Inconsistent signal detect rules

(comments I-433, I-436, R1-289, R1-290, R1-291, R1-292, R1-293, R2-197)

Piers Dawe

NVIDIA

Comment R2-197

- R2-197 SC 180.5.1 P 468 L 27 TR
- As previously pointed out, Figure 180-2 (block diagram) indicates that signal detect is disconnected from SIGNAL_OK, contrary to all optical PMDs. 180.5.5 says "This standard imposes no response time requirements on the generation of the PMD_signal_detect_i variable." It turns out that they are connected and timing is important: see 178B.8.3.1 as revised: rx_ok "it shall be set to false when an input signal corresponding to tx_disable = true in the peer is received for 100 ms (the minimum duration of the quiet_timer) or longer". As the criterion for tx_disable = true is Average launch power of OFF transmitter <= -15 dBm, which is the criterion in Table 180-5, Signal detection criteria for each lane: "Average optical power at TP3 <= -15 dBm", signal detect IS required, AND its timing is required (which is new and would surprise optical engineers if they read it).
- I am aware that other things have to be right for SIGNAL_OK at the service interface to be OK. Also I am aware that both 180 and 178A say that various implementations of the signal detect function are permitted by this standard / the specific conditions for setting this variable to true or false are implementation dependent. Also that the timings for reporting to management might trail behind a state machine. Nevertheless, Clause 180 is seriously misleading. Part of the problem may be a belief that IMDD signal detect must be based on average optical power, or that now that 178B post-processes traditional optical power detection with other criteria, we may be using one name for two things. It is useful to distinguish between "is there a signal?" and "is everything ready?" so we should try to be clear.
- Revise Figure 180-2 and 180.5.5 to be consistent with 178B.8.3.1, with appropriate cross-references. Similarly in other optical clauses.

What is normal 1/3

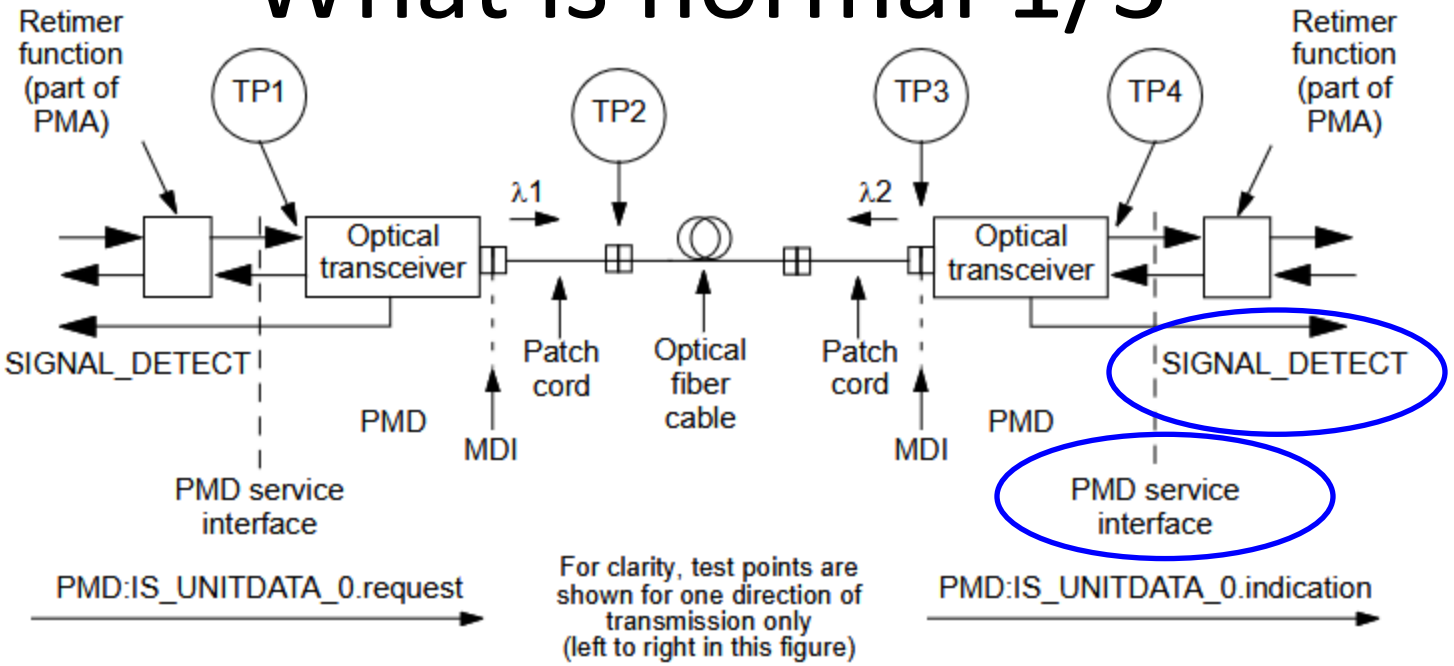

Figure 168–2—Block diagram for 100GBASE-BRx transmit/receive paths

- Signal detect is a function, a primitive and a status reported to management
  - Names are quite complicated – they are not the point
- The function and primitive is required
- There is a truth table that it shall obey
  - As permissive as possible
  - Consistent with definition of transmit disable
- There are no timing requirements
- Various implementations are allowed
- Signal Indication Logic (SIL) in the PMA immediately above the PMD acts on it

What is normal 2/3

168.5.4 PMD global signal detect function

The PMD global signal detect function shall report the state of SIGNAL_DETECT via the PMD service interface. The SIGNAL_DETECT parameter is signaled continuously, while the PMD:IS_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the inter-sublayer service interface primitives defined in 80.3.

SIGNAL_DETECT is a global indicator of the presence of the optical signal. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 168–4. The PMD receiver is not required to verify whether a compliant 100GBASE-BRx signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

...

Various implementations of the Signal Detect function are permitted by this standard, including implementations that generate the SIGNAL_DETECT parameter values in response to the amplitude of the modulation of the optical signal and implementations that respond to the average optical power of the
modulated optical signal.

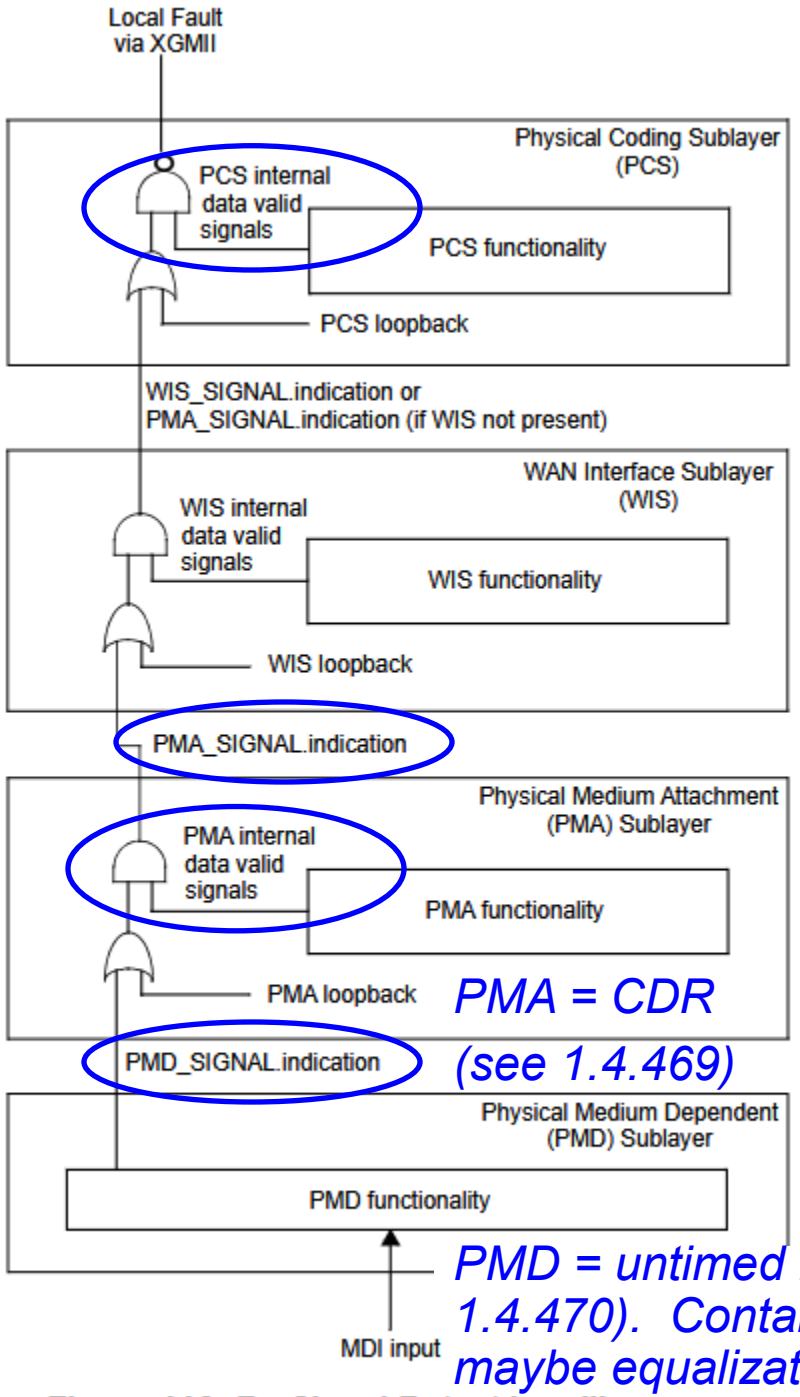

Figure 44A–7—Signal Detect handling across sublayers

What is normal 3/3

- The signal OK information is gated at each sublayer up the stack
  - A CDR recovering timing from a chattering Rx front end is not allowed to claim that there is an OK signal when no light is coming in
- PMAs can be stacked

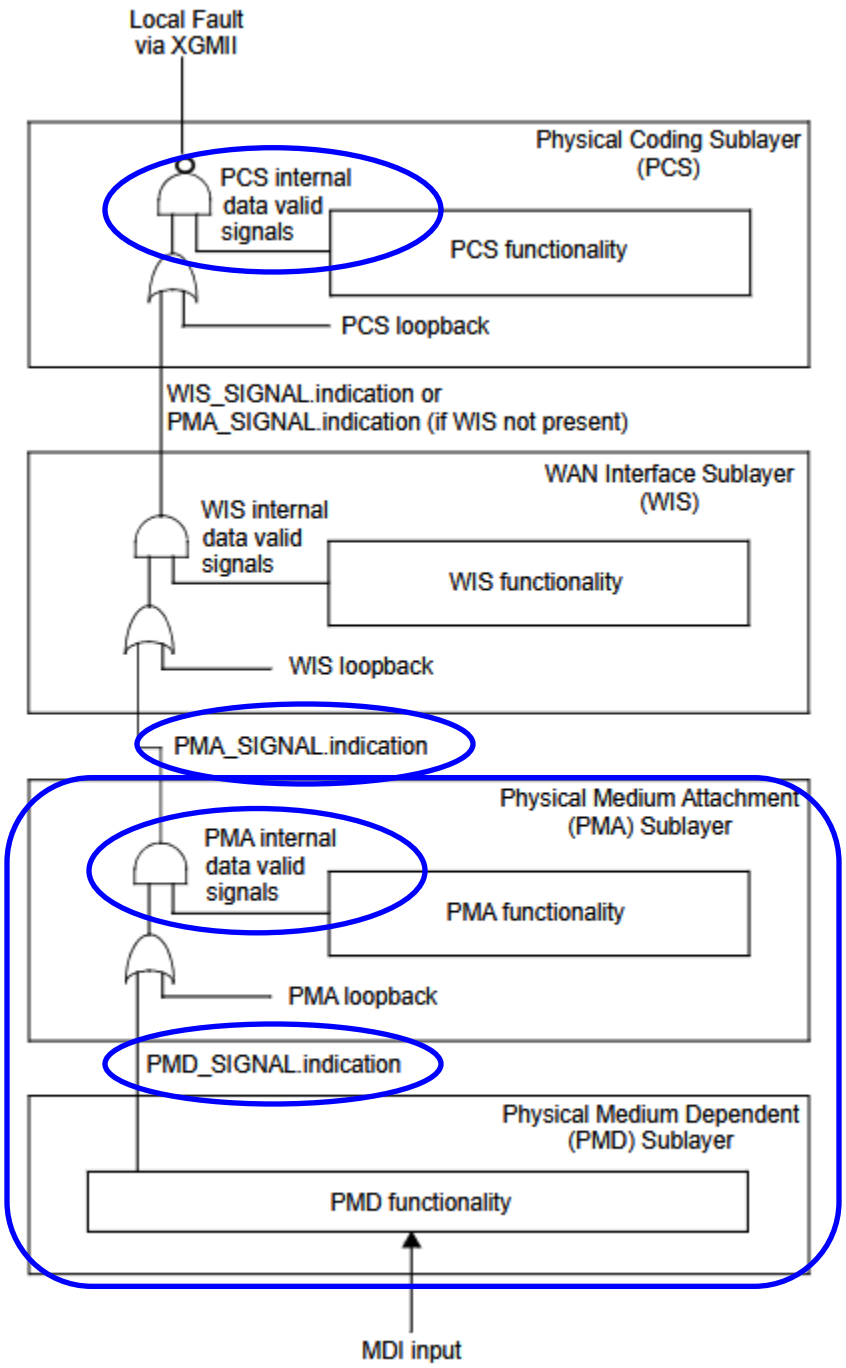

Figure 44A–7—Signal Detect handling across sublayers

New in 3dj

- P802.3dj is written as if the PMD sublayer contains a PMA
- Signal detect still exists but as the traditional PMD-PMA interface is not the boundary between two sublayers, it is not a primitive
- It as an input to APSU (see 178B), which replaces SIL functionality related to "PMA internal data valid signals" and generates the SIGNAL_OK parameter / IS_SIGNAL.indication primitive

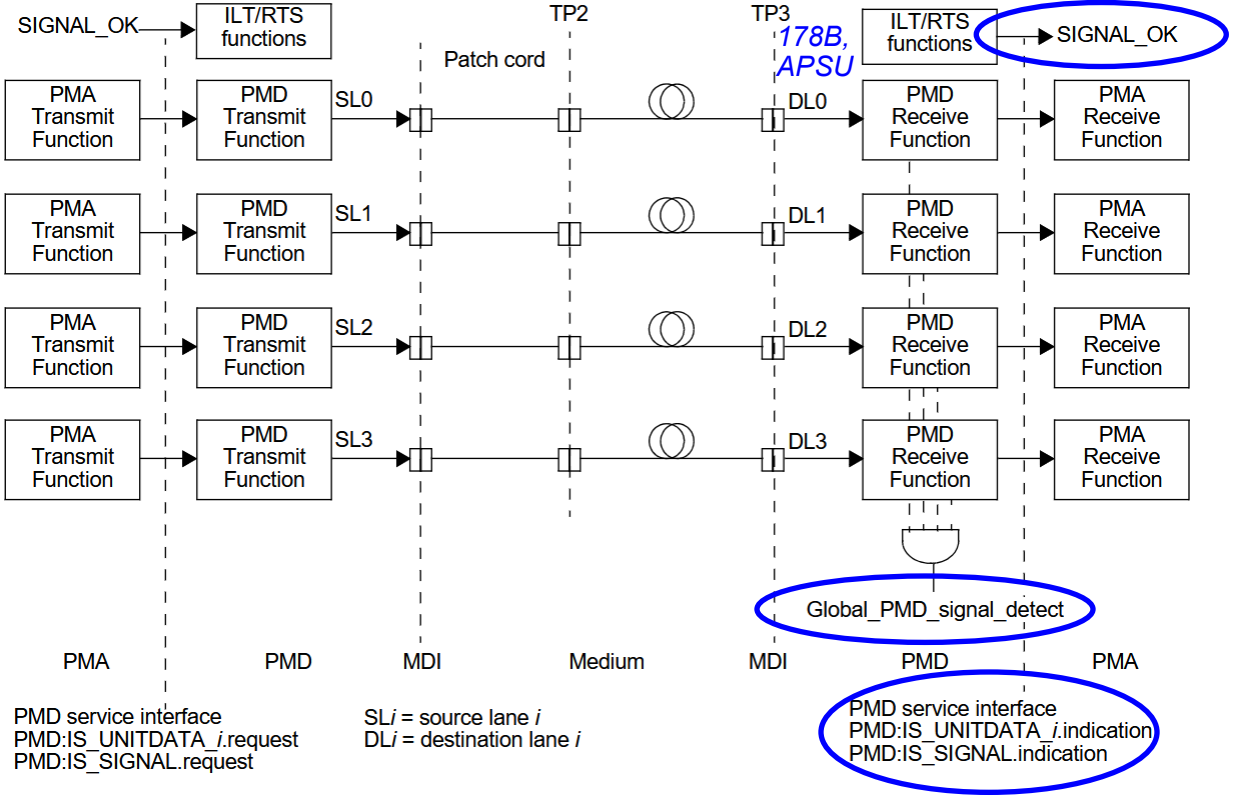

NOTE: For clarity, only one direction of transmission is shown

Figure 180–2—Block diagram for 800GBASE-DR4 transmit/receive paths

180.5.4 PMD global signal detect function

The variable Global_PMD_signal_detect is a global indicator of the presence of optical signals on all *n* lanes. For the 200GBASE-DR1 PMD, the variable Global_PMD_signal_detect is set to PMD_signal_detect_0 (see 180.5.5). For the 400GBASE-DR2, 800GBASE-DR4, and 1.6TBASE-DR8 PMDs, the variable Global_PMD_signal_detect is set to the logical-and of PMD_signal_detect_*i* from all of the PMD lanes.

180.5.5 PMD lane-by-lane signal detect function

The PMD lane-by-lane signal detect function is used by the PMD to indicate sufficient optical power is detected at the receiver input on each lane. For each lane *i*, where *i* is in the range 0 to *n*–1, a receiver shall set the variable PMD_signal_detect_*i* to either 0 or 1 according to the criteria in Table 180–5.

Various implementations of the signal detect function are permitted by this standard. This standard imposes no response time requirements on the generation of the PMD_signal_detect_*i* variable.

- In the draft Clause 180,
- SIGNAL_OK is a primitive that is transferred across the "PMD" service interface as PMD:IS_SIGNAL.indication
- According to 180, Global_PMD_signal_detect (the logical-and of PMD_signal_detect_i from all the PMD lanes) goes nowhere but to management (Register 1.10, see 45.2.1.9, as usual)
- This is wrong
  - Not how it should work, hence the older comments
  - Not how it does work according to 178B; see next slide
- Cl. 180 says that various implementations are allowed
- Cl. 180 says that there are no response time requirements

178B tells a different story

From 178B.8.3 ILT per-lane state diagram definitions, 178B.8.3.1 Variables

rx_ok

Boolean variable that is set to true when the receiver on a lane of the interface has determined that the peer interface transmitter is transmitting the expected signal format, that the remote transmit has been optimized if training = true, that the local receive equalizers have been optimized, and that no further adjustments are required for normal data transmission. The specific conditions for setting this variable to true or false are implementation dependent. However, it shall be set to false when an input signal corresponding to tx_disable = true in the peer is received for 100 ms (the minimum duration of the quiet_timer) or longer.

- rx_ok appears several times in Figure 178B–10a and 178B–10b, training control state diagram. It is normative, it is used by the signal indication logic in APSU, which determines SIGNAL_OK at the PMD service interface
  - It is mapped to 45.2.1.169e PMA/PMD interface lane training status 2 (Register 1.1478)
- The criterion for setting rx_ok to true is the same as 180 has for its status-only signal detect (compliant and acceptable signal, at receiver's discretion, no timing requirement)
  - It goes without saying that a receiver presented with a signal that's adequate in strength and compliant in form, will train to it and accept it as OK
- The criterion for setting rx_ok to false is the same -15 dBm as in 180 has for its status-only signal detect, and there is a timing requirement
  - Clause 180 says that there are no response time requirements
- 178B says that various implementations are allowed, agreeing with 180

rx_los and signal detect are intimately connected

Table 180–5—Signal detection criteria for each lane

Receive conditions	PMD_signal_detect_*i*
Average optical power at TP3 ≤ –15 dBm	0
[The input signal is a compliant 200GBASE-R, 400GBASE-R, 800GBASE-R or 1.6TBASE-R signal][a] AND [Optical power at TP3 ≥ average receive power, each lane (min) in Table 180–8]	1
All other conditions	Unspecified, may be 0 or 1

[a] This is a condition for testing of the receiver, not a requirement for the receiver to detect that the signal is valid.

Table 180–7—200GBASE-DR1, 400GBASE-DR2, 800GBASE-DR4, and 1.6TBASE-DR8 transmit characteristics *(continued)*

Description	Reference	200GBASE-DR1	400GBASE-DR2 800GBASE-DR4 1.6TBASE-DR8	Unit
Extinction ratio, each lane (min)	180.9.11	3.5		dB
Transmitter transition time, each lane (max)	180.9.12	8		ps
Average launch power of OFF transmitter, each lane (max)	—	–15		dBm

178B.8.3.1 Variables

rx_ok

Boolean variable that is set to true when the receiver on a lane of the interface has determined that the peer interface transmitter is transmitting the expected signal format, that the remote transmit has been optimized if training = true, that the local receive equalizers have been optimized, and that no further adjustments are required for normal data transmission. The specific conditions for setting this variable to true or false are implementation dependent. However, it shall be set to false when an input signal corresponding to tx_disable = true in the peer is received for 100 ms (the minimum duration of the quiet_timer) or longer.

- They are almost the same (nested) apart from timing
  - Traditional RSSI style signal detect can be slow to respond, either way
  - A signal might have adequate average power but not enough net usable modulation amplitude (NUMA = OMA – ECQ)

If SD is not OK, in steady state, rx_los must not transition to OK. SD gates rx_los

- That's OK, the implementer is allowed to choose the SD threshold with wide discretion

If SD is OK, rx_los might stay at not OK

- Useful to have separate management registers, as in the draft

If the signal disappears, rx_los could change before SD does

- Typically, a PMA or PCS can detect long strings of zeros. Here, there is a PMA inside the PMD

If a signal appears, rx_los *might* change to OK before RSSI-based SD does

- Whether this is allowed or disallowed hardly matters unless it goes on for so long that CMIS reports the signal is trained but doesn't exist

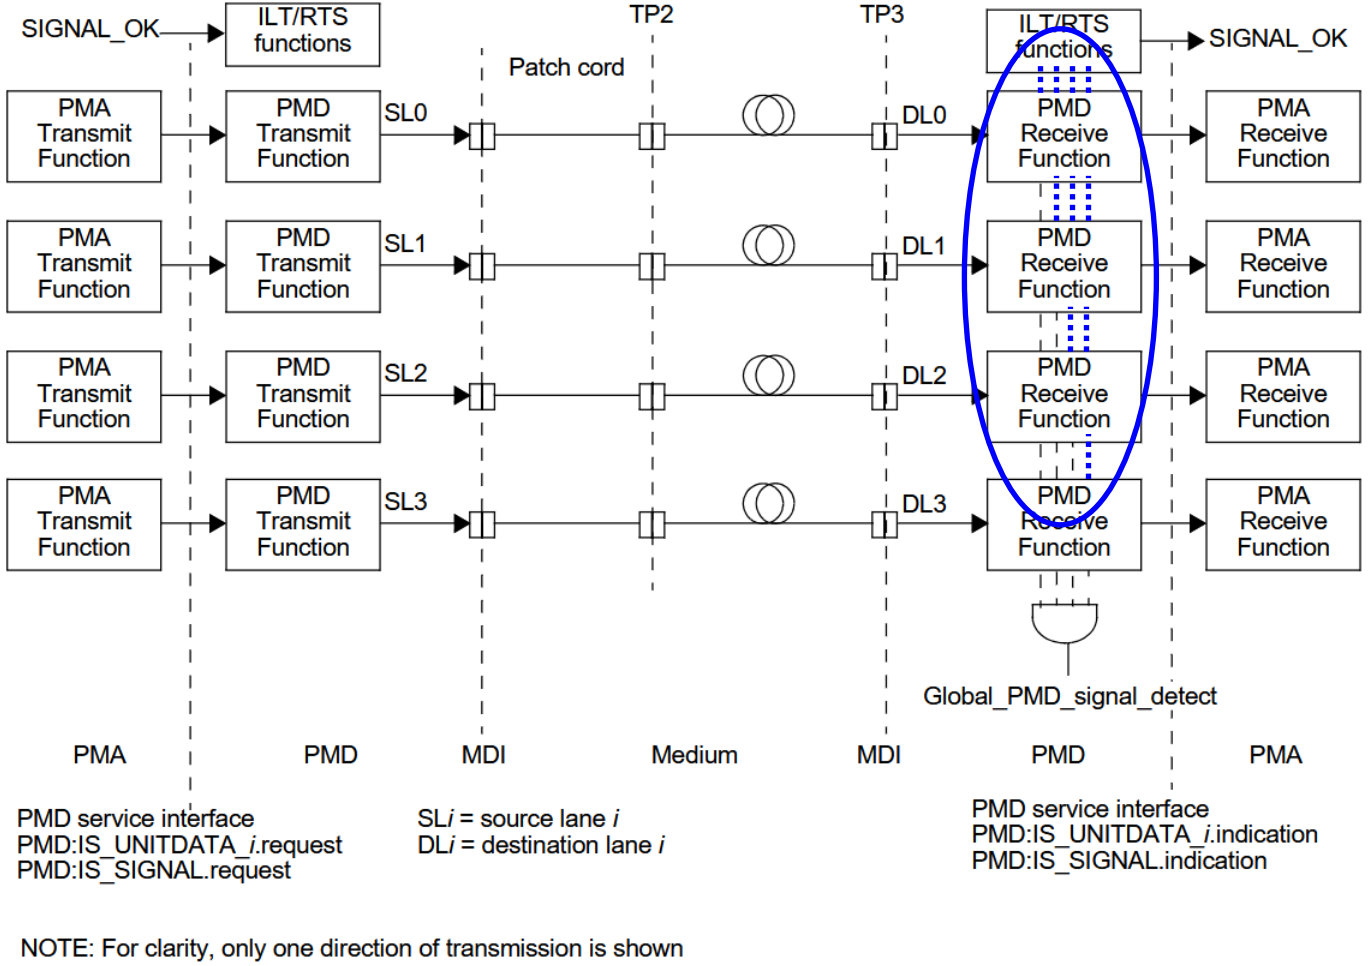

Figure 180–2—Block diagram for 800GBASE-DR4 transmit/receive paths

- Add lines to show that lane-by-lane SD goes to APSU
- Or, remove the lines, the gate and Global_PMD_signal_detect, to match Tx , and because we don't usually draw individual management status variables in block diagrams

180.5.5 PMD lane-by-lane signal detect function

The PMD lane-by-lane signal detect function ~~is used by the PMD to~~ indicates if insufficient optical power is detected at the receiver input on each lane. It is used by ILT (see 178B.8.3). For each lane i, where i is in the range 0 to n–1, a receiver shall set the variable PMD_signal_detect_i to either 0 or 1 according to the criteria in Table 180–5.

Various implementations of the signal detect function are permitted by this standard. See the definition of rx_ok in 178B.8.3.1 for the timing requirement for the signal detect function. However, t~~T~~his standard imposes no response time requirements on the generation of the PMD_signal_detect_i variable.

Implementations need to provide adequate margin between the input optical power level at which the PMD_signal_detect_i variable is set to 1 and the inherent noise level of the PMD, including the effects of crosstalk, power supply noise, etc.

180.5.12 Autonomous path startup (APSU) functions

A PMD shall provide the ready-to-send (RTS) function and the ILT function with a training frame of O1 format, specified in Annex 178B.

When the variable mr_training_enable is true, the ILT function is used to request changes to the peer transmitter state (modulation, training pattern, and precoder state), indicate the receiver state, and coordinate the transition to DATA mode.

When tx_mode = LOCAL_PATTERN (see 178B.8.3.1) the PMD transmits PRBS31Q (see 176.7.4.2).

The default value of max_wait_timer_duration is 60, corresponding to a duration of 60 seconds for max_wait_timer (see 178B.8.3.3).

- *All this is transmitter-related. Add a sentence to indicate that APSU involves the receiver too*
- *Similarly in other clauses*