# De-correlating 100GBASE-KR4/CR4 training sequences between lanes

Adee Ran, Kent Lusted

Intel Corporation

# Supported by

- Andre Szczepanek, Inphi
- Dariush Dabiri, Applied Micro
- Arash Farhood, Cortina Systems
- Stephen Alie, Cortina Systems

# Background

- PMD training in 10GBASE-KR and similar PHYs specify PRBS11 as training pattern, with “random seed” requirements:
  - **72.6.10.2.6:** “The pseudo-random generator shall have a random seed at the start of the training pattern”
  - **84.7.12 and 85.7.12:** “The random seed for the training pattern described in 72.6.10.2.6 shall be different for each of the lanes”
- The reasoning for this requirement seems to be de-correlating training patterns between lanes, although this is not explicitly stated in the text.
  - Responses to comments show this intent (see backup)
- Reasoning seems even more important with multi-lane PHYs
  - Coupled routing, synchronized signaling and training periods
  - And in 802.3bj, budgets are tight
- Draft 1.2 of P802.3bj states in 92.7.12 and 93.7.12:

  “The seeds of the training pattern described in 72.6.10.2.6 should be chosen to minimize the correlation between lanes on the medium.”

# What's wrong?

- True random number generators are hard to design; reasonable implementations are pseudo-random at best
- There is no definition of required "randomness"
  - Are the seeds "320, 641, 1282, 516, 1032, 17" random enough?
    - Seems so, but these seeds create patterns which are 1 bit shift apart
  - Can "2, 3, 4, 5, 6" be used as random seed sequence?
    - Not random at all, but resulting patterns are quite distant
  - How about "x, x, x, …" where x is chosen once at random?
    - Random, but has high correlation at a fixed offset
    - What if x is actually deterministic? Effect is the same
  - Can 0 be part of the random sequence? ar1
  - No criteria to validate/qualify implementations
- Pseudo-random seed doesn't guarantee de-correlated crosstalk…

# Cross-correlation example

Here, 4 lanes use seeds which are 512 UI (¼ of the PRBS11 length) apart. Graphs show cross-correlation between the whole training frame waveforms (marker and control channel included).

Results are clearly correlated – crosstalk poses as ISI with some offset (in this case, n*512).

This is not a safe distance for CR4 (see backup).

If de-correlation is important, another way to achieve it is required.

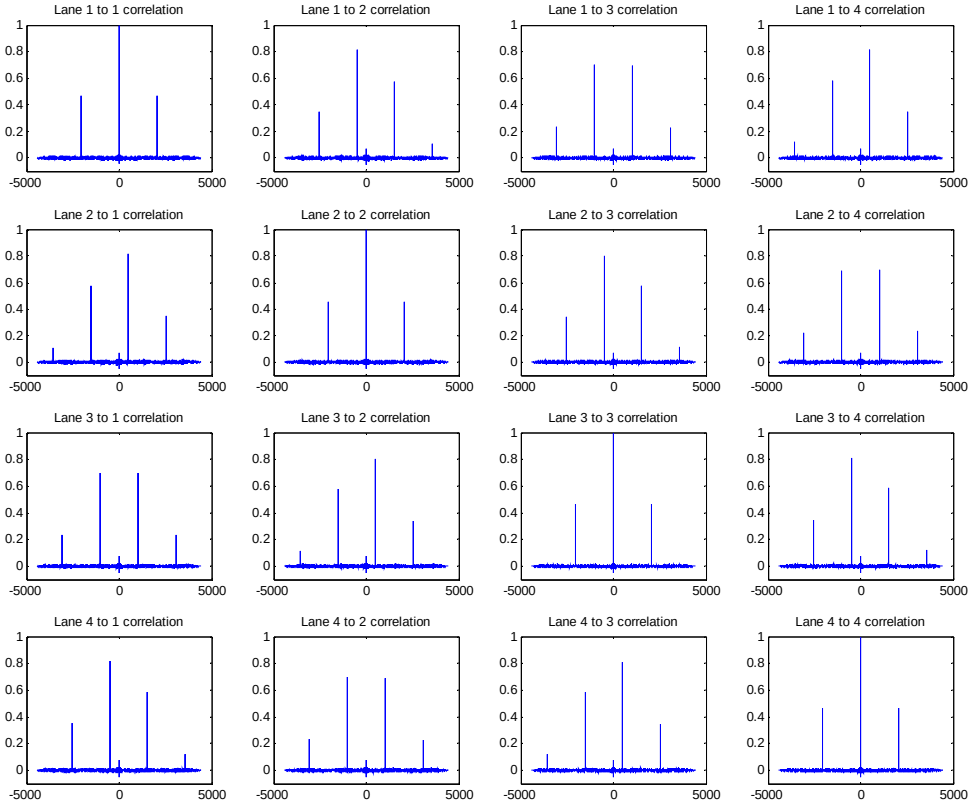

# Is it enough to require different and varying seeds?

- A reasonable implementation will have some pseudo-random “seed variation pattern”
- Such a pattern might maintain a constant offset between lanes…
  - That would be difficult to detect, as each lane seems random
  - Crosstalk will appear as “phantom ISI” at a fixed location
  - “phantom ISI” location depends on lane skew (which may be large for cables) and training frame offset between lanes (not specified at all) – so can practically be anywhere
  - See backup
- Specifying criteria for “good variation pattern” is difficult and will unnecessarily burden design and validation.

# Proposal

1. **Use different PRBS11 polynomial per lane**
   - There are 178 suitable (primitive) polynomials…
   - PRBS11 sequences with different polynomials are practically uncorrelated
   - Adding two “0” bits as in 72.6.10.2.6 maintains DC balance and rich spectrum
   - Incurs some implementation complexity (relatively mild)
2. **Specify initial output of each PRBS**
   - Values taken from Galois implementation with “1” for all initial states
   - Locations of the longest run of “1” bits are separated

# Proposal details

| Lane | Polynomial | Initial output (left to right) |
|---|---|---|
| 0 | $x^{11}+x^{10}+x^6+x^5+1$ | 11111011111100011100101100111110 |
| 1 | $x^{11}+x^9+x^6+x^5+1$ | 11111011101100011110011001100101 |
| 2 | $x^{11}+x^8+x^6+x^4+1$ | 11110011111111011010111001000110 |
| 3 | $x^{11}+x^7+x^6+x^4+1$ | 11110010111111111010010001101011 |

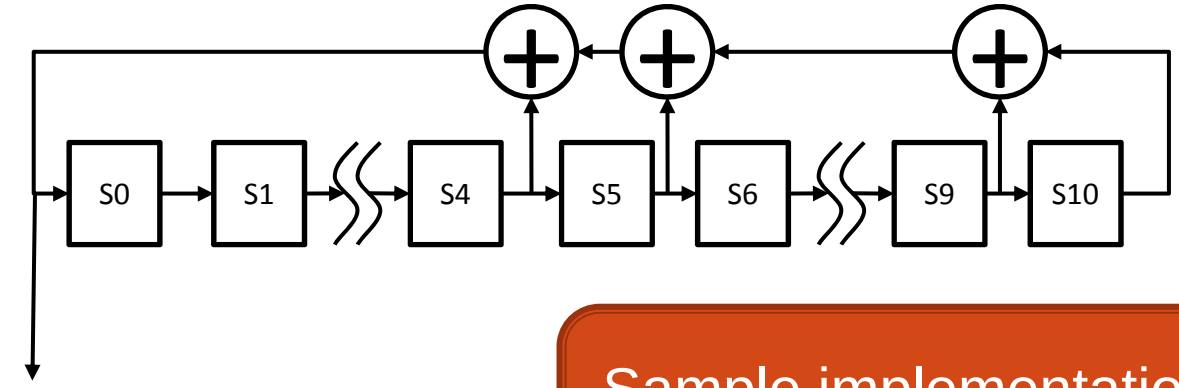


Sample implementation (lane 0)

# Cross-correlation with proposed PRBS polynomials

Here each lane uses a different PRBS11 polynomial with the proposed initial values.

Marker and control channels with logical “0” cells are included.

The generated sequences are uncorrelated.

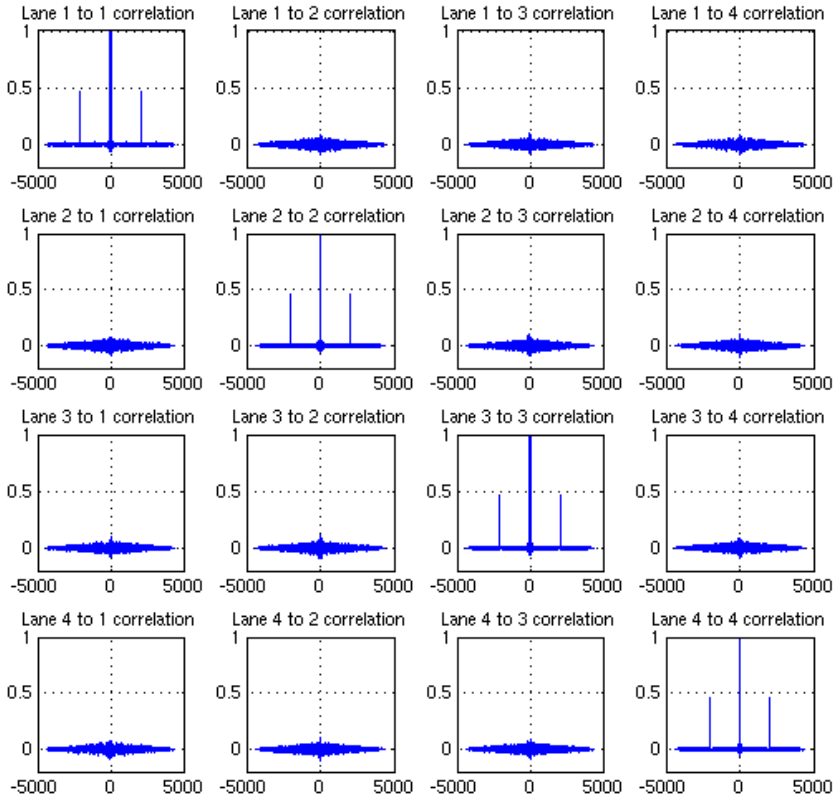

# Summary

- Current requirements
  - Are not clearly verifiable
  - Do not guarantee de-correlation of lanes
- Remedy is completely specified (so easily verifiable) and solves the problem
- Proposed for both 100GBASE-KR4 and 100GBASE-CR4. ar2

# Backup

# Comment #5 on 802.3ap D2.3

*C/* **72** *SC* **72.6.10.2.6** *P* **109** *L* **21** # 5

Andre, Szczepanek Texas Instruments

*Comment Type* **TR** *Comment Status* **A**

The problem highlighted by Comment #130 on the previous draft regarding aligned training patterns is a real problem that must be addressed, however the solution implemented in the current draft is inappropriate.

1) Random seeding of the PRBS must be mandated (Whatever PRBS we use)

2) The change from PRBS11 to PRBS58 is unnecessary and detrimental

A PRBS58 sequence has a cycle time of 1 year at 10Gbps !.
With random initialization we have no guarantee of DC-Balance except over extremely long time scales. We went to a lot of trouble to ensure DC balance in the choice of both our previous training sequences, but now we have changed to a sequence with completely unknown DC balance during any reasonable training time.

Also the ability of the equalizer to converge will be very dependant on the section of PRBS58 sequence sent. With such a long sequence some sections of the sequence may have very little useful timing information for the equalizer to use. The time taken for equalizer convergence will be unpredictable and unrepeatable. The convergence point could also be off for the real traffic that the link will carry meaning the TX remains sub-optimal and could even stay sub-optimal if re-trained.

*SuggestedRemedy*

Return to the previous training sequence of two PRBS11 cycles plus two zero bits, but mandate random seeding of the PRBS11 register before the first training frame.
Subsequent frames can either use a rolling PRBS11 (that continues to shift through the 2 zero bits, frame marker and control channel), or re-use the same initial seed.

*Response* *Response Status* **W**

ACCEPT.

Reference #37
Return to the previous training sequence of two PRBS11 cycles plus two zero bits, but mandate random seeding of the PRBS11 register before every training frame.

# Comment #10 on 802.3ap D2.4

Cl 72 SC 6.10.2.6 P 105 L 51 # 10

Andre, Szczepanek Texas Instruments

*Comment Type* **T** *Comment Status* **R** *Complete*

""The pseudo-random generator shall have a random seed at the start of the training pattern"

My reading of this is that the generator must be reseeded for every training pattern, and it is not acceptable to free-run the generator between sucessive patterns.
Was this what was agreed ?.

*SuggestedRemedy*

*Response* *Response Status* **C**

REJECT.

It is intended that a random seed be used for each training frame.

The intent of the random seed is to make synchronization between victim and aggressors les: likely, thereby decorrelating crosstalk. A free running PRBS makes synchronization between victim and aggressors more likely.

Intent was not captured by the text; some implementations use constant or slowly-varying seeds instead

# Does correlation at a large offset matter?

- With constant seed selection of PRBS11, an offset of 512 UI can be created between lanes. This is larger than a reasonable equalizer length. Is it safe?
- This offset is created at the transmitter, but lane skew changes the offset seen by the receiver
- Clause 92 skew limit is still TBD; clause 85 allows 80 ns between SP3 and SP4 (cable only)
  - > 1900 UI at 25.7812 Gbaud
  - Should probably be reduced for clause 92
- In addition, there is no requirement in clause 92 that the four PMD lanes align their training frames
  - Any delay is possible
- In fact, no constant offset is safe.

# Should receiver identify Lanes?

- With possible lane re-ordering, any RX lane can get 4 possible PRBS sequences
- RX can identify each lane by the sequence, although it is not required to do so
  - Implementations may or may not benefit from that
- If identified, lanes can be re-aligned at the PMD
  - Valid implementation, but not required and need not be specified
- Both choices are beyond the scope of the standard.