Restricted muxing update

Vancouver, March 12th

Contributors: M. Gustlin - Xilinx, D. Ofelt - Juniper, G. Nicholl - Cisco,
K. Shrikhande - Innovium, S. Trowbridge – Nokia

Restricted muxing initial solution

- Restricted muxing as a potential fix to clock content issue presented
  - PMA 16:8 to mux natural pair of PCSLs on to 50G lanes
  - PMA 8:4 to mux any two sets of natural pairs of PCSLs on to 100G lanes
  - Avoids rogue combinations on 100G lanes
  - Blind demux on 4:8 and 8:16 since we've avoided issue on 100G lanes
- But was pointed out that assuming 802.3bs PCS/FEC can work for 100G per lane AUIs, the above does not cover cases like 400GAUI-4 + 4:8 + 8:4 PMA
  - 4:8 PMA not guaranteed to maintain natural pairs on 50G output lanes
  - Subsequent 8:4 PMA could lead to a rogue PCSL combination on 100G lanes

Restricted muxing updated with options A, B

- Additional rules for 8:4 and 4:8 PMA mux that preserves natural pair of PCSLs on 50G lanes presented, that avoid rogue combinations through 4:8 + 8:4 PMA chain
  - Two options A and B were presented, as follows
  - Option B was seen as a specific way to do PMA bit-muxing – seemed to be favored

- Option A: Keep natural pairs together in each PAM4 symbol: 0+1 in one PAM4 symbol, 2+3 in 2nd PAM4 symbol, then back to 0+1

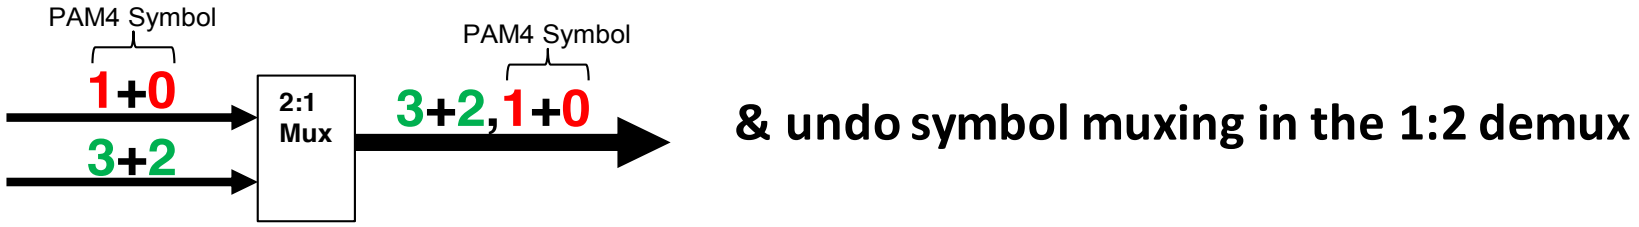

- Option B: Bit mux the natural pairs with each other: 0+2 in one PAM4 symbol, or 1+3 in the next PAM4 symbol (0+3 and 1+2)

Restricted muxing proposal

- PMA 16:8 rule as before
- PMA 8:4 rule (option B)
- PMA 4:8 rule (option B)
- PMA 16:4 is a combination of 16:8 + 8:4 and follows rules of both
- PMA 8:16 – no rules (but ...)

(Content of original submission to 802.3bs interim meeting in Vancouver)

Revisit the 8:16 PMA rule

- No rule associated with PMA 8:16? OK if followed by PCS Rx
- But was pointed out that if PMA 8:16 is followed by another PMA 16:8, e.g. 400GAUI-8 + PMA 8:16 + 400G-DR4 module with 16:4 PMA inside the module ?
  - Cannot guarantee natural pairs of PCSLs at output of PMA 8:4
- To solve the above - natural pair of PCSLs should be mapped to a natural pair of 26.5625G PMA lanes

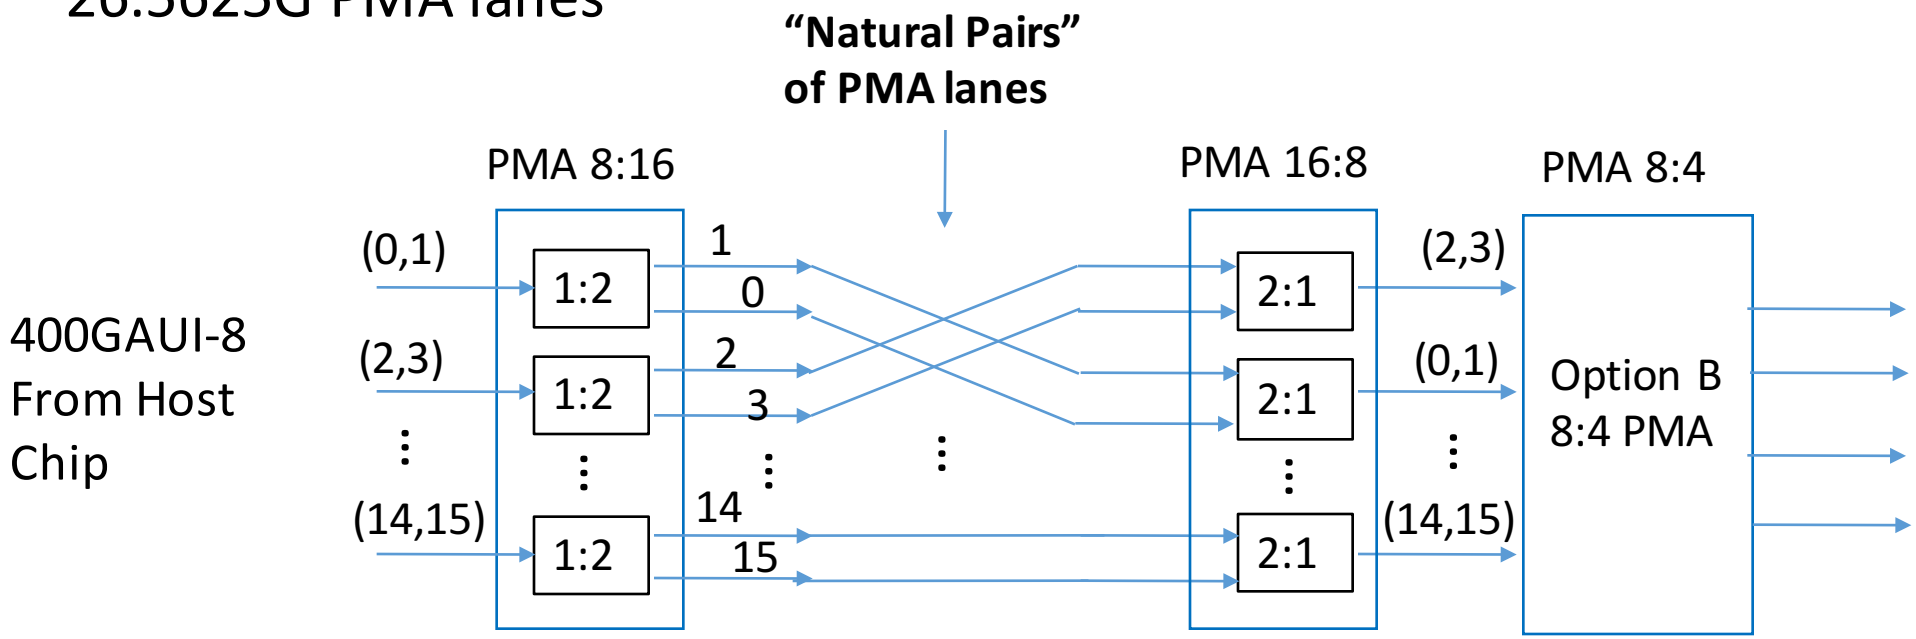

Restricted muxing proposal - revised

- PMA 16:8 rule as before
- PMA 8:4 rule (option B)
- PMA 4:8 rule (option B)
- PMA 16:4 is a combination of 16:8 + 8:4 and follows rules of both
- PMA 8:16 rule
  - Natural pair of PCSLs to fall on natural pair of 26.5625 Gb/s PMA lanes

Restricted muxing described as invariant conditions

- Conditions that are held invariant across every multiplexing stage

1. Every natural pair of physical lanes operating at 26.5625 Gb/s carries a natural pair of PCSLs. Not necessarily the same pair, and not necessarily in the same order. For example, physical lanes 7,8 could carry PCSLs 3,2.

2. Every physical lane operating at 53.125 Gb/s carries a natural pair of PCSLs, with one PCSL encoded as the A bit of each PAM4 symbol and the other PCSL encoded as the B bit of each PAM4 symbol.

3. Every physical lane operating at 106.25 Gb/s carries two natural pairs of PCSLs, with one natural pair encoded on the A bits of two consecutive PAM4 symbols, and the other natural pair encoded on the B bits of two consecutive PAM4 symbols.

(courtesy Steve Trowbridge)

Restricted muxing proposal – revised submission

- First list the conditions held invariant across every multiplexing stage
  - Prior slide
- Then define each PMA m:n function
  - PMA 16:8 rule
  - PMA 8:4 rule (option B)
  - PMA 4:8 rule (option B)
  - PMA 16:4 as a combination of 16:8 + 8:4
  - PMA 8:16 rule (new)

(Content of revised submission to 802.3bs interim meeting in Vancouver)

SR16 issue

- Was pointed out rules not sufficient to cover all SR16 use-cases
- SR16 with 400GAUI-4 can run into rogue combinations
  - Example: 400GAUI-8 + 8:16 + SR16 + 16:4 + 400GAUI-16
- Natural pairs of PMA lanes at output of the 8:16 may not fall on natural pairs at input of the 16:4

SR16 use-case: diagram

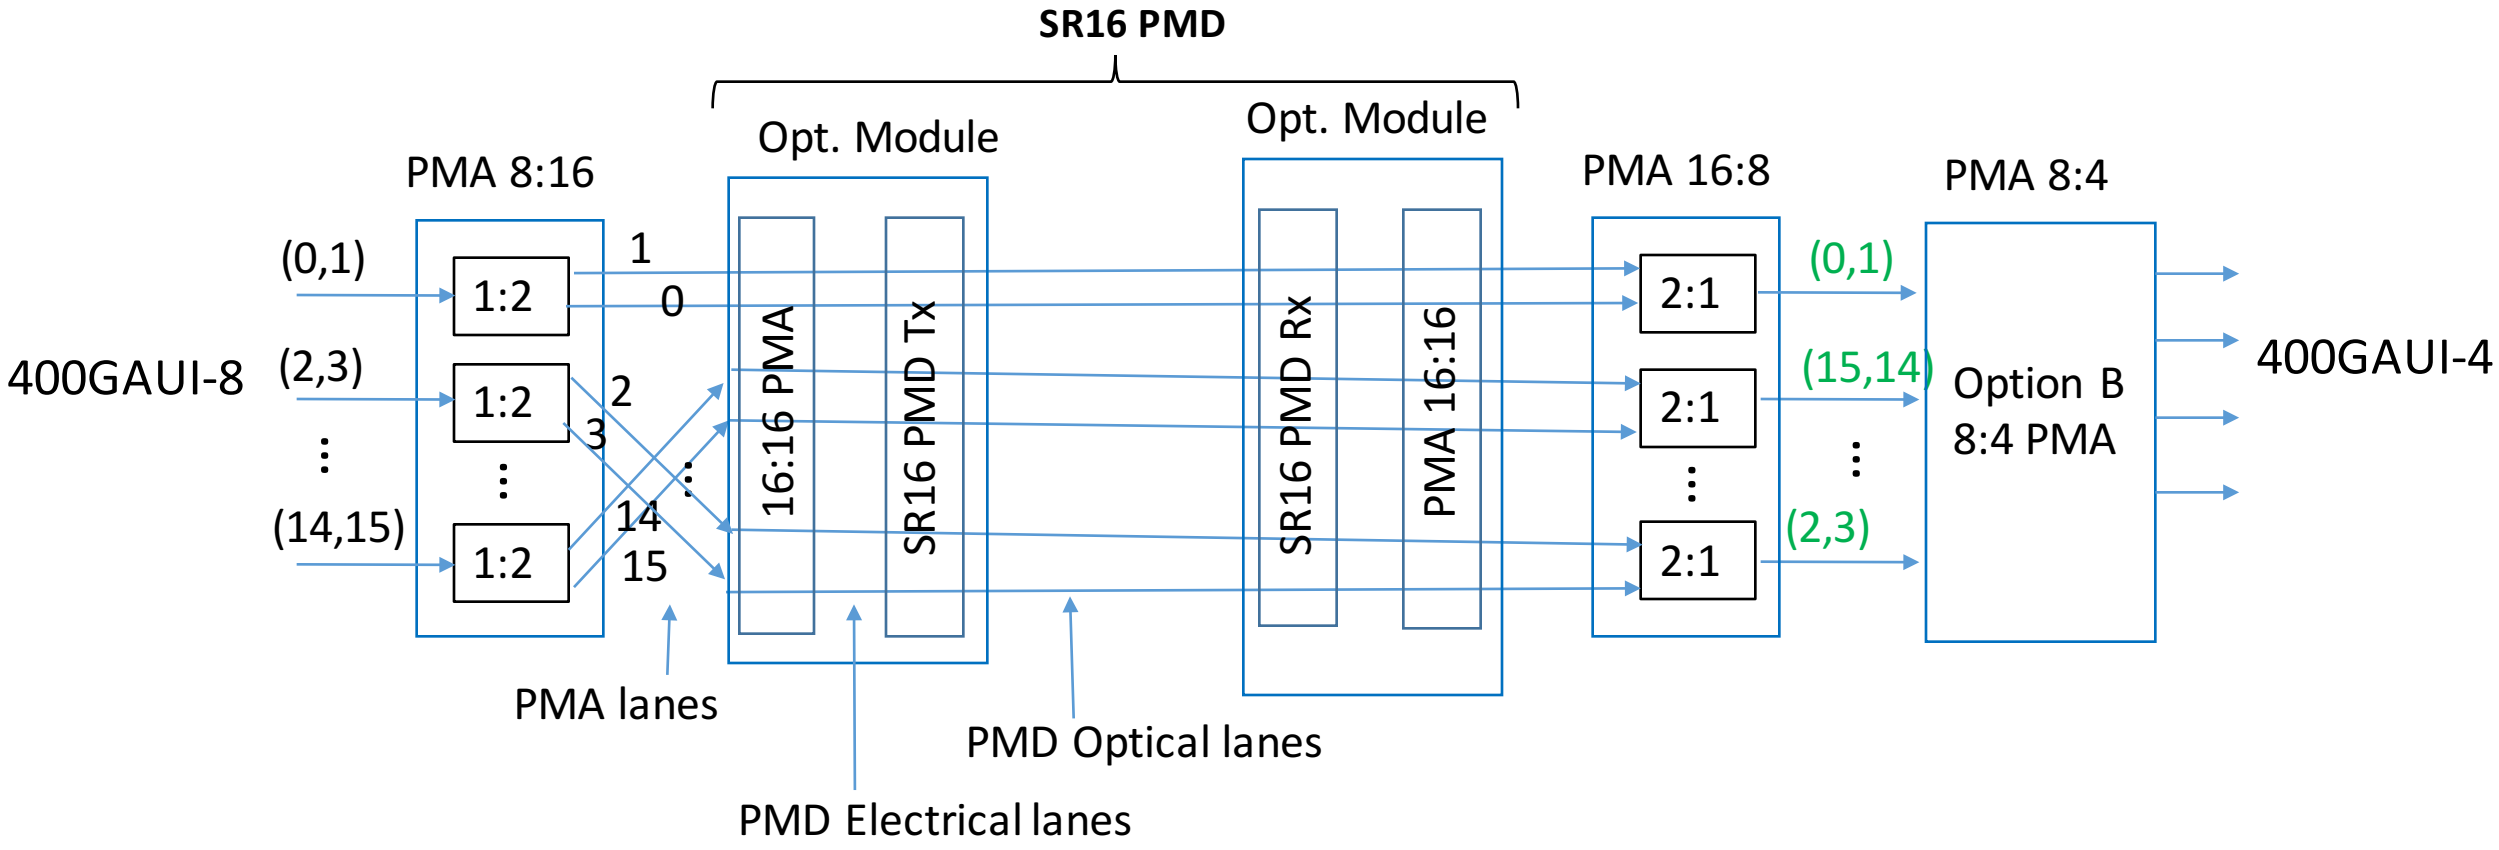

No guarantee that SR16 will map electrical lane i to optical lane i and back

SR16 use-case: potential solution

- Amend SR16 sub-clause 123.6 Lane Assignments?
  - State that bits on electrical lane i are sent on optical lane i at a transmitter, and bits received on optical lane i are sent on electrical lane i at the receiver

- Does this need any MSA definition (beyond IEEE)?

Alternate Option: 100G Slices

- Stay with the concept of generic bit muxes, but keep 100G slices together
- A 400G 16:8 mux consists of four generic 4:2 muxes from four consecutive input lanes to two consecutive output lanes. A 16:4 mux is four 4:1 muxes. Similar for 200G, where it is partitioned into two 100G slices (not using that terminology in the draft).
- Net effect is that any 100G lane is composed of four consecutive PCSLs in some order, and none of these are pathologically low clock content options.

(courtesy Steve Trowbridge)

Next steps

- Is the Restricted muxing proposal acceptable to the Task Force as a Fix to the clock content issue?

- Is there interest and consensus in developing an alternate proposal using 100G slices?
  - Different pros/cons from current proposal