

RS(544,514) FEC performance including precoding

Pete Anslow, Ciena

IEEE P802.3cd Task Force, 6 July 2016

Introduction

*

The IEEE 802.3bs Task Force has adopted RS(544,514) FEC with interleaving of FEC symbols from two FEC codewords to give good burst error tolerance.

Concerns over the latency of this scheme has led to proposals for either non-multiplexed or symbol multiplexed FEC schemes for 50 Gb/s and next generation 100 Gb/s Ethernet.

[anslow_3cd_01_0516](#) analysed the performance of such schemes using a development of the principles explained for the NRZ case in Annex 1 of [anslow_3bs_02_1114](#).

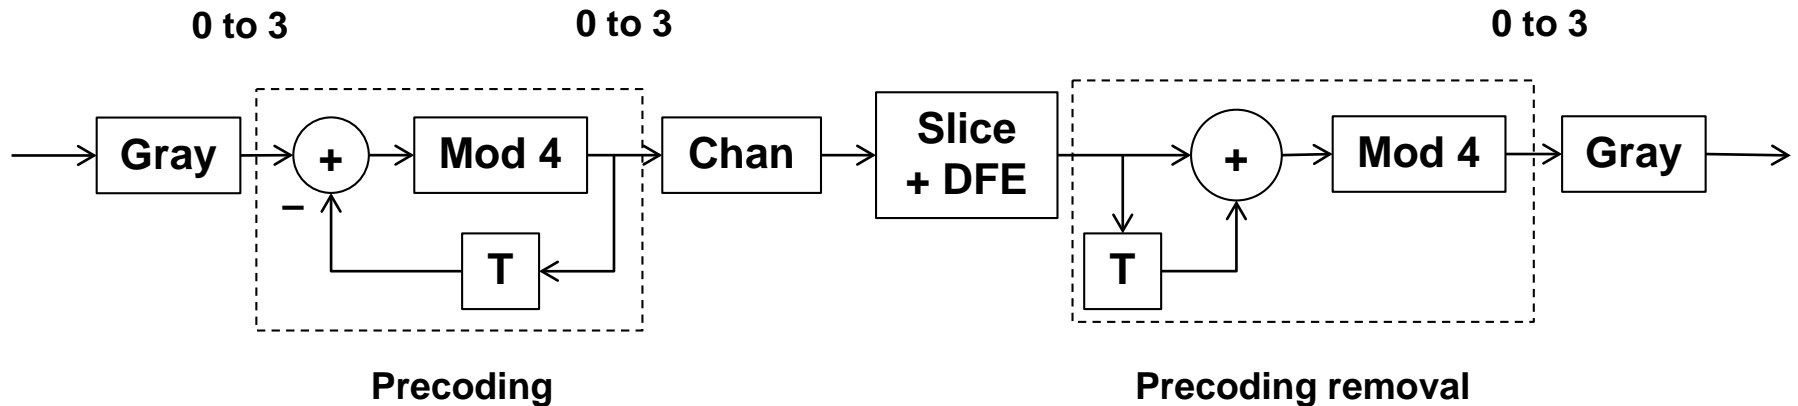
This presentation adds analysis of the performance with precoding.

Note: slides with a “*” in the top right corner are new or have changed compared to [anslow_3cd_01_0516](#). For changed slides, the changes are shown in red. For new slides, the whole title is red.

Precoding

*

Precoding as defined in 94.2.2.6 for 100GBASE-KP4 was assumed. This is performed as illustrated below.



See page 5 of [parthasarathy_01_0911](#) for a worked example.

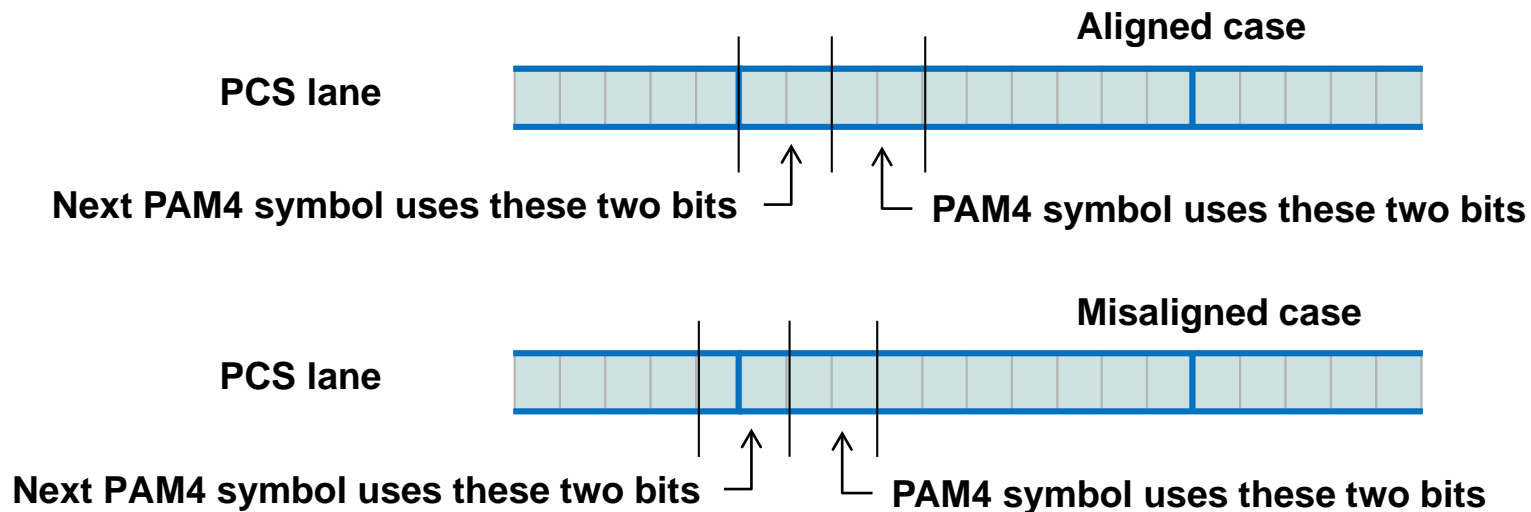
A “feature” of this precoding process is that a single random errored PAM4 symbol at the slicer output turns into two errored PAM4 symbols after the precoding is removed.

Signal structure

Assuming:

- A single PCS lane or multiple PCS lanes formed by round robin distribution of FEC symbols to the PCS lanes
- RS(544,514) FEC (which has 10-bit symbols)
- Gray coding (see P802.3bs D1.3 120.5.7)

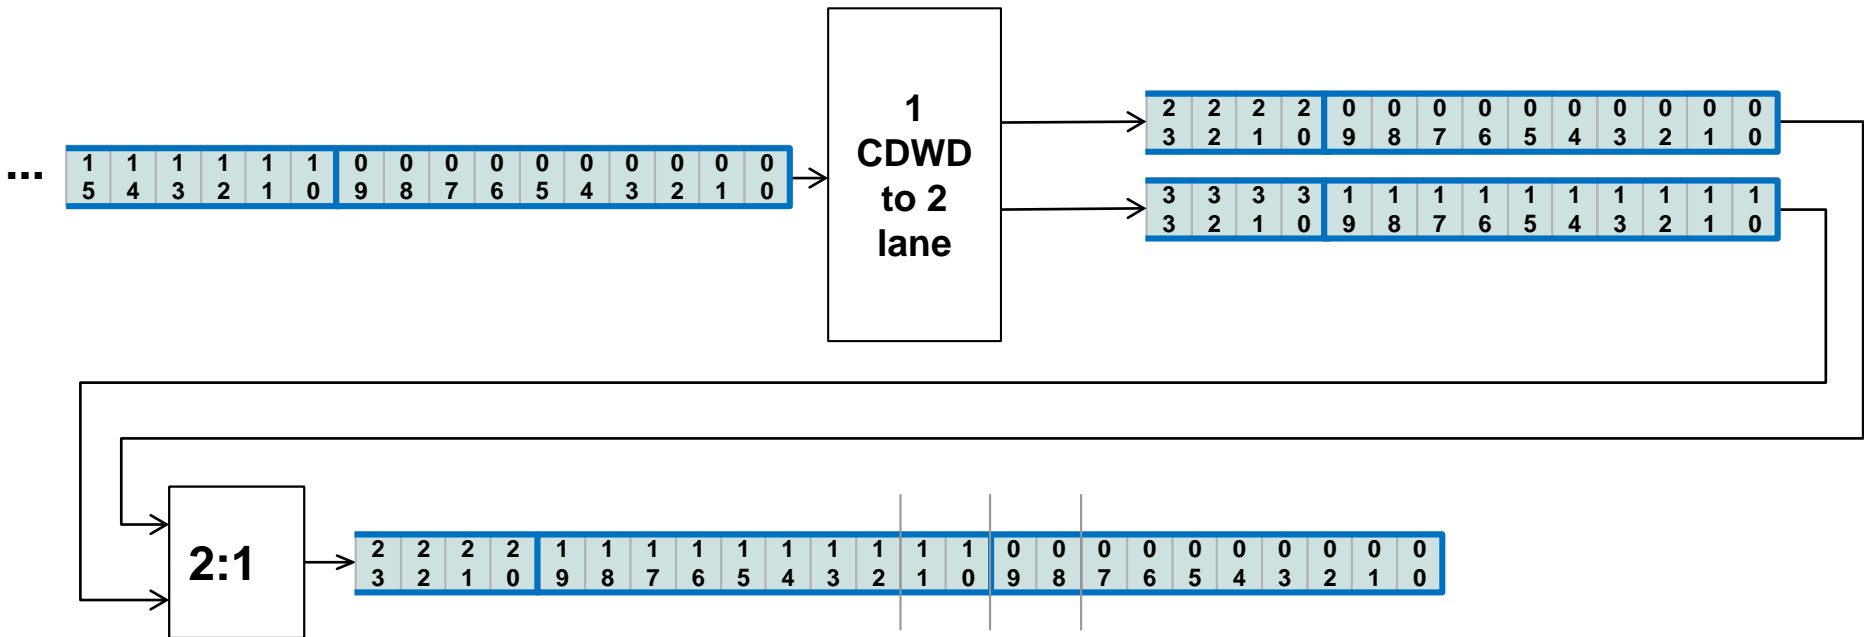
There are two ways that the PAM4 coding can occur:



The analysis in [anslow_3bs_03_0515](#) showed no difference in performance between the two cases.

Symbol multiplexing


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



The 2:1 PMA must find FEC symbol boundaries. If not totally de-skewed, the symbol order may be changed, but performance is the same as without multiplexing.

Gray coding

Assume the use of Gray coding (see P802.3bs D1.3 120.5.7) 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 $3.8\text{E-}4^*$ due to single level errors, then the probability of that noise causing both bits to be in error is $2\text{E-}24$.


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	✓, ✗	✓, ✓

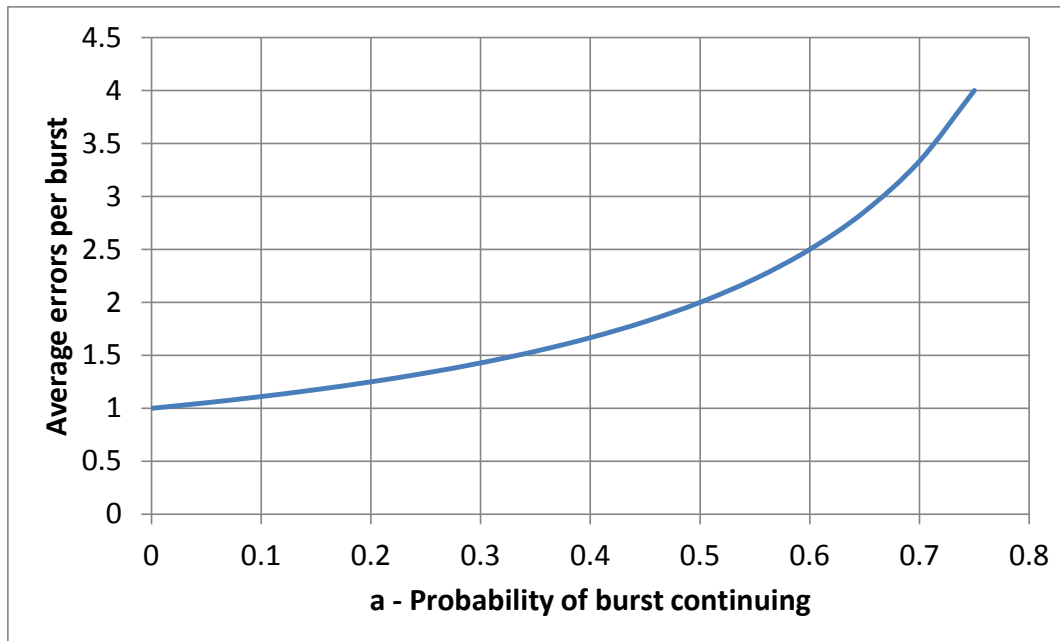
The analysis in the remainder of this contribution therefore assumes that if a given symbol is in error, the probability of a bit error in the first bit is $1/3$ and in the second bit is $2/3$.

Burst error model 3

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

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

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 $0.75^{74} = 5.7\text{E-}10$. 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:

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

For $a = 0.75$, this evaluates to $8.6\text{E-}10$.


This bound is plotted as a dashed line on the next page.

RS(544,514) no mux or symbol mux



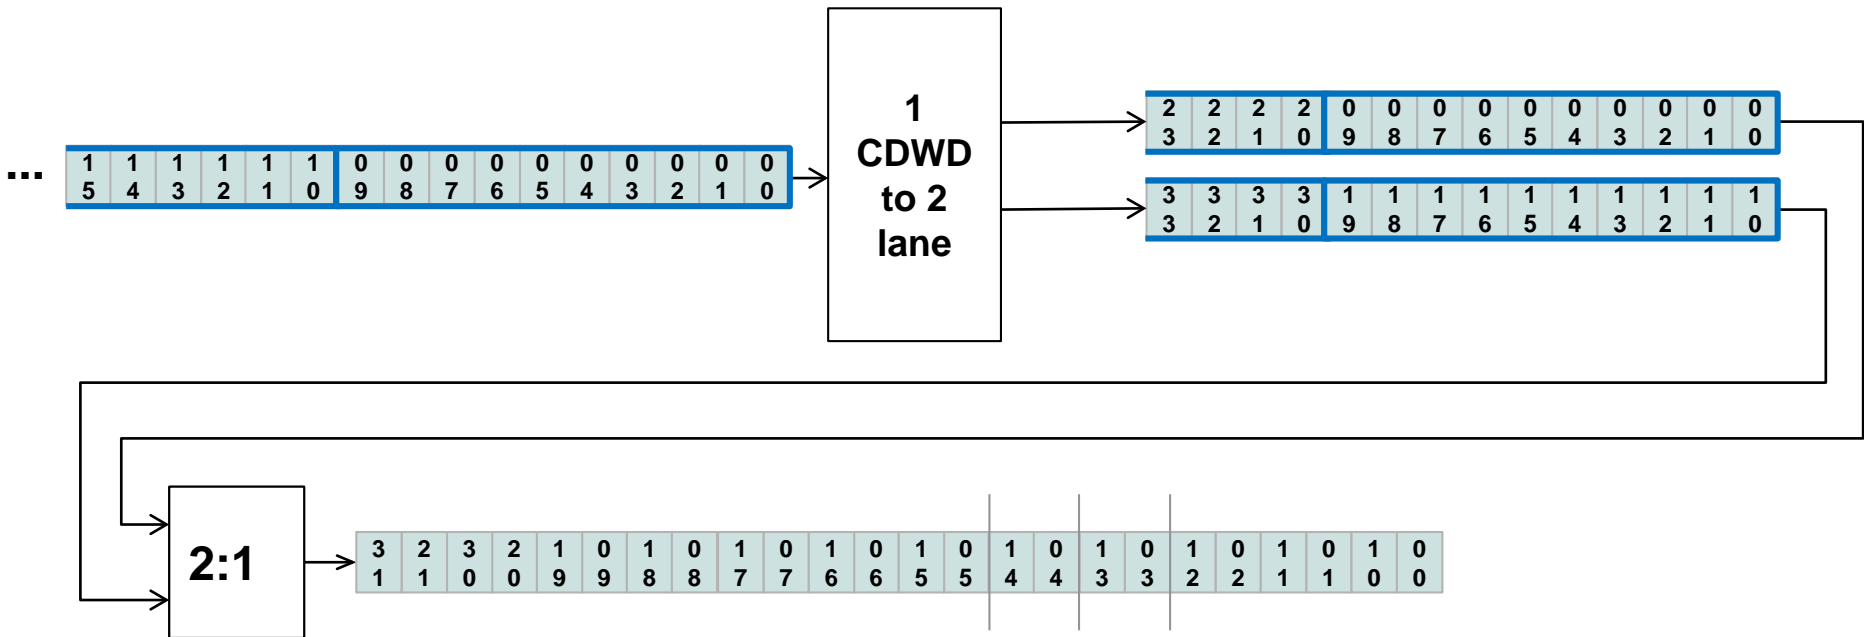
RS(544,514) no mux or s. mux with precoding

*




Bit multiplexing

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




Here a burst that is only 2 PAM symbols long is likely to hit 2 FEC symbols from the same codeword.

RS(544,514) 2:1 bit mux



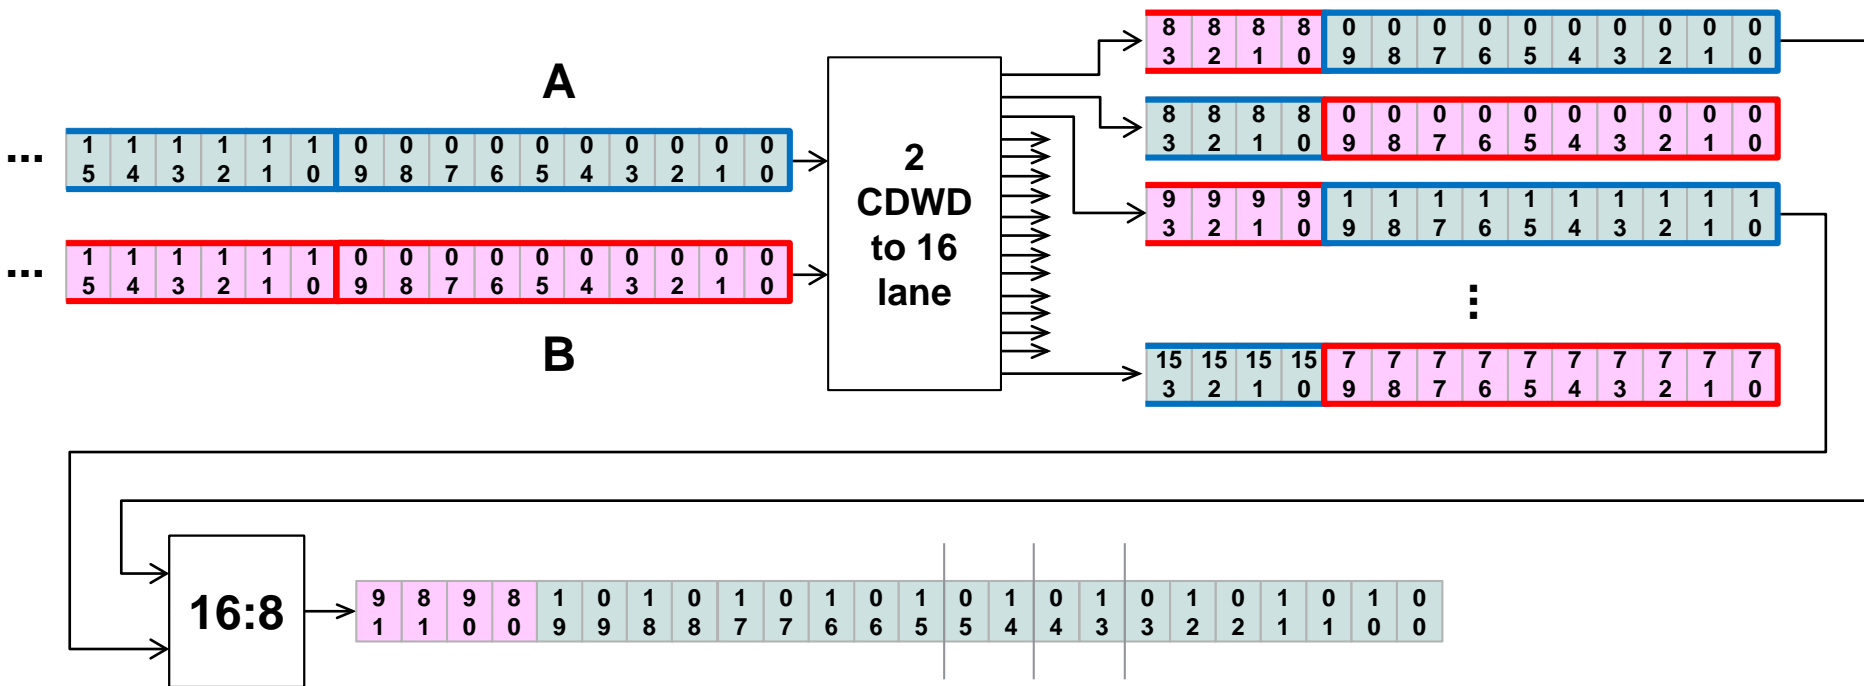
RS(544,514) 2:1 bit mux with precoding

*




P802.3bs D1.5 scheme

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



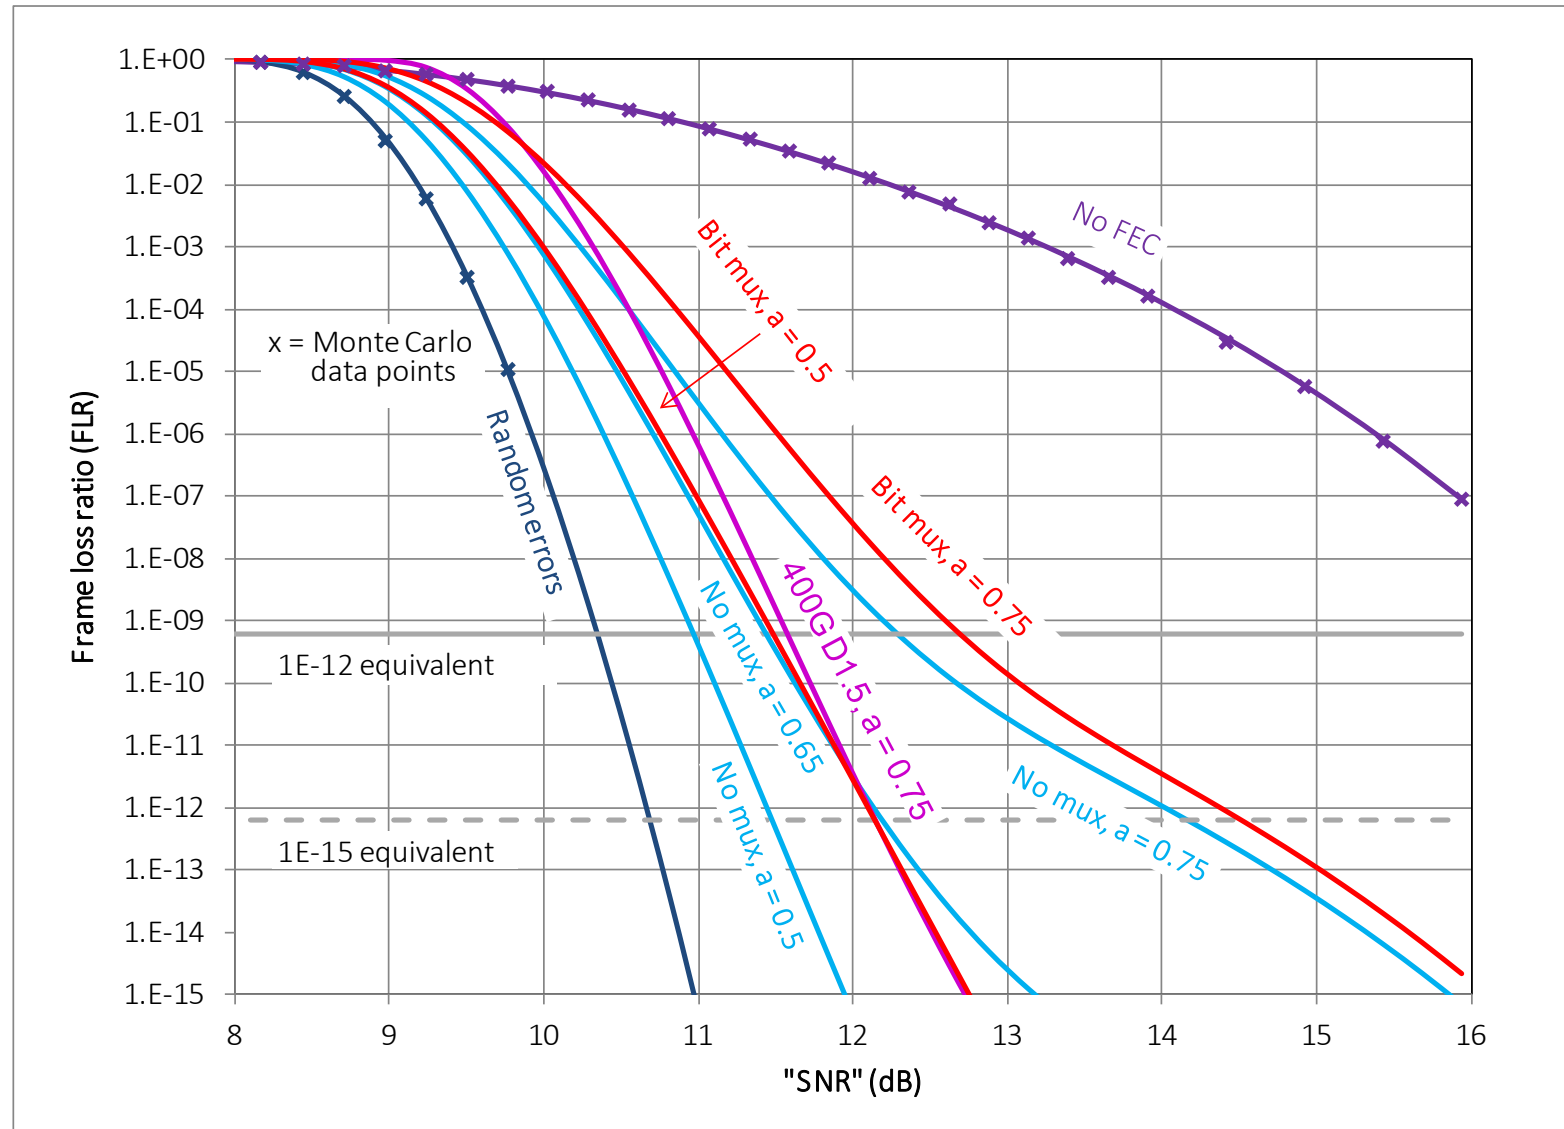
If one codeword is uncorrectable, the other is marked bad also.

P802.3bs D1.5 performance




All curves **without** precoding

*



Precoding vs 400G

*



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

*

From the curves shown on **page 18 (without precoding)**, if all of the coding gain were to be used for the PAM4 link, the BERs at the FEC input required to give FLRs equivalent to a BER of 1E-12 and 1E-15 are:

	RS(544,514)	
	FLR = 6.2E-10	FLR = 6.2E-13
No FEC	1E-12	1E-15
2:1 bit mux, a = 0.75	2.5E-5*	1.6E-7*
No mux, a = 0.75	5.9E-5*	4.9E-7*
2:1 bit mux, a = 0.5	1.3E-4*	3.9E-5*
No mux, a = 0.65	2.1E-4*	5.1E-5*
P802.3bs D1.5, a = 0.75	2.3E-4*	7.8E-5*
No mux, a = 0.5	3.1E-4*	1.3E-4*
Random errors	3.8E-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 when a = 0.75, 2.9 when a = 0.65, 2 when a = 0.5.

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

*

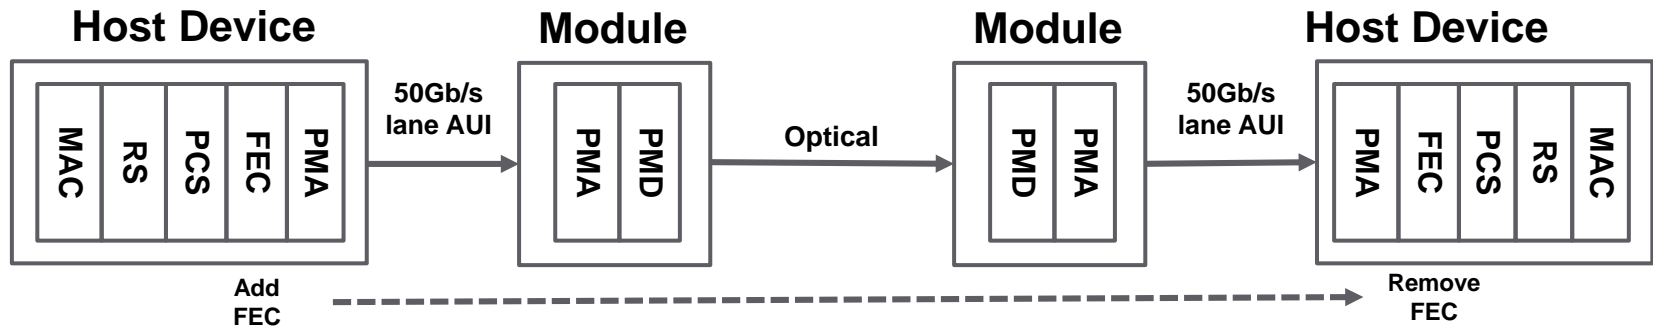
From the curves with precoding, if all of the coding gain were to be used for the PAM4 link, the BERs at the FEC input required to give FLRs equivalent to that of a BER of $1\text{E-}12$ and $1\text{E-}15$ are:

	At slicer output		At FEC input	
	FLR = $6.2\text{E-}10$	FLR = $6.2\text{E-}13$	FLR = $6.2\text{E-}10$	FLR = $6.2\text{E-}13$
No FEC	$1\text{E-}12$	$1\text{E-}15$	$1\text{E-}12$	$1\text{E-}15$
1:2 bit mux precoded, random	$9.3\text{E-}5$	$4\text{E-}5$	$1.9\text{E-}4$	$8\text{E-}5$
Precoded, random	$1.7\text{E-}4$	$7.9\text{E-}5$	$3.3\text{E-}4$	$1.6\text{E-}4$
P802.3bs D1.5, $a = 0.75$	$2.3\text{E-}4^*$	$7.8\text{E-}5^*$	$2.3\text{E-}4$	$7.8\text{E-}5$
1:2 bit mux precoded, $a=0.75$	$3\text{E-}4^*$	$1.3\text{E-}4^*$	$1.5\text{E-}4$	$6.3\text{E-}5$
Precoded, $a=0.75$	$3.2\text{E-}4^*$	$1.3\text{E-}4^*$	$1.6\text{E-}4$	$6.7\text{E-}5$
Random errors	$3.8\text{E-}4$	$2.3\text{E-}4$	$3.8\text{E-}4$	$2.3\text{E-}4$

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 when $a = 0.75$, 2.9 when $a = 0.65$, 2 when $a = 0.5$.

Multi-part links with FEC

If the FEC bytes are added at the source FEC sublayer and then the correction is applied only at the destination FEC sublayer as in:



Then the worst case input BER for the FEC decoder must be met by the concatenation of all of the sub-links.

In the case of CDAUI-8 -> FR8 -> CDAUI-8, the worst case BER for each lane of the electrical sub-links is $1E-5$. Even though there may be two additional CDAUI-8 C2C sub-links, this is tolerated on the basis that it is extremely unlikely that all four sub-links will be at the worst case BER at the same time given that each sub-link BER is averaged over 8 lanes.

The results for multiple sub-links sharing the same RS(544,514) protection is shown on the next slide.

Multi-part link results 1 (no precoding) *

The BER of the electrical sub-links for a BER of $2.4\text{E-}4$ in the optical sub-link are shown in the table below (0.16 dB optical penalty).

	RS(544,514) FLR = $6.2\text{E-}10$			
	Electrical		Optical	
2:1 bit mux, $a = 0.75$	Burst	$2.5\text{E-}6^*$	Random	$2.4\text{E-}4$
No mux, $a = 0.75$	Burst	$6.3\text{E-}6^*$	Random	$2.4\text{E-}4$
2:1 bit mux, $a = 0.5$	Burst	$3.8\text{E-}5^*$	Random	$2.4\text{E-}4$
No mux, $a = 0.65$	Burst	$5.7\text{E-}5^*$	Random	$2.4\text{E-}4$
P802.3bs D1.5, $a = 0.75$	Burst	$6.3\text{E-}5^*$	Random	$2.4\text{E-}4$
No mux, $a = 0.5$	Burst	$1\text{E-}4^*$	Random	$2.4\text{E-}4$
Random errors	Random	$1.4\text{E-}4$	Random	$2.4\text{E-}4$

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 when $a = 0.75$, 2.9 when $a = 0.65$, 2 when $a = 0.5$.

Multi-part link results 2 (precoded)

*

The BER of the electrical sub-links for a BER of $2.4\text{E-}4$ in the optical sub-link are shown in the table below (0.16 dB optical penalty).

	RS(544,514) At slicer output for FLR = $6.2\text{E-}10$			
	Electrical		Optical	
1:2 bit mux precoded, random	Random	$3.1\text{E-}5$	Random	$2.4\text{E-}4$
Precoded, random	Random	$5.7\text{E-}5$	Random	$2.4\text{E-}4$
P802.3bs D1.5, $a = 0.75$	Burst	$6.3\text{E-}5^*$	Random	$2.4\text{E-}4$
1:2 bit mux precoded, $a=0.75$	Burst	$1\text{E-}4^*$	Random	$2.4\text{E-}4$
Precoded, $a=0.75$	Burst	$1.1\text{E-}4^*$	Random	$2.4\text{E-}4$
Random errors	Random	$1.4\text{E-}4$	Random	$2.4\text{E-}4$

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 when $a = 0.75$.

Multi-part link results 3

*

For the two cases where the electrical BER is below $2\text{E-}5$ (does not allow $1\text{E-}5$ for each of two AUI sub-links) the table below shows what the optical BER would have to be reduced to.

	RS(544,514) FLR = $6.2\text{E-}10$			
	Electrical		Optical	
2:1 bit mux, $a = 0.75$	Burst	$2\text{E-}5^*$	Random	$3.6\text{E-}5$
No mux, $a = 0.75$	Burst	$2\text{E-}5^*$	Random	$1.4\text{E-}4$

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 when $a = 0.75$.

Multi-part link results 4

*

The BER of the sub-links for an FLR equivalent to $1\text{E-}15$ BER and the same optical penalty are shown in the table below.

	RS(544,514) FLR = $6.2\text{E-}13$			
	Electrical		Optical	
2:1 bit mux, $a = 0.75$	Burst	$1.3\text{E-}7^*$	Random	$1.9\text{E-}5$
No mux, $a = 0.75$	Burst	$1.2\text{E-}7^*$	Random	$8.1\text{E-}5$
2:1 bit mux, $a = 0.5$	Burst	$1.1\text{E-}5^*$	Random	$1.4\text{E-}4$
No mux, $a = 0.65$	Burst	$1.1\text{E-}5^*$	Random	$1.4\text{E-}4$
P802.3bs D1.5, $a = 0.75$	Burst	$2.2\text{E-}5^*$	Random	$1.4\text{E-}4$
No mux, $a = 0.5$	Burst	$4.7\text{E-}5^*$	Random	$1.4\text{E-}4$
Random errors	Random	$9.4\text{E-}5$	Random	$1.4\text{E-}4$

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 when $a = 0.75$, 2.9 when $a = 0.65$, 2 when $a = 0.5$.

Multi-part link results 5 (precoded)

*

The BER of the sub-links for an FLR equivalent to $1\text{E-}15$ BER and the same optical penalty are shown in the table below.

	RS(544,514) At slicer output for FLR = $6.2\text{E-}13$			
	Electrical		Optical	
1:2 bit mux precoded, random	Random	$1.5\text{E-}5$	Random	$1.4\text{E-}4$
P802.3bs D1.5, $a = 0.75$	Burst	$2.2\text{E-}5^*$	Random	$1.4\text{E-}4$
Precoded, random	Random	$2.9\text{E-}5$	Random	$1.4\text{E-}4$
1:2 bit mux precoded, $a=0.75$	Burst	$4.7\text{E-}5^*$	Random	$1.4\text{E-}4$
Precoded, $a=0.75$	Burst	$5\text{E-}5^*$	Random	$1.4\text{E-}4$
Random errors	Random	$9.4\text{E-}5$	Random	$1.4\text{E-}4$

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 when $a = 0.75$.

Conclusion 1

*

Solutions for RS(544,514) FEC schemes for 50 Gb/s and next generation 100 Gb/s Ethernet considered here are:

- Allow the probability of a burst continuing “a” to be as high as 0.75 and restrict the electrical sub-link BER to be $\sim 3\text{E}-6$ (or $\sim 1\text{E}-6$ for bit mux) see [results 1](#)
 - Reduced capability for electrical links and large margin needed to reach $1\text{E}-15$ ([results 4](#))
- Allow “a” to be as high as 0.75 and restrict the optical sub-link BER to be $\sim 1.4\text{E}-4$ (or $\sim 3.6\text{E}-5$ for bit mux) with optical penalty 0.34 dB (or 0.74 dB for bit mux) see [results 3](#)
 - Reduced capability for optical links and large margin needed to reach $1\text{E}-15$
- Restrict “a” to 0.65 (or 0.5 for bit mux) see [results 1](#)
 - Reduced capability for electrical links and not clear how a limit for “a” is achieved

Conclusion 2

*

- Apply precoding to bursty electrical sub-links (remove at electrical receiver) see [results 2](#)
 - Good solution as long as practical Rx DFE produces an error pattern that becomes two errored symbols after decoding
 - If the above condition is true then there is no significant difference for bit mux
- For links where the majority of the errors are made in the bursty sublink (e.g. backplane or copper cable) text such as in 121.1.1 can be added:

“The bit error ratio (BER) when processed according to Clause 120 shall be less than 2.4×10^{-4} provided that the error statistics are sufficiently random that this results in a frame loss ratio (see 1.4.223) of less than 1.7×10^{-12} for 64-octet frames with minimum interpacket gap when processed according to Clause 120 and then Clause 119. For a complete Physical Layer, the frame loss ratio may be degraded to 6.2×10^{-11} for 64-octet frames with minimum interpacket gap due to additional errors from the electrical interfaces.”
- For links where the majority of the errors are made in a different sublink from the bursty one (e.g. chip-to-chip) would need a revised requirement that the frame loss ratio has to be met with additional random errors.

Thanks!