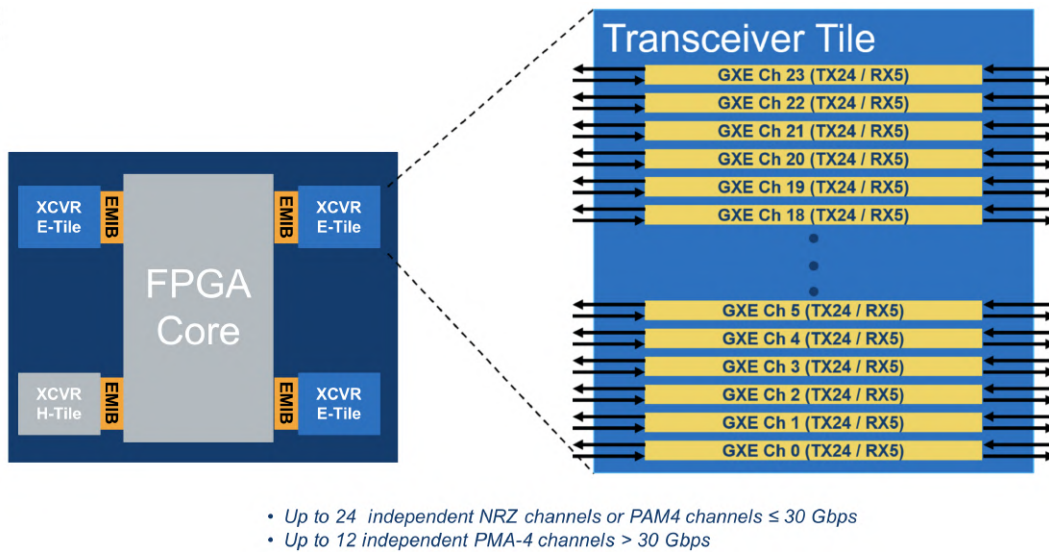


Figure B.5.- E-tile xcvr internal GXE channels(tx/RX) set (courtesy of Intel [16]).

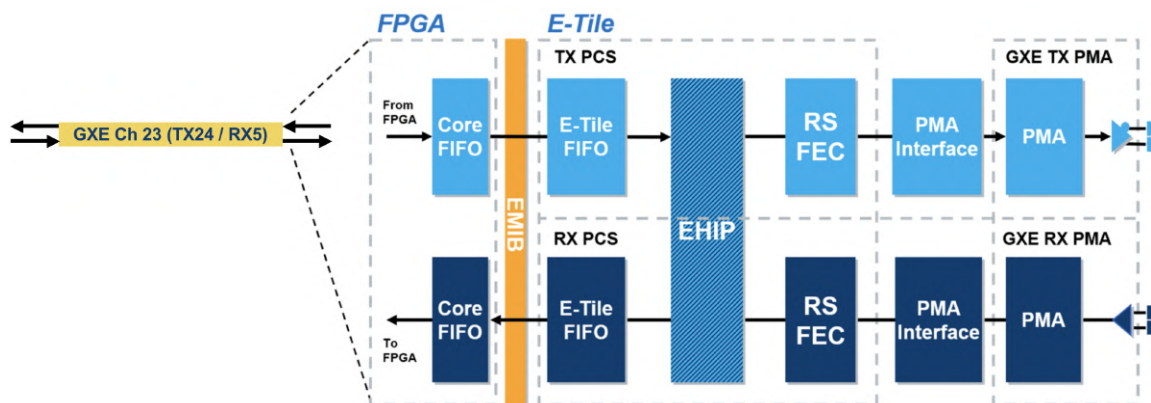


Figure B.6.- E-Tile GXE channel xcvr simplified internal structure block diagram(courtesy of Intel [16]).



GXE channels include separate TX,RX equalization structures to improve signal integrity(compensate link losses) in high-speed transmissions (note that equalization structure is included in PMA block B.6 , so equalization parameters can be configured separately for each GXE channel. B.3.1.1 sub-section provides further detail on PMA internal structure).

24 GXE channels on each e-tile are lay on the same die, and so share the same reference clock input; but baudrate is configurable on per-channel basis

B.2.2.2 E-TILE RESOURCES

B.2.2.2.1 S10 E-TILE BANKS

S10 E-tiles' dies connected to S10 FPGA core are referred to as E-tile banks (as each one of the 5 E-tile xcvr's B.7 connected to Stratix 10TX FPGA is implemented over a separate die, its resources are limited to the associated die. E-tile's die is referred to as E-tile's bank when including the input resources that are available within the die (clocking resources , connectors).)

B.7 figure show a block diagram for Stratix 10TX transceivers placement, where each block of 24 GXE channels belong to same E-tile xcvr, thus to same E-tile bank.

B.7 associates each E-tile bank (thus each E-tile xcvr) to the corresponding connectors over S10 board (note that each E-tile xcvr is physically connected to the indicated connector available on S10 board; so that each FPGA design will aim a certain E-tile xcvr depending on the connector to be used). 9A and 9C E-tile xcvr's (E-tile banks) are used in PAM4 serial link testing system to provide signal input/output over QSFP-DD¹²³ 1x2, and SMA 2.4mm RF connectors .

GXE channels within each E-tile can be mapped to any pin belonging to one of the connectors associated with its E-tile.

B.2 shows connectors' location over S10 board (note that naming in B.7 matches names used in figure B.2).

S10 E-tile bank used also limits the available clock to the source clocks available within the associated die , B.2.2.2.2 sub-section provides detail on tile clocking.

B.2.2.2.2 S10 E-TILE CLOCKING

S10 clocking structure provides dedicated reference clock sources to each E-tile xcvr, from an I2C programmable oscillator or PLL (note that each E-tile bank receives other input clock signal, thus having different clock domains within each E-tile xcvr). This section will cover only the input clocks generated by the si5341 PLL included in S10 board, that is used as reference clock to drive the E-tile xcvr's used. [35] gives further info. on other clock sources on S10 devices.

B.8 shows the reference input clocks available within each S10 bank, where highlighted inputs are the clocks generated by the si5341 PLL as shown in figure B.9 (note that in signal naming used in both B.8,B.9 as 'CLK_bank_PLL_freq_p/n' , 'freq' refers to the default clock frequency at PLL startup).

In PAM4 serial link testing system developed, si5341 PLL is configured so that both CLK_9C_PLL_322M_p/n and CLK_9A_PLL_176M_p/n will be 280 MHz¹²⁴ reference clocks used to drive GXE xcvr's channels within each S10 bank¹²⁵ .

CLK_9C_PLL_322M_p/n and CLK_9A_PLL_176M_p/n configured as 280 MHz reference clocks are used within each E-tile xcvr used in PAM4 serial link testing system to derive baudrates for each of the 12 used channels within the bank as:

$$\text{baudrate} = 280 \text{ MHz} \cdot \text{TX_clk_divider}^{126}$$

, where TX_clk_divider is a configurable parameter on per-channel basis.

CDR (Clock and Data Recovery) structure is available within each E-tile bank for each RX-channel. CDR's output clock is used to drive FPGA core

¹²³ note that QSFP-DD modules are also connected to FPGA core through and I2C bus, thus allowing for direct control communication(not involving the associated E-tile bank or EMIB interface)

¹²⁴ each clock signal from si5341 PLL can be configured to output clock frequencies in 125 – 700 MHz

¹²⁵ note that when a 'native-phy-ipcore' is used to control an E-tile xcvr from FPGA design , input reference clock signal (referred to as 'pll_ref-clock0') and GXE-channel serial output/input signals (referred to as TX_serial_data and RX_parallel_data) must belong to the same S10 bank (so that all signals must be mapped to pins belonging to the same S10 bank for each native-phy-ipcore , as each E-tile xcvr is implemented over a separated die)

¹²⁶ note that bitrate will be obtained as 2xbaudrate when using PAM4

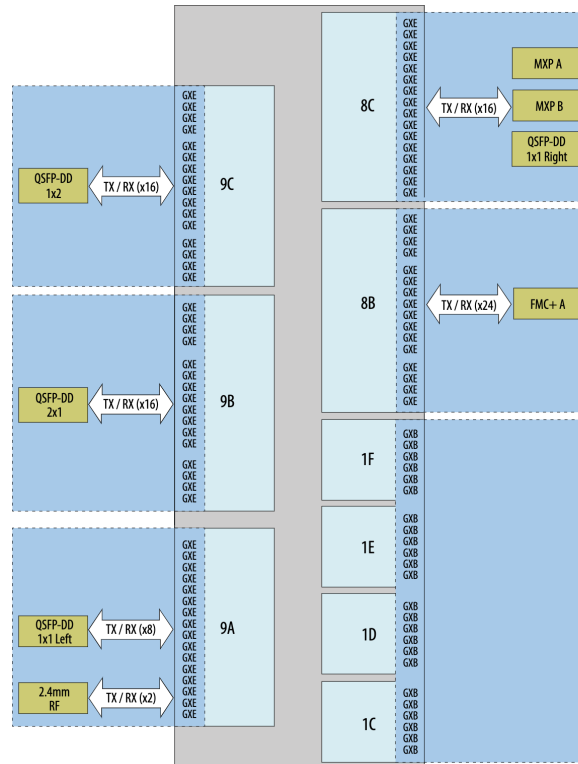


Figure B.7.- E-tile transceiver channels block diagram , including E-tiles transceivers and connector placement among S10 E-tile banks (courtesy of Intel [17]).

logic processing received data (refer to B.2.2.3 sub-section for detail on CDR function).

The core clock network for Intel Stratix 10TX devices support FPGA core architecture working at frequencies up to 1GHz (thus, supporting the maximum frequency used in the PAM4 serial link testing system designed of 437.5 MHz)

B.2.2.3 CDR(CLOCK AND DATA RECOVERY)

CDR structure recovers a signal clock at the carrier frequency of the received data for each RX-channel within each E-tile bank; thus allowing to synchronize the receiver to detect correctly every single bit supplied by the receiver under test, regardless of the datarate , incoming data pattern and jitter affecting the received signal [2].

'native-phy-ipc' instances placed in the PAM4 serial link testing system developed, outputs a clock signal derived from CDR recovered clock as
 $'RX_clkout' = CDRclock / data-width$

, where CDRclock is the recovered clock at received signal's bitrate , and data-width is de data-density used to exchange data through EMIB interface (data-width is 128 for the PAM4 serial link system described, as double density is required for both PAM4 and NRZ transmissions). (note that 'RX_clkout' naming matches signal names used for system functional division in figure 5.4)

B.2.3. DOUBLE WIDTH TRANSFER MODE

1x 64-b EMIB interface is assigned to each of the 24 channels, for each E-tile xcvr.

If PAM4 high-data rate is used within an E-tile bank , each channel performing PAM4 transmission requires a data-parallel interface to FPGA core wider than the 64-b interface available through its EMIB interface, so as to keep the transfer frequency over the parallel interface within the supported FPGA core's frequencies (thus to keep the both 'tx_clkout' and 'RX_clkout' clocks exposed by each 'native-phy-ipc' instance at supported frequencies; 5.4 shows usage of TX_clkout' and 'RX_clkout' clocks).

So that, when PAM4 high-datarate transmissions are performed, EMIB re-assigns channel's physical 64-b parallel data-interfaces such that 12 out of the

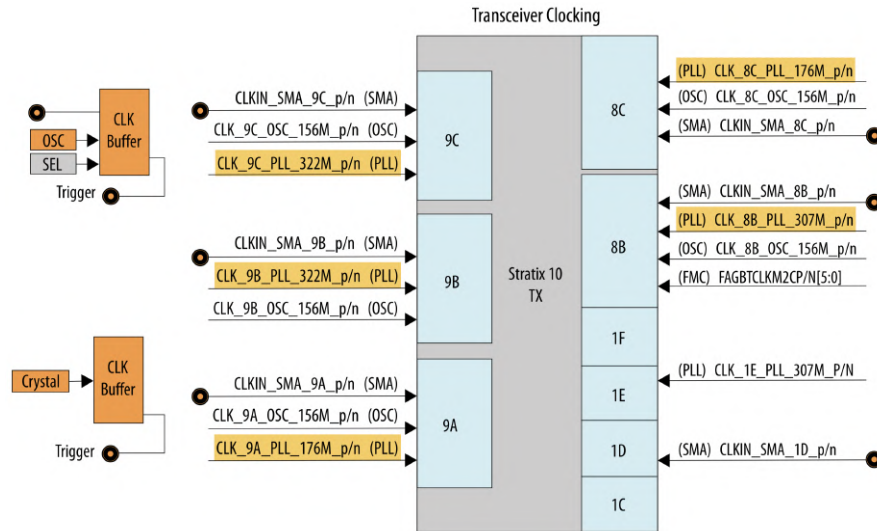


Figure B.8.- E-tile transceiver dedicated input reference clocks (courtesy of Intel [17]).

24 channels will use 2x64-b EMIB interfaces, thus disconnecting the remaining 12 GXE channels in the E-tile.

So that, 12 GXE channels will be capable of actively transmitting/receiving PAM4 signals using 128-b parallel data-interfaces through EMIB to/from FPGA core; as 2 EMIB 64-b parallel interfaces will be allocated for the data transmitted/received over 1 single serial channel.

So that when double-width interfaces are used through E-tile's EMIB interface, only 12 out of the 24 GXE channels (numbered 0 – 23) within the E-tile can be used.

EMIB does the re-assignment by keeping even channels active and re-assigning the adjacent 64-b odd-interface (corresponding to each odd-numbered channel) to each active even channel (the described re-assignment is shown in B.11, where odd, unavailable GXE channels are shaded; and their 64-b interfaces are re-assigned to the adjacent even, active GXE channel to build a double-width 128-b parallel data-interface, B.10).

So that, in PAM4 serial link testing system the 0-11 numbered channels actually correspond to GXE 0, GXE 2, GXE 4 .. GXE 22 channels.

If using 2 EMIB interfaces (thus double-width data interfaces through EMIB), both 'tx_clkout' and 'RX_clkout' must be reconfigured so that generated clocks will be adjusted to the doubled-density data-interface (i.e. to output clocks with half the frequency than used when 64-b data interfaces are assigned). B.1 table shows selectable frequencies for 'tx_clkout' (same applies for 'RX_clkout'), where PMA-width = 64 is a fixed value reflecting the width of physical data-interfaces through EMIB.

TX-EQUALIZATION PARAM	RANGE	STEP SIZE	DESCRIPTION
vod (ATTN)	$0 \leq \text{ATTN} \leq 26$ Increment decrement by 1	18.5 mV/step (0.17 dB)	de-emphasis tx-equalization FIR filter's main cursor
pre-tap 1 (PRE1)	$-10 \leq \text{PRE1} \leq 10$ Increment decrement by 2	18.5 mV/step (0.34 dB)	de-emphasis tx-equalization FIR filter's 1st pre-cursor
pre-tap 2 (PRE2)	$-15 \leq \text{PRE2} \leq 15$ Increment decrement by 1	9.25 mV/step	de-emphasis tx-equalization FIR filter's 2nd pre-cursor
pre-tap 3 (PRE3)	$-1 \leq \text{PRE3} \leq 1$ Increment decrement by 1	9.25 mV/step	de-emphasis tx-equalization FIR filter's 3rd pre-cursor
post-tap 1 (POST)	$-18 \leq \text{POST} \leq 18$ Increment decrement by 2	18.5 mV/step (0.34 dB)	de-emphasis tx-equalization FIR filter's 1st post-cursor

Table B.1.- Parallel clock (TX_clkout) definitions, [35]

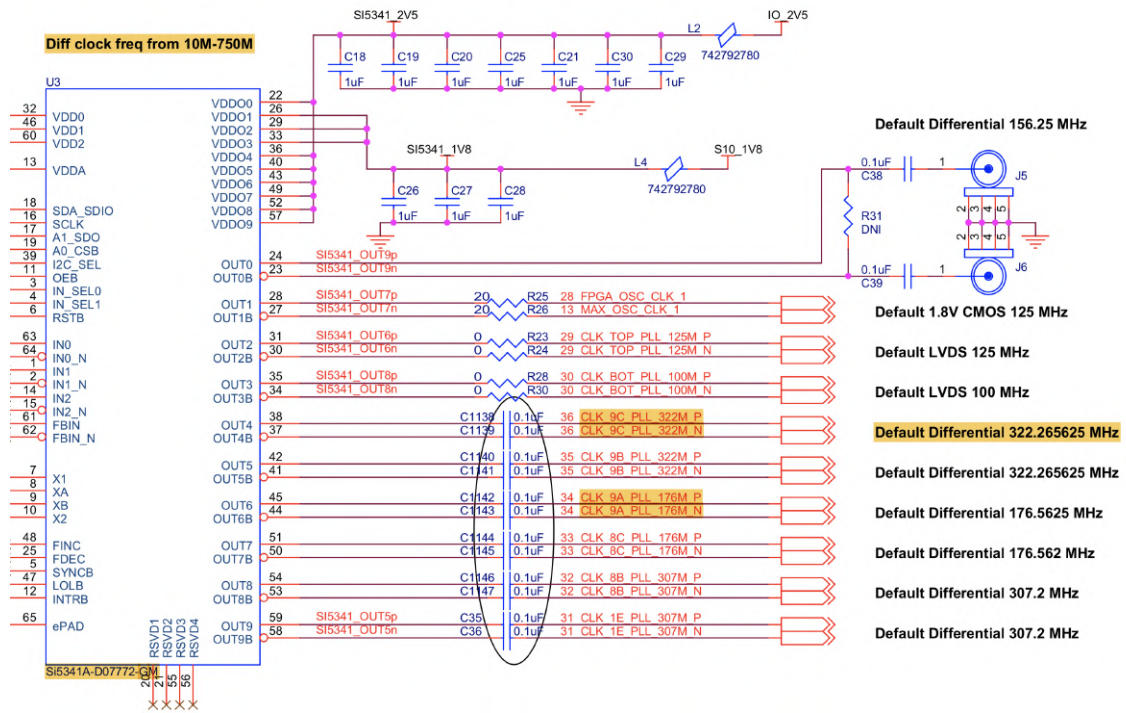


Figure B.9.- E-tile transceiver dedicated input reference clocks generation schema on Stratix 10TX devices (courtesy of Intel [36]).

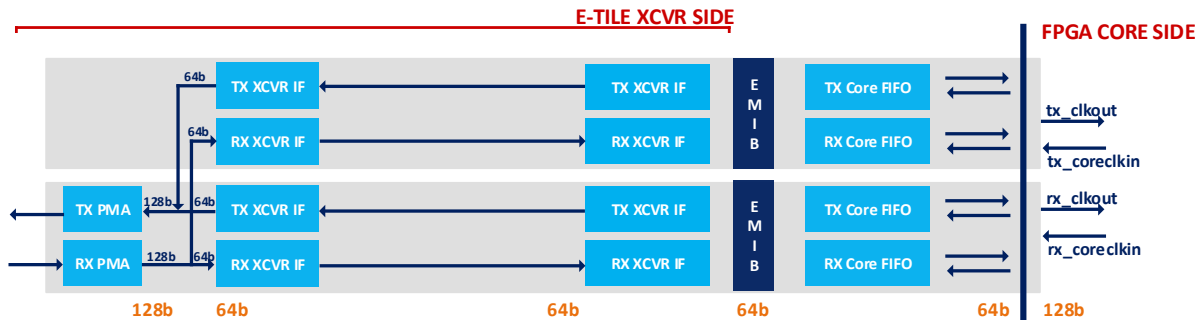


Figure B.10.- Data interfaces width on E-Tile - FPGA core signal path when double width transfer selected (EMIB 64-b parallel interfaces re-allocated) (courtesy of Intel [35]).

B.2.3.1 DATA ALIGNMENT (DE-SKEW)

If 2 EMIB 64-b width data-interfaces are used to support PAM4 high-datarate is used within an E-tile xcvr, skew can occur between the data transferred to/from FPGA core over each EMIB interface (as the 2 EMIB interfaces involved become populated at different times, 64 UIs apart, where 1UI = 1/bitrate).

Native-phy-ipcres implement a de-skew function to mitigate that effect on RX side, which aligns the data received over the two EMIBs' 64-b data-interfaces.

EMIB interfaces are by default 80-b width for each GXE channel, from which up to 64-b are populated with data from the incoming signal being deserialized in 1:64 de-serializer (converting high-speed serial data to 64-b parallel buses over the 2 64-b EMIB data-interfaces used).

B.12 shows how deserialization is done for each of the 2 EMIB interfaces involved, keeping 79:72 and 39:32 lanes reserved for control info transfer over the EMIB interface (note that left-most 80-b interface is the interface between FPGA-core and e-tile bak; and right-most lane is the serial single physical link under test. B.3.1 sub-section gives thorough info on PMA named block.).




Figure B.11.- S10 E-Tile transceiver GXE used even channels on used E-Tile PHYs (courtesy of Intel [35]).




Figure B.12.- Parallel data interface bit assignment in 80-b EMIB parallel interface (with 64-b data out of 80-b) (E-tile - S10 FPGA interface) (courtesy of Intel [37]).

80-64 remaining bits over the EMIB interface contain control info, for each of the 2 EMIB interfaces.¹²⁷

EMIB's control info include 1 de-skew lane (on each of the 2 EMIB 80-b interfaces used for PAM4 high-speed GXE channels), over which transmitter sends a de-skew pulse used to align data at RX-side.¹²⁸

B.13 figure [11] shows how de-skew pulses are used on RX-side to align data received over the 2 EMIB interfaces involved (note that ln0_datain[33]




Figure B.13.- Double width transfer data alignment between odd-, even- channels 64-b parallel data lanes (de-skew logic performance) (courtesy of Intel [11]).

and In1_datain[33] are the control lanes in the 80-b EMIB interface that are used to send de-skew pulses).

B.3. SIGNAL PATH

B.3.1. GXE TRANSCEIVER DATAPATH

B.14 shows the internal structure for each GXE channel where PCS and PMA segments feature 'hard-coded' channel dedicated resources (note that B.14 is a copy of B.6 included in this sub-section for the sake of clarity).

PCS(Physical Coding Sub-Layer) is xcvr segment that typically prepare parallel data for transmission across the channel (and reverse on RX-side), handling encoding/decoding, scrambling/descrambling. PMA(Physical Medium Attachment) converts digital signals to analog domain and reverse, providing interfacing capabilities to TX/RX physical channels.

Intel Stratix 10 TX PMA interfaces with core logic through configurable and by-passable PCS interface layer (for each GXE channel). 'PMA direct high data rate PAM4' xcvr configuration is used in the PAM4 serial link testing system so that data is transferred directly between PMA interface and FPGA fabric through EMIB over 64-b data-interfaces by-passing PCS, as shown in figure B.15.

GXE channels can be configured in PMA direct + FEC mode, thus by-passing every block within PCS segment but RS-FEC encoder/decoder. PMA direct + FEC mode cannot be enabled from the PAM4 serial link testing system implemented, as CEI-OIF specification provides pre-FEC BER limits.

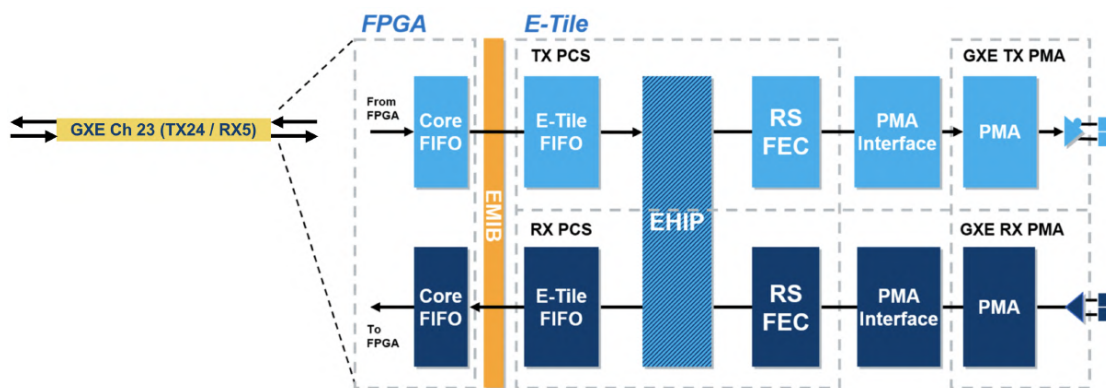


Figure B.14.- E-Tile GXE channel xcvr simplified internal structure block diagram (courtesy of Intel [16]), copy of B.14 included here for the sake of clarity..

¹²⁷ note that EMIB interfaces on diagrams in this appendix are marked as 64-b or 128-b to ease explanation by considering only data-lanes

¹²⁸ de-skew pulse is generated by PRBS31 generator in PAM4 serial link testing, [11] includes step-by-step explanation on how de-skew pulse must be generated.

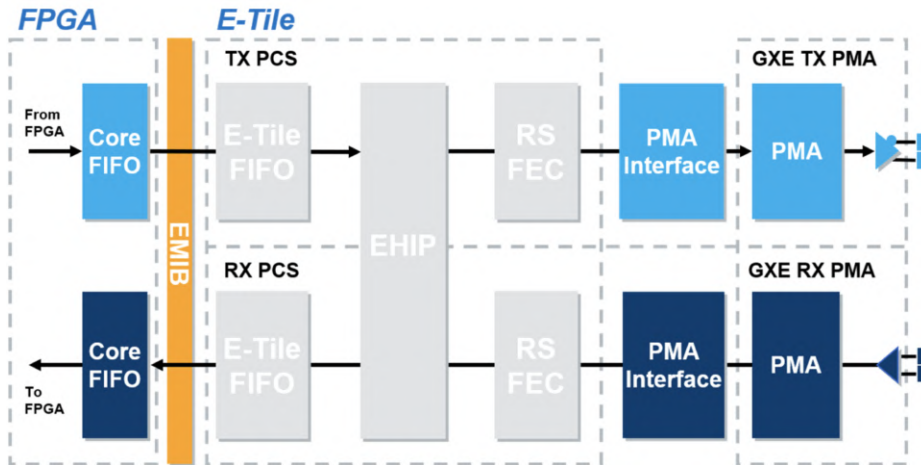



Figure B.15.- E-Tile GXE channel xcvr simplified internal structure block diagram in PMA direct mode, PCS block by-passed (courtesy of Intel [16]).

B.3.1.1 PMA DATAPATH


B.16a ,B.16b figures show internal signal path for both TX and RX PMAs, respectively (PMA signal path is not explained in this appendix as is depicted in 5.1.3 sub-section for both TX and RX).

B.4 sub-section provides more detail on PMA equalization settings configuration (note that CDR(clock data recovery) described in B.2.2.3 is included within the equalization structure in the PMA block for each GXE channel; thus , each PMA has an independent channel PLL that allows analog tracking for clock-data recovery).

B.17 shows complete PMA architecture block diagram



(a)



(b)

Figure B.16.- TX(a) / RX(b) PMA internal signal path simplified block diagram for each GXE channel in S10 E-Tile xcvs (courtesy of Intel [16]).

B.3.2. STRATIX 10 E-TILE XCVRS LOOPBACK MODES

'loopback modes' are 'hard-coded' DFT(design for test) signal paths placed over S10 board for each GXE channels, that can be used to verify different

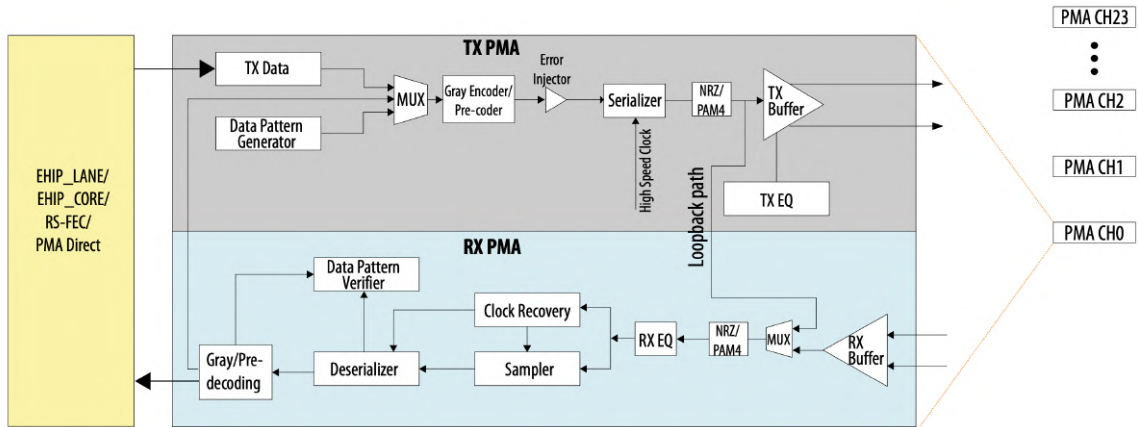


Figure B.17.- PMA internal architecture complete block diagram (for each GXE channel in S10 E-Tile xcvr) (courtesy of Intel [11]).

blocks of the transceiver PMA (note that 'external loopback' is not considered as loopback mode, as both TX/RX are not aware about the loopback condition; thus performing the same as if the serial link were established to/from another channel).

'internal-serial loopback' mode sets CDR to recover data from TX-serializer's output instead of the receiver serial input pin; thus performing an in-chip loopback (keeping signal inside S10 board traces, thus avoiding distortion introduced by both connectors and serial link).

'reverse-loopback' is supported by MUX highlighted in B.19 that feeds transmitter with data from FPGA-core, 'hard-coded' PRBS generator or the data recovered from the incoming signal on channel's RX.

B.18, B.19 figures show internal signal path within channel's PMA for internal-serial loopback and reverse-loopback respectively. 'internal-serial loopback' is used in PAM4 serial link testing system to verify 'native-phy-ipcore' configuration sequences (thus e-tile xcvr configuration); by confirming error-free PRBS pattern reception is achieved in under zero-distortion conditions, guaranteed by 'internal-serial loopback' mode.

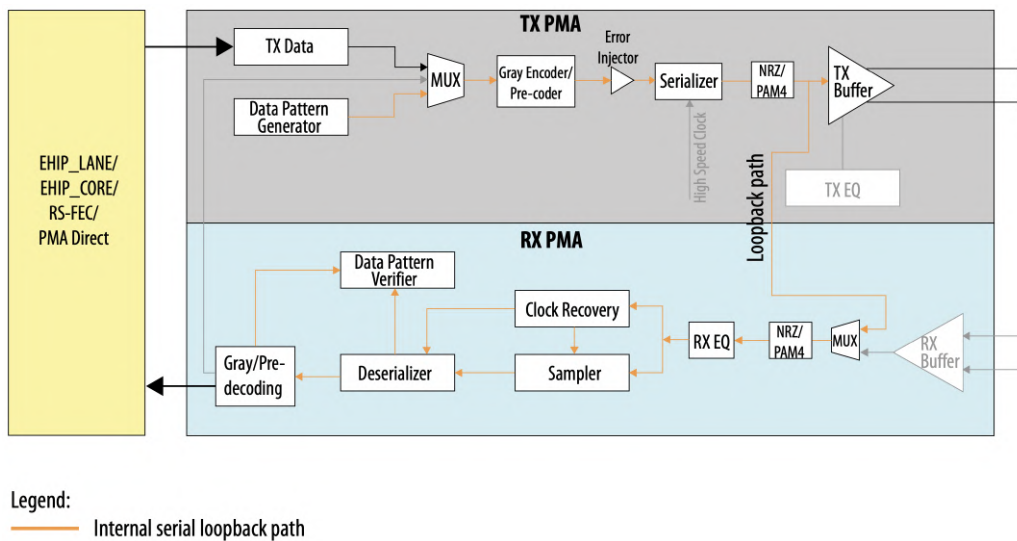


Figure B.18.- Internal serial loopback path within PMA internal architecture on S10 E-Tile xcvr GXE channel (courtesy of Intel [11]).

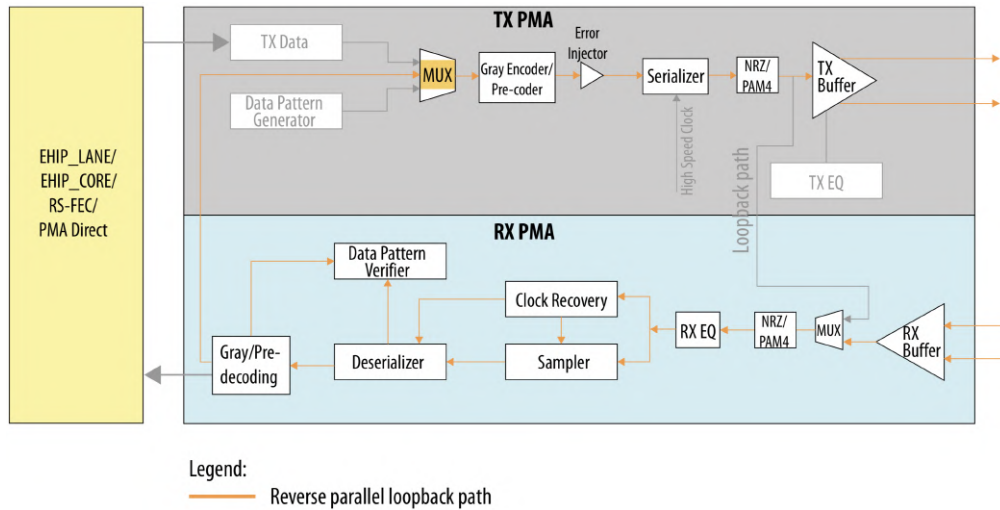


Figure B.19.- Reverse serial loopback path within PMA internal architecture on S10 E-Tile xcvr GXE channel (courtesy of Intel [11]).

B.4. PMA EQUALIZATION

This section aims to describe PMA supported adaptation modes for RX-equalization settings adjustment; and provide an insight on equalization aspects not covered in 3 section (note that RX-equalization structure, B.20 is described in 3 section, thus the contents in this section will assume RX-equalization structure understanding) ¹²⁹.

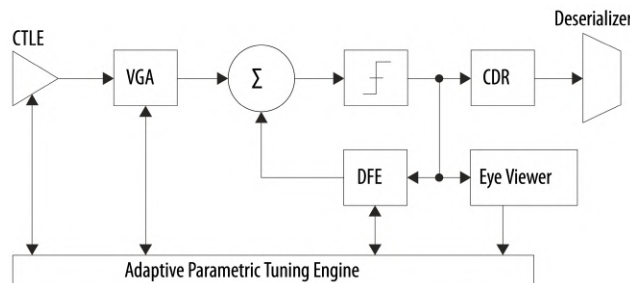


Figure B.20.- Intel S10 TX receiver equalization architecture block diagram including on-die instrumentation (courtesy of Intel [7]).

B.4.1. RX-EQUALIZER ADAPTATION MODES

'RX-equalization adaption' refers to PMA parameters tuning to adjust RX-equalization parameters to channel condition, aiming mainly to overcome channel response variations across temperature range.

If temperature increases during an ongoing transmission, channel distortion increases causing signal quality degradation and so, eye-diagram horizontal and vertical closure (2.3 sub-section describes eye-diagram based signal integrity analysis).

RX-equalization structure shown in B.20 includes an 'adaptive parametric tuning engine' that adapts RX-equalization parameters (DFE tap weights, adjust CTLE parameters, and optimize VGA gain ¹³⁰ and decision threshold voltage) for optimal performance under changing environmental conditions.

¹²⁹ note that B.20 corresponds to the S10 RX-equalization schema included in [7], that adds on-die instrumentation (eye-view, adaptive parametric tuning engine) to the RX-equalization structure depicted in 3 section

¹³⁰ note that VGA(variable gain amplifiers) are common active elements included at receivers' input for signal conditioning



'adaptive parametric tuning engine' can be used to perform RX-equalization parameters adjustment for:

case 1: optimum link performance at static temperatures (by running an initial adaptation when invariant temperature conditions are expected during the transmission to be performed, thus in controlled temperature conditions). (initial-adaptation-mode)

case 2: optimum link performance when temperature sweeps are expected.(continuous-adaptation-mode)

So that, RX-PMA can be configured in one of the following adaptation modes ¹³¹ :

initial-adaptation mode

during initial adaptation all RX-equalization parameters are adjusted to optimize eye opening at RX-side and adjust both vertical and horizontal sample point location properly (6.1.2 sub-section explains the implications of sampling 'point' location). 'initial-adaptation' disrupts ongoing transmissions as all equalization settings are optimized, thus initial-adaptation shall be run only when there is any change on physical channel requiring or transmission re-start [11].

continuous-adaptation mode

is a continuously low-impact adaptation process that runs continuously in background without disrupting ongoing transmissions. 'continuous-adaptation' tracks temperature over time changing RX-equalization settings in small steps to overcome slight temperature variations (note that its low-impact nature – based on varying RX-equalization settings in small steps to avoid disruptions on going transmission– makes continuous-adaptation process incapable of adapting to major temperature changes).

So that, when continuous-adaptation mode is selected, an initial-adaptation is run over the channel to set an initial optimized RX-equalizer configuration from which continuous-adaptation process can keep adjusting RX-equalization parameters [11].

'adaptive parametric tuning engine' performs RX-equalizer adaptation adjusting settings to obtain minimum achievable BER (for both initial and continuous-adaptation).

B.2 table indicates which RX-equalization parameters can be adjusted using RX-adaptation modes [38]('firmware default' valued indicated the default parameter valued when PMA is started for each GXE RX-channel).

Parameter	Min	Max	Initial Adaptation	Continuous Adaptation	Manual Optimization Possible	Firmware Default
GainLF	0	15	Yes	Yes	Yes	8
CTLE LF Min	0	15	N/A	N/A	Yes	0
CTLE LF Max	0	15	N/A	N/A	Yes	15
GainHF	0	15	Yes	Yes	Yes	0
CTLE HF min	0	15	N/A	N/A	Yes	0
CTLE HF max	0	15	N/A	N/A	Yes	15
GS1	0	3	No	No	Yes	0
GS2	0	3	No	No	Yes	0
RF_P2	-10	10	Yes	No	No	0
RF_P2_MIN	-10	10	N/A	N/A	Yes	-10
RF_P2_MAX	-10	10	N/A	N/A	Yes	10
RF_P1	0	15	Yes	Yes	No	0
RF_P1_MIN	0	15	N/A	N/A	Yes	0
RF_P1_MAX	0	15	N/A	N/A	Yes	15
RF_P0	-15	15	Yes	Yes	No	0
RF_B1	0	8	Yes	Yes	Yes	0
RF_B0	0	5	Yes	Yes	Yes	0
RF_B0T	10	50	No	No	Yes	20
RF_A - NRZ	100	160	No	No	Yes	160
RF_A - PAM4	100	160	No	No	Yes	130

Table B.2.- PMA RX-equalization parameters supported ranges (definition for parameters in table can be found in 5.4)[11]

¹³¹ note that only continuous-adaptation,initial-adaptation modes are available for RX-equalization adjustment on e-tile transceivers

B.4.2. EQUALIZATION

B.4.2.1 SLICER

RX-slicer (shown in figure B.20 right before CDR block) can be understood as receiver's decisor, that determines certain signal level thresholds used to decide which modulation symbol corresponds to each sample of the received signal (thus the element on charge of decode the input signal samples to output the corresponding modulation symbols).

So that, if channel response were known and invariant, slicer thresholds could be adjusted so that every signal-sample level were decode without error (as the distortion suffered by each symbol sent – signal-level– would be known).

'RX-decision-blocks' included on e-tile xcvsr consist of 3 slicers (required to determine 3 different threshold levels between PAM4 signal levels) with varying voltage thresholds to detect each PAM4 signal value.

B.21 figure shows a theoretical example indicating how the 3 thresholds should ideally¹³² be set.

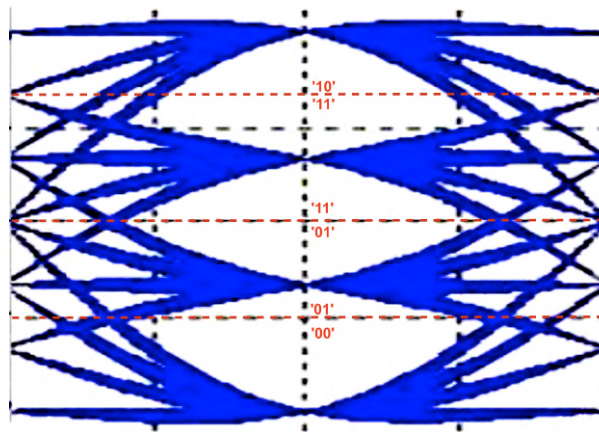


Figure B.21.- Example signal level thresholds for PAM4 RX-slicer measuring distortion-,ISI- free eye-diagram (3 symmetric inner-eyes without misalignment).

B.5. XCVR CONTROLLER INTERFACE (NPDM AND NATIVE-PHY IP ARBITRATOR)

NPDM(native phy master debug endpoint) interface is added to each native-phy-ipcore when ipcore is instanced. NPDM connects to the internal avalon-mm slave that can be used for dynamic re-configuration of the e-tile xcvr.

The exposed avalon-mm slave (that is controlled from Nios-2 in PAM4 serial link testing system) is actually connected to an arbitrator (together with NPDM interface) instead of directly to internal avalon-mm slave; thus the arbitrator controls whether re-configuration requests over the external avalon-mm slave are delivered to the internal 'true' avalon-mm slave for xcvr re-configuration , or not [11].

B.22 figure shows a diagram of native-phy-ipcore's avalon-mm slave interface connection for xcvr re-configuration , with the structure described.

NPDM interface enables Intel's debugging tools to access native-phy-ipcore's avalon-mm slave for xcvr reconfiguration and control (thus, allowing TTK(transceiver tool-kit) to measure on-die eye-diagram over xcvr's channels).

The arbitration logic supports multiple avalon-mm master connections aiming for xcvr re-configuration:

NPDM (used by Intel's debugging tools)

user re-configuration logic (avalon-mm slave exposed when 'native-phy-ipcore' is instantiated; used by Nios-2 in the PAM4 serial link testing system)

So that, these feature arbitrates the control over the internal avalon-mm slave for xcvr re-configuration/control; thus deciding which avalon-mm master delivers its requests to the internal slave to exert control over the xcvr.

If both NPDM and user-reconfiguration-logic try to make a re-configuration/control request at the same time, the arbitrator gives access to the re-

¹³² considering a distortion-,ISI- free eye-diagram , and 3 aligned symmetric-inner-eyes symmetric

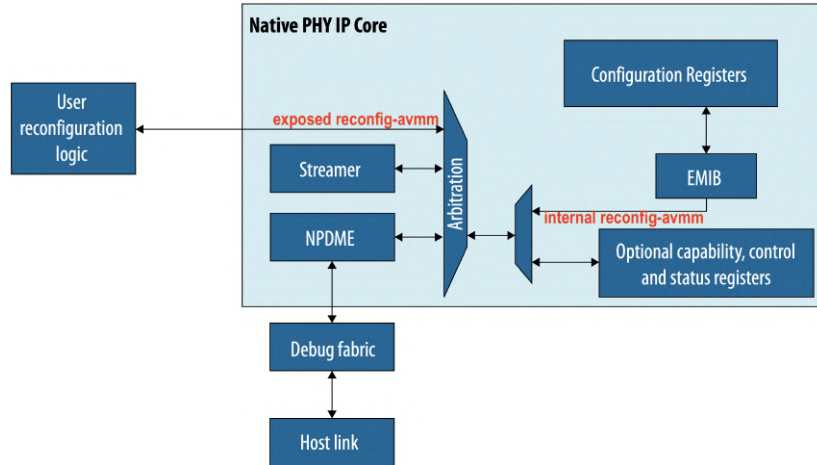


Figure B.22.- Arbitration logic structure supporting multiplexation on native-phy-ipcore reconfig-avmm internal reconfiguration interface (courtesy of Intel [11]).

configuration request coming from the highest priority block.

NPDME is the lower priority re-configuration block connected to the arbitrator; thus xcvr reconfiguration/control requests issued by user-reconfiguration-logic will be performed before orders received from NPDME (note that when one re-configuration block issues a request while a lower priority block is performing a re-configuration operation, it must wait until the ongoing operation is completed; though the ongoing operation owner has lower priority).



C. PRBS31 PARALLEL GENERATOR/CHECKER

This appendix describes the generation of 128-b parallel QPRBS31 generator and checker using an LFSR + xor-tree structure [39], to support high-speed PRBS31 sequence generation (note that besides references placed throughout this appendix, QPRBS31 generator/checker designs described in this section are inspired in U.S patents [43],[44] and designs in [45],[46]).

Stratix 10TX e-tile xcvr include 'hard-coded' multi-PRBS both generator and checker that can be enabled to feed the transmitter and check the received sequence at RX-side, only the error count over the received sequence is exposed.

As this work aims for further error statistical analysis over the received sequence, the error pattern detected by the PRBS31 checker is required, thus justifying the development described in this appendix.

Furthermore, the 'hard-coded' PRBS generator and checker stop the error analysis when a subsequence of certain length is affected by transmission errors, thus missing error data for error-prone transmissions with poor signal quality at RX-side (further detail on this aspect is given in sub-sequent sections).

So that, both QPRBS31 generator and checker are developed to expose the 'complete' error pattern affecting the received sequence, thus enabling for bitwise error-position tracking as required for further statistical analysis QPRBS31 checker developed also exposes both error-count and bit-count.

QPRBS31 generator/checker pair described in this appendix are designed to generate/receive QPRBS31 sequences over 128-b parallel interfaces, thus supporting the generation/checking at bitrates up to (and beyond) 57.8Gbps though working at a lower frequencies, supported by FPGA core¹³³. 128-b width was selected for consistency with the data-interfaces by the e-tile xcvr.

C.1. PRBS31 (QPRBS31) TEST PATTERN

PRBS sequences are used for serial link testing as non-deterministic data-signals cause the 'worst-case' channel stress¹³⁴.

PRBS sequences are 'pseudo-random' in the sense that the occurrence frequency for '1' and '0' are close to 50%¹³⁵.

PRBS sequences while generated with deterministic algorithm (so that can be re-produced and thus checked), preserves statistical behavior similar to random sequences (occurrence probabilities for '1', '0' symbols are close to 0.5) thus causing greater stress on channels.

C.2. LFSR-BASED SERIAL TEST PATTERN GENERATOR

PRBSn sequences can be generated using synchronous LFSR(linear feedback shift register) with n DFFs(d-latch) connected in series (see figure C.1) where the 1st DFF's input is fed from a linear combination of the outputs of every DFF in the LFSR.

PRBS generation is based in binary sequences' properties and can be generated. LFSR's feedback expression can be expressed as polynomial $g(x)$ of degree n:

$$g(x) = 1 + \sum_{j=1}^n c_j \cdot x^j = 1 + c_1x + c_2 \cdot x^2 + \dots + c_n \cdot x^n \quad (7)$$

where n is the number of DFFs in the LFSR and c_j the associated coefficients¹³⁶ defined over GF(2) (binary-subspace with modulus 2 arithmetic)¹³⁷ (note that $x_1 \dots x_n$ value is referred to as LFSR state).


¹³³ the maximum frequency used in FPGA core for the system developed is below $56E9/128 = 437.5$ MHz

¹³⁴ non-deterministic signals cause greater stress on channels under test as its non-deterministic nature makes them noise-alike

¹³⁵ for ideal random sequences '1', '0' symbols turn out to be equally probable when a long enough sequence is considered.

¹³⁶ when $c_j = '1'$ the output of DFFj is 'xored' in LFSR's feedback

¹³⁷ in GF(2), $\text{xor} \equiv ('+' \text{ mod } 2)$



'1+' term in $g(x)$ stands for 'not-gate' in figure C.1, as 'xor 1' \equiv 'not'.

C.1 figure shows the schema for a n-DFF LFSR based PRBSn generator

If $g(x)$ is generator polynomial for $GF(2^n)$ (subspace of n-cipher numbers in base 2), the LFSR generates a cyclic binary sequence with periodicity $2^{31}-1$ (meaning that $2^{31}-1$ bits are generated before any periodicity pattern is found in the generated sequence) [38].

PRBS31 is generated using $g(x) = 1 + x^{28} + x^{31}$ (as indicated in CEI-OIF[8], IEEE 802.3[40]. C.2a ,C.2a figures show the structure for PRBS31 serial generator,checker , respectively (PRBS checker has the same structure as the generator, futher explanation is given in C.3 sub-section for the parallel checker structure)




Figure C.1.- Generic schema for synchronous n-LFSR(n-DFF LFSR)-based PRBSn generator (c_j defined in $GF(2)$ for $j=1, \dots, n$; and $c_n = 1$)

PRBS31 pattern parallel generator described in C.3 sub-section shall produce the same result as the implementation shown in figure C.2a for the polynomial generator [41]:

$$g(x) = 1 + x^{28} + x^{31} \quad (8)$$

C.3. PRBS PARALLEL GENERATOR/CHECKER

C.3.1. PRBS GENERATOR

PRBS serial generator/checker (as shown in figures C.2a , C.2b) can be turned into parallel structures to generate m consecutive bits of the PRBS sequence at a time,on each clock edge (over a m-b parallel interface); using transition matrixes.

C.3.1.1 TRANSITION MATRIX SERIAL-PARALLEL TRANSFORMATION

PRBS generation using a n-DFF LFSR can be expressed using matrix formulation. PRBS9 generation example will be used to explain serial-parallel transformation using transition matrixes for simplicity.

T matrix in 9 is the transition matrix for a PRBS9 LFSR generator with generator polynomial $g(x) = 1 + x^5 + x^9$, with the 9-DFF LFSR structure in figure C.3 [42].

$$M = \begin{matrix} & \begin{matrix} d_1 & d_2 & d_3 & d_4 & d_5 & d_6 & d_7 & d_8 & d_9 \end{matrix} \\ \begin{matrix} s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \\ s_6 \\ s_7 \\ s_8 \\ s_9 \end{matrix} & \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \end{matrix} \quad (9)$$




Figure 49-9—PRBS31 pattern generator

(a)




Figure 49-11—PRBS31 pattern checker

(b)

Figure C.2.- PRBS31 test pattern serial generator (a); checker (b), (generator polynomial $1+x^2+8+x^3+1$), [41].




Figure C.3.- PRBS9 serial generator using 9-DFF LFSR structure (generator polynomial $g(x) = 1 + x^5 + x^9$)

where s_i is input of DFF $_i$ and d_i is output of DFF $_i$.

So that, each row in T gives the feedback expression applied to each DFF $_i$: each s_i ($s_i = \text{DFF}_i$'s input value) is calculated by 'xor-ing' every d_j such that $(s_i, d_j) = '1'$ in T; thus:

$$s_i = \sum_{j=1}^n d_j \cdot T(s_i, d_j) \quad (10)$$

, where $n=9$ for PRBS9. ¹³⁹

s_1 for PRBS9 example is given by $d_5 \text{ xor } d_9$ (as shown in figure C.3); and $s_j = d_{j-1}$ (with $j \neq 1$) as the rest of latches in the LFSR are fed from the output of the previous DFF in the row.

So that, the state evolution of the LFSR (being $d(n) = [d_1 \dots d_9]$ referred to as LFSR's state) can be obtained as: $d(n) = T \cdot [s_1 \dots s_9]$, where n is the number of clock cycles elapsed from the start of the generation (thus, the number of PRBS sequence's bits generated); and $[s_1 \dots s_9]$ the current values at DFFs' inputs.

$d(n)$ matrix formulation allows to calculate the state of the LFSR (the outputs of each DFF) at any point of the PRBS sequence generation (i.e. when any number of bits have been already generated), as :

$d(l) = T^l \cdot d(0)$, where $1 \leq l \leq N$, $l > 0$ and $d(0)$ is the initial state of the LFSR (the set of values with which LFSR's DFFs outputs, d_j , are loaded initially;

¹³⁸ note that the longer the sequence, the longer the transmission time elapsed until being repeated; thus periodicity can be disregarded and the sequence generated can be considered non-deterministic

¹³⁹ note that '+' equals to xor in GF(2)

referred to as LFSR's seed)¹⁴⁰.

So that T^l can be used to implement an LFSR structure that will 'advance' by l states on each clock edge, thus generating an output $o(n)$ (PRBS9_OUT in C.3) = $\{o(0), o(1), o(2), \dots\}$ (instead of $\{o(1), o(2), o(3), \dots\}$) where n is the number of clock edges arrives to the LFSR from the start of the generation.

Such LFSR is implemented using T^l to determine the feedback expression for each DFF.

If $l = 4$ is selected for the PRBS9 generator, [42]:

$$M^4 = \begin{matrix} & d_1 & d_2 & d_3 & d_4 & d_5 & d_6 & d_7 & d_8 & d_9 \\ \begin{matrix} s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \\ s_6 \\ s_7 \\ s_8 \\ s_9 \end{matrix} & \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix} \end{matrix} \quad (11)$$

so that $s_1 = d_2 \text{ xor } d_6$ (e.g.), meaning that DFF₁ will be fed by xor-ing the outputs of DFF₂, DFF₆.

PRBS9 4-b parallel LFSR structure corresponding to transition matrix ('feedback matrix') in 11 is shown in figure C.4¹⁴¹.




Figure C.4.- PRBS9 4-b LFSR-based parallel generator (generator polynomial $g(x) = 1 + x^5 + x^9$)

C.3.1.2 128-b PARALLEL GENERATOR FEEDBACK EQUATIONS

PRBS31 128-b parallel generator is implemented from the transformation of the LFSR shown in figure C.2a into a LFSR structure that produces PRBS31 pattern (referred to as $\{o(n)\}$, where $n = 0, 1 \dots 2^{31}-1$ from here on, within this appendix) in 128-b steps (outputting $o'(n) = o(0) o(127) o(255) \dots$).

12 shows the transition matrix T associated to the serial PRBS31 generator in figure C.2a, T .

T^{128} used to transform the serial LFSR generator is shown in 13 (T^{128} (31x31) matrix provides feedback equations for the 31 DFFs in the target PRBS31 LFSR producing $o'(n) = o(0) o(127) o(255) \dots$ instead of $o(n) = o(0) o(1) o(2) \dots$ (where $o(n)$ is n -th bit of the PRBS31 sequence and $n > 0$).

' $g(x) = 1 + x^{28} + x^{31}$ generator polynomial is used to obtain the intermediate $[o(128k) \dots o(128k+127)]$ with $k = 0, 1, 2, 3, \dots$, from each state of the LFSR $d(n) = [d_0 d_1 d_2 \dots d_{30}] = \text{not}([o(128n) \dots o(128n+30)])$.

So that $[o(128n) \dots o(128n+30)]$ are generated using feedback equations given by transition matrix 13; and remaining $[o(128n+31) \dots o(128n+127)]$ are generated from the current LFSR state, on each clock edge (thus generating 128-b PRBS31 bits on each clock cycle.), using equations derived from generator polynomial $g(x)$.

¹⁴⁰ if $d(1) = T d(0)$ then $d(2) = T d(1) = T(T d(0)) = T^2 d(0)$, and so on

¹⁴¹ note that DFFs in figure C.4 are not connected in 'chain', but 'isolated' as each DFF has its own feedback structure

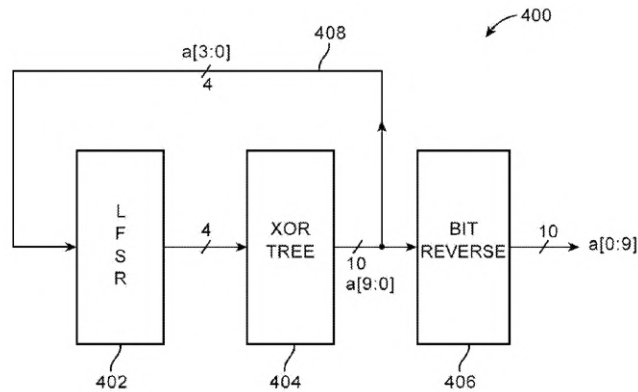


FIG. 4

Figure C.5.- LFSR + XOR-TREE structure used for PRBS31 parallel generator implementation , [39].

pattern, so that the generating a test pattern with the statistical properties of the PRBS31 sequence, but with even $(2) \cdot (2^{31} - 1)$ bit length.

C.3.2. PRBS CHECKER

128-b parallel PRBS31 checker has the same LFSR + xor-tree structure as the generator; so that it will be able to generate the same PRBS31 128-b 'sub-sequence' on each clock cycle .

So that, the PRBS31 checker uses each generated 128-b PRBS31 'sub-sequent' to verify that the received data (from the 128-b parallel data input port) matches the expected 128-b subsequence.

PRBS31 verification works as explained, assuming that PRBS31 checker is generating the 'inside-reference' PRBS31 sequence used for comparison 'at the same point within the PRBS31 sequence' as the received sequence.

So that if the received sequence is not affected by errors, both the received 128-b sequence and the 'inside-reference' 128-b sequence (generated inside the checker), correspond to the same $o(k) \dots o(k+127)$ sub-sequence of the PRBS31 pattern.

DFFs 'pre-load' logic is modified for PRBS31 checker so that when the checker is connected to a 128-b parallel received signal , the checking can be started with independence on the point (within the cyclic PRBS31 pattern) at which the received signal is.

If considering LFSR's feedback equations, it can be seen that the 'next' LFSR's state is obtained as $d(n+1) = [d_0 \dots d_{30}] = o(128n + 98) \dots o(128n + 127)$, where $o(128n + 98) \dots o(128n + 127)$ are the last 30 bits of the received 128-b PRBS31 'sub-sequence'.¹⁴²

So that, the 'next-expected' 128-b sub-sequence of the PRBS31 pattern can calculated from the last 30-bit of the last received 128-b sub-sequence; thus enabling the PRBS31 checker to be started at any point in the PRBS31 received sequence and though being able to 'auto-synchronize' from the 1st 128-b sub-sequence received, to start checking for the 2nd 128-b subsequence received (assuming that the las 30 bits of the 1st 128-b sub-sequence received are not affected by errors).

PRBS31 checker 'auto-synchronization' is thus achieved by pre-loading the LFSR's DFFs with the last 30-bit of the 1st received 128-b sub-sequence when the checker is started; from then on, the PRBS31 checker keeps generating the reference PRBS31 sequence.

PRBS received pattern is checked on each clock edge by performing a bitwise xor with the reference PRBS31 sequence; so that the PRBS31 checker will be able to detect every single error bit , and also its position within PRBS31 pattern¹⁴³ .

PRBS31 checker outputs each 128-b error pattern obtained from the xor operation over a 128-b parallel interface, for further statistical error analysis

¹⁴² note that only 30 bits of the incoming PRBS31 received pattern must be error-free so that the developed PRBS31 can auto-synchronize (and thus start accurately checking the incoming sequence); while 'hard-coded' PRBS checker in the e-tile xcvr requires (40)x(128) non-error bits before start accounting for error in the received sequence, thus missing error related data in comparison to the developed PRBS31 checker.

¹⁴³ xor result is '1' for every mismatching bit 'position' (within the 128-b sub-sequence), thus for every errored bit in the 128-b received sub-sequence



(see D appendix for detail on further statistical analysis over error pattern).

PRBS31 checker also outputs an error counter that is incremented by one for each incoming bit error in the PRBS31 pattern for isolated single bit errors , as required by IEEE 802.3 standard [41].

PRBS31 checker design described is based in the 'PSEUDO-RANDOM BINARYSEQUENCE CHECKER WITH AUTOMATIC SYNCHRONIZATION' defined in [10].

C.3.2.1 LOCKED CONDITION

In addition to the error counter, the 'hard-coded' PRBS checker in the e-tile xcvr also outputs a control/status signal referred to as 'prbs_locked'. 'prbs_locked' signal indicates when the number of errors in the received PRBS pattern is lower enough to guarantee that PRBS verification is accurate.

'hard-coded' PRBS checker requires 40x128 non-error received bits to start checking the received sequence (thus, 40 128-b 'sub-sequences' must be received without errors for the 'prbs_locked' signal to be asserted)

PRBS31 checker described here differs in the sense that only 40x30 bits are required to assert an analogous 'prbs_locked' output status signal (being that '40' a configurable parameter set to 40 just for consistency).

Besides the behavior when 'prbs_locked' signal is de-asserted was modified for the designed checker.

PRBS31 checker developed starts by pre-loading the LFSR with the last 30 bits from the 1st received 128-b sub-sequence; and calculating the expected 128-b sequence in the next clock edge.

- If the last 30 bits of the new 128-b sub-sequence received match with the expected values the 'prbs_locked' will keep asserted ('hard-coded' PRBS checker requires all 128-b to be error-free)
- If there is a mismatch in those 30 last bits, for 40 consecutive clock edges, 'prbs_locked' is de-asserted and the LFSR will be pre-loaded again in every clock-edge until the last 30 bits received match again the expected values ('hard-coded' PRBS checker does not pre-load again its LSFR, thus not adapting to changes in the received pattern if not reset).
- If that match occurs, the PRBS31 checker waits 40 clock edges before re-asserting 'prbs_locked' signal:
If all 40 128-b sub-sequences are received without errors affecting the las 30 bits of each sub-sequence, 'prbs_locked' signal is re-asserted; otherwise the LFSR's DFFs are pre-loaded again from the received sequence ('hard-coded' PRBS checker requires all (40x)128-b to be error-free).

PRBS31 generator and checker described in this appendix were implemented in matlab in 1st instance , so that the design could be conceptually verified by comparing the 1st segment of the sequence generated with a sample provided in [41],[47]; and then 'translated' into VHDL (note that verified matlab implementations of the PRBS31 generator and checker was used later for verification purposes , see C.3.3 sub-section for detail on PRBS31 generator/checker verification).

C.3.3. VERIFICATION

PRBS31 generator and checker described were validated both in simulation and over FPGA , using altera_avalon_data_pattern_generator, altera_avalon_data_pattern_checker Intel's ipcores [50] as reference for comparison (note that altera_avalon_data_pattern_generator, altera_avalon_data_pattern_checker behaves equal to 'hard-coded' PRBS generator and checker included in e-tile xcvr).

C.3.3.1 SIMULATION BASED VERIFICATION

PRBS31 generator and checker behavior was verified ins simulation environment thus avoiding real-hw related issues that could affect the functional verification.

Both PRBS31 generator and checker , and altera_avalon_data_pattern_generator, altera_avalon_data_pattern_checker ipcores were instantiated; and stimulus signals were generated so that generators and checkers will be synchronized and start aligned (thus starting the generation of the PRBS31 pattern at the same point within the PRBS sequences)¹⁴⁴ .

¹⁴⁴ for alignment purposes, PRBS31 generator developed is initialized with 'all ones' seed for LFSR's DFFs




So that, a native-phy-ipcore instance was configured to keep performing a 56Gbps transmission , thus outputting a 437.5MHz 'TX_clkout' clock signal that was used to drive both PRBS31 generator and checker developed , and altera_avalon_data_pattern_generator, altera_avalon_data_pattern_checker.

Signal-tap was configured to take 128 consecutive samples of the PRBS31 patterns generated by both PRBS31 generator developed and altera_avalon_data_pattern_generator; and also of the error counters exposed by both PRBS31 checker developed and altera_avalon_data_pattern_checker. Signal-tap sampling throw identical results as shown in simulation-based verification results (see figure C.6) thus verifying design's performance.

Signal-tap samples for PRBS31 patterns generated by PRBS31 generator developed a were also exported to .csv format; and imported in matlab to confirm that sequence generated by the 128-b parallel generator matched the sequence generated by the already verified matlab implementation for the PRBS31 serial generator (as aforementioned in C.3.2.1).

PRBS31 sequences' match was verified for the 1st 512 bits in the PRBS31 test pattern, thus validating the design described in this appendix.



D. STATISTICAL ERROR ANALYZER, RS(544,514) FEC SPECIFIC

This appendix aims to describe the implemented VHDL entity for high-speed statistical error analysis from a 'continuous' error pattern received over a 128-b parallel interface.

The designed error analyzer is aimed to provide data related to statics that enable for analyzing the impact of using RS-FEC(544,514) in PAM4 serial links¹⁴⁵.

D.1 sub-section gives a brief into RS-FEC(544,514) random error correction and burst error correction capabilities to justify the approach used in the design of the error analyzer described in sub-section D.2 (note that this appendix focuses only on error correction capabilities, as system's goal is to keep the measured BER below the target limit).

D.1. RS(544,514) FORWARD ERROR CORRECTION (FEC)

D.1.1. RS-FEC IMPACT ON PAM4 BER

RS-FEC is supposed to be supported on PAM4 capable xcvr to be used for serial transmissions over 18GBs [8].

RS-FEC used relax the maximum BER limit required up to $BER < 1E-4$, as RS-FEC(544,514) is reported to reduce BER at RX-side down to $1E-15$ after error correction is applied (figure D.1 shows BER reduction reported in [8] for some FEC algorithms).

RS-FEC strategies are considered sufficient if able to keep measured BER below system's target BER, assuming a normal, uniform random error distribution, plus implements burst randomization to reduce the probability of uncorrectable error bursts using interleaving, see D.1.1.2 sub-section for explanation in interleaving techniques.




Figure D.1.- Input to output BER when implementing FEC on serial link transmissions for common FEC codes [8].

D.1.1.1 FEC GAIN

BER reduction capabilities of a FEC algorithm can be expressed in terms of supported signal distortion.

FEC coding gain is used to refer to the reduction in SNR that can be accommodated when FEC is used; while keeping BER below the target limit.

¹⁴⁵ FEC (forward error correction) strategies introduce redundancy to data blocks such that each block of k symbols becomes a codeword of n symbols where the added $n-k$ symbols are used to correct errors occurred during the transmission, at RX-side, to a certain extent

If assuming normal link operation conditions ¹⁴⁶, RS(544,514) presents about 7 – 8 dB coding gain.

D.1.1.2 BURST CORRECTION ENHANCEMENT (INTERLEAVING)

Interleaving techniques spread error bursts on n consecutive codewords, thus turning error bursts of length l into n shorter bursts of length l/n. So that, the maximum correctable burst length for a FEC code capable of correcting bursts up to k bits, will be increased up to nk when FEC is combined with interleaving (refer to [51] for analysis on RS(544,514) performance when inter-leaving techniques are used).

D.2 figure shows example of typical n-branch interleaving structure with n=4 that reflects how sent symbols x_j are spread over the signal sent such that consecutive symbols are spaced by n symbols (note that analogous structure is required at receiver side to revert the interleaving and recover the original symbol sequence).




Figure D.2.- Typical n-branch interleaving structure for n=4 (Δ represents 1UI delay synchronized latches).

D.1.2. FEC ERROR CORRECTION CAPABILITIES

FEC introduces redundancy in data by adding a number of overhead bits that determine the extent of the correction capabilities. This sub-section aims to explain how the error correction capabilities are determined from the no. of redundancy bits added, for binary sequences (the no. of redundancy symbols are referred to as 'c' within this section).

D.1.2.1 CORRECTION OF UP TO M RANDOM ERRORS

In order to correct up to m bit-errors within a block of length L, all possible error patterns must be recognizable. So that the number of possible combinations for the redundancy bits added, 2^c , must be at least equal to number of possible error patterns of length $\leq m$.

So that, the lower bound for 'c' can be obtained as [20]:

$$c_{rec}(m) \geq \log_2 \sum_{j=0}^m \binom{L}{j} \quad (14)$$

where $c_{rec}(m)$ is the no. of redundancy bits required to correct up to m random error bits, and the binomial coefficient gives the no. of possible error patterns of length L containing j errors ¹⁴⁷ (note that $c_{rec}(m)$ bits are assumed to be included in the L sequence).

A different approach eases the calculation, the distance between codewords must be at least 2m + 1 bits (meaning that at least 2m + 1 must be different between every 2 valid codewords); so that in case m errors occur (or less), there will always be codeword differing with the received word in less number of bits than the rest of valid codewords (shorter distance).

So that, a shorter distance base decoding can correct the errored codeword without ambiguity. 14 equation and 2m+1 expression, provide the same lower

¹⁴⁶ no signal fading or other unusual phenomenon

¹⁴⁷ note that 1 out of the 2^c possible combinations generated for $c_{rec}(m)$ redundancy bits must identify the absence of errors; thus m = 0 is considered in equation 14



bound for m.

D.1.2.2 CORRECTION OF BURSTS UP TO LENGTH λ

In order to correct error bursts up to length λ , all possible bursts of length ≥ 1 (and up to λ) must be recognizable (as explained for random error patterns in previous sub-section). (note that the no. of possible bursts of length j is given by $(L-j+1) \cdot 2^{j-2}$, where 2^{j-2} are the no. of possible burst generated with j bits¹⁴⁸; and $(L-j+1)$ are all possible start positions for a burst of length j within a L -bit sequence¹⁴⁹)

15 equation gives the lower bound for the no. of redundancy bits required to correct bursts of length up to λ [20]:

$$c_{bec}(\lambda) \geq \log_2(L + \sum_{j=2}^{\lambda} (L-j+1) \cdot 2^{j-2}) \quad (15)$$

where $c_{bec}(\lambda)$ is the no. of redundancy bits required to correct bursts of length up to λ ;

$\sum_{j=2}^{\lambda} (L-j+1) \cdot 2^{j-2}$ gives the no. of possible bursts of length j within a L -bit sequence, $\forall j \leq \lambda$;

and L is no. of single bit errors that can occur within the L -bit sequence (note that c_{bec} bits are assumed to be included in the L sequence).

D.1.3. RS-FEC(544,514) ERROR CORRECTION CAPABILITIES

RS are non-binary codes defined over $GF(2^k)$ ¹⁵⁰, instead of $GF(2)$ but can be analyzed similarly to binary codes if considering symbols as k -bits groups instead single bits.

RS-FEC (544, 514) used in PAM4 serial link transmissions is defined in [40](IEEE Std 802.3TM-2015 clause 91.5.2.7 (section six)) as a non-binary code over $GF(2^{10})$.

If a RS code over $GF(2^{10})$ is able to correct 1 errored symbol over a received block, all the 10-b will be replaced by the correct value 10-b symbols with independency on what 'error-symbol' occurred¹⁵¹; thus being able to correct all possible 10-b error sequences regardless if there is only 1 error bit, or all the 10 bits in the symbol (thus, being able to correct up to 10 single bit errors).

This improves RS codes' behavior for burst error correction in comparison to binary codes [52],[53].

D.1.3.1 CORRECTION OF UP TO M RANDOM ERRORS

$2m+1$ expression given D.1.2.1 sub-section is also valid for symbol-based calculation so it is used to determine the no. of correctable random symbols, m as:

$$m < d-1/2 \quad 152$$

where d is the minimum distance between 2 valid codewords.

$d = n-k+1 = 31$ for RS-FEC(544,514)¹⁵³, so up to $m = 15$ random error-symbols (10-b error symbols over $GF(2^{10})$) can be corrected for each codeword¹⁵⁴.

In $GF(2^{10})$ each error symbol(10-b symbol $\neq 0$) can contain 1 – 10 error bits, thus 15 error-symbols can contain 15 – 150 errored bits.

¹⁴⁸ note that error bursts of length λ are defined as $(\lambda-2)$ length sequences limited by 2 error bits (thus sequences with edge bits = '1' and the central $(\lambda-2)$ bits can take any value)

¹⁴⁹ note that in order to correct an error burst of length j within a L -bit sequence $\forall j < L$, burst position must be known

¹⁵⁰ thus its coefficients for both message symbols ($m_{k-1} \dots m_0$), and generation polynomial ($g_{2t} \dots g_1 g_0$), belong to $GF(2^k)$, instead of $GF(2)$

¹⁵¹ note that if each symbol is a 10-group defined over $GF(2^{10})$, error symbols can take any possible value in $GF(2^{10})$ except the 'all zeros' 10-b sequence reserved the absence of errors.

¹⁵² derived from $d > 2m+1$ expression given in D.1.2.1 sub-section

¹⁵³ reed-solomon codes have min. distance(no. of different 10-b symbols) $d = n-k + 1$, where $n-k$ is the no. of redundancy symbols in each codeword. [54]

¹⁵⁴ note that error correction capabilities of RS-FEC(544,514) depend on the location of the errors within the 5440-bit sequence



But so, an errored codeword will be uncorrectable if 16 error bits spread into 16 different 10-b symbols (worst correction case).

If considering the worst correction case, a high bound for BER measured when RS-FEC(544,514) can be derived as a BER 15/5440 = 2.757E-3, where 15 is the maximum no. of bits that can be corrected independently on their position within the codeword (as in worst case where they spread over 15 different 10-b symbols, m is not exceeded) [55].

D.1.3.2 CORRECTION OF BURSTS UP TO LENGTH λ

14 equation can be adapted for RS(544,514) over GF(2^{10}) as:

$$c_{bec}(\lambda) \geq [\log_2(L \cdot (2^{10} - 1) + \sum_{j=2}^{\lambda} (L - j + 1) \cdot (2^{10})^{j-2})] / 10 \quad (16)$$

where $L \cdot (2^{10} - 1)$ is the no. of possible values for bursts of L symbols, where each symbol is 10-b length and can take any possible value $\neq 0$; and $\sum_{j=2}^{\lambda} (L - j + 1) \cdot (2^{10})^{j-2}$ was modified such that the no. of possible bursts of length λ (the no. of possible combinations with the inner ($\lambda - 2$) symbols) is $(2^{10})^{(j-2)}$ for 10-b symbols defined over GF(2^{10}).

'log₂' in 14 equation gives the no. of bits required to generate a different redundancy value for each correctable burst; thus the added '/10' divider gives $c_{bec}(m)$ expressed in terms of (10-b symbols).

$c_{bec}(m) = 544 - 415$ for RS(544,514) is used to find λ (max. correctable burst length when burst's symbols are defined in GF(2^{10})); thus determining burst error correction capabilities of RS(544,514).

16 equation was evaluated for increasing λ values until the condition is not met (for pre-known value $c_{bec}(m) = 30$ (10-b symbols)):

$$\log_2(L \cdot (2^{10} - 1) + \sum_{j=2}^{\lambda} (L - j + 1) \cdot (2^{10})^{j-2}) = 299.00 \text{ bits, for } j = 31$$

$$\log_2(L \cdot (2^{10} - 1) + \sum_{j=2}^{\lambda} (L - j + 1) \cdot (2^{10})^{j-2}) = 309.00 \text{ bits, for } j = 32$$

thus obtaining a max correctable length burst $\lambda = 31$ (10-b symbols) = 310 bits.

D.2. STATISTICAL ERROR ANALYZER (RS(544,514) SPECIFIC)

The high speed statistical error analyzer designed operates over the error pattern affecting the PRBS31 sequence received (thus over the 128-b parallel error-pattern output from the PRBS31 checker described in C appendic) to measure burst length (for each burst detected), the no. of error-bits within each burst, and gap length (for each gap detected) (detail on how measurements are performed and variable definitions are given in this appendix)¹⁵⁵.

D.2.1. RS-FEC(544,514) BURST DENSITY

$\Delta 0$ denotes burst density defined as:

$$\Delta 0 = \text{error-bits-in-burst} / \text{burst-length} \quad 156$$

where $\Delta 0$ usually $\in [0.5 - 1]$.

$\Delta 0 = 0.5$ is commonly used for PAM4 transmissions as higher burst density bound, as errored symbols increase burst length by 2, but introduce mostly only 1 error-bit (assuming gray coding is used, thus adjacent symbols differ in only 1 single bit).

D.2.1.1 END CRITERIA (MNZ)

If burst definition given in previous sub-section c is considered, there is no criteria to determine burst's end for ongoing transmissions, as inner ($\lambda - 2$) bits can even be all zero (error-free).

¹⁵⁵ note that overall statistic data are measured (considering the received sequence as a continuous error pattern); but statistics are also measured in per 5440-b block basis, enabling to analyze how if RS(544,514) correction capabilities are sufficient for each serial link under test

¹⁵⁶ where an error burst of length λ is defined as any combination of ($\lambda - 2$) bits (errored or not), delimited by '1's (errored-bits) on the edges[56]



MNZ(minimum number of zeros) is the no. of zeros that must be accounted over an error pattern after the last 1(last errored bit) to consider that the ongoing burst is over (thus the last received '1' marked the end of the burst).

MNZ is calculated as [56]:

$$MNZ > (\text{int})(\text{error_bits} + 1) / \Delta_0 - (\text{bits_in_burst} + 1) \quad (17)$$

where Δ_0 is burst density.

So that, when accounting a detected burst, if the no. of zeros after the last '1' exceeds MNZ, the burst must be considered to have ended at that last '1' received (note that MNZ must be updated on each error received¹⁵⁷).

D.2.2. GAP DEFINITION

A gap within a bit-error pattern is defined as the sequence of zeros (non-errored bits) that occurs between 2 bursts (the error statistical analyzer implemented uses MNZ criteria to determine the end of last burst, and so gap's start).

The gap length distribution is the probability that a gap of length 'm' or greater will occur, which can be written as:

$$P(0^m | 1) \text{ for } m = 1, 2, 3, \dots$$

, where $P(0^m | 1)$ is the probab. of finding m-zeros when the last bit was a '1' (an errored-bit).

The gap length distribution can be obtained by analyzing the gap length measures taken for each gap encountered in an ongoing transmission (note that another distribution of interest is the burst length distribution that can be analogously defined as $P(1^m | 0)$ for $m = 1, 2, 3, \dots$, where $P(1^m | 0)$ is the burst length distribution that can be also obtained from the burst length measurements taken for each burst encountered).

D.2.3. GAP, BURST LENGTH MEASUREMENT

The statistical error analyzer developed performs a bit-by-bit analysis of the received error pattern measuring gap and burst lengths under MNZ criteria as:


1. when 1st '1' is found 'in_burst' signal is asserted ('in_gap' signal is initially de-asserted), both biterror_counter and burst_lenght are increased
(note that after resetting, no measurement is taken until 1st error is received; thus in case an initial run of zeros occurs will not be accounted¹⁵⁸)

MNZ is als initialized as $MNZ = 1/1 = 1$.

2. when in_burst asserted
If a '1' (errored-bit) appears both biterror_counter and burst_lenght are increased; and MNZ updated;
In case of '0' (non-error bit), mnz_counter is initialized to '0' and 'in_gap' signal is asserted (keeping both 'in_burst' and 'in_gap' asserted)
3. when both 'in_burst' and 'in_gap' asserted (a run of zeros after the last '1' that has not exceeded MNZ)
In case of '0', mnz_counter is increased:
If $mnz_counter > MNZ$ last burst is considered over
both biterror_counter and burst_lenght are stored; and 'in_burst' is de-asserted
gap_lenght is initialized to mnz_counter
In case of '1':
burst_lenght is increased by $(mnz_counter + 1)$; biterror_counter is increased; and MNZ updated;
'in_gap' is de-asserted
4. when 'in_gap' asserted
In case of '0': gap_lenght is increased
In case of '1': gap_lenght is stored and error statistical analyzer goes back to (1)

¹⁵⁷ more detail on how the accounting is performed is given in sub-section D.2.3

¹⁵⁸ initial gaps are not accounted as gap is defined as non-error run between bursts; and there is no guarantee of previous errors before error statistical analyzer reset



MNZ is updated only on reception of error-bit in the error statistical analyzer implemented; and bursts are not considered over until MNZ exceeded.

'gap_lenght', 'burst_lenght' and 'biterror_counter' are stored for every gap/burst accounted thus providing data required to obtain gap-lenght, burst-lenght and burst-density distributions (note that burst-density can be calculator from 'burst_lenght', 'biterror_counter' values).

'burst_lenght' and 'biterror_counter' are measured for the overall sequence; but also on per 5440-b block basis to enable for RS-FEC specific analysis (D.2.4 sub-section gives detail error statistics data measured by error analyzer).

D.2.4. STATISTICAL ERROR ANALYZER STRUCTURE

D.3 figure shows the internal structure of the statistical error analyzer implemented organized in 4 entities (note that 4 blocks are synchronized using the signal clock referred to as 'clk_mgmt' that will be the clk_mgmt clock used for system's synchronization in PAM4 serial link testing system (see 5.4 figure in 5.1 section); and the input pattern is analyzed using the input clock 'RX_clkout' that should correspond to the clock output from native-phy-ipc core (see 5.4 figure in 5.1 section) in the system designed).

/D.2.5 sub-section describes how measured data is output for further processing.

4 blocks in the statistical error analyzer designed are depicted below.

- per_word_coun:
 - analyzes the incoming error pattern monitoring for burst/gap occurrence as described in D.2.3 sub-section (thus accounting for 'gap_lenght', 'burst_lenght' and 'biterror_counter'); in a per 5440-b block basis (thus, resetting internal counters every 5440 bits received).
 - If last bit of the 5440 in the block is received when accounting for an 'uncomplete' burst, per_word_coun will report a burst of length 5440.
 - If burst length exceeds 310-b (thus 31 10-b symbols), burst_uncorrectable signal is asserted, indicating that the 'current' codeword cannot be corrected by RS(544,514) burst correction capabilities.
 - If otherwise the 5440 block ends while accounting a gap, per_word_coun will not output gaps's length (note that no more than 1 burst should occur within the same codeword, thus in case of non-fatal transmission no gap_lenght should be reported from per_word_coun)
 - If more than 1 burst is detected within the same codeword, burst_lenght and biterror_counter will be reported, but burst_uncorrectable signal will be asserted (as no more than 1 burst can be corrected per 5440-b word when using RS(544,514); use of burst_uncorrectable signal is described below).
 - per_word_coun also accounts for the no. of error-free received codewords.
- overall_coun:
 - performs overall error pattern analysis thus considering the incoming error pattern as a 'continuous' error sequence instead of per block analysis.
 - The error analysis is equal to the performance described for per_word_coun, differing only in internal counters are not reset unless an external reset is issued; and there is no burst_uncorrectable output
- symbased_coun:
 - accounts for errored 10-b symbols within the incoming error pattern.
 - The accounting for error-affected 10-b symbols is done from both overall and per codeword approach (thus per-codeword counter is reset every 5440 received bits).
 - If the no. of errored 10-b symbols within the 'current' codeword is exceeded, random_uncorrectable signal is asserted, indicating that the 'current' codeword cannot be corrected by RS(544,514) random error correction capabilities.
- uncorrectable_coun:
 - accounts for the no. of words that cannot be corrected by RS(544,514); thus monitors both burst_uncorrectable and random_uncorrectable at the end of each 5440-b received block
 - If both burst_uncorrectable, random_uncorrectable are asserted by the end of the current codeword, the uncorrectable-word-counter is increased

D.2.5. STATISTICS DATA OUTPUT

'statistical error analyzer' developed outputs 'gap_lenght', 'burst_lenght' and 'biterror_counter' to the system designed for each gap/burst encountered within the received signal; so that enabling to obtain gap_lenght, burst_lenght and burst_density distributions (note that burst density can be calculated

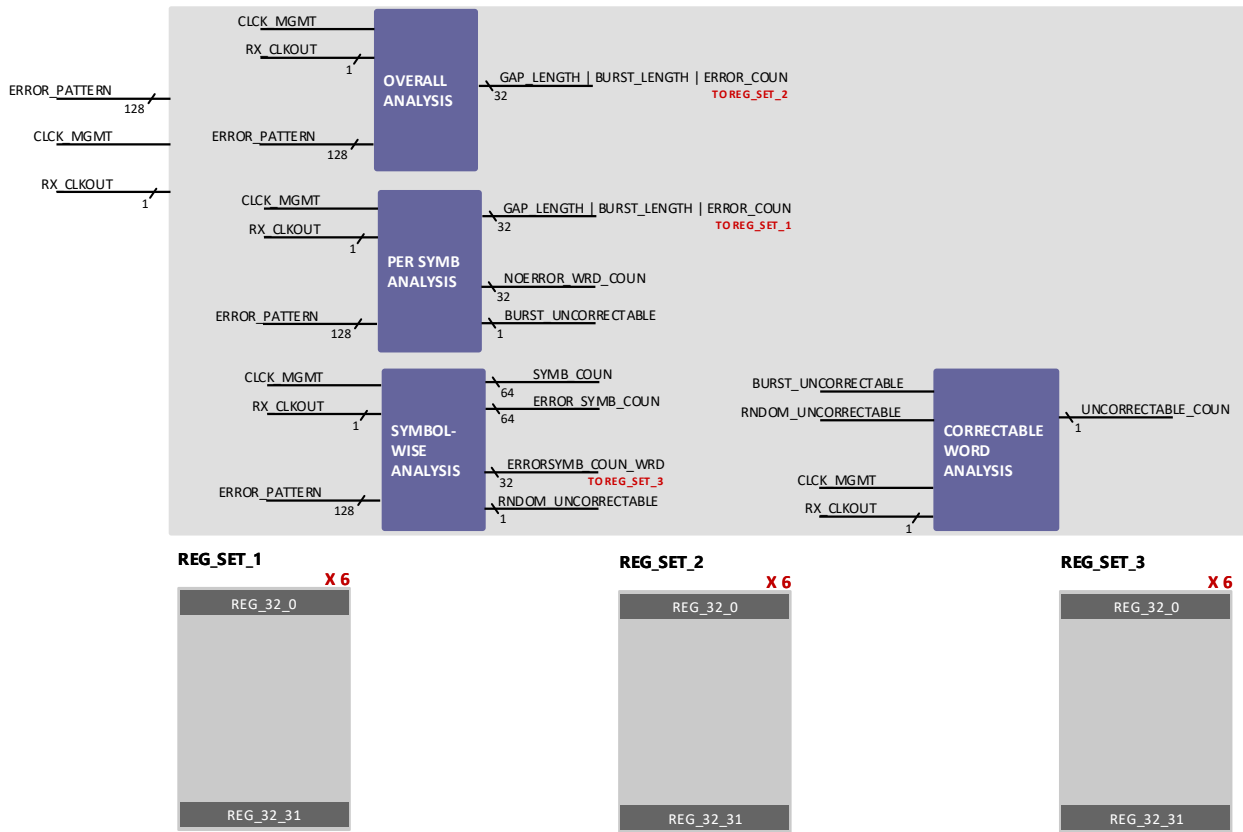


Figure D.3.- Developed VHDL statistical error analyzer internal structure.

as $\text{biterror_counter/burst_length}$ for every burst reported)¹⁵⁹.

Such great amount of data is output using 3 register sets (1 per 'accounting entity': per_word_coun ,overall_coun , symbased_coun), (as shown in D.3 figure); where each regset consists of 6 blocks of 32 registers each (each register is 32-bit width).

So that, whenever error data must be output by an entity, the current block of the associated regset is written (the 32 registers within the block are written form 'reg_0' – 'reg_31').

If all the 32 registers within the 'current' block were already written, a control register is used to indicate the 32-register block can be read; and next data output is done over the next block (once the 6 block are all written once, the entity starts storing the output in block_0 again)¹⁶⁰.

D.3. VALIDATION

The statistical error analyzer developed was functionally verified in simulation basis. So that a testbench including burst generators was implemented to confirm error analyzer's performance.

¹⁵⁹ further analysis on RS impact on BER measured for the serial link under test can be done once distributions are obtained. If burst length distribution shows that burst length is below 31 for most cases (or n-31 when interleaving is used, see D.1.1.2 sub-section for interleaving explanation), then RS-FEC(544,514) will succeed in correcting most bursts

¹⁶⁰ note that in the designed PAM4 serial link testing system , Nios-2 will be on charge of reading each 32-register block when the control signal is asserted (indicating it is ready to be read).

Nios-2 has access to an additional control register used to confirm when a 'register-block' read is completed. If this indicator is asserted, the statistical error analyzer de-asserts its 'ready-to-read' signal; and when noticing the de-assertion, Nios-2 also de-asserts the 'already-read' indicator.

So that, if the statistical error analyzer writes all the 6 register-block of 1 register-set and goes back to block_0 to continue outputting error-data, information loss can be detected (when the 'already-read' indicator is still de-asserted and so, the 'ready-to-read' signal asserted).

In case of information loss, the error statistical analyzer will overwrite the register-block (thus causing older data to be loss), but will inform Nios-2 of that data loss.

D.4 figure shows the structure implemented for the validation testbench.

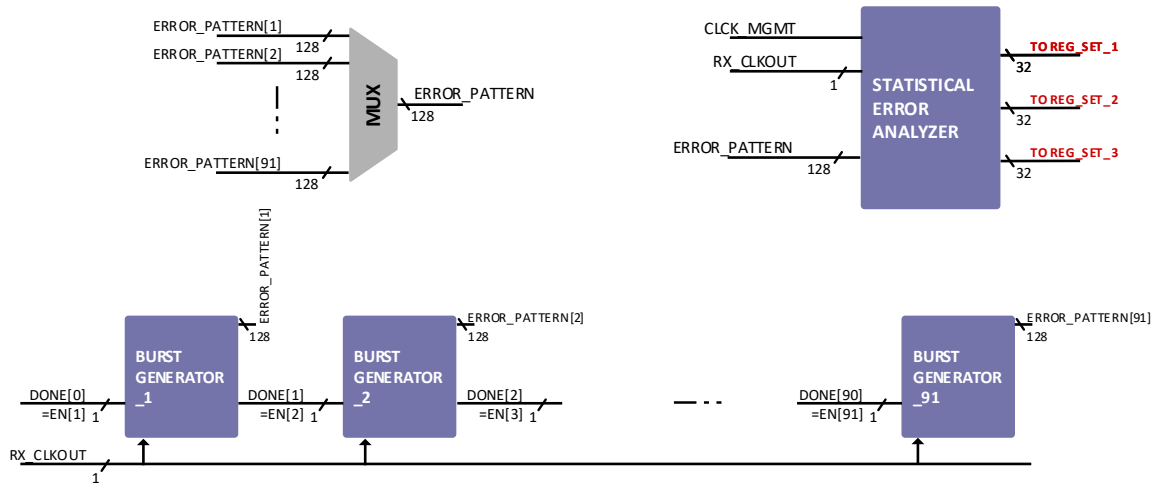


Figure D.4.- Simplified schema for simulation testbench used to validate the statistical error analyzer.

'error-patter-simulation' was done by chaining 91 burst generators (each generating a $(40)(91)$ -bit length burst over a 128-b parallel output interface). 'burst-generators' developed include an output signal that connects assert the enable input of the next generator in the row (note that generators' enable signals are used to determine which of the 91 error-pattern outputs is used to feed the statistical error analyzer).

'burst-generators' are parametrized with the no. of errors to be generated and burst length (note that the 91 chained generators will produce 91 different burst lengths, but all use $\Delta 0 = 0.5$ as set for the statistical error analyzer to ensure MNZ match).

So that, each generator produces an error burst of the indicated length, with the amount of bit-errors indicated

(note that each generator monitors MNZ while producing the error burst using burst-density $\Delta 0 = 0.5$ (same set for the statistical error analyzer).

If the no. of errors indicated to a generator is not enough to reach the required burst length without causing the error analyzer to consider burst end, then the no. of error-bits within the burst generated will be greater than the indicated in parameter value¹⁶¹.

Once the burst is generated, burst-generator keeps outputting zeroes until MNZ is exceeded, thus ensuring that the burst generated by the next generator in the chain will be considered a different burst (when processed by the error analyzer).

D.5 figure shows the results obtained from modelsim simulation, where it can be observed:

- that register-sets are written as expected,
- both in_burst and in_gap signals behave as intended
- symbol-error counters performs accurate accounting
- uncorrectable-word counter gives accurate count

¹⁶¹ the burst-generator developed keeps calculating MNZ for the burst being generated, thus inserting errored bits if required to ensure the no. of zeroes does not exceed MNZ. So that guaranteeing that the error statistical analyzer will analyze the burst produce as only 1 burst)

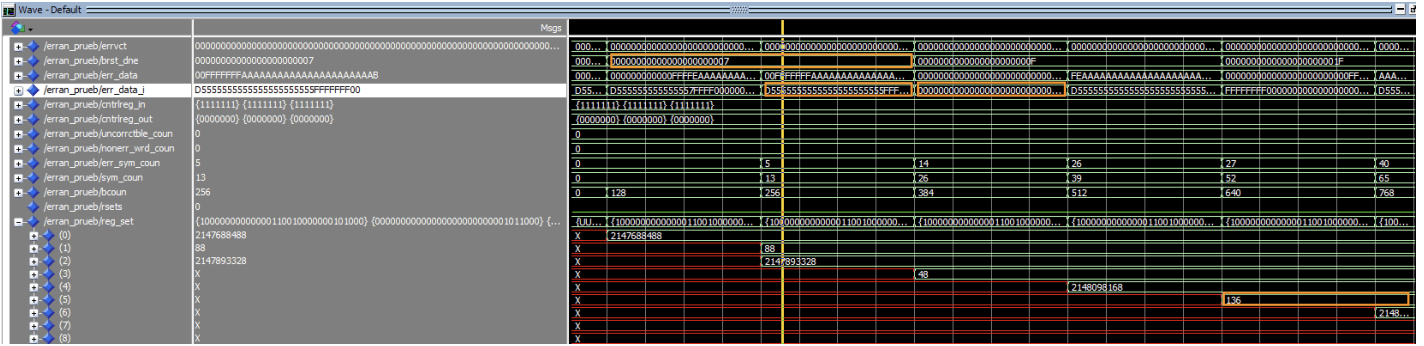



Figure D.5.- Modelsim simulation results for statistical error analyzer validation testbench (note that 128-b width signals are shown in hexadecimal format).

D.5 figure shows simulation results partially, the yellow cursor is placed over the 1st clock cycle for which error pattern (highlighted 'err_data_i') is produced by 3rd generator in the row.

'brst_done' signal is used as burst-generation-finished indicator by each generator. So that each generator has 1 bit associated that is connected to next generator in the row, to enable next burst generation. brst_done = 7 (3 bits asserted) indicates 3rd generator is active (as 1st bit is asserted initially to start the 1st generator in the row)¹⁶².

3rd generator produced 120-b length burst ('err_data_i' value below cursor) followed by 136-b length gap to guarantee MNZ is exceeded (during the gap in_gap and in_burst signals keep asserted as expected¹⁶³).

136 value written in 'reg_set'(highlighted in orange) corresponds to the 136-b gap length observed.

¹⁶² note that bit 'n' in 'brst_done' is asserted 1 cycle before generator n starts for signal stability purposes

¹⁶³ note that in_gap,in_burst signal values are updated on clock edges, so the exact bit for which gap length exceeds MNZ cannot be observed



E. NIOS-2 APP MENU

select action :

- 1 config. test
 - 1 temperature control
 - 1 keep constant temperature (max. deviation before cooling = 10° (default))
 - 2 keep temp. below threshold (temp. threshold = 80° (default))
 - 3 no temp. cntrol
 - 2 stop|end condition
 - 1 enable time limit (max test time = (s) (default))
 - 2 enable numbbits limit (max bits sent = (default))
 - 3 enable time|numbbits limits
 - 4 set min measurable BER
 - x show test config
- 2 run test
 - 1 single channel test
 - 2 run test on all selected phy channels
 - 3 run test on all active channels of active phys
 - x show test progress|status
- 3 etile|channel config
 - 1 set tx PMA equ. settings (vod|pre-tap-(1-3)|post-tap)
 - 1 config single channel PMA tx-equ. settings
 - 2 config PMA tx-equ. settings for all channels in selected phy
 - 2 set PMA rx-equ. params for all channels in selected phy
 - 1 select PMA rx-equ. pre-set
 - 2 manual
 - 3 NRZ|PAM4
 - 4 re-run current adaptaion mode (run initial adaptation + re-set adaptation mode)
 - 0 on selected channel (selected phy)
 - 1 on active channels (selected phy)
 - 2 all channel (selected phy)
 - 3 on active channels (all phys)
 - 5 control serial lpback in selected phy
 - 6 control reverse lpback in selected phy
 - 7 enable|disable gray encoding
 - 8 enable|disable 1/1+D
 - 9 enable|disabñe swizzle
 - a invert tx polarity
 - b invert tx polarity
 - c show eye-h
 - 1 selected channel (selected phy)
 - 2 active channels (selected phy)
 - 3 all channels (selected phy)
 - 4 active channels (active phys)
- 4 reset selected phy
- 5 set selected phy
- 6 set selected channel
- 7 set active channel (for further actions)
- 8 start xcvr for selected channel in selected phy
- 9 start active channels on all phys



- a reset error countr
 - 1 for selected channel on selected phy
 - 2 for active channels on selected phy
 - 3 for active channels on active phys
- b show detailed err countr
- c show continuous BER|err-coun update (only applicable when xcvr running (but not in tst mode))
- d set BER interval (only applicable when xcvr running (but not in tst mode))
- e PMA reset
 - 1 selected channel (selected phy)
 - 2 active channels (selected phy)
 - 3 all channels (selected phy)
 - 4 active channels (active phys)
- f set adaptation mode (runs initial adaptation before setting adaptation mode)
- g sweep PMA param (configs PMA with optim value)(on selected channel,phy)
- h search valid signal (periodic iADP until valid signal found)(on selected channel,phy)
- i dump i2c
- j select tx|rx_clk_dvder (baudrate = 280E6 x divider) (on selected phy)
- k measure BER for supported tx_clk_dvder values 50-100