Improvements to Training Patterns and Protocol

Adee Ran, Cisco

Kent Lusted, Intel

Intro

- In-band training was adopted in January for AUIs (segments of the link) as well as PMD-to-PMD electrical links.
- The training method will be based on the Clause 136 PMD control, which has been available for two generations of PAM4 transceivers.
- Lessons learned from usage of Clause 136 PMD control with 100 Gb/s per lane links devices give rise to some desired enhancements, for either PMD or PMA training.
- The proposed enhancements are quite simple but may require new design. Therefore, they should be considered early in the project.
- Enabling training over multi-segment links is not covered in this presentation; will be described in a future contribution.

Handling polarity inversion

Problem statement

- Allowing p/n swapping adds flexibility to PCB and package design
  - This can be applied either in the Tx direction or in the Rx direction.
  - Can be applied in a chip-to-chip connection (KR, C2C) and in routing to a pluggable connection (CR, C2M).
- To enable this swapping, polarity inversion capability is required in each PMA (at the input and/or output depending on where swapping occurs).
  - This is not a specified PMA function, but it is very simple to implement, and exists in many devices.
- In high-radix devices, managing the locations that need to be polarity-inverted by the PMA can be very tedious.
  - This configuration varies between boards / products.
  - With modular systems with pluggable boards, there is no flexibility at the interface.
  - Once configured correctly this is not seen by the end user – but it takes a lot of effort behind the scene!

The opportunity

- Automatic polarity correction is part of many specifications, but is absent in the BASE-R architecture.
- CR ports usually double as optical ports, with AUI-C2M between two PMAs.
- Detecting polarity inversion in large systems with the existing specifications is not practical
  - There is no known pattern that can be identified by the receiving PMA
  - Trial-and-error on multi-lane, multi-segment links is virtually impossible
- It becomes very easy when a PMA uses training, since training has known patterns.
- Allowing polarity inversion will greatly simplify system integration.

Proposed solution

- When training is used, polarity detection is possible simply by inspecting the received frame marker:
  - Sixteen “3” followed by sixteen “0” è not inverted
  - Sixteen “0” followed by sixteen “3” è inverted

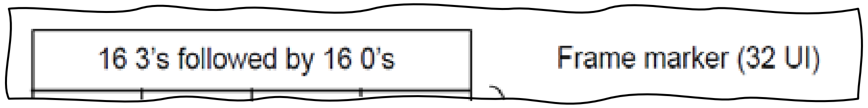

- A receiver that uses training can easily identify and correct the polarity for the data path.
  - If training is not enabled on an AUI, this “feature” will not be available, and inversion needs to be configured manually (as done today).
- **To allow polarity inversion on Tx pairs, detection and automatic correction must be a mandatory function in the Rx.**
- Proposal: **Define detection and automatic correction of polarity by the receiver as a mandatory function in the PMA/PMD training.**

Training patterns

Problem statement

- Currently the training patterns are created from two full cycles of PRBS13Q followed two 0 symbols (zero pad)
- This makes the training pattern consistent across frames
- However, it can create issues in prevalent time-interleaved (polyphase) ADC implementations…
  - For example, with a 64-phase ADC, the pattern seen on each phase of the ADC repeats itself every two frames (it has a period of 16672/64*2=521 samples)
  - These patterns are not PRBS of any kind, and are unfriendly for calibration/adaptation algorithms
  - Notably, these patterns are very unbalanced (see next slide)
- The training pattern spectrum and statistics are not representative of "mission" data
  - Once data mode is entered, the statistics change considerably
- Moving to 200G/lane, accurate calibration and training will be more important…
  - **We should get this fixed!**

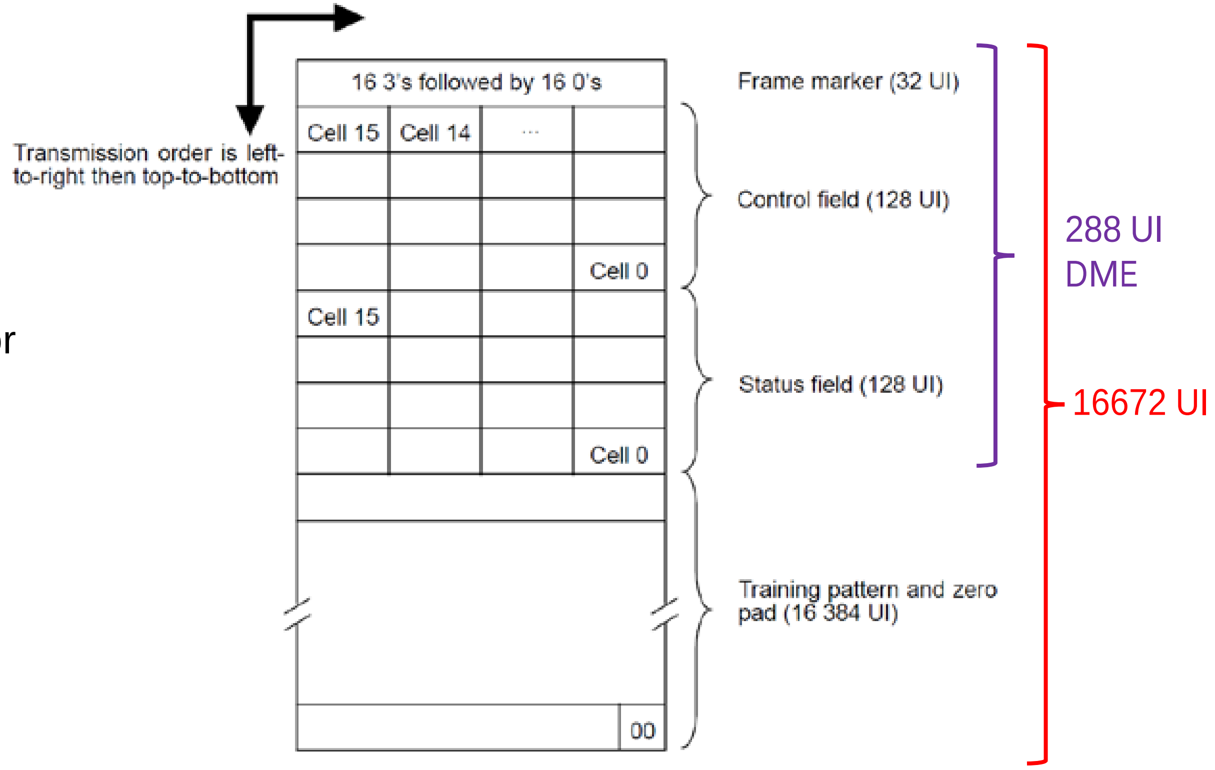

Figure 136–3—Training frame structure

Means of the clause 136 training patterns on each phase of a 64-UI subsampled pattern

Full training frame

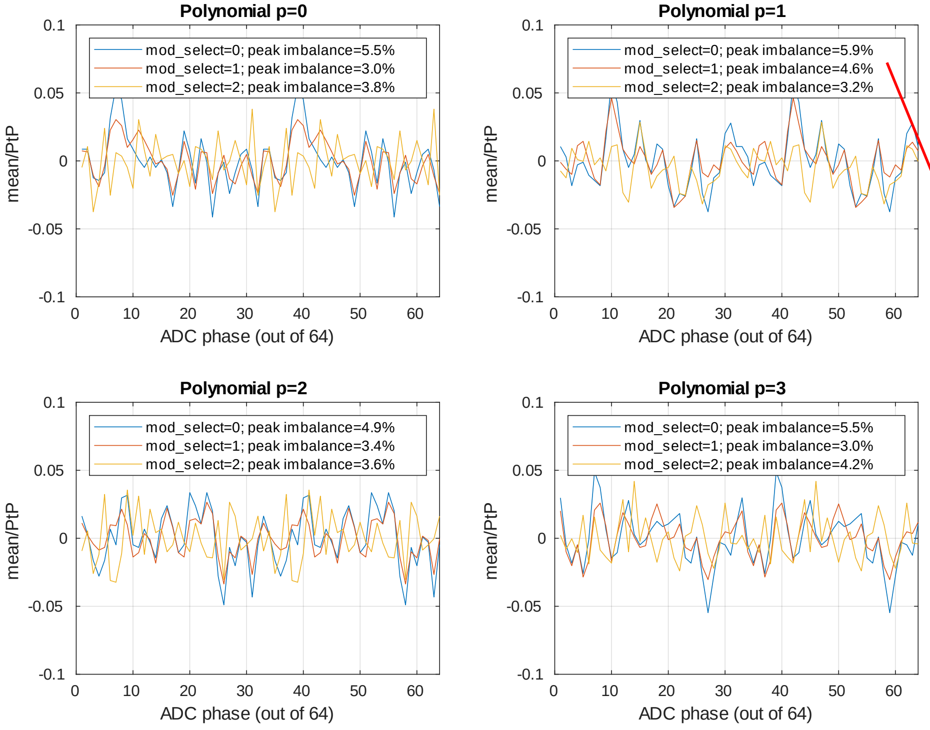

Training pattern only

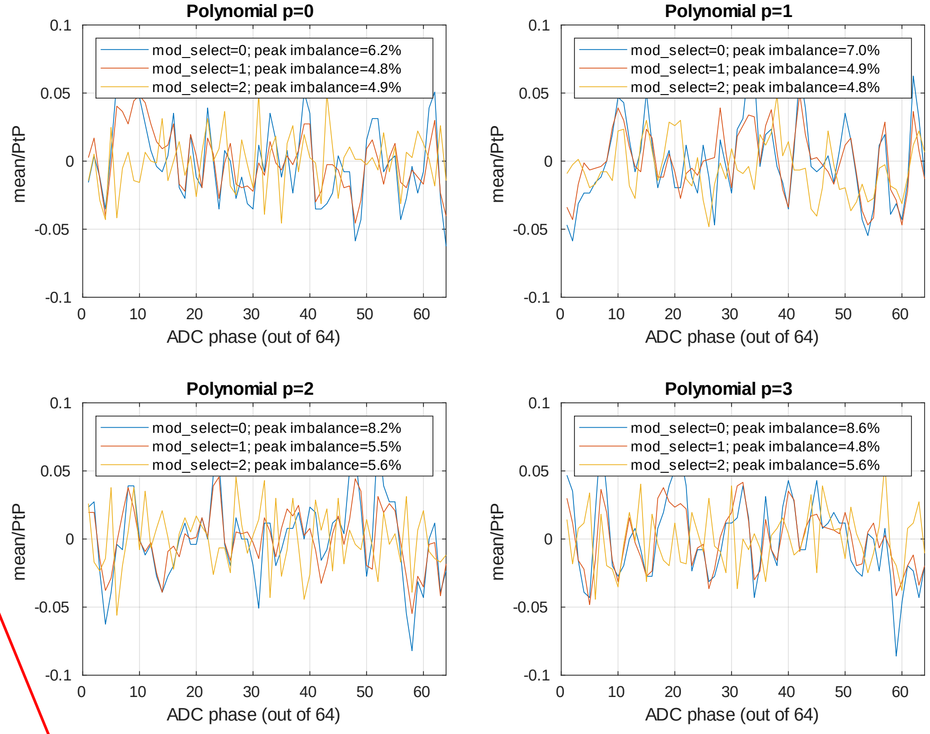

Worst case is offset is ~6% of PtP

mod_select: 0=PAM2, 1=PAM4, 2=PAM4 with precoding

Proposal

- Add another training sequence created from a **free-running PRBS13 generator** (with same choice of polynomials), without the zero pad symbols
  - The Marker/Control/Status DME portion periodically overrides the PRBS13 generator output (288 UI every 16672 UI – same as in clause 136)
  - The PRBS generator is not stopped or reset – decoupling the DME logic from the pattern generation
  - This creates per-phase sequences that have good properties – balanced (see next slide) and spectrally flat
- Add a similar training sequence using **free-running PRBS31** instead of PRBS13, without changing the frame structure
  - In this case, the same polynomial is used in all lanes
  - Two options when used with PAM4 encoding: with/without precoding
  - A receiver can ask for one of these patterns to possibly improve the training
- These new patterns are very simple to implement and test
  - Simpler than the current training pattern generators
  - Free-running PRBS13 and PRBS31 generators are available in most SerDes
- The existing patterns will be retained; the new patterns will be selectable in training using an extended "modulation and precoding" selector in the control and status fields.

Means of the modified training patterns (free-running PRBS13), same subsampling

Full training frame

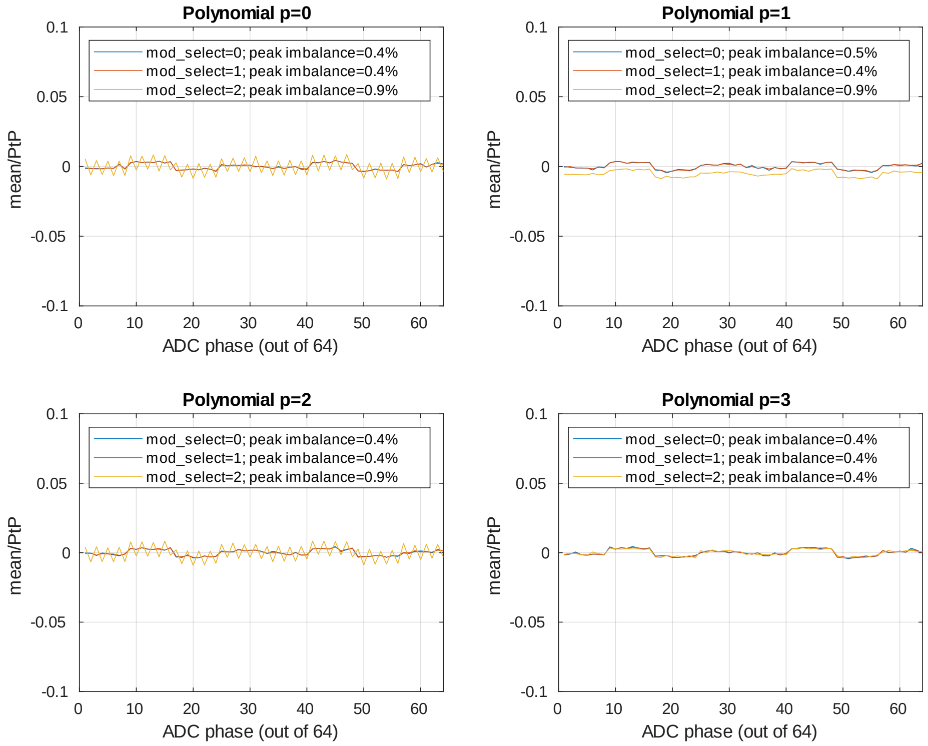

Training pattern only

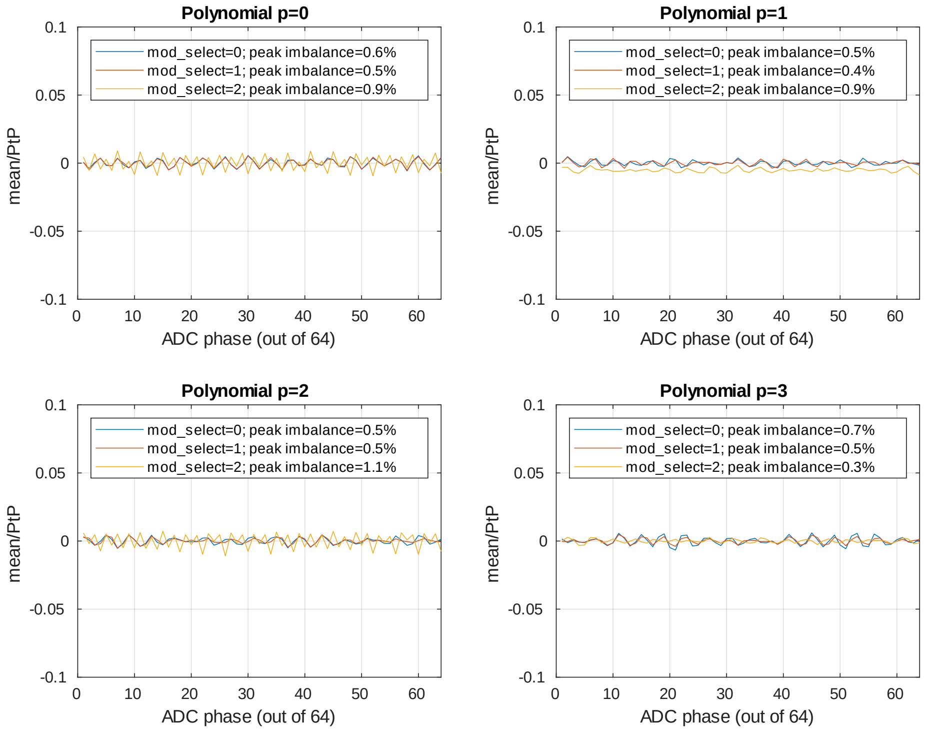

Worst case is offset is ~1% of PtP

mod_select: 0=PAM2, 1=PAM4, 2=PAM4 with precoding

Tx swing control

Problem statement

- The PMD control function of Clause 136 (50 Gb/s per lane) enables a receiver to use either **individual coefficient control** or choose one of the defined **presets**
  - Coefficient control: increment/decrement one of 4 coefficients (index from -2 to +1), encoded by 3 bits
  - Presets: 3 possible settings (index 1 to 3), encoded by 2 bits
  - See bits 4:2 and 13:2 in the control field, Table 136–9, and corresponding bits the status field, Table 136–10
- In Clause 162 (100 Gb/s per lane) the same control method is used, with one additional coefficient (index -3) and 2 additional presets
  - “Coefficient select” field kept as 3 bits, with three “reserved” values, 100, 010 and 011
  - “Initial condition” field extended to 3 bits (13:11 in Table 162–9), with two “reserved” values, 111 and 101
- With individual coefficient control, **coefficients other than the requested one are expected to be unchanged**
- Often, it is desirable to just **scale the Tx output** of the partner, without changing the equalization.
  - But the receiver does not know the exact coefficient values and step sizes of the transmitter.
  - With the current method, changing one coefficient at a time, this requires multiple requests, and complicated algorithms which can take a long time to execute.
  - The accuracy of the result can be sub-optimal due to coefficient step size quantization.

Potential improvement

- The output swing could be changed in the Tx more easily, and with better accuracy, compared to the Rx using the current training
  - With digital FFE, all coefficients can be scaled by simple calculation
  - Analog control may also be available
- The missing piece is enabling it in training protocol…
- Adding swing control to the protocol can simplify training algorithm implementation in the receiver, reduce training time, and improve performance.
- We should consider transmitter implementations that implement swing control by either digital or analog methods.

Suggested change

- Allocate one of the reserved values of the **coefficient select** field to “swing control”.
  - Suggested value is 011 (corresponding to value +3; c(+3) is not expected to be useful)
- When coefficient select is set to “swing control”:
  - “Increment” and “Decrement” in the **coefficient request** field mean: change the output swing without changing the equalization.
  - “No equalization” means: use the default output swing (not necessarily the maximum).
  - The response to these requests can be either “updated” or “coefficient at limit”.
  - Other than that, the protocol is identical to individual coefficient control (coefficient select echo, state diagrams, timing, etc.)
- An initial condition request that selects a “preset” setting shall reset the swing to the default.

Effect on electrical specifications

- Use the electrical specification “Transmitter output waveform” with the exiting linear fit method (136.9.3.1.1)
  - This method outputs a vector of coefficients c(i) relative to the “preset 1” setting
- The required effect of a swing control request is:
  - Scale (multiply/divide) c(0) by a factor k, specified with some tolerance; suggested to be between 1.03 (~0.25 dB) and 1.12 (~1 dB).
  - All other coefficients c(i) should scale by the same factor k as c(0), but are allowed deviation of ±0.025 from the expected values (e.g., due to quantization).
- The swing control range is defined as the ratio between values of c(0) in “maximum” and “minimum” conditions.
  - Suggested specification is a range of least 1.5 (~3.5 dB).

Summary of suggested enhancements

Feature	Benefits	Design considerations
Mandatory polarity detection and correction in Rx	Simplify host configuration	Required function when training is used; very simple (assuming inversion capability exists)
Improved training patterns	Better match of training to mission data	Simple design change
Swing control	Simplified training	Requires Tx swing control, simple implementation in FW with digital FFE

- The proposed changes in training frame content are summarized in the next slides.
- These changes are intended to enable backward compatibility with the existing training protocol.
  - For example, a retimer chip with 200GAUI-1 and 200GBASE-CR2 interfaces should be able to use the new training features on both interfaces, if supported by the partner.
  - This should not be specified in the standard (scope…), but it is easy to do, and can enable more applications.

Suggested changes to the control field

Current control field structure (clause 162)

9:8	Modulation and precoding request	9 8 1 1 = PAM4 with precoding 1 0 = PAM4 0 1 = Reserved 0 0 = PAM2
7:5	Reserved	Transmit as 0, ignore on receipt
4:2	Coefficient select	4 3 2 1 0 0 = Reserved 1 0 1 = c(–3) 1 1 0 = c(–2) 1 1 1 = c(–1) 0 0 0 = c(0) 0 0 1 = c(1) 0 1 x = Reserved

Proposed change

Bit 7 can be interpreted as "PRBS31 enable", and PRBS31 is always free-running.

9:7	Pattern request	9 8 7 1 1 1 = PAM4 free-running PRBS31 with precoding 1 0 1 = Reserved 0 1 1 = PAM4 free-running PRBS31 0 0 1 = PAM2 free-running PRBS31 1 1 0 = PAM4 PRBS13 with precoding 1 0 0 = PAM4 PRBS13 0 1 0 = PAM4 free-running PRBS13 0 0 0 = PAM2 PRBS13
6:5	Reserved	Transmit as 0, ignore on receipt
4:2	Coefficient select	4 3 2 1 0 0 = Reserved 1 0 1 = c(–3) 1 1 0 = c(–2) 1 1 1 = c(–1) 0 0 0 = c(0) 0 0 1 = c(1) 0 1 0 = Reserved 0 1 1 = Swing control

Bits 9:8 retain their current meaning, except for the currently reserved combination "0 1" that becomes "PAM4 free-running" (with PRBS13).

Suggested changes to the status field

Current status field structure (Clause 162)

14:12	Reserved	Transmit as 0, ignore on receipt
11:10	Modulation and precoding Status	11 10 1 1 = PAM4 with precoding 1 0 = PAM4 0 1 = Reserved 0 0 = PAM2

Proposed change

14	One	Transmit as 1
13	Reserved	Transmit as 0, ignore on receipt
12:10	Pattern status	12 11 10 1 1 1 = PAM4 free-running PRBS31 with precoding 1 0 1 = Reserved 0 1 1 = PAM4 free-running PRBS31 0 0 1 = PAM2 free-running PRBS31 1 1 0 = PAM4 PRBS13 with precoding 1 0 0 = PAM4 PRBS13 0 1 0 = PAM4 free-running PRBS13 0 0 0 = PAM2 PRBS13

Same values as in "Pattern request" in the control field.

Backup

Backward compatibility

- Not intended to be specified in the standard, since changes to 100G training are not within scope – but should be considered.
- Existing (“legacy”) devices always transmit 0 in bit 14, new devices always transmit 1 in bit 14.
- A new device can tell by the value of received bit 14 whether its partner supports the new features (training patterns, swing control) or not, and decide whether to ask for any of them.