# RS-FEC with error bursts at 200G/lane

## compared to earlier rates

Adee Ran, Cisco

# Supporters

- Will Bliss — Broadcom
- Vasudevan Parthasarathy — Broadcom

# Introduction

- A popular architecture in 100G SerDes includes a single-tap DFE and a longer FFE. This architecture is likely to be used with some 200G SerDes as well.
  - Typically assumed for high-loss electrical links, but also used in some optical receivers.
- When using PAM4 modulation, a single-tap DFE can cause error bursts (contiguous sequences of PAM4 symbol errors) due to error propagation.
- RS-FEC is considered resilient to burst errors (no impact if a burst stays within one FEC symbol) but with bit-muxing PMAs, error bursts can corrupt increasing number of FEC symbols (depending on muxing ratio) and degrade the FEC performance.
- Precoding (as defined in 135.5.7.2) can be used to mitigate error bursts.
  - Not specified with optical and C2M.
- This presentation examines the effect of error propagation with bit-muxing to 200G/lane for the 200GBASE-R, 400GBASE-R, and the proposed 800GBASE-R PCSs.

# Some terminology used in this presentation

- PAM4 symbol (2 bits)
  - NRZ symbol (1 bit)
- Burst of length L – a sequence of L symbols which are received incorrectly
- FEC symbol (10 bits)
- FEC codeword or CW (n FEC symbols) – correctable if up to t symbol errors occur, otherwise uncorrectable, and its whole content is discarded.
  - In clause 119 n=544 and t=15.
  - Codewords are distributed to PCS lanes symbol-by-symbol in a checkerboard pattern.
- DER = detector error ratio (PAM4 symbol errors)
  - SER = FEC symbol error ratio
  - CER = FEC codeword error ratio
- FLR = Frame loss ratio

# Error model

- Every burst starts with a random error that occurs with probability $p_{initial} = DER$.
- Receiver with 1-tap DFE
  - Gilbert model; probability of propagation denoted *a*
  - With PAM4, $0 \leq a \leq 0.75$
- Consider two cases:
  1. Without precoding
     - Error propagation creates bursts of length L with probability $p_{initial} \cdot a^{L}(1-a)$
     - Every PAM4 symbol in the burst has either incorrect MSB or incorrect LSB (ignoring low probability event that both are incorrect)
     - Errors are twice as likely to occur in the LSB than in the MSB (see anslow_3ck_adhoc_01_072518, slide 11)
  2. With precoding
     - Every burst creates two incorrect PAM4 symbols; the initial random error and an additional error with distance L UI from the initial error, with probability $p_{initial} \cdot a^{L-1}(1-a)$
     - Errors occur equally likely in the LSB or the MSB (due to the removal of precoding)

# Recall: 200GBASE-R PCS (Clause 119) with 8 physical lanes (25G NRZ)

When data is transmitted over 8 physical lanes (NRZ), the bits on each lane comprise FEC symbols, and bursts usually corrupt only 1 FEC symbol (or 1+1 due to CW interleaving)

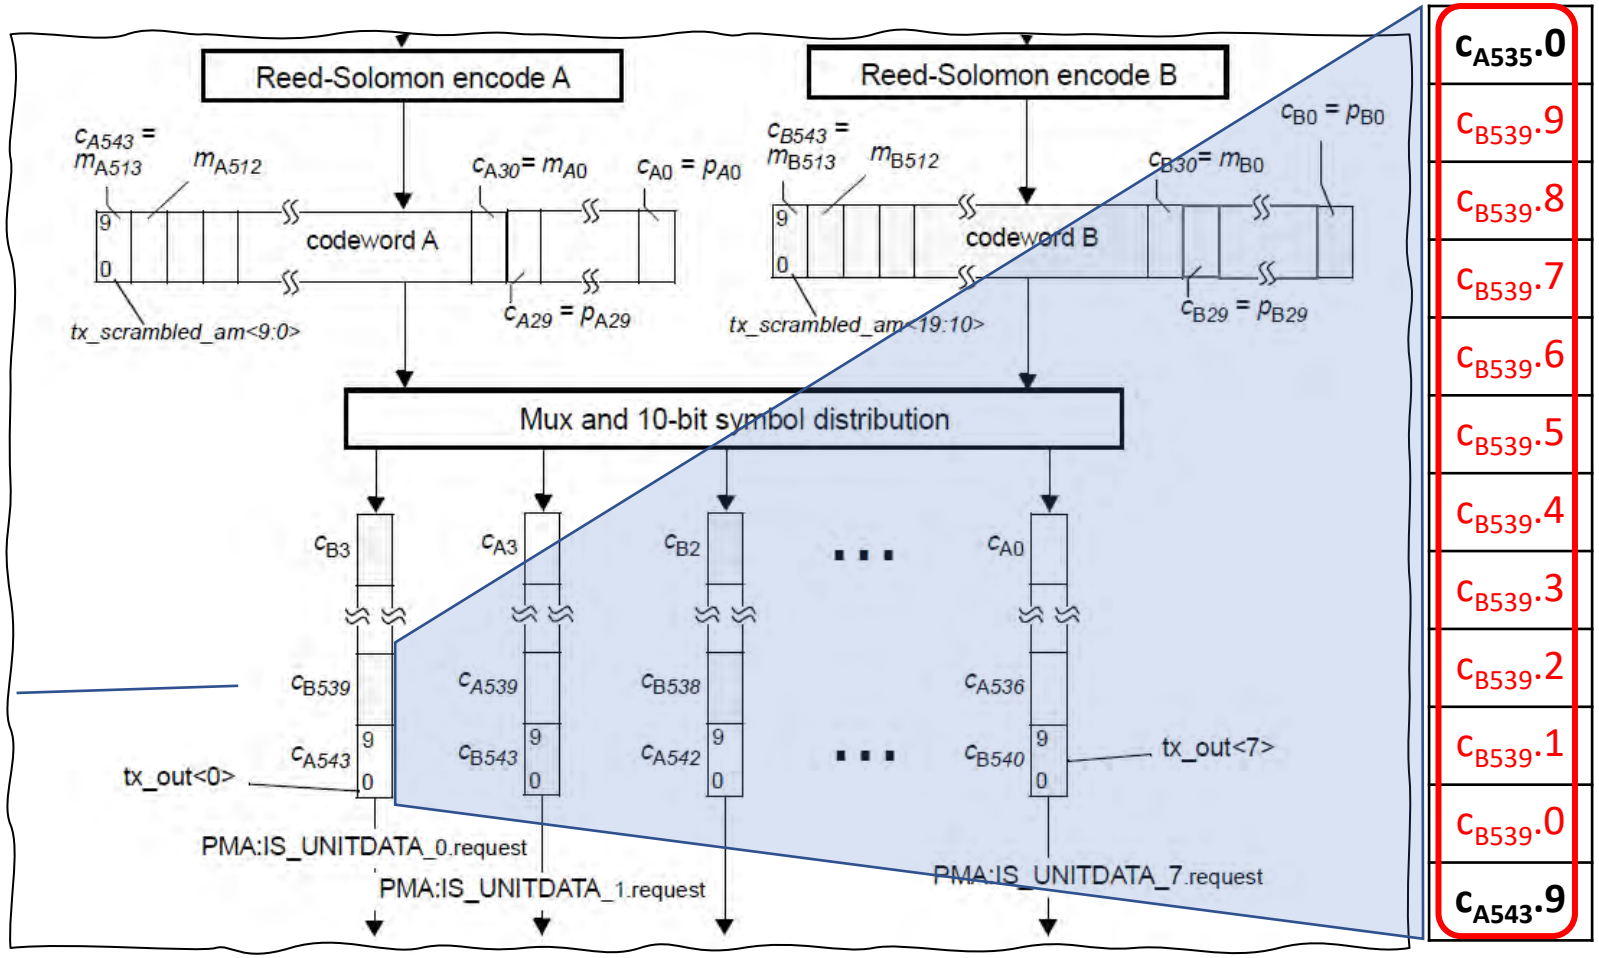





To corrupt 2 symbols of the same CW, a burst has to be at least 12 UI long.
3 symbols in one CW require 32 UI, etc.

Excerpt from Figure 119–10—200GBASE-R Transmit bit ordering and distribution

# Recall: 200GBASE-R PMA (Clause 120)

- Uses bit-muxing of multiple PCS lanes, as originally introduced in 802.3ba (clause 83 PMA and clause 82 PCS)
  - At the time, physical lanes used NRZ modulation
  - Bit-muxing enables handling of skew (static and dynamic) between lanes and reordering of lanes when using multiple physical lanes
  - Also works with gearboxes when higher signaling rate is used on the wire
  - Also works when pairs of PCS lanes are muxed into PAM4 physical lanes.
- Many newer devices implement bit-muxing internally without skew and with known order
  - For example, a 50G/100G PAM4 SerDes co-located with the PCS always transmits the same group of PCS lanes in the same order
  - Also, PAM4 bit grouping is deterministic, and gearboxes are likely to keep bit-pairs together.
  - In the following slides, muxing order and PAM4 pairing are shown as if bit-muxing is co-located with the PCS. The results with other choices would be equivalent.

# 200GBASE-R with 50G/lane signaling

## z=8 PCSLs, 2-way CW interleaving (color indicates CW A or B)

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|---|---|---|---|---|---|---|---|
| 89 | 99 | 109 | 119 | 129 | 139 | 149 | 159 |
| 88 | 98 | 108 | 118 | 128 | 138 | 148 | 158 |
| 87 | 97 | 107 | 117 | 127 | 137 | 147 | 157 |
| 86 | 96 | 106 | 116 | 126 | 136 | 146 | 156 |
| 85 | 95 | 105 | 115 | 125 | 135 | 145 | 155 |
| 84 | 94 | 104 | 114 | 124 | 134 | 144 | 154 |
| 83 | 93 | 103 | 113 | 123 | 133 | 143 | 153 |
| 82 | 92 | 102 | 112 | 122 | 132 | 142 | 152 |
| 81 | 91 | 101 | 111 | 121 | 131 | 141 | 151 |
| 80 | 90 | 100 | 110 | 120 | 130 | 140 | 150 |
| 9 | 19 | 29 | 39 | 49 | 59 | 69 | 79 |
| 8 | 18 | 28 | 38 | 48 | 58 | 68 | 78 |
| 7 | 17 | 27 | 37 | 47 | 57 | 67 | 77 |
| 6 | 16 | 26 | 36 | 46 | 56 | 66 | 76 |
| 5 | 15 | 25 | 35 | 45 | 55 | 65 | 75 |
| 4 | 14 | 24 | 34 | 44 | 54 | 64 | 74 |
| 3 | 13 | 23 | 33 | 43 | 53 | 63 | 73 |
| 2 | 12 | 22 | 32 | 42 | 52 | 62 | 72 |
| 1 | 11 | 21 | 31 | 41 | 51 | 61 | 71 |
| 0 | 10 | 20 | 30 | 40 | 50 | 60 | 70 |

When bit-muxed to n=4 physical lanes (50G PAM4)

| | 0 | 1 | 2 | 3 |
|---|---|---|---|---|
| 41 | (81,91) | (101,111) | (121,131) | (141,151) |
| 40 | (80,90) | (100,110) | (120,130) | (140,150) |
| 39 | (9,19) | (29,39) | (49,59) | (69,79) |
| 38 | (8,18) | (28,38) | (48,58) | (68,78) |
| | ⋮ | ⋮ | ⋮ | ⋮ |
| 3 | (3,13) | (23,33) | (43,53) | (63,73) |
| 2 | (2,12) | (22,32) | (42,52) | (62,72) |
| 1 | (1,11) | (21,31) | (41,51) | (61,71) |
| 0 | (0,10) | (20,30) | (40,50) | (60,70) |

**Without precoding:**

- A single error corrupts one FEC symbol
- Bursts with $2 \leq L \leq 11$ corrupt up to 2 FEC symbols per CW (2 if FEC symbol boundary is crossed)
- Bursts with $11 < L \leq 20$ corrupt up to 3 FEC symbols per CW…

**With precoding:**

- A single error corrupts:
  - Two FEC symbols in the same CW w.p. $\frac{1}{2} \cdot \frac{1}{10} = 0.05$
  - One FEC symbol in each CW w.p. $0.5$
  - Only one FEC symbol w.p. $1 - \frac{1}{2} - \frac{1}{20} = 0.45$
- Longer bursts corrupt 1 or 2 FEC symbols in the “starting” CW, and 0 or 1 in the other CW.

# 200GBASE-R with 100G/lane signaling

## z=8 PCSLs, 2-way CW interleaving (color indicates CW A or B)

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|---|---|---|---|---|---|---|---|
| 89 | 99 | 109 | 119 | 129 | 139 | 149 | 159 |
| 88 | 98 | 108 | 118 | 128 | 138 | 148 | 158 |
| 87 | 97 | 107 | 117 | 127 | 137 | 147 | 157 |
| 86 | 96 | 106 | 116 | 126 | 136 | 146 | 156 |
| 85 | 95 | 105 | 115 | 125 | 135 | 145 | 155 |
| 84 | 94 | 104 | 114 | 124 | 134 | 144 | 154 |
| 83 | 93 | 103 | 113 | 123 | 133 | 143 | 153 |
| 82 | 92 | 102 | 112 | 122 | 132 | 142 | 152 |
| 81 | 91 | 101 | 111 | 121 | 131 | 141 | 151 |
| 80 | 90 | 100 | 110 | 120 | 130 | 140 | 150 |
| 9 | 19 | 29 | 39 | 49 | 59 | 69 | 79 |
| 8 | 18 | 28 | 38 | 48 | 58 | 68 | 78 |
| 7 | 17 | 27 | 37 | 47 | 57 | 67 | 77 |
| 6 | 16 | 26 | 36 | 46 | 56 | 66 | 76 |
| 5 | 15 | 25 | 35 | 45 | 55 | 65 | 75 |
| 4 | 14 | 24 | 34 | 44 | 54 | 64 | 74 |
| 3 | 13 | 23 | 33 | 43 | 53 | 63 | 73 |
| 2 | 12 | 22 | 32 | 42 | 52 | 62 | 72 |
| 1 | 11 | 21 | 31 | 41 | 51 | 61 | 71 |
| 0 | 10 | 20 | 30 | 40 | 50 | 60 | 70 |

When bit-muxed to n=2 physical lanes (100G PAM4)

| | 0 | 1 |
|---|---|---|
| 43 | (121,131) | (141,151) |
| 42 | (81,91) | (101,111) |
| 41 | (120,130) | (140,150) |
| 40 | (80,90) | (100,110) |
| 39 | (49,59) | (69,79) |
| 38 | (9,19) | (29,39) |
| 37 | (48,58) | (68,78) |
| 36 | (8,18) | (28,38) |
| | ⋮ | ⋮ |
| 5 | (42,52) | (62,72) |
| 4 | (2,12) | (22,32) |
| 3 | (41,51) | (61,71) |
| 2 | (1,11) | (21,31) |
| 1 | (40,50) | (60,70) |
| 0 | (0,10) | (20,30) |

**Without precoding:**

- A single error corrupts one FEC symbol
- Bursts with $L \leq 8$ can corrupt L FEC symbols, up to 4 symbols in each CW (considering FEC symbol boundary crossing)
- Bursts with $8 < L \leq 20$ corrupt up to 4 FEC symbols per CW
- To corrupt $n > 4$ symbols in one CW, $L > 20 \left\lfloor \frac{n}{4} \right\rfloor$ is required

**With precoding:**

- A single error corrupts:
  - Two FEC symbols in the same CW w.p. $\frac{1}{2}$
  - One FEC symbol in each CW w.p. $\frac{1}{2}$
  - The probability of corrupting only one FEC symbol in one CW is 0
- Longer bursts corrupt 1 or 2 FEC symbols in the "starting" CW, and 0 or 1 in the other CW.

# 200GBASE-R with 200G/lane signaling

## z=8 PCSLs, 2-way CW interleaving (color indicates CW A or B)

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|---|---|---|---|---|---|---|---|
| 89 | 99 | 109 | 119 | 129 | 139 | 149 | 159 |
| 88 | 98 | 108 | 118 | 128 | 138 | 148 | 158 |
| 87 | 97 | 107 | 117 | 127 | 137 | 147 | 157 |
| 86 | 96 | 106 | 116 | 126 | 136 | 146 | 156 |
| 85 | 95 | 105 | 115 | 125 | 135 | 145 | 155 |
| 84 | 94 | 104 | 114 | 124 | 134 | 144 | 154 |
| 83 | 93 | 103 | 113 | 123 | 133 | 143 | 153 |
| 82 | 92 | 102 | 112 | 122 | 132 | 142 | 152 |
| 81 | 91 | 101 | 111 | 121 | 131 | 141 | 151 |
| 80 | 90 | 100 | 110 | 120 | 130 | 140 | 150 |
| 9 | 19 | 29 | 39 | 49 | 59 | 69 | 79 |
| 8 | 18 | 28 | 38 | 48 | 58 | 68 | 78 |
| 7 | 17 | 27 | 37 | 47 | 57 | 67 | 77 |
| 6 | 16 | 26 | 36 | 46 | 56 | 66 | 76 |
| 5 | 15 | 25 | 35 | 45 | 55 | 65 | 75 |
| 4 | 14 | 24 | 34 | 44 | 54 | 64 | 74 |
| 3 | 13 | 23 | 33 | 43 | 53 | 63 | 73 |
| 2 | 12 | 22 | 32 | 42 | 52 | 62 | 72 |
| 1 | 11 | 21 | 31 | 41 | 51 | 61 | 71 |
| 0 | 10 | 20 | 30 | 40 | 50 | 60 | 70 |

When bit-muxed to n=1 physical lanes (200G PAM4)

**Without precoding:**

- Single error corrupts one FEC symbol
- Bursts with $L \leq 16$ can corrupt L FEC symbols, up to 8 symbols in each CW
- Bursts with $16 < L \leq 40$ corrupt up to 8 FEC symbols per CW
- To corrupt $n > 8$ symbols in one CW, $L > 40 \left\lfloor \frac{n}{8} \right\rfloor$ is required

**With precoding:**

- For a single error:
  - The probability of corrupting two FEC symbols is in the same CW is $\frac{1}{2}$
  - The probability of corrupting one FEC symbol in each CW is $\frac{1}{2}$
  - The probability of corrupting only one FEC symbol in one CW is 0
- Longer bursts corrupt 1 or 2 FEC symbols in the “starting” CW, and 0 or 1 in the other CW

| | 0 |
|---|---|
| 47 | (141,151) |
| 46 | (121,111) |
| 45 | (101,111) |
| 44 | (81,91) |
| 43 | (140,150) |
| 42 | (120,100) |
| 41 | (100,110) |
| 40 | (80,90) |
| 39 | (69,79) |
| 38 | (49,59) |
| 37 | (29,39) |
| 36 | (9,19) |
| 35 | (68,78) |
| 34 | (48,58) |
| 33 | (28,38) |
| 32 | (8,18) |
| | ⋮ |
| 3 | (60,70) |
| 2 | (40,50) |
| 1 | (20,30) |
| 0 | (0,10) |

# 400GBASE-R

## Essentially the same as 200G: z=16 PCSLs and 2× physical lanes

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 169 | 179 | 189 | 199 | 209 | 219 | 229 | 239 | 249 | 259 | 269 | 279 | 289 | 299 | 309 | 319 |
| 168 | 178 | 188 | 198 | 208 | 218 | 228 | 238 | 248 | 258 | 268 | 278 | 288 | 298 | 308 | 318 |
| 167 | 177 | 187 | 197 | 207 | 217 | 227 | 237 | 247 | 257 | 267 | 277 | 287 | 297 | 307 | 317 |
| 166 | 176 | 186 | 196 | 206 | 216 | 226 | 236 | 246 | 256 | 266 | 276 | 286 | 296 | 306 | 316 |
| 165 | 175 | 185 | 195 | 205 | 215 | 225 | 235 | 245 | 255 | 265 | 275 | 285 | 295 | 305 | 315 |
| 164 | 174 | 184 | 194 | 204 | 214 | 224 | 234 | 244 | 254 | 264 | 274 | 284 | 294 | 304 | 314 |
| 163 | 173 | 183 | 193 | 203 | 213 | 223 | 233 | 243 | 253 | 263 | 273 | 283 | 293 | 303 | 313 |
| 162 | 172 | 182 | 192 | 202 | 212 | 222 | 232 | 242 | 252 | 262 | 272 | 282 | 292 | 302 | 312 |
| 161 | 171 | 181 | 191 | 201 | 211 | 221 | 231 | 241 | 251 | 261 | 271 | 281 | 291 | 301 | 311 |
| 160 | 170 | 180 | 190 | 200 | 210 | 220 | 230 | 240 | 250 | 260 | 270 | 280 | 290 | 300 | 310 |
| 9 | 19 | 29 | 39 | 49 | 59 | 69 | 79 | 89 | 99 | 109 | 119 | 129 | 139 | 149 | 159 |
| 8 | 18 | 28 | 38 | 48 | 58 | 68 | 78 | 88 | 98 | 108 | 118 | 128 | 138 | 148 | 158 |
| 7 | 17 | 27 | 37 | 47 | 57 | 67 | 77 | 87 | 97 | 107 | 117 | 127 | 137 | 147 | 157 |
| 6 | 16 | 26 | 36 | 46 | 56 | 66 | 76 | 86 | 96 | 106 | 116 | 126 | 136 | 146 | 156 |
| 5 | 15 | 25 | 35 | 45 | 55 | 65 | 75 | 85 | 95 | 105 | 115 | 125 | 135 | 145 | 155 |
| 4 | 14 | 24 | 34 | 44 | 54 | 64 | 74 | 84 | 94 | 104 | 114 | 124 | 134 | 144 | 154 |
| 3 | 13 | 23 | 33 | 43 | 53 | 63 | 73 | 83 | 93 | 103 | 113 | 123 | 133 | 143 | 153 |
| 2 | 12 | 22 | 32 | 42 | 52 | 62 | 72 | 82 | 92 | 102 | 112 | 122 | 132 | 142 | 152 |
| 1 | 11 | 21 | 31 | 41 | 51 | 61 | 71 | 81 | 91 | 101 | 111 | 121 | 131 | 141 | 151 |
| 0 | 10 | 20 | 30 | 40 | 50 | 60 | 70 | 80 | 90 | 100 | 110 | 120 | 130 | 140 | 150 |

Serialized to n=8 physical lanes (50G PAM4)

| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|---|---|---|---|---|---|---|---|---|
| 11 | (161, 171) | (181, 191) | (201, 211) | (221, 231) | (241, 251) | (261, 271) | (281, 291) | (301, 311) |
| 10 | (160, 170) | (180, 190) | (200, 210) | (220, 230) | (240, 250) | (260, 270) | (280, 290) | (300, 310) |
| 9 | (9, 19) | (29, 39) | (49, 59) | (69, 79) | (89, 99) | (109, 119) | (129, 139) | (149, 159) |
| 8 | (8, 18) | (28, 38) | (48, 58) | (68, 78) | (88, 98) | (108, 118) | (128, 138) | (148, 158) |
| | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ |
| 2 | (2, 12) | (22, 32) | (42, 52) | (62, 72) | (82, 92) | (102, 112) | (122, 132) | (142, 152) |
| 1 | (1, 11) | (21, 31) | (41, 51) | (61, 71) | (81, 91) | (101, 111) | (121, 131) | (141, 151) |
| 0 | (0, 10) | (20, 30) | (40, 50) | (60, 70) | (80, 90) | (100, 110) | (120, 130) | (140, 150) |

Serialized to n=4 physical lanes (100G PAM4)

| | 0 | 1 | 2 | 3 |
|---|---|---|---|---|
| 23 | (241, 251) | (261, 271) | (281, 291) | (301, 311) |
| 22 | (161, 171) | (181, 191) | (201, 211) | (221, 231) |
| 21 | (240, 250) | (260, 270) | (280, 290) | (300, 310) |
| 20 | (160, 170) | (180, 190) | (200, 210) | (220, 230) |
| 19 | (89, 99) | (109, 119) | (129, 139) | (149, 159) |
| 18 | (9, 19) | (29, 39) | (49, 59) | (69, 79) |
| 17 | (88, 98) | (108, 118) | (128, 138) | (148, 158) |
| 16 | (8, 18) | (28, 38) | (48, 58) | (68, 78) |
| | ⋮ | ⋮ | ⋮ | ⋮ |
| 2 | (1, 11) | (21, 31) | (41, 51) | (61, 71) |
| 1 | (80, 90) | (100, 110) | (120, 130) | (140, 150) |
| 0 | (0, 10) | (20, 30) | (40, 50) | (60, 70) |

# 400GBASE-R

Essentially the same as 200G: z=16 PCSLs and 2× physical lanes

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 169 | 179 | 189 | 199 | 209 | 219 | 229 | 239 | 249 | 259 | 269 | 279 | 289 | 299 | 309 | 319 |
| 168 | 178 | 188 | 198 | 208 | 218 | 228 | 238 | 248 | 258 | 268 | 278 | 288 | 298 | 308 | 318 |
| 167 | 177 | 187 | 197 | 207 | 217 | 227 | 237 | 247 | 257 | 267 | 277 | 287 | 297 | 307 | 317 |
| 166 | 176 | 186 | 196 | 206 | 216 | 226 | 236 | 246 | 256 | 266 | 276 | 286 | 296 | 306 | 316 |
| 165 | 175 | 185 | 195 | 205 | 215 | 225 | 235 | 245 | 255 | 265 | 275 | 285 | 295 | 305 | 315 |
| 164 | 174 | 184 | 194 | 204 | 214 | 224 | 234 | 244 | 254 | 264 | 274 | 284 | 294 | 304 | 314 |
| 163 | 173 | 183 | 193 | 203 | 213 | 223 | 233 | 243 | 253 | 263 | 273 | 283 | 293 | 303 | 313 |
| 162 | 172 | 182 | 192 | 202 | 212 | 222 | 232 | 242 | 252 | 262 | 272 | 282 | 292 | 302 | 312 |
| 161 | 171 | 181 | 191 | 201 | 211 | 221 | 231 | 241 | 251 | 261 | 271 | 281 | 291 | 301 | 311 |
| 160 | 170 | 180 | 190 | 200 | 210 | 220 | 230 | 240 | 250 | 260 | 270 | 280 | 290 | 300 | 310 |
| 9 | 19 | 29 | 39 | 49 | 59 | 69 | 79 | 89 | 99 | 109 | 119 | 129 | 139 | 149 | 159 |
| 8 | 18 | 28 | 38 | 48 | 58 | 68 | 78 | 88 | 98 | 108 | 118 | 128 | 138 | 148 | 158 |
| 7 | 17 | 27 | 37 | 47 | 57 | 67 | 77 | 87 | 97 | 107 | 117 | 127 | 137 | 147 | 157 |
| 6 | 16 | 26 | 36 | 46 | 56 | 66 | 76 | 86 | 96 | 106 | 116 | 126 | 136 | 146 | 156 |
| 5 | 15 | 25 | 35 | 45 | 55 | 65 | 75 | 85 | 95 | 105 | 115 | 125 | 135 | 145 | 155 |
| 4 | 14 | 24 | 34 | 44 | 54 | 64 | 74 | 84 | 94 | 104 | 114 | 124 | 134 | 144 | 154 |
| 3 | 13 | 23 | 33 | 43 | 53 | 63 | 73 | 83 | 93 | 103 | 113 | 123 | 133 | 143 | 153 |
| 2 | 12 | 22 | 32 | 42 | 52 | 62 | 72 | 82 | 92 | 102 | 112 | 122 | 132 | 142 | 152 |
| 1 | 11 | 21 | 31 | 41 | 51 | 61 | 71 | 81 | 91 | 101 | 111 | 121 | 131 | 141 | 151 |
| 0 | 10 | 20 | 30 | 40 | 50 | 60 | 70 | 80 | 90 | 100 | 110 | 120 | 130 | 140 | 150 |

Serialized to n=2 physical lanes (200G PAM4)

| | 0 | 1 |
|---|---|---|
| 47 | (281, 291) | (301, 311) |
| 46 | (241, 251) | (261, 271) |
| 45 | (201, 211) | (221, 231) |
| 44 | (161, 171) | (181, 191) |
| 43 | (280, 290) | (300, 310) |
| 42 | (240, 250) | (260, 270) |
| 41 | (200, 210) | (220, 230) |
| 40 | (160, 170) | (180, 190) |
| 39 | (129, 139) | (149, 159) |
| 38 | (89, 99) | (109, 119) |
| 37 | (49, 59) | (69, 79) |
| 36 | (9, 19) | (29, 39) |
| 35 | (128, 138) | (148, 158) |
| 34 | (88, 98) | (108, 118) |
| 33 | (48, 58) | (68, 78) |
| 32 | (8, 18) | (28, 38) |
| | | |
| 2 | (80, 90) | (100, 110) |
| 1 | (40, 50) | (60, 70) |
| 0 | (0, 10) | (20, 30) |

# 800GBASE-R (assumed)

Assuming two z=16 PCSs, with a similar analysis as for 400G:

- If each PCS uses a separate set of physical lanes – same results as in 400G (both 100G/lane and 200G/lane)
- If the PMA bit-muxes all 32 PCSLs across physical lanes (4-way codeword interleaving) on each, the results are still essentially the same
  - Higher interleaving helps with long bursts (>20 UI with 100G/lane)
  - However, shorter bursts can still cause several FEC symbol errors in the same CW (up to 4+4 with 100G/lane, up to 8+8 with 200G/lane)

# Summary

Number of FEC symbols that can be corrupted by a single burst of length L

| Physical lanes and Signaling rate | Without precoding | | With precoding | |
|---|---|---|---|---|
| | L=1 | L=16 | L=1 | L=16 |
| 8×25G NRZ | 1 | Up to 2 per codeword | Up to 1 per codeword | Up to 1 per codeword |
| 4×50G PAM4 or 8×50G PAM4 | 1 | Up to 3 per codeword | Up to 2 per codeword (usually 1+0 or 1+1) | Up to 2 per codeword (2+0, 1+1 or 1+0) |
| 2×100G PAM4, 4×100G PAM4, or 8×100G PAM4 | 1 | Up to 4 per codeword | Up to 2 per codeword (always 2+0 or 1+1) | Up to 2 per codeword (2+0, 1+1 or 1+0) |
| 1×200G PAM4, 2×200G PAM4, or 4×200G PAM4 | 1 | Up to 8 per codeword | Up to 2 per codeword (always 2+0 or 1+1) | Up to 2 per codeword (2+0, 1+1 or 1+0) |

The effect of long bursts without precoding becomes more significant at 200G/lane – it is similar to that of single-lane muxing with clause 91 FEC (not analyzed here).

When error propagation is expected, precoding is likely preferred over no precoding.

# Future work

- Analyze FLR vs. FEC input BER for various values of *a* (error propagation probability)
  - With and without precoding
- Should we make any changes?

# Thank you

Questions?