Annex 24A Energy Efficient Ethernet for Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer, type 100BASE-X

(normative)

24A.1 Scope

The 100BASE-X may support the capability of Energy Efficient Ethernet as described in Clause 78. When a transmitting station of a link with this capability does not need the full bandwidth, the LPI agent can put the local PHY transmitter and the link partner's receiver into low power idle mode to conserve energy. Energy is conserved by deactivating some or all functional blocks. The transmit and receive paths can enter and exit low power states independently. A 100BASE-X PHY with support for Energy Efficient Ethernet is exactly a 100BASE-X PHY as defined in Clause 24, with the additions and changes described in this annex. The only 100BASE-X PHY that supports this capability is 100BASE-TX.

24A.2 Objectives

In addition to the objectives listed in 24.1.2, the following objective is included for Energy Efficient Ethernet:

> Support Energy Efficient Ethernet with the optional function of Low Power Idle as described in Clause 78 for the embodiment of 100BASE-TX.

24A.3 Physical Coding Sublayer (PCS)

In addition to the services listed in 24.1.4.1, the following service is included for Energy Efficient Ethernet:

> Optionally, interpreting and generating MII opcodes to enable or disable the low power idle mode.

24A.4 Physical Medium Attachment (PMA) sublayer

In addition to the PMA functions listed in 24.1.4.2, the following function is included for Energy Efficient Ethernet:

> Optionally, receiving and processing low power idle state control signals from the PCS

24A.5 Functional block diagram

Figure 24A–1 provides a functional block diagram of the 100BASE-X PHY with support for Energy Efficient Ethernet. Signals or functions shown with dashed lines are optional.

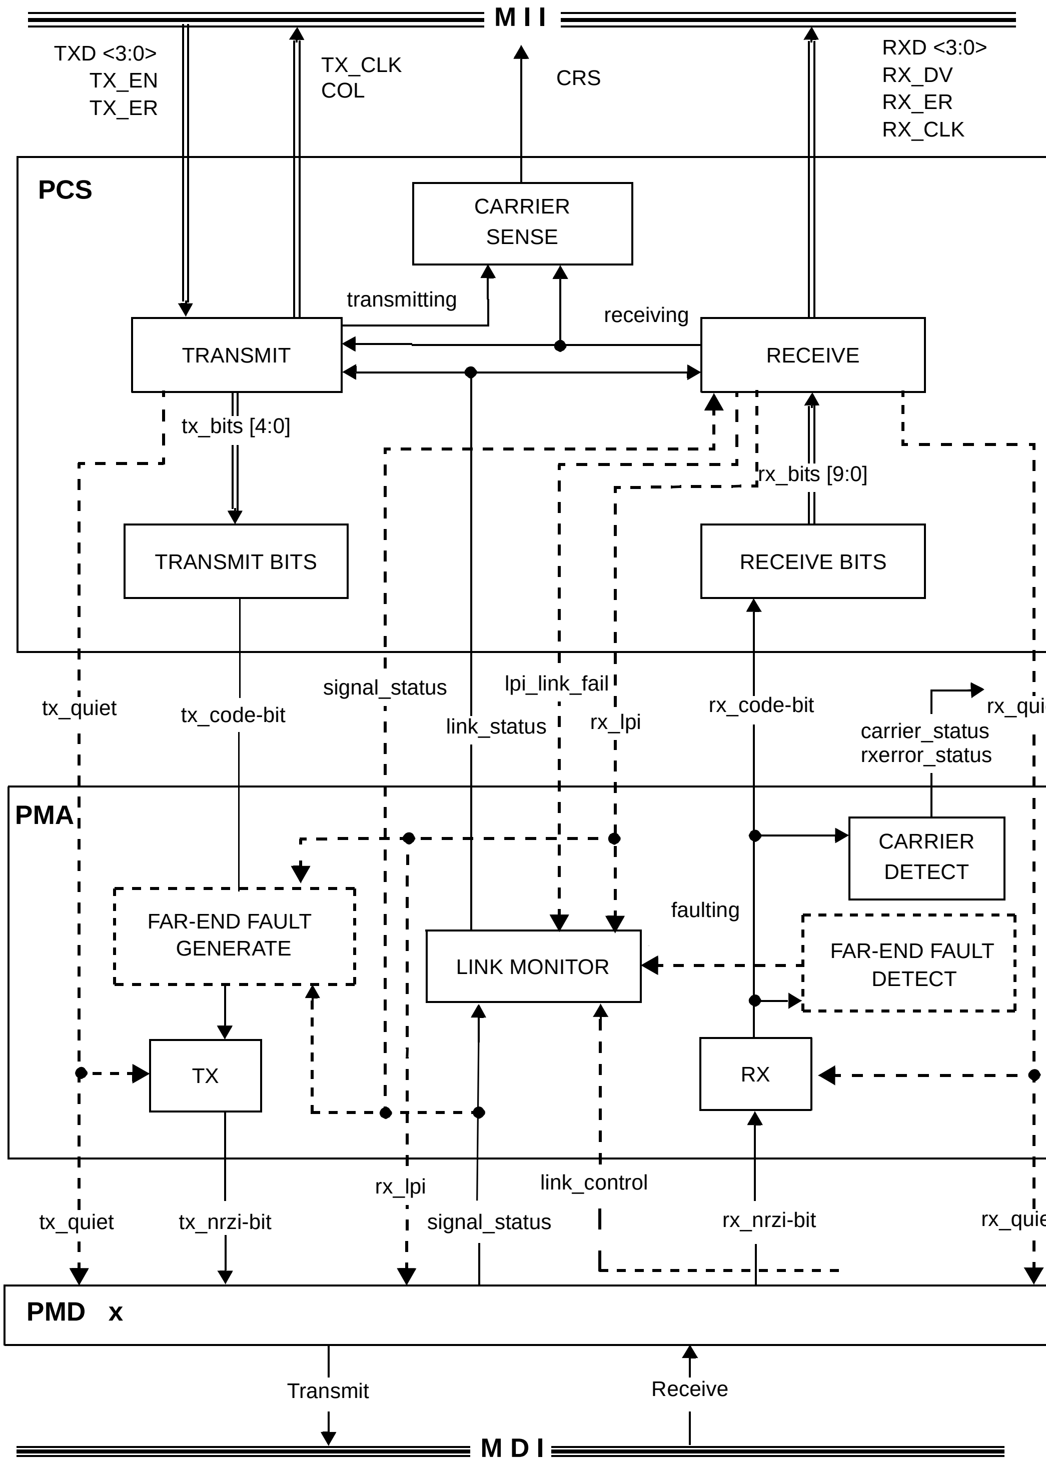

Figure 24A–1—Functional block diagram

24A.6 Physical Coding Sublayer (PCS)

24A.6.1 Functional requirements

In addition to the functional requirements stated in 24.2.2, the additional PCS functional requirements for Energy Efficient Ethernet are as follows:

The Receive process may support the low power idle by deactivating all or part of receive functional blocks of PCS, PMA, and PMD to conserve energy during low link utilization upon receiving proper code-groups via rx_code_bits from the link partner as described in 24.2.2.1.6, and generate proper commands sending through MII as described in 22.2.2.7. By interacting with Link Monitor of PMA, a link failure detection mechanism is included to differentiate two conditions of link failure due to signal off: the loss of channel signal during normal operation and the loss of refresh signal in low power idle.

The Transmit process may support the low power idle by deactivating all or part of transmit functional blocks of PCS, PMA, and PMD to conserve energy for low link utilization upon receiving the proper command from MII as described in 22.2.2.4. In this mode, the Transmit process is periodically activated to transmit refresh signal through tx_code_bits in order to allow remote receiver to keep track of the long term variation of channel characteristics and the clock drift between link partners.

24A.6.1.1 Code-groups

Code groups are defined in exactly the same way as described in 24.2.2.1, with one additional code group defined for use with Energy Efficient Ethernet:

The /P/ code-group is used to start a low power state and to refresh the link during the LPI mode.

This additional code group is shown in Table 24A–1.

The SLEEP code-group (/P/) is used to delineate the boundary of a low power idle sequence and to deliver a refresh signal to maintain clock synchronization and verify the link status. The SLEEP code-groups are emitted from, and interpreted by, the PCS.

24A.6.2 Low Power Idle

The 100BASE-X PCS accepts LPI commands from the Reconciliation Sublayer and MII (Table 22–1) to start low power transmit state. The PCS returns to the normal state when it detects the termination of the LPI command. Upon receiving LPI command, it replaces the continuous IDLE code-groups with a signal stream comprising several intermediate line states as described below and shown in Table 24A–2. The timing parameter of each line state is defined with a fixed value within a specified range.

a) *Sleep state.* The start of a LPI state is indicated by a series of SLEEP code-groups for fixed amount of time denoted by Ts as defined in Table 24A–2. Upon reception, SLEEP is interpreted by the PCS as a request to transit to low power idle mode.
b) *Quiet state.* Following SLEEP code-groups, the PCS sends a control signal to indicate the start of the Quiet state, which is consuming less power than the normal state. During the Quiet state, the PMD may cease the transmission by turning the output to a low power steady level (DC 0 volt). This state is not allowed to last longer than a fixed amount of time Tq before a Refresh or Wake state must present.
c) *Wake state.* At the end of the LPI state, the stream is terminated by a series of IDLE code-groups for default or negotiated amount of time denoted by Tw. Upon reception, IDLE is triggering the wakeup process of PMD and is interpreted by the PCS as a request to exit the low power idle mode.

Table 24A–1—4B/5B code-groups

	PCS code-group [4:0] 4 3 2 1 0	Name	MII (TXD/RXD) <3:0> 3 2 1 0	Interpretation
DATA	1 1 1 1 0	0	0 0 0 0	Data 0
	0 1 0 0 1	1	0 0 0 1	Data 1
	1 0 1 0 0	2	0 0 1 0	Data 2
	1 0 1 0 1	3	0 0 1 1	Data 3
	0 1 0 1 0	4	0 1 0 0	Data 4
	0 1 0 1 1	5	0 1 0 1	Data 5
	0 1 1 1 0	6	0 1 1 0	Data 6
	0 1 1 1 1	7	0 1 1 1	Data 7
	1 0 0 1 0	8	1 0 0 0	Data 8
	1 0 0 1 1	9	1 0 0 1	Data 9
	1 0 1 1 0	A	1 0 1 0	Data A
	1 0 1 1 1	B	1 0 1 1	Data B
	1 1 0 1 0	C	1 1 0 0	Data C
	1 1 0 1 1	D	1 1 0 1	Data D
	1 1 1 0 0	E	1 1 1 0	Data E
	1 1 1 0 1	F	1 1 1 1	Data F
	1 1 1 1 1	I	undefined	IDLE; used as inter-stream fill code
	0 0 0 0 0	P	0 0 0 1	SLEEP; Low Power Idle code; refer to Table 22–1 and Table 22–2
CONTROL	1 1 0 0 0	J	0 1 0 1	Start-of-Stream Delimiter, Part 1 of 2; always used in pairs with K
	1 0 0 0 1	K	0 1 0 1	Start-of-Stream Delimiter, Part 2 of 2; always used in pairs with J
	0 1 1 0 1	T	undefined	End-of-Stream Delimiter, Part 1 of 2; always used in pairs with R
	0 0 1 1 1	R	undefined	End-of-Stream Delimiter, Part 2 of 2; always used in pairs with T
INVALID	0 0 1 0 0	H	Undefined	Transmit Error; used to force signaling errors
	0 0 0 0 1	V	Undefined	Invalid code
	0 0 0 1 0	V	Undefined	Invalid code
	0 0 0 1 1	V	Undefined	Invalid code
	0 0 1 0 1	V	Undefined	Invalid code
	0 0 1 1 0	V	Undefined	Invalid code
	0 1 0 0 0	V	Undefined	Invalid code
	0 1 1 0 0	V	Undefined	Invalid code
	1 0 0 0 0	V	Undefined	Invalid code
	1 1 0 0 1	V	Undefined	Invalid code

Upon successfully receiving SLEEP code-groups, the 100BASE-X PCS enters low power receive state. It then sends LPI commands to the Reconciliation Sublayer and MII (Table 22–2) to notify the upper layer the change of operation mode. It returns to normal mode and ceases the transmission of LPI commands on MII

Table 24A–2—Timing Parameters and Signals of Low Power Idle line state

Line State	Symbol	Timing Parameters (TX)	Timing Parameters (RX)	Line Signal
Sleep	Ts	200 us - 220 us	240 us - 260 us	4b5b code-group /P/
Quiet	Tq	20 ms - 22 ms	24 ms - 26 ms	Differential DC zero volt
Wake	Tw	30 us - 36 us	30 us - 36 us	4b5b code-group /I/

if consecutive IDLE code-groups are received. The refresh function, which is used to maintain some internal parameters of the receiver, such as those necessary for timing recovery and signal equalization, is achieved by re-entering Sleep state from Quiet state periodically.

24A.6.3 State variables

24A.6.3.1 Constants

In addition to the constants defined in 24.2.3.1, three new constants are defined for Energy Efficient Ethernet as follows:

SLEEP

The SLEEP code-group (/P/) used for Low Power Idle state delineator, as specified in 24A.6.1.1.

TX_LP_IDLE

A value 0001 of transmit nibble-wide Data signals (TXD), together with the deassertion of TX_EN and the assertion of TX_ER on the MII, used to indicate "assert low power idle", as specified in 22.2.2.

RX_LP_IDLE

A value 0001 of receive nibble-wide Data signals (RXD), together with the deassertion of RX_DV and the assertion of RX_ER on the MII, used to indicate "receive low power idle", as specified in 22.2.2.

24A.6.3.2 Variables

In addition to the variables defined in 24.2.3.2, five new variables are defined for Energy Efficient Ethernet as follows:

lpi_link_fail

A Boolean set by the Receive process to control the transition to a Link Down state during the low power receive state. Used by the Link Monitor process of PMA as communicated through the PMA_LPILINKFAIL.request primitive.

Values: TRUE; Local receiver has detected a link failure status during low power idle state
FALSE; Local receiver is functioning normally during low power idle state

rx_lpi

A Boolean set by the Receive process to indicate the low power receive state. Used by the Link Monitor process of PMA as communicated through the PMA_RXLPI.request primitive. This parameter is used to alter the signal detection time as shown in Table 25–3. It can also be used to halt the clock RXC of MII as described in Clause 22.

Values: TRUE; Local receiver is in low power receive state
FALSE; Local receiver is in normal state

rx_quiet

A Boolean set by the Receive process to indicate the quiet line state of low power receive state as communicated through PMD_RXQUIET.request primitive. Also may be used to control the power saving function of various receiver blocks (PCS, PMA, and PMD).

Values: TRUE; The local receiver is in Quiet state
FALSE; The local receiver is not in Quiet state

tx_quiet

A Boolean set by the Transmit process to indicate the quiet line state of low power transmit state as communicated through PMD_TXQUIET.request primitive. Also may be used to control the power saving function of various transmit blocks (PCS, PMA, and PMD).

Values: TRUE; The local transmitter is in Quiet state
FALSE; The local transmitter is not in Quiet state

signal_status

The signal_status parameter as communicated by the PMD_SIGNAL.indicate primitive.

Values: ON; the quality and level of the received signal is satisfactory
OFF; the quality and level of the received signal is not satisfactory

24A.6.3.3 Timers

In addition to the timers defined in 24.2.3.4, seven new timers are defined for Energy Efficient Ethernet as follows:

lpi_link_fail_timer

In low power receive state, the receiver in Wake state is checking if valid symbols are properly received. This timer defines the maximum time allowed for PHY between entry into the Wake state and subsequent entry into the Quiet, Sleep, or Idle states before assuming a link failure. The timer shall have a period between 90 us to 110 us.

lpi_rx_ti_timer

In low power receive state, the receiver can move to Idle state when it receives consecutive IDLE symbols. In order to distinguish the intended IDLE symbols sent by link partner from ones falsely decoded during the transition from Sleep state to Quiet state before the signal status is deasserted, this receiver timer counts the minimum duration of received IDLE symbols. During this period of time, the receiver stays in an intermediate state. The timer shall have a period between 0.8 us to 0.9 us.

lpi_rx_tq_timer

In low power receive state, this receiver timer counts the maximum duration PHY stays in Quiet state before it expects a Refresh signal. If the PHY fails to receive a valid Refresh signal or Wake signal before this timer expires, the receiver shall assume a link failure when the timer has expired. The timer shall have a period between 24 ms to 26 ms.

lpi_rx_ts_timer

In low power receive state, this receiver timer counts the maximum duration PHY is allowed to stay in Sleep state before assuming a link failure. The timer shall have a period between 240 us to 260 us.

lpi_rx_tw_timer

In low power receive state, the receiver in Quiet state is woken up by the receiving signal. This receiver timer counts the expected duration for PHY to identify if valid SLEEP symbols for Refresh state or valid IDLES for Wake state have been properly received. If none of the SLEEP or IDLE symbols are received when the timer is expired, the wake error counter as defined in MDIO manageable device (MMD) register 3.22 (refer to Table 45-1) shall be incremented.. The timer shall have a period between 30 us to 36 us.

lpi_tx_tq_timer

In low power transmit state, this transmitter timer counts the duration PHY remains in Quiet state before it must wake for refresh signal. The timer shall have a period between 20 ms to 22 ms.

lpi_tx_ts_timer

In low power transmit state, this transmitter timer counts the duration PHY is sending continuous SLEEP symbols in Sleep state before going into Quiet state. The timer shall have a period between 200 us to 220 us.

24A.6.3.4 Transmit

The transmit state diagram shown in Figure 24-8 is supplanted by a new transmit state diagram for Energy Efficient Ethernet, as shown in Figure 24A–2.

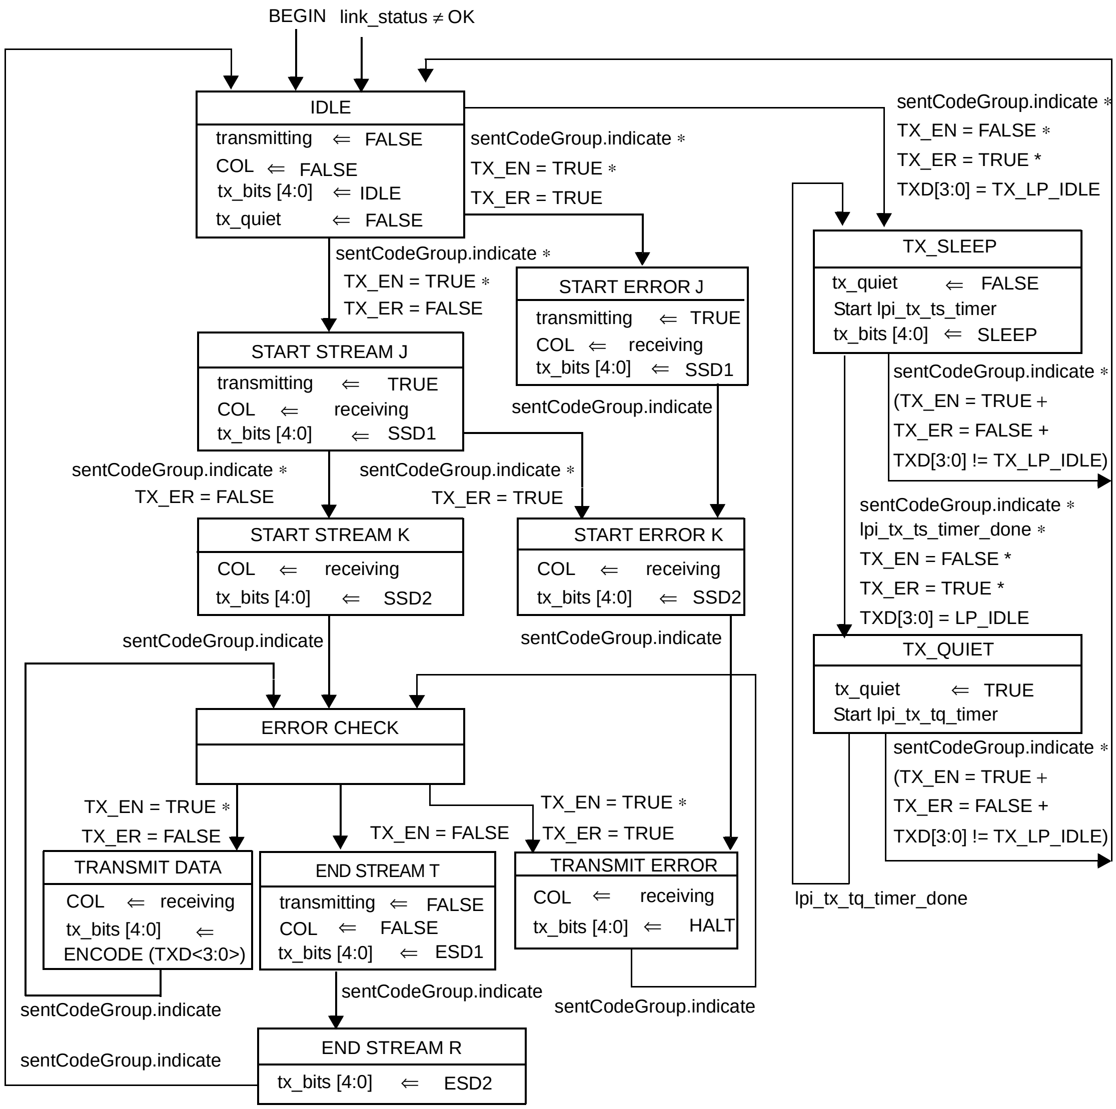

Figure 24A–2—Transmit state diagram

24A.6.3.5 Receive

The receive state diagram shown in Figure 24-11 is supplanted by the receive state diagram for Energy Efficient Ethernet as shown in Figure 24A-3.

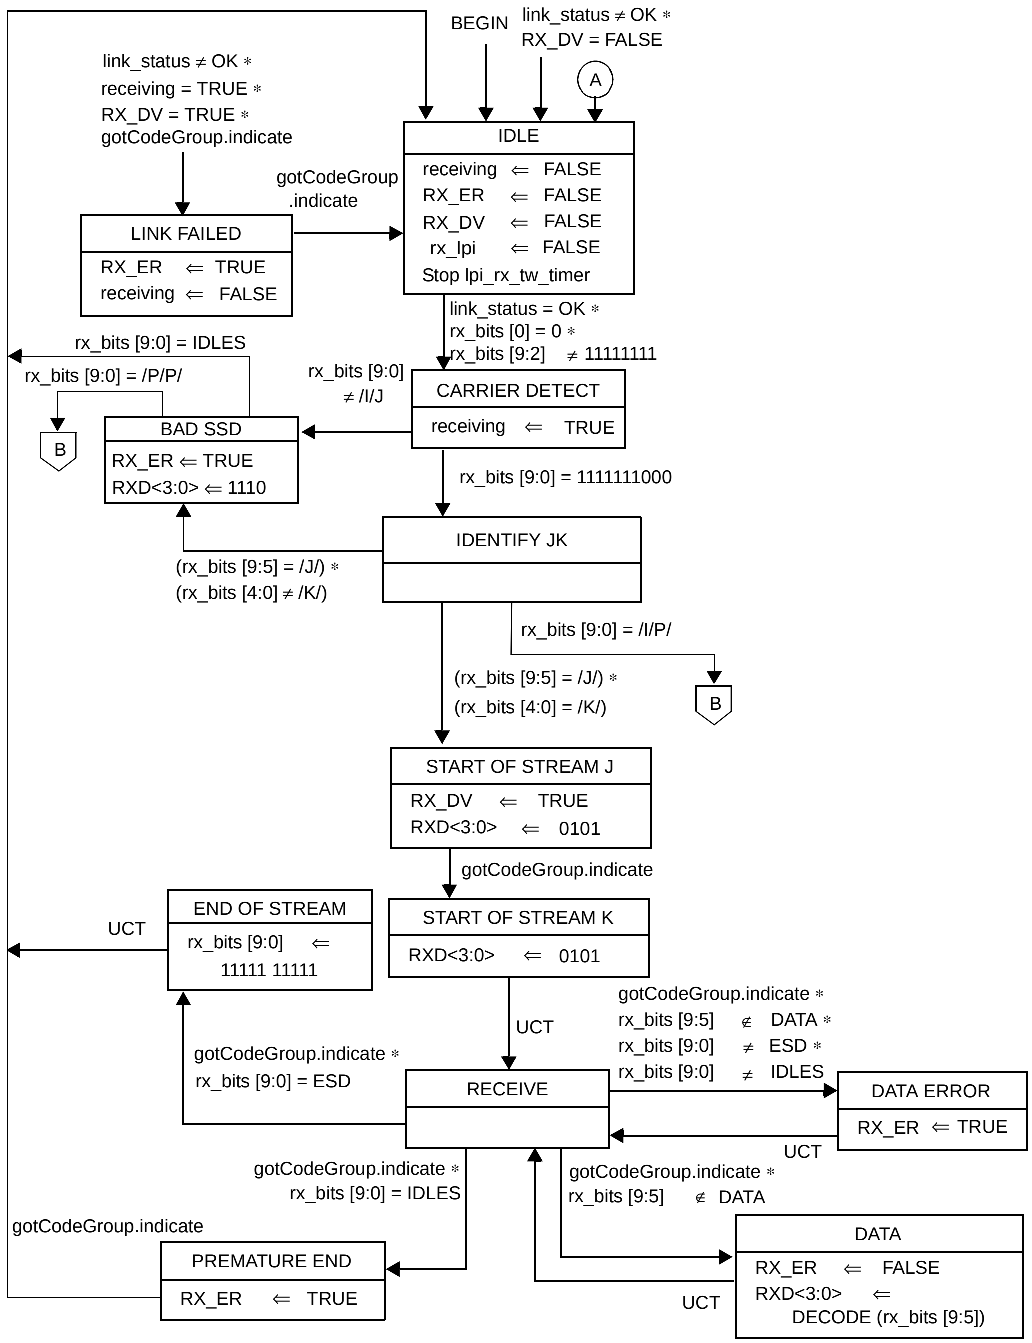

Figure 24A–3a—Receive state diagram, part a

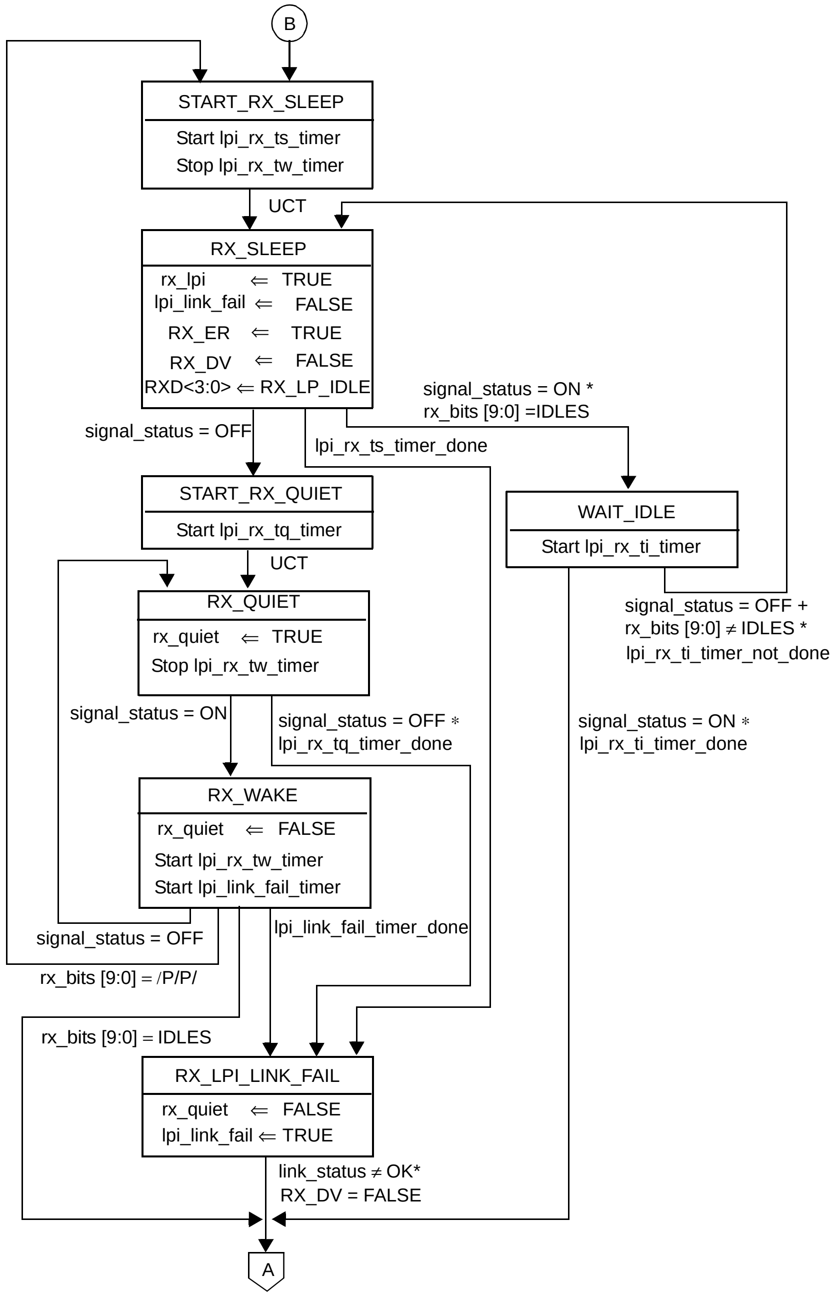

Figure 24A–3b—Receive state diagram, part b

24A.7 Physical Medium Attachment (PMA) sublayer

24A.7.1 Service Interface

In addition to the service primitives defined in 24.3.1, two new service primitives are defined for Energy Efficient Ethernet:

PMA_LPILINKFAIL.request

PMA_RXLPI.request

24A.7.1.1 PMA_LPILINKFAIL.request

This primitive is generated by the Receive Process of PCS, when Low Power Idle mode is implemented, to control one of link failure conditions of the Link Monitor of the PMA. See 24A.6.3.5 and Figure 24A–3b. When Low Power Idle mode is not implemented, the primitive is never invoked and the PMA behaves as if lpi_link_fail = FALSE.

24A.7.1.1.1 Semantics of the service primitive

PMA_LPILINKFAIL.request (lpi_link_fail)

The lpi_link_fail parameter takes on one of two values: TRUE or FALSE. The value of TRUE during low power receive state sets link_status of Link Monitor to FAIL. See 24A.7.4 and Figure 24A–4.

24A.7.1.1.2 When generated

The PCS generates this primitive to indicate a link failure condition caused by the loss of Refresh signal during low power receive state.

24A.7.1.1.3 Effect of receipt

This primitive affects operation of the PMA Link Monitor function as described in 24A.7.4.

24A.7.1.2 PMA_RXLPI.request

This primitive is generated by the Receive Process of PCS, when LPI mode is implemented, to indicate that the receiver is in low power state. See 24A.6.3.5 and Figure 24A–3b. When LPI mode is not implemented, the primitive is never invoked and the PMA behaves as if rx_lpi = FALSE.

24A.7.1.2.1 Semantics of the service primitive

PMA_RXLPI.request (rx_lpi)

The rx_lpi parameter takes on one of two values: TRUE or FALSE.

24A.7.1.2.2 When generated

The PCS generates this primitive to indicate the low power receive state.

24A.7.1.2.3 Effect of receipt

This primitive affects operation of the PMA Link Monitor function as described in 24A.7.4. Other use of receipt of this primitive by the client is unspecified by the PMA sublayer.

Change the sixth paragraph of 24.3.2.1 Far-End fault as shown below:

24A.7.2 Far-End Fault generation

Far-End fault is not generated during the low power idle mode.

24A.7.3 State Variables

24A.7.3.1 Variables

In addition to the variables defined in 24.3.3.2, two new variables are defined for Energy Efficient Ethernet as follows:

lpi_link_fail

The lpi_link_fail parameter is communicated by the PMA_LPILINKFAIL.request primitive. When low power idle mode is executed, this variable is generated by the Receive process of PCS to control the transition to a Link Down state during the low power receive state.

Values: TRUE; Local receiver has detected a link failure status during low power idle state
FALSE; Local receiver is functioning normal during low power idle state

rx_lpi

The rx_lpi parameter is communicated by the PMA_RXLPI.request primitive. This variable is from the Receive process of PCS to control the transition to indicate the low power receive state.

Values: TRUE; Local receiver is in low power receive state
FALSE; Local receiver is in normal Active state

24A.7.4 Link Monitor

The Link Monitor behavior for Energy Efficient Ethernet is the same as described in 24.3.4.4, with the following modifications:

The Link Monitor process monitors signal_status, setting link_status to FAIL whenever signal_status is OFF during normal operation or when Auto-Negotiation sets link_control to DISABLE. If the low power idle mode is implemented and the receiver is in low power state the assertion of lpi_link_fail sets the link_status to FAIL and eventually exits the low power idle mode. The link is deemed to be reliably operating when signal_status has been continuously ON for a period of time. This period is implementation dependent but not less than 330 µs or greater than 1000 µs. If so qualified, Link Monitor sets link_status to READY in order to synchronize with Auto-Negotiation, when implemented. Auto-Negotiation permits full operation by setting link_control to ENABLE. When Auto-Negotiation is not implemented, Link Monitor operates with link_control always set to ENABLE.

The Link Monitor diagram shown in Figure 24–15 is supplanted by the link monitor diagram for Energy Efficient Ethernet as shown in Figure 24A–4.

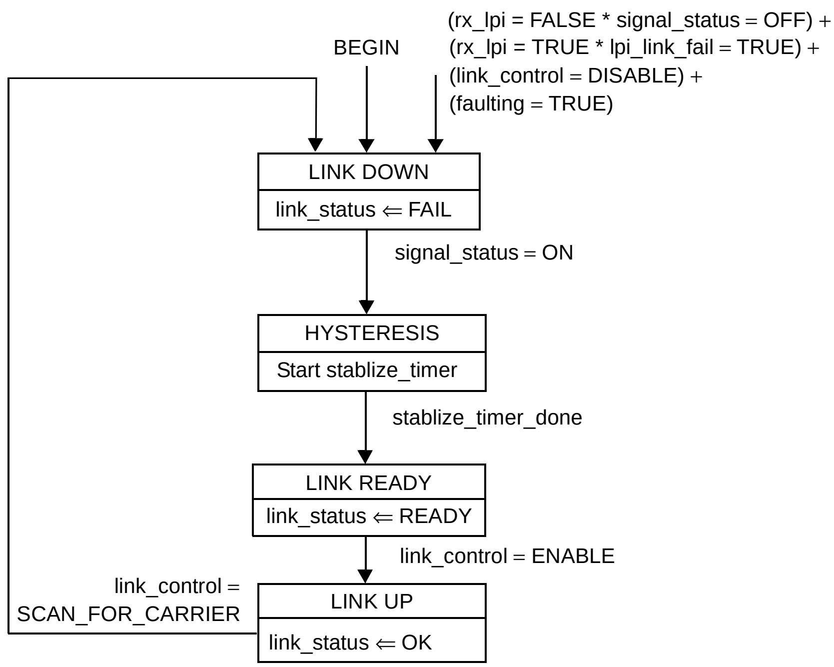

NOTE—The variables link_control and link_status are designated as link_control_[TX] and link_status_[TX], respectively, by the Auto-Negotiation Arbitration state diagram (Figure 28–18).

Figure 24A–4—Link Monitor State Diagram

24A.7.5 Far-End Fault Generation

Far-End Fault Generation for Energy Efficient Ethernet is the same as described in 24.3.4.5, with the following modifications:

Far-End Fault Generate simply passes tx_code-bits to the TX process when signal_status=ON. When signal_status=OFF and not in the low power receive state, it repetitively generates each cycle of the Far-End Fault Indication until signal_status is reasserted.

The Far-End fault diagram shown in Figure 24–16 is supplanted by the Far-End fault diagram for Energy Efficient Ethernet as shown in Figure 24A–5.

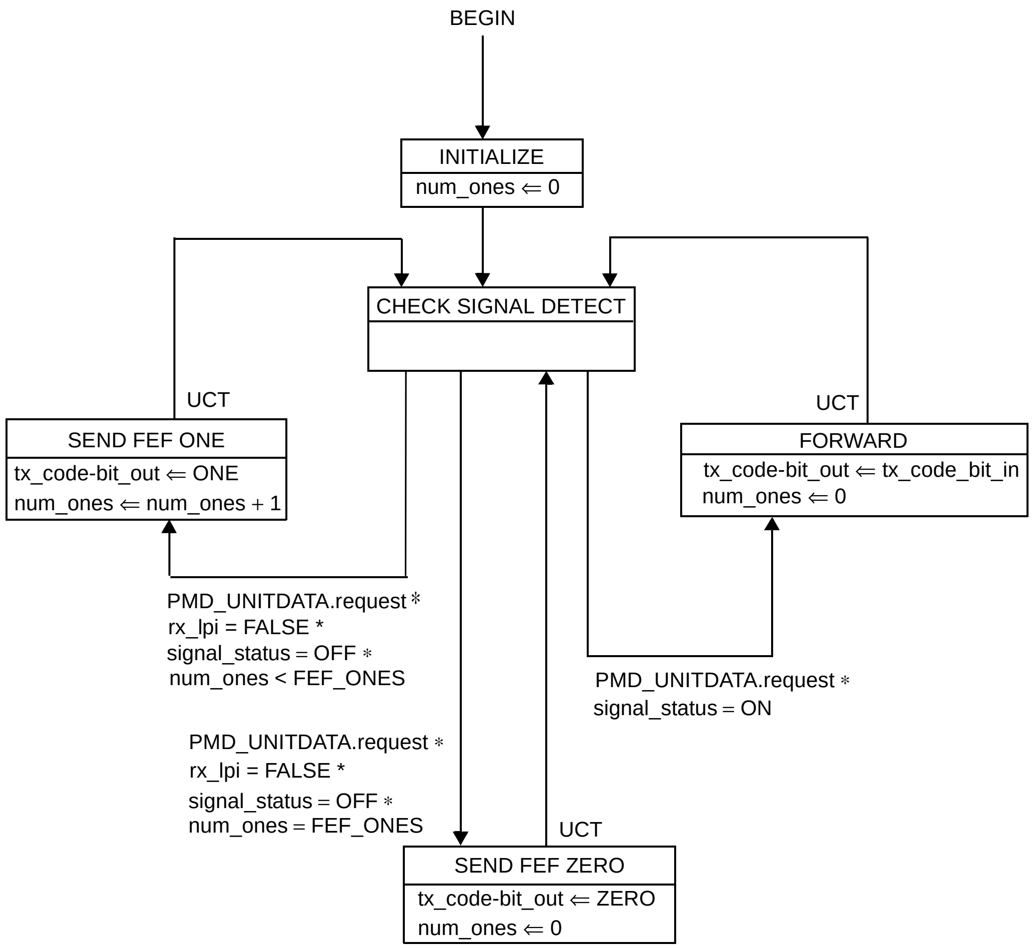

Figure 24A–5—Far-End Fault Generation state diagram

24A.8 Physical Medium Dependent (PMD) sublayer

24A.8.1 PMD Service Interface

In addition to the services listed in 24.4.1, an additional PMD service is provided for Energy Efficient Ethernet:

> 100BASE-X supports Low Power Idle mode if Energy Efficient Ethernet is implemented and low power idle mode is utilized.

In addition to the service primitives defined in 24.4.1, two additional PMD service primitives are provided for Energy Efficient Ethernet:

PMD_RXQUIET.request

PMD_TXQUIET.request

24A.8.1.1 PMD_RXQUIET.request

This primitive is generated by the Receive Process of PCS, when low power idle mode is implemented, to indicate that the receiver is in low power receive state and the line is in Quiet state. See 24A.6.3.5 and Figure 24A–3b. When low power idle mode is not implemented, the primitive is never invoked and the PMD behaves as if rx_quiet = FALSE.

24A.8.1.1.1 Semantics of the service primitive

PMD_RXQUIET.request(rx_quiet)

The rx_quiet parameter takes on one of two values: TRUE or FALSE.

24A.8.1.1.2 When generated

The PCS generates this primitive to indicate the Quiet line of low power receive state.

24A.8.1.1.3 Effect of receipt

This primitive affects operation of the PMD function of type 100BASE-TX as described in 25A.2.2.2.1.1. Other use of receipt of this primitive by the client is unspecified by the PMD sublayer.

24A.8.1.2 PMD_TXQUIET.request

This primitive is generated by the Transmit Process of PCS, when low power idle mode is implemented, to indicate that the transmitter is in low power transmit state and the line is in Quiet state. See 24A.6.3.4 and Figure 24A–2. When low power idle mode is not implemented, the primitive is never invoked and the PMD behaves as if rx_quiet = FALSE.

24A.8.1.2.1 Semantics of the service primitive

PMD_TXQUIETrequest(tx_quiet)

The tx_quiet parameter takes on one of two values: TRUE or FALSE.

24A.8.1.2.2 When generated

The PCS generates this primitive to indicate the Quiet line of low power transmit state.

24A.8.1.2.3 Effect of receipt

This primitive affects operation of the PMD function of type 100BASE-TX as described in 25A.2.2.1.1.1. Other use of receipt of this primitive by the client is unspecified by the PMD sublayer.

24A.9 Protocol implementation conformance statement (PICS) proforma for 24A, Energy Efficient Ethernet for Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer, type 100BASE-X[1]

24A.9.1 Introduction

The supplier of a protocol implementation that is claimed to conform to 24A, Energy Efficient Ethernet for Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer, type 100BASE-X, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

24A.9.2 Identification

24A.9.2.1 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2005, 24A, Energy Efficient Ethernet for Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer, type 100BASE-X
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [ ] Yes [ ] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2005.)	

Date of Statement	

[1]*Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

24A.9.3 Major capabilities/options

Item	Feature	Subclause	Status	Support	Value/Comment
*LPI	Supports LPI function	24A.6.2	O		

24A.9.4 PICS proforma tables for the Energy Efficient Ethernet for Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer, type 100BASE-X

24A.9.4.1 LPI timers

Item	Feature	Subclause	Status	Support	Value/Comment
LT1	lpi_rx_ti_timer	24A.6.3.3	LPI:M		Expired between 0.8-0.9 us after being started
LT2	lpi_rx_tq_timer	24A.6.3.3	LPI:M		Expired between 24-26 ms after being started
LT3	lpi_rx_ts_timer	24A.6.3.3	LPI:M		Expired between 240-260 us after being started
LT4	lpi_rx_tw_timer	24A.6.3.3	LPI:M		Expired between 30-36 us after being started
LT5	lpi_tx_tq_timer	24A.6.3.3	LPI:M		Expired between 20-22 ms after being started
LT6	lpi_tx_ts_timer	24A.6.3.3	LPI:M		Expired between 200-220 us after being started
LT7	lpi_link_fail_timer	24A.6.3.3	LPI:M		Expired between 90-110 us after being started