SOLARFLARE Communications


# Fixing the 10GBASE-T EEE state diagrams

9/22/2009 Gavin Parnaby

.

# State diagram comments

- A large proportion of the technical comments were directed at the state diagrams
  - **15 technical, 6 editorial**
- E: 79, (ER), 80 (ER), 81 (ER), 144, 372, 376 (ER)
- T: 102 (TR), 93 (TR), 94 (TR), 95 (TR), 142, 143, 96 (TR), 89, 84, 380 (TR), 378 (TR), 90, 91, 377, 379 (TR)
- This presentation shows the resulting state diagrams if the proposed comments/remedies are accepted

# Main issues

1. Lpi_tx_mode description does not match state diagram.
2. Tx_coded is assigned in two state machines and in the text. Description is ambiguous. Similar with tx_alert_symb.
3. Transitions from TX_T to TX_L need to be added, as in Clause 49 (also required for the Rx side)
4. Incorrect names for 65B blocks. Replace /LF/ and IDLE with new block names.
5. LII should be defined as a new type. It is valid only before transitions from LP_IDLE to IDLE.
6. SEND_WAKE to SEND_ERROR transition; related to 2)
7. Comment #377: Pass XGMII signaling for sleep/wake

# Problem 1) - Lpi_tx_mode 55.3.5.2.2

lpi_tx_mode

A variable indicating the signaling to be used from the PCS to the PMA across the PMA_UNITDATA.request (tx_symb_vector) interface.

~~lpi_tx_mode controls tx_symb_vector only when tx_mode is set to SEND_N.~~

The variable is set to NORMAL when tx_lpi_active is false, indicating that the PCS is in the normal mode of operation and will encode code-groups from the XGMII as described in 55.3.2.2.

The variable is set to REFRESH_A when tx_lpi_active * tx_active_pair==PAIR_A * tx_refresh_active.

The variable is set to REFRESH_B when tx_lpi_active * tx_active_pair==PAIR_B * tx_refresh_active.

The variable is set to REFRESH_C when tx_lpi_active * tx_active_pair==PAIR_C * tx_refresh_active.

The variable is set to REFRESH_D when tx_lpi_active * tx_active_pair==PAIR_D * tx_refresh_active.

The variable is set to QUIET when tx_lpi_active * (!tx_refresh_active + tx_lpi_initial_quiet)

- Comment #380; Tx_lpi_active is used ambiguously
- Tx_lpi_active is true from after sleep until just before alert - it does not include ALERT or WAKE
- But during ALERT and WAKE we DO NOT encode code-groups from the XGMII, so the boxed definition is not accurate and inconsistent with the state diagrams

# Problem 1) - Fig 55-16b

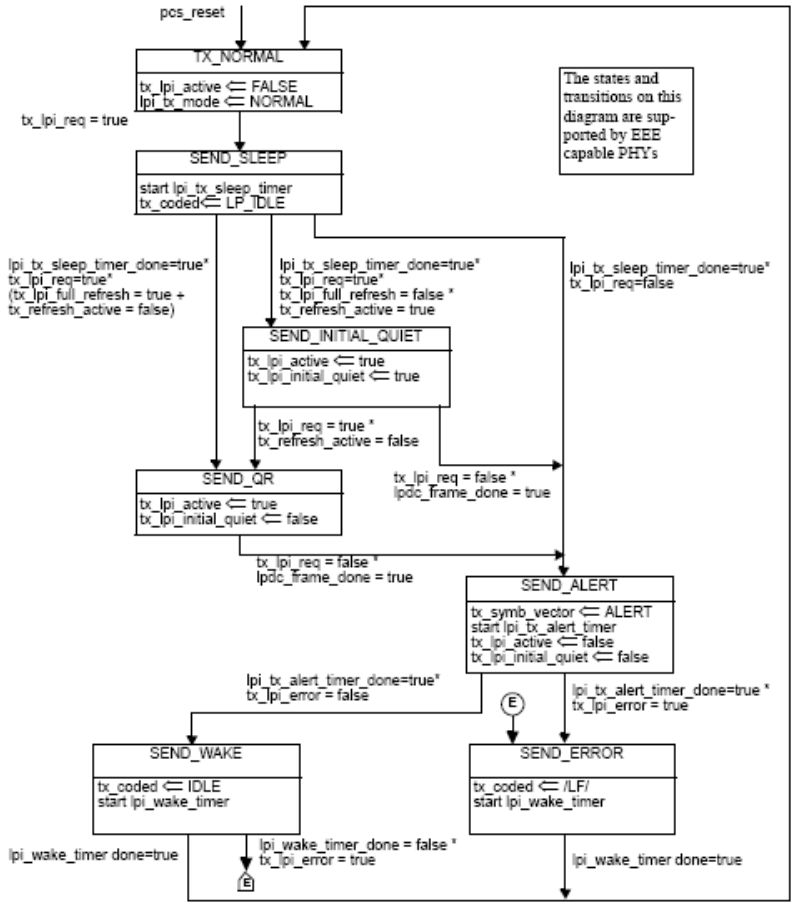


Figure 55–16b—EEE transmit state diagram

Tx_lpi_active is set false in SEND_ALERT

Tx_lpi_active==false means that lpi_tx_mode is NORMAL. From the previous text,
'the PCS… will encode codewords from the XGMII'
This does not match the state diagram which has tx_coded assignments



SOLARFLARE Communications

# Problem 2) - Figure 55-15a

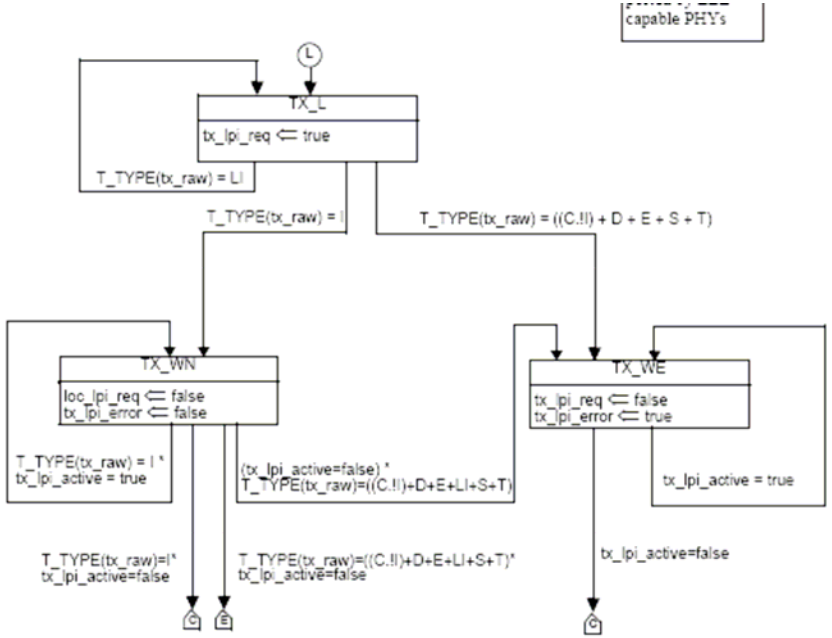


Figure 55–15a—PCS 64B/65B Transmit state diagram part b)

The new LPI 64B/65B Tx states do not control tx_coded

tx_coded is handled in a different state machine (55-16b) during Wake

It is desirable to make all tx_coded assignments in 55-15 and 55-15a (a single state machine)

Part of comment #380

# Problem 3) - Figure 55-15

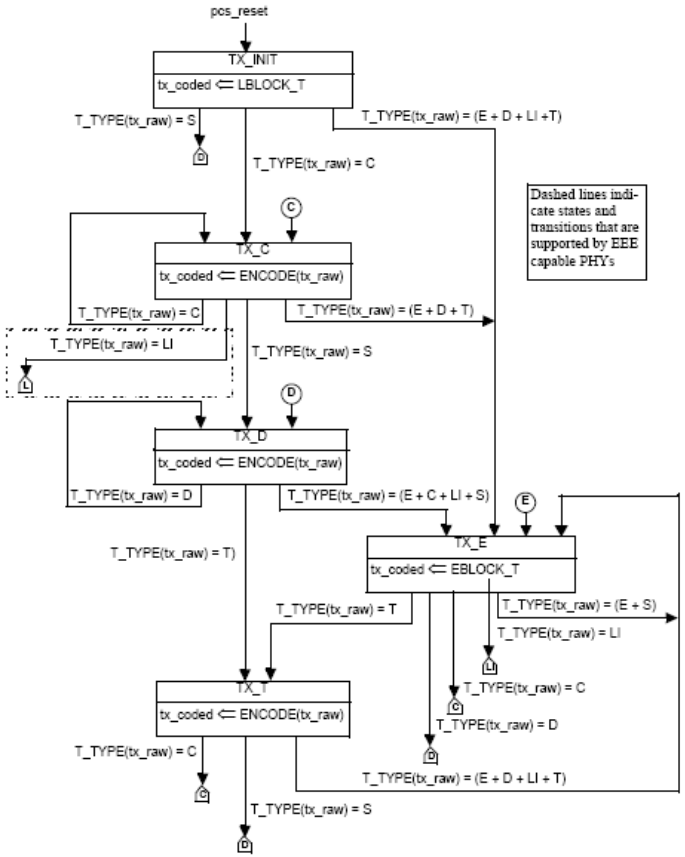


Figure 55–15—PCS 64B/65B Transmit state diagram part a)

The 64B/65B transmit state diagram with LI transition from TX_C

**Comments #142, #143**

**Problem:**

No transition from TX_T to TX_L

# Problem 2) and 6) - Fig 55-16b

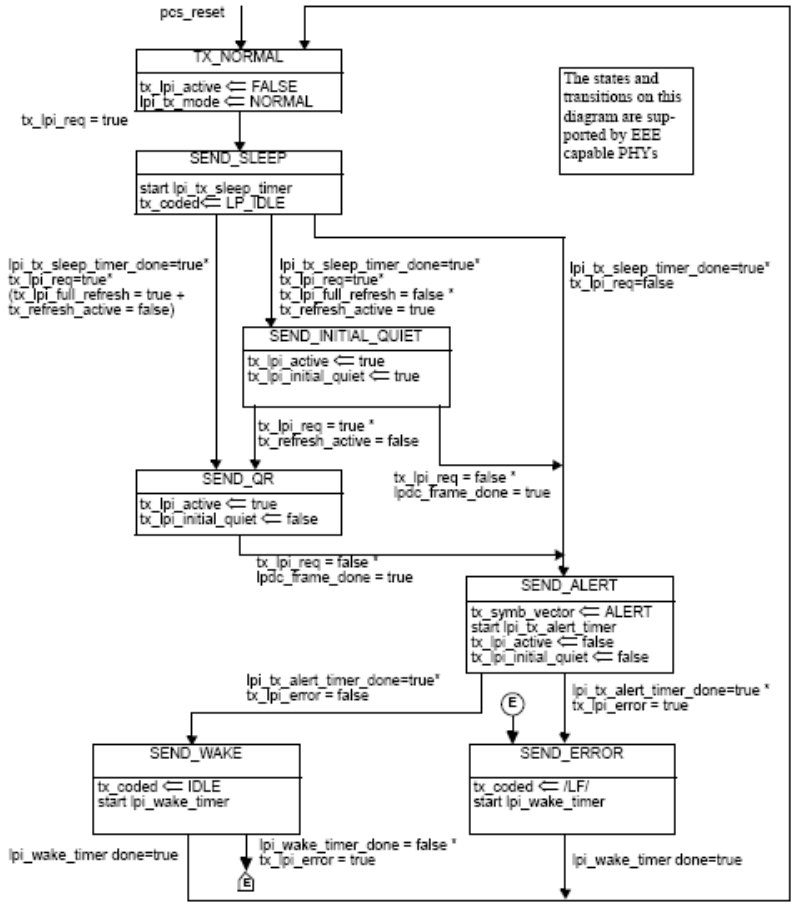


Figure 55–16b—EEE transmit state diagram

- tx_symb_vector is used in SEND_ALERT but not set anywhere else

•SEND_WAKE transitions into SEND_ERROR

Both states start lpi_wake_timer

Therefore 9 blocks of LF will be sent if any wake symbol is decoded incorrectly; it could result in 8 Wake Idles, followed by 9 LF blocks if an error is detected in the last block.

Is this desired behaviour?



SOLARFLARE Communications

# Other problems (#4)

- Block names are not correct [XGMII codeword names used instead of 65B block types]
  - **Use LP_BLOCK_T in preference to LP_IDLE**
  - **Same issue for local fault blocks and idle blocks**
    - Define new block types in 55.3.5.2.1
- Editorial
  - **Avoid boolean values in tests**

# Tx fixes : new definitions
# For problem #4

**Add the following definitions to 55.3.5.2.1**

LP_BLOCK_T<64:0>

65 bit vector to be sent to the LDPC encoder containing /LI/ in all the eight character locations

I_BLOCK_T<64:0>

65 bit vector to be sent to the LDPC encoder containing /I/ in all the eight character locations

LP_BLOCK_R<71:0>

72 bit vector to be sent to the XGMII interface containing /LI/ in all the eight character locations

I_BLOCK_R<71:0>

72 bit vector to be sent to the XGMII interface containing /I/ in all the eight character locations

# Tx fixes : Add new type, edit existing C / LI definitions For problem #5

- C: The vector meets one of the following criteria:

  a) contains eight valid control characters other than /O/, /S/, /T/ and /E/ and, if the low power idle function is supported, the vector is not LI or LII;

  ….

- LII: If the optional Low Power Idle function is supported then the LII type occurs when the vector contains four /LI/ control characters followed by four /I/ characters.
- LI : If the optional Low Power Idle function is supported then the LI type occurs when the vector contains eight /LI/ control characters.

# Tx fixes : Add new variable to 55.3.5.2.2 For problem #1

- Lpi_qr_active: set true when the PHY is transmitting quiet/refresh signaling. Set false otherwise

# Tx fixes : Amend lpi_tx_mode description I For problem #1

**Subclause 55.3.5.2.2; new text in red**

lpi_tx_mode

A variable indicating the signaling to be used from the PCS to the PMA across the PMA_UNITDATA.request (tx_symb_vector) interface.

lpi_tx_mode controls tx_symb_vector only when tx_mode is set to SEND_N.

The variable is set to NORMAL when !tx_lpi_qr_active * lpi_tx_alert_time_done indicating that the PCS is in the normal mode of operation and will encode code-groups as specified in Figures 55-15 and 55-15a.

The variable is set to REFRESH_A when tx_lpi_qr_active * (tx_active_pair==PAIR_A) * tx_refresh_active.

The variable is set to REFRESH_B when tx_lpi_qr_active * (tx_active_pair==PAIR_B) * tx_refresh_active.

The variable is set to REFRESH_C when tx_lpi_qr_active * (tx_active_pair==PAIR_C) * tx_refresh_active.

The variable is set to REFRESH_D when tx_lpi_qr_active * (tx_active_pair==PAIR_D) * tx_refresh_active.

The variable is set to QUIET when tx_lpi_qr_active * (!tx _refresh_active + tx_lpi_initial_quiet)

The variable is set to ALERT when tx_lpi_active * !tx_lpi_qr_active * !lpi_tx_alert_timer_done

# Tx fixes : Amend lpi_tx_mode description II For problems #1 and #3

- **Change 55.3.2.2.21, page 164; new text in red**

- **While the PMA asserts SEND_N, the lpi_tx_mode variable shall control the transmit signal through the PMA_UNITDATA.request primitive as described below:**

  - **During PMA training the lpi_tx_mode variable is ignored.**
  - **When the lpi_tx_mode variable takes the value NORMAL and the PMA asserts SEND_N the PCS passes coded data to the PMA via the PMA_UNITDATA.request primitive as described in 55.3.2.2.**
  - **When the lpi_tx_mode variable takes the value QUIET and the PMA asserts SEND_N the PCS passes zeros to the PMA through the PMA_UNITDATA.request primitive.**
  - **When the lpi_tx_mode variable takes the value REFRESH_A and the PMA asserts SEND_N the PCS passes the PMA training signal to the PMA on pair A, to allow both the local and remote PHY to refresh adaptive filters and timing loops. The PCS passes zeros to all other pairs in this condition. REFRESH_B, REFRESH_C and REFRESH_D operate in a analogous manner for the other pairs.**
  - **When the lpi_tx_mode variable takes the value ALERT and the PMA asserts SEND_N the PCS passes the ALERT vector to the PMA.**

# Tx fixes : change 64B/65B Tx state diagram 55-15 For problems #2 and #5

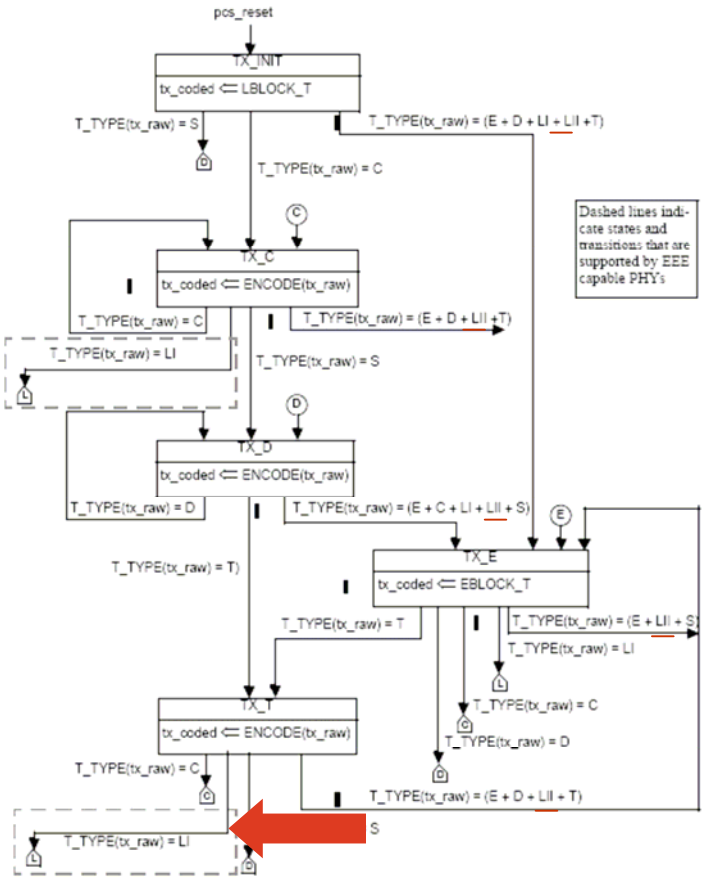


Figure 55–15—PCS 64B/65B Transmit state diagram part a)

- Added LII transitions
- Added transition from TX_T to TX_L
- Fixed connection from TX_E to TX_L
- Removed /LI/ transition from TX_T to TX_E



SOLARFLARE Communications

# Tx fixes : change 55-16b EEE Transmit state diag Problems #1, #3, #6

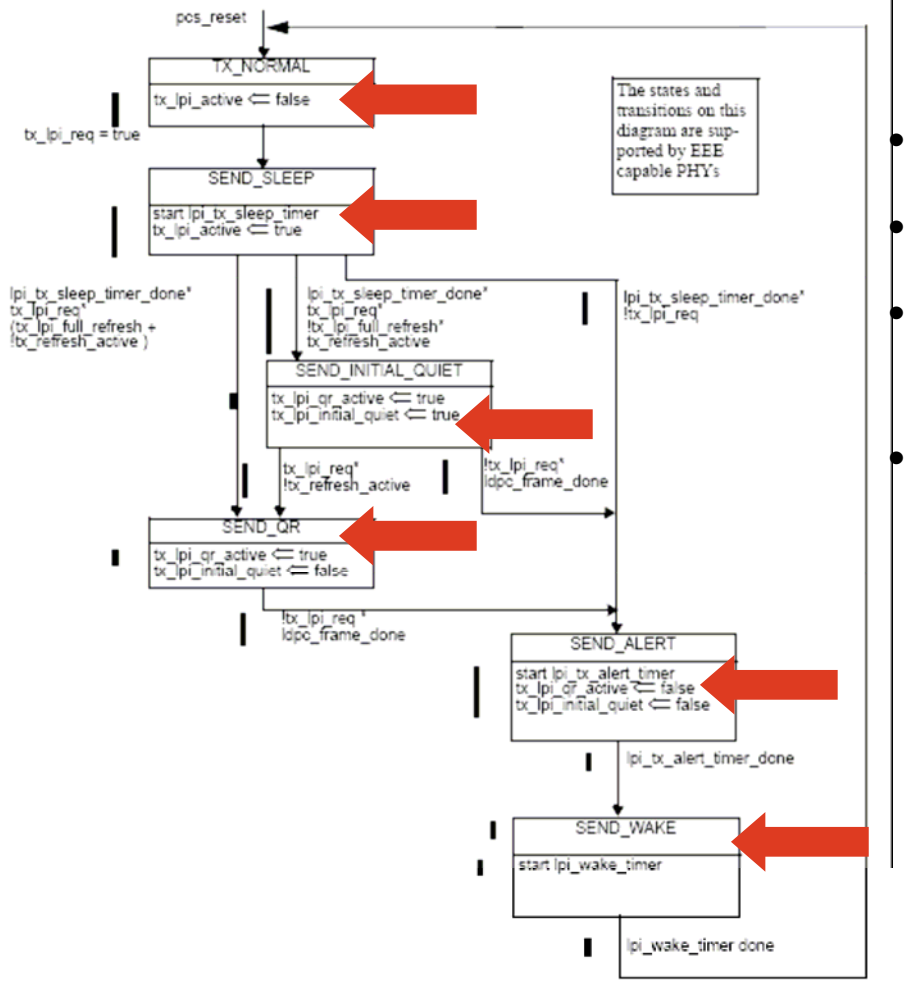


Figure 55–16b—EEE transmit state diagram

- Added tx_lpi_qr_active variable
- Fixed tx_lpi_active assignments
- Removed tx_coded assignments
- Removed tx_symb_vector <= ALERT
  - Ambiguous: tx_symb_vector is not set anywhere else in this state machine so it is not clear when tx_symb_vector is reset
  - Use lpi_tx_mode instead



SOLARFLARE Communications

# Comment #377

- Comment #377 may be controversial and stimulate debate
- Updates to 55-15a with and without that change are shown on the next two slides
- Note also that 55-15 and 55-15a can be consolidated into one figure if #377 is accepted

# If comment #377 is accepted:
## Tx fixes : change 64B/65B Tx state diagram 55-15a

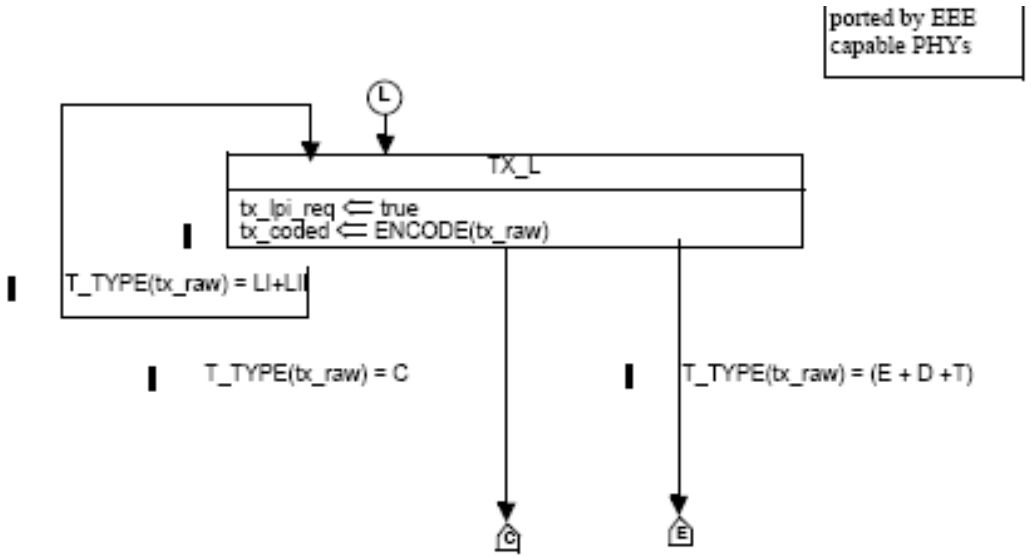


Figure 55–15a—PCS 64B/65B Transmit state diagram part b)

- Pass codewords on XGMII through PCS datapath
- Does not force values onto tx_coded through state machine
- Removes wake states and simplifies state machine



SOLARFLARE Communications

# If comment #377 is accepted:
## Tx fixes : change 64B/65B Tx state diagram 55-15

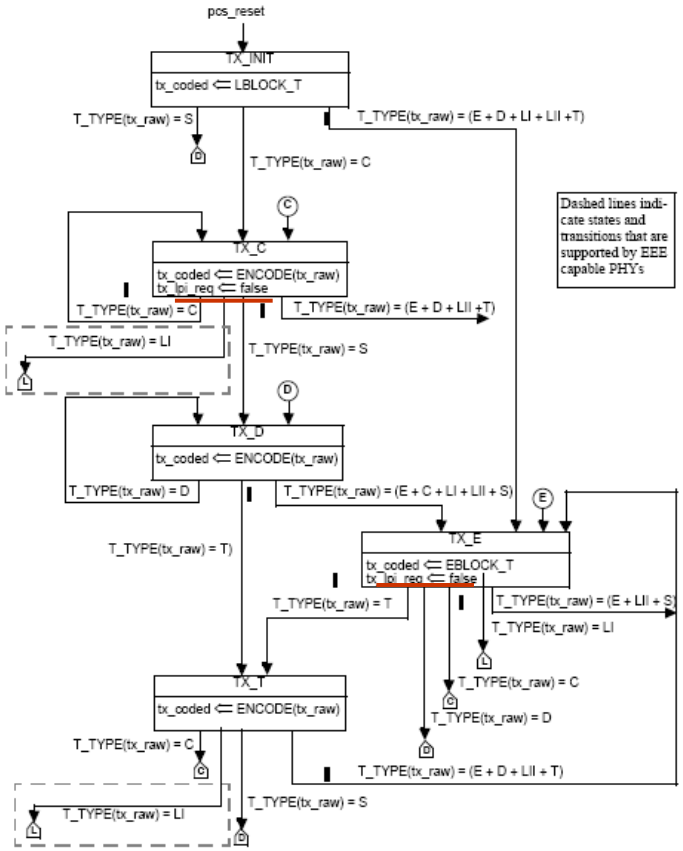


Figure 55–15—PCS 64B/65B Transmit state diagram part a)

- Reset tx_lpi_req in TX_C and TX_E
- Make it clear that these assignments are only required for EEE capable PHYs [box the assignments?]



SOLARFLARE Communications

# If comment #377 is rejected
## Tx fixes : change 64B/65B Tx state diagram 55-15a

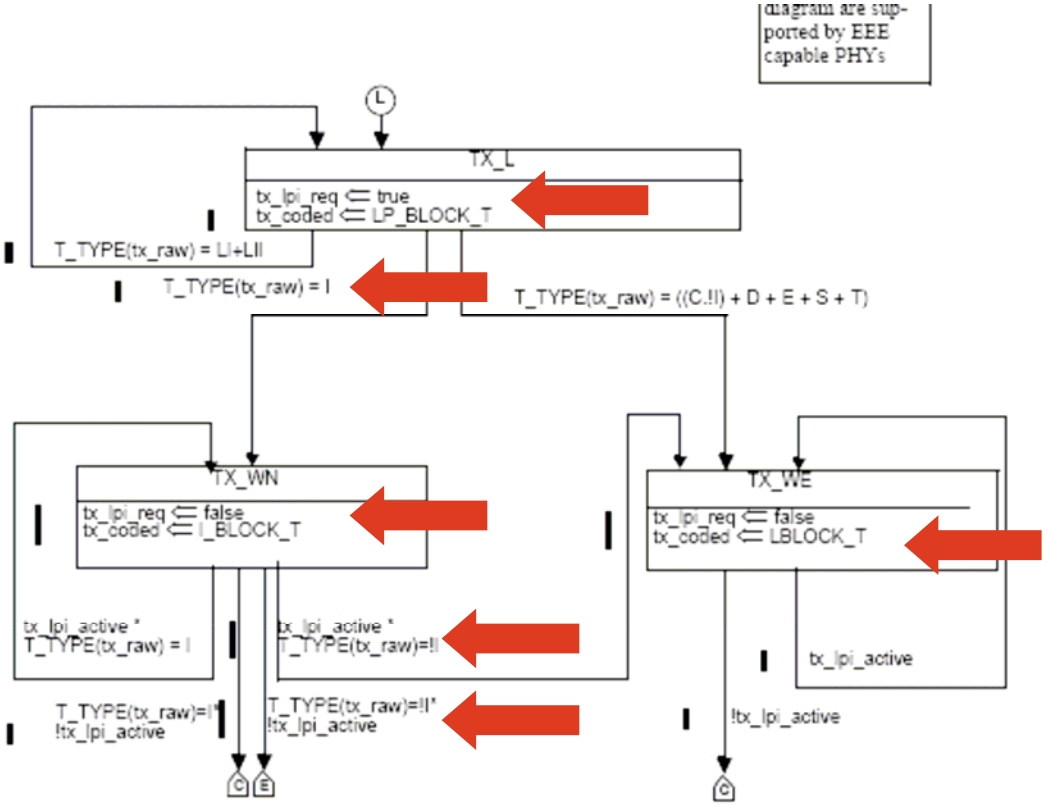


Figure 55–15a—PCS 64B/65B Transmit state diagram part b)

- Added tx_coded assignments to these states
- Corrected block names / functions
- Added LII transitions
- Fixed tx_lpi_active on TX_WN to TX_WE branch



SOLARFLARE Communications

# Conclusion – Tx side

- Fixed ambiguity in transmit state machines
- Addresses comments
  - **Use XGMII signaling: #377; accept/reject options are included in presentation**
  - **tx_lpi_active used incorrectly : #380**
  - **LP_IDLE_4+4 : #102, #93, #94, #95**
  - **/LF/ usage: #84, #379**
  - **Transition from TX_T to TX_L: #142, #143**
  - **loc_lpi_req typo for tx_lpi_req: #376, #89**
  - **Typo TX_WN -> TX_WE (tx_lpi_active): #378**
  - **Typo TX_E box should go to TX_L : #90**
  - **Editorials: #79, #376**

# Receive Side

- E: 80 (ER), 144
- T: 143, 96 (TR), 91, 379 (TR)

- Technical issues:
  - **PMA rx state machine is redundant**
  - **Missing transition from RX_T to RX_L**
  - **Names used in assignments to rx_raw are not correct**

# Rx fixes: Edit 55-16

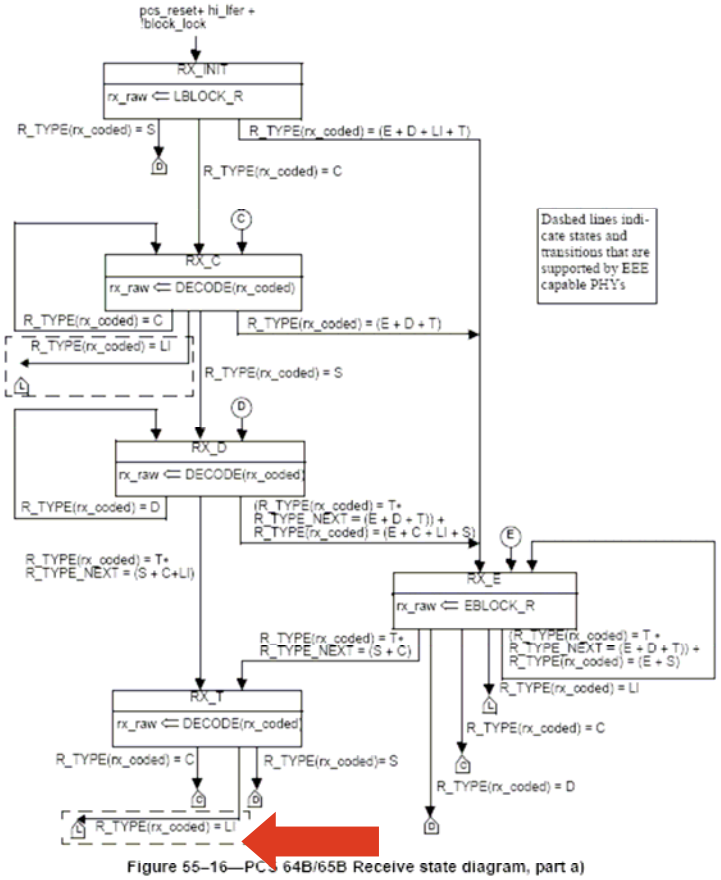


Figure 55–16—PCS 64B/65B Receive state diagram, part a)

- Added LI transition to RX_T state
  - Analogous transition added to tx side

# Rx fixes: Edit 55-16a and delete 55-27a

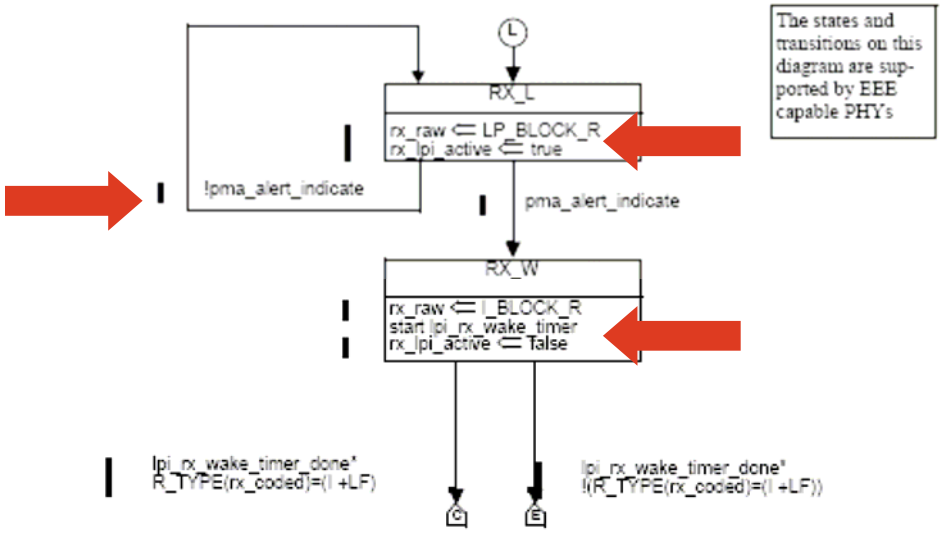


Figure 55–16a—PCS 64B/65B Receive state diagram, part b)

- Delete PMA Rx state diagram
  - Pass pma_alert_indicate to PCS
  - Delete references to rx_lpi_req
  - Simplifies state diagrams without changing functionality
- Fix rx_raw assignments
  - Incorrect names were used
- **Functionally exactly the same as draft 2.0**



SOLARFLARE Communications

# Conclusions – Rx side

- Simplified & corrected receive state machines
- Addresses comments
  - **RX_T transitions (96)**
  - **Remove redundant Boolean variables (80, 81)**
  - **Remove PMA state diagram (91)**
  - **Typo (144)**
  - **Incorrect block assigments (379)**