IEEE 802.3cb PCS compatibility to 1000BASE-X PCS

2017-09

Yong Kim

V1

Contents

- Concern:
  - 1000BASE-X PCS running at 2.5X speed (i.e. 2.5 Gb/s) interoperating with 2.5GBASE-X PCS (with XGMII lane 0 start & implicit lane 3 SFD alignments).
- Related Comments (copies on following slides):
  - # i-125, Law, D on D3.0 – XGMII that added a note to CL46 46.3.3.3, and then
  - # r01-10, Anslow, P on D3.1 – 46.3.3.3 changes TBD text to describe and make the behavior normative, and
    - and a consequence -- # r01-16, Anslow, P on D3.1 – External reference Green → Xref, since CL46 is now (due to note) in .3cb.
- Related Clauses:
  - 127.2 802.3cb PCS
  - Annex 127A, informative, Compatibility of …
- Related Facts: MAC has no concept of any size alignment – just bit-serial protocol that is converted to parallel protocol via RS.
- XGMII alignment discussion (not a part of the bluebook but precursor of): http://grouper.ieee.org/groups/802/3/ae/public/jul00/haddock_1_0700.pdf

Copy of comment db - Law

- i-126 Law

- As stated in annex 127A 'It is permissible for a compliant 1000BASE-X PCS transmit process to truncate the first byte of a preamble in order to align the start of packet on the EVEN boundary.'. In the 2.5GBASE-X receive path the WALIGN function called by the RX_XGMII state of the Figure 127-9 'Octets-to-Word and Decode state diagram' performs alignment according to subclause 127.2.4.4. Based on the rules described in that subclause the first packet received will set the Deficit Idle Count to place the first Data symbol, in this case the SPD replaced by a preamble octet by the PCS, on wd_rpd<7:0> of the 2.5GPII. This in turn will be encoded as a XGMII 'start' (RXC = 1, RXD = 0xFB) on lane 0 as required by Clause 46. As noted above the first octet of preamble may be discarded on transmit by a 1000BASE-X PCS. This results in the transmission of, and therefore reception of, a 7 octet preamble. With the first octet of this 7 octet preamble aligned by the WALIGN function on XGMII lane 0, the SFD will be received on Lane 2 of XGMII, not Lane 3 as illustrated in IEEE Std 802.3-2015 Figure 46-8 and 46-9. IEEE Std 802.3-2015 subclause 46.3.3.3 'Response to received invalid frame sequences' states 'Error free 10 Gb/s operation will not change the SFD alignment in lane 3' and 'A 10 Gb/s MAC/RS implementation is not required to process a packet that has an SFD in a position other than lane 3 of the column following the column containing the Start control character.'. There appears to be no changes to this text as a result of IEEE Std 802.3bz-2016 amending the XGMII specification to support operation at 2.5 Gb/s and 5 Gb/s as well as 10 Gb/s. As a result the above text only applies to XGMII 10 Gb/s operation and IEEE 802.3 is silent in this respect for 2.5 Gb/s and 5 Gb/s XGMII operation. That being said, there may be an assumption made that a 10 Gb/s MAC/RS/XGMII implementation may also support 2.5 Gb/s operation through quarter rate clocking. This however is not the case if the implementation took the option of not processing packets that have an SFD in a position other than lane 3 as is permitted by IEEE Std 802.3-2015 subclause 46.3.3.3. If that option is implemented all packets received from a 2.5GBASE-X would not be processed as the SFD will always be received in lane 2.

- While strictly speaking IEEE Std 802.3-2015 subclause 46.3.3.3 only applies to a 10 Gb/s MAC/RS/XGMMI, to avoid any incorrect assumptions, suggest that:
  [1] The text '...to be able to accept a seven byte preamble on the XGMII.' in the penultimate paragraph of Annex 127A be changed to read '...to be able to accept a seven byte preamble on the XGMII with the SFD positioned on lane 2.'.
  [2] A note that reads 'Note: To support 2.5GBASE-X compatibility with a 1000BASE-X PCS/PMA running 2.5 times faster, a 2.5Gb/s MAC/RS implementation has to support a Start control character received on either lane 2 or lane 3.' Be added to the end of subclause 46.3.3.3

46.3.3.3

Observation statement

46.3.3.3 Response to received invalid frame sequences

The 10 Gb/s PCS is required to either preserve the column alignment of the transmitting RS, or align the Start control character to lane 0. The RS shall not indicate DATA_VALID to the MAC for a Start control character received on any other lane. Error free 10 Gb/s operation will not change the SFD alignment in lane

3. A 10 Gb/s MAC/RS implementation is not required to process a packet that has an SFD in a position other than lane 3 of the column following the column containing the Start control character.

Optional behavior. Not a normative statement.

MAC has no concept of any lane – just bit-serial.
RS is aware of lanes.

Copy of comment db - Anslow

- r01-10 Anslow
  - *Comment Type* **TR**
  - The normative behaviour of the RS defined in 46.3.3.3 cannot be changed with the addition of a note (which is defined to be informative). The relevant part of 46.3.3 for this is 46.3.3.3, so that is where the change should be made. The headings for 46.3 and 46.3.3 are missing.
  - *SuggestedRemedy*
  - Insert the headings for 46.3, 46.3.3, and 46.3.3.3. Change the editing instruction to: "Change 46.3.3.3 as follows:“ Show appropriate changes to the text of 46.3.3.3 that preserve the behavior for the existing PHYs that use this Clause and add normative support for a Start control character received on either lane 2 or lane 3 for 2.5GBASE-X.
  - REJECT.
  - The Comment Resolution Group could not come to a consensus to make changes necessary to address the comment because the solution is unknown and needs further research.

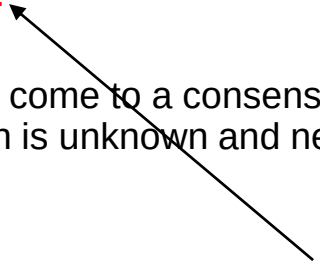

So, should we accept in principle, and insert proper statements for 2.5GBASE-X?

Problem Statement (AGAIN!)

- CL46 XGMII lane 0 start and lane 3 SFD compatibility, when CL36 PCS sends extra idle bytes that replaces preamble (“EVEN” alignment function”).

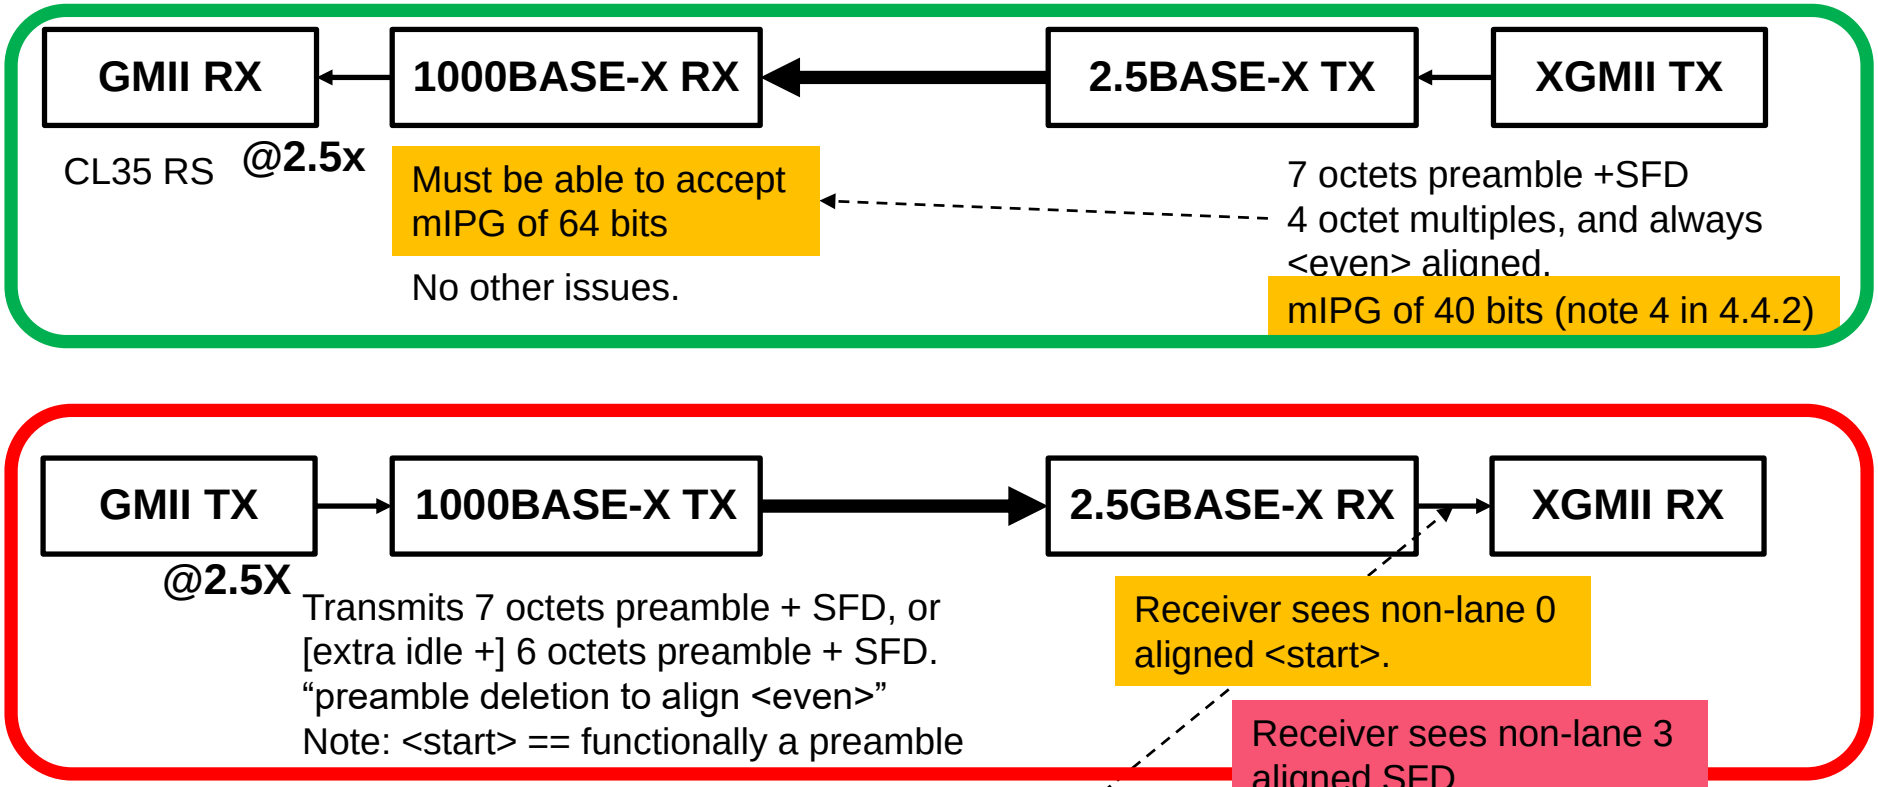

BROADCOM

XGMII TX ➔ 2.5x-GMII RX, no issues.

- IPG range on XGMII, from purely lane 0 alignment function, is 9~15 octets (72~120 bits), with long term average of 12 octets (96 bits), but IPG @ MAC is always 12 octets.
- IPG range on XGMII, from other PHYs, particularly multilane delay skew and block (FEC, codes, etc) – the comfortable lower bound is set to a minimum of 40 bits.
- Single lane PHY like 2.5GBASE-X have 72~120 bit IPG range, well above the mIPG of 64 bits – inferred.

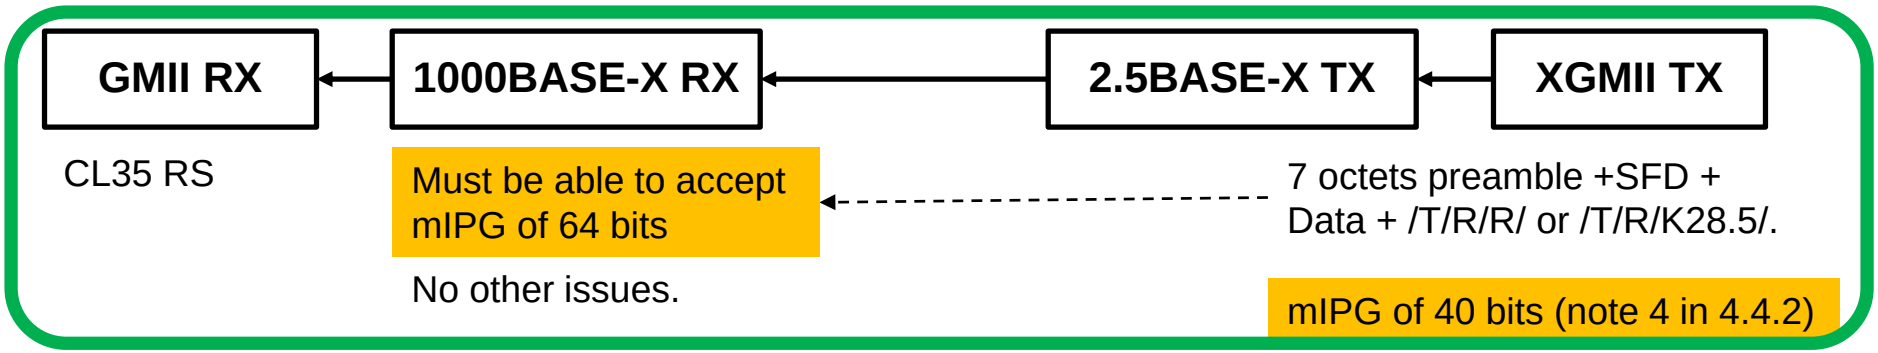

BROADCOM

2.5x-GMII TX ➔ XGMII RX, Issue – Lane 3 SFD

- CL46 XGMII lane 0 <start> and lane 3 SFD compatibility, when CL36 PCS sends extra idle bytes that replaces preamble (“odd-even alignment function”).

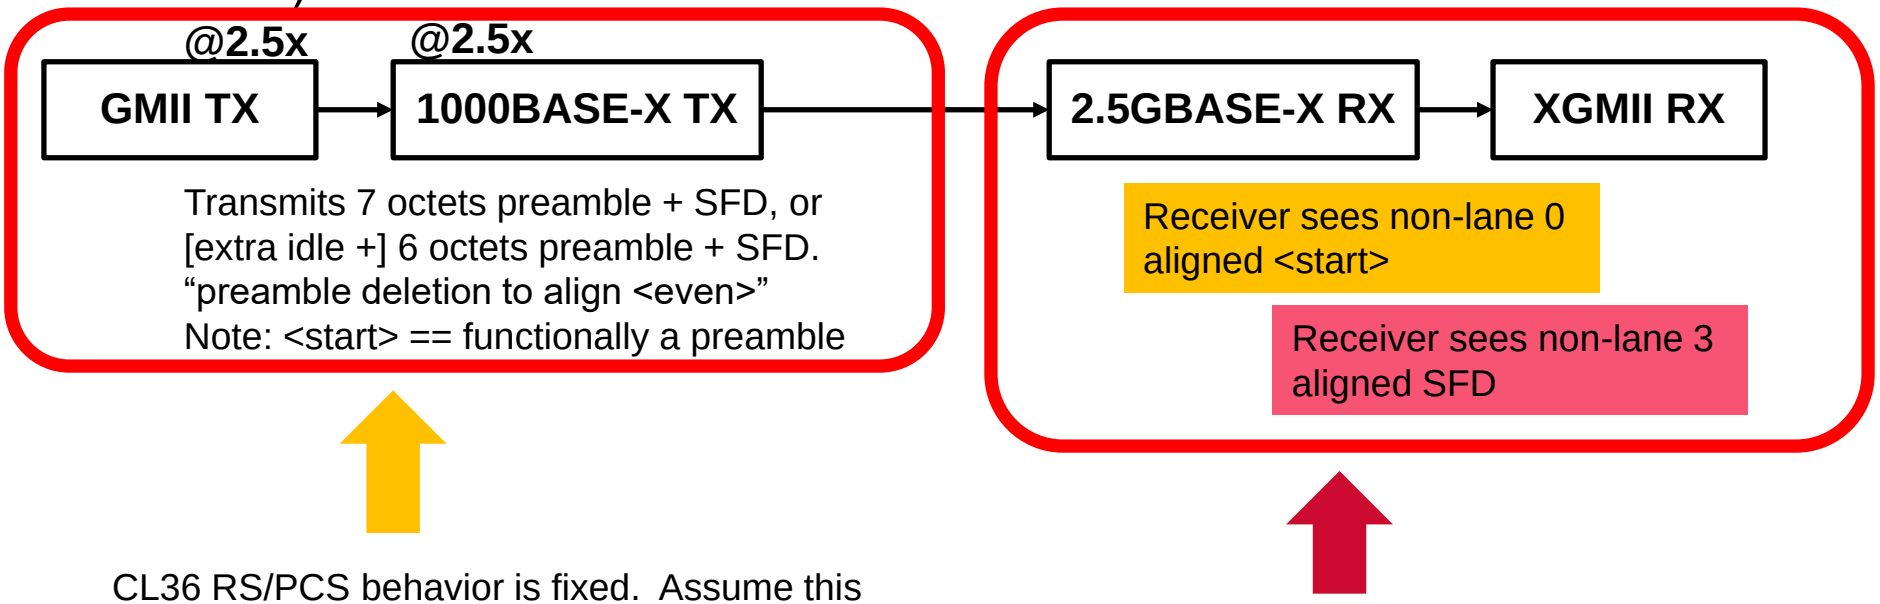

CL36 RS/PCS behavior is fixed. Assume this cannot be changed.

1. Any “repair work” could be done local to 2.5GBASE-X receive path. Any “repair work” should be done here.

2. Let’s see what it does.

3. Let’s see what it should do.

BROADCOM

Fig 127-2

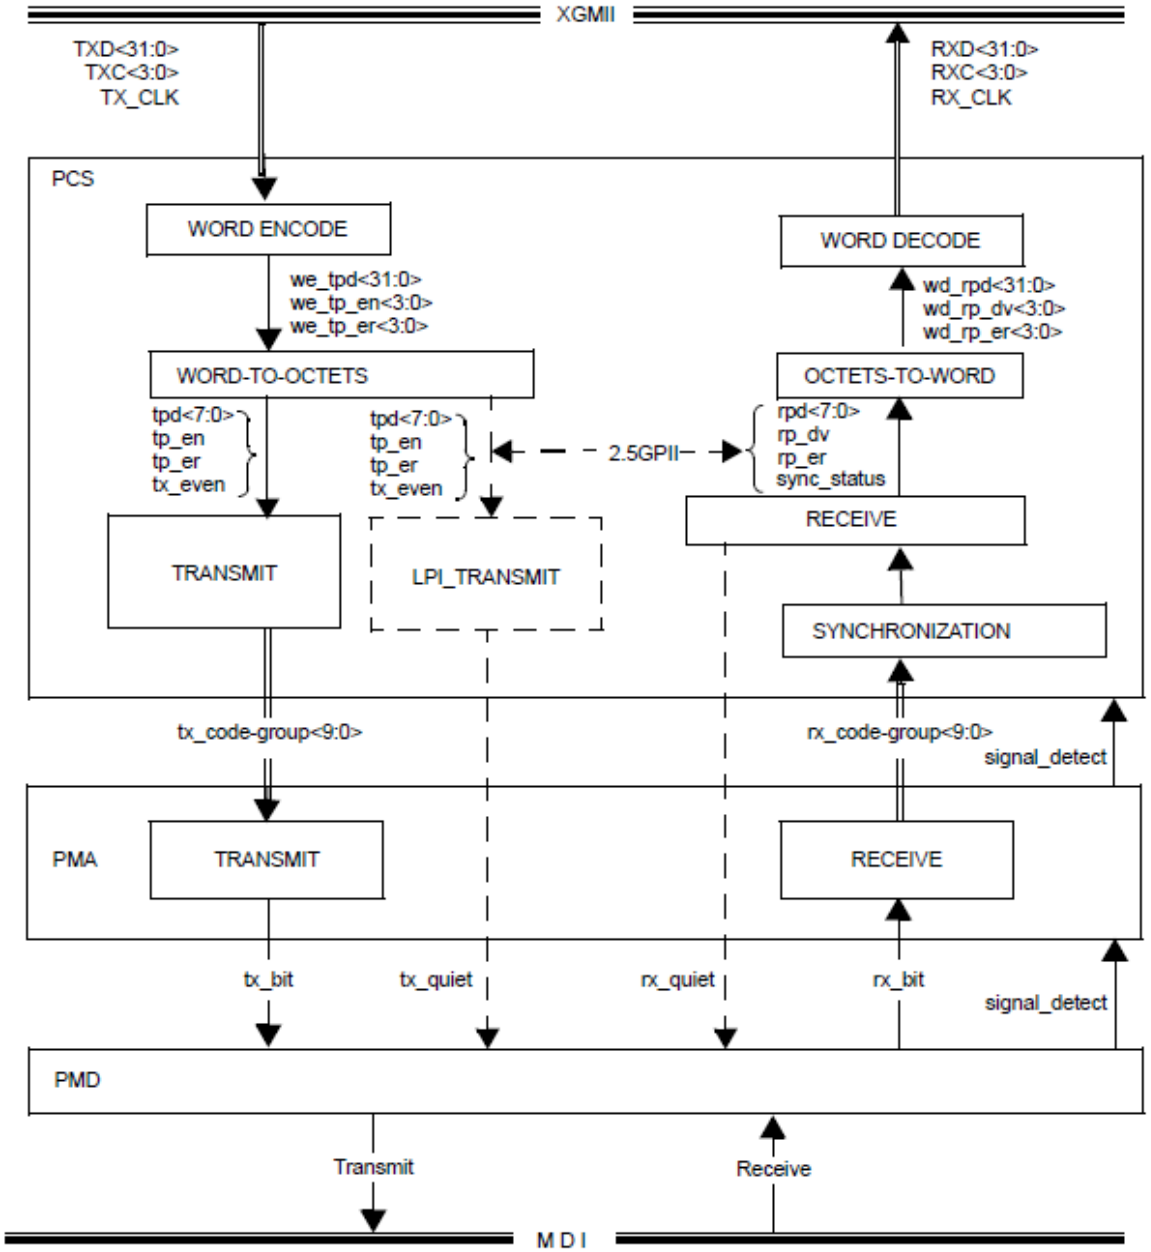

Figure 127–2—Functional block diagram

BROADCOM

2.5GBASE-X Receive Path (MDI → XGMII)

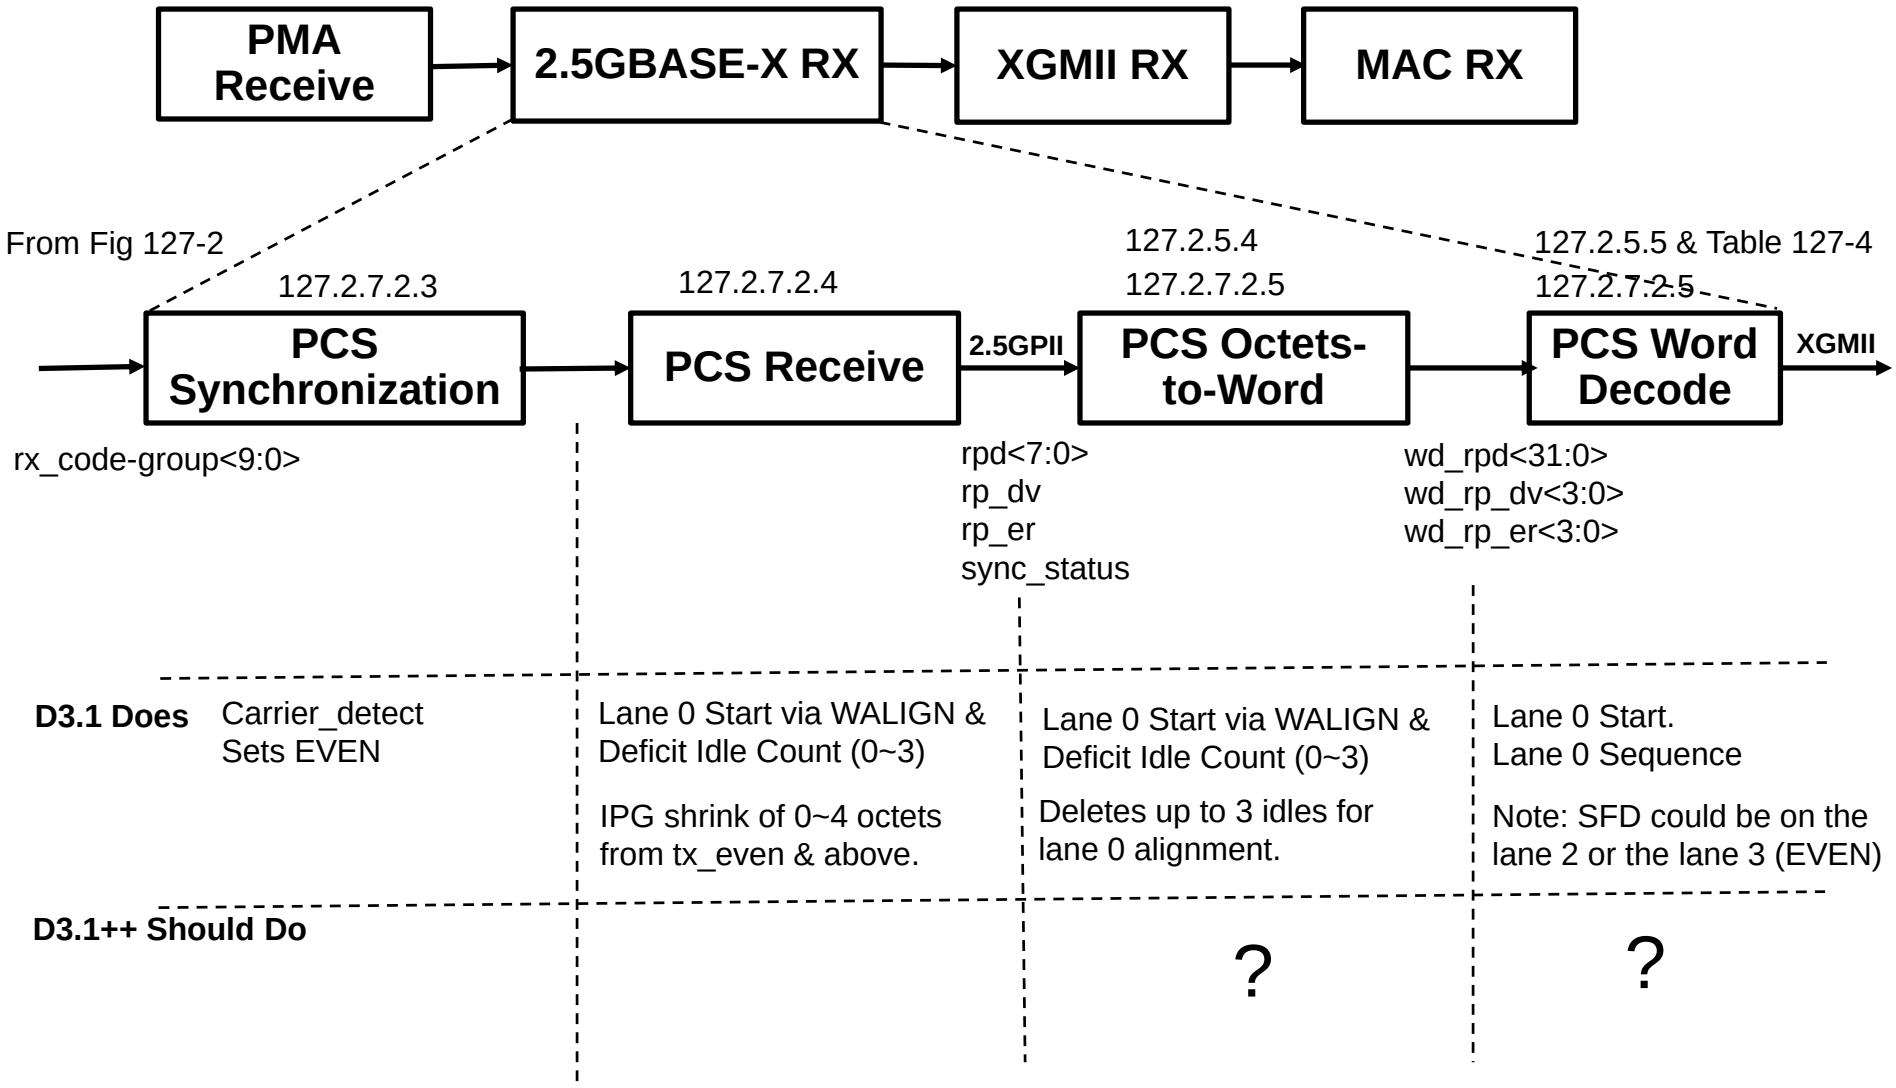

BROADCOM

Observations WRT possible remedy

- **46.3 XGMII functional specifications**
- **46.3.2 Receive**

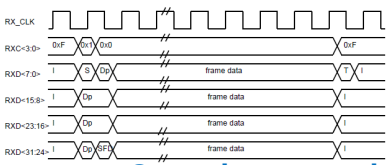

Figure 46–8—Basic frame reception

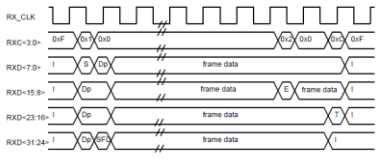

Figure 46–9—Reception with error

Should state (explicitly) that SFD could be on lane 2 in addition to lane 3 in Receive clause

Error and fault clause.
SFD on lane 2 is not an error.

Similar to 10M ignored pkts w/ "00" in SFD

Text in 127A works well

Note In 46.3.3.3 is helpful!

- **46.3.3 Error and fault handling**
- **46.3.3.3 Response to received invalid frame sequences**
- The 10 Gb/s PCS is required to either preserve the column alignment of the transmitting RS, or align the Start control character to lane 0. The RS shall not indicate DATA_VALID to the MAC for a Start control character received on any other lane. Error free 10 Gb/s operation will not change the SFD alignment in lane 3. A 10 Gb/s MAC/RS implementation is not required to process a packet that has an SFD in a position other than lane 3 of the column following the column containing the Start control character.
- Suggested Remedy: [1] The text '...to be able to accept a seven byte preamble on the XGMII.' in the penultimate paragraph of Annex 127A be changed to read '...to be able to accept a seven byte preamble on the XGMII with the SFD positioned on lane 2.'.
  [2] A note that reads 'Note: To support 2.5GBASE-X compatibility with a 1000BASE-X PCS/PMA running 2.5 times faster, a 2.5Gb/s MAC/RS implementation has to support a Start control character received on either lane 2 or lane 3.' Be added to the end of subclause 46.3.3.3

Proposed resolution to comments & the issue

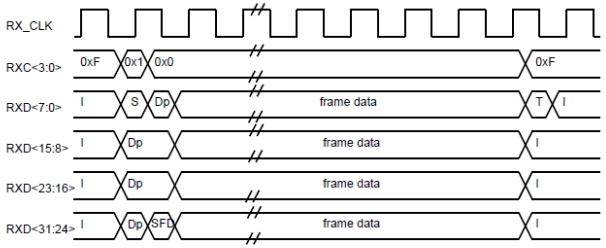

I: Idle control character, S: Start control character, Dp: preamble Data octet, T: Terminate control character

Figure 46–8—Basic frame reception

- **46.3.2 Receive**
- **46.3.2.1 RX_CLK (receive clock)**
- **46.3.2.2 RXC<3:0> (receive control)**
- WRT "…Figure 46–8 shows the behavior of RXC<3:0> during an example frame reception with no errors. …",
  No need to read into the illustration of the SFD on lane 3. "T" shows up on lane 0 only. Proposing no action.
- **46.3.2.3 RXD (receive data)**
- At the end of the 1st paragraph "…RXD<8> the least significant bit of lane 1, RXD<16> the least significant bit of lane 2, and RXD<24> the least significant bit of lane 3. Figure 46–8 shows the behavior of RXD<31:0> during frame reception."
  Propose to add "Error free 10 Gb/s operation will not change the transmitted SFD alignment in lane 3. ~~Error free 2.5 Gb/s operation with 2.5GBASE-X~~ Operation described in CL127A may result in received SFD alignment in lane 2 or lane 3."
  ← perhaps move this back to error handling section?
- And use the above as the Accept-in-principle to the comment # r01-10 (Anslow).

BROADCOM