



# Update to bit-muxing constraints for 8:8 PMA (in support of comment 83)

Gary Nicholl- Cisco

P802.3df, Interim meeting, May 2023

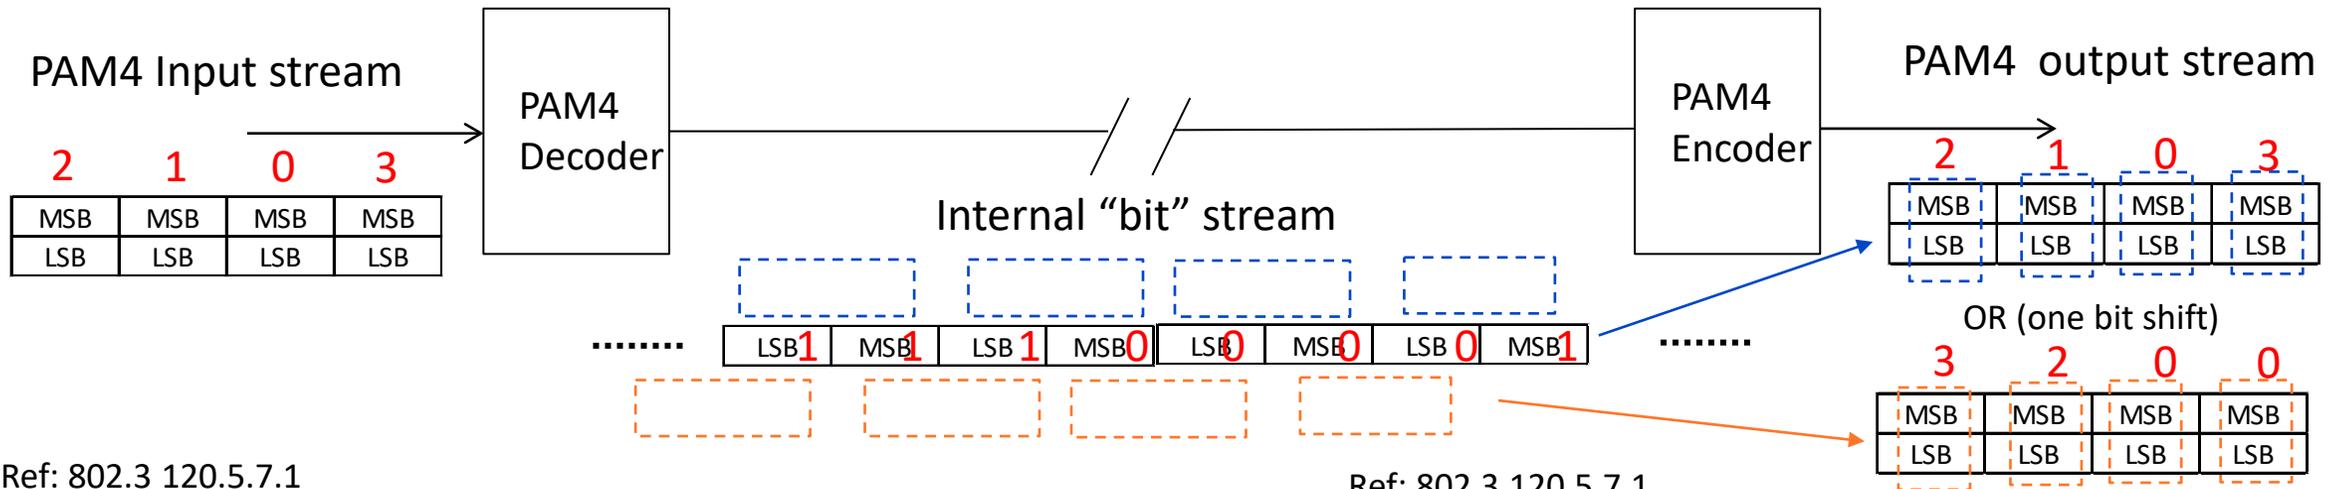
# Background

- The 8:8 PMA multiplexing rules in 173.4.2.3 (along with the multiplexing rules for 32:8 PMA and 8:32 PMA in 173.4.2.1 and 173.4.2.2 respectively) were updated based on comment #27 against D1.1 and supporting presentation ([ran\\_3df\\_01b\\_230130](#))
- As captured in slide 3 of [ran\\_3df\\_01b\\_230130](#), the primary motivation for the proposed change was to avoid the situation "*where one of two flows always gets the LSB of the PAM4 symbols*"
- The changes to the multiplexing rules for the 32:8 PMA (173.4.2.1) and PM8:32 PMA (173.4.2.2) achieve this goal.
- However, the change to the multiplexing rules for 8:8 PMA (173.4.2.3) went one step further and included an additional constraint that does not appear to be absolutely necessary.

# Background (cntd)

- The additional constraint is the requirement that "the Gray mapped PAM4 symbol sequence on the output lane is identical to the Gray mapped PAM4 symbol sequence on the input lane".
- This means that the PAM4 output must be MSB/LSB aligned to the PAM4 input. It is not clear that this would always be the case with existing 100G/lane PAM4 retimer devices, and it is something that was certainly not required for such devices (see Clause 120).
- It is also not fully consistent with the description of the PAM4 Encoding described in 173.4.7.1 (which essentially references the PAM4 encoding rules from Clause 120, which do not require PAM4 outputs to be MSB/LSB aligned to PAM4 inputs).
- This additional restriction step is not required in order to meet the intent captured in slide 3 of ran\_3df\_01b\_230130.

# Recap - 400GbE PAM4 Retimer options (Clause 120)



Ref: 802.3 120.5.7.1

For input lanes encoded as PAM4 (for 200GBASE-R, where the number of input lanes 4, or for 400GBASE-R, where the number of input lanes is 4 or 8), the PMA receive process shall map Gray-coded PAM4 symbols to pairs of bits {A, B} where A is considered to be the first bit as follows:

- 0 maps to {0, 0},
- 1 maps to {0, 1},
- 2 maps to {1, 1}, and
- 3 maps to {1, 0},

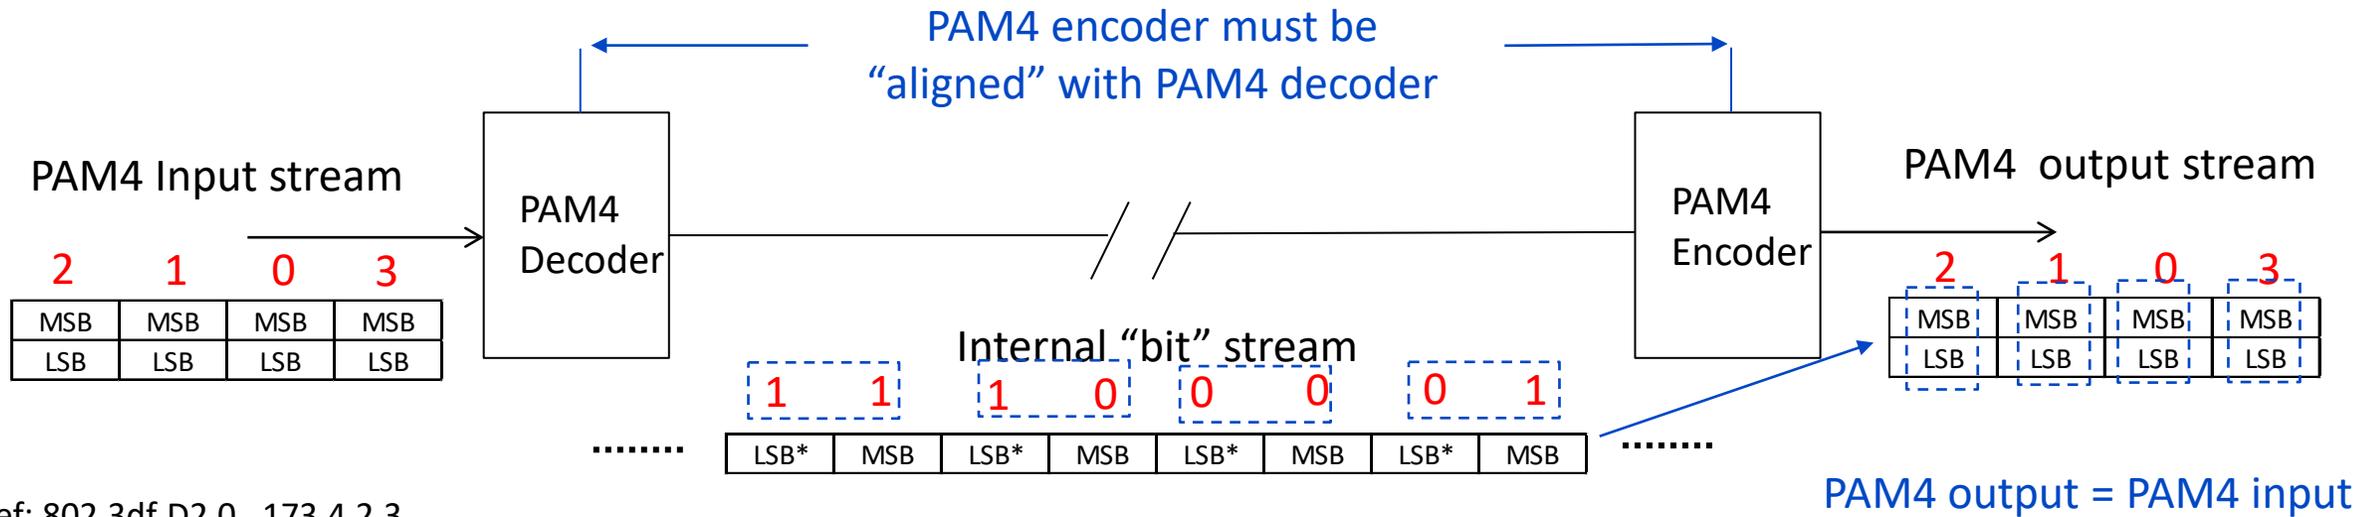
Ref: 802.3 120.5.7.1

For output lanes encoded as PAM4 (for 200GBASE-R, where the number of output lanes is 4, or for 400GBASE-R, where the number of output lanes is 4 or 8), the PMA transmit process shall map **consecutive pairs of bits {A, B}, where A is the bit arriving first**, to a Gray-coded symbol as follows:

- {0, 0} maps to 0,
- {0, 1} maps to 1,
- {1, 1} maps to 2, and
- {1, 0} maps to 3.

- No requirement for LSB/MSB alignment between PAM4 decoder and PAM4 encoder (can operate independently)
- Depending on how the PAM4 encoder samples the internal "bit" stream there could be "one bit shift" in the PAM4 output
- Note, in the encoder description above "A" is simply defined as the first bit arriving at the encoder (independent of whether it was the first bit output from decoder or not)

# Recap - PAM4 Retimer (based on 802.3cw D2.0)



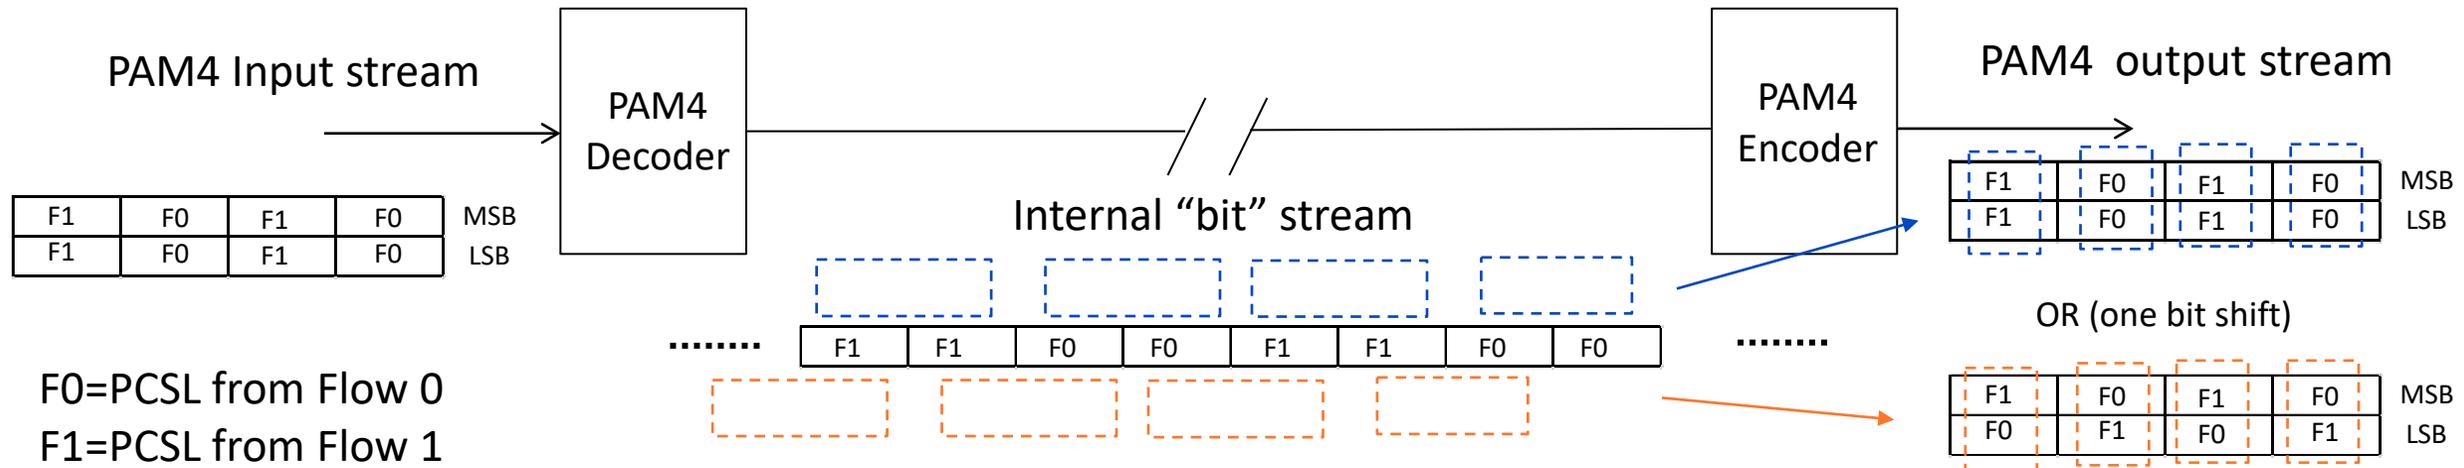
Ref: 802.3df D2.0, 173.4.2.3

In both the transmit and receive directions, the bit-level multiplexing function is identical to that specified in 120.5.2, with the following exceptions:

- The number of PCSLs is 32.
- The 4 PCSLs received on an input lane shall be mapped to an output lane such that the Gray mapped PAM4 symbol sequence on the output lane is identical to the Gray mapped PAM4 symbol sequence on the input lane, except for possible swapping of each bit pair (see 173.4.7.1).

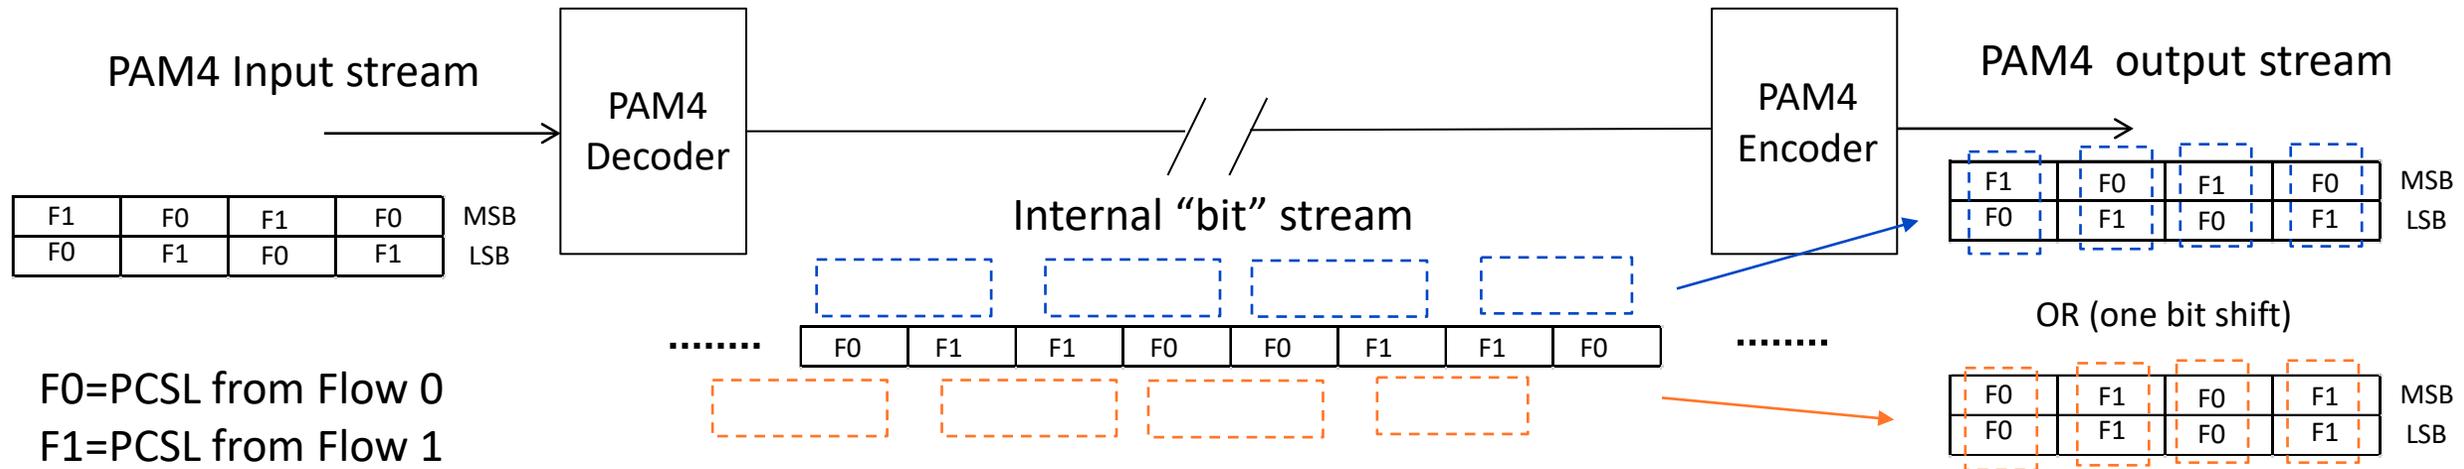
- PAM4 encoder must be PAM4 symbol aligned to the PAM4 decoder
  - how this alignment is achieved is an implementation choice (there are many possible options)
- PAM4 output symbol stream is always identical to the PAM4 input symbol stream (assuming no errors)
- But this forces an additional constraint which was not required for 400GbE PAM4 retimer chips
  - though it is possible many implementations work like this anyway ?

# Proposal - revert to “400GbE based approach” ...



- Does not require the PAM4 encoder and PAM4 decoder to be “aligned”
  - thereby consistent with 802.3bs (400GbE) implementations based on Clause 120
- Allows a “one bit shift” between PAM4 encoder and PAM4 decoder
- Produces two possible PAM4 output streams for a given PAM4 input stream, but in both cases the LSB is split equally between the two flows (meeting the requirement of slide 3 of [ran\\_3df\\_01b\\_230130](#))

# Proposal - alternate PAM4 input (/w “one bit shift”)



- Note, 802.3df D2.0 allows a “one bit shift” for the 32:8 PMA and 8:32 PMA, so it is possible to get the alternate PAM4 stream at the input to the retimer shown above
- In this case the two possible PAM4 output streams from the retimer are simply reversed from the ones shown on the previous slide
- This essentially means there is no value on the additional constraint on the 8:8PMA, as you can still get the same two possible PAM4 output streams anyway (as the 32:8 PMA and 8:32 PMA can both generate the same two options for the PAM4 stream)

# Summary

- 802.3df D2.0 added a constraint to the 8:8 PMA multiplexing rules (PAM4 output must equal PAM4 input).
- This constraint goes beyond the constraints adopted for the 32:8 and 8:32 PMA multiplexing rules and appears to be unnecessary.
- In general, we should seek to minimize any constraints on implementations, unless they are absolutely necessary (and this one does not appear to be).
- The proposal is to remove this constraint and bring the 8:8 multiplexing rules into alignment with the multiplexing rules for the 32:8 PMA and the 8:32 PMA.
- Detailed changes to the draft in support of this proposal are captured on the next slide.

# Proposal - detailed changes to the draft

Change 173.4.2.3 (page 233, line 7) from:

“The 4 PCSLs received on an input lane shall be mapped to an output lane such that the Gray mapped PAM4 symbol sequence on the output lane is identical to the Gray mapped PAM4 symbol sequence on the input lane, except for possible swapping of each bit pair (see 173.4.7.1).”

To:

“The 4 PCSLs received on an input lane shall be mapped to an output lane such that the order of PCSLs is maintained from input lane to output lane.”

Thanks

# Backup

# 802.3df D2.0 - Comment #83

Cl 173 SC 173.4.2.3 P 233 L 7 # 83

Nicholl, Gary

Cisco Systems

Comment Type T Comment Status X

The multiplexing rules in this section (along with the multiplexing rules in 173.4.2.1 and 173.4.2.1) were updated based on comment #27 against D1.1 and supporting presentation "[https://www.ieee802.org/3/df/public/23\\_01/0130/ran\\_3df\\_01b\\_230130.pdf](https://www.ieee802.org/3/df/public/23_01/0130/ran_3df_01b_230130.pdf)".

As captured in slide 3 of [ran\\_3df\\_01b\\_230130](#) the motivation of the proposed change was to avoid the situation "where one of two flows always gets the LSB of the PAM4 symbols"

The changes to the multiplexing rules for PMA 32:8 (173.4.2.1) and PMA 8:32 (173.4.2.2) achieve this goal.

However the change to the multiplexing rules for the PMA 8:8 (173.4.3) goes one step further than the changes to the PMA 32:8 and PMA 8:32. This additional restriction is unnecessary (as the situation this step is trying to avoid can be caused by both the PMA 32:8 and PMA 8:32 anyway), and it may make some existing 100G PAM4 retimer implementations non-compliant.

The additional step is the requirement that "the Gray mapped PAM4 symbol sequence on the output lane is identical to the Gray mapped PAM4 symbol sequence on the input lane" This means the PAM4 output must be MSB/LSB aligned to the PAM4 input. It is not clear that this would always be the case, and is something that is not required for the 400GbE generation of PAM4 retimer chips. It is also not fully consistent with the description of the PAM4 Encoding described in 173.4.7.1 (which essentially references the PAM4 encoding rules from Clause 120, which do not require PAM4 outputs to be MSB/LSB aligned to PAM4 inputs).

This step is not required in order to meet the intent captured in slide 3 of [ran\\_3df\\_01b\\_230130.pdf](#).

If the PAM4 input is decoded to a serial bit stream, then in order to meet the intent of [ran\\_3df\\_01b\\_230130.pdf](#), the only requirement is that the bit stream be sent in the same order (no rearrangement of bits) to the PAM4 output encoder. The output encoder just has to take two bits at a time and encode into a PAM4 symbol (consistent with the description in 173.4.7.1). There is no need for the PAM4 encoder to be MSB/LSB aligned to the bit stream coming from the PAM4 receiver.

It should also be noted that this section only describes the bit level multiplexing functions of a serial bit stream (in keeping with Figure 173-5), and the PAM4 decoding and encoding rules are described in a different section (173.4.7.1).

## Suggested Remedy

Change from:

"The 4 PCSLs received on an input lane shall be mapped to an output lane such that the Gray mapped PAM4 symbol sequence on the output lane is identical to the Gray mapped PAM4 symbol sequence on the input lane, except for possible swapping of each bit pair

(see 173.4.7.1)."

to:

"The 4 PCSLs received on an input lane shall be mapped to an output lane such that the order of PCSLs is maintained from input lane to output lane, except for possible swapping of each bit pair (see 173.4.7.1)."

Proposed Response

Response Status