RS(544,514) FEC performance with 4:1 interleaving (updated 2)

Pete Anslow, Ciena
Including data kindly provided by Rich Mellitz

IEEE P802.3ck Task Force, August 2018

Introduction

*

The IEEE P802.3ck Task Force has objectives to define 100 Gb/s lanes for AUI interfaces, backplanes and twin-axial copper cables for 100 Gb/s, 200 Gb/s, and 400 Gb/s Ethernet. If the FEC sublayer in Clause 91 is re-used for 100 Gb/s Ethernet and the Clause 119 PCS is re-used for 200 Gb/s and 400 Gb/s Ethernet, the 100 Gb/s lanes will have to be formed from interleaving four 25 Gb/s lanes.

healey_100GEL_01_0318 contained some analysis of a variety of interleaving schemes and proposed that pre-coding should be used in 100 Gb/s per lane electrical PHYs as a tool to improve error correction performance.

gustlin_3ck_01_0718 contained some concerns about whether the error pattern seen within bursts from practical receivers matches the error pattern where pre-coding turns each burst into two isolated symbol errors and suggested investigating RS symbol multiplexing as an alternative.

This presentation analyses the performance of a variety of 4:1 interleaving schemes using a development of the principles explained for the NRZ case in Annex 1 of anslow_3bs_02_1114 and prompted by discussions with Geoff Zhang from Xilinx now includes some multi-tap DFE results for both 100G and 400G.

Slides that are changed from those in anslow_3ck_adhoc_01_082918 have a “*” in the top right corner.

Precoding

The results with precoding in anslow_3ck_adhoc_01_072518 were for a 1 tap DFE where it is expected that precoding will have the effect of converting an error burst of any length into just two PAM4 symbol errors.

To try to help answer the question of what the performance will be for a realistic multi-tap DFE with precoding, the Monte Carlo model used for previous presentations was extended to include a 5-tap DFE as below.

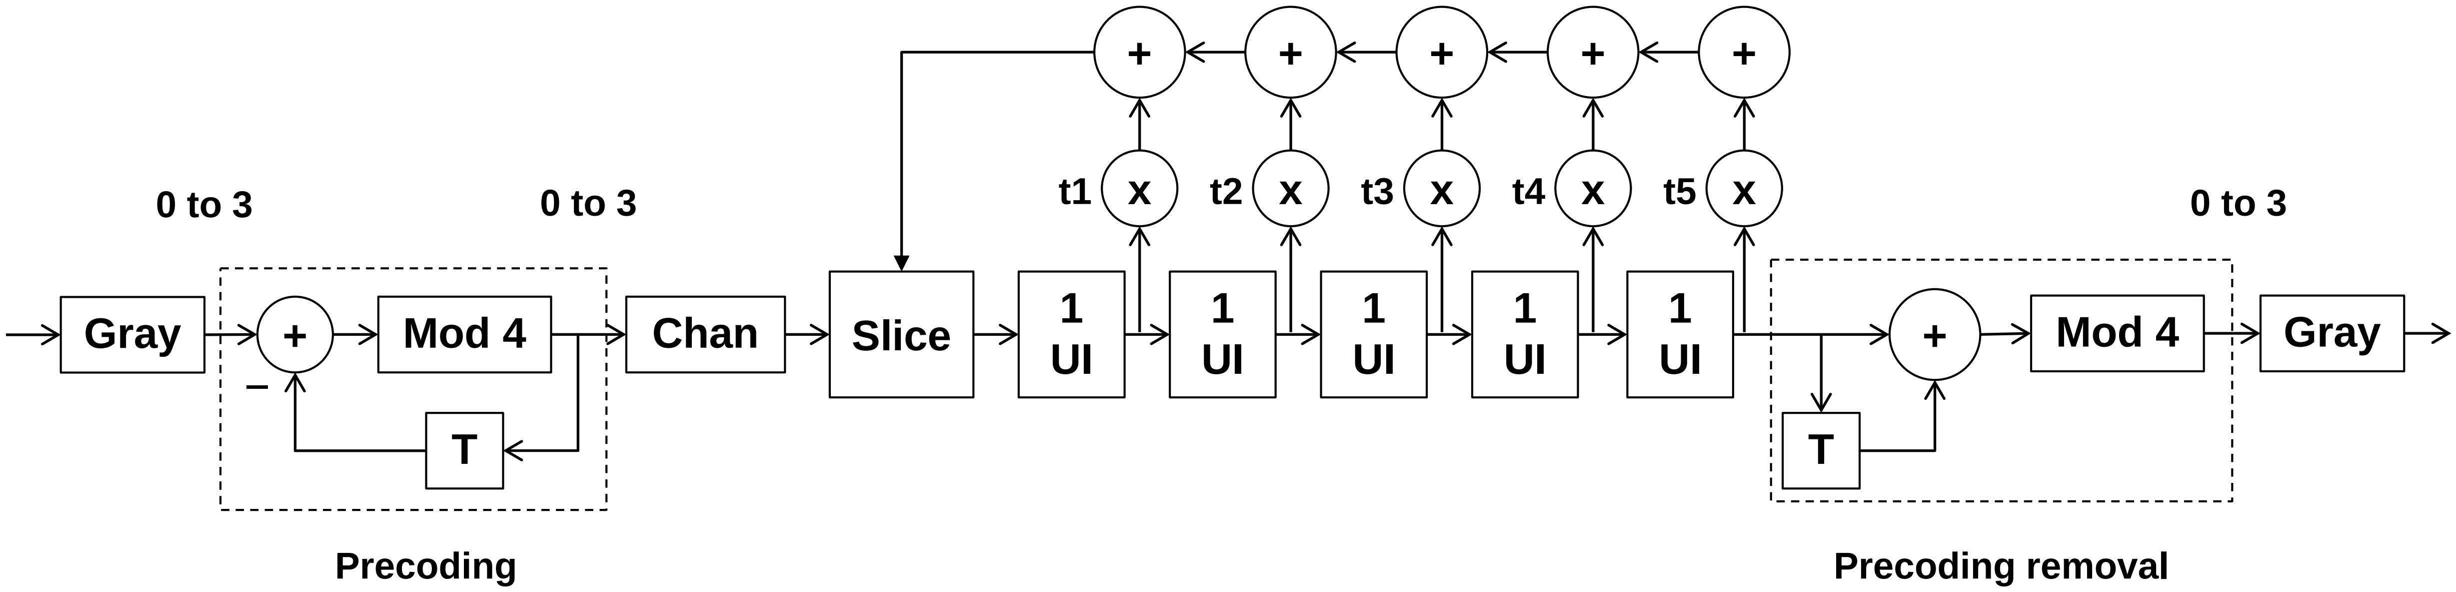

Model validation

The new Monte Carlo model was then used to re-create the previously simulated (1-tap DFE) results.

The error statistics (probability of n FEC symbol errors per burst) were also extracted and used to create dashed curves at lower FLR than could be simulated directly by the Monte Carlo program.

The Monte Carlo values (x markers) and dashed curves generated as above are plotted on the next slide together with the solid curves generated using probability analysis as per the principles explained for the NRZ case in Annex 1 of anslow_3bs_02_1114.

100G previous result validation (1-tap DFE)

*

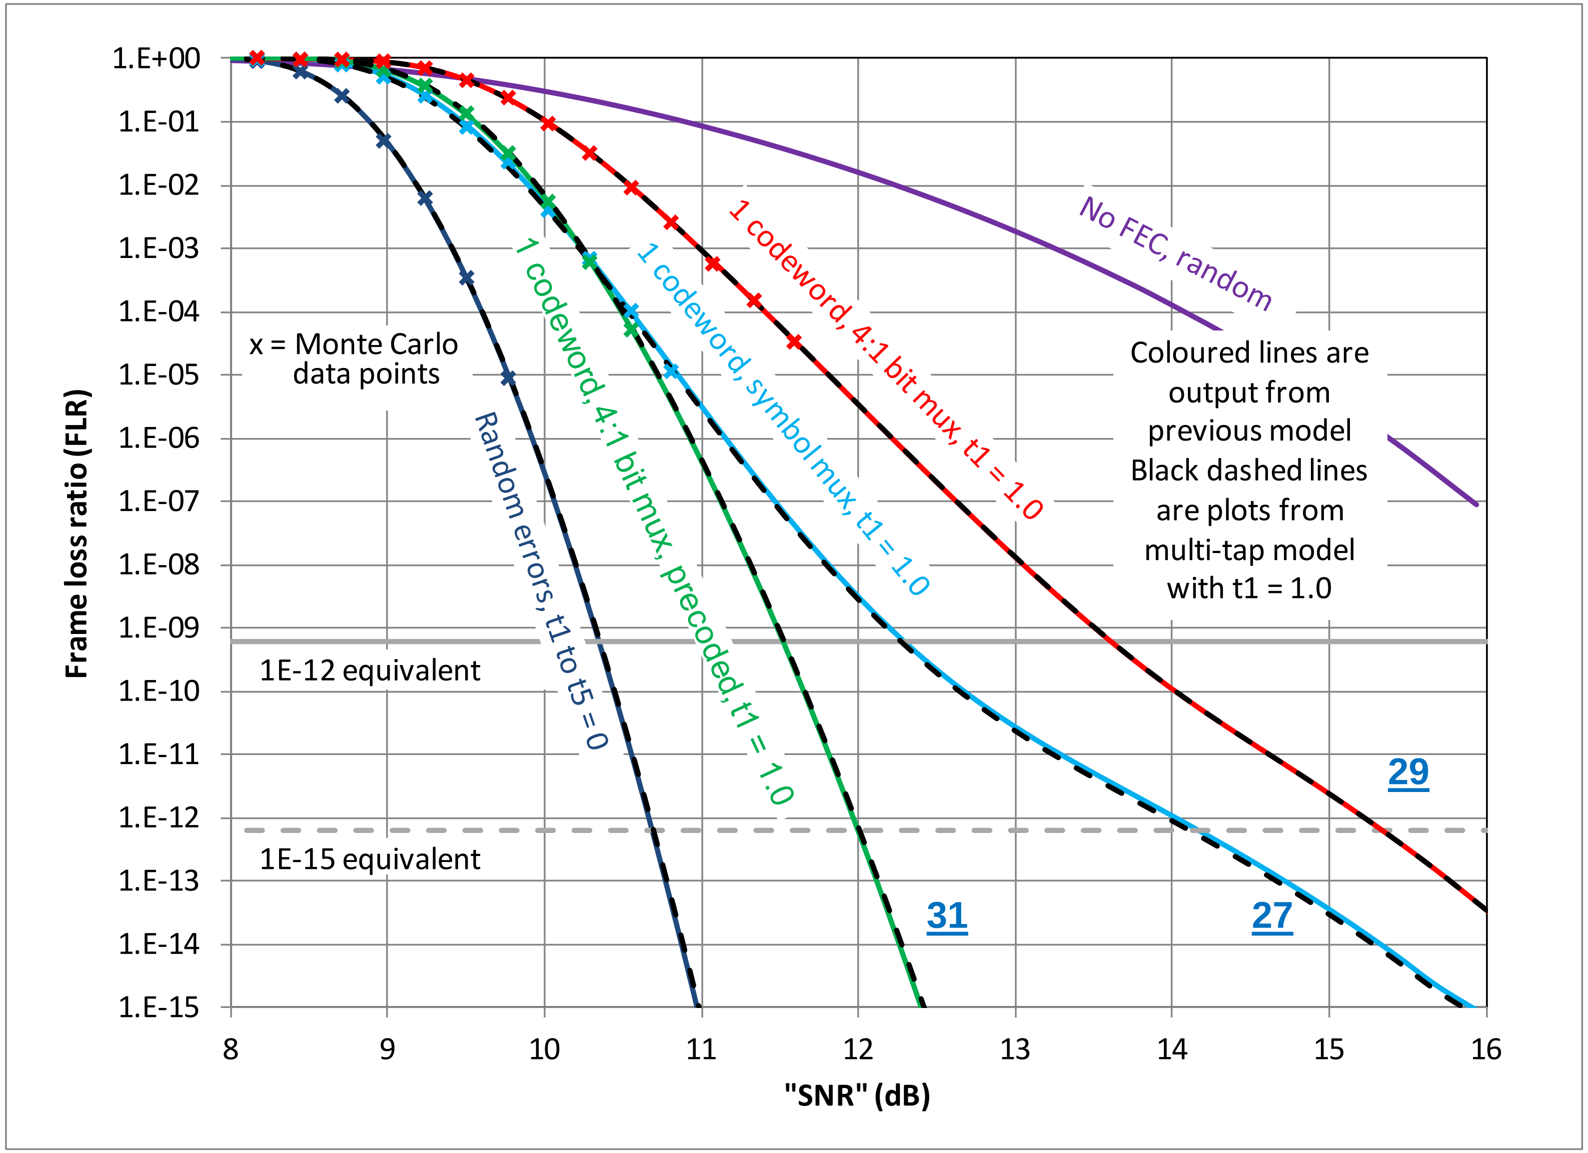

DFE tap weights

Rich Mellitz kindly provided the normalized DFE tap weights as calculated by COM for some channels submitted to IEEE 802.3. For details of the values see 32

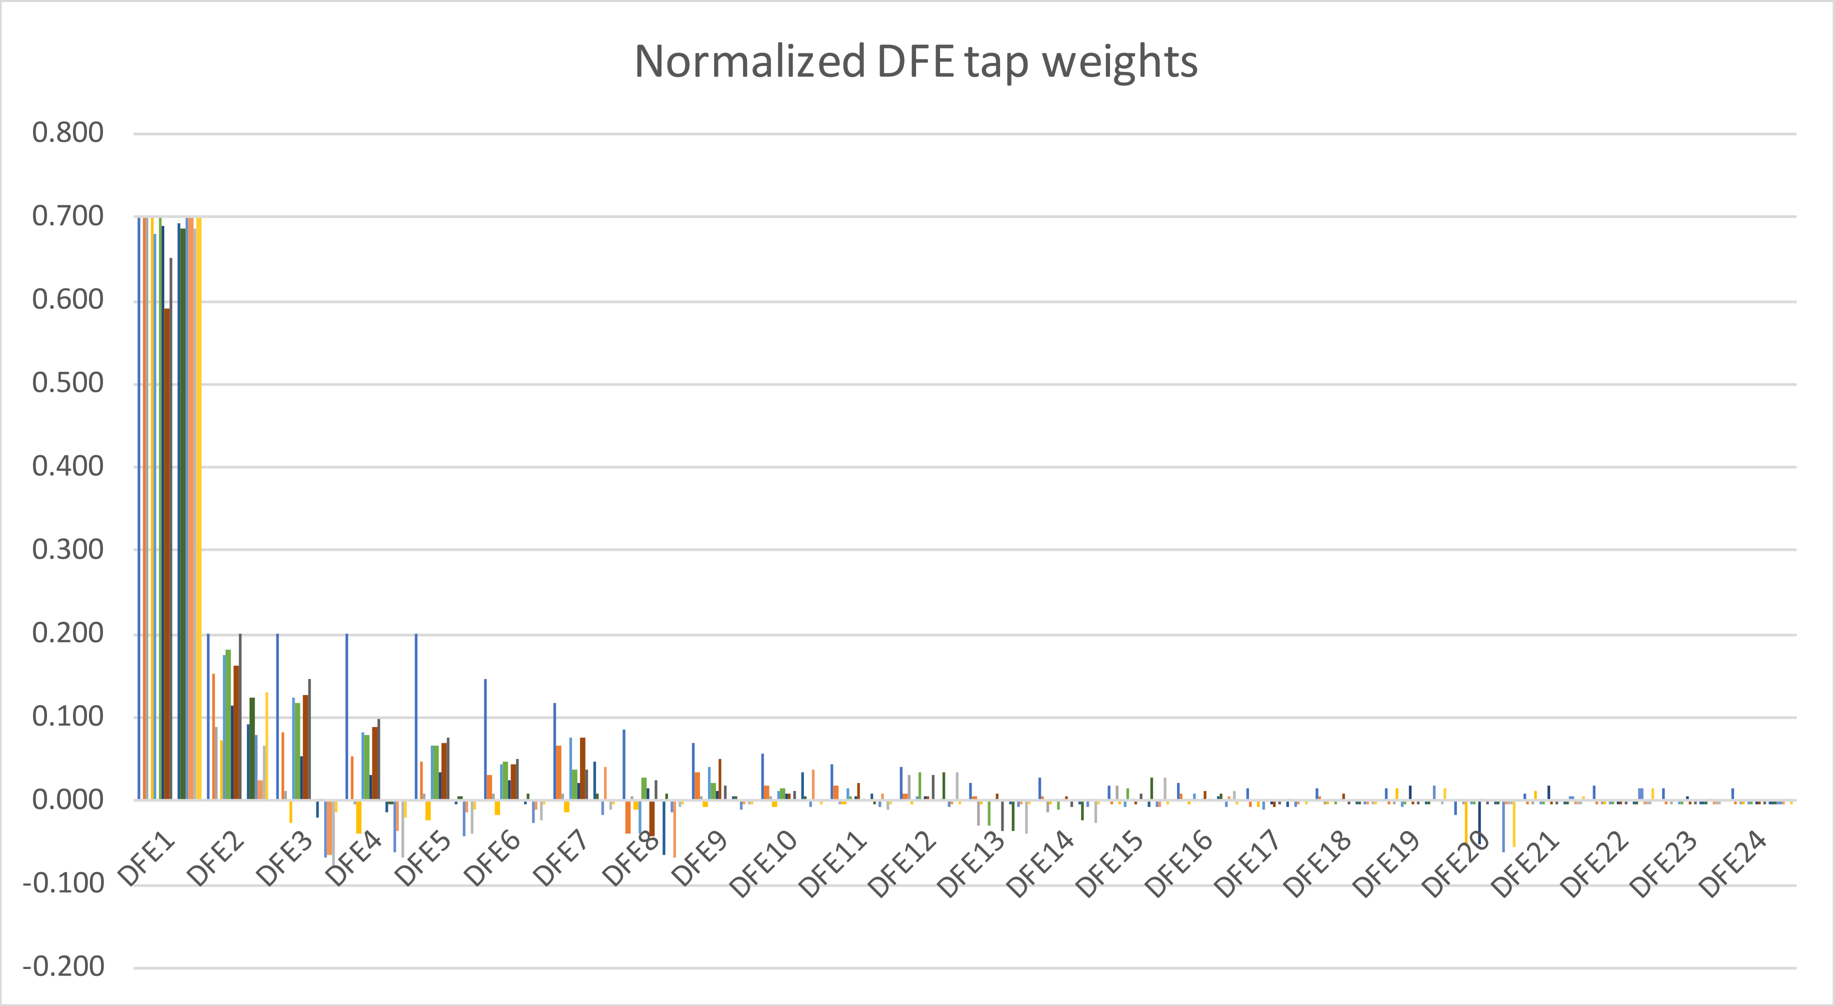

Multi-tap DFE

In the P802.3cd draft, Table 136-18 for 50GBASE-CR, 100GBASE-CR2, and 200GBASE-CR4 and Table 137-6 for 50GBASE-KR, 100GBASE-KR2, and 200GBASE-KR4 have a normalized DFE coefficient magnitude limit of 0.7 for tap 1 and 0.2 for all other taps.

This, together with the tap information shown on the previous slide (which has the same limits) led to the selection of:

t1 = 0.7, t2 = 0, t3 = 0.2, t4 = 0, t5 = 0.2

t1 = 0.7, t2 = –0.1, t3 = 0.1, t4 = –0.1, t5 = 0.1

as “realistic” but possibly near worst case candidate sets of tap weights.

The Monte Carlo values (x markers) and solid curves generated from the error statistics for these sets of tap weights with precoding turned on for both bit and FEC symbol interleaving are plotted on the next two slides together with the 1-tap DFE with precoding result shown earlier.

As a check, slide 10 shows the results for t1 = 0.7, t2 = –0.1, t3 = 0.1, t4 = –0.1, t5 = 0.1 with precoding turned off.

100G 5-tap DFE results (0.7, 0, 0.2, 0, 0.2) with precoding

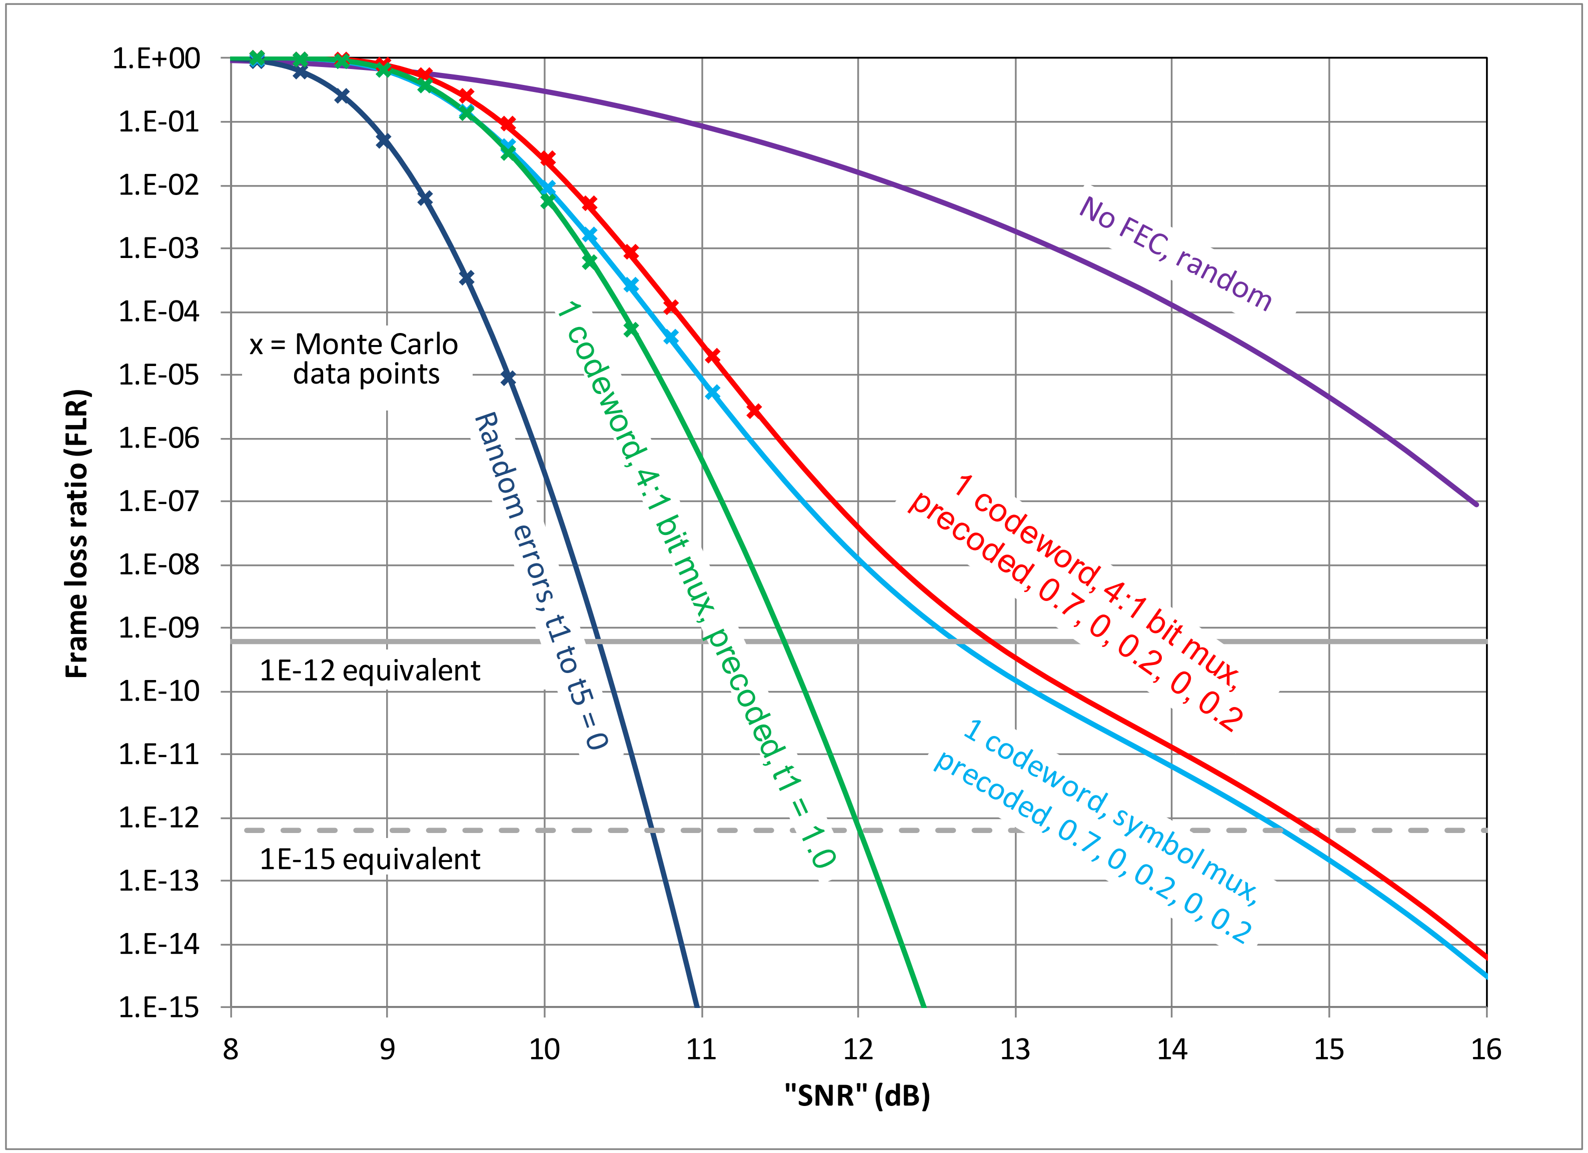

100G 5-tap DFE results (0.7, -0.1, 0.1, -0.1, 0.1) with precoding

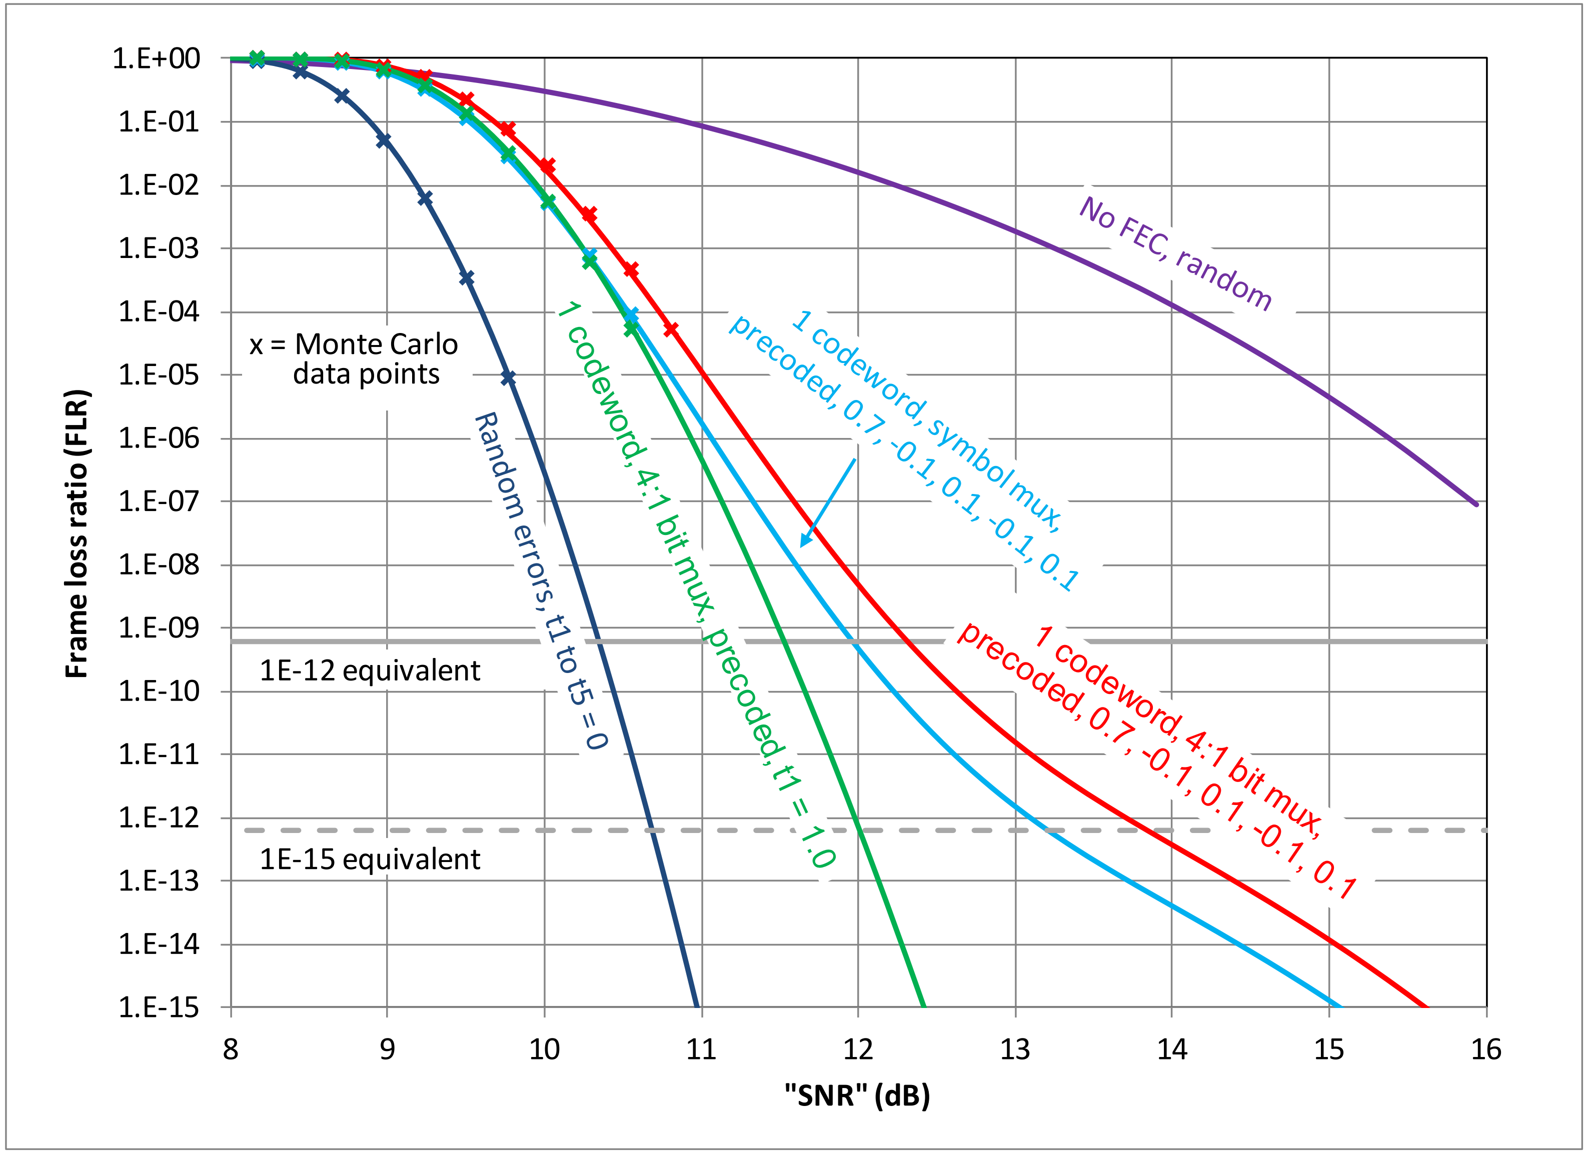

100G 5-tap DFE results (0.7, -0.1, 0.1, -0.1, 0.1) no precoding

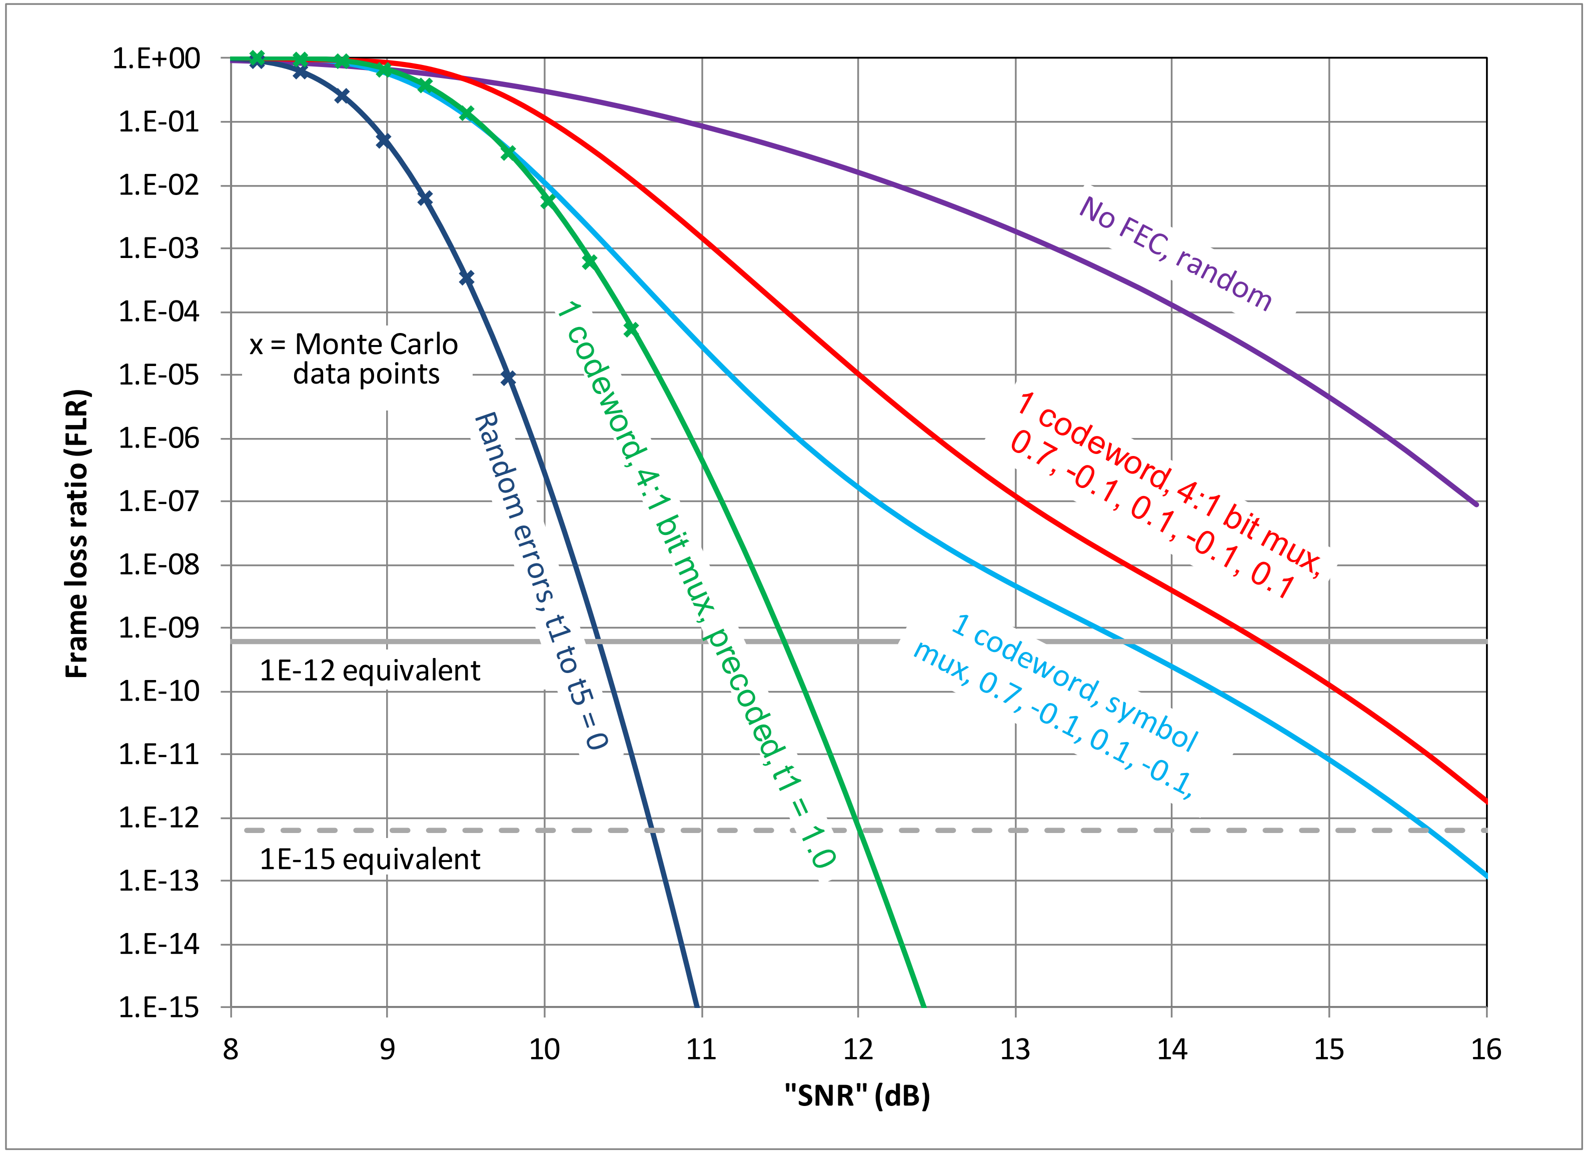

400G Validation

*

The new Monte Carlo model was then modified to re-create the previously simulated (1-tap DFE) results for 2 codeword interleaving as per Clause 119.

For the 4:1 bit mux and symbol mux results, there is good agreement with the previous model and the results are shown using the same method as before on slide 12.

For the 4:1 bit mux with precoding result the new model produces results that are significantly different from the old model as shown on slide 13. This is due to an error in the old model in the equation used to relate SNR and BER.

This change brings the curve for 2 codeword 4:1 bit mux with precoding in this presentation in line with that from Adam Healey and Cathy Liu as shown on slide 41.

400G previous result validation (1-tap DFE) part 1

*

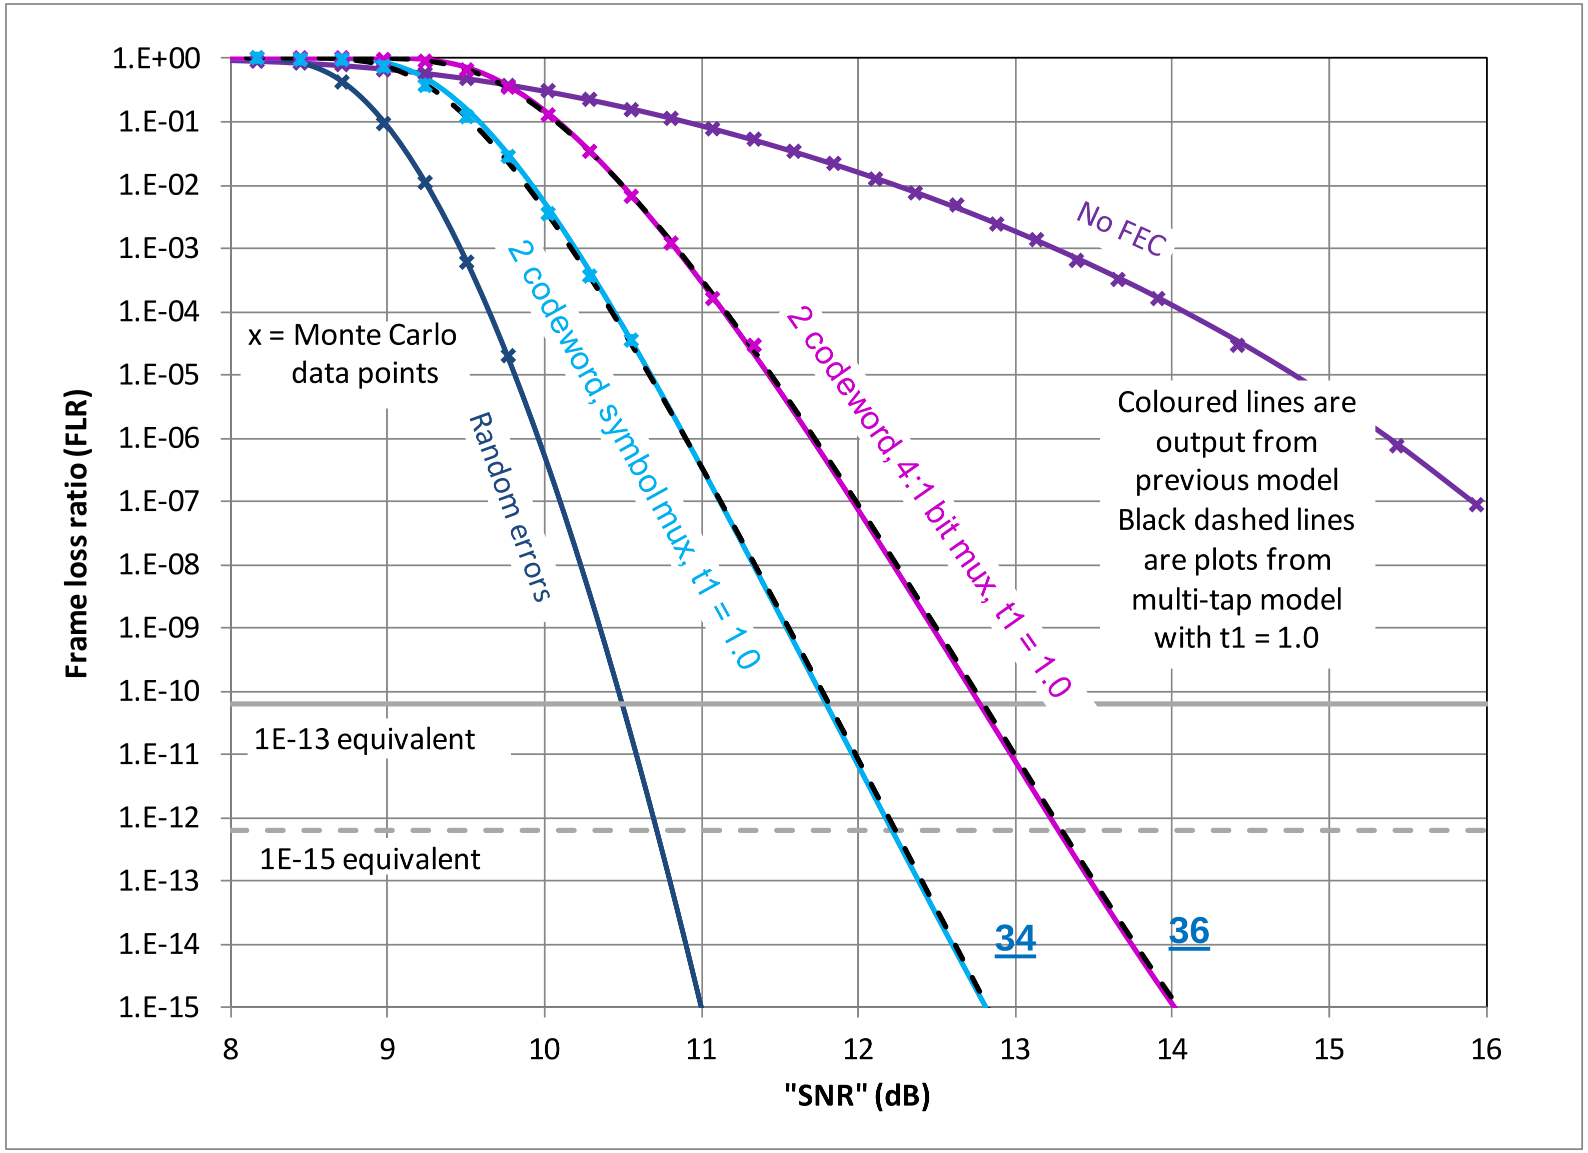

400G previous result validation (1-tap DFE) part 2

*

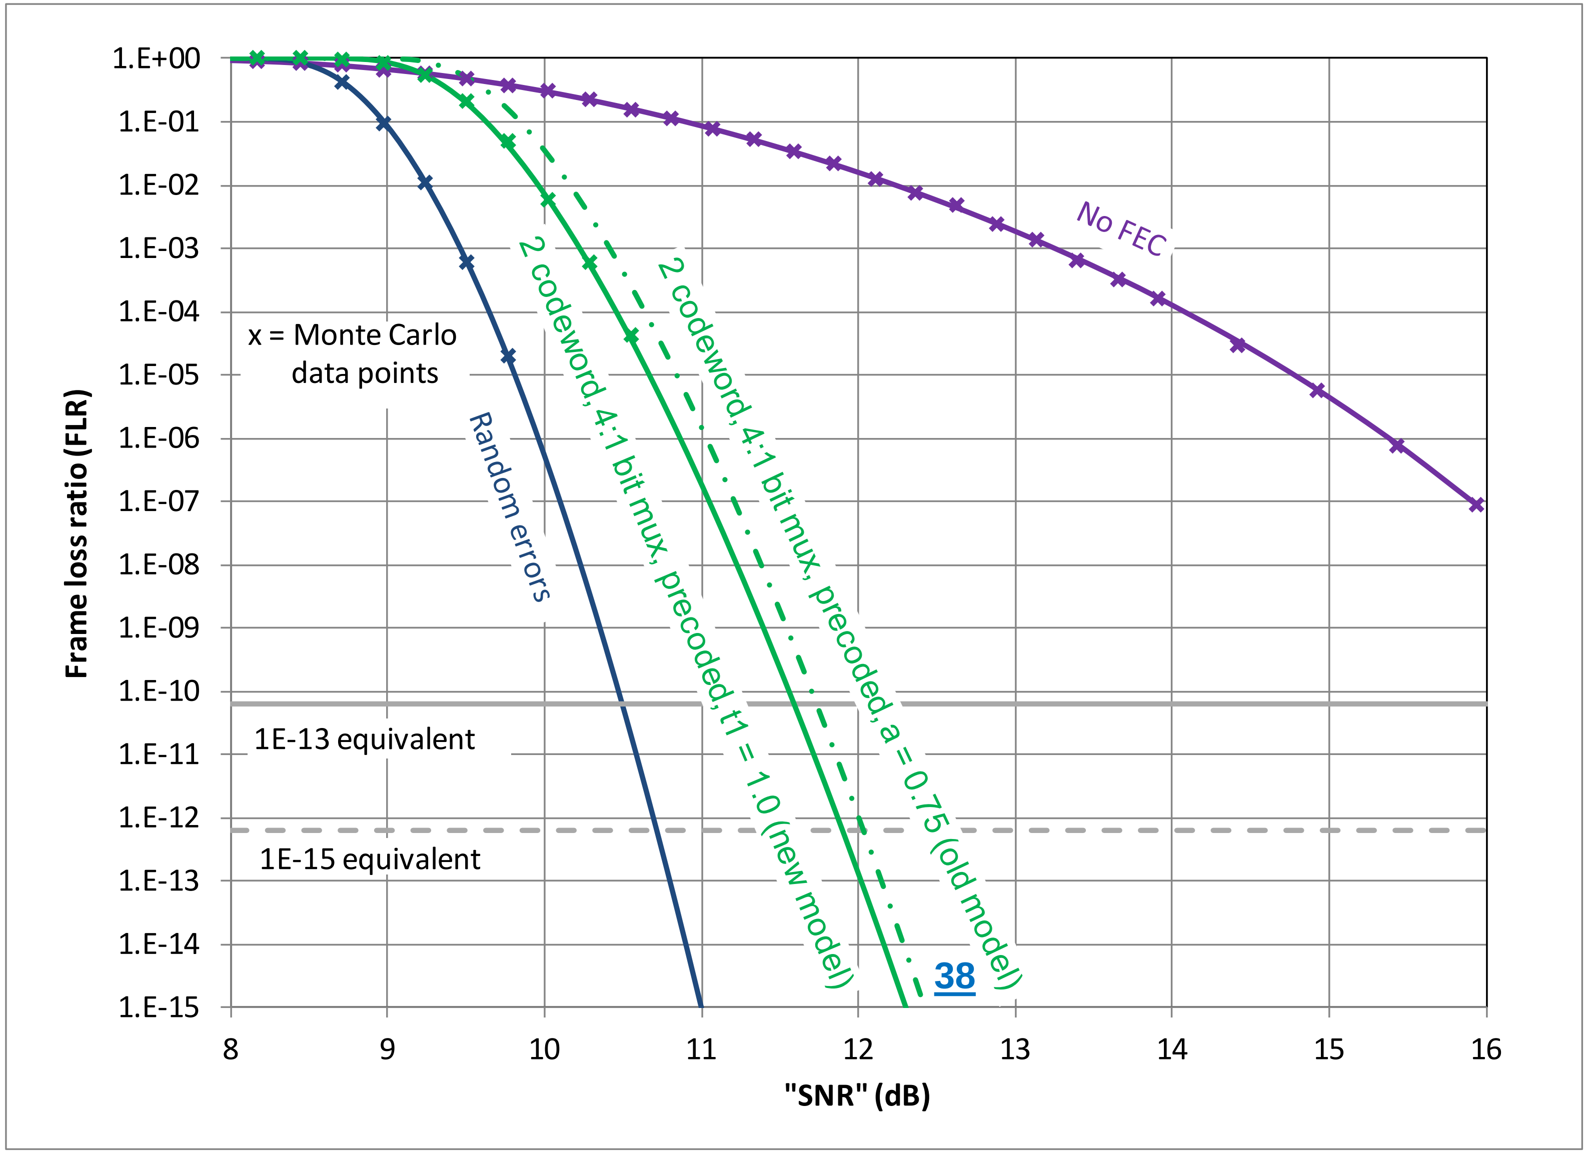

Multi-tap DFE 400G

*

Since the two codeword interleaving as per Clause 119 has better performance than the 100G case, only the more severe set of taps as for the 100G case was simulated:

t1 = 0.7, t2 = 0, t3 = 0.2, t4 = 0, t5 = 0.2

The Monte Carlo values (x markers) and solid curves generated from the error statistics for these sets of tap weights with precoding turned on for both bit interleaving (red curve) and FEC symbol interleaving (grey curve) are plotted on the next slide together with the 1-tap DFE with precoding result shown earlier.

As a check, slide 16 shows the results for t1 = 0.7, t2 = 0, t3 = 0.2, t4 = 0, t5 = 0.2 with precoding turned off.

400G 5-tap DFE results (0.7, 0, 0.2, 0, 0.2) with precoding

*

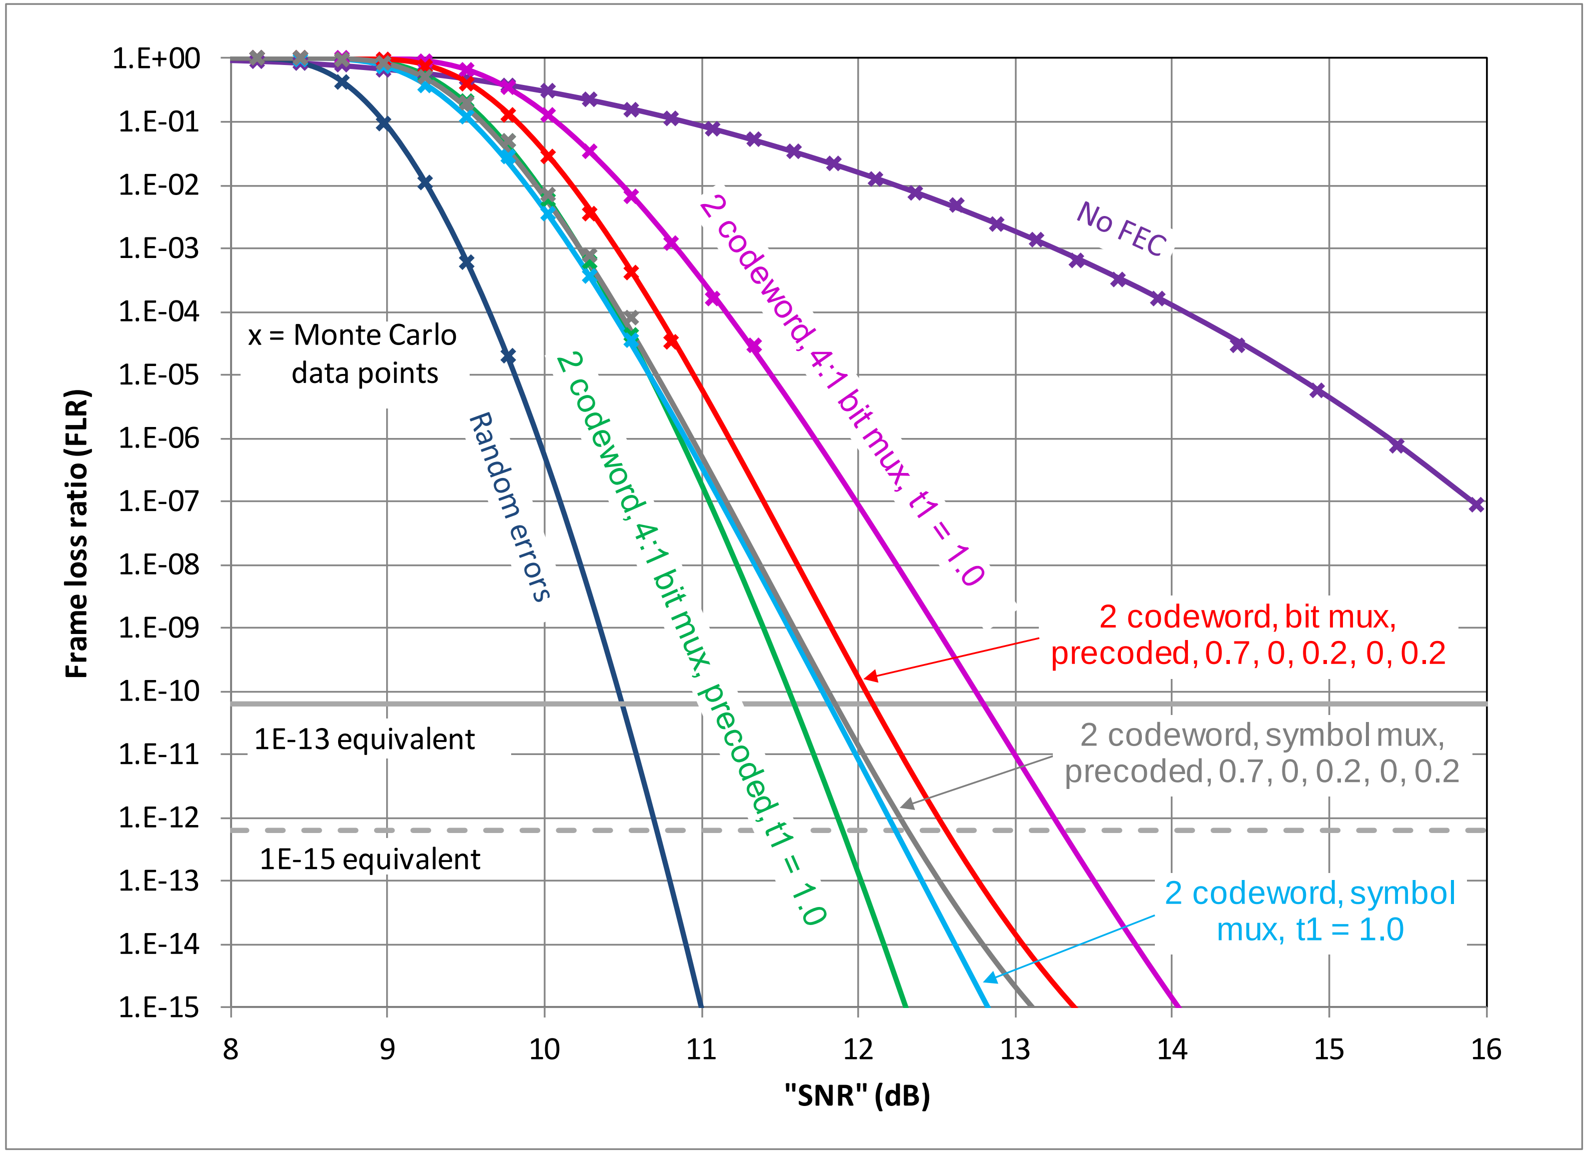

400G 5-tap DFE results (0.7, 0, 0.2, 0, 0.2) no precoding

*

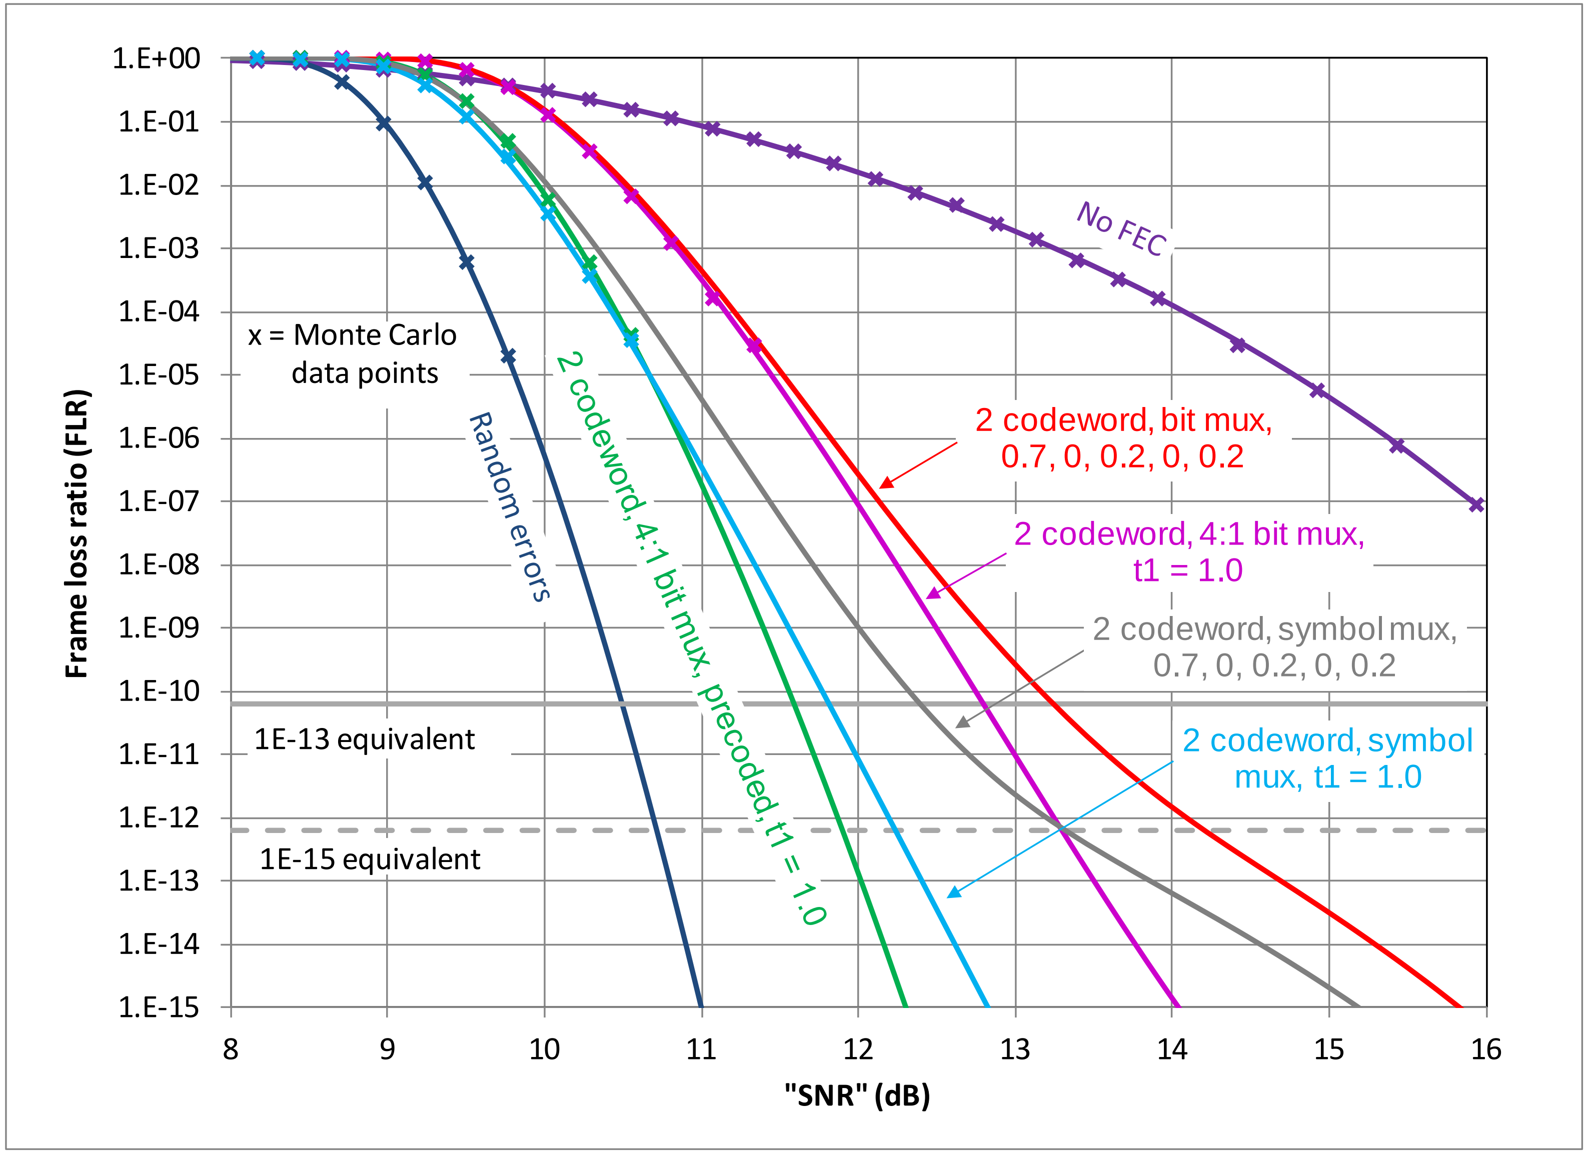

Multi-tap DFE burst error characteristics

*

The graph on the right shows the probability of occurrence vs burst length for the 1-tap DFE (t1 = 1.0) and the two multi-tap DFEs simulated. For the multi-tap case a burst is defined as a sequence of symbols where the first and last symbol are errored and the burst contains no sequence of correct symbols of length equal to the number of taps.

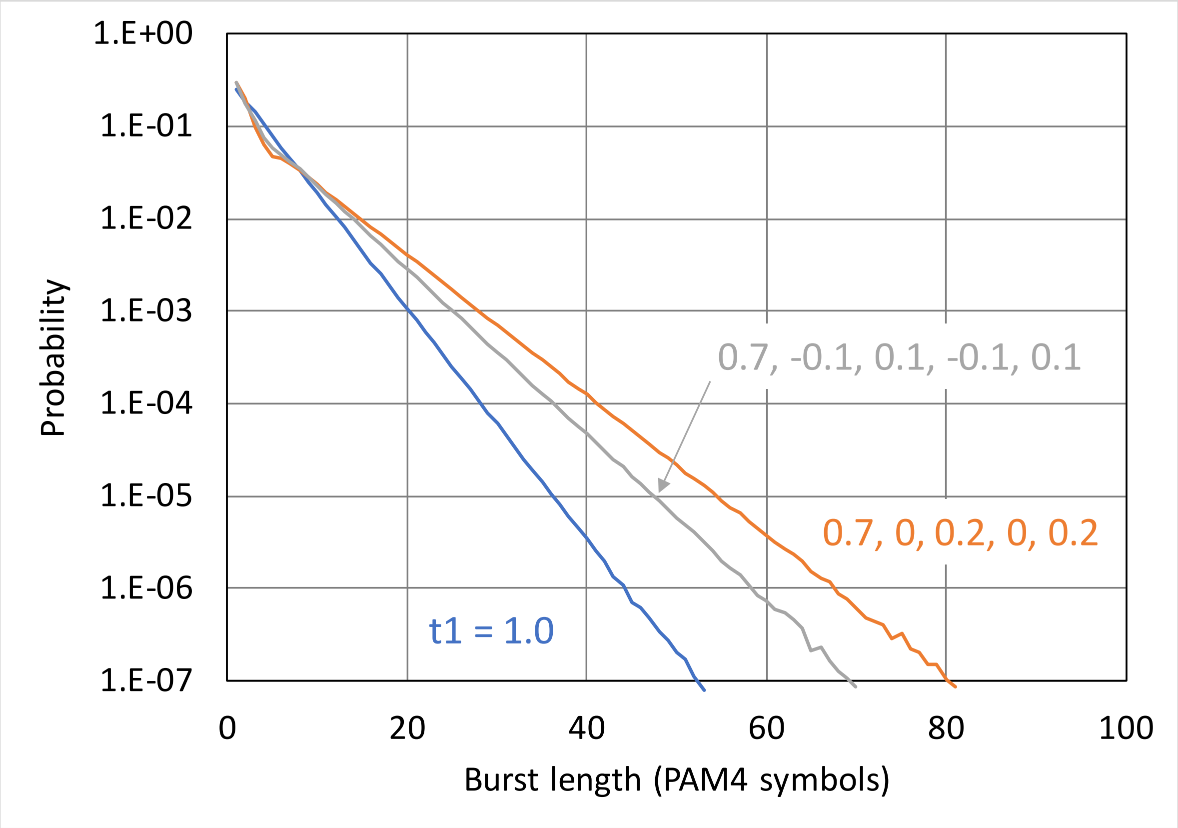

The average number of errors in a burst for the three cases are:

t1 = 1.0	4.0
0.7, -0.1, 0.1, -0.1, 0.1	4.25
0.7, 0, 0.2, 0, 0.2	4.42

For the 1-tap DFE case the bursts are of the form:
-1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1

An example burst for the 0.7, -0.1, 0.1, -0.1, 0.1 case is:
-1 1 0 0 -1 1 -1 0 0 0 -1 1 -1 1 -1

*

Results for RS(544,514) 100G all gain used for PAM4 part 1

From the curves shown on the previous pages, if all of the coding gain were to be used for the PAM4 link, the BERs required to give FLRs equivalent to that of a BER of 1E-12 (for 100G) or 1E-13 (for 400G) and 1E-15 are:

	At slicer output		At FEC input	
100G	FLR = 6.2E-10	FLR = 6.2E-13	FLR = 6.2E-10	FLR = 6.2E-13
No FEC	1E-12	1E-15		
1 codeword, 4:1 bit mux, t1 = 1.0	2.5E-6*	7.6E-9*		
1 codeword, symbol mux, t1 = 1.0	5.8E-5*	3.8E-7*		
1 codeword, 4:1 bit mux, precoded, t1 = 1.0	2.5E-4*	1.0E-4*	1.2E-4	5.1E-5
1 codeword, symbol mux, precoded, 0.7, 0, 0.2, 0, 0.2	3.2E-5*	9.0E-8*	1.7E-5	4.9E-8
Random errors	3.8E-4	2.3E-4		
400G	FLR = 6.2E-11	FLR = 6.2E-13	FLR = 6.2E-11	FLR = 6.2E-13
No FEC	1E-13	1E-15		
2 codeword, 4:1 bit mux, t1 = 1.0	2.3E-5*	6.9E-6*		
2 codeword, symbol mux, t1 = 1.0	1.7E-4*	7.3E-5*		
2 codeword, 4:1 bit mux, precoded, t1 = 1.0	2.4E-4*	1.4E-4*	1.2E-4	6.8E-5
2 codeword, symbol mux, precoded, 0.7, 0, 0.2, 0, 0.2	1.7E-4*	7.0E-5*	8.5E-5	3.5E-5
Random errors	3.2E-4	2.3E-4		

Note – these values are the BER **including** the additional errors due to the bursts. To account for burst errors, the values marked with “*” have been multiplied by 4 for a 1-tap DFE and 4.42 for the 0.7, 0, 0.2, 0, 0.2 case.

*

Results for RS(544,514) 100G all gain used for PAM4 part 2

From the curves shown on the previous pages, if all of the coding gain were to be used for the PAM4 link, the SER_{in} and SNR required to give FLRs equivalent to that of a BER of 1E-12 (for 100G) or 1E-13 (for 400G) and 1E-15 are:

100G	For FLR = 6.2E-10			For FLR = 6.2E-13		
	SER_{in}	SNR (dB)	SNR + 6.99	SER_{in}	SNR (dB)	SNR + 6.99
1 codeword, 4:1 bit mux, t1 = 1.0	1.3E-6	13.6	20.6	3.8E-9	15.3	22.3
1 codeword, symbol mux, t1 = 1.0	2.9E-5	12.3	19.3	1.9E-7	14.2	21.2
1 codeword, 4:1 bit mux, precoded, t1 = 1.0	1.2E-4	11.5	18.5	5.1E-5	12.0	19.0
1 codeword, symbol mux, precoded, t1 = 0.7, t3 & t5 = 0.2	1.7E-5	12.6	19.6	4.9E-8	14.7	21.7

400G	For FLR = 6.2E-11			For FLR = 6.2E-13		
	SER_{in}	SNR (dB)	SNR + 6.99	SER_{in}	SNR (dB)	SNR + 6.99
2 codeword, 4:1 bit mux, t1 = 1.0	1.1E-5	12.7	19.7	3.5E-6	13.2	20.2
2 codeword, symbol mux, t1 = 1.0	8.3E-5	11.7	18.7	3.7E-5	12.2	19.2
2 codeword, 4:1 bit mux, precoded, t1 = 1.0	1.2E-4	11.5	18.5	6.8E-5	11.9	18.8
2 codeword, symbol mux, precoded, t1 = 0.7, t3 & t5 = 0.2	9.2E-5	11.8	18.8	3.8E-5	12.2	19.2

Where:

SER_{in} is the symbol error ratio due to noise only (does not include bursts)
SNR (dB) is the "SNR" in equation (1) on page 25
SNR + 6.99 is the SNR as defined on page 5 of healey_100GEL_01_0318

Conclusion

*

For 400G using a Clause 119 PCS:
The performance with a multi-tap DFE (0.7, 0, 0.2, 0, 0.2) at 1E-13 equivalent BER with precoding and worst case RS symbol multiplexing is 0.3 dB worse than the performance of precoding and bit multiplexing with a 1-tap DFE.
At 1E-15 equivalent BER the difference is also about 0.3 dB.

Bit multiplexing rather than symbol multiplexing is about 0.2 dB worse.

For 100G using a Clause 91 FEC sublayer:
The performance with a multi-tap DFE (0.7, 0, 0.2, 0, 0.2) at 1E-12 equivalent BER with precoding and worst case RS symbol multiplexing is 1.1 dB worse than the performance of precoding and bit multiplexing with a 1-tap DFE.
At 1E-15 equivalent BER the difference is about 2.7 dB.

Bit multiplexing rather than symbol multiplexing is about 0.2 dB worse.

Annex

Gray coding

Assume the use of Gray coding (see IEEE Std 802.3-2018 120.5.7) as illustrated below:

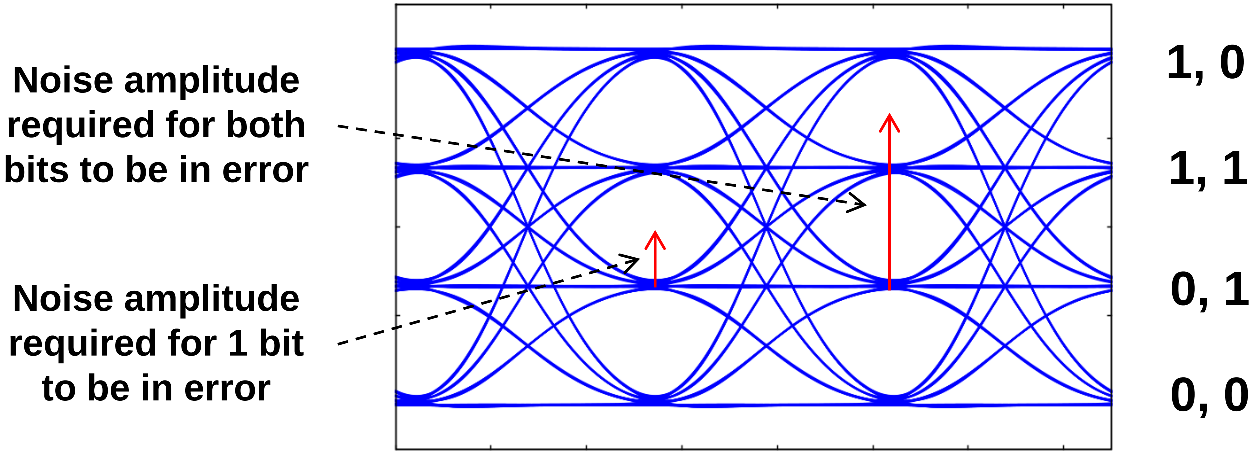

If noise causes any of the 4 levels to be mistaken for an adjacent level, this causes one of the two bits to be in error.

If there is just enough Gaussian noise to cause a BER of 3.8E-4* due to single level errors, then the probability of that noise causing both bits to be in error is 2.8E-23.

This analysis therefore assumes that only one of the two bits is in error.

* FLR = 6.2E-10 (equivalent to BER = 1E-12 with random errors) after RS(544,514) FEC

Burst error model 1

The NRZ burst analysis in anslow_3bs_02_1114 page 12 assumed that if a **bit** is in error, the worst case probability that the next **bit** is also in error is 0.5. If we assume, for Gray coded PAM4, that an error in a particular symbol only causes the decision on the next symbol to move up or down one level, then the possibilities are:

Correct level	Received level		Error pattern	
	One up	One down	One up	One down
3	3	2	✓, ✓	✓, ✗
2	3	1	✓, ✗	✗, ✓
1	2	0	✗, ✓	✓, ✗
0	1	0	✓, ✗	✓, ✓

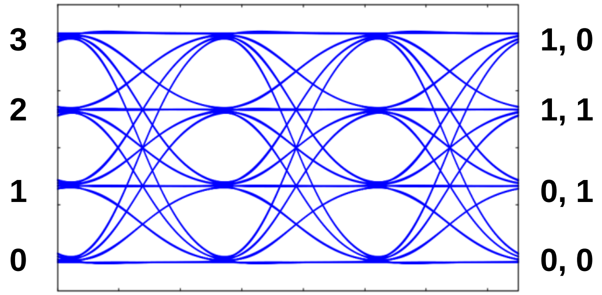

Since two of the eight possibilities result in both bits being correct, these states terminate the burst. Therefore for Gray coded PAM4, if a **symbol** is in error, the worst case probability that the next **symbol** is also in error is 0.75.

Burst error model 2

The second aspect of this table is that of the six possibilities giving bits in error, two have errors in the first bit while four have errors in the second bit.

Correct level	Received level		Error pattern	
	One up	One down	One up	One down
3	3	2	✓, ✓	✓, ✗
2	3	1	✓, ✗	✗, ✓
1	2	0	✗, ✓	✓, ✗
0	1	0	✓, ✗	✓, ✓

The analysis in the remainder of this contribution therefore assumes that if a given symbol is in error, the probability of a bit error in the first bit is 1/3 and in the second bit is 2/3.

Burst error model 3

The "SNR" shown on the X axis of the results slides is related to the noise induced input SER via the following equation:

$$SER_{in} = \frac{3}{4} erfc\left(\sqrt{\frac{SNR}{2}}\right) \qquad (1)$$

Which does not include the additional errors due to the bursts. The average number of errors in a burst is related to the probability of the burst continuing "a" as shown below:

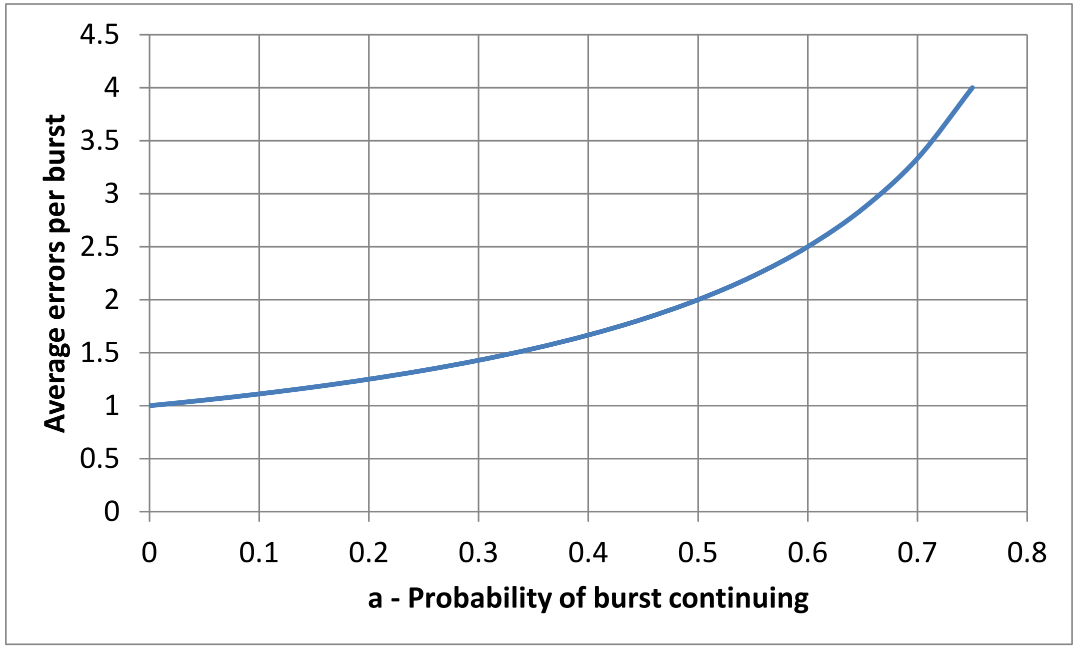

For a = 0.75, the BER_{in} including bursts is 4 x the BER_{in} due to noise.

Single burst bound

As pointed out in anslow_01_0815_logic, for a non-interleaved scheme, a single burst that lasts for ~74 PAM4 symbols has a high probability of causing errors in 16 FEC symbols (which is uncorrectable). With a = 0.75, the probability of a burst this long is 0.75^74 = 5.7E-10. When this is combined with the probability that the codeword has at least one error in it, a simple lower bound for the FLR can be calculated.

If a is the probability of the burst continuing, a more accurate calculation for the probability that a single burst is uncorrectable is:

$$P_{uncorr} = 1/5*a^{71}*(1-a) + 2/5*a^{72}*(1-a) + 3/5*a^{73}*(1-a) + 4/5*a^{74}*(1-a) + a^{75}*(1-a) + a^{76}*(1-a) + a^{77}*(1-a) + \ldots$$

For a = 0.75, this evaluates to 8.2E-10.

This bound is plotted as a dotted line on page 28.

Clause 91 100G with symbol mux PMA

Round robin distribution of FEC symbols to the FEC lanes. Symbol multiplex in the PMA.

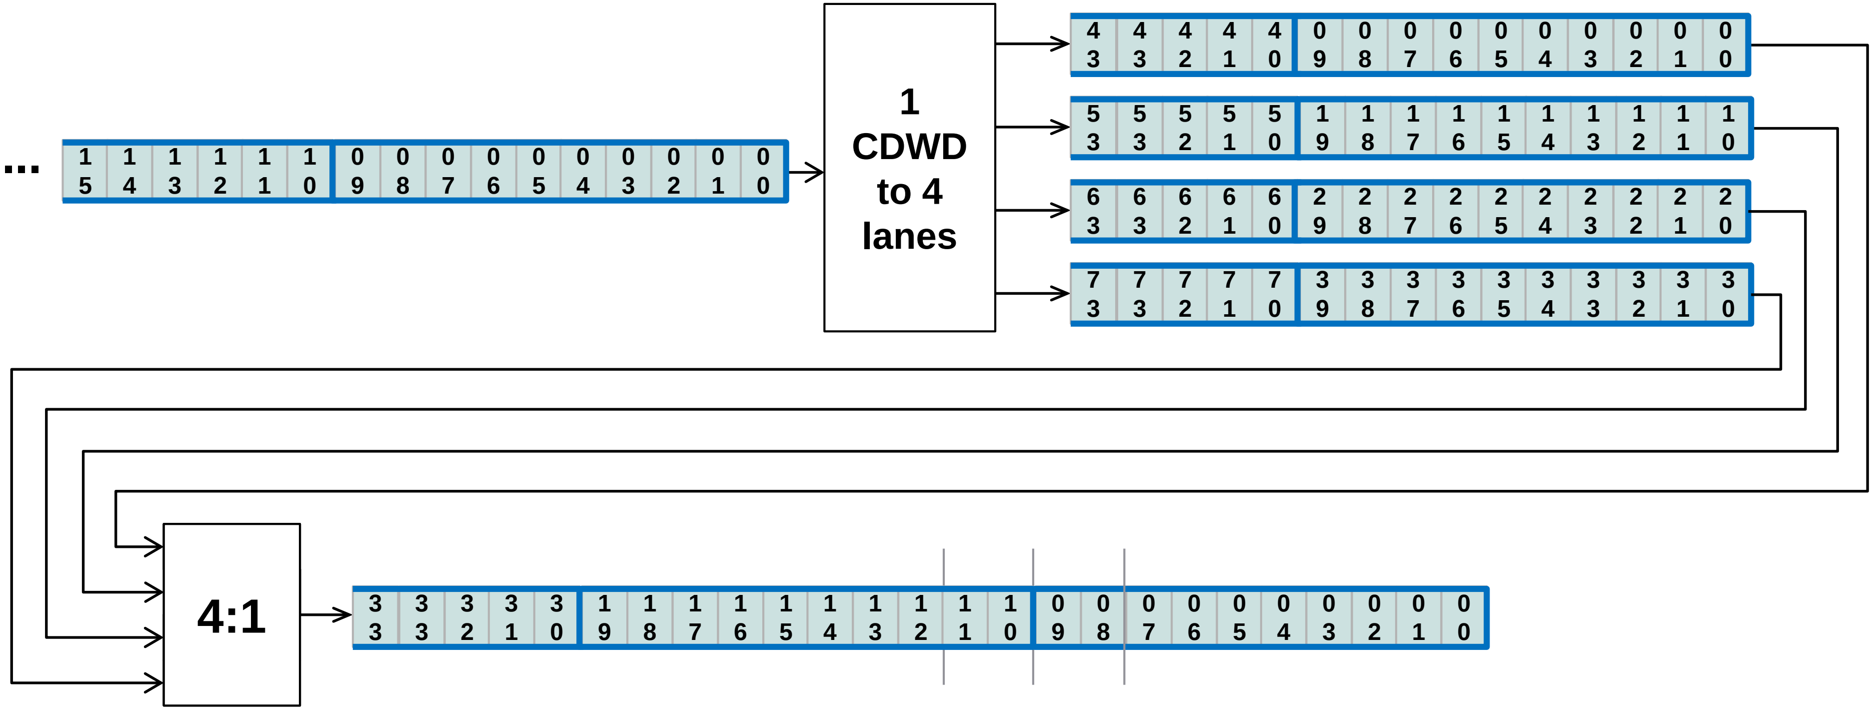

PMA must find FEC symbol boundaries.

Back

Clause 91 100G with symbol mux PMA

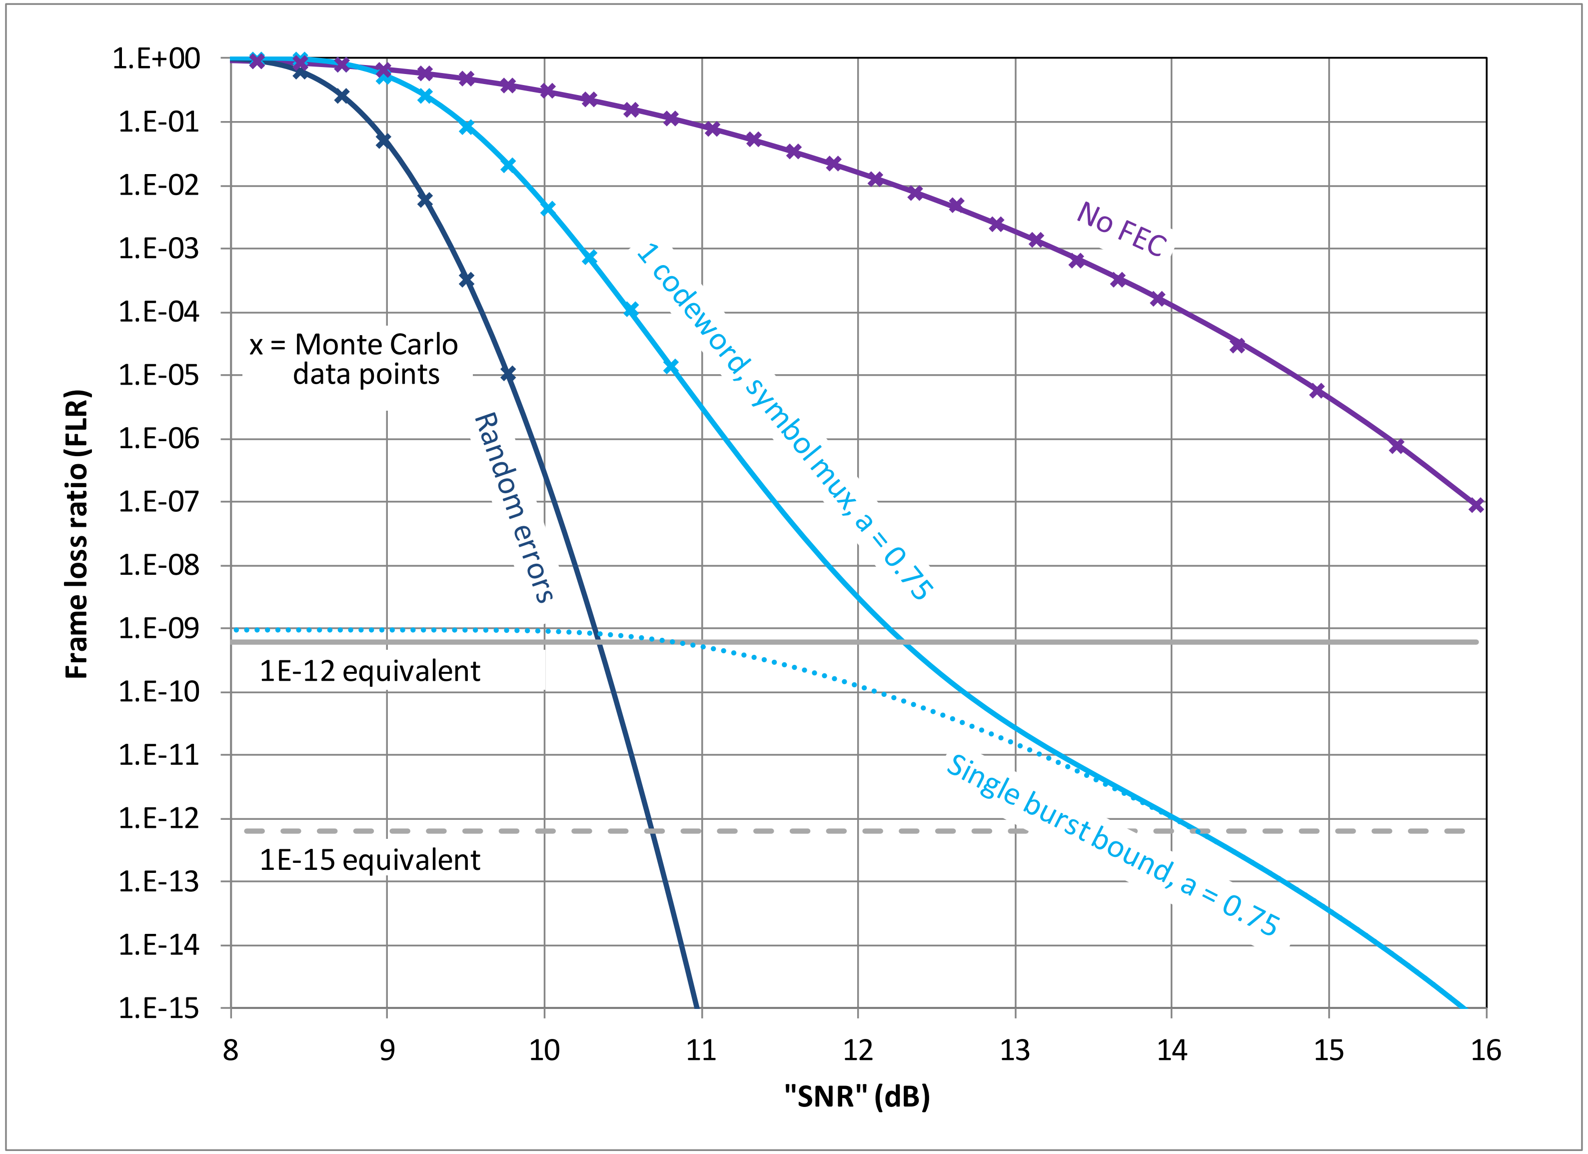

Back

Clause 91 100G with bit mux PMA

Round robin distribution of FEC symbols to the FEC lanes. Bit multiplex in the PMA.

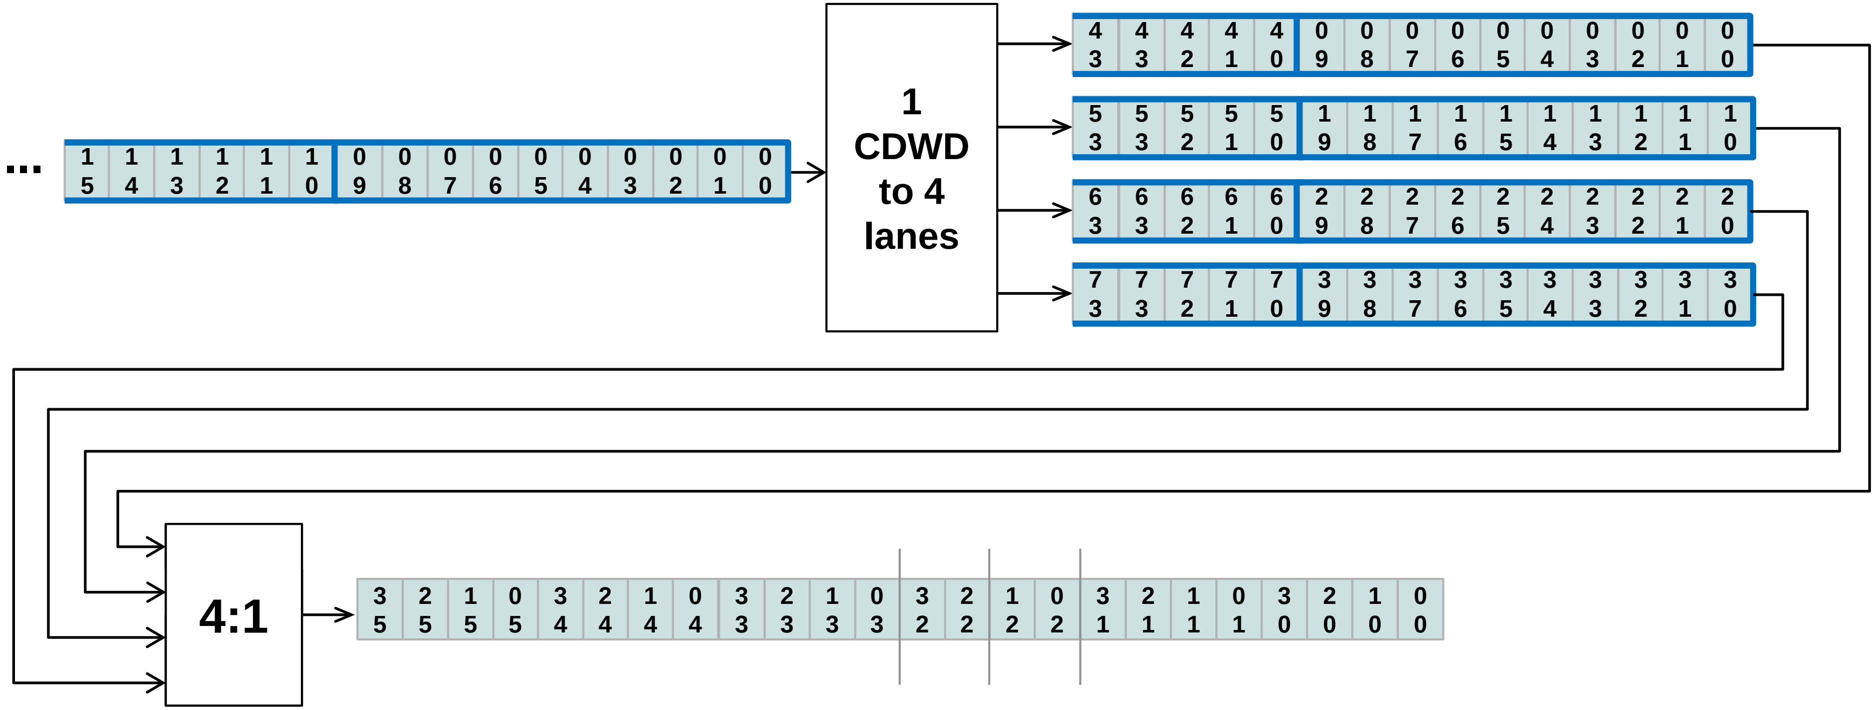

Clause 91 100G with bit mux PMA

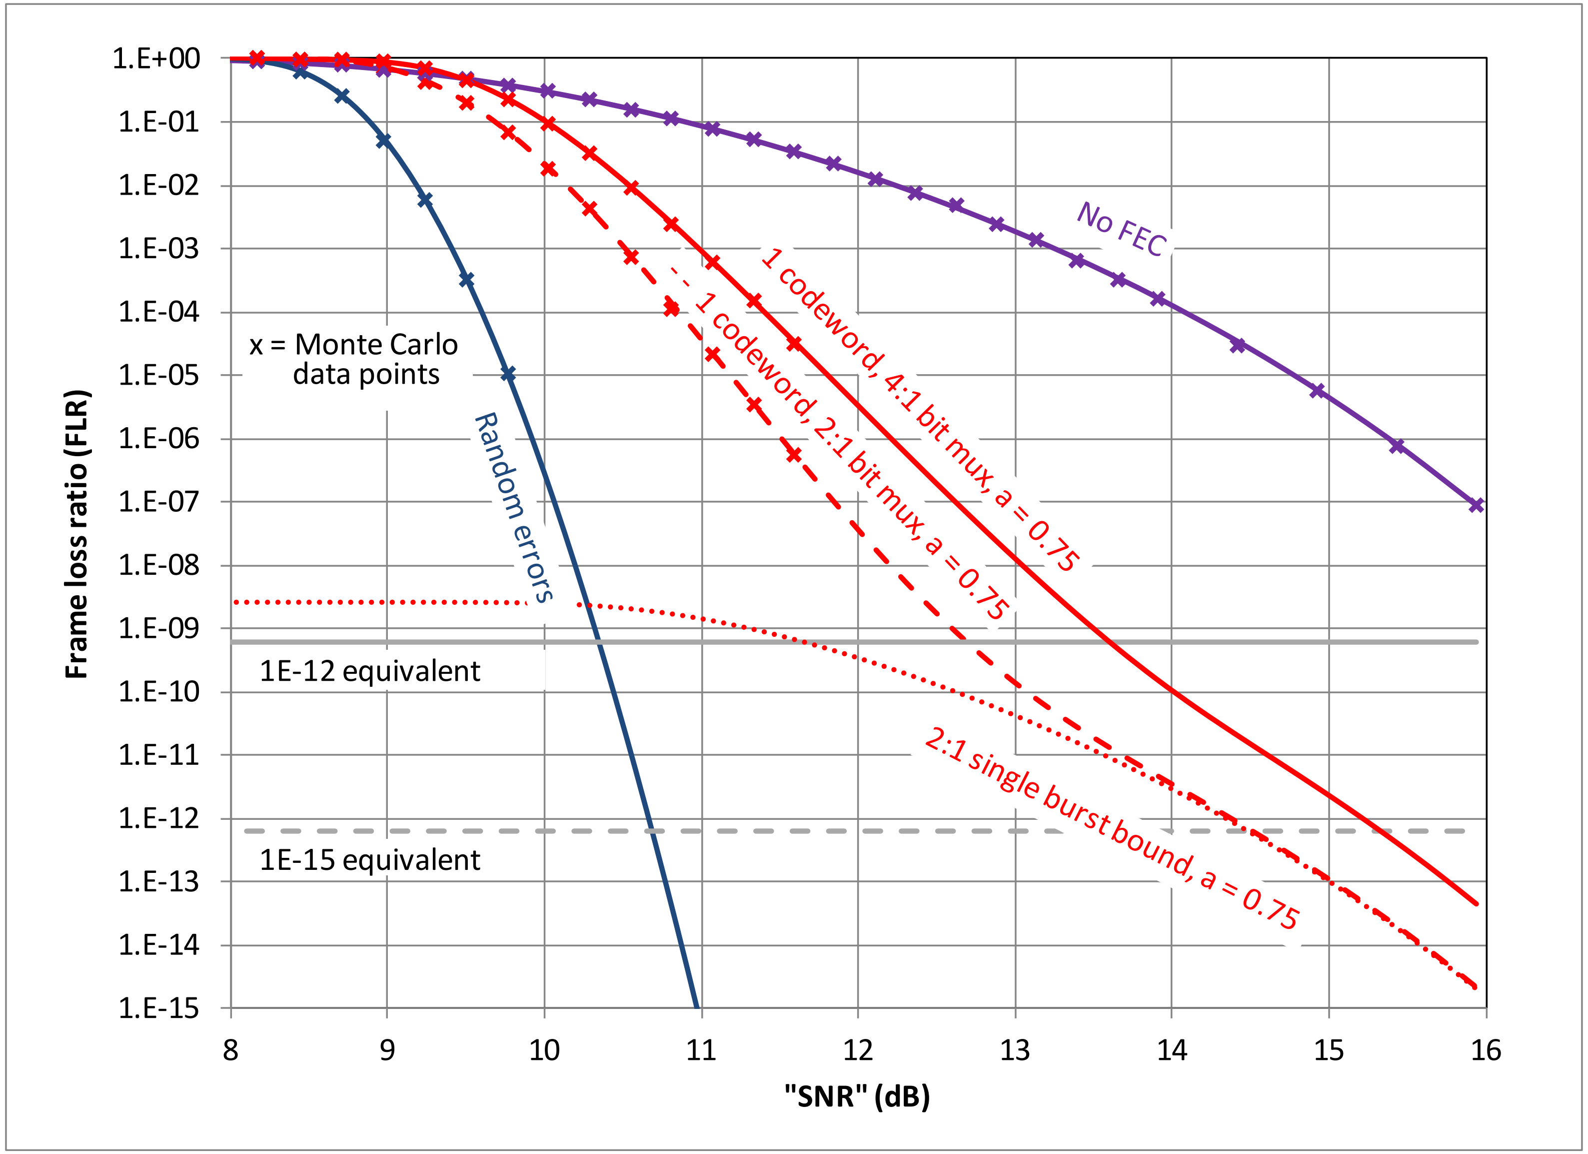

Back

Precoding

Precoding as defined in 802.3cd 120.5.7.2 was assumed. This is performed as illustrated below.

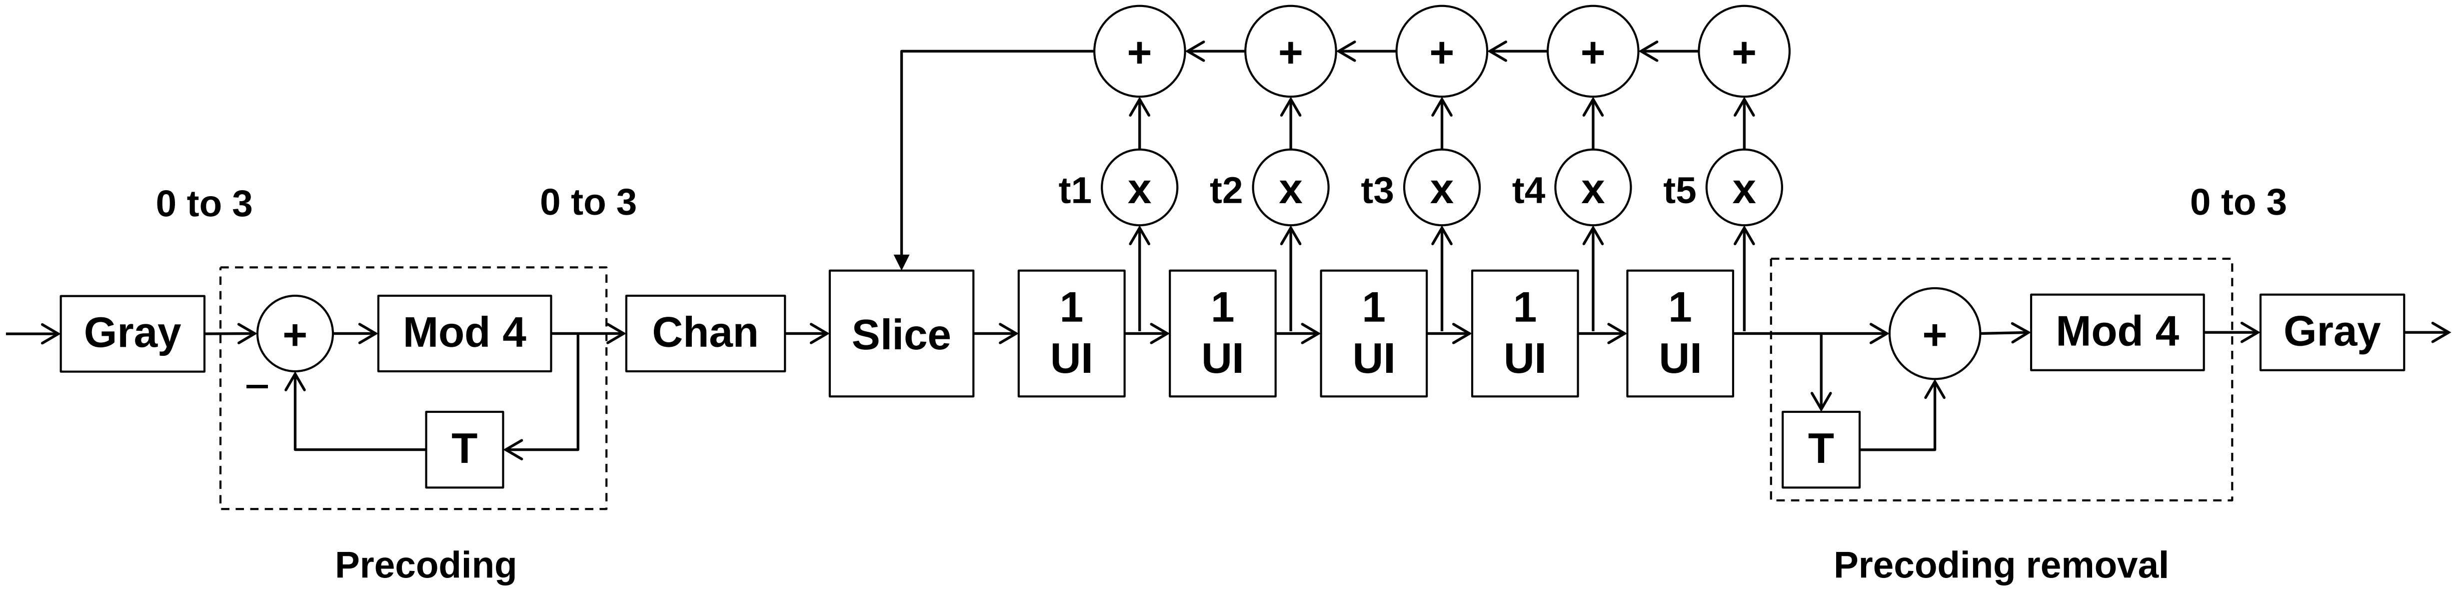

See page 5 of parthasarathy_01_0911 for a worked example.

A “feature” of this precoding process is that a single random errored PAM4 symbol at the slicer output turns into two errored PAM4 symbols after the precoding is removed.

Back

DFE tap weights

Rich Mellitz kindly provided the normalized DFE tap weights as calculated by COM for some channels submitted to IEEE 802.3.

channel	pkg len	Il (dB)	DFE1	DFE2	DFE3	DFE4	DFE5	DFE6	DFE7	DFE8	DFE9	DFE10	DFE11	DFE12	DFE13	DFE14	DFE15	DFE16	DFE17	DFE18	DFE19	DFE20	DFE21	DFE22	DFE23	DFE24
tracy_100GEL_05_0118\B56_Thru_CblBP	30 mm	23.641	0.700	0.200	0.200	0.200	0.200	0.147	0.116	0.086	0.071	0.056	0.044	0.042	0.023	0.027	0.019	0.022	0.014	0.017	0.017	-0.018	0.010	0.020	0.016	0.017
mellitz_3ck_02_081518_CBP\CaBP_BGAVia_Opt2_32dB	12 mm	30.146	0.700	0.154	0.083	0.055	0.048	0.033	0.066	-0.041	0.035	0.017	0.017	0.007	0.005	0.005	-0.004	0.009	-0.007	0.004	-0.005	0.000	-0.002	-0.002	-0.001	-0.003
mellitz_3ck_02_081518_CBP\CaBP_BGAVia_Opt2_32dB	20 mm	30.146	0.700	0.089	0.014	-0.001	0.009	0.009	0.008	0.007	0.007	0.006	-0.002	0.031	-0.029	-0.014	0.018	0.003	0.003	-0.001	-0.001	-0.001	-0.001	-0.002	-0.001	-0.001
mellitz_3ck_02_081518_CBP\CaBP_BGAVia_Opt2_32dB	30 mm	30.146	0.700	0.072	-0.027	-0.039	-0.023	-0.017	-0.012	-0.009	-0.006	-0.006	-0.005	-0.005	-0.005	-0.005	-0.005	-0.004	-0.007	-0.001	0.014	-0.051	0.013	-0.002	0.003	-0.003
mellitz_3ck_02_081518_CBP/CaBP_BGAVia_Opt2_28dB	12 mm	26.347	0.680	0.176	0.123	0.083	0.067	0.043	0.074	-0.040	0.041	0.013	0.016	0.005	0.004	0.003	-0.006	0.008	-0.010	0.003	-0.007	-0.001	-0.004	-0.004	-0.003	-0.005
mellitz_3ck_02_081518_CBP/CaBP_BGAVia_Opt2_28dB	20 mm	26.347	0.699	0.181	0.117	0.080	0.066	0.048	0.037	0.027	0.021	0.015	0.005	0.035	-0.028	-0.009	0.014	0.003	0.002	-0.002	-0.002	-0.002	-0.002	-0.003	-0.003	-0.002
mellitz_3ck_02_081518_CBP/CaBP_BGAVia_Opt2_28dB	30 mm	26.347	0.691	0.116	0.053	0.032	0.033	0.026	0.021	0.016	0.013	0.010	0.007	0.005	0.003	0.002	0.001	0.001	-0.002	0.002	0.018	-0.050	0.018	-0.002	0.005	-0.002
mellitz_3ck_02_081518_CBP\CaBP_BGAVia_Opt2_24dB	12 mm	22.569	0.590	0.161	0.127	0.087	0.069	0.043	0.077	-0.043	0.051	0.010	0.020	0.007	0.007	0.006	-0.003	0.012	-0.008	0.007	-0.005	0.002	-0.001	-0.001	0.000	-0.002
mellitz_3ck_02_081518_CBP\CaBP_BGAVia_Opt2_24dB	20 mm	22.569	0.650	0.200	0.145	0.097	0.075	0.051	0.037	0.025	0.018	0.011	0.001	0.033	-0.035	-0.007	0.008	0.000	-0.001	-0.004	-0.004	-0.004	-0.004	-0.005	-0.004	-0.004
mellit...072518/Z0d_100_206in_0p13dBpi_twinax26_smooth	12 mm	27.980	0.692	0.091	-0.020	-0.013	-0.001	-0.003	0.047	-0.066	0.005	0.034	0.009	0.000	-0.002	0.000	-0.008	0.005	-0.006	-0.005	-0.002	-0.003	-0.003	-0.003	-0.002	-0.004
mellit...072518/Z0d_100_206in_0p13dBpi_twinax26_smooth	20 mm	27.980	0.688	0.125	0.002	-0.002	0.005	0.008	0.009	0.008	0.006	0.005	-0.003	0.033	-0.037	-0.025	0.027	0.010	0.001	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002
mellit...072518/Z0d_100_206in_0p13dBpi_twinax26_smooth	30 mm	27.980	0.700	0.079	-0.067	-0.062	-0.041	-0.026	-0.017	-0.012	-0.010	-0.008	-0.007	-0.007	-0.007	-0.006	-0.006	-0.006	-0.007	-0.003	0.018	-0.060	0.005	0.016	0.003	-0.003
mellitz...72518_channels--Z0d_100_14p25in_2dBpi_meg6_rtf	12 mm	27.980	0.700	0.026	-0.065	-0.037	-0.015	-0.011	0.041	-0.067	-0.003	0.038	0.008	0.000	-0.001	0.001	-0.006	0.006	-0.003	-0.003	0.000	-0.001	-0.001	-0.001	0.000	-0.002
mellitz...72518_channels--Z0d_100_14p25in_2dBpi_meg6_rtf	20 mm	27.980	0.685	0.067	-0.082	-0.068	-0.040	-0.022	-0.011	-0.006	-0.004	0.002	-0.012	0.035	-0.040	-0.025	0.029	0.013	0.002	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	0.000
mellitz...72518_channels--Z0d_100_14p25in_2dBpi_meg6_rtf	30 mm	27.980	0.700	0.129	-0.015	-0.021	-0.011	-0.005	-0.002	-0.001	-0.001	-0.001	-0.002	-0.002	-0.003	-0.003	-0.004	-0.002	-0.005	0.000	0.016	-0.055	0.006	0.016	0.004	-0.001

Back

Clause 91 100G with bit mux PMA and pre-coding

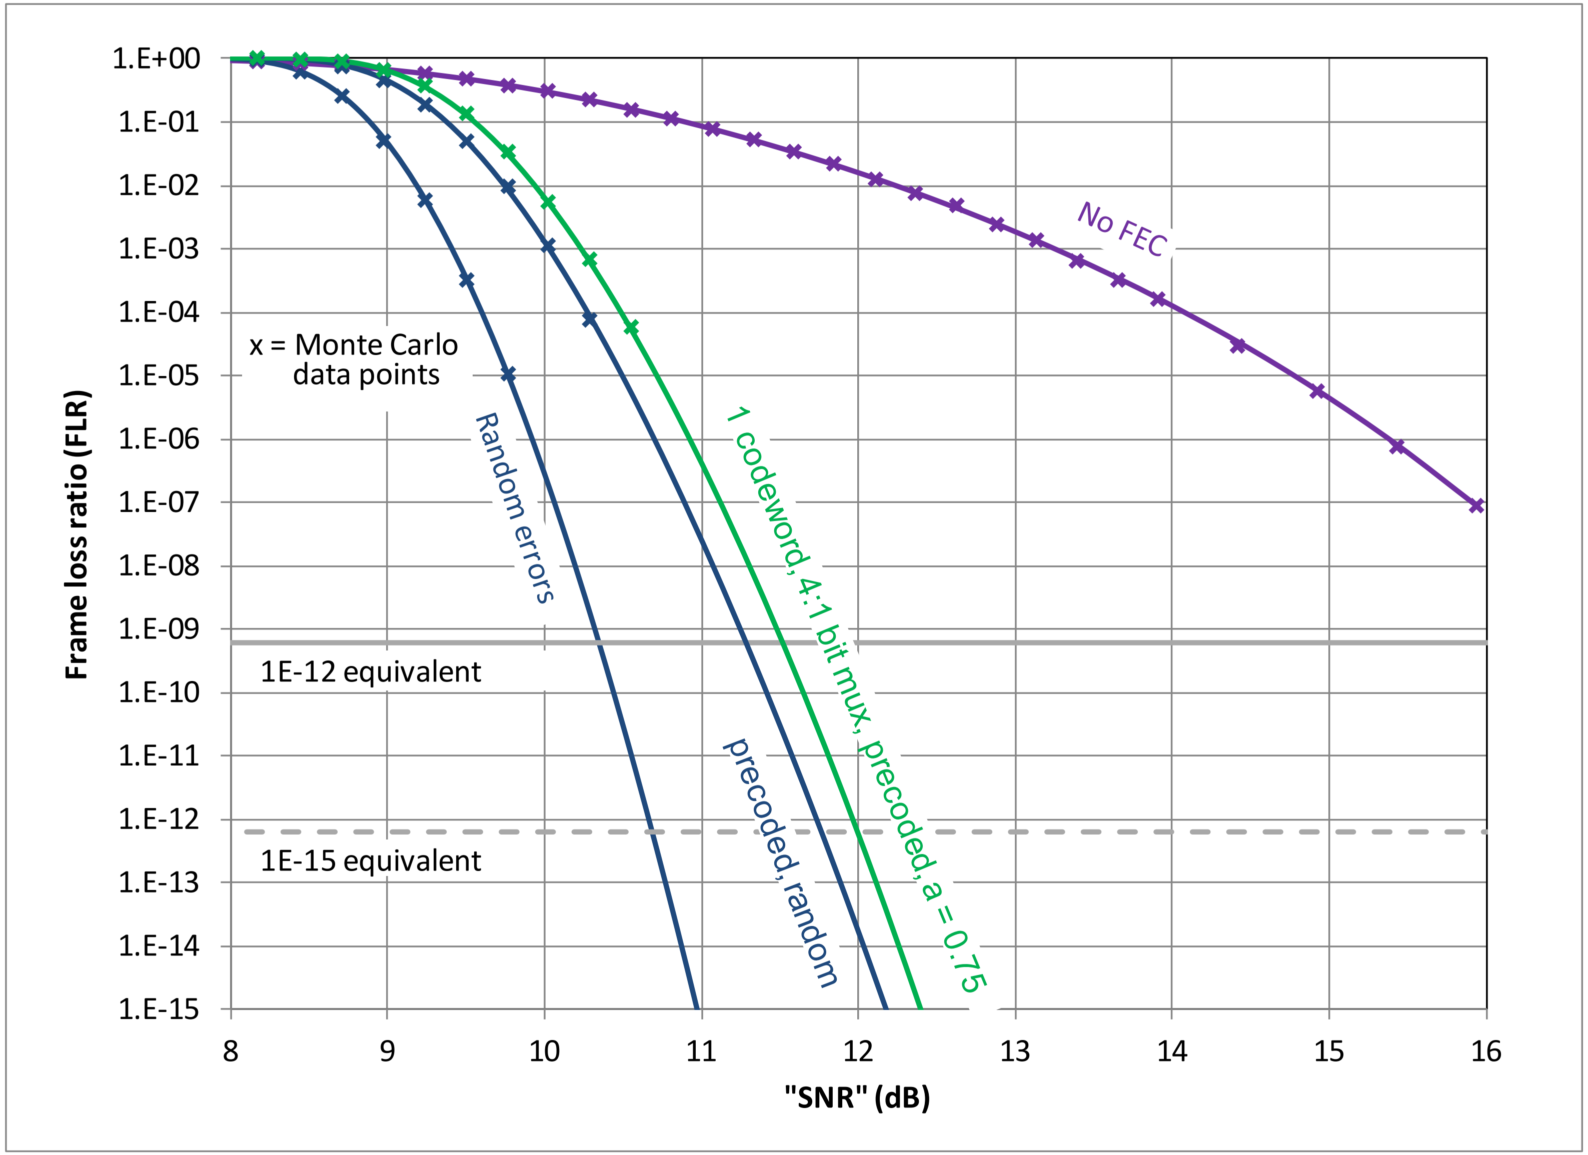

Back

Clause 119 400G with symbol mux PMA

Symbol interleave from 2 FEC codewords. Symbol multiplex in the PMA.

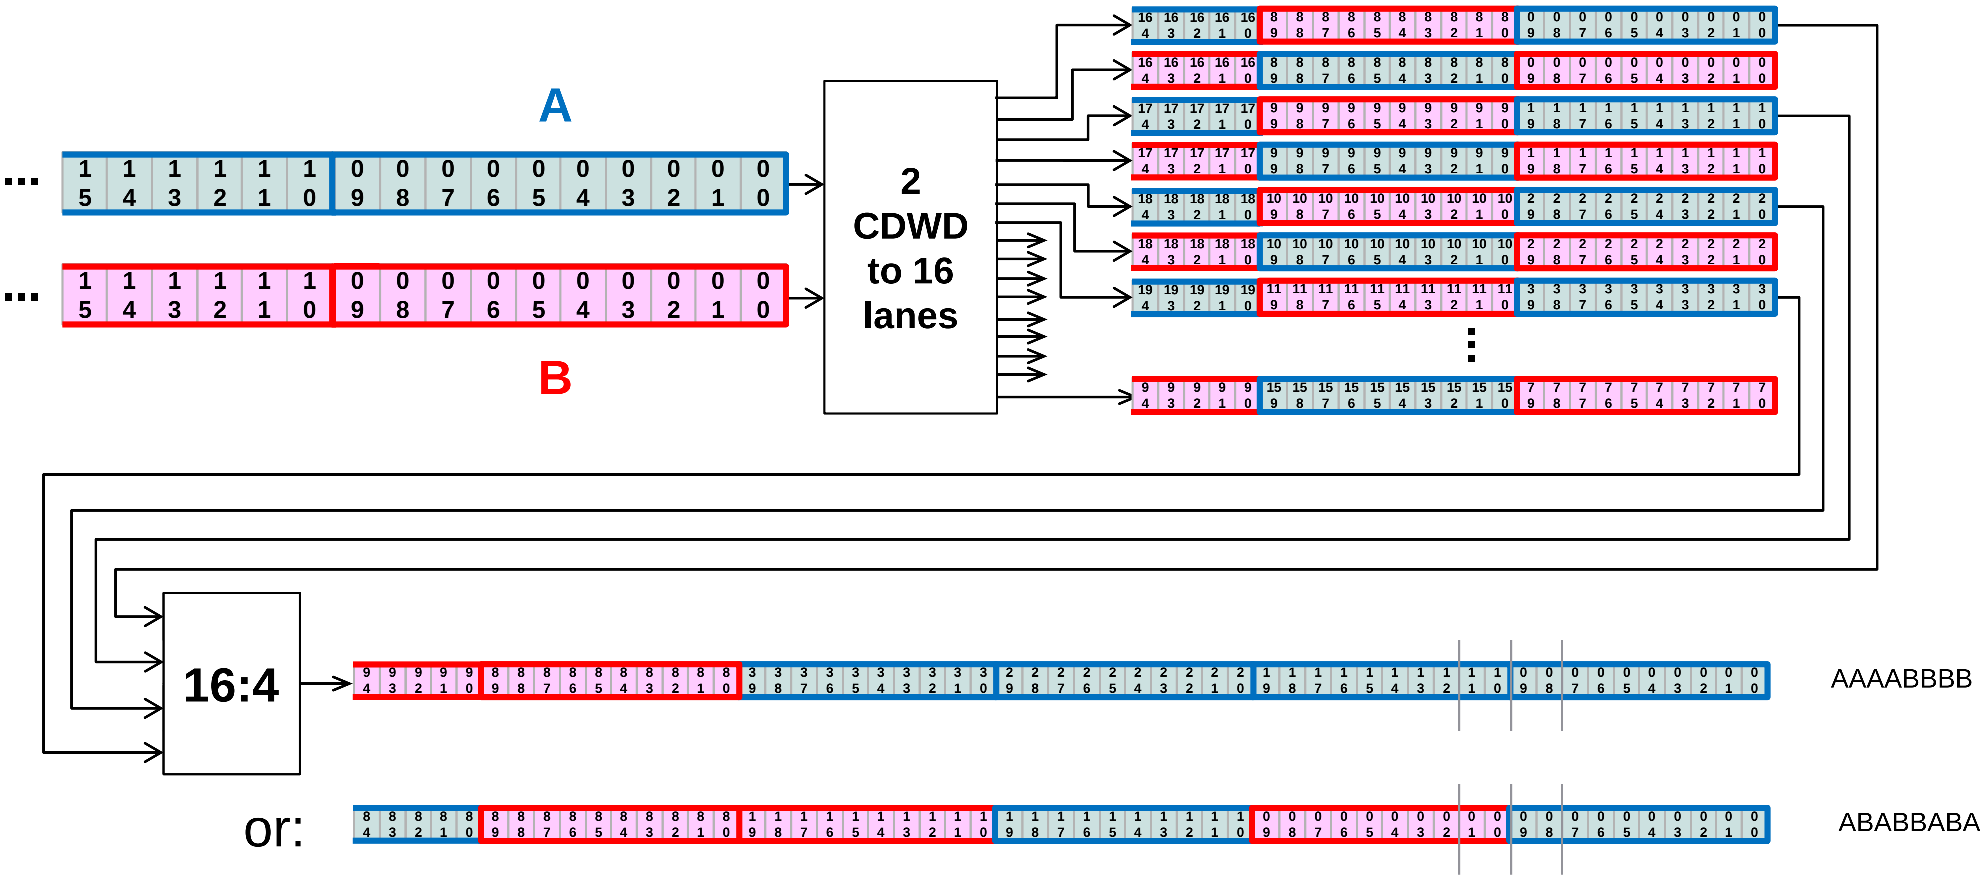

If one codeword is uncorrectable, the other is marked bad also.

Back

Clause 119 400G with symbol mux PMA

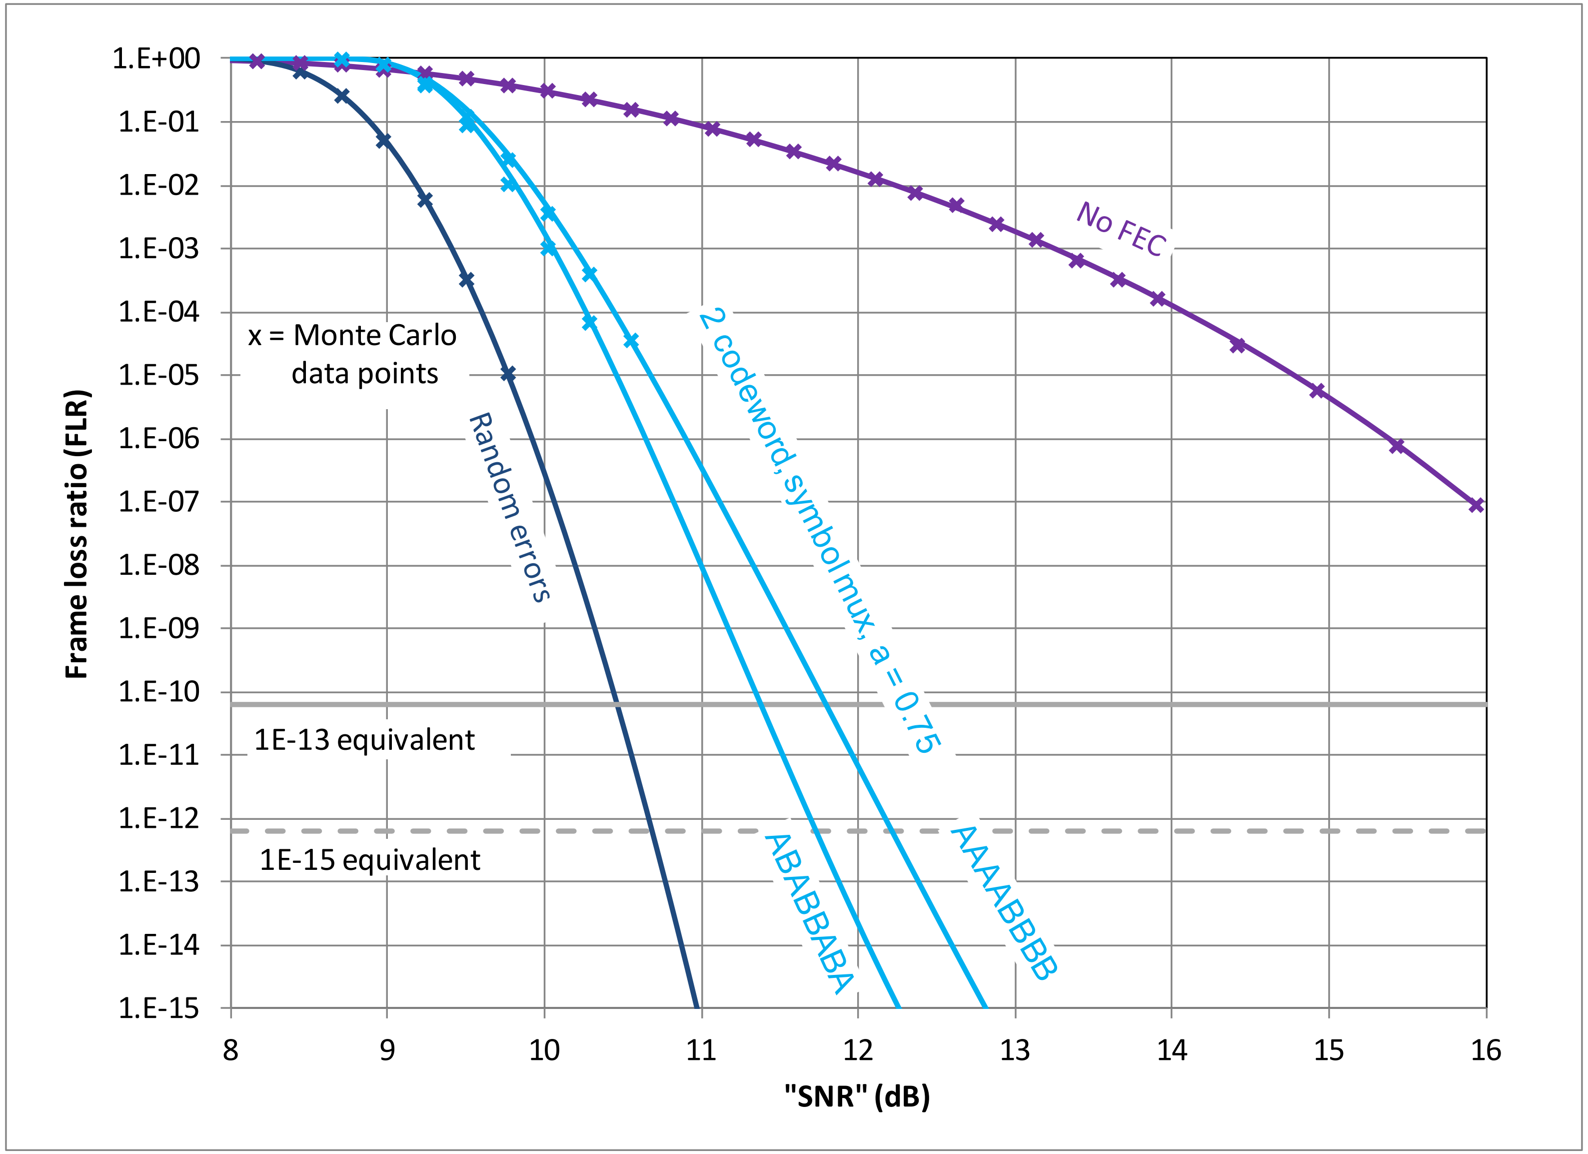

[Back](#)

Clause 119 400G with bit mux PMA

Symbol interleave from 2 FEC codewords. Bit multiplex in the PMA.

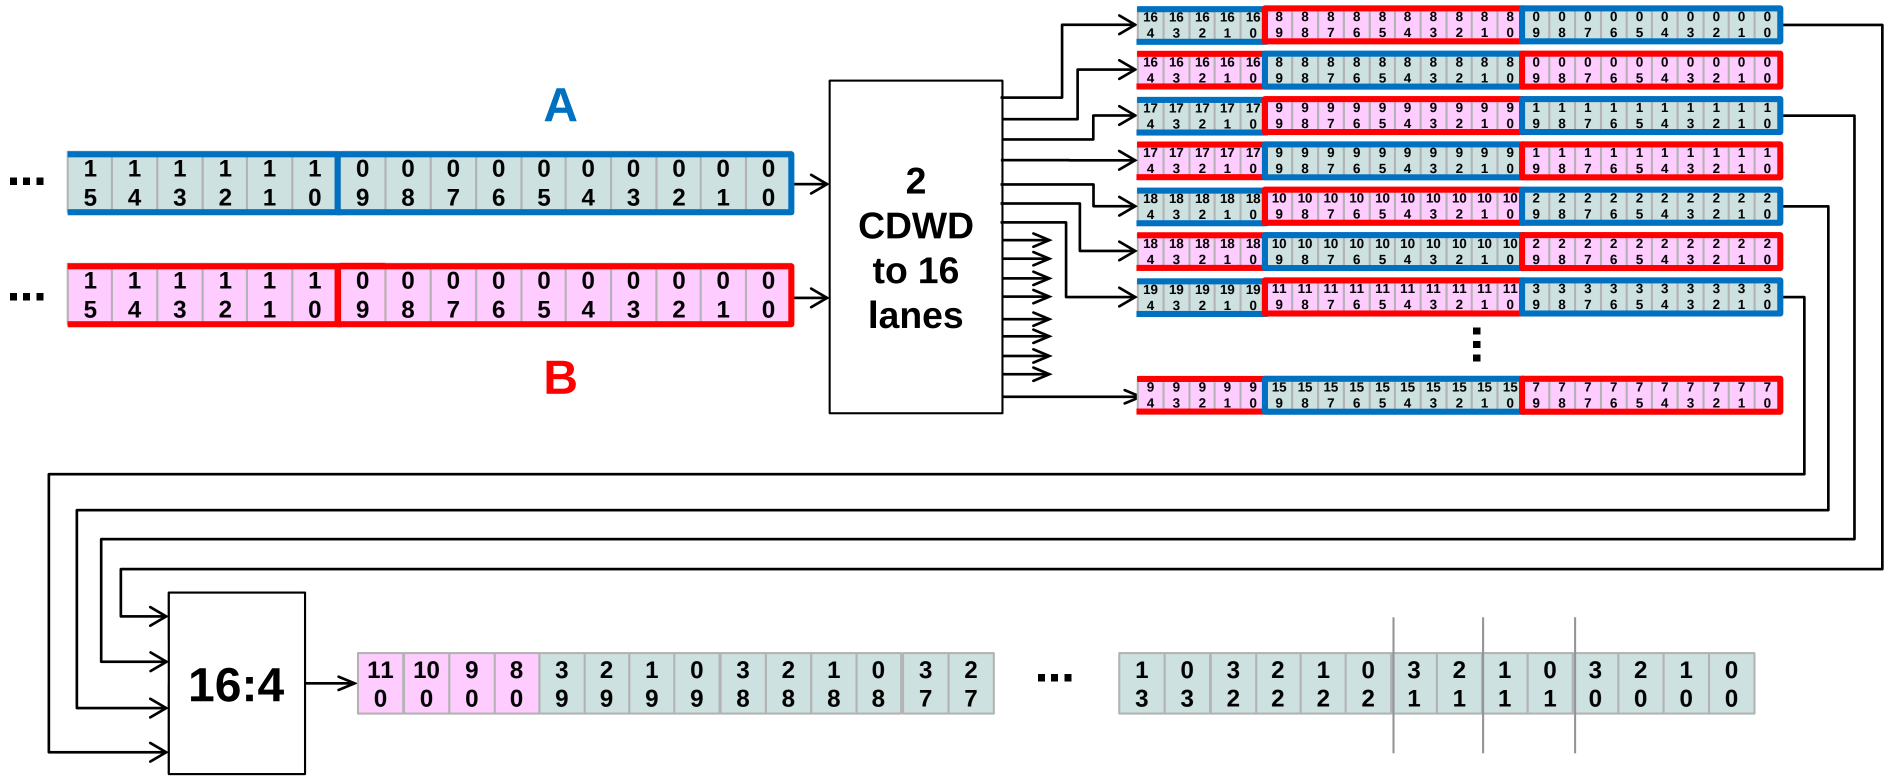

If one codeword is uncorrectable, the other is marked bad also.

Back

Clause 119 400G with bit mux PMA

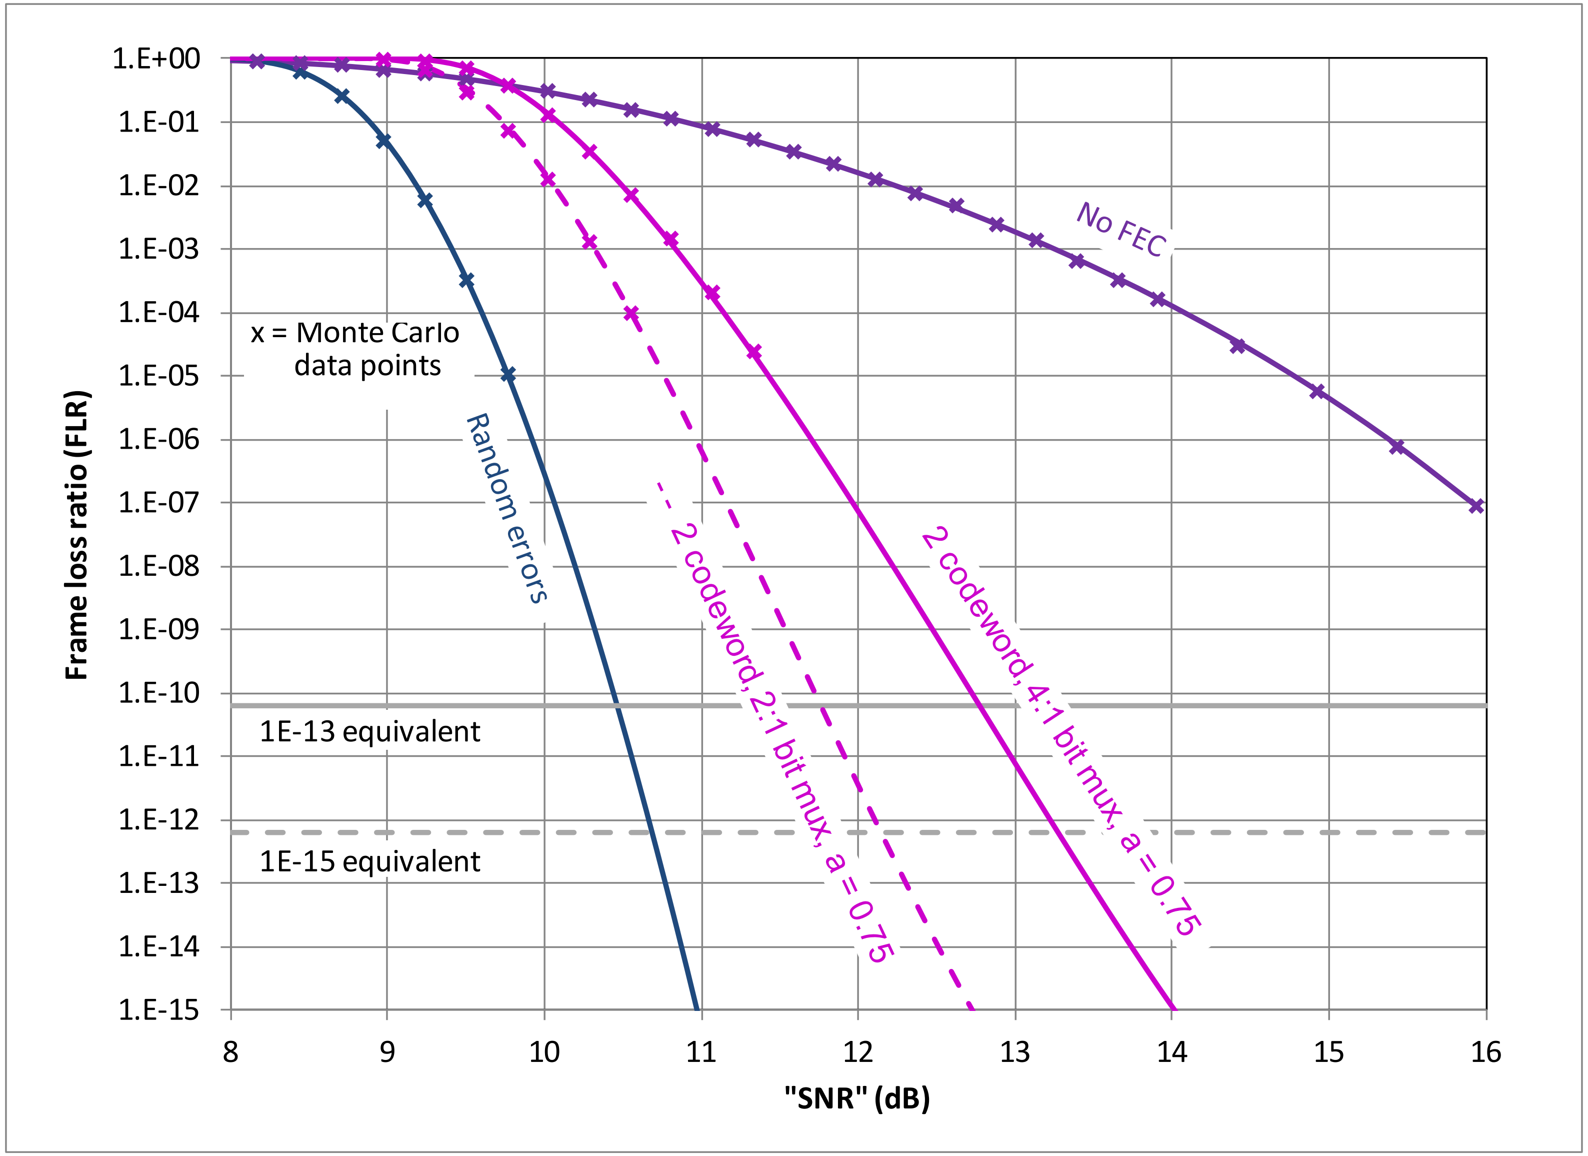

[Back]

Clause 119 400G with bit mux PMA and pre-coding

*

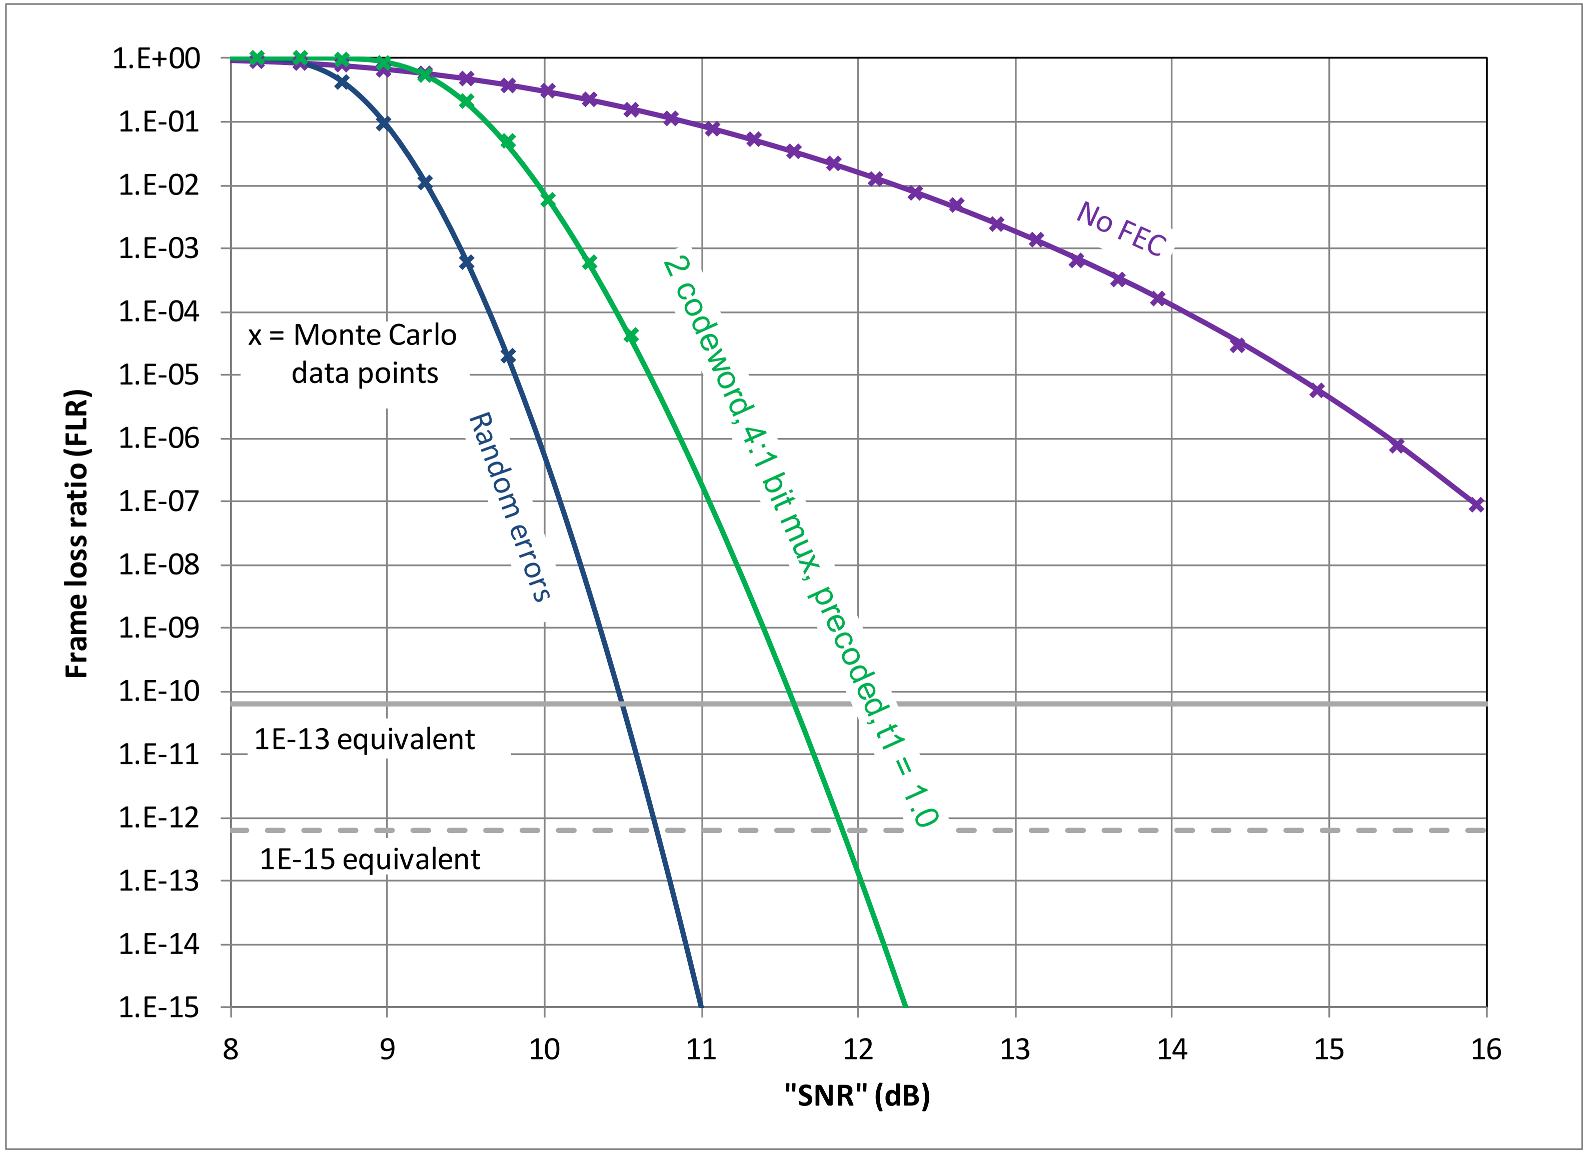

Back

Comparison with healey_100GEL_01_0318

100G curves re-plotted on same axes as healey_100GEL_01_0318

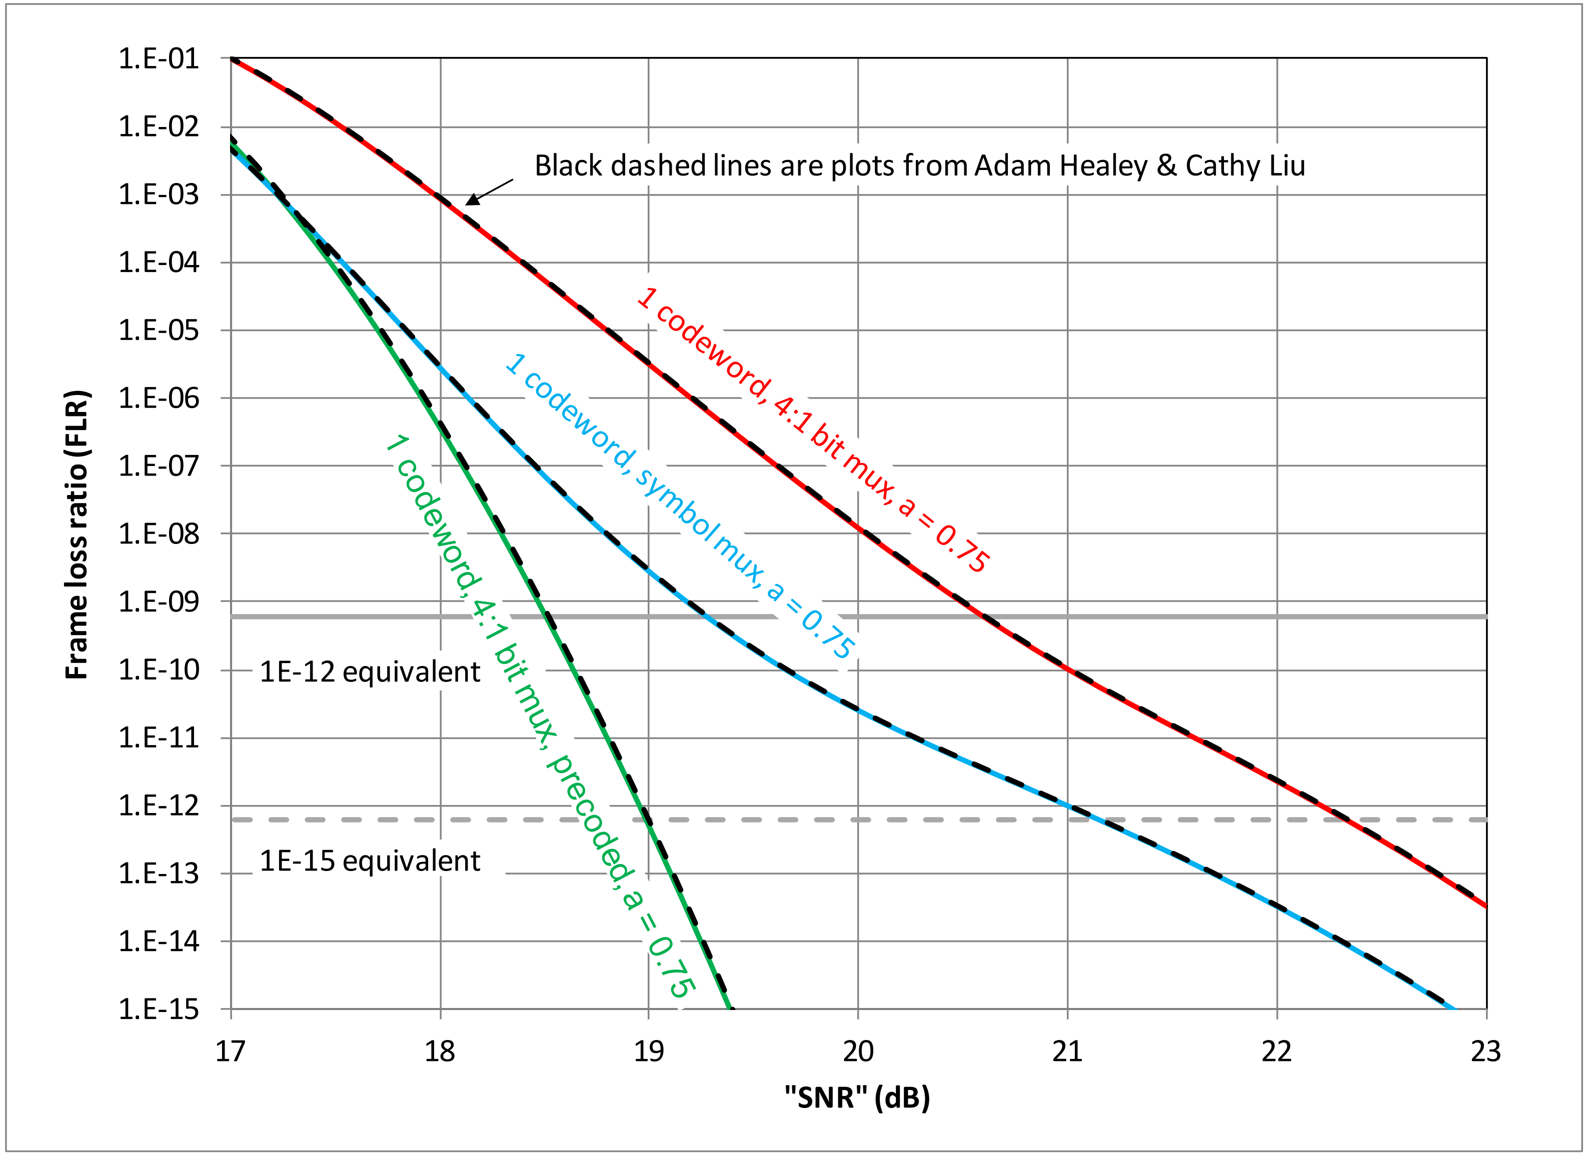

400G curves re-plotted on same axes as healey_100GEL_01_0318

*

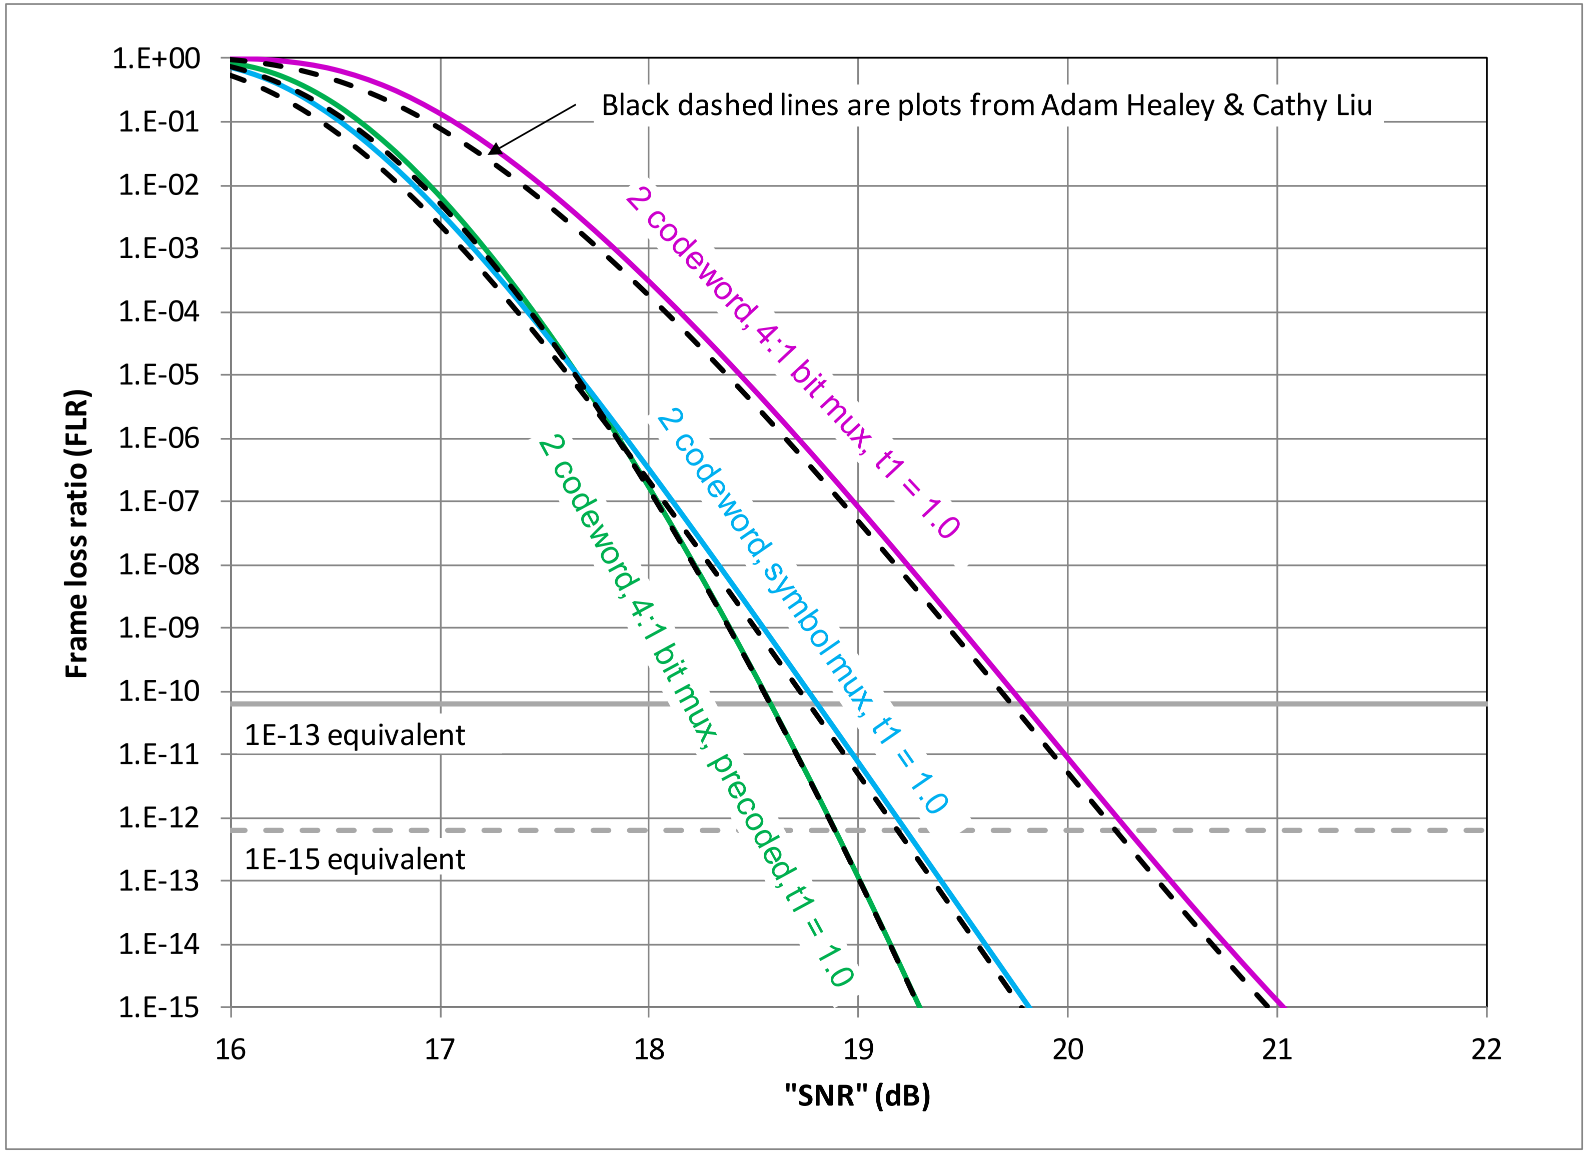

Back

Thanks!