ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

Updates to the Text for the Draft for Energy Efficient Ethernet and LPI Signalling

Philip Curran
Brian Murray
Jacobo Riesco

Introduction

- This presentation presents an update to the text for Energy Efficient Ethernet and LPI Signalling required for the draft
  - There have been a number of presentations on EEE, all of which generally follow the scheme of clauses 149 and 165
    - Including a description of the EEE abilities advertised in the InfoField
    - A description of the alternating quiet and refresh periods, and the LPI signalling and synchronization
    - And a description of the LPI Timing Parameters
- We have adopted the following motion on EEE
  - Nov 2024 Motion #4:
    Move that the IEEE P802.3dg Task Force adopt slides 4 to 7 of Murray_3dg_03a_11132024
    - EEE LPI quiet-refresh timing and definition of the auxiliary bit for signaling insufficient LPI refresh
- However, much of the detail is missing and there are major gaps in the text
  - The section 190.3.3.12 EEE capability is empty
  - The section 190.3.8 Detailed functions and state diagrams, including the sub-sections on constants, variables, timers, functions, messages and state diagrams are all empty

Energy Efficient Ethernet – Transmit Direction

- The optional 100BASE-T1L EEE capability allows PHYs to transition to an LPI mode of operation when link utilization is low in either direction
  - In the transmit direction, the transition to the LPI transmit mode begins when 2N + 8 consecutive MII transfers represent Assert LPI
  - The PHY transmits the sleep signal to indicate to the link partner that it is transitioning to the LPI transmit mode
  - The sleep signal is composed of 8 PCS partial frame periods within which each (8N + 1)B block is constructed by encoding N /LI/ control symbols
    - Transmission of the sleep signal may start at the beginning of any multiple of 16 PCS partial frame periods offset to allow transmission of the alert signal to start as soon as transmission of the sleep signal finishes
  - Following the transmission of the sleep signal, quiet-refresh signaling begins
  - The quiet-refresh cycle is repeated until an MII transfer does not represent Assert LPI
    - Or until the eee_low_snr variable is set TRUE, or the rem_eee_low_snr variable is set TRUE
  - Following any of these events, the PHY transmits the alert signal to indicate to the link partner that it is transitioning back to the normal operational mode
    - Transmission of the alert signal may start at the beginning of any multiple of 16 PCS partial frame periods starting at the beginning of the PCS partial frame that follows the refresh period
  - After the transmission of the alert signal, the PCS completes the transition from LPI transmit mode back to the normal operational mode by sending a wake signal
    - The wake signal is composed of 8 PCS partial frame periods (wake_time) within which each (8N + 1)B block is constructed by encoding N /I/ control symbols

Energy Efficient Ethernet – LPI Signaling

- The signal timing is shown in Figure 190–11—LPI signal timing
  - The PHY uses a repeating quiet-refresh cycle using the LPI timing parameters are in Table 190–8
  - The timing parameters are all integer multiples of the PCS partial frame period
  - The LEADER and FOLLOWER refresh periods and alert windows are offset

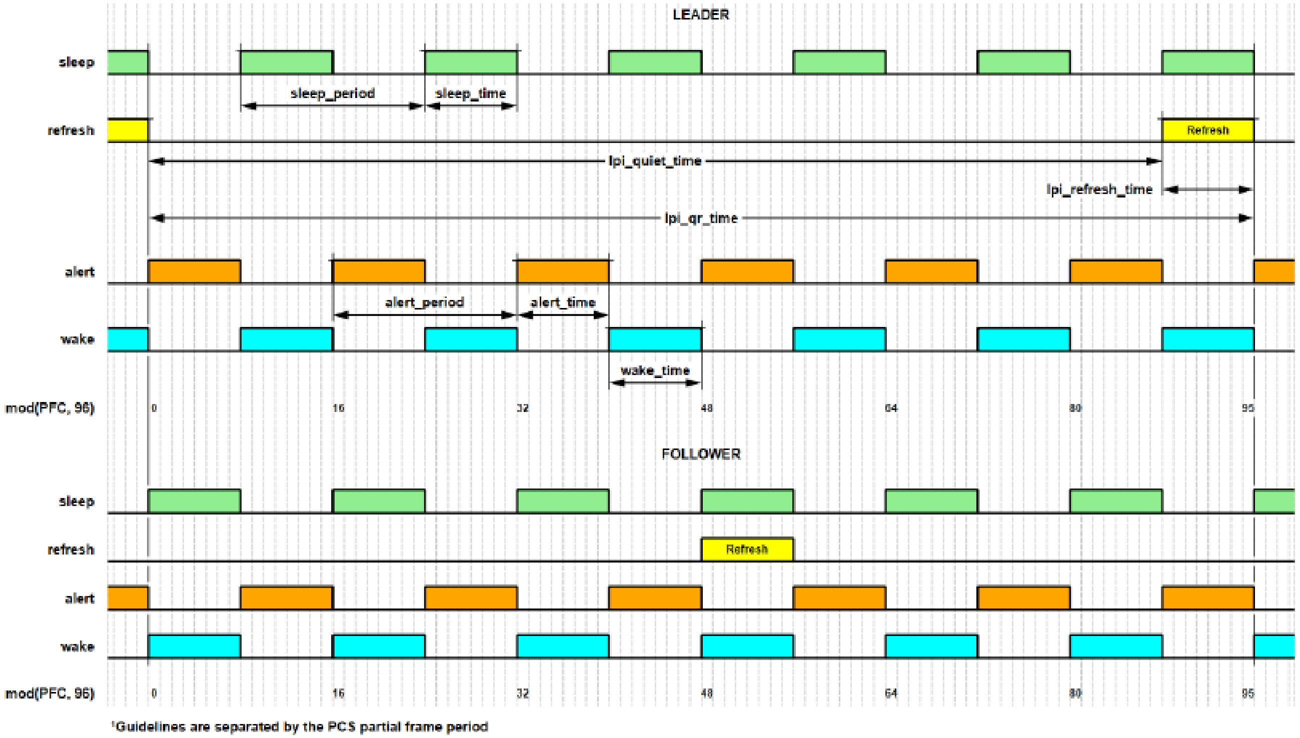

Figure 190–11—LPI signal timing

Table 190–8—LPI timing parameters

Parameter	Number of PCS partial frame periods	Duration (µs)
sleep_period	16	38.4
sleep_time	8	19.2
lpi_offset	56	–
lpi_qr_time	96	230.4
lpi_quiet_time	88	211.2
lpi_refresh_time	8	19.2
alert_period	16	38.4
alert_time	8	19.2
wake_time	8	19.2

Energy Efficient Ethernet – Alignment

- The FOLLOWER must align its partial frame count (PFC) modulo 96 with that of the LEADER
  - As the lpi_qr_time is 96 times the PCS partial frame period
- Alignment is achieved as follows
  - The start of the training frame transmitted by the FOLLOWER shall be delayed by not more than 1 PCS partial frame with reference to the start of the training frame received from the LEADER, as seen at the MDI of the FOLLOWER
    - This aligns the PFC modulo 16
    - To achieve the remaining alignment the LEADER includes a formatted training frame count (FTFC) in its InfoField

$$\mathrm{FTFC} = \mathrm{mod}(\mathrm{PFC},\ \mathrm{lpi\_qr\_time}) \gg 4$$

    - The FTFC is in Octet 7 of the InfoField

Octet 1 <7:0>	Octet 2 <7:0>	Octet 3 <7:0>	Octet 4 <7:0>	Octet 5 <7:0>	Octet 6 <7:0>	Octet 7 <7:0>	Octet 8 <7:0>	Octet 9 <7:0>	Octet 10 <7:0>	Octet 11 <7:0>	Octet 12 <7:0>
0xEE	0xA7	0x00	Reserved			FTFC	PHY Capability			CRC16	

Energy Efficient Ethernet – Alert Signal

- When RS FEC is enabled there is a large propagation delay through the PCS from transmit to receive
- To avoid the situation where this large propagation delay causes excessive wake time we introduced an Alert signal
  - This signal is injected directly at the interface to the PMA and can be detected directly at the interface to the PMA
    - This is done by injecting 1's into the scrambler
    - The Alert signal can be readily distinguished from the Refresh signal, which is done by injecting 0's into the scrambler

Energy Efficient Ethernet – Transmit State Diagram

- PCS (8N)B/(8N + 1)B Transmit state diagram
  - This state diagram determines tx_coded

    Vector containing the output from the (8N)B/(8N+1)B encoder
  - The variable tx_lpi_req is defined differently from clause 149

    tx_lpi_req
    Boolean variable that is set TRUE when EEE is enabled for the link, and eee_low_snr is FALSE, and rem_eee_low_snr is FALSE, and 2N + 8 consecutive MII transfers, including the 2N transfers represented by the elements of the tx_mii array, indicate Assert LPI. It is set FALSE otherwise.
    - ➢ The variables eee_low_snr and rem_eee_low_snr are included in the generation of tx_lpi_req
  - The variable tx_lpi_enable now drives the operation of the EEE transmit state diagram
  - No /LI/ control symbols are sent to the line until this state diagram enters TX_SLEEP
    - Once this state diagram enters TX_SLEEP a complete sleep signal will be sent
  - In previous clauses the state diagram can transition from TX_SLEEP to TX_WAKE as soon as transmission of the sleep signal begins
    - This has been resolved by introducing the WAIT_ALERT state

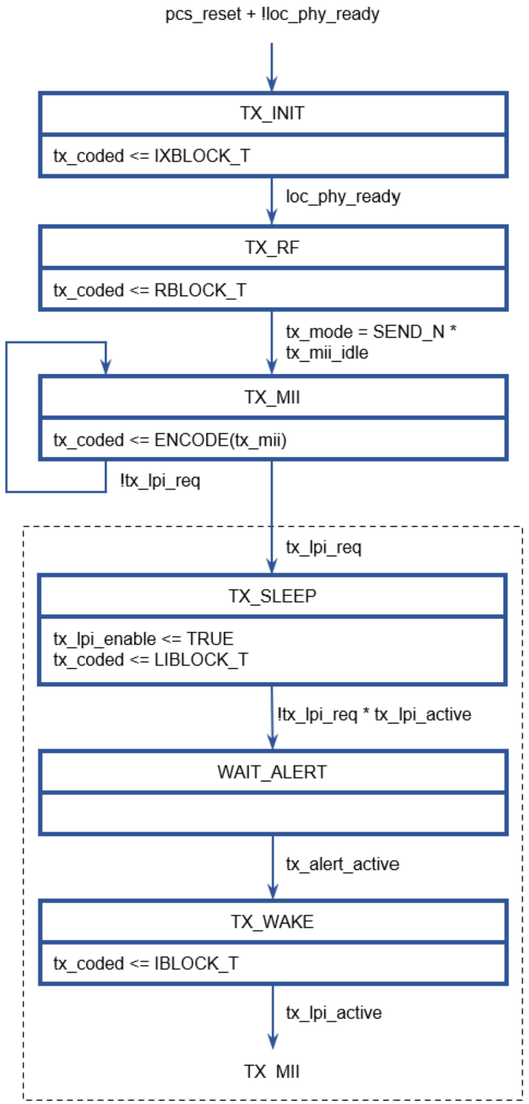

NOTE—Signals and functions shown with dashed lines are only required when EEE is enabled for the link.

Figure 190–12—PCS (8N)B/(8N + 1)B Transmit state diagram

Energy Efficient Ethernet – EEE Transmit State Diagram

- EEE transmit state diagram
  - Added SYNC_SLEEP state to deal with transmit latency
    - All previous clauses have specified that the full sleep signal is transmitted once entry to LPI is initiated
    - But latency in the transmit path means that this is never implemented in practice

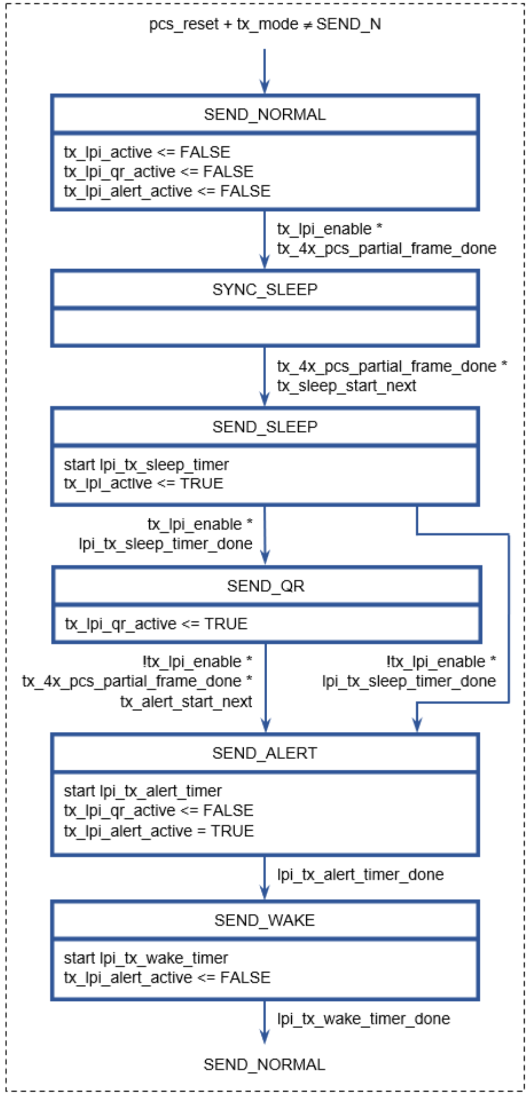

NOTE—This figure is mandatory when EEE is enabled for the link.

Figure 190–13A—EEE transmit state diagram

Energy Efficient Ethernet – Receive Direction

- The receive side is similar to what is done in clauses 149 / 165 with the exception that we have added the RX_ALERT state
  - If we go directly from RX_LPI to RX_WAKE then by the time the wake_timer expires the wake signal from the link partner may not have propagated through the receive path
    - When RS-FEC is enabled the receive path latency is very large
  - We detect that the alert signal at the PMA interface has passed and we infer that the wake signal has arrived at the PMA to PCS interface
  - The whole of the lpi_rx_wake_timer period is available to allow for latency in the receive path

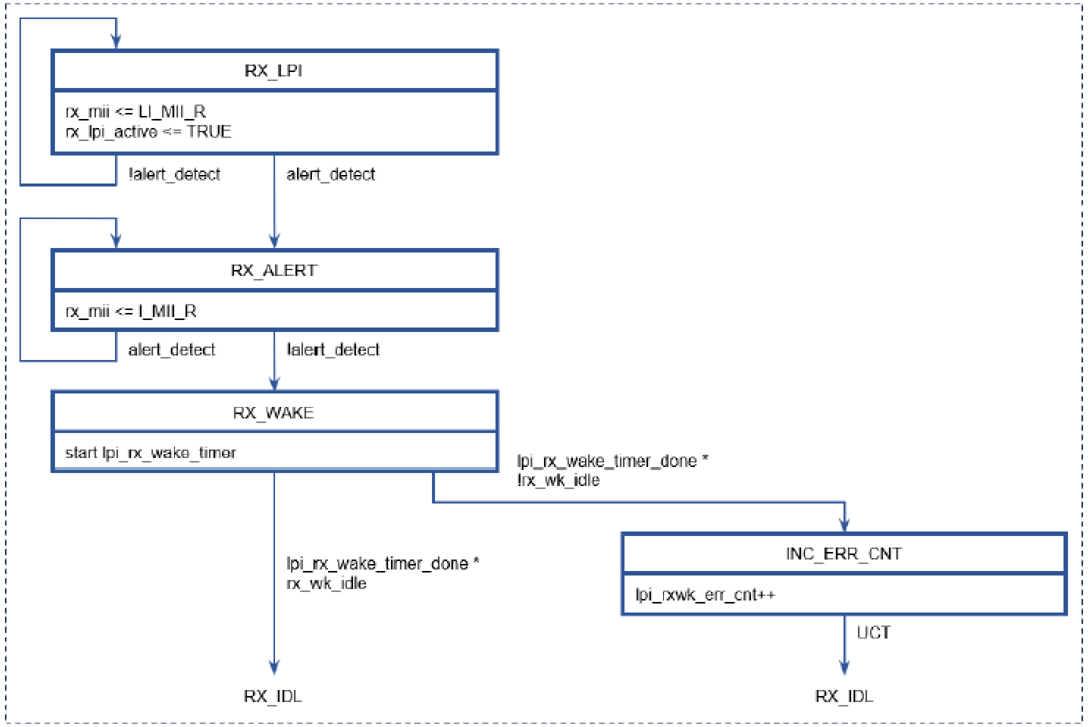

NOTE—This figure is mandatory when EEE is enabled for the link.

Figure 190–13B—PCS Receive state diagram, part b

Energy Efficient Ethernet – Receive Direction

- Entry to RX_LPI happens when rx_lpi_sleep becomes TRUE
  - The variable rx_lpi_sleep is defined as follows

    rx_lpi_sleep
    Boolean variable that is set TRUE when 32 consecutive rx_char values each represent /LI/. It is set FALSE otherwise.
  - Hence the PCS Receive state diagram enters RX_LPI before the sleep signal from the link partner has finished
  - The sleep signal and the alert signal cannot be confused with each other, this is just by the way they have been designed

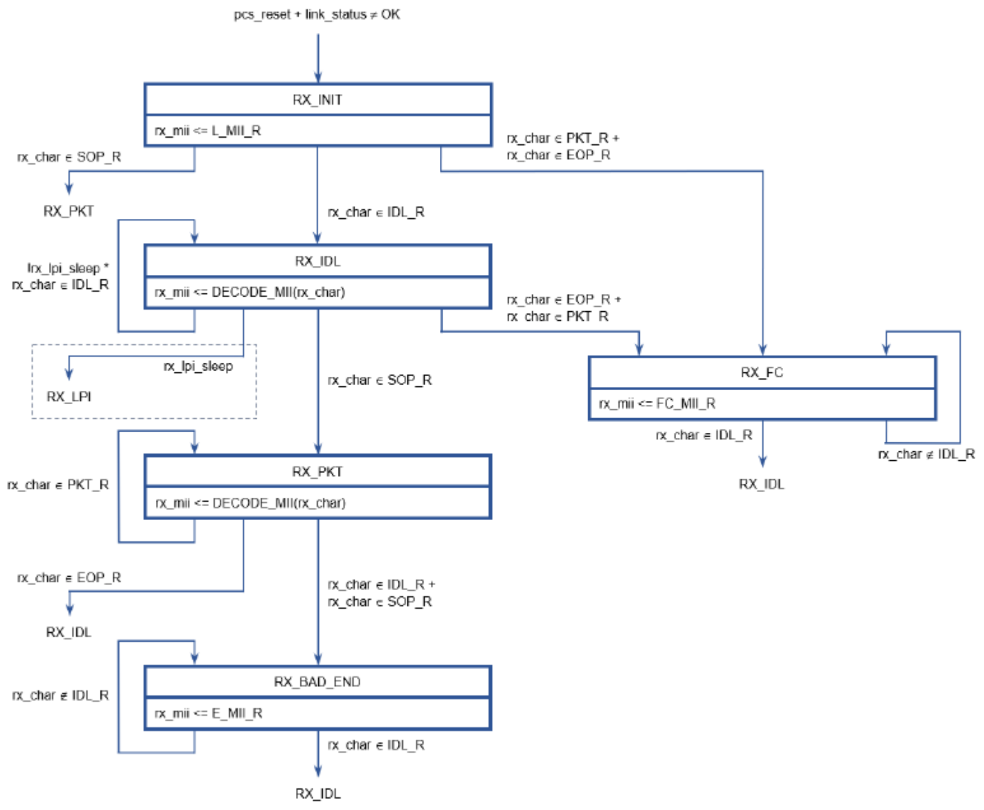

NOTE—Signals and functions shown with dashed lines are only required when EEE is enabled for the link.

Figure 190–13B—PCS Receive state diagram, part a

Energy Efficient Ethernet – Wake Time

▶ Maximum Wake Time

- If the PCS Transmit function initiates the transition back to the normal operational mode before the sleep signal has been completely transmitted
  - The maximum duration of the PHY wake time increases to 44 PCS partial frame periods (4 PCS partial frame periods + sleep_period + sleep_time + alert_time + wake_time = 105.6 µs)
- This is comparable to the maximum wait time for 1000BASE-T1 in bit times

Table 78–4—Summary of the LPI timing parameters for supported PHYs or interfaces

PHY or interface type	Case	$T_{w\_sys\_tx}$ (min) (µs)	T_{w_phy} (min) (µs)	$T_{phy\_shrink\_tx}$ (max) (µs)	$T_{phy\_shrink\_rx}$ (max) (µs)	$T_{w\_sys\_rx}$ (min) (µs)
10BASE-T1L		270	250.5	10	240	20
100BASE-TX		30	20.5	5	15	10
1000BASE-KX		13.26	11.25	5	6.5	1.76
1000BASE-T1		10.8	10.8	10.8	0[a]	0[a]
1000BASE-RHC 1000BASE-RHA 1000BASE-RHB		25	25	25	0	0
1000BASE-T	Case-1	16.5	16.5	5	2.5	1.76
	Case-2	16.5	16.5	12.24	9.74	1.76

Proposed Text for the Draft

- The following slides include the proposed changes to the text of draft 1.0 required for the draft
  - A number of sections, tables and state diagrams are completely new and are shown inside a blue outline text box ☐ with a 'New Text' or 'New Figure' or 'New Table' pointer
  - Some sections have existing text that has been completely rewritten, in these cases the new text (inside ☐) is shown side by side with the original text inside a red outline text box ☐
  - Original text boxes have a header *IEEE P802.3dg™/D1.0, 29th April 2025*
    - Original text included for reference that is being kept, is shown inside a brown outline text box ☐
- Note that there are additional Tables and Figures
  - The same numbering is used for existing tables and figures as used in draft 1.0 so that reference to existing draft 1.0 tables and figures is easier
  - New proposed figures have been given numbers 13A, 13B and 13C, so that figures 14 and above have the same figure numbers as draft 1.0
  - The tables and figures will be renumbered in the draft
  - Reference to table or figures whose number will change in the draft are show in red

Proposed Text for the Draft - Transmit Process

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

- Text change for section 190.3.3.6 Transmit process
  - Add sentence on EEE between lines 16 and 17 on page 66

IEEE P802.3dg™/D1.0, 29th April 2025

The first (8N)B/(8N + 1)B block transmitted is tx_coded<0>.

New Text

When EEE is not enabled for the link, the aux bit is zero. When EEE is enabled for the link, the aux bit is used to communicate the value of the eee_low_snr parameter, as described in 190.3.3.12.

When RS-FEC is disabled, N is 2, and the group of 15N + 2 octets form a PCS frame of 256 bits.

When RS-FEC is enabled, N is 8, and the group of 15N + 2 octets are encoded by the RS-FEC encoder which adds 6 parity octets. The resulting 128 octets form a PCS frame of 1,024 bits.

Proposed Text for the Draft – EEE Capability

- Text for section 190.3.3.12 EEE capability
  - This section is empty, insert the following text in section 190.3.3.12

IEEE P802.3dg™/D1.0, 29th April 2025

190.3.3.12 EEE capability

New Text

The optional 100BASE-T1L EEE capability allows PHYs to transition to an LPI mode of operation when link utilization is low in either direction of transmission.

When EEE is enabled for the link, the PHY shall implement the PCS (8N)B/(8N + 1)B Transmit state diagram including the EEE portion, noted by dotted lines in Figure 190–12, and shall conform to the EEE transmit state diagram, shown in Figure 190–13A.

All EEE timing parameters are expressed as integer multiples of the PCS partial frame period as shown in Table 190–8. The PCS partial frame is defined in 190.3.3.1.

The eee_low_snr parameter communicated through the PMA_EEE_LOW_SNR.indication primitive indicates whether the SNR is too low to maintain reliable operation in LPI receive mode. The aux bit of every group of transmit bits, tx_group, is set to 1 when eee_low_snr is TRUE and is set to 0 otherwise. The variable rem_eee_low_snr indicates the value of the eee_low_snr variable communicated by the remote PHY.

In the transmit direction, the transition to the LPI transmit mode begins when the eee_low_snr variable is FALSE, and the rem_eee_low_snr variable is FALSE, and when 2N + 8 consecutive MII transfers represent Assert LPI. Following this event, the PHY transmits the sleep signal to indicate to the link partner that it is transitioning to the LPI transmit mode. The sleep signal is composed of 8 PCS partial frame periods (sleep_time) within which each (8N + 1)B block is constructed by encoding N /LI/ control symbols. Once the PCS Transmit function initiates the transition to the LPI transmit mode, the complete sleep signal is transmitted.

Proposed Text for the Draft – EEE Capability

- Continue new text for section 190.3.3.12 EEE capability

New Text

Following the transmission of the sleep signal, quiet-refresh signaling, as described in 190.3.6, begins.

When the tx_lpi_active variable is TRUE, the tx_lpi_qr_active, tx_refresh_active and tx_alert_active variables shall control the transmit signal as follows:

> When the tx_lpi_qr_active variable is FALSE and the tx_alert_active variable is FALSE, the PCS passes coded data to the PMA through the PMA_UNITDATA.request primitive as described in 190.3.2.
>
> When the tx_lpi_qr_active variable is TRUE and the tx_refresh_active variable is FALSE, the PHY transmits the quiet signal as described in 190.3.7.2.
>
> When the tx_lpi_qr_active variable is TRUE and the tx_refresh_active variable is TRUE, the PHY transmits the refresh signal as described in 190.3.7.3.
>
> When the tx_alert_active variable is TRUE, the PHY transmits the alert signal as described in 190.3.7.4.

The quiet-refresh cycle is repeated until the eee_low_snr variable is set TRUE, or the rem_eee_low_snr variable is set TRUE, or an MII transfer does not represent Assert LPI. Following any of these events, the PHY transmits the alert signal, as described in 190.3.7.4, to indicate to the link partner that it is transitioning back to the normal operational mode.

After the transmission of the alert signal, the PCS completes the transition from LPI transmit mode back to the normal operational mode by sending a wake signal. The wake signal is composed of 8 PCS partial frame periods (wake_time) within which each (8N + 1)B block is constructed by encoding N /I/ control symbols.

Proposed Text for the Draft – EEE Capability

- Continue new text for section 190.3.3.12 EEE capability

New Text

Transmission of the sleep signal may start at the beginning of any multiple of 16 PCS partial frame periods (sleep_period), offset to allow transmission of the alert signal to start as soon as transmission of the sleep signal finishes. Transmission of the alert signal may start at the beginning of any multiple of 16 PCS partial frame periods (alert_period) starting at the beginning of the PCS partial frame that follows the refresh period. The maximum duration of the PHY wake time (T_{w_phy} as defined by Clause 78) is 32 PCS partial frame periods (alert_period + alert_time + wake_time = 76.8 µs), if the PCS Transmit function initiates the transition back to the normal operational mode after the sleep signal has been completely transmitted. The maximum duration of the PHY wake time increases to 44 PCS partial frame periods (4 PCS partial frame periods + sleep_period + sleep_time + alert_time + wake_time = 105.6 µs), if the PCS Transmit function initiates the transition back to the normal operational mode before the sleep signal has been completely transmitted.

Proposed Text for the Draft - PMA Training

- Continue changes for section 190.3.5 PMA Training
  - Replace text on page 75, lines 26 to 27 and update Figure 190–9—InfoField format

IEEE P802.3dg™/D1.0, 29th April 2025

The InfoField shall include the fields in 190.3.5.2.2 through 190.3.5.2.4 as shown in Figure 190–9.

Octet 1 <7:0>	Octet 2 <7:0>	Octet 3 <7:0>	Octet 4 <7:0>	Octet 5 <7:0>	Octet 6 <7:0>	Octet 7 <7:0>	Octet 8 <7:0>	Octet 9 <7:0>	Octet 10 <7:0>	Octet 11 <7:0>	Octet 12 <7:0>
0xEE	0xA7	0x00	Reserved				PHY Capability			CRC16	

Figure 190–9—InfoField format

New Text

The InfoField shall include the fields in 190.3.5.2.2 through 190.3.5.2.5, also shown in Figure 190–9.

New Figure

Octet 1 <7:0>	Octet 2 <7:0>	Octet 3 <7:0>	Octet 4 <7:0>	Octet 5 <7:0>	Octet 6 <7:0>	Octet 7 <7:0>	Octet 8 <7:0>	Octet 9 <7:0>	Octet 10 <7:0>	Octet 11 <7:0>	Octet 12 <7:0>
0xEE	0xA7	0x00	Reserved			FTFC	PHY Capability			CRC16	

Figure 190–9—InfoField format

Proposed Text for the Draft - PMA Training

- Text changes for section 190.3.5 PMA Training
  - Add paragraph in section 190.3.5.2 PCS frame alignment and advertisement of PHY capabilities on page 75 between lines 16 and 17 on page 66

IEEE P802.3dg™/D1.0, 29th April 2025

The start of the training frame transmitted by the FOLLOWER shall be delayed by not more than 1 PCS partial frame with reference to the start of the training frame received from the LEADER, as seen at the MDI of the FOLLOWER.

New Text

When the config parameter is LEADER and EEE is supported, the PHY incorporates a formatted training frame count (FTFC) into the transmitted InfoField to indicate the alignment of the formatted training frame within the quiet-refresh cycle described in 190.3.6. When the config parameter is FOLLOWER and EEE is enabled for the link, the FOLLOWER shall use the FTFC value received from the LEADER to align its quiet-refresh cycle to that of the LEADER as specified in 190.3.6.

At each code-group time n, bits $Sd_n[3:0]$ are generated from scrambler bits $Sd_n[3:0]$ as follows

$$Sd_n[3] = \begin{cases} Sy_n[3] \wedge InfoField\,[(4n+3) \bmod 128] & \text{if } (480 \le n \bmod 512 \le 503) \\ Sy_n[3] & \text{else} \end{cases}$$

Proposed Text for the Draft - PMA Training

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

- Continue text changes for section 190.3.5 PMA Training
  - Add new section 190.3.5.2.3 Formatted Training Frame Count between sections 190.3.5.2.2 Start Delimiter and 190.3.5.2.3 PHY capability bits on page 76

IEEE P802.3dg™/D1.0, 29th April 2025

190.3.5.2.2 Start Delimiter

The Start Delimiter consists of three octets [Octet 1<7:0>, Octet 2<7:0>, Octet 3<7:0>] and shall use the hexadecimal value 0xEEA700. 0xEE corresponds to Octet 1 <7:0> and so forth.

New Section

190.3.5.2.3 Formatted Training Frame Count

When the config parameter is LEADER and EEE is supported, Octet 7<7:0> shall be set equal to the value of FTFC, which indicates the alignment of the formatted training frame within the quiet-refresh cycle. FTFC is defined as follows:

FTFC = mod(PFC, lpi_qr_time) >> 4

When the config parameter is FOLLOWER or EEE is not supported, bits Octet 7<7:0> shall be set to zero.

Inc to 190.3.5.2.4

190.3.5.2.3 PHY capability bits

The PHY capabilities are advertised in the three octets [Oct8<7:0>, Oct9<7:0>, Oct10<7:0>]. See Table 190–6 for details.

Proposed Text for the Draft – LPI Signaling

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

- Text for section 190.3.7 LPI signaling
  - Replace text on page 79, lines 37 to 54 and line 1 of page 79

New Text

190.3.7 LPI signaling

PHYs with EEE capability have transmit and receive functions that can enter and leave the LPI mode independently. The PHY can transition to the LPI mode when the PHY has successfully completed training, and the tx_mode parameter communicated through the PMA_TXMODE.indication primitive is SEND_N. The transmit function of the PHY initiates a transition to the LPI transmit mode by generating the sleep signal as described in 190.3.3.12. When transmission of the sleep signal begins, the PCS Transmit function asserts tx_lpi_active and the transmit function enters the LPI transmit mode. The sleep signal is restricted to starting at predetermined PFC values as described in 190.3.7.1.

Within the LPI mode, a PHY use a repeating quiet-refresh cycle (see Figure 190–11). The LPI timing parameters are shown in Table 190–8. One part of this cycle is known as the quiet period and lasts for a time lpi_quiet_time. The quiet period signaling is defined in 190.3.7.2. The other part of this cycle is known as the refresh period and lasts for a time lpi_refresh_time. The refresh period signaling is defined in 190.3.7.3. A cycle composed of one quiet period and one refresh period is known as an LPI cycle and lasts for a time lpi_qr_time.

The parameters sleep_period, sleep_time, lpi_offset, lpi_qr_time, lpi_quiet_time, lpi_refresh_time, alert_period, alert_time, and wake_time are timing parameters that are integer multiples of the PCS partial frame period. The parameter lpi_offset is a fixed value equal to lpi_qr_time / 2 + alert_period / 2 (56) PCS partial frames. This offsets the LEADER and FOLLOWER refresh periods and alert windows as shown in Figure 190–11.

The end of LPI transmit mode occurs at the transmission of the alert signal indicating the end of the quiet-refresh cycle. The alert signal is restricted to starting at predetermined PFC values as described in 190.3.7.1.

IEEE P802.3dg™/D1.0, 29th April 2025

190.3.7 LPI signaling

A PHY with EEE capability has transmit and receive functions that can enter and leave the LPI mode independently. The PHY can transition to the LPI mode when the PHY has successfully completed training and pcs_data_mode is TRUE. The transmit function of the PHY initiates a transition to the LPI transmit mode by generating the sleep signal composed of TBD, each composed entirely of LPI control characters, as described in 199.3.2.3.17. When the transmitter begins to send the sleep signal, it asserts tx_lpi_active and the transmit function enters the LPI transmit mode.

Within the LPI mode PHYs use a repeating quiet-refresh cycle (see Figure 190–11). The first part of this cycle is known as the quiet period and lasts for a time lpi_quiet_time equal to 88 partial PHY frame periods. The quiet period is defined in 190.3.3.2. The second part of this cycle is known as the refresh period and lasts for a time lpi_refresh_time equal to 8 partial PHY frame periods. The refresh period is defined in 190.3.3.3. A cycle composed of one quiet period and one refresh period is known as an LPI cycle and lasts for an lpi_qr_time equal to 96 partial PHY frame periods.

lpi_offset, lpi_quiet_time, lpi_refresh_time, and lpi_qr_time are timing parameters that are integer multiples of the partial PHY frame period. lpi_offset is a fixed value of 56 partial PHY frame period.

Each direction of the link can enter and exit LPI mode independently.

Proposed Text for the Draft – LPI Signaling

▸ Update Figure 190–11—LPI signal timing

New Figure

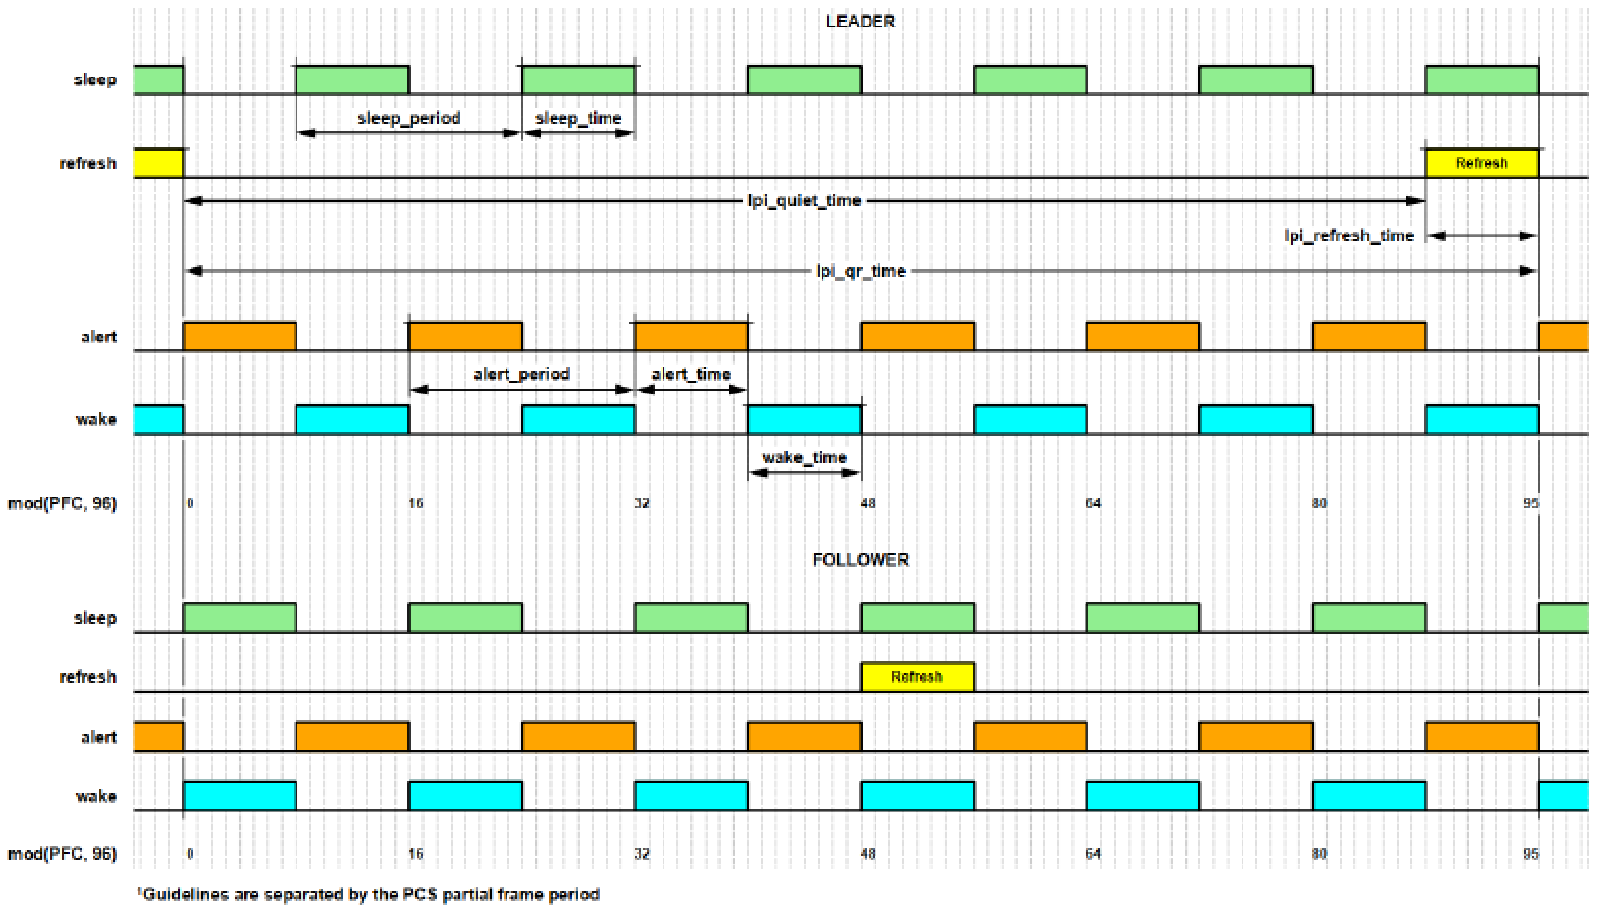

Figure 190–11—LPI signal timing

Proposed Text for the Draft – LPI Signaling

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

- Update Table 190–8—LPI timing parameters

New Table

Table 190–8—LPI timing parameters

Parameter	Number of PCS partial frame periods	Duration (µs)
sleep_period	16	38.4
sleep_time	8	19.2
lpi_offset	56	–
lpi_qr_time	96	230.4
lpi_quiet_time	88	211.2
lpi_refresh_time	8	19.2
alert_period	16	38.4
alert_time	8	19.2
wake_time	8	19.2

IEEE P802.3dg™/D1.0, 29th April 2025

Table 190–8—LPI timing parameters

Parameter	Number of partial frame periods (*Values may change)	µs
lpi_offset	56	
lpi_qr_time	96	230.4
lpi_quiet_time	88	211.2
lpi_refresh_time	8	19.2
sleep	8	19.2
wake_period	16	

Proposed Text for the Draft – LPI Synchronization

- Text for section 190.3.7.1 LPI Synchronization
  - Replace text on page 81, lines 3 to 5

IEEE P802.3dg™/D1.0, 29th April 2025

190.3.7.1 LPI Synchronization

The quiet-refresh cycle is established from the LEADER partial PHY frame Count (PFC24) during PMA Training. FOLLOWER PHY shall synchronize its PFC24 to the leader's during training. The synchronization for the alert signaling is described in Table 190–9 and Table 190–10.

New Text

190.3.7.1 LPI synchronization

An EEE-capable PHY shall synchronize refresh intervals during the LPI mode. A PHY in FOLLOWER mode is responsible for synchronizing its PFC to the PFC of the LEADER during PAM2 training. For the requirements on the FOLLOWER alignment with reference to the LEADER, see 190.3.5.2. Following the transition to PAM3, the PCS continues to increment the PFC and uses the count to generate sleep, refresh, alert, and wake signals for the transmit functions.

The sleep signal is a sequence lasting eight PCS partial frame periods (sleep_time). The refresh signal is a sequence lasting eight PCS partial frame periods (lpi_refresh_time). The alert signal is a sequence lasting eight PCS partial frame periods (alert_length). Transmission of the sleep signal may start at the beginning any multiple of 16 PCS partial frame periods (sleep_period), offset to allow transmission of the alert signal to start as soon as transmission of the sleep signal finishes. The alert signal may start at the beginning any multiple of 16 PCS partial frame periods (alert_period) starting at the beginning of the PCS partial frame that follows the refresh period. The synchronization of sleep, refresh, and alert signaling is described in Table 190–9 and Table 190–10.

Proposed Text for the Draft – LPI Synchronization

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

- Update Table 190–9—Synchronization signals derived from LEADER PFC

New Table

Table 190–9—Synchronization signals derived from LEADER PFC

LEADER variable	Condition
tx_refresh_active = TRUE	mod(PFC, lpi_qr_time) ≥ lpi_qr_time - lpi_refresh_time
tx_sleep_start_next = TRUE	mod(PFC, sleep_time) = sleep_period / 2 - 1
tx_alert_start_next = TRUE	mod(PFC, alert_period) = alert_period - 1

IEEE P802.3dg™/D1.0, 29th April 2025

Table 190–9—Synchronization logic derived from FOLLOWER signal partial PHY frame count

FOLLOWER-side Variable	u=tx_pfc
tx_refresh_active=true	lpi_offset-lpi_refresh_time ≤ mod(PFC24, lpi_qr_time) < lpi_offset
tx_wake_start = true	(PFC24, wake_period) = wake_period/2

- Update Table 190–10—Synchronization signals derived from FOLLOWER PFC

New Table

Table 190–10—Synchronization signals derived from FOLLOWER PFC

FOLLOWER variable	Condition
tx_refresh_active = TRUE	lpi_offset – lpi_refresh_time ≤ mod(PFC, lpi_qr_time) < lpi_offset
tx_sleep_start_next = TRUE	mod(PFC, sleep_time) = sleep_period - 1
tx_alert_start_next = TRUE	mod(PFC, alert_period) = alert_period / 2 - 1

IEEE P802.3dg™/D1.0, 29th April 2025

Table 190–10—Synchronization logic derived from LEADER signal partial PHY frame count

LEADER-side Variable	u=tx_pfc
tx_refresh_active=true	lpi_quiet_time ≤ mod(PFC24, lpi_qr_time)
tx_wake_start = true	mod(PFC24, wake_period) = wake_period/2

Proposed Text for the Draft – Quiet Period Signaling

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

- Text for section 190.3.7.2 Quiet period signaling
  - Replace text on page 81, lines 28 to 29

IEEE P802.3dg™/D1.0, 29th April 2025

190.3.7.2 Quiet period signaling

During quiet periods, the PCS transmitter passes zero data encoded symbols to the PMA.

New Text

190.3.7.2 Quiet period signaling

During the quiet period the transmitter shall pass zeros to the PMA via the PMA_UNITDATA.request primitive.

Proposed Text for the Draft – Refresh Period Signaling

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

- Text for section 190.3.7.3 Refresh period signaling
  - Replace text on page 81, lines 32 to 41

IEEE P802.3dg™/D1.0, 29th April 2025

190.3.7.3 Refresh period signaling

During the staggered out of phase refresh periods, the PCS transmitter operates as in normal mode, with PCS transmit data (TB_n[0:7]) set to zero.

During normal data, LPI refresh is insufficient is sent using the auxiliary bit and During LPI, LPI refresh is insufficient is sent using TB_n[0].

During wake-up, the PCS transmitter operates as in normal mode, with the PCS transmit data (tx_coded) containing (8N)B/(8N+1)B encoded normal inter-frame symbols. There is no alert signal (same as Clause 97).

New Text

190.3.7.3 Refresh period signaling

While operating in the LPI transmit mode, the PHY periodically transmits refresh signaling to allow the receiver in the link partner to maintain reliable operation. The refresh signal shall be formed by setting all of the bits of each transmit octet, Txb_n<0:7>, which is shown in Figure 190–5, to zero. Each transmit octet is then scrambled and converted into a code-group consisting of 6 PAM3 symbols as in the normal operational mode. The resulting PAM3 symbols are passed to the PMA through the PMA_UNITDATA.request primitive.

Proposed Text for the Draft – Alert Signaling

- New section 190.3.7.4 Alert signaling

New Text

190.3.7.4 Alert signaling

The PHY transmits the alert signal to indicate to the link partner that it is transitioning from the LPI transmit mode back to the normal operational mode. The alert signal shall be formed by setting all of the bits of each transmit octet, $Txb_n<0:7>$, which is shown in Figure 190–5, to one. Each transmit octet is then scrambled and converted into a code-group consisting of 6 PAM3 symbols as in the normal operational mode. The resulting PAM3 symbols are passed to the PMA through the PMA_UNITDATA.request primitive.

Proposed Text for the Draft – Functions & State Diagrams

- Text for section 190.3.8 Detailed functions and state diagrams
  - Replace text on page 81, lines 42 to 54 and page 82, lines 1 to 13 with new sections

IEEE P802.3dg™/D1.0, 29th April 2025

190.3.8 Detailed functions and state diagrams

190.3.8.1 State diagrams parameters

190.3.8.1.1 Constants

190.3.8.1.2 Variables

pcs_reset
: The pcs_reset parameter set by the PCS reset function.
: Values: TRUE or FALSE

tx_mode
: Variable set by the PHY control function and communicated through the PMA_TXMODE.indication primitive. See 190.2.2.2.

190.3.8.1.3 Timers

190.3.8.1.4 Functions

190.3.8.1.5 Counters

190.3.8.1.6 Messages

190.3.8.2 State diagrams

Proposed Text for the Draft – Functions & State Diagrams

- Text for section 190.3.8 Detailed functions and state diagrams
  - Content for section 190.3.8.1.1 Constants

190.3.8.1.1 Constants

E_MII_R<0:1><0:5>
Array containing 2 elements, each representing the MII encoding RX_DV = 1, RX_ER = 1, and RXD<3:0> = 0000.

FC_MII_R<0:1><0:5>
Array containing 2 elements, each representing the MII encoding RX_DV = 0, RX_ER = 1, and RXD<3:0> = 1110.

I_MII_R<0:1><0:5>
Array containing 2 elements, each representing the MII encoding RX_DV = 0, RX_ER = 0, and RXD<3:0> = 0000.

L_MII_R<0:1><0:5>
Array containing 2 elements, each representing the MII encoding RX_DV = 0, RX_ER = 1, and RXD<3:0> = 0101.

Editor's Note

The definition of L_MII_R assumes that Assert Local Fault is defined in clause 22.

LI_MII_R<0:1><0:5>
Array containing 2 elements, each representing the MII encoding RX_DV = 0, RX_ER = 1, and RXD<3:0> = 0001.

The following constants represent sets that are used to classify the contents of the rx_char structure. In the PCS Receive state diagram, the set operators $\in$ and $\notin$ are used to represent set membership and non-membership, respectively.

DAT_R
The set of characters that represent data.

EOP_R = {/Tp/, /Tu0/, /Tu1/, /Tu2/, /Tu3/, /Tu4/, /Tu5/, /Tu6/, /Tu7/, /Tu8/, /Tu9/, /TuA/, /TuB/, /TuC/, /TuD/, /TuE/, /TuF/}
The set of characters that represent end-of-packet.

IDL_R = {/I/, /LI/, /R/}
The set of characters that may occur between packets.

PKT_R = DAT_R $\cup$ {/E/}
The set of characters that may occur within a packet.

SOP_R = {/Sp/, /Su/}
The set of characters that represent start-of-packet.

Proposed Text for the Draft – Functions & State Diagrams

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

- Continue text for section 190.3.8 Detailed functions and state diagrams
  - Continue content for section 190.3.8.1.1 Constants / content for section 190.3.8.1.2 Variables

The constants listed below are parameterized by the value N but are fixed once N is set. Since N is known before tx_mode takes the value SEND_I, these constants are fixed whenever the PCS (8N)B/(8N + 1)B Transmit state diagram shown in Figure 190–12 is in use.

IBLOCK_T<0:8N>
Vector constructed by (8N)B/(8N + 1)B encoding of N /I/ control symbols.

IXBLOCK_T<0:8N>
Vector constructed by (8N)B/(8N + 1)B encoding of N /Ix/ control symbols.

LIBLOCK_T<0:8N>
Vector constructed by (8N)B/(8N + 1)B encoding of N /LI/ control symbols.

RBLOCK_T<0:8N>
Vector constructed by (8N)B/(8N + 1)B encoding of N /R/ control symbols.

The following constants are required when RS-FEC is enabled for the link:

RFER_CNT_LIMIT
TYPE: Integer
VALUE: 16
Number of RS-FEC frames with uncorrectable errors.

RFRX_CNT_LIMIT
TYPE: Integer
VALUE: 88
Number of RS-FEC frames in the RFER monitoring interval.

190.3.8.1.2 Variables

pcs_reset
Variable used by PCS Reset to initialize all PCS functions.

rx_char
Structure representing one of the N characters that are output by the (8N)B/(8N+1)B decoder. The structure is comprised of a Boolean value and an 8-bit numerical value. The Boolean value indicates whether the numerical value represents data or control.

rx_mii<0:1><0:5>
Array containing 2 6-bit elements representing MII transfers. Each element includes the values of RX_DV, RX_ER and RXD<3:0> for the corresponding MII transfer.

tx_coded<0:8N>
Vector containing the output from the (8N)B/(8N+1)B encoder as described in 190.3.3.4.

tx_mii<0:(2N - 1)><0:5>
Array containing 2N 6-bit elements representing MII transfers. Each element includes the values of TX_EN, TX_ER and TXD<3:0> for the corresponding MII transfer.

tx_mii_idle
Boolean variable that is set TRUE when all 2N MII transfers represented by the elements of the tx_mii array indicate either Normal inter-frame or Assert remote fault.

tx_mode
Variable set by the PHY control function and communicated through the PMA_TXMODE.indication primitive. See 199.2.2.2.

Proposed Text for the Draft – Functions & State Diagrams

- Continue text for section 190.3.8 Detailed functions and state diagrams
  - Continue content for section 190.3.8.1.2 Variables

The following variables are required when EEE is enabled for the link:

alert_detect
Boolean variable that is set TRUE when the PCS Receive state diagram of Figure 190–13B is in the RX_LPI state or the RX_ALERT state and a sequence of symbols is received from the PMA via the rx_symb parameter that is consistent with alert signaling, as specified in 190.3.7.4. It is set FALSE otherwise.

eee_low_snr
Parameter set by the PMA Receive function and communicated through the PMA_EEE_LOW_SNR.indication primitive. See 190.2.2.17.

rem_eee_low_snr
Variable set by the PMA Receive function to indicate whether to SNR of the remote PHY is too low to maintain reliable operation in LPI mode.

rx_lpi_active
Parameter set by the PMA Receive function and communicated through the PMA_PCS_RX_LPI_STATUS primitive. See 190.2.2.15. The parameter is set to its default value (FALSE) in each state of the PCS Receive state diagram of Figure 190–14 where it is not explicitly set TRUE.

rx_lpi_sleep
Boolean variable that is set TRUE when 32 consecutive rx_char values each represent /LI/. It is set FALSE otherwise.

rx_wk_idle
Boolean variable that is set TRUE when the last 8 rx_char values received each represent /I/. It is set FALSE otherwise.

tx_4x_pcs_partial_frame_done
Boolean variable that is set TRUE when the final symbol of each PCS partial frame is transmitted where the associated PFC satisfies the condition mod(PFC, 4) = 3. It is set FALSE otherwise.

tx_alert_active
Boolean variable that is set TRUE in the LPI transmit mode, when the PHY is transmitting alert signaling. It is set FALSE otherwise.

tx_alert_start_next
Boolean variable that is set TRUE on the PCS partial frame before any PCS partial frame on which the alert transmission can start. It is set FALSE otherwise. The precise conditions under which tx_alert_start_next is set TRUE are specified in 190.3.7.1.

tx_lpi_active
Boolean variable that is set TRUE when the PHY transmit function is operating in the LPI transmit mode and during transitions to and from the LPI transmit mode (i.e., at any time when the PHY is transmitting sleep, quiet-refresh, or wake signaling). It is set FALSE otherwise.

tx_lpi_enable
Boolean variable that is set TRUE by the PCS (8N)B/(8N + 1)B Transmit state diagram of Figure 190–12 to enable operation in the LPI transmit. It is set to its default value (FALSE) in each state of the PCS (8N)B/(8N + 1)B Transmit state diagram where it is not explicitly set TRUE.

tx_lpi_qr_active
Boolean variable that is set TRUE in the LPI transmit mode, when the PHY is transmitting quiet-refresh signaling. It is set FALSE otherwise.

Proposed Text for the Draft – Functions & State Diagrams

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

- Continue text for section 190.3.8 Detailed functions and state diagrams
  - Continue content for section 190.3.8.1.2 Variables / content for section 190.3.8.1.3 Timers

tx_lpi_req
Boolean variable that is set TRUE when EEE is enabled for the link, and eee_low_snr is FALSE, and rem_eee_low_snr is FALSE, and 2N + 8 consecutive MII transfers, including the 2N transfers represented by the elements of the tx_mii array, indicate Assert LPI. It is set FALSE otherwise.

tx_refresh_active
Boolean variable that determines the signaling to be used from the PCS to the PMA across the PMA_UNITDATA.request interface when the variable tx_lpi_qr_active is TRUE, as described in 190.3.3.11. The tx_refresh_active variable is set TRUE on any PCS partial frame during which refresh signaling occurs. It is set FALSE otherwise. The precise conditions under which tx_refresh_active is set TRUE are specified in 190.3.7.1.

tx_sleep_start_next
Boolean variable that is set TRUE on the PCS partial frame prior to any PCS partial frame on which the sleep transmission can start. It is set FALSE otherwise. The precise conditions under which tx_sleep_start_next is set TRUE are specified in 190.3.7.1.

The following variables are required when RS-FEC is enabled for the link:

hi_rfer
Boolean variable that is set TRUE when the RS-FEC frame error ratio is too high. It is set FALSE otherwise.

rf_valid
Boolean indication that is set TRUE if the received RS-FEC is valid. The RS-FEC frame is valid if the message as defined in 190.3.3.7 can be decoded.

190.3.8.1.3 Timers

The following timers are required when EEE is enabled for the link:

lpi_rx_wake_timer
This timer determines how long the PCS Receive function sends idle signaling to the MII after the wake signal is detected.
Values: The condition lpi_rx_wake_timer_done becomes TRUE upon timer expiration.
Duration: This timer shall have a period equal to 8 PCS partial frame periods.

lpi_tx_alert_timer
This timer defines the duration of the alert signal.
Values: The condition lpi_tx_alert_timer_done becomes TRUE upon timer expiration.
Duration: This timer shall have a period equal to 8 PCS partial frame periods (alert_time).

lpi_tx_sleep_timer
This timer defines the duration of the sleep signal that the PHY transmits to indicate to the link partner that it is transitioning to the LPI transmit mode.
Values: The condition lpi_tx_sleep_timer_done becomes TRUE upon timer expiration.
Duration: This timer shall have a period equal to 8 PCS partial frame periods (sleep_time).

lpi_tx_wake_timer
This timer defines the duration of the wake signal.
Values: The condition lpi_tx_wake_timer_done becomes TRUE upon timer expiration.
Duration: This timer shall have a period equal to 8 PCS partial frame periods (wake_time).

Proposed Text for the Draft – Functions & State Diagrams

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

- Continue text for section 190.3.8 Detailed functions and state diagrams
  - Content for section 190.3.8.1.4 Functions / content for section 190.3.8.1.5 Counters/Messages

190.3.8.1.4 Functions

DECODE_MII(rx_char)
Decodes the received character represented by the rx_char structure, returning the rx_mii<0:1><0:5> array, the elements of which represent MII transfers. The DECODE_MII function shall generate the MII transfers as specified in 190.3.4.3.

ENCODE(tx_mii<0:(2N-1)><0:5>)
Encodes the 2N MII transfers represented by the elements of the tx_mii array, returning the (8N + 1)-bit vector tx_coded. The ENCODE function shall encode the MII transfers as specified in 190.3.3.4.

190.3.8.1.5 Counters

The following counters are required when EEE is enabled for the link:

lpi_rxwk_err_cnt
An integer value that counts the number of wake errors. This counter is reset to zero when link_status is not OK and is reflected in register 3.22 (see 45.2.3.12).

rfer_cnt
Count up to a maximum of RFER_CNT_LIMIT of the number of invalid RS-FEC frames received within the current RFER monitoring interval.

rfrx_cnt
Count of the number of RS-FEC frames received within the current RFER monitoring interval.

190.3.8.1.6 Messages

The following message is required when RS-FEC is enabled for the link:

RX_FRAME
A signal sent to PCS Receive indicating that a full RS-FEC frame has been decoded and the variable rf_valid is updated. This signal is not generated when the PCS Receive state diagram of Figure 190–13B is in the RX_LPI or the RX_WAKE states.

Proposed Text for the Draft –State Diagrams

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

- Section 190.3.8.2 State diagrams
  - State Diagrams and text for the PCS transmit
  - Add intro text

New Text

190.3.8.2 State diagrams

The PCS (8N)B/(8N + 1)B Transmit state diagram shown in Figure 190–12 controls the encoding of (8N + 1)B transmitted blocks. It evaluates the exit conditions for the current state exactly once for every 2N cycles of TX_CLK.

  - Replace placeholder Figure 190–12—PCS (8N)B/(8N + 1)B Transmit state diagram

- The variable tx_lpi_req is defined differently from clause 149

tx_lpi_req
Boolean variable that is set TRUE when EEE is enabled for the link, and eee_low_snr is FALSE, and rem_eee_low_snr is FALSE, and 2N + 8 consecutive MII transfers, including the 2N transfers represented by the elements of the tx_mii array, indicate Assert LPI. It is set FALSE otherwise.

- The variable tx_lpi_enable takes the place of the variable that was called tx_lpi_req in clause 149
- The variables eee_low_snr and rem_eee_low_snr are included in the generation of tx_lpi_req

New Figure

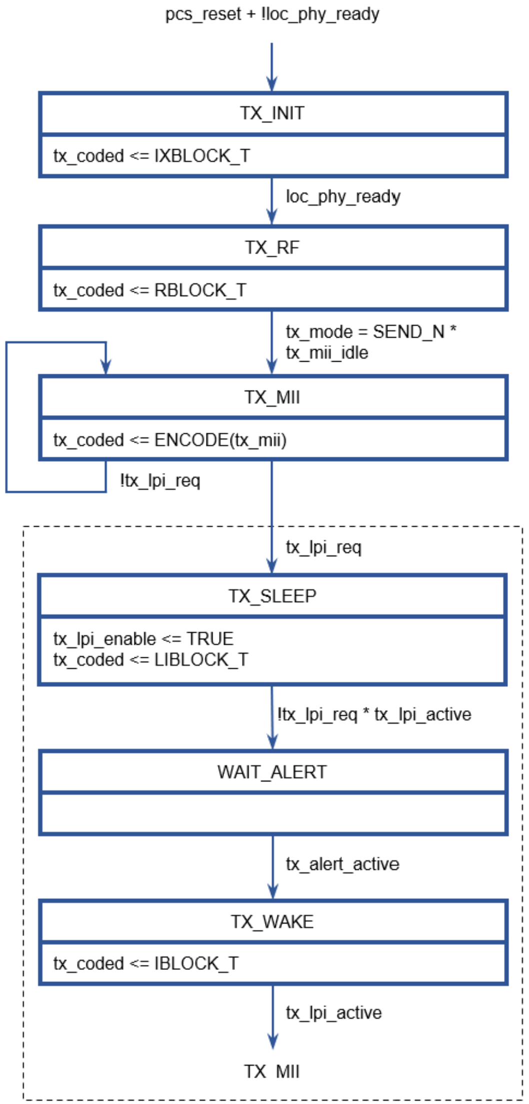

NOTE—Signals and functions shown with dashed lines are only required when EEE is enabled for the link.

Figure 190–12—PCS (8N)B/(8N + 1)B Transmit state diagram

Proposed Text for the Draft –State Diagrams

- Continue section 190.3.8.2 State diagrams
  - State Diagrams for the EEE transmit
  - New Figure 190–13A—EEE transmit state diagram

New Figure

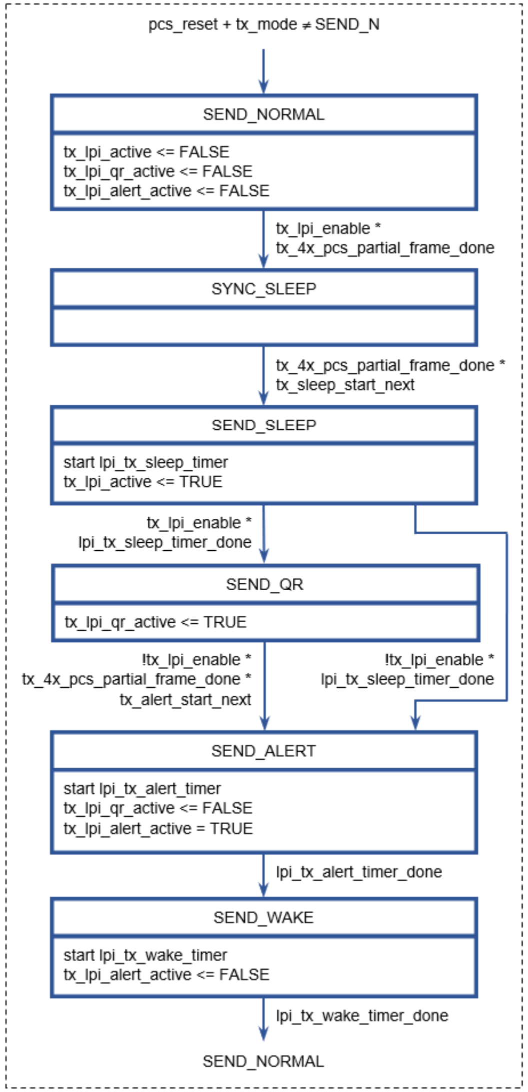

NOTE—This figure is mandatory when EEE is enabled for the link.

Figure 190–13A—EEE transmit state diagram

Proposed Text for the Draft –State Diagrams

- Continue section 190.3.8.2 State diagrams
  - Continue State Diagrams for the PCS receive
  - New Figure 190–13B—PCS Receive state diagram, part b

New Figure

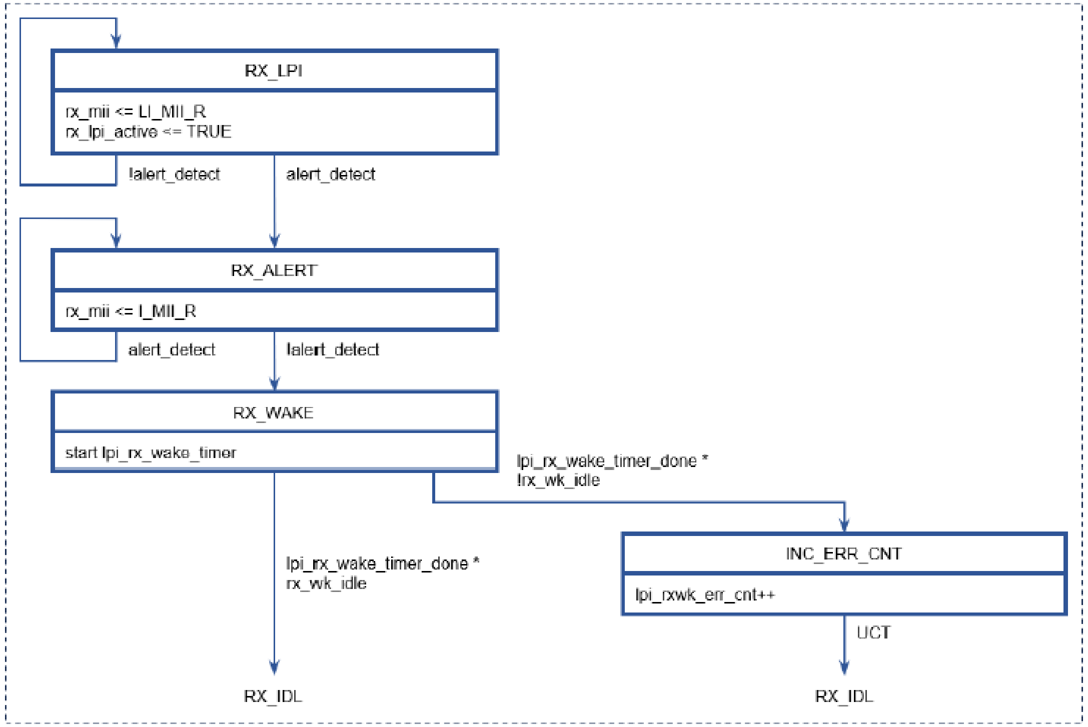

NOTE—This figure is mandatory when EEE is enabled for the link.

Figure 190–13B—PCS Receive state diagram, part b

Conclusions

- There have been a number of presentations on EEE, all of which generally follow the scheme of clauses 149 and 165
- However, most of the detail is missing with major gaps in the text
- This presentation provides new text and tables for Energy Efficient Ethernet and LPI Signalling for the draft including the necessary state diagrams

Questions ?