Restricted PMA muxing (option B) revised: proposed changes to the Draft

Co-authors:

M. Gustlin, Xilinx
G. Nicholl, Cisco
D. Ofelt, Juniper
K. Shrikhande, Innovium

802.3bs interim meeting, Vancouver, Canada
March 12, 2017

120.2 PMA interfaces

200GBASE-R PMA variants have 8 or 4 input/output lanes, while 400GBASE-R PMA variants have 16, 8, or 4 input/output lanes, per direction. The number of input PMA lanes can be different from the number of output PMA lanes. Therefore, PMA variants for 200GBASE-R and 400GBASE-R employ one or more bit muxes to adapt from the number of input lanes to the number of output lanes in each direction.

A PMA bit mux function with m input lanes and n output lanes is shown in Figure 120-3. Conceptually, the PMA bit mux operates in one direction of transmission by demultiplexing PCSLs from m PMA input lanes and remultiplexing them into n PMA output lanes. Depending on the values of m and n, the bit mux function follows a specific set of rules that restricts the combination of PCSLs mapped to a PMA lane. For example, a 400GBASE-R PMA with 16 input lanes and 8 output lanes follows a different set of rules than a 400GBASE-R PMA with 8 input lanes and 4 output lanes. This is different from 802.3 PMA specifications such as Clause 83 where the bit mux function is generic and independent of m and n. These rules are defined in 120.5.2.1 for specific values of m and n.

Figure 120-5 provides a functional block diagram of a PMA. The parameters of a PMA include the following:

- The number of input and output lanes in each direction.
- Whether the PMA is adjacent to a physically instantiated interface (200GAUI-n or 400GAUI-n above or below).
- Whether the PMA is adjacent to the PCS or DTE XS.
- Whether the PMA is adjacent to the PMD or PHY XS.

120.5.2 Bit-level multiplexing

The PMA provides bit-level multiplexing in both the Tx and Rx directions. In the Tx direction, the function is performed among the bits received from the PMA client via the PMA:IS_UNITDATA_*i*.request primitives (for PMA client lanes *i* = 0 to p–1) with the result sent to the service interface below the PMA using the *inst*:IS_UNITDATA_*i*.request primitives (for service interface lanes *i* = 0 to q–1), referencing the functional block diagram shown in Figure 120–5. ~~The bit multiplexing behavior is illustrated in Figure 120–4.~~

The aggregate signal carried by the group of input lanes or the group of output lanes is arranged as a set of PCSLs denoted by z. The number of PCSLs z is 8 for 200GBASE-R interfaces and 16 for 400GBASE-R interfaces. The nominal bit rate ~~R~~ of each PCSL is 26.5625 Gb/s.

The 200GBASE-R and 400GBASE-R PMA introduces the concept of natural pairs of PCSLs and natural pairs of PMA lanes. Natural pairs of PCSLs are defined as PCSLs (j, j+1); j = 0,2,4,6 for 200GBASE-R and j = 0,2,4,6,8,10,12,14 for 400GBASE-R. A natural pair of PMA lanes carries bits from a natural pair of PCSLs and is defined as PMA lanes (k, k+1) where k=0,2,4,6 for 200GBASE-R and k = 0,2,4,6,8,10,12,14 for 400GBASE-R. It is not necessary for k to be equal to j.

For a PMA with m input lanes (Tx or Rx direction), each input lane carries, bit multiplexed, z/m PCSLs. Each input lane has a nominal bit rate of 26.5625 z/m Gb/s. Note that the signaling (Baud) rate is equal to the bit rate when the number of physical lanes is 8 for 200GBASE-R or 16 for 400GBASE-R (bits are sent or received on the lanes). The Baud rate is equal to half of the bit rate when the number of physical lanes is 4 for 200GBASE-R or the number of physical lanes is 8 or 4 for 400GBASE-R (PAM4 symbols are sent or received on the lanes). If necessary, PAM4 symbols are converted to pairs of bits on the input lanes and/or pairs of bits are converted to PAM4 symbols on the output lanes. If bit *x* received on an input lane belongs to a particular PCSL, the next bit of that same PCSL is received on the same input lane at bit position *x+*(z/m). The z/m PCSLs may arrive in any sequence on a given input lane.

For a PMA with n output lanes (Tx or Rx direction), each output lane carries, bit multiplexed, z/n PCSLs. Each output lane has a nominal signaling rate of 26.5625 z/n Gb/s. Each PCSL is mapped from a position in the sequence on one of the m input lanes to a position in the sequence on one of the n output lanes. If bit *x* sent on an output lane belongs to a particular PCSL, the next bit of that same PCSL is sent on the same output lane at bit position *x +* (z/n). The PMA shall maintain the chosen sequence of PCSLs on all output lanes while it is receiving a valid stream of bits on all input lanes.

~~Each PCSL received in any temporal position on an input lane is transferred into a temporal position on an output lane. As the PCS (see Clause 119) has fully flexible receive logic, an implementation is free to perform the mapping of PCSLs from input lanes to output lanes without constraint. Figure 120–6 illustrates one possible bit ordering for a 400GBASE-R 8:4 PMA bit mux. Other bit orderings are also valid.~~

Note that since the number of input lanes and output lanes for a 200GBASE-R or 400GBASE-R PMA is always a power of two, many PMAs converting between different numbers of lanes normally simply multiplex two or four input lanes to one output lane, or demultiplex two or four output lanes from one input lane. However, any PMA implementation which produces an allowable order of bits from all PCSLs on the output lanes is valid.

120.5.2.1 Bit multiplexing rules

The PMA bit muxing is specified to hold the following set of conditions invariant across every multiplexing stage:

1) Every natural pair of PMA lanes operating at 26.5625 Gb/s carries a natural pair of PCSLs. A natural pair of PMA lanes can carry any natural pair of PCSLs, and a PMA lane can carry either PCSL within a natural pair. For example, PMA lanes (6,7) could carry PCSLs (3,2).

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

3) Every PMA 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.

The PMA m:n muxing functions for different values of m and n are defined below:

A 400GBASE-R PMA with 16 input and 8 output lanes (400Gb/s PMA 16:8) and a 200GBASE-R PMA with 8 input and 4 output lanes (200Gb/s PMA 8:4) multiplexes bits from a natural pair of PCSLs or from a natural pair of PMA lanes on to an output PMA lane. This ensures that bits from a natural pair of PCSLs form the {A, B} bits of the PAM4 symbol transmitted on an output PMA lane. There is no restriction on which output lane a natural pair of PCSLs or natural pair of input PMA lanes are multiplexed on to.

For a 400GBASE-R PMA with 8 input and 4 output lanes (400Gb/s PMA 8:4), an input lane carries bits from a natural pair of PCSLs encoded as the {A, B} bits of PAM4 symbols. The PMA 8:4 bit mux takes the A bits from two input lanes to create a PAM4 symbol on the output PMA lane, and the B bits from the same two input lanes to create the subsequent PAM4 symbol on the output lane. While doing so, the A bits of the output PAM4 symbols are taken from one input PMA lane and the B bits are taken from the other input PMA lane. There is no restriction on which two input PMA lanes are used to multiplex on to an output PMA lane.

A 400GBASE-R PMA 4:8 reverses the PMA 8:4 function by taking the A bits of two consecutive PAM4 symbols received on an input lane to form the {A, B} bits of a PAM4 symbol on an output PMA lane, and it takes the B bits of the PAM4 symbols on the two input lanes to form the {A, B} of a PAM4 symbol on the second output PMA lane. While doing so, there is no restriction on which of the A bits (or B bits) on the input lanes become the {A, B} bits of the output PAM4 symbol(s).

A 400Gb/s PMA 16:4 is considered to be a combination of two bit muxes, a 16:8 followed by 8:4 mux, and follows the definition of these individual bit muxes described earlier.

~~There is no rule associated with a 400Gb/s PMA 8:16 and a 200Gb/s PMA 4:8. These PMAs can map the two bits encoded on PAM4 symbols received on an input lane to any two output lanes.~~

A 400Gb/s PMA 8:16 and a 200Gb/s PMA 4:8 demultiplexes the {A, B} bits of the PAM4 symbol on an input PMA lane to a natural pair of output PMA lanes. There is no restriction on which output lane (of the natural pair) carries the A bit and which carries the B bit. This does not require the PMA 8:16 to be PCSL aware. When the 400Gb/s PMA 8:16 is adjacent to a PCS Rx, the PMA demultiplexes the {A, B} bits on each input lane to two receive PCS lanes, and the PCS Rx can reorder the lanes as necessary.