802.3ch PCS + FEC Design

Dr. Paul Langner

Goal: Create a PCS+FEC that transports XGMII reliably across the desired channel in the presence of burst and Gaussian noise

Current Plan of Record:

- RS1024 with a 9/8 symbol-rate bit-rate to bit-rate ratio, and some 2^mB2^m+1 transcoding (i.e. 64B65, 128B129, 256B257, 512B513)

Constraints:

- FEC + Interleaving must be designed to handle the worst case error burst and background Gaussian noise
- PCS structure should contain an integer number of XGMII frames
  - Referred to as transcoding
- Frame structure should be integrated with FEC so that an integer number of PCS frames are contained within a FEC frame so that the PCS frames can be delineated

FEC Constraints

- For PAM4 line-coding with a DFE at 10G, we need to be able to correct a 110ns burst every 100us, which translates into 1100 bits every 1,000,000 bits
- With a 9/8 rate,10-bit symbols, and a we can write the relationship between the number of symbols in a RS frame (N) and the number of payload symbols (K) as:

$$N / K = 9(2^m) / 8(2^m+1)$$

- If we need to correct a burst of 124 symbols (9/8 x 110), which requires 248 check symbols (two check symbols can correct one errored symbol), we have a second relationship between N and K as:

$$(\#\text{RS Frames in Superframe}) \times (N-K) >= 248$$

- Combining these we get:

$$N >= 248 / ((\#\text{Frames in Superframe}) \times (1 – (8(2^m+1) / 9(2^m)))) \qquad (1)$$

* The transcoding overhead must also be factored in, hence ~

FEC Constraints (continued)

- Tabulating Eq. 1 for each transcoding type, shows the effect of PCS transcoding on the minimum frame size required to correct a 110ns
  - 10-bit symbol duration is 8/9 ns = 124 symbols, thus need 248 symbols to correct

Burst Length	Transcoding	Superframe Duration (symbols)	time (ns)
248	64B65	2551	2267
248	128B129	2381	2116
248	256B257	2304	2048
248	512B513	2267	2015

➢ All transcoders provide the same service and increasing the block size increases the efficiency

FEC Delay and Complexity

- In the Tx direction, there is no delay associated with RS, as it is systematic and the check symbols are calculated "on-the-fly" and sent at the end of an RS frame
- In the Rx direction, decoder delay is roughly the number of check symbols + the 2x frame duration
  - Decoding can't start until the last symbol in the frame is received, and checking can't end until the last symbol has been passed through the decoder (Chien search)
- Complexity of the RS decoder is proportional to:
  - The field size (ours is 1024)
  - The number of check symbols
- Examples of RS decoder sizes in current process node:

Code	Datapath width in 10-bit symbols	Number of gates	Std cell area / mm2
528,514	4	126k	0.021
544,514	4	204k	0.034
576,514	4	444k	0.074

Interleaving Delay

- If our channel is not Gaussian noise limited, we can achieve the same burst-error correcting capability by interleaving many smaller length decoders (all with the same effective 9/8 rate) to create a “Superframe”
  - However this adds fixed delay to the overall system

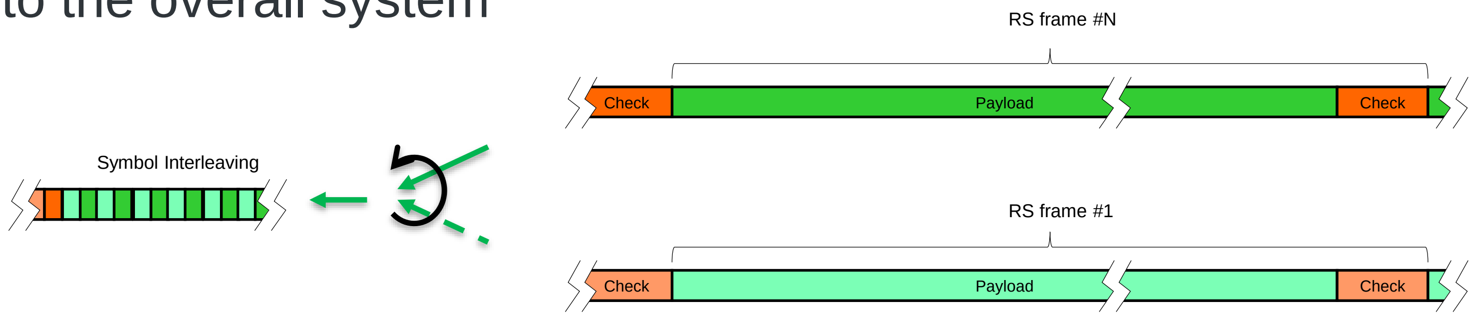

- In Tx direction, transmission cannot start until data for the last block starts, adding a delay of (N-1), where N is the amount of interleaving
- Similarly in the Rx direction, decoding cannot start until the last symbol of the first RS frame is received, which again adds (N-1) frame duration delay
- The result is that for xN interleaving, you add an additional delay of 2(N-1) frames

Interleaving Delay Versus N and Transcoding Type

- The following table shows the total delay at a 9/8 rate in RS1024 to deal with a 110ns burst (248 check symbols required) with different interleaving and different transcoding

Interleaving Depth	Transcoding	Transcode Blocks per RS Frame	N	K	Required N-K	Actual N-K	Superframe Duration (symbols)	Decode Delay (ns)	Interleave Delay (ns)	Total Delay (ns)	~# kGates
3	64B65	120	864	780	83	84	2592	1625	3072	4697	602
4	64B65	90	648	585	62	63	2592	1222	3456	4678	451
5	64B65	80	576	520	50	56	2880	1088	4096	5184	401
6	64B65	60	432	390	42	42	2592	820	3840	4660	301
7	64B65	60	432	390	36	42	3024	820	4608	5428	301
8	64B65	50	360	325	31	35	2880	685	4480	5165	251
3	128B129	60	864	774	83	90	2592	1630	3072	4702	645
4	128B129	50	720	645	62	75	2880	1361	3840	5201	537
5	128B129	40	576	516	50	60	2880	1092	4096	5188	430
6	128B129	30	432	387	42	45	2592	822	3840	4662	322
3	256B257	30	864	771	83	93	2592	1633	3072	4705	666
4	256B257	20	576	514	62	62	2304	1093	3072	4165	444
4	512B513	10	576	513	62	63	2304	1094	3072	4166	451

Notes: Interleave depth x (N-K) >= 248
N,K chosen so there is an integer number of transcoding blocks per frame
4-symbol wide decoder with total delay = (2N+2T+16) *8/9
10-bit symbol duration = 8/9ns

Conclusions

- The optimal combination of complexity versus total delay appears to be around an RS1024 (576,514) with 256B257 transcoding or RS1024(576,513) with 512B513 transcoding
  - Transcoder complexity of 256B257 and 512B513 is essentially the same, and in the “noise” relative to the RS decoder complexity (both require ~4k gates)
- Since N-K in the RS1024(576,513) code is odd, it has the same error correction capability as RS1024(576,514)

➢ Consequently, the recommendation is to use RS1024(576,514) with 512B513 transcoding + 10-bit vendor reserved symbols per frame

Thank you.

AQUANTIA®

ACCELERATING CONNECTIVITY