The PMA PHY Control function generates signals that control the PCS and PMA sublayer operations. PHY Control begins following the completion of Auto-Negotiation and provides the start-up functions required for successful 1000BASE-T operation. It determines whether the PHY operates in a normal state, enabling data transmission over the link segment, or whether the PHY sends special code-groups that represent the idle mode. The latter occurs when either one or both of the PHYs that share a link segment are not operating reliably.

PMA functions and state diagrams are specified in 40.4. PMA electrical specifications are given in 40.6.

40.1.4 Signaling

1000BASE-T signaling is performed by the PCS generating continuous code-group sequences that the PMA transmits over each wire pair. The signaling scheme achieves a number of objectives including

a) Forward Error Correction (FEC) coded symbol mapping for data.
b) Algorithmic mapping and inverse mapping from octet data to a quartet of quinary symbols and back.
c) Uncorrelated symbols in the transmitted symbol stream.
d) No correlation between symbol streams traveling both directions on any pair combination.
e) No correlation between symbol streams on pairs BI_DA, BI_DB, BI_DC, and BI_DD.
f) Idle mode uses a subset of code-groups in that each symbol is restricted to the set {2, 0, –2}to ease synchronization, start-up, and retraining.
g) Ability to rapidly or immediately determine if a symbol stream represents data or idle or carrier extension.
h) Robust delimiters for Start-of-Stream delimiter (SSD), End-of-Stream delimiter (ESD), and other control signals.
i) Ability to signal the status of the local receiver to the remote PHY to indicate that the local receiver is not operating reliably and requires retraining.
j) Ability to automatically detect and correct for pair swapping and unexpected crossover connections.
k) Ability to automatically detect and correct for incorrect polarity in the connections.
l) Ability to automatically correct for differential delay variations across the wire-pairs.

The PHY operates in two basic modes, normal mode or training mode. In normal mode, PCS generates code-groups that represent data, control, or idles for transmission by the PMA. In training mode, the PCS is directed to generate only idle code-groups for transmission by the PMA, which enable the receiver at the other end to train until it is ready to operate in normal mode. (See the PCS reference diagram in 40.2.)

40.1.5 Inter-sublayer interfaces

All implementations of the balanced cabling link are compatible at the MDI. Designers are free to implement circuitry within the PCS and PMA in an application-dependent manner provided that the MDI and GMII (if the GMII is implemented) specifications are met. When the PHY is incorporated within the physical bounds of a single-port device or a multiport device, implementation of the GMII is optional. System operation from the perspective of signals at the MDI and management objects are identical whether the GMII is implemented or not.

40.1.6 Conventions in this clause

The body of this clause contains state diagrams, including definitions of variables, constants, and functions. Should there be a discrepancy between a state diagram and descriptive text, the state diagram prevails.

The notation used in the state diagrams follows the conventions of 21.5.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 152 of 379

The values of all components in test circuits shall be accurate to within ±1% unless otherwise stated.

Default initializations, unless specifically specified, are left to the implementor.

40.2 1000BASE-T Service Primitives and Interfaces

1000BASE-T transfers data and control information across the following four service interfaces:

a) Gigabit Media Independent Interface (GMII)
b) PMA Service Interface
c) Medium Dependent Interface (MDI)
d) Technology-Dependent Interface

The GMII is specified in Clause 35; the Technology-Dependent Interface is specified in Clause 28. The PMA Service Interface is defined in 40.2.2 and the MDI is defined in 40.8.

40.2.1 Technology-Dependent Interface

1000BASE-T uses the following service primitives to exchange status indications and control signals across the Technology-Dependent Interface as specified in Clause 28:

PMA_LINK.request (link_control)

PMA_LINK.indicate (link_status)

40.2.1.1 PMA_LINK.request

This primitive allows the Auto-Negotiation algorithm to enable and disable operation of the PMA as specified in 28.2.6.2.

40.2.1.1.1 Semantics of the primitive

PMA_LINK.request (link_control)

The link_control parameter can take on one of three values: SCAN_FOR_CARRIER, DISABLE, or ENABLE.

SCAN_FOR_CARRIER	Used by the Auto-Negotiation algorithm prior to receiving any fast link pulses. During this mode the PMA reports link_status=FAIL.PHY processes are disabled.
DISABLE	Set by the Auto-Negotiation algorithm in the event fast link pulses are detected. PHY processes are disabled. This allows the Auto-Negotiation algorithm to determine how to configure the link.
ENABLE	Used by Auto-Negotiation to turn control over to the PHY for data processing functions.

40.2.1.1.2 When generated

Auto-Negotiation generates this primitive to indicate a change in link_control as described in Clause 28.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 153 of 379

40.2.1.1.3 Effect of receipt

This primitive affects operation of the PMA Link Monitor function as defined in 40.4.2.5.

40.2.1.2 PMA_LINK.indicate

This primitive is generated by the PMA to indicate the status of the underlying medium as specified in 28.2.6.1. This primitive informs the PCS, PMA PHY Control function, and the Auto-Negotiation algorithm about the status of the underlying link.

40.2.1.2.1 Semantics of the primitive

PMA_LINK.indicate (link_status)

The link_status parameter can take on one of three values: FAIL, READY, or OK.

FAIL	No valid link established.
READY	The Link Monitor function indicates that a 1000BASE-T link is intact and ready to be established.
OK	The Link Monitor function indicates that a valid 1000BASE-T link is established. Reliable reception of signals transmitted from the remote PHY is possible.

40.2.1.2.2 When generated

The PMA generates this primitive continuously to indicate the value of link_status in compliance with the state diagram given in Figure 40–16.

40.2.1.2.3 Effect of receipt

The effect of receipt of this primitive is specified in 40.3.3.1.

40.2.2 PMA Service Interface

1000BASE-T uses the following service primitives to exchange symbol vectors, status indications, and control signals across the service interfaces:

PMA_TXMODE.indicate (tx_mode)

PMA_CONFIG.indicate (config)

PMA_UNITDATA.request (tx_symb_vector)

PMA_UNITDATA.indicate (rx_symb_vector)

PMA_SCRSTATUS.request (scr_status)

PMA_RXSTATUS.indicate (loc_rcvr_status)

PMA_REMRXSTATUS.request (rem_rcvr_status)

The use of these primitives is illustrated in Figure 40–4.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 154 of 379

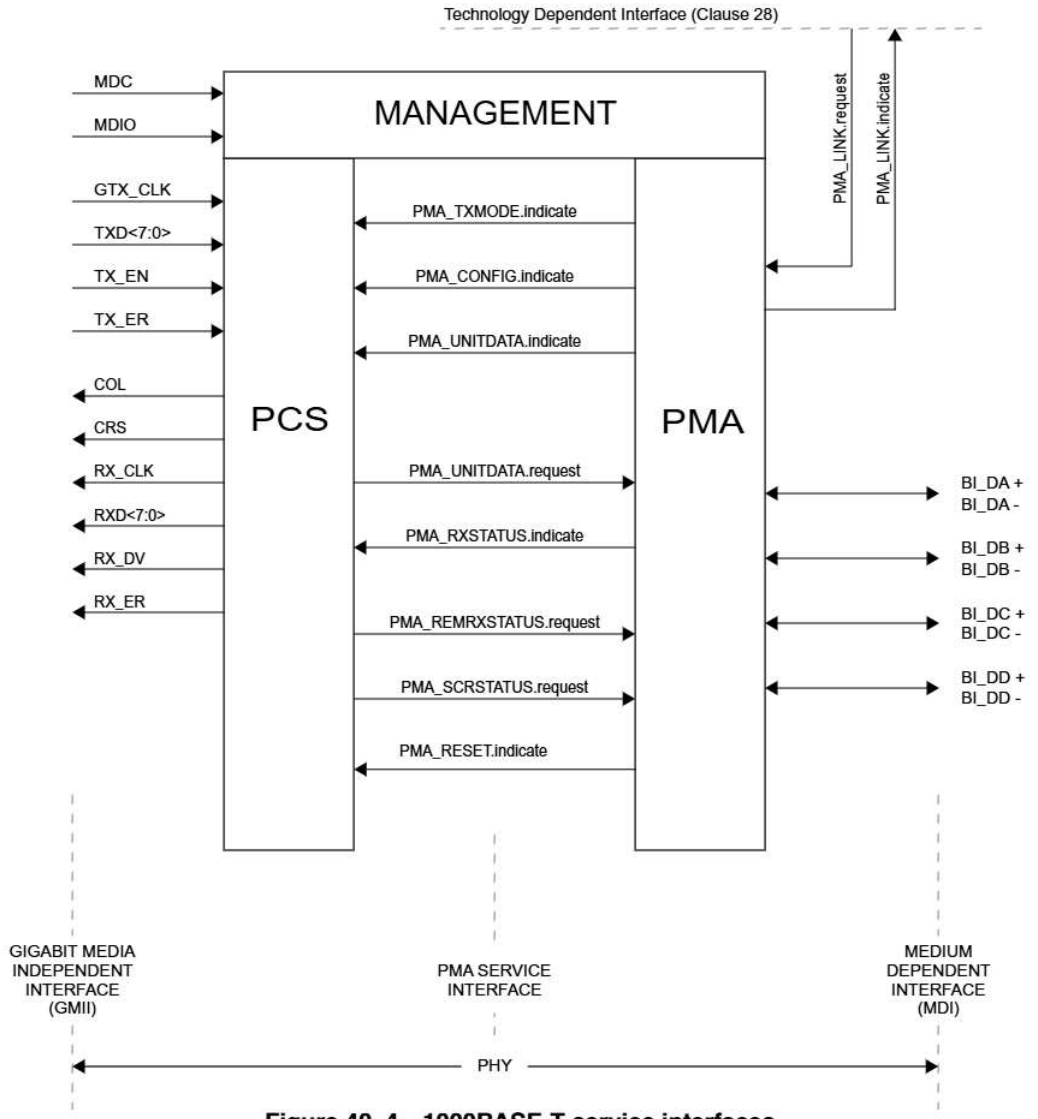


Figure 40–4—1000BASE-T service interfaces

40.2.3 PMA_TXMODE.indicate

The transmitter in a 1000BASE-T link normally sends over the four pairs, code-groups that can represent a GMII data stream, control information, or idles.

40.2.3.1 Semantics of the primitive

PMA_TXMODE.indicate (tx_mode)

PMA_TXMODE.indicate specifies to PCS Transmit via the parameter tx_mode what sequence of code-groups the PCS should be transmitting. The parameter tx_mode can take on one of the following three values of the form:

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 155 of 379

SEND_N	This value is continuously asserted when transmission of sequences of code-groups representing a GMII data stream (data mode), control mode or idle mode is to take place.
SEND_I	This value is continuously asserted in case transmission of sequences of code-groups representing the idle mode is to take place.
SEND_Z	This value is continuously asserted in case transmission of zeros is required.

40.2.3.2 When generated

The PMA PHY Control function generates PMA_TXMODE.indicate messages continuously.

40.2.3.3 Effect of receipt

Upon receipt of this primitive, the PCS performs its Transmit function as described in 40.3.1.3.

40.2.4 PMA_CONFIG.indicate

Each PHY in a 1000BASE-T link is capable of operating as a MASTER PHY and as a SLAVE PHY. MASTER-SLAVE configuration is determined during Auto-Negotiation (40.5). The result of this negotiation is provided to the PMA.

40.2.4.1 Semantics of the primitive

PMA_CONFIG.indicate (config)

PMA_CONFIG.indicate specifies to PCS and PMA Transmit via the parameter config whether the PHY must operate as a MASTER PHY or as a SLAVE PHY. The parameter config can take on one of the following two values of the form:

MASTER	This value is continuously asserted when the PHY must operate as a MASTER PHY.
SLAVE	This value is continuously asserted when the PHY must operate as a SLAVE PHY.

40.2.4.2 When generated

PMA generates PMA_CONFIG.indicate messages continuously.

40.2.4.3 Effect of receipt

PCS and PMA Clock Recovery perform their functions in MASTER or SLAVE configuration according to the value assumed by the parameter config.

40.2.5 PMA_UNITDATA.request

This primitive defines the transfer of code-groups in the form of the tx_symb_vector parameter from the PCS to the PMA. The code-groups are obtained in the PCS Transmit function using the encoding rules defined in 40.3.1.3 to represent GMII data streams, an idle mode, or other sequences.

40.2.5.1 Semantics of the primitive

PMA_UNITDATA.request (tx_symb_vector)

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 156 of 379

During transmission, the PMA_UNITDATA.request simultaneously conveys to the PMA via the parameter tx_symb_vector the value of the symbols to be sent over each of the four transmit pairs BI_DA, BI_DB, BI_DC, and BI_DD. The tx_symb_vector parameter takes on the form:

SYMB_4D	A vector of four quinary symbols, one for each of the four transmit pairs BI_DA, BI_DB, BI_DC, and BI_DD. Each quinary symbol may take on one of the values –2, –1, 0, +1, or +2.

The quinary symbols that are elements of tx_symb_vector are called, according to the pair on which each will be transmitted, tx_symb_vector[BI_DA], tx_symb_vector[BI_DB], tx_symb_vector[BI_DC], and tx_symb_vector[BI_DD].

40.2.5.2 When generated

The PCS generates PMA_UNITDATA.request (SYMB_4D) synchronously with every transmit clock cycle.

40.2.5.3 Effect of receipt

Upon receipt of this primitive the PMA transmits on the MDI the signals corresponding to the indicated quinary symbols. The parameter tx_symb_vector is also used by the PMA Receive function to process the signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DD.

40.2.6 PMA_UNITDATA.indicate

This primitive defines the transfer of code-groups in the form of the rx_symb_vector parameter from the PMA to the PCS.

40.2.6.1 Semantics of the primitive

PMA_UNITDATA.indicate (rx_symb_vector)

During reception the PMA_UNITDATA.indicate simultaneously conveys to the PCS via the parameter rx_symb_vector the values of the symbols detected on each of the four receive pairs BI_DA, BI_DB, BI_DC, and BI_DD. The rx_symbol_vector parameter takes on the form:

SYMB_4D	A vector of four quinary symbols, one for each of the four receive pairs BI_DA, BI_DB, BI_DC, and BI_DD. Each quinary symbol may take on one of the values –2, –1, 0, +1, or +2.

The quinary symbols that are elements of rx_symb_vector are called, according to the pair upon which each symbol was received, rx_symbol_vector[BI_DA], rx_symbol_vector[BI_DB], rx_symbol_vector[BI_DC], and rx_symb_vector[BI_DD].

40.2.6.2 When generated

The PMA generates PMA_UNITDATA.indicate (SYMB_4D) messages synchronously with signals received at the MDI. The nominal rate of the PMA_UNITDATA.indicate primitive is 125 MHz, as governed by the recovered clock.

40.2.6.3 Effect of receipt

The effect of receipt of this primitive is unspecified.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 157 of 379

40.2.7 PMA_SCRSTATUS.request

This primitive is generated by PCS Receive to communicate the status of the descrambler for the local PHY. The parameter scr_status conveys to the PMA Receive function the information that the descrambler has achieved synchronization.

40.2.7.1 Semantics of the primitive

PMA_SCRSTATUS.request (scr_status)

The scr_status parameter can take on one of two values of the form:

OK	The descrambler has achieved synchronization.
NOT_OK	The descrambler is not synchronized.

40.2.7.2 When generated

PCS Receive generates PMA_SCRSTATUS.request messages continuously.

40.2.7.3 Effect of receipt

The effect of receipt of this primitive is specified in 40.4.2.3, 40.4.2.4, and 40.4.6.1.

40.2.8 PMA_RXSTATUS.indicate

This primitive is generated by PMA Receive to indicate the status of the receive link at the local PHY. The parameter loc_rcvr_status conveys to the PCS Transmit, PCS Receive, PMA PHY Control function, and Link Monitor the information on whether the status of the overall receive link is satisfactory or not. Note that loc_rcvr_status is used by the PCS Receive decoding functions. The criterion for setting the parameter loc_rcvr_status is left to the implementor. It can be based, for example, on observing the mean-square error at the decision point of the receiver and detecting errors during reception of symbol streams that represent the idle mode.

40.2.8.1 Semantics of the primitive

PMA_RXSTATUS.indicate (loc_rcvr_status)

The loc_rcvr_status parameter can take on one of two values of the form:

OK	This value is asserted and remains true during reliable operation of the receive link for the local PHY.
NOT_OK	This value is asserted whenever operation of the link for the local PHY is unreliable.

40.2.8.2 When generated

PMA Receive generates PMA_RXSTATUS.indicate messages continuously on the basis of signals received at the MDI.

40.2.8.3 Effect of receipt

The effect of receipt of this primitive is specified in Figure 40–15 and in subclauses 40.2 and 40.4.6.2.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 158 of 379

40.2.9 PMA_REMRXSTATUS.request

This primitive is generated by PCS Receive to indicate the status of the receive link at the remote PHY as communicated by the remote PHY via its encoding of its loc_rcvr_status parameter. The parameter rem_rcvr_status conveys to the PMA PHY Control function the information on whether reliable operation of the remote PHY is detected or not. The criterion for setting the parameter rem_rcvr_status is left to the implementor. It can be based, for example, on asserting rem_rcvr_status is NOT_OK until loc_rcvr_status is OK and then asserting the detected value of rem_rcvr_status after proper PCS receive decoding is achieved.

40.2.9.1 Semantics of the primitive

PMA_REMRXSTATUS.request (rem_rcvr_status)

The rem_rcvr_status parameter can take on one of two values of the form:

OK	The receive link for the remote PHY is operating reliably.
NOT_OK	Reliable operation of the receive link for the remote PHY is not detected.

40.2.9.2 When generated

The PCS generates PMA_REMRXSTATUS.request messages continuously on the basis on signals received at the MDI.

40.2.9.3 Effect of receipt

The effect of receipt of this primitive is specified in Figure 40–15.

40.2.10 PMA_RESET.indicate

This primitive is used to pass the PMA Reset function to the PCS (pcs_reset=ON) when reset is enabled.

The PMA_RESET.indicate primitive can take on one of two values:

TRUE	Reset is enabled.
FALSE	Reset is not enabled.

40.2.10.1 When generated

The PMA Reset function is executed as described in 40.4.2.1.

40.2.10.2 Effect of receipt

The effect of receipt of this primitive is specified in 40.4.2.1.

40.3 Physical Coding Sublayer (PCS)

The PCS comprises one PCS Reset function and four simultaneous and asynchronous operating functions. The PCS operating functions are: PCS Transmit Enable, PCS Transmit, PCS Receive, and PCS Carrier Sense. All operating functions start immediately after the successful completion of the PCS Reset function.

The PCS reference diagram, Figure 40–5, shows how the four operating functions relate to the messages of the PCS-PMA interface. Connections from the management interface (signals MDC and MDIO) to other

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 159 of 379

layers are pervasive, and are not shown in Figure 40–5. Management is specified in Clause 30. See also Figure 40–7, which defines the structure of frames passed from PCS to PMA.

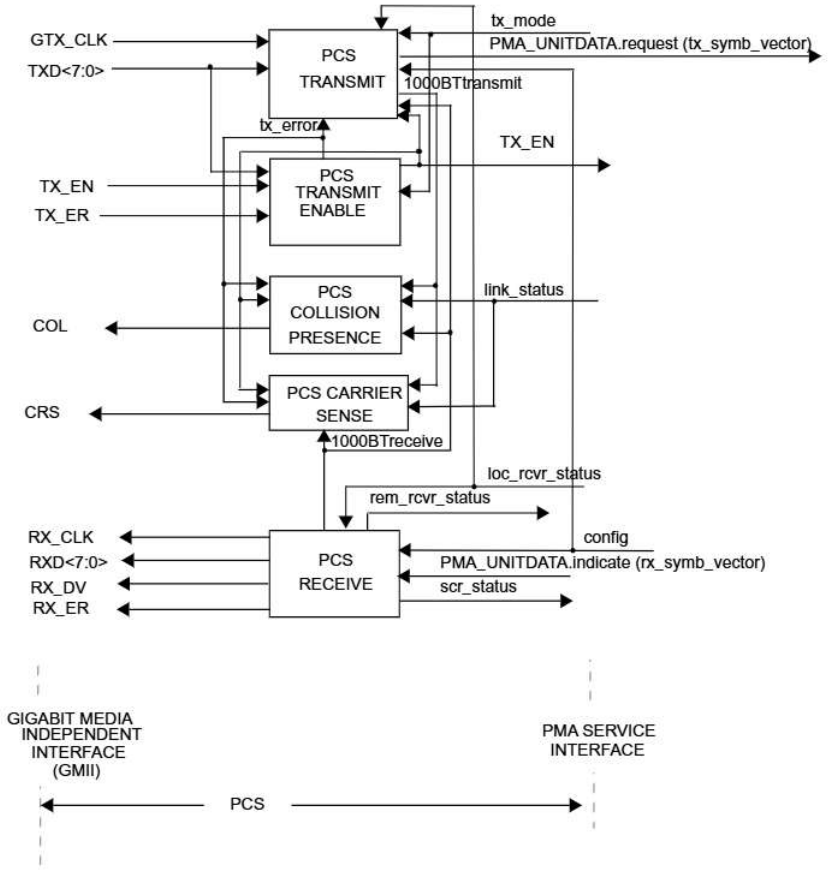


Figure 40–5—PCS reference diagram

40.3.1 PCS functions

40.3.1.1 PCS Reset function

PCS Reset initializes all PCS functions. The PCS Reset function shall be executed whenever one of the following conditions occur:

a) Power on (see 36.2.5.1.3).
b) The receipt of a request for reset from the management entity.

PCS Reset sets pcs_reset=ON while any of the above reset conditions hold true. All state diagrams take the open-ended pcs_reset branch upon execution of PCS Reset. The reference diagrams do not explicitly show the PCS Reset function.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 160 of 379

40.3.1.2 PCS Data Transmission Enable

The PCS Data Transmission Enabling process generates the signals tx_enable and tx_error, which PCS Transmit uses for data and carrier extension encoding. The process uses logical operations on tx_mode, TX_ER, TX_EN, and TXD<7:0>. The PCS shall implement the Data Transmission Enabling process as depicted in Figure 40–8 including compliance with the associated state variables as specified in 40.3.3.

40.3.1.3 PCS Transmit function

The PCS Transmit function shall conform to the PCS Transmit state diagram in Figure 40–9.

The PCS Transmit function generates the GMII signal COL based on whether a reception is occurring simultaneously with transmission. The PCS Transmit function is not required to generate the GMII signal COL in a 1000BASE-T PHY that does not support half duplex operation.

In each symbol period, PCS Transmit generates a code-group (A_n, B_n, C_n, D_n) that is transferred to the PMA via the PMA_UNITDATA.request primitive. The PMA transmits symbols A_n, B_n, C_n, D_n over wire-pairs BI_DA, BI_DB, BI_DC, and BI_DD respectively. The integer, n, is a time index that is introduced to establish a temporal relationship between different symbol periods. A symbol period, T, is nominally equal to 8 ns. In normal mode of operation, between streams of data indicated by the parameter tx_enable, PCS Transmit generates sequences of vectors using the encoding rules defined for the idle mode. Upon assertion of tx_enable, PCS Transmit passes a SSD of two consecutive vectors of four quinary symbols to the PMA, replacing the first two preamble octets. Following the SSD, each TXD<7:0> octet is encoded using an 4D-PAM5 technique into a vector of four quinary symbols until tx_enable is de-asserted. If TX_ER is asserted while tx_enable is also asserted, then PCS Transmit passes to the PMA vectors indicating a transmit error. Note that if the signal TX_ER is asserted while SSD is being sent, the transmission of the error condition is delayed until transmission of SSD has been completed. Following the de-assertion of tx_enable, a Convolutional State Reset (CSReset) of two consecutive code-groups, followed by an ESD of two consecutive code-groups, is generated, after which the transmission of idle or control mode is resumed.

If a PMA_TXMODE.indicate message has the value SEND_Z, PCS Transmit passes a vector of zeros at each symbol period to the PMA via the PMA_UNITDATA.request primitive.

If a PMA_TXMODE.indicate message has the value SEND_I, PCS Transmit generates sequences of code-groups according to the encoding rule in training mode. Special code-groups that use only the values {+2, 0, –2} are transmitted in this case. Training mode encoding also takes into account the value of the parameter loc_rcvr_status. By this mechanism, a PHY indicates the status of its own receiver to the link partner during idle transmission.

In the normal mode of operation, the PMA_TXMODE.indicate message has the value SEND_N, and the PCS Transmit function uses an 8B1Q4 coding technique to generate at each symbol period code-groups that represent data, control or idle based on the code-groups defined in Table 40–1 and Table 40–2. During transmission of data, the TXD<7:0> bits are scrambled by the PCS using a side-stream scrambler, then encoded into a code-group of quinary symbols and transferred to the PMA. During data encoding, PCS Transmit utilizes a three-state convolutional encoder.

The transition from idle or carrier extension to data is signalled by inserting a SSD, and the end of transmission of data is signalled by an ESD. Further code-groups are reserved for signaling the assertion of TX_ER within a stream of data, carrier extension, CSReset, and other control functions. During idle and carrier extension encoding, special code-groups with symbol values restricted to the set {2, 0, –2} are used. These code-groups are also generated using the transmit side-stream scrambler. However, the encoding rules for the idle, SSD, and carrier extend code-groups are different from the encoding rules for data, CSReset, CSExtend, and ESD code-groups. During idle, SSD, and carrier extension, the PCS Transmit function reverses the sign of the transmitted symbols. This allows, at the receiver, sequences of code-groups that represent data,

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 161 of 379

CSReset, CSExtend, and ESD to be easily distinguished from sequences of code-groups that represent SSD, carrier extension, and idle.

PCS encoding involves the generation of the four-bit words $Sx_n[3:0]$, $Sy_n[3:0]$, and $Sg_n[3:0]$ from which the quinary symbols (A_n, B_n, C_n, D_n) are obtained. The four-bit words $Sx_n[3:0]$, $Sy_n[3:0]$, and $Sg_n[3:0]$ are determined (as explained in 40.3.1.3.2) from sequences of pseudorandom binary symbols derived from the transmit side-stream scrambler.

40.3.1.3.1 Side-stream scrambler polynomials

The PCS Transmit function employs side-stream scrambling. If the parameter config provided to the PCS by the PMA PHY Control function via the PMA_CONFIG.indicate message assumes the value MASTER, PCS Transmit shall employ

$$g_M(x) = 1 + x^{13} + x^{33}$$

as transmitter side-stream scrambler generator polynomial. If the PMA_CONFIG.indicate message assumes the value of SLAVE, PCS Transmit shall employ

$$g_S(x) = 1 + x^{20} + x^{33}$$

as transmitter side-stream scrambler generator polynomial. An implementation of master and slave PHY side-stream scramblers by linear-feedback shift registers is shown in Figure 40–6. The bits stored in the shift register delay line at time n are denoted by $Scr_n[32:0]$. At each symbol period, the shift register is advanced by one bit, and one new bit represented by $Scr_n[0]$ is generated. The transmitter side-stream scrambler is reset upon execution of the PCS Reset function. If PCS Reset is executed, all bits of the 33-bit vector representing the side-stream scrambler state are arbitrarily set. The initialization of the scrambler state is left to the implementor. In no case shall the scrambler state be initialized to all zeros.

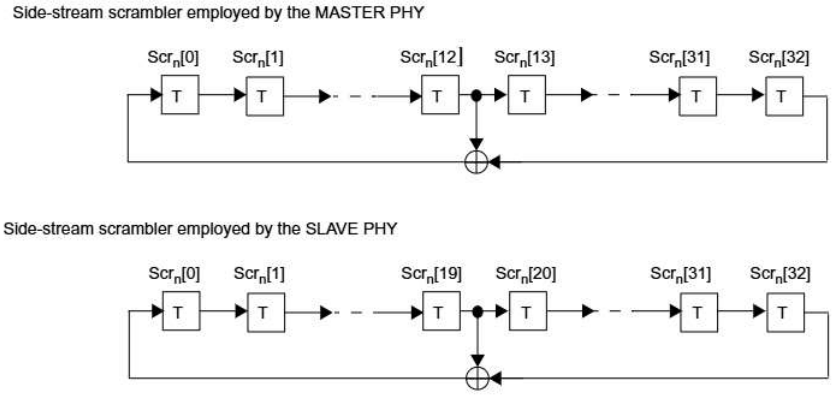


Figure 40–6—A realization of side-stream scramblers by linear feedback shift registers

40.3.1.3.2 Generation of bits $Sx_n[3:0]$, $Sy_n[3:0]$, and $Sg_n[3:0]$

PCS Transmit encoding rules are based on the generation, at time n, of the twelve bits $Sx_n[3:0]$, $Sy_n[3:0]$, and $Sg_n[3:0]$. The eight bits, $Sx_n[3:0]$ and $Sy_n[3:0]$, are used to generate the scrambler octet $Sc_n[7:0]$ for decorrelating the GMII data word TXD<7:0> during data transmission and for generating the idle and training symbols. The four bits, $Sg_n[3:0]$, are used to randomize the signs of the quinary symbols (A_n, B_n, C_n,

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 162 of 379

D_n) so that each symbol stream has no dc bias. These twelve bits are generated in a systematic fashion using three bits, X_n, Y_n, and $Scr_n[0]$, and an auxiliary generating polynomial, g(x). The two bits, X_n and Y_n, are mutually uncorrelated and also uncorrelated with the bit $Scr_n[0]$. For both master and slave PHYs, they are obtained by the same linear combinations of bits stored in the transmit scrambler shift register delay line. These two bits are derived from elements of the same maximum-length shift register sequence of length $2^{33} - 1$ as $Scr_n[0]$, but shifted in time. The associated delays are all large and different so that there is no short-term correlation among the bits $Scr_n[0]$, X_n, and Y_n. The bits X_n and Y_n are generated as follows:

$$X_n = Scr_n[4] \wedge Scr_n[6]$$

$$Y_n = Scr_n[1] \wedge Scr_n[5]$$

where ^ denotes the XOR logic operator. From the three bits X_n, Y_n, and $Scr_n[0]$, further mutually uncorrelated bit streams are obtained systematically using the generating polynomial

$$g(x) = x^3 \wedge x^8$$

The four bits $Sy_n[3:0]$ are generated using the bit $Scr_n[0]$ and g(x) as in the following equations:

$$Sy_n[0] = Scr_n[0]$$

$$Sy_n[1] = g(Scr_n[0]) = Scr_n[3] \wedge Scr_n[8]$$

$$Sy_n[2] = g^2(Scr_n[0]) = Scr_n[6] \wedge Scr_n[16]$$

$$Sy_n[3] = g^3(Scr_n[0]) = Scr_n[9] \wedge Scr_n[14] \wedge Scr_n[19] \wedge Scr_n[24]$$

The four bits $Sx_n[3:0]$ are generated using the bit X_n and g(x) as in the following equations:

$$Sx_n[0] = X_n = Scr_n[4] \wedge Scr_n[6]$$

$$Sx_n[1] = g(X_n) = Scr_n[7] \wedge Scr_n[9] \wedge Scr_n[12] \wedge Scr_n[14]$$

$$Sx_n[2] = g^2(X_n) = Scr_n[10] \wedge Scr_n[12] \wedge Scr_n[20] \wedge Scr_n[22]$$

$$Sx_n[3] = g^3(X_n) = Scr_n[13] \wedge Scr_n[15] \wedge Scr_n[18] \wedge Scr_n[20] \wedge Scr_n[23] \wedge Scr_n[25] \wedge Scr_n[28] \wedge Scr_n[30]$$

The four bits $Sg_n[3:0]$ are generated using the bit Y_n and g(x) as in the following equations:

$$Sg_n[0] = Y_n = Scr_n[1] \wedge Scr_n[5]$$

$$Sg_n[1] = g(Y_n) = Scr_n[4] \wedge Scr_n[8] \wedge Scr_n[9] \wedge Scr_n[13]$$

$$Sg_n[2] = g^2(Y_n) = Scr_n[7] \wedge Scr_n[11] \wedge Scr_n[17] \wedge Scr_n[21]$$

$$Sg_n[3] = g^3(Y_n) = Scr_n[10] \wedge Scr_n[14] \wedge Scr_n[15] \wedge Scr_n[19] \wedge Scr_n[20] \wedge Scr_n[24] \wedge Scr_n[25] \wedge Scr_n[29]$$

By construction, the twelve bits $Sx_n[3:0]$, $Sy_n[3:0]$, and $Sg_n[3:0]$ are derived from elements of the same maximum-length shift register sequence of length $2^{33}-1$ as $Scr_n[0]$, but shifted in time by varying delays. The associated delays are all large and different so that there is no apparent correlation among the bits.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 163 of 379

40.3.1.3.3 Generation of bits $Sc_n[7:0]$

The bits $Sc_n[7:0]$ are used to scramble the GMII data octet TXD[7:0] and for control, idle, and training mode quartet generation. The definition of these bits is dependent upon the bits $Sx_n[3:0]$ and $Sy_n[3:0]$ that are specified in 40.3.1.3.2, the variable tx_mode that is obtained through the PMA Service Interface, the variable tx_enable_n that is defined in Figure 40–8, and the time index n.

The four bits $Sc_n[7:4]$ are defined as

$$Sc_n[7:4] = \begin{cases} Sx_n[3:0] & \text{if } (tx\_enable_{n-2} = 1) \\ [0\ 0\ 0\ 0] & \text{else} \end{cases}$$

The bits $Sc_n[3:1]$ are defined as

$$Sc_n[3:1] = \begin{cases} [0\ 0\ 0] & \text{if } (tx\_mode = SEND\_Z) \\ Sy_n[3:1] & \text{else if } (n\text{-}n_0) = 0 \pmod 2 \\ (Sy_{n-1}[3:1] \wedge [1\ 1\ 1]) & \text{else} \end{cases}$$

where n_0 denotes the time index of the last transmitter side-stream scrambler reset.

The bit $Sc_n[0]$ is defined as

$$Sc_n[0] = \begin{cases} 0 & \text{if } (tx\_mode = SEND\_Z) \\ Sy_n[0] & \text{else} \end{cases}$$

40.3.1.3.4 Generation of bits $Sd_n[8:0]$

The PCS Transmit function generates a nine-bit word $Sd_n[8:0]$ from Sc_n that represents either a convolutionally encoded stream of data, control, or idle mode code-groups. The convolutional encoder uses a three-bit word $cs_n[2:0]$, which is defined as

$$cs_n[1] = \begin{cases} Sd_n[6] \wedge cs_{n-1}[0] & \text{if } (tx\_enable_{n-2} = 1) \\ 0 & \text{else} \end{cases}$$

$$cs_n[2] = \begin{cases} Sd_n[7] \wedge cs_{n-1}[1] & \text{if } (tx\_enable_{n-2} = 1) \\ 0 & \text{else} \end{cases}$$

$$cs_n[0] = cs_{n-1}[2]$$

from which $Sd_n[8]$ is obtained as

$$Sd_n[8] = cs_n[0]$$

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 164 of 379

The convolutional encoder bits are non-zero only during the transmission of data. Upon the completion of a data frame, the convolutional encoder bits are reset using the bit $csreset_n$. The bit $csreset_n$ is defined as

$$csreset_n = (tx\_enable_{n-2}) \text{ and } (\text{not } tx\_enable_n)$$

The bits $Sd_n[7:6]$ are derived from the bits $Sc_n[7:6]$, the GMII data bits $TXD_n[7:6]$, and from the convolutional encoder bits as

$$Sd_n[7] = \begin{cases} Sc_n[7] \wedge TXD_n[7] & \text{if } (csreset_n = 0 \text{ and } tx\_enable_{n-2} = 1) \\ cs_{n-1}[1] & \text{else if } (csreset_n = 1) \\ Sc_n[7] & \text{else} \end{cases}$$

$$Sd_n[6] = \begin{cases} Sc_n[6] \wedge TXD_n[6] & \text{if } (csreset_n = 0 \text{ and } tx\_enable_{n-2} = 1) \\ cs_{n-1}[0] & \text{else if } (csreset_n = 1) \\ Sc_n[6] & \text{else} \end{cases}$$

The bits $Sd_n[5:3]$ are derived from the bits $Sc_n[5:3]$ and the GMII data bits $TXD_n[5:3]$ as

$$Sd_n[5:3] = \begin{cases} Sc_n[5:3] \wedge TXD_n[5:3] & \text{if } (tx\_enable_{n-2} = 1) \\ Sc_n[5:3] & \text{else} \end{cases}$$

The bit $Sd_n[2]$ is used to scramble the GMII data bit $TXD_n[2]$ during data mode and to encode loc_rcvr_status otherwise. It is defined as

$$Sd_n[2] = \begin{cases} Sc_n[2] \wedge TXD_n[2] & \text{if } (tx\_enable_{n-2} = 1) \\ Sc_n[2] \wedge 1 & \text{else if } (loc\_rcvr\_status = \text{OK}) \\ Sc_n[2] & \text{else} \end{cases}$$

The bits $Sd_n[1:0]$ are used to transmit carrier extension information during tx_mode=SEND_N and are thus dependent upon the bits $cext_n$ and $cext\_err_n$. These bits are dependent on the variable tx_error_n, which is defined in Figure 40–8. These bits are defined as

$$cext_n = \begin{cases} tx\_error_n & \text{if } ((tx\_enable_n = 0) \text{ and } (TXD_n[7:0] = 0\text{x}0\text{F})) \\ 0 & \text{else} \end{cases}$$

$$cext\_err_n = \begin{cases} tx\_error_n & \text{if } ((tx\_enable_n = 0) \text{ and } (TXD_n[7:0] \neq 0\text{x}0\text{F})) \\ 0 & \text{else} \end{cases}$$

$$Sd_n[1] = \begin{cases} Sc_n[1] \wedge TXD_n[1] & \text{if } (tx\_enable_{n-2} = 1) \\ Sc_n[1] \wedge cext\_err_n & \text{else} \end{cases}$$

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 165 of 379

$$Sd_n[0] = \begin{cases} Sc_n[0] \wedge TXD_n[0] & \text{if } (tx\_enable_{n-2} = 1) \\ Sc_n[0] \wedge cext_n & \text{else} \end{cases}$$

40.3.1.3.5 Generation of quinary symbols TA_n, TB_n, TC_n, TD_n

The nine-bit word $Sd_n[8:0]$ is mapped to a quartet of quinary symbols (TA_n, TB_n, TC_n, TD_n) according to Table 40–1 and Table 40–2 shown as $Sd_n[6:8] + Sd_n[5:0]$.

Encoding of error indication:

If tx_error_n=1 when the condition (tx_enable_n * tx_enable_{n-2}) = 1, error indication is signaled by means of symbol substitution. In this condition, the values of $Sd_n[5:0]$ are ignored during mapping and the symbols corresponding to the row denoted as "xmt_err" in Table 40–1 and Table 40–2 shall be used.

Encoding of Convolutional Encoder Reset:

If tx_error_n=0 when the variable $csreset_n$ = 1, the convolutional encoder reset condition is normal. This condition is indicated by means of symbol substitution, where the values of $Sd_n[5:0]$ are ignored during mapping and the symbols corresponding to the row denoted as "CSReset" in Table 40–1 and Table 40–2 shall be used.

Encoding of Carrier Extension during Convolutional Encoder Reset:

If tx_error_n=1 when the variable $csreset_n$ = 1, the convolutional encoder reset condition indicates carrier extension. In this condition, the values of $Sd_n[5:0]$ are ignored during mapping and the symbols corresponding to the row denoted as "CSExtend" in Table 40–1 and Table 40–2 shall be used when TXD_n = 0x'0F, and the row denoted as "CSExtend_Err" in Table 40–1 and Table 40–2 shall be used when $TXD_n \neq$ 0x'0F. The latter condition denotes carrier extension with error. In case carrier extension with error is indicated during the first octet of CSReset, the error condition shall be encoded during the second octet of CSReset, and during the subsequent two octets of the End-of-Stream delimiter as well. Thus, the error condition is assumed to persist during the symbol substitutions at the End-of-Stream.

Encoding of Start-of-Stream delimiter:

The Start-of-Stream delimiter (SSD) is related to the condition SSD_n, which is defined as (tx_enable_n) * ($!tx\_enable_{n-2}$) = 1, where "*" and "!" denote the logic AND and NOT operators, respectively. For the generation of SSD, the first two octets of the preamble in a data stream are mapped to the symbols corresponding to the rows denoted as SSD1 and SSD2 respectively in Table 40–1. The symbols corresponding to the SSD1 row shall be used when the condition (tx_enable_n) * ($!tx\_enable_{n-1}$) = 1. The symbols corresponding to the SSD2 row shall be used when the condition (tx_enable_{n-1}) * ($!tx\_enable_{n-2}$) = 1.

Encoding of End-of-Stream delimiter:

The definition of an End-of-Stream delimiter (ESD) is related to the condition ESD_n, which is defined as ($!tx\_enable_{n-2}$) * (tn_enable_{n-4}) = 1. This occurs during the third and fourth symbol periods after transmission of the last octet of a data stream.

If carrier extend error is indicated during ESD, the symbols corresponding to the ESD_Ext_Err row shall be used. The two conditions upon which this may occur are

(tx_error_n) * (tx_error_{n-1}) * (tx_error_{n-2}) * ($TXD_n \neq$ 0x0F) = 1, and

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 166 of 379

$(tx\_error_n) * (tx\_error_{n-1}) * (tx\_error_{n-2}) * (tx\_error_{n-3}) * (TXD_n \neq 0x0F) = 1$.

The symbols corresponding to the ESD1 row in Table 40–1 shall be used when the condition $(!tx\_enable_{n-2}) * (tx\_enable_{n-3}) = 1$, in the absence of carrier extend error indication at time n.

The symbols corresponding to the ESD2_Ext_0 row in Table 40–1 shall be used when the condition $(!tx\_enable_{n-3}) * (tx\_enable_{n-4}) * (!tx\_error_n) * (!tx\_error_{n-1}) = 1$.

The symbols corresponding to the ESD2_Ext_1 row in Table 40–1 shall be used when the condition $(!tx\_enable_{n-3}) * (tx\_enable_{n-4}) * (!tx\_error_n) * (tx\_error_{n-1}) * (tx\_error_{n-2}) * (tx\_error_{n-3}) = 1$.

The symbols corresponding to the ESD2_Ext_2 row in Table 40–1 shall be used when the condition $(!tx\_enable_{n-3}) * (tx\_enable_{n-4}) * (tx\_error_n) * (tx\_error_{n-1}) * (tx\_error_{n-2}) * (tx\_error_{n-3}) * (TXD_n = 0x0F) = 1$, in the absence of carrier extend error indication.

NOTE—The ASCII for Table 40–1 and Table 40–2 is available from http://www.ieee802.org/3/publication/index.html.[8]

Table 40–1—Bit-to-symbol mapping (even subsets)

		$Sd_n[6:8] = [000]$	$Sd_n[6:8] = [010]$	$Sd_n[6:8] = [100]$	$Sd_n[6:8] = [110]$
Condition	**$Sd_n[5:0]$**	**TA_n,TB_n,TC_n,TD_n**	**TA_n,TB_n,TC_n,TD_n**	**TA_n,TB_n,TC_n,TD_n**	**TA_n,TB_n,TC_n,TD_n**
Normal	000000	0, 0, 0, 0	0, 0,+1,+1	0,+1,+1, 0	0,+1, 0,+1
Normal	000001	–2, 0, 0, 0	–2, 0,+1,+1	–2,+1,+1, 0	–2,+1, 0,+1
Normal	000010	0,–2, 0, 0	0,–2,+1,+1	0,–1,+1, 0	0,–1, 0,+1
Normal	000011	–2,–2, 0, 0	–2,–2,+1,+1	–2,–1,+1, 0	–2,–1, 0,+1
Normal	000100	0, 0,–2, 0	0, 0,–1,+1	0,+1,–1, 0	0,+1,–2,+1
Normal	000101	–2, 0,–2, 0	–2, 0,–1,+1	–2,+1,–1, 0	–2,+1,–2,+1
Normal	000110	0,–2,–2, 0	0,–2,–1,+1	0,–1,–1, 0	0,–1,–2,+1
Normal	000111	–2,–2,–2, 0	–2,–2,–1,+1	–2,–1,–1, 0	–2,–1,–2,+1
Normal	001000	0, 0, 0,–2	0, 0,+1,–1	0,+1,+1,–2	0,+1, 0,–1
Normal	001001	–2, 0, 0,–2	–2, 0,+1,–1	–2,+1,+1,–2	–2,+1, 0,–1
Normal	001010	0,–2, 0,–2	0,–2,+1,–1	0,–1,+1,–2	0,–1, 0,–1
Normal	001011	–2,–2, 0,–2	–2,–2,+1,–1	–2,–1,+1,–2	–2,–1, 0,–1
Normal	001100	0, 0,–2,–2	0, 0,–1,–1	0,+1,–1,–2	0,+1,–2,–1
Normal	001101	–2, 0,–2,–2	–2, 0,–1,–1	–2,+1,–1,–2	–2,+1,–2,–1
Normal	001110	0,–2,–2,–2	0,–2,–1,–1	0,–1,–1,–2	0,–1,–2,–1
Normal	001111	–2,–2,–2,–2	–2,–2,–1,–1	–2,–1,–1,–2	–2,–1,–2,–1
Normal	010000	+1,+1,+1,+1	+1,+1, 0, 0	+1, 0, 0,+1	+1, 0,+1, 0
Normal	010001	–1,+1,+1,+1	–1,+1, 0, 0	–1, 0, 0,+1	–1, 0,+1, 0
Normal	010010	+1,–1,+1,+1	+1,–1, 0, 0	+1,–2, 0,+1	+1,–2,+1, 0

[8]*Copyright release for symbol codes:* Users of this standard may freely reproduce the symbol codes in this subclause so it can be used for its intended purpose.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 167 of 379

Table 40–1—Bit-to-symbol mapping (even subsets) *(continued)*

		$Sd_n[6:8]$ = [000]	$Sd_n[6:8]$ = [010]	$Sd_n[6:8]$ = [100]	$Sd_n[6:8]$ = [110]
Condition	**$Sd_n[5:0]$**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**
Normal	010011	–1,–1,+1,+1	–1,–1,0,0	–1,–2,0,+1	–1,–2,+1,0
Normal	010100	+1,+1,–1,+1	+1,+1,–2,0	+1,0,–2,+1	+1,0,–1,0
Normal	010101	–1,+1,–1,+1	–1,+1,–2,0	–1,0,–2,+1	–1,0,–1,0
Normal	010110	+1,–1,–1,+1	+1,–1,–2,0	+1,–2,–2,+1	+1,–2,–1,0
Normal	010111	–1,–1,–1,+1	–1,–1,–2,0	–1,–2,–2,+1	–1,–2,–1,0
Normal	011000	+1,+1,+1,–1	+1,+1,0,–2	+1,0,0,–1	+1,0,+1,–2
Normal	011001	–1,+1,+1,–1	–1,+1,0,–2	–1,0,0,–1	–1,0,+1,–2
Normal	011010	+1,–1,+1,–1	+1,–1,0,–2	+1,–2,0,–1	+1,–2,+1,–2
Normal	011011	–1,–1,+1,–1	–1,–1,0,–2	–1,–2,0,–1	–1,–2,+1,–2
Normal	011100	+1,+1,–1,–1	+1,+1,–2,–2	+1,0,–2,–1	+1,0,–1,–2
Normal	011101	–1,+1,–1,–1	–1,+1,–2,–2	–1,0,–2,–1	–1,0,–1,–2
Normal	011110	+1,–1,–1,–1	+1,–1,–2,–2	+1,–2,–2,–1	+1,–2,–1,–2
Normal	011111	–1,–1,–1,–1	–1,–1,–2,–2	–1,–2,–2,–1	–1,–2,–1,–2
Normal	100000	+2,0,0,0	+2,0,+1,+1	+2,+1,+1,0	+2,+1,0,+1
Normal	100001	+2,–2,0,0	+2,–2,+1,+1	+2,–1,+1,0	+2,–1,0,+1
Normal	100010	+2,0,–2,0	+2,0,–1,+1	+2,+1,–1,0	+2,+1,–2,+1
Normal	100011	+2,–2,–2,0	+2,–2,–1,+1	+2,–1,–1,0	+2,–1,–2,+1
Normal	100100	+2,0,0,–2	+2,0,+1,–1	+2,+1,+1,–2	+2,+1,0,–1
Normal	100101	+2,–2,0,–2	+2,–2,+1,–1	+2,–1,+1,–2	+2,–1,0,–1
Normal	100110	+2,0,–2,–2	+2,0,–1,–1	+2,+1,–1,–2	+2,+1,–2,–1
Normal	100111	+2,–2,–2,–2	+2,–2,–1,–1	+2,–1,–1,–2	+2,–1,–2,–1
Normal	101000	0,0,+2,0	+1,+1,+2,0	+1,0,+2,+1	0,+1,+2,+1
Normal	101001	–2,0,+2,0	–1,+1,+2,0	–1,0,+2,+1	–2,+1,+2,+1
Normal	101010	0,–2,+2,0	+1,–1,+2,0	+1,–2,+2,+1	0,–1,+2,+1
Normal	101011	–2,–2,+2,0	–1,–1,+2,0	–1,–2,+2,+1	–2,–1,+2,+1
Normal	101100	0,0,+2,–2	+1,+1,+2,–2	+1,0,+2,–1	0,+1,+2,–1
Normal	101101	–2,0,+2,–2	–1,+1,+2,–2	–1,0,+2,–1	–2,+1,+2,–1
Normal	101110	0,–2,+2,–2	+1,–1,+2,–2	+1,–2,+2,–1	0,–1,+2,–1
Normal	101111	–2,–2,+2,–2	–1,–1,+2,–2	–1,–2,+2,–1	–2,–1,+2,–1
Normal	110000	0,+2,0,0	0,+2,+1,+1	+1,+2,0,+1	+1,+2,+1,0
Normal	110001	–2,+2,0,0	–2,+2,+1,+1	–1,+2,0,+1	–1,+2,+1,0

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 168 of 379

Table 40–1—Bit-to-symbol mapping (even subsets) *(continued)*

		$Sd_n[6:8]$ = [000]	$Sd_n[6:8]$ = [010]	$Sd_n[6:8]$ = [100]	$Sd_n[6:8]$ = [110]
Condition	**$Sd_n[5:0]$**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**
Normal	110010	0,+2,–2,0	0,+2,–1,+1	+1,+2,–2,+1	+1,+2,–1,0
Normal	110011	–2,+2,–2,0	–2,+2,–1,+1	–1,+2,–2,+1	–1,+2,–1,0
Normal	110100	0,+2,0,–2	0,+2,+1,–1	+1,+2,0,–1	+1,+2,+1,–2
Normal	110101	–2,+2,0,–2	–2,+2,+1,–1	–1,+2,0,–1	–1,+2,+1,–2
Normal	110110	0,+2,–2,–2	0,+2,–1,–1	+1,+2,–2,–1	+1,+2,–1,–2
Normal	110111	–2,+2,–2,–2	–2,+2,–1,–1	–1,+2,–2,–1	–1,+2,–1,–2
Normal	111000	0,0,0,+2	+1,+1,0,+2	0,+1,+1,+2	+1,0,+1,+2
Normal	111001	–2,0,0,+2	–1,+1,0,+2	–2,+1,+1,+2	–1,0,+1,+2
Normal	111010	0,–2,0,+2	+1,–1,0,+2	0,–1,+1,+2	+1,–2,+1,+2
Normal	111011	–2,–2,0,+2	–1,–1,0,+2	–2,–1,+1,+2	–1,–2,+1,+2
Normal	111100	0,0,–2,+2	+1,+1,–2,+2	0,+1,–1,+2	+1,0,–1,+2
Normal	111101	–2,0,–2,+2	–1,+1,–2,+2	–2,+1,–1,+2	–1,0,–1,+2
Normal	111110	0,–2,–2,+2	+1,–1,–2,+2	0,–1,–1,+2	+1,–2,–1,+2
Normal	111111	–2,–2,–2,+2	–1,–1,–2,+2	–2,–1,–1,+2	–1,–2,–1,+2
xmt_err	XXXXXX	0,+2,+2,0	+1,+1,+2,+2	+2,+1,+1,+2	+2,+1,+2,+1
CSExtend_Err	XXXXXX	–2,+2,+2,–2	–1,–1,+2,+2	+2,–1,–1,+2	+2,–1,+2,–1
CSExtend	XXXXXX	+2,0,0,+2	+2,+2,+1,+1	+1,+2,+2,+1	+1,+2,+1,+2
CSReset	XXXXXX	+2,–2,–2,+2	+2,+2,–1,–1	–1,+2,+2,–1	–1,+2,–1,+2
SSD1	XXXXXX	+2,+2,+2,+2	—	—	—
SSD2	XXXXXX	+2,+2,+2,–2	—	—	—
ESD1	XXXXXX	+2,+2,+2,+2	—	—	—
ESD2_Ext_0	XXXXXX	+2,+2,+2,–2	—	—	—
ESD2_Ext_1	XXXXXX	+2,+2,–2,+2	—	—	—
ESD2_Ext_2	XXXXXX	+2,–2,+2,+2	—	—	—
ESD_Ext_Err	XXXXXX	–2,+2,+2,+2	—	—	—
Idle/Carrier Extension	000000	0,0,0,0	—	—	—
Idle/Carrier Extension	000001	–2,0,0,0	—	—	—
Idle/Carrier Extension	000010	0,–2,0,0	—	—	—

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 169 of 379

Table 40–1—Bit-to-symbol mapping (even subsets) *(continued)*

		$Sd_n[6:8]$ = [000]	$Sd_n[6:8]$ = [010]	$Sd_n[6:8]$ = [100]	$Sd_n[6:8]$ = [110]
Condition	**$Sd_n[5:0]$**	**TA_n,TB_n,TC_n,TD_n**	**TA_n,TB_n,TC_n,TD_n**	**TA_n,TB_n,TC_n,TD_n**	**TA_n,TB_n,TC_n,TD_n**
Idle/Carrier Extension	000011	–2,–2,0,0	—	—	—
Idle/Carrier Extension	000100	0,0,–2,0	—	—	—
Idle/Carrier Extension	000101	–2,0,–2,0	—	—	—
Idle/Carrier Extension	000110	0,–2,–2,0	—	—	—
Idle/Carrier Extension	000111	–2,–2,–2,0	—	—	—
Idle/Carrier Extension	001000	0,0,0,–2	—	—	—
Idle/Carrier Extension	001001	–2,0,0,–2	—	—	—
Idle/Carrier Extension	001010	0,–2,0,–2	—	—	—
Idle/Carrier Extension	001011	–2,–2,0,–2	—	—	—
Idle/Carrier Extension	001100	0,0,–2,–2	—	—	—
Idle/Carrier Extension	001101	–2,0,–2,–2	—	—	—
Idle/Carrier Extension	001110	0,–2,–2,–2	—	—	—
Idle/Carrier Extension	001111	–2,–2,–2,–2	—	—	—

Table 40–2—Bit-to-symbol mapping (odd subsets)

		$Sd_n[6:8]$ = [001]	$Sd_n[6:8]$ = [011]	$Sd_n[6:8]$ = [101]	$Sd_n[6:8]$ = [111]
Condition	**$Sd_n[5:0]$**	**TA_n,TB_n,TC_n,TD_n**	**TA_n,TB_n,TC_n,TD_n**	**TA_n,TB_n,TC_n,TD_n**	**TA_n,TB_n,TC_n,TD_n**
Normal	000000	0,0,0,+1	0,0,+1,0	0,+1,+1,+1	0,+1,0,0
Normal	000001	–2,0,0,+1	–2,0,+1,0	–2,+1,+1,+1	–2,+1,0,0
Normal	000010	0,–2,0,+1	0,–2,+1,0	0,–1,+1,+1	0,–1,0,0
Normal	000011	–2,–2,0,+1	–2,–2,+1,0	–2,–1,+1,+1	–2,–1,0,0
Normal	000100	0,0,–2,+1	0,0,–1,0	0,+1,–1,+1	0,+1,–2,0

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 170 of 379

Table 40–2—Bit-to-symbol mapping (odd subsets) *(continued)*

		$Sd_n[6:8]$ = [001]	$Sd_n[6:8]$ = [011]	$Sd_n[6:8]$ = [101]	$Sd_n[6:8]$ = [111]
Condition	**$Sd_n[5:0]$**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**
Normal	000101	–2, 0, –2, +1	–2, 0, –1, 0	–2, +1, –1, +1	–2, +1, –2, 0
Normal	000110	0, –2, –2, +1	0, –2, –1, 0	0, –1, –1, +1	0, –1, –2, 0
Normal	000111	–2, –2, –2, +1	–2, –2, –1, 0	–2, –1, –1, +1	–2, –1, –2, 0
Normal	001000	0, 0, 0, –1	0, 0, +1, –2	0, +1, +1, –1	0, +1, 0, –2
Normal	001001	–2, 0, 0, –1	–2, 0, +1, –2	–2, +1, +1, –1	–2, +1, 0, –2
Normal	001010	0, –2, 0, –1	0, –2, +1, –2	0, –1, +1, –1	0, –1, 0, –2
Normal	001011	–2, –2, 0, –1	–2, –2, +1, –2	–2, –1, +1, –1	–2, –1, 0, –2
Normal	001100	0, 0, –2, –1	0, 0, –1, –2	0, +1, –1, –1	0, +1, –2, –2
Normal	001101	–2, 0, –2, –1	–2, 0, –1, –2	–2, +1, –1, –1	–2, +1, –2, –2
Normal	001110	0, –2, –2, –1	0, –2, –1, –2	0, –1, –1, –1	0, –1, –2, –2
Normal	001111	–2, –2, –2, –1	–2, –2, –1, –2	–2, –1, –1, –1	–2, –1, –2, –2
Normal	010000	+1, +1, +1, 0	+1, +1, 0, +1	+1, 0, 0, 0	+1, 0, +1, +1
Normal	010001	–1, +1, +1, 0	–1, +1, 0, +1	–1, 0, 0, 0	–1, 0, +1, +1
Normal	010010	+1, –1, +1, 0	+1, –1, 0, +1	+1, –2, 0, 0	+1, –2, +1, +1
Normal	010011	–1, –1, +1, 0	–1, –1, 0, +1	–1, –2, 0, 0	–1, –2, +1, +1
Normal	010100	+1, +1, –1, 0	+1, +1, –2, +1	+1, 0, –2, 0	+1, 0, –1, +1
Normal	010101	–1, +1, –1, 0	–1, +1, –2, +1	–1, 0, –2, 0	–1, 0, –1, +1
Normal	010110	+1, –1, –1, 0	+1, –1, –2, +1	+1, –2, –2, 0	+1, –2, –1, +1
Normal	010111	–1, –1, –1, 0	–1, –1, –2, +1	–1, –2, –2, 0	–1, –2, –1, +1
Normal	011000	+1, +1, +1, –2	+1, +1, 0, –1	+1, 0, 0, –2	+1, 0, +1, –1
Normal	011001	–1, +1, +1, –2	–1, +1, 0, –1	–1, 0, 0, –2	–1, 0, +1, –1
Normal	011010	+1, –1, +1, –2	+1, –1, 0, –1	+1, –2, 0, –2	+1, –2, +1, –1
Normal	011011	–1, –1, +1, –2	–1, –1, 0, –1	–1, –2, 0, –2	–1, –2, +1, –1
Normal	011100	+1, +1, –1, –2	+1, +1, –2, –1	+1, 0, –2, –2	+1, 0, –1, –1
Normal	011101	–1, +1, –1, –2	–1, +1, –2, –1	–1, 0, –2, –2	–1, 0, –1, –1
Normal	011110	+1, –1, –1, –2	+1, –1, –2, –1	+1, –2, –2, –2	+1, –2, –1, –1
Normal	011111	–1, –1, –1, –2	–1, –1, –2, –1	–1, –2, –2, –2	–1, –2, –1, –1
Normal	100000	+2, 0, 0, +1	+2, 0, +1, 0	+2, +1, +1, +1	+2, +1, 0, 0
Normal	100001	+2, –2, 0, +1	+2, –2, +1, 0	+2, –1, +1, +1	+2, –1, 0, 0
Normal	100010	+2, 0, –2, +1	+2, 0, –1, 0	+2, +1, –1, +1	+2, +1, –2, 0

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 171 of 379

Table 40–2—Bit-to-symbol mapping (odd subsets) *(continued)*

		$Sd_n[6:8]$ = [001]	$Sd_n[6:8]$ = [011]	$Sd_n[6:8]$ = [101]	$Sd_n[6:8]$ = [111]
Condition	**$Sd_n[5:0]$**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**
Normal	100011	+2,–2,–2,+1	+2,–2,–1, 0	+2,–1,–1,+1	+2,–1,–2, 0
Normal	100100	+2, 0, 0,–1	+2, 0,+1,–2	+2,+1,+1,–1	+2,+1, 0,–2
Normal	100101	+2,–2, 0,–1	+2,–2,+1,–2	+2,–1,+1,–1	+2,–1, 0,–2
Normal	100110	+2, 0,–2,–1	+2, 0,–1,–2	+2,+1,–1,–1	+2,+1,–2,–2
Normal	100111	+2,–2,–2,–1	+2,–2,–1,–2	+2,–1,–1,–1	+2,–1,–2,–2
Normal	101000	0, 0,+2,+1	+1,+1,+2,+1	+1, 0,+2, 0	0,+1,+2, 0
Normal	101001	–2, 0,+2,+1	–1,+1,+2,+1	–1, 0,+2, 0	–2,+1,+2, 0
Normal	101010	0,–2,+2,+1	+1,–1,+2,+1	+1,–2,+2, 0	0,–1,+2, 0
Normal	101011	–2,–2,+2,+1	–1,–1,+2,+1	–1,–2,+2, 0	–2,–1,+2, 0
Normal	101100	0, 0,+2,–1	+1,+1,+2,–1	+1, 0,+2,–2	0,+1,+2,–2
Normal	101101	–2, 0,+2,–1	–1,+1,+2,–1	–1, 0,+2,–2	–2,+1,+2,–2
Normal	101110	0,–2,+2,–1	+1,–1,+2,–1	+1,–2,+2,–2	0,–1,+2,–2
Normal	101111	–2,–2,+2,–1	–1,–1,+2,–1	–1,–2,+2,–2	–2,–1,+2,–2
Normal	110000	0,+2, 0,+1	0,+2,+1, 0	+1,+2, 0, 0	+1,+2,+1,+1
Normal	110001	–2,+2, 0,+1	–2,+2,+1, 0	–1,+2, 0, 0	–1,+2,+1,+1
Normal	110010	0,+2,–2,+1	0,+2,–1, 0	+1,+2,–2, 0	+1,+2,–1,+1
Normal	110011	–2,+2,–2,+1	–2,+2,–1, 0	–1,+2,–2, 0	–1,+2,–1,+1
Normal	110100	0,+2, 0,–1	0,+2,+1,–2	+1,+2, 0,–2	+1,+2,+1,–1
Normal	110101	–2,+2, 0,–1	–2,+2,+1,–2	–1,+2, 0,–2	–1,+2,+1,–1
Normal	110110	0,+2,–2,–1	0,+2,–1,–2	+1,+2,–2,–2	+1,+2,–1,–1
Normal	110111	–2,+2,–2,–1	–2,+2,–1,–2	–1,+2,–2,–2	–1,+2,–1,–1
Normal	111000	+1,+1,+1,+2	0, 0,+1,+2	+1, 0, 0,+2	0,+1, 0,+2
Normal	111001	–1,+1,+1,+2	–2, 0,+1,+2	–1, 0, 0,+2	–2,+1, 0,+2
Normal	111010	+1,–1,+1,+2	0,–2,+1,+2	+1,–2, 0,+2	0,–1, 0,+2
Normal	111011	–1,–1,+1,+2	–2,–2,+1,+2	–1,–2, 0,+2	–2,–1, 0,+2
Normal	111100	+1,+1,–1,+2	0, 0,–1,+2	+1, 0,–2,+2	0,+1,–2,+2
Normal	111101	–1,+1,–1,+2	–2, 0,–1,+2	–1, 0,–2,+2	–2,+1,–2,+2
Normal	111110	+1,–1,–1,+2	0,–2,–1,+2	+1,–2,–2,+2	0,–1,–2,+2
Normal	111111	–1,–1,–1,+2	–2,–2,–1,+2	–1,–2,–2,+2	–2,–1,–2,+2
xmt_err	XXXXXX	+2,+2, 0,+1	0,+2,+1,+2	+1,+2,+2, 0	+2,+1,+2, 0

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 172 of 379

Table 40–2—Bit-to-symbol mapping (odd subsets) *(continued)*

		$Sd_n[6:8]$ = [001]	$Sd_n[6:8]$ = [011]	$Sd_n[6:8]$ = [101]	$Sd_n[6:8]$ = [111]
Condition	**$Sd_n[5:0]$**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**	**TA_n,TB_n,TC_n, TD_n**
CSExtend_Err	XXXXXX	+2,+2, –2,–1	–2,+2,–1,+2	–1,+2,+2,–2	+2,–1,+2,–2
CSExtend	XXXXXX	+2, 0,+2,+1	+2, 0,+1,+2	+1, 0,+2,+2	+2,+1, 0,+2
CSReset	XXXXXX	+2,–2,+2,–1	+2,–2,–1,+2	–1,–2,+2,+2	+2,–1,–2,+2

40.3.1.3.6 Generation of A_n, B_n, C_n, D_n

The four bits $Sg_n[3:0]$ are used to randomize the signs of the quinary symbols (A_n, B_n, C_n, D_n) so that each symbol stream has no dc bias. The bits are used to generate binary symbols (SnA_n, SnB_n, SnC_n, SnD_n) that, when multiplied by the quinary symbols (TA_n, TB_n, TC_n, TD_n), result in (A_n, B_n, C_n, D_n).

PCS Transmit ensures a distinction between code-groups transmitted during idle mode plus SSD and those transmitted during other symbol periods. This distinction is accomplished by reversing the mapping of the sign bits when the condition (tx_enable_{n-2} + tx_enable_{n-4}) = 1. This sign reversal is controlled by the variable $Srev_n$ defined as

$$Srev_n = tx\_enable_{n-2} + tx\_enable_{n-4}$$

The binary symbols SnA_n, SnB_n, SnC_n, and SnD_n are defined using $Sg_n[3:0]$ as

$$SnA_n = \begin{cases} +1 \text{ if } [(Sg_n[0] \wedge Srev_n) = 0] \\ -1 \text{ else} \end{cases}$$

$$SnB_n = \begin{cases} +1 \text{ if } [(Sg_n[1] \wedge Srev_n) = 0] \\ -1 \text{ else} \end{cases}$$

$$SnC_n = \begin{cases} +1 \text{ if } [(Sg_n[2] \wedge Srev_n) = 0] \\ -1 \text{ else} \end{cases}$$

$$SnD_n = \begin{cases} +1 \text{ if } [(Sg_n[3] \wedge Srev_n) = 0] \\ -1 \text{ else} \end{cases}$$

The quinary symbols (A_n, B_n, C_n, D_n) are generated as the product of (SnA_n, SnB_n, SnC_n, SnD_n) and (TA_n, TB_n, TC_n, TD_n) respectively.

$$A_n = TA_n \times SnA_n$$

$$B_n = TB_n \times SnB_n$$

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 173 of 379

$$C_n = TC_n \times SnC_n$$

$$D_n = TD_n \times SnD_n$$

40.3.1.4 PCS Receive function

The PCS Receive function shall conform to the PCS Receive state diagram in Figure 40–10a including compliance with the associated state variables as specified in 40.3.3.

The PCS Receive function accepts received code-groups provided by the PMA Receive function via the parameter rx_symb_vector. To achieve correct operation, PCS Receive uses the knowledge of the encoding rules that are employed in the idle mode. PCS Receive generates the sequence of vectors of four quinary symbols (RA_n, RB_n, RC_n, RD_n) and indicates the reliable acquisition of the descrambler state by setting the parameter scr_status to OK. The sequence (RA_n, RB_n, RC_n, RD_n) is processed to generate the signals RXD<7:0>, RX_DV, and RX_ER, which are presented to the GMII. PCS Receive detects the transmission of a stream of data from the remote station and conveys this information to the PCS Carrier Sense and PCS Transmit functions via the parameter 1000BTreceive.

40.3.1.4.1 Decoding of code-groups

When the PMA indicates that correct receiver operation has been achieved by setting the loc_rcvr_status parameter to the value OK, the PCS Receive continuously checks that the received sequence satisfies the encoding rule used in idle mode. When a violation is detected, PCS Receive assigns the value TRUE to the parameter 1000BTreceive and, by examining the last two received vectors (RA_{n-1}, RB_{n-1}, RC_{n-1}, RD_{n-1}) and (RA_n, RB_n, RC_n, RD_n), determines whether the violation is due to reception of SSD or to a receiver error.

Upon detection of SSD, PCS Receive also assigns the value TRUE to the parameter 1000BTreceive that is provided to the PCS Carrier Sense and Collision Presence functions. During the two symbol periods corresponding to SSD, PCS Receive replaces SSD by preamble bits. Upon the detection of SSD, the signal RX_DV is asserted and each received vector is decoded into a data octet RXD<7:0> until ESD is detected.

Upon detection of a receiver error, the signal RX_ER is asserted and the parameter rxerror_status assumes the value ERROR. De-assertion of RX_ER and transition to the IDLE state (rxerror_status=NO_ERROR) takes place upon detection of four consecutive vectors satisfying the encoding rule used in idle mode.

During reception of a stream of data, PCS Receive checks that the symbols RA_n, RB_n, RC_n, RD_n follow the encoding rule defined in 40.3.1.3.5 for ESD whenever they assume values ± 2. PCS Receive processes two consecutive vectors at each time n to detect ESD. Upon detection of ESD, PCS Receive de-asserts the signal RX_DV on the GMII. If the last symbol period of ESD indicates that a carrier extension is present, PCS Receive will assert the RX_ER signal on the GMII. If no extension is indicated in the ESD2 quartet, PCS Receive assigns the value FALSE to the parameter receiving. If an extension is present, the transition to the IDLE state occurs after detection of a valid idle symbol period and the parameter receiving remains TRUE until check_idle is TRUE. If a violation of the encoding rules is detected, PCS Receive asserts the signal RX_ER for at least one symbol period.

A premature stream termination is caused by the detection of invalid symbols during the reception of a data stream. Then, PCS Receive waits for the reception of four consecutive vectors satisfying the encoding rule used in idle mode prior to de-asserting the error indication. Note that RX_DV remains asserted during the symbol periods corresponding to the first three idle vectors, while RX_ER=TRUE is signaled on the GMII. The signal RX_ER is also asserted in the LINK FAILED state, which ensures that RX_ER remains asserted for at least one symbol period.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 174 of 379

40.3.1.4.2 Receiver descrambler polynomials

The PHY shall descramble the data stream and return the proper sequence of code-groups to the decoding process for generation of RXD<7:0> to the GMII. For side-stream descrambling, the MASTER PHY shall employ the receiver descrambler generator polynomial $g'_M(x) = 1 + x^{20} + x^{33}$ and the SLAVE PHY shall employ the receiver descrambler generator polynomial $g'_S(x) = 1 + x^{13} + x^{33}$.

40.3.1.5 PCS Carrier Sense function

The PCS Carrier Sense function generates the GMII signal CRS, which the MAC uses for deferral in half duplex mode. The PCS shall conform to the Carrier Sense state diagram as depicted in Figure 40–11 including compliance with the associated state variables as specified in 40.3.3. The PCS Carrier Sense function is not required in a 1000BASE-T PHY that does not support half duplex operation.

40.3.2 Stream structure

The tx_symb_vector and rx-symb_vector structure is shown in Figure 40–7.



Figure 40–7—The tx_symb_vector and rx-symb_vector structure

40.3.3 State variables

40.3.3.1 Variables

CEXT

A vector of four quinary symbols corresponding to the code-group generated in idle mode to denote carrier extension, as specified in 40.3.1.3.

CEXT_Err

A vector of four quinary symbols corresponding to the code-group generated in idle mode to denote carrier extension with error indication, as specified in 40.3.1.3.

COL

The COL signal of the GMII as specified in 35.2.2.10.

config

The config parameter set by PMA and passed to the PCS via the PMA_CONFIG.indicate primitive.
Values: MASTER, SLAVE.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 175 of 379

CRS

The CRS signal of the GMII as specified in 35.2.2.9.

CSExtend

A vector of four quinary symbols corresponding to the code-group indicating convolutional encoder reset condition during carrier extension, as specified in 40.3.1.3.

CSExtend_Err

A vector of four quinary symbols corresponding to the code-group indicating convolutional encoder reset condition during carrier extension with error indication, as specified in 40.3.1.3.

CSReset

A vector of four quinary symbols corresponding to the code-group indicating convolutional encoder reset condition in the absence of carrier extension, as specified in 40.3.1.3.

DATA

A vector of four quinary symbols corresponding to the code-group indicating valid data, as specified in 40.3.1.3.

ESD1

A vector of four quinary symbols corresponding to the first code-group of End-of-Stream delimiter, as specified in 40.3.1.3.

ESD2_Ext_0

A vector of four quinary symbols corresponding to the second code-group of End-of-Stream delimiter in the absence of carrier extension over the two ESD symbol periods, as specified in 40.3.1.3.

ESD2_Ext_1

A vector of four quinary symbols corresponding to the second code-group of End-of-Stream delimiter when carrier extension is indicated during the first symbol period of the End-of-Stream delimiter, but not during the second symbol period, as specified in 40.3.1.3.

ESD2_Ext_2

A vector of four quinary symbols corresponding to the second code-group of End-of-Stream delimiter when carrier extension is indicated during the two symbol periods of the End-of-Stream delimiter, as specified in 40.3.1.3.

ESD_Ext_Err

A vector of four quinary symbols corresponding to either the first or second code-group of End-of-Stream delimiter when carrier extension with error is indicated during the End-of-Stream delimiter, as specified in 40.3.1.3.

IDLE

A sequence of vectors of four quinary symbols representing the special code-group generated in idle mode in the absence of carrier extension or carrier extension with error indication, as specified in 40.3.1.3.

link_status

The link_status parameter set by PMA Link Monitor and passed to the PCS via the PMA_LINK.indicate primitive.
Values: OK or FAIL

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 176 of 379

loc_rcvr_status
: The loc_rcvr_status parameter set by the PMA Receive function and passed to the PCS via the PMA_RXSTATUS.indicate primitive.
Values: OK or NOT_OK

pcs_reset
: The pcs_reset parameter set by the PCS Reset function.
Values: ON or OFF

(RA_n, RB_n, RC_n, RD_n)
: The vector of the four correctly aligned most recently received quinary symbols generated by PCS Receive at time n.

1000BTreceive
: The receiving parameter generated by the PCS Receive function.
Values: TRUE or FALSE

rem_rcvr_status
: The rem_rcvr_status parameter generated by PCS Receive.
Values: OK or NOT_OK

repeater_mode
: See 36.2.5.1.3

Rx_n
: Alias for rx_symb_vector (a vector RA_n, RB_n, RC_n, RD_n) at time n.

rxerror_status
: The rxerror_status parameter set by the PCS Receive function.
Values: ERROR or NO_ERROR

RX_DV
: The RX_DV signal of the GMII as specified in 35.2.2.6.

RX_ER
: The RX_ER signal of the GMII as specified in 35.2.2.8.

rx_symb_vector
: A vector of four quinary symbols received by the PMA and passed to the PCS via the PMA_UNITDATA.indicate primitive.
Value: SYMB_4D

RXD[7:0]
: The RXD<7:0> signal of the GMII as specified in 35.2.2.7.

SSD1
: A vector of four quinary symbols corresponding to the first code-group of the Start-of-Stream delimiter, as specified in 40.3.1.3.5.

SSD2
: A vector of four quinary symbols corresponding to the second code-group of the Start-of-Stream delimiter, as specified in 40.3.1.3.5.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 177 of 379

1000BTtransmit
: A boolean used by the PCS Transmit Process to indicate whether a frame transmission is in progress. Used by Carrier Sense process.
Values: TRUE: The PCS is transmitting a stream
FALSE: The PCS is not transmitting a stream

TXD[7:0]
: The TXD<7:0> signal of the GMII as specified in 35.2.2.4.

tx_enable
: The tx_enable parameter generated by PCS Transmit as specified in Figure 40–8.
Values: TRUE or FALSE

tx_error
: The tx_error parameter generated by PCS Transmit as specified in Figure 40–8.
Values: TRUE or FALSE

TX_EN
: The TX_EN signal of the GMII as specified in 35.2.2.3.

TX_ER
: The TX_ER signal of the GMII as specified in 35.2.2.5.

tx_mode
: The tx_mode parameter set by the PMA PHY Control function and passed to the PCS via the PMA_TXMODE.indicate primitive.
Values: SEND_Z, SEND_N, or SEND_I

Tx_n
: Alias for tx_symb_vector at time n.

tx_symb_vector
: A vector of four quinary symbols generated by the PCS Transmit function and passed to the PMA via the PMA_UNITDATA.request primitive.
Value: SYMB_4D

xmt_err
: A vector of four quinary symbols corresponding to a transmit error indication during normal data transmission or reception, as specified in 40.3.1.3.

40.3.3.2 Functions

check_end
: A function used by the PCS Receive process to detect the reception of valid ESD symbols. The check_end function operates on the next two rx_symb_vectors, (Rx_{n+1}) and (Rx_{n+2}), available via PMA_UNITDATA.indicate, and returns a boolean value indicating whether these two consecutive vectors contain symbols corresponding to a legal ESD encoding or not, as specified in 40.3.1.3.

check_idle
: A function used by the PCS Receive process to detect the reception of valid idle code-groups after an error condition during the process. The check_idle function operates on the current rx_symb_vector and the next three rx_symb_vectors available via PMA_UNITDATA.indicate and returns a boolean value indicating whether the four consecutive vectors contain symbols corresponding to the idle mode encoding or not, as specified in 40.3.1.3.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 178 of 379

DECODE
: In the PCS Receive process, this function takes as its argument the value of rx_symb_vector and returns the corresponding GMII RXD<7:0> octet. DECODE follows the rules outlined in 40.2.6.1.

ENCODE
: In the PCS Transmit process, this function takes as its argument GMII TXD <7:0> and returns the corresponding tx_symb_vector. ENCODE follows the rules outlined in 40.2.5.1.

40.3.3.3 Timer

symb_timer
: Continuous timer: The condition symb_timer_done becomes true upon timer expiration.

  Restart time: Immediately after expiration; timer restart resets the condition symb_timer_done.

  Duration: 8 ns nominal. (See clock tolerance in 40.6.1.2.6.)

  Symb-timer shall be generated synchronously with TX_TCLK. In the PCS Transmit state diagram, the message PMA_UNITDATA.request is issued concurrently with symb_timer_done.

40.3.3.4 Messages

PMA_UNITDATA.indicate (rx_symb_vector)
: A signal sent by PMA Receive indicating that a vector of four quinary symbols is available in rx_symb_vector. (See 40.2.6.)

PMA_UNITDATA.request (tx_symb_vector)
: A signal sent to PMA Transmit indicating that a vector of four quinary symbols is available in tx_symb_vector. (See 40.2.5.)

PUDI
: Alias for PMA_UNITDATA.indicate (rx_symb_vector).

PUDR
: Alias for PMA_UNITDATA.request (tx_symb_vector).

STD
: Alias for symb_timer_done.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 179 of 379

40.3.4 State diagrams

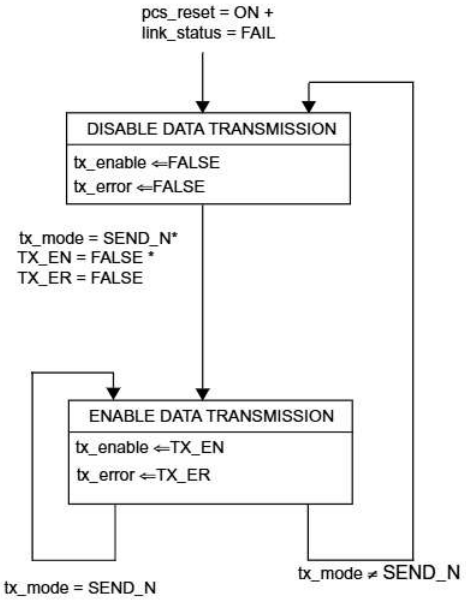


Figure 40–8—PCS Data Transmission Enabling state diagram

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 180 of 379

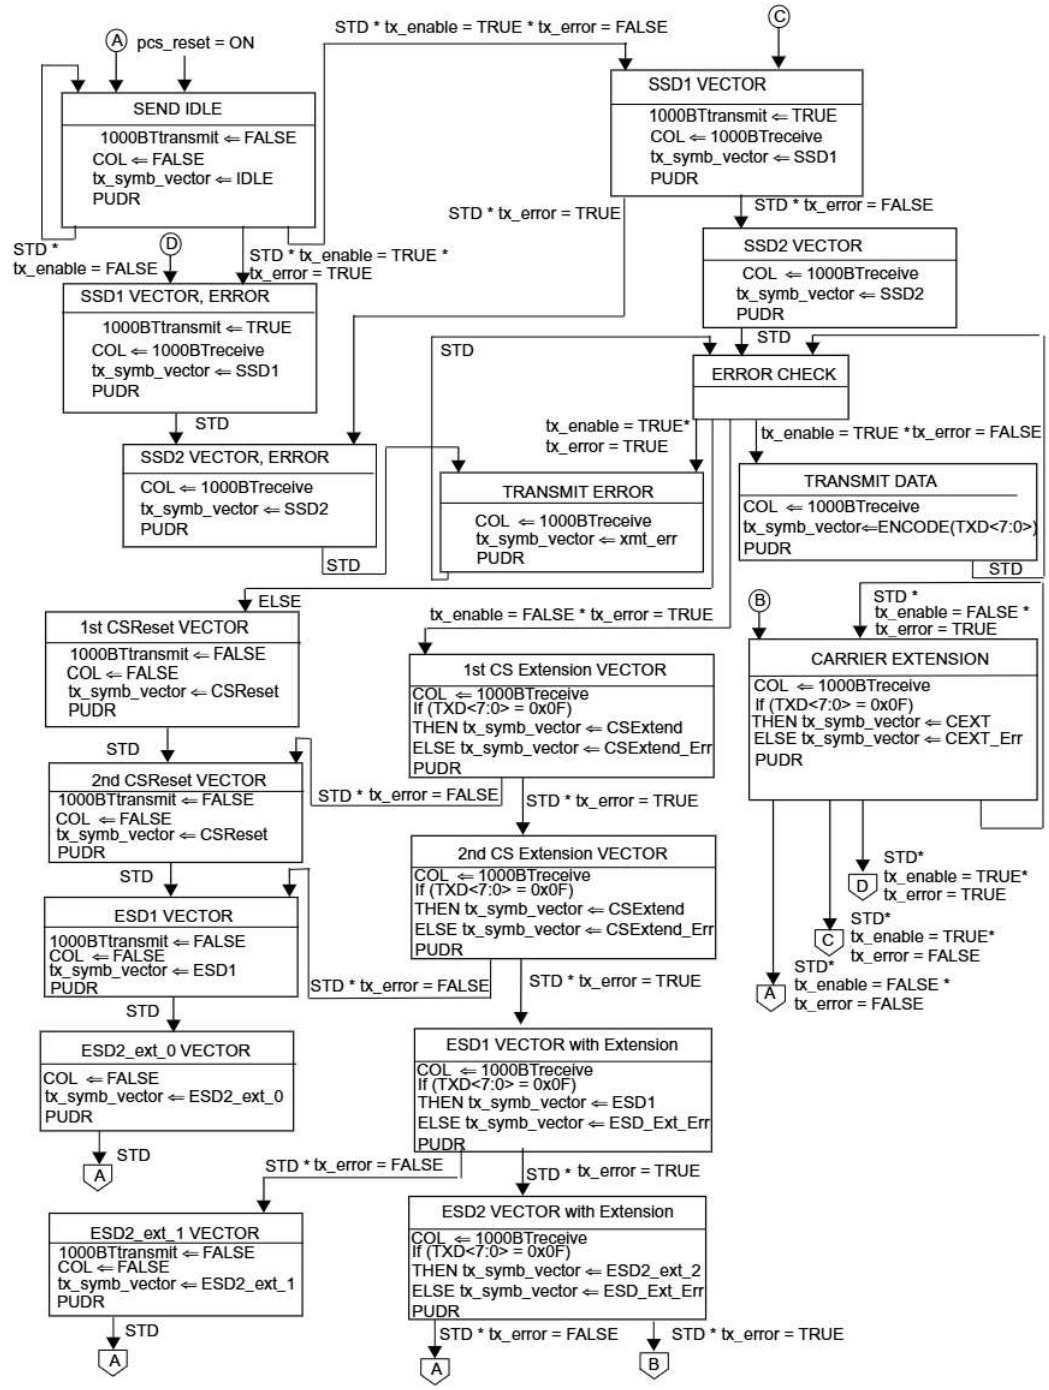


Figure 40–9—PCS Transmit state diagram

Copyright © 2002 IEEE. All rights reserved.

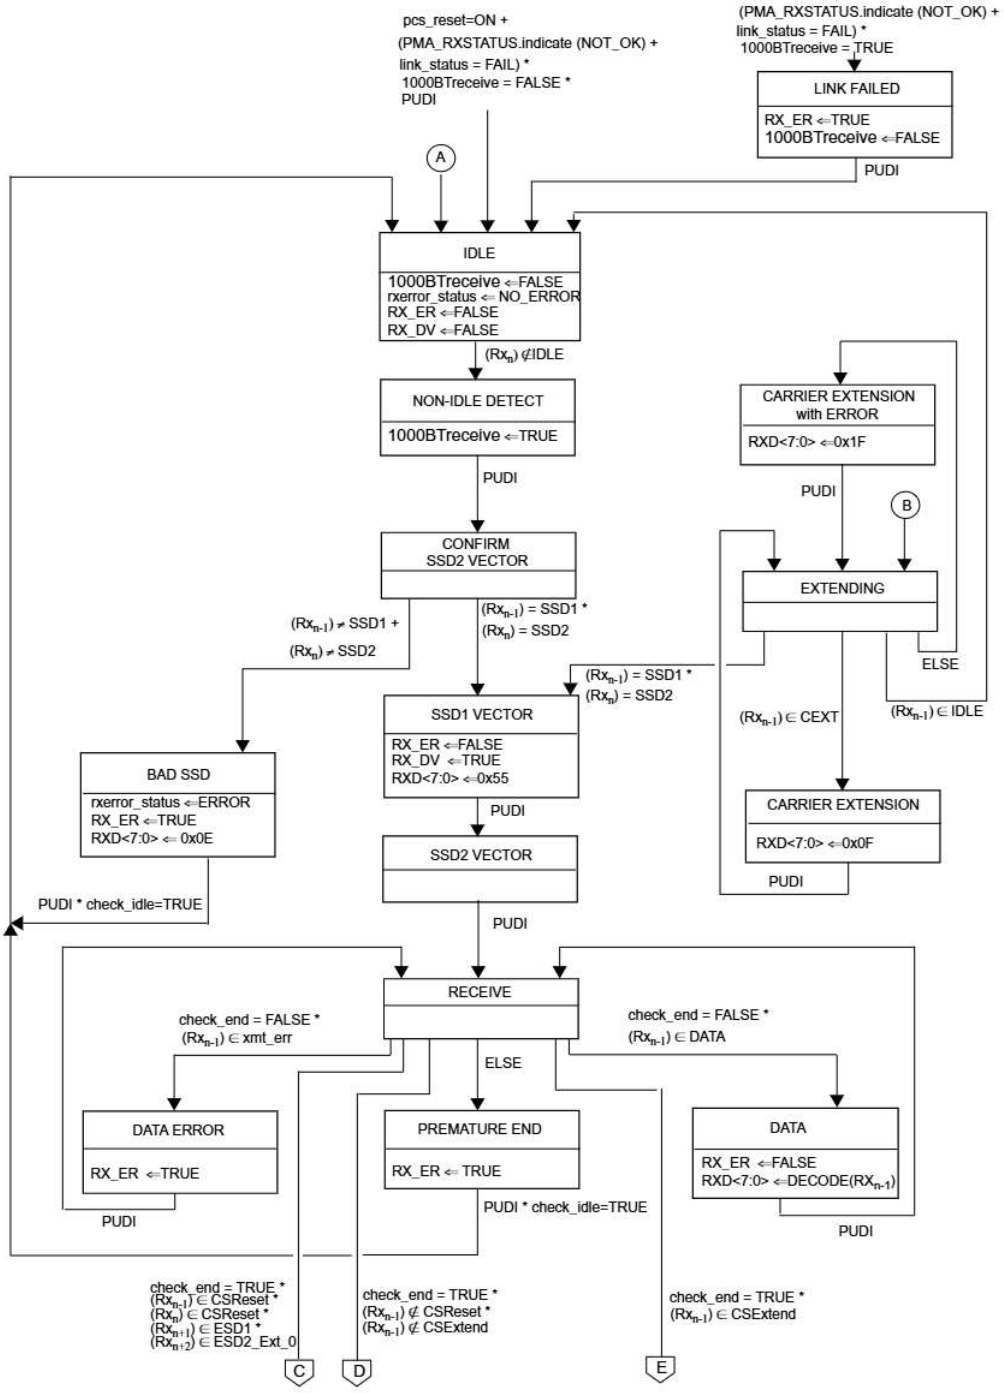


Figure 40–10a—PCS Receive state diagram, part a

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 182 of 379

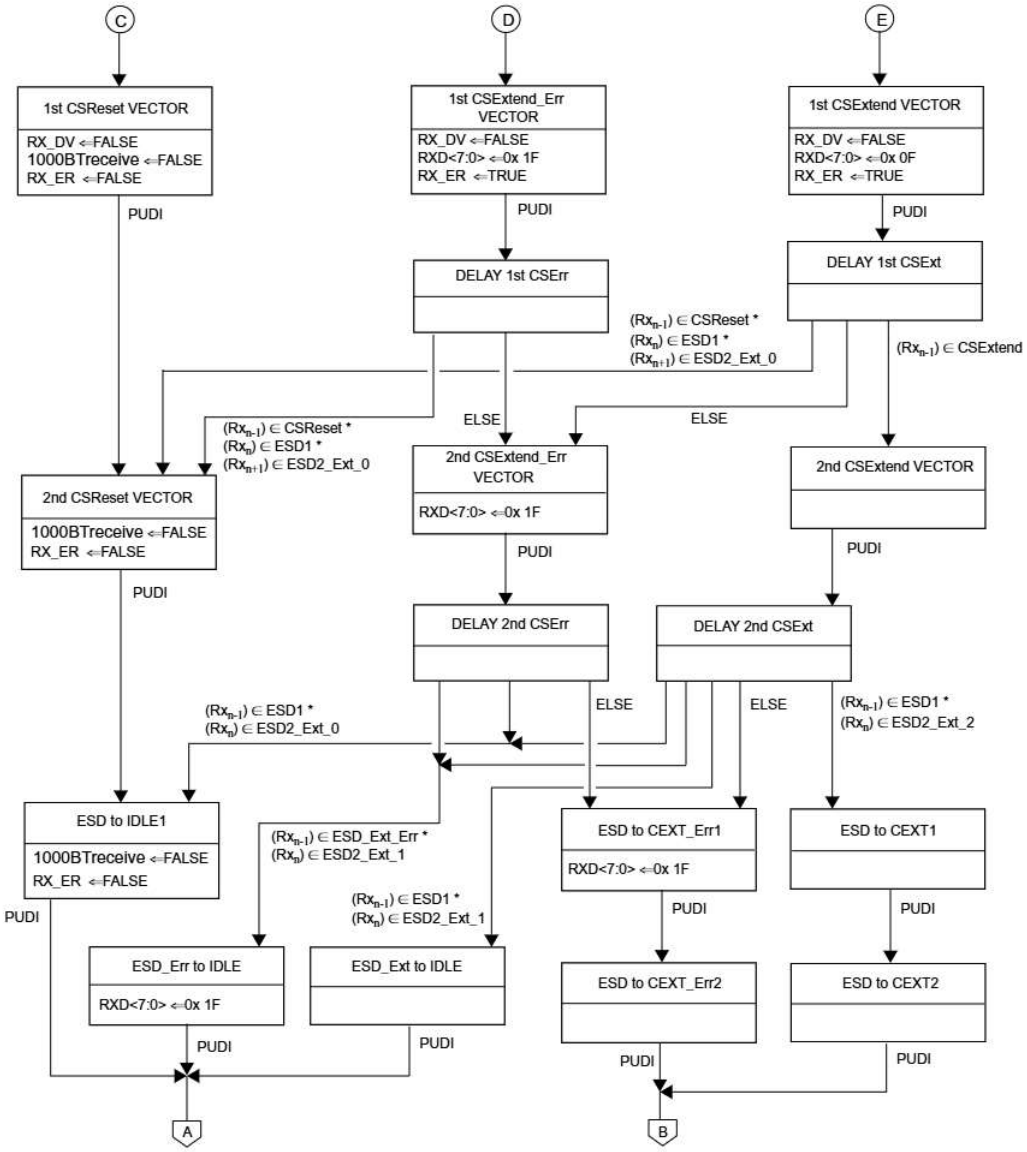


Figure 40–10b—PCS Receive state diagram, part b

Copyright © 2002 IEEE. All rights reserved.

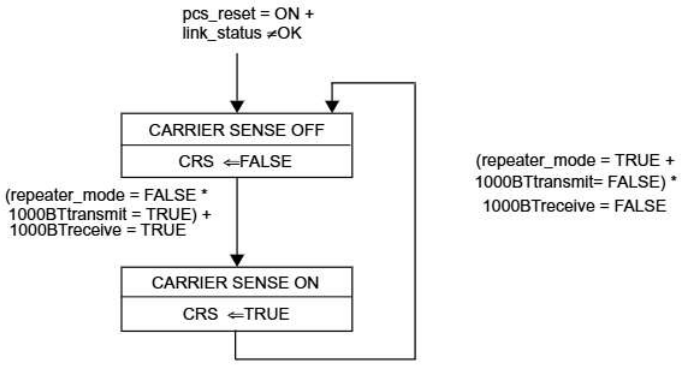


Figure 40–11—PCS Carrier Sense state diagram

40.3.4.1 Supplement to state diagram

Figure 40–12 reiterates the information shown in Figure 40–9 in timing diagram format. It is informative only. Time proceeds from left to right in the figure.

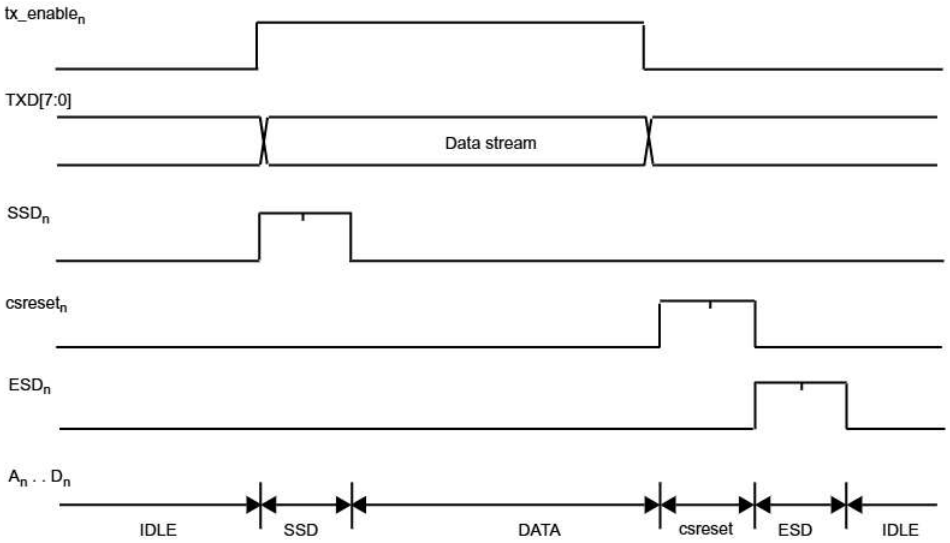


Figure 40–12—PCS sublayer to PMA timing

Copyright © 2002 IEEE. All rights reserved.

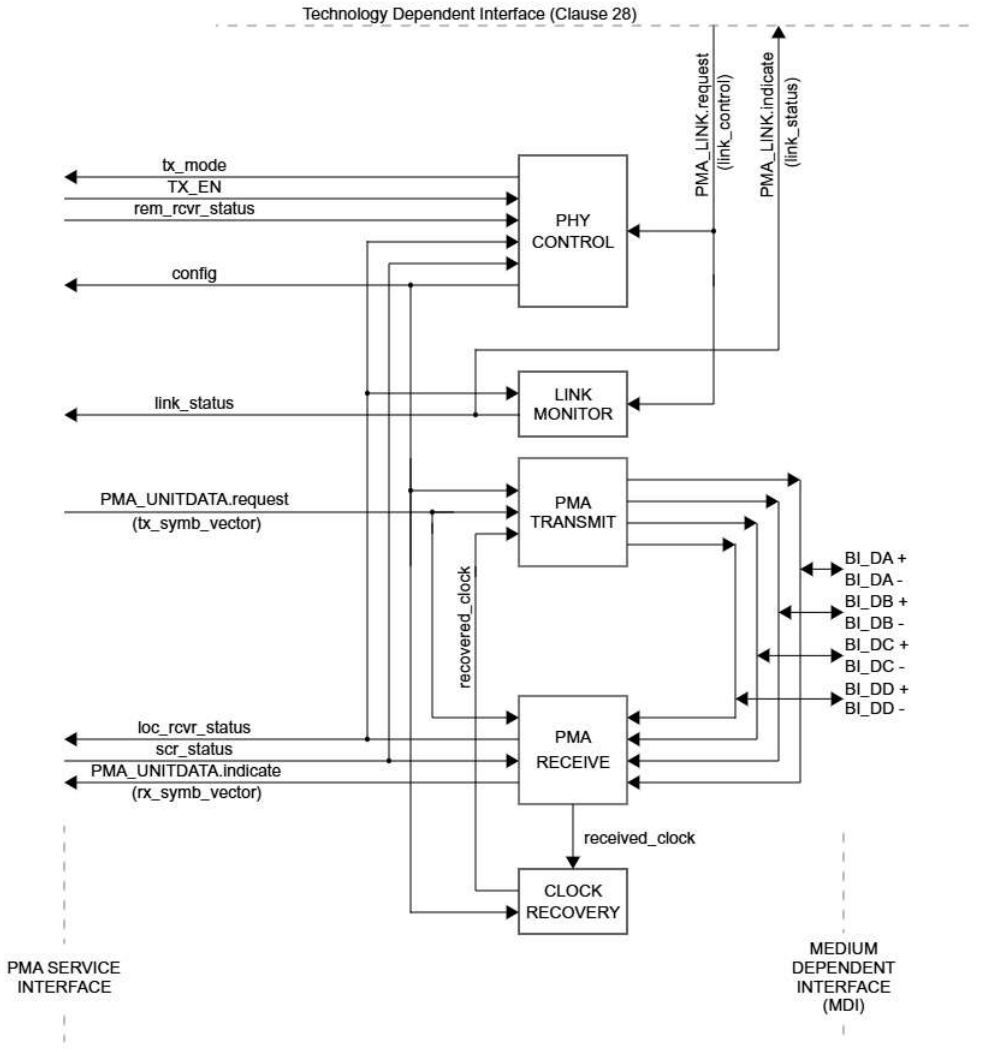
LINKPLAY EXHIBIT 1017
Page 184 of 379

40.4 Physical Medium Attachment (PMA) sublayer

40.4.1 PMA functional specifications

The PMA couples messages from a PMA service interface specified in 40.2.2 to the 1000BASE-T baseband medium, specified in 40.7.

The interface between PMA and the baseband medium is the Medium Dependent Interface (MDI), which is specified in 40.8.



NOTE The recovered_clock arc is shown to indicate delivery of the recovered clock signal back to PMA TRANSMIT for loop timing.

Figure 40–13—PMA reference diagram

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 185 of 379

40.4.2 PMA functions

The PMA sublayer comprises one PMA Reset function and five simultaneous and asynchronous operating functions. The PMA operating functions are PHY Control, PMA Transmit, PMA Receive, Link Monitor, and Clock Recovery. All operating functions are started immediately after the successful completion of the PMA Reset function.

The PMA reference diagram, Figure 40–13, shows how the operating functions relate to the messages of the PMA Service interface and the signals of the MDI. Connections from the management interface, comprising the signals MDC and MDIO, to other layers are pervasive and are not shown in Figure 40–13. The management interface and its functions are specified in Clause 22.

40.4.2.1 PMA Reset function

The PMA Reset function shall be executed whenever one of the two following conditions occur:

a) Power on (see 36.2.5.1.3)
b) The receipt of a request for reset from the management entity

PMA Reset sets pcs_reset=ON while any of the above reset conditions hold true. All state diagrams take the open-ended pma_reset branch upon execution of PMA Reset. The reference diagrams do not explicitly show the PMA Reset function.

40.4.2.2 PMA Transmit function

The PMA Transmit function comprises four synchronous transmitters to generate four 5-level pulse-amplitude modulated signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD. PMA Transmit shall continuously transmit onto the MDI pulses modulated by the quinary symbols given by tx_symb_vector[BI_DA], tx_symb_vector[BI_DB], tx_symb_vector[BI_DC] and tx_symb_vector[BI_DD], respectively. The four transmitters shall be driven by the same transmit clock, TX_TCLK. The signals generated by PMA Transmit shall follow the mathematical description given in 40.4.3.1, and shall comply with the electrical specifications given in 40.6.

When the PMA_CONFIG.indicate parameter config is MASTER, the PMA Transmit function shall source TX_TCLK from a local clock source while meeting the transmit jitter requirements of 40.6.1.2.5. When the PMA_CONFIG.indicate parameter config is SLAVE, the PMA Transmit function shall source TX_TCLK from the recovered clock of 40.4.2.6 while meeting the jitter requirements of 40.6.1.2.5.

40.4.2.3 PMA Receive function

The PMA Receive function comprises four independent receivers for quinary pulse-amplitude modulated signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD. PMA Receive contains the circuits necessary to both detect quinary symbol sequences from the signals received at the MDI over receive pairs BI_DA, BI_DB, BI_DC, and BI_DD and to present these sequences to the PCS Receive function. The signals received at the MDI are described mathematically in 40.4.3.2. The PMA shall translate the signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DB into the PMA_UNITDATA.indicate parameter rx_symb_vector with a symbol error rate of less than 10^{-10} over a channel meeting the requirements of 40.7.

To achieve the indicated performance, it is highly recommended that PMA Receive include the functions of signal equalization, echo and crosstalk cancellation, and sequence estimation. The sequence of code-groups assigned to tx_symb_vector is needed to perform echo and self near-end crosstalk cancellation.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 186 of 379

The PMA Receive function uses the scr_status parameter and the state of the equalization, cancellation, and estimation functions to determine the quality of the receiver performance, and generates the loc_rcvr_status variable accordingly. The precise algorithm for generation of loc_rcvr_status is implementation dependent.

40.4.2.4 PHY Control function

PHY Control generates the control actions that are needed to bring the PHY into a mode of operation during which frames can be exchanged with the link partner. PHY Control shall comply with the state diagram description given in Figure 40–15.

During Auto-Negotiation PHY Control is in the DISABLE 1000BASE-T TRANSMITTER state and the transmitters are disabled. When the Auto-Negotiation process asserts link_control=ENABLE, PHY Control enters the SLAVE SILENT state. Upon entering this state, the maxwait timer is started and PHY Control forces transmission of zeros by setting tx_mode=SEND_Z. The transition out of the SLAVE SILENT state depends on whether the PHY is operating in MASTER or SLAVE mode. In MASTER mode, PHY Control transitions immediately to the TRAINING state. In SLAVE mode, PHY Control transitions to the TRAINING state only after the SLAVE PHY converges its decision feedback equilizer (DFE), acquires timing, and acquires its descrambler state, and sets scr_status=OK.

For the SLAVE PHY, the final convergence of the adaptive filter parameters is completed in the TRAINING state. The MASTER PHY performs all its receiver convergence functions in the TRAINING state. Upon entering the TRAINING state, the minwait_timer is started and PHY Control forces transmission into the idle mode by asserting tx_mode=SEND_I. After the PHY completes successful training and establishes proper receiver operations, PCS Transmit conveys this information to the link partner via transmission of the parameter loc_rcvr_status. (See $Sd_n[2]$ in 40.3.1.3.4.) The link partner's value for loc_rcvr_status is stored in the local device parameter rem_rcvr status. When the minwait_timer expires and the condition loc_rcvr_status=OK is satisfied, PHY Control transitions into either the SEND IDLE OR DATA state if rem_rcvr_status=OK or the SEND IDLE state if rem_rcvr_status=NOT_OK. On entry into either the SEND IDLE or SEND IDLE OR DATA states, the maxwait_timer is stopped and the minwait_timer is started.

The normal mode of operation corresponds to the SEND IDLE OR DATA state, where PHY Control asserts tx_mode=SEND_N and transmission of data over the link can take place. In this state, when no frames have to be sent, idle transmission takes place.

If unsatisfactory receiver operation is detected in the SEND IDLE OR DATA or SEND IDLE states (loc_rcvr_status=NOT_OK) and the minwait_timer has expired, transmission of the current frame is completed and PHY Control enters the SLAVE SILENT state. In the SEND IDLE OR DATA state, whenever a PHY that operates reliably detects unsatisfactory operation of the remote PHY (rem_rcvr_status=NOT_OK) and the minwait_timer has expired, it enters the SEND IDLE state where tx_mode=SEND_I is asserted and idle transmission takes place. In this state, encoding is performed with the parameter loc_rcvr_status=OK. As soon as the remote PHY signals satisfactory receiver operation (rem_rcvr_status=OK) and the minwait_timer has expired, the SEND IDLE OR DATA state is entered.

PHY Control may force the transmit scrambler state to be initialized to an arbitrary value by requesting the execution of the PCS Reset function defined in 40.3.1.1.

40.4.2.5 Link Monitor function

Link Monitor determines the status of the underlying receive channel and communicates it via the variable link_status. Failure of the underlying receive channel typically causes the PMA's clients to suspend normal operation.

The Link Monitor function shall comply with the state diagram of Figure 40–16.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 187 of 379

Upon power on, reset, or release from power down, the Auto-Negotiation algorithm sets link_control=SCAN_FOR_CARRIER and, during this period, sends fast link pulses to signal its presence to a remote station. If the presence of a remote station is sensed through reception of fast link pulses, the Auto-Negotiation algorithm sets link_control=DISABLE and exchanges Auto-Negotiation information with the remote station. During this period, link_status=FAIL is asserted. If the presence of a remote 1000BASE-T station is established, the Auto-Negotiation algorithm permits full operation by setting link_control=ENABLE. As soon as reliable transmission is achieved, the variable link_status=OK is asserted, upon which further PHY operations can take place.

40.4.2.6 Clock Recovery function

The Clock Recovery function couples to all four receive pairs. It may provide independent clock phases for sampling the signals on each of the four pairs.

The Clock Recovery function shall provide clocks suitable for signal sampling on each line so that the symbol-error rate indicated in 40.4.2.3 is achieved. The received clock signal must be stable and ready for use when training has been completed (loc_rcvr_status=OK). The received clock signal is supplied to the PMA Transmit function by received_clock.

40.4.3 MDI

Communication through the MDI is summarized in 40.4.3.1 and 40.4.3.2.

40.4.3.1 MDI signals transmitted by the PHY

The quinary symbols to be transmitted by the PMA on the four pairs BI_DA, BI_DB, BI_DC, and BI_DD are denoted by tx_symb_vector[BI_DA], tx_symb_vector[BI_DB], tx_symb_vector[BI_DC], and tx_symb_vector[BI_DD], respectively. The modulation scheme used over each pair is 5-level Pulse Amplitude Modulation. PMA Transmit generates a pulse-amplitude modulated signal on each pair in the following form:

$$s(t) = \sum_k a_k h_1(t - kT)$$

In the above equation, a_k represents the quinary symbol from the set $\{2, 1, 0, -1, -2\}$ to be transmitted at time kT, and $h_1(t)$ denotes the system symbol response at the MDI. This symbol response shall comply with the electrical specifications given in 40.6.

40.4.3.2 Signals received at the MDI

Signals received at the MDI can be expressed for each pair as pulse-amplitude modulated signals that are corrupted by noise as follows:

$$r(t) = \sum_k a_k h_2(t - kT) + w(t)$$

In this equation, $h_2(t)$ denotes the impulse response of the overall channel between the transmit symbol source and the receive MDI and $w(t)$ is a term that represents the contribution of various noise sources. The four signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DD shall be processed within the PMA Receive function to yield the quinary received symbols rx_symb_vector[BI_DA], rx_symb_vector[BI_DB], rx_symb_vector[BI_DC], and rx_symb_vector[BI_DD].

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 188 of 379

40.4.4 Automatic MDI/MDI-X Configuration

Automatic MDI/MDI-X Configuration is intended to eliminate the need for crossover cables between similar devices. Implementation of an automatic MDI/MDI-X configuration is optional for 1000BASE-T devices. If an automatic configuration method is used, it shall comply with the following specifications. The assignment of pin-outs for a 1000BASE-T crossover function cable is shown in Table 40–12 in 40.8.

40.4.4.1 Description of Automatic MDI/MDI-X state machine

The Automatic MDI/MDI-X state machine facilitates switching the BI_DA(C)+ and BI_DA(C)– with the BI_DB(D)+ and BI_DB(D)– signals respectively prior to the auto-negotiation mode of operation so that FLPs can be transmitted and received in compliance with Clause 28 Auto-Negotiation specifications. The correct polarization of the crossover circuit is determined by an algorithm that controls the switching function. This algorithm uses an 11-bit Linear Feedback Shift Register (LFSR) to create a pseudo-random sequence that each end of the link uses to determine its proposed configuration. Upon making the selection to either MDI or MDI-X, the node waits for a specified amount of time while evaluating its receive channel to determine whether the other end of the link is sending link pulses or PHY-dependent data. If link pulses or PHY-dependent data are detected, it remains in that configuration. If link pulses or PHY-dependent data are not detected, it increments its LFSR and makes a decision to switch based on the value of the next bit. The state machine does not move from one state to another while link pulses are being transmitted.

40.4.4.2 Pseudo-random sequence generator

One possible implementation of the pseudo-random sequence generator using a linear-feedback shift register is shown in Figure 40–14. The bits stored in the shift register delay line at time n are denoted by S[10:0]. At each sample period, the shift register is advanced by one bit and one new bit represented by S[0] is generated. Switch control is determined by S[10].

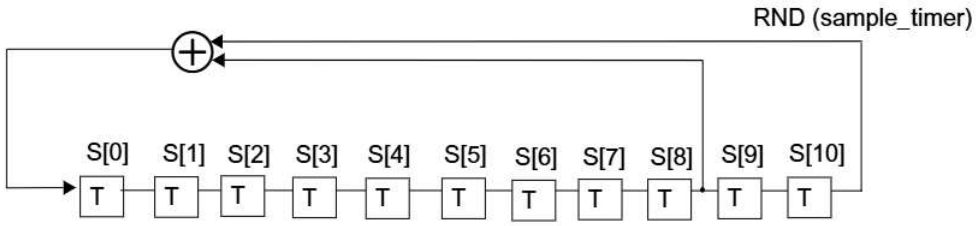


Figure 40–14—Automatic MDI/MDI-X linear-feedback shift register

40.4.5 State variables

40.4.5.1 State diagram variables

config

The PMA shall generate this variable continuously and pass it to the PCS via the PMA_CONFIG.indicate primitive.

Values: MASTER or SLAVE

link_control

This variable is defined in 28.2.6.2.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 189 of 379

Link_Det

This variable indicates linkpulse = true or link_status = READY has occurred at the receiver since the last time sample_timer has been started.

Values: TRUE: linkpulse = true or link_status = READY has occurred since the last time sample_timer has been started.

FALSE: otherwise

linkpulse

This variable is defined in 28.2.6.3.

link_status

This variable is defined in 28.2.6.1.

loc_rcvr_status

Variable set by the PMA Receive function to indicate correct or incorrect operation of the receive link for the local PHY.

Values: OK: The receive link for the local PHY is operating reliably.
NOT_OK: Operation of the receive link for the local PHY is unreliable.

MDI_Status

This variable defines the condition of the Automatic MDI/MDI-X physical connection.

Values: MDI: The BI_DA, BI_DB, BI_DC, and BI_DD pairs follow the connections as described in the MDI column of Table 40–12.

MDI-X: The BI_DA, BI_DB, BI_DC, and BI_DD pairs follow the connections as described in the MDI-X column of Table 40–12.

pma_reset

Allows reset of all PMA functions.

Values: ON or OFF

Set by: PMA Reset

rem_rcvr_status

Variable set by the PCS Receive function to indicate whether correct operation of the receive link for the remote PHY is detected or not.

Values: OK: The receive link for the remote PHY is operating reliably.

NOT_OK: Reliable operation of the receive link for the remote PHY is not detected.

RND (sample_timer)

This variable is defined as bit S[10] of the LSFR described in 40.4.4.2

scr_status

The scr_status parameter as communicated by the PMA_SCRSTATUS.request primitive.

Values: OK: The descrambler has achieved synchronization.

NOT_OK: The descrambler is not synchronized.

T_Pulse

This variable indicates that a linkpulse is being transmitted to the MDI.

Values: TRUE: Pulse being transmitted to the MDI

FALSE: Otherwise

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 190 of 379

tx_enable

The tx_enable parameter generated by PCS Transmit as specified in Figure 40–8.

Values: TRUE or FALSE as per 40.3.3.1.

tx_mode

PCS Transmit sends code-groups according to the value assumed by this variable.

Values: SEND_N: This value is continuously asserted when transmission of sequences of code-groups representing a GMII data stream, control information, or idle mode is to take place.

SEND_I: This value is continuously asserted when transmission of sequences of code-groups representing the idle mode is to take place.

SEND_Z: This value is asserted when transmission of zero code-groups is to take place.

40.4.5.2 Timers

All timers operate in the manner described in 14.2.3.2 with the following addition. A timer is reset and stops counting upon entering a state where "stop timer" is asserted.

A_timer

An asynchronous (to the Auto-Crossover State Machine) free-running timer that provides for a relatively arbitrary reset of the state machine to its initial state. This timer is used to reduce the probability of a lock-up condition where both nodes have the same identical seed initialization at the same point in time.

Values: The condition A_timer_done becomes true upon timer expiration.

Duration: This timer shall have a period of 1.3 s ± 25%.

Initialization of A_timer is implementation specific.

maxwait_timer

A timer used to limit the amount of time during which a receiver dwells in the SLAVE SILENT and TRAINING states. The timer shall expire 750 ± 10 ms if config = MASTER or 350 ± 5 ms if config = SLAVE. This timer is used jointly in the PHY Control and Link Monitor state diagrams. The maxwait_timer is tested by the Link Monitor to force link_status to be set to FAIL if the timer expires and loc_rcvr_status is NOT_OK. See Figure 40–15.

minwait_timer

A timer used to determine the minimum amount of time the PHY Control stays in the TRAINING, SEND IDLE, or DATA states. The timer shall expire 1 ± 0.1μs after being started.

sample_timer

This timer provides a long enough sampling window to ensure detection of Link Pulses or link_status, if they exist at the receiver.

Values: The condition sample_timer_done becomes true upon timer expiration.

Duration: This timer shall have a period of 62 ± 2 ms.

stabilize_timer

A timer used to control the minimum time that loc_rcvr_status must be OK before a transition to Link Up can occur. The timer shall expire 1 ± 0.1 μs after being started.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 191 of 379

40.4.6 State Diagrams

40.4.6.1 PHY Control state diagram

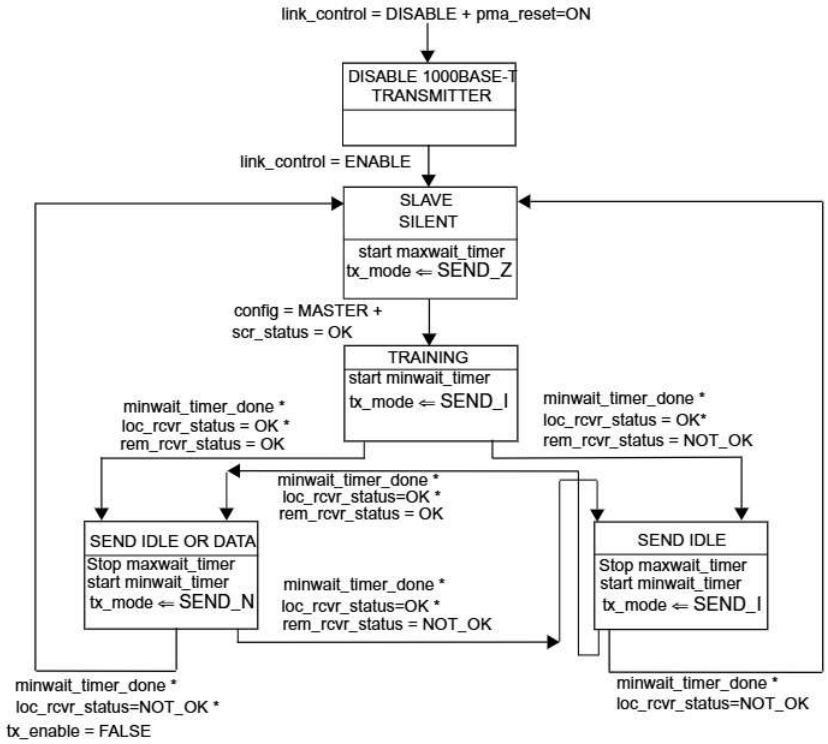
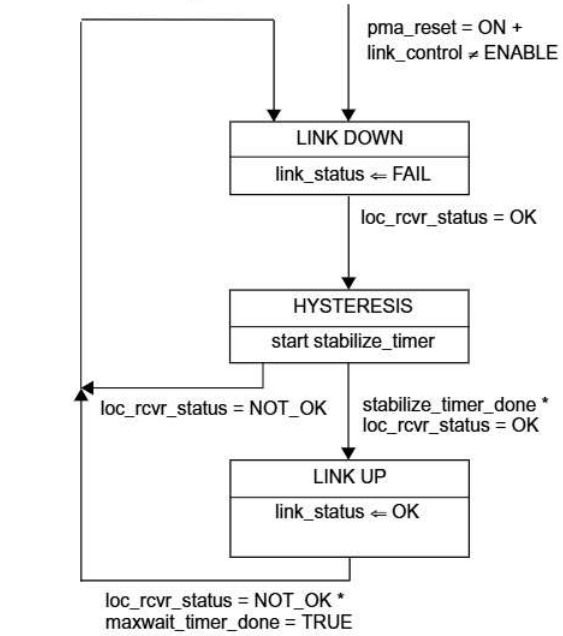


Figure 40–15—PHY Control state diagram

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 192 of 379

40.4.6.2 Link Monitor state diagram



NOTES
1 maxwait_timer is started in PHY Control state diagram (see Figure 40—15).≥
2 The variables link_control and link_status are designated as link_control_(1GigT) and link_status_(1GigT), respectively, by the Auto-Negotiation Arbitration state diagram (Figure 28—16)

Figure 40–16—Link Monitor state diagram

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 193 of 379

40.4.6.2.1 Auto Crossover state diagram

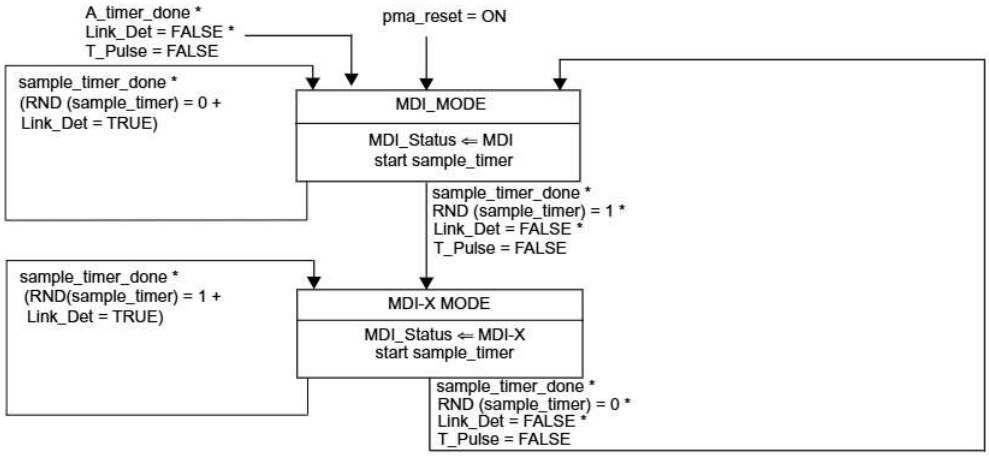


Figure 40–17—Auto Crossover state diagram

40.5 Management interface

1000BASE-T makes extensive use of the management functions provided by the MII Management Interface (see 22.2.4), and the communication and self-configuration functions provided by Auto-Negotiation (Clause 28.)

40.5.1 Support for Auto-Negotiation

All 1000BASE-T PHYs shall provide support for Auto-Negotiation (Clause 28) and shall be capable of operating as MASTER or SLAVE.

Auto-Negotiation is performed as part of the initial set-up of the link, and allows the PHYs at each end to advertise their capabilities (speed, PHY type, half or full duplex) and to automatically select the operating mode for communication on the link. Auto-negotiation signaling is used for the following two primary purposes for 1000BASE-T:

a) To negotiate that the PHY is capable of supporting 1000BASE-T half duplex or full duplex transmission.
b) To determine the MASTER-SLAVE relationship between the PHYs at each end of the link.

This relationship is necessary for establishing the timing control of each PHY. The 1000BASE-T MASTER PHY is clocked from a local source. The SLAVE PHY uses loop timing where the clock is recovered from the received data stream.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 194 of 379

40.5.1.1 1000BASE-T use of registers during Auto-Negotiation

A 1000BASE-T PHY shall use the management register definitions and values specified in Table 40–3.

Table 40–3—1000BASE-T Registers

Register	Bit	Name	Description	Type [a]
0	0.15:0	MII control register	Defined in 28.2.4.1.1	RO
1	1.15:0	MII status register	Defined in 28.2.4.1.2	RO
4	4.15:0	Auto-Negotiation advertisement register	The Selector Field (4.4:0) is set to the appropriate code as specified in Annex 28A. The Technology Ability Field bits 4.12:5 are set to the appropriate code as specified in Annexes 28B and 28D. Bit 4.15 is set to logical one to indicate the desired exchange of Next Pages describing the gigabit extended capabilities.	R/W
5	5.15:0	Auto-Negotiation link partner ability register	Defined in 28.2.4.1.4. 1000BASE-T implementations do not use this register to store Auto-Negotiation Link Partner Next Page data.	RO
6	6.15:0	Auto-Negotiation expansion register	Defined in 28.2.4.1.5	RO
7	7.15:0	Auto-Negotiation Next Page transmit register	Defined in 28.2.4.1.6	R/W
8	8.15:0	Auto-Negotiation link partner Next Page register	Defined in 28.2.4.1.8	RO
9	9.15:13	Test mode bits	Transmitter test mode operations are defined by bits 9.15:13 as described in 40.6.1.1.2 and Table 40–7. The default values for bits 9.15:13 are all zero.	R/W
9	9.12	MASTER-SLAVE Manual Config Enable	1=Enable MASTER-SLAVE Manual configuration value 0=Disable MASTER-SLAVE Manual configuration value Default bit value is 0.	R/W
9	9.11	MASTER-SLAVE Config Value	1=Configure PHY as MASTER during MASTER-SLAVE negotiation, only when 9.12 is set to logical one. 0=Configure PHY as SLAVE during MASTER-SLAVE negotiation, only when 9.12 is set to logical one. Default bit value is 0.	R/W
9	9.10	Port type	Bit 9.10 is to be used to indicate the preference to operate as MASTER (multiport device) or as SLAVE (single-port device) if the MASTER-SLAVE Manual Configuration Enable bit, 9.12, is not set. Usage of this bit is described in 40.5.2. 1=Multiport device 0=single-port device	R/W
9	9.9	1000BASE-T Full Duplex	1 = Advertise PHY is 1000BASE-T full duplex capable. 0 = Advertise PHY is not 1000BASE-T full duplex capable.	R/W
9	9.8	1000BASE-T Half Duplex	1 = Advertise PHY is 1000BASE-T half duplex capable. 0 = Advertise PHY is not 1000BASE-T half duplex capable.	R/W

[a] R/W = Read/Write, RO = Read Only, SC = Self Clearing, LH = Latch High

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 195 of 379

Table 40–3—1000BASE-T Registers *(continued)*

Register	Bit	Name	Description	Type [a]
9	9.7:0	Reserved	Write as 0, ignore on read.	R/W
10	10.15	MASTER-SLAVE configuration fault	Configuration fault, as well as the criteria and method of fault detection, is PHY specific. The MASTER-SLAVE Configuration Fault bit will be cleared each time register 10 is read via the management interface and will be cleared by a 1000BASE-T PMA reset. This bit will self clear on Auto-Negotiation enable or Auto-Negotiation complete.This bit will be set if the number of failed MASTER-SLAVE resolutions reaches 7. For l000BASE-T, the fault condition will occur when both PHYs are forced to be MASTERs or SLAVEs at the same time using bits 9.12 and 9.11. Bit 10.15 should be set via the MASTER-SLAVE Configuration Resolution function described in 40.5.2. 1 = MASTER-SLAVE configuration fault detected 0 = No MASTER-SLAVE configuration fault detected	RO/LH/SC
10	10.14	MASTER-SLAVE configuration resolution	1 = Local PHY configuration resolved to MASTER 0 = Local PHY configuration resolved to SLAVE	RO
10	10.13	Local Receiver Status	1 = Local Receiver OK (loc_rcvr_status=OK) 0 = Local Receiver not OK (loc_rcvr_status=NOT_OK) Defined by the value of loc_rcvr_status as per 40.4.5.1.	RO
10	10.12	Remote Receiver Status	1 = Remote Receiver OK (rem_rcvr_status=OK) 0 = Remote Receiver not OK (rem_rcvr_status=NOT_OK) Defined by the value of rem_rcvr_status as per 40.4.5.1.	RO
10	10.11	LP 1000T FD	1 = Link Partner is capable of 1000BASE-T full duplex 0 = Link Partner is not capable of 1000BASE-T full duplex This bit is guaranteed to be valid only when the Page received bit (6.1) has been set to 1.	RO
10	10.10	LP 1000T HD	1 = Link Partner is capable of 1000BASE-T half duplex 0 = Link Partner is not capable of 1000BASE-T half duplex This bit is guaranteed to be valid only when the Page received bit (6.1) has been set to 1.	RO
10	10.9:8	Reserved	Reserved	RO
10	10.7:0	Idle Error Count	Bits 10.7:0 indicate the Idle Error count, where 10.7 is the most significant bit. These bits contain a cumulative count of the errors detected when the receiver is receiving idles and PMA_TXMODE.indicate is equal to SEND_N (indicating that both local and remote receiver status have been detected to be OK). The counter is incremented every symbol period that rxerror_status is equal to ERROR. These bits are reset to all zeros when the error count is read by the management function or upon execution of the PCS Reset function and are to be held at all ones in case of overflow (see 30.5.1.1.11).	RO/SC
15	15.15:12	Extended status register	See 22.2.4.4	RO

[a] R/W = Read/Write, RO = Read Only, SC = Self Clearing, LH = Latch High

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 196 of 379

40.5.1.2 1000BASE-T Auto-Negotiation page use

1000BASE-T PHYs shall exchange one Auto-Negotiation Base Page, a 1000BASE-T formatted Next Page, and two 1000BASE-T unformatted Next Pages in sequence, without interruption, as specified in Table 40–4. Additional Next Pages can be exchanged as described in Annex 40C.

Note that the Acknowledge 2 bit is not utilized and has no meaning when used for the 1000BASE-T message page exchange.

Table 40–4—1000BASE-T Base and Next Pages bit assignments

Bit	Bit definition	Register location
BASE PAGE		
D15	1 (to indicate that Next Pages follow)	
D14:D1	As specified in 28.2.1.2	Management register 4
PAGE 0 (Message Next Page)		
M10:M0	8	
PAGE 1 (Unformatted Next Page)		
U10:U5	Reserved transmit as 0	
U4	1000BASE-T half duplex (1 = half duplex and 0 = no half duplex)	GMII register 9.8 (MASTER-SLAVE Control register)
U3	1000BASE-T full duplex (1 = full duplex and 0 = no full duplex)	GMII register 9.9 (MASTER-SLAVE Control register)
U2	1000BASE-T port type bit (1 = multiport device and 0 = single-port device)	GMII register 9.10 (MASTER-SLAVE Control register)
U1	1000BASE-T MASTER-SLAVE Manual Configuration value (1 = MASTER and 0 = SLAVE.) This bit is ignored if 9.12 = 0.	GMII register 9.11 (MASTER-SLAVE Control register)
U0	1000BASE-T MASTER-SLAVE Manual Configuration Enable (1 = Manual Configuration Enable.) This bit is intended to be used for manual selection in a particular MASTER-SLAVE mode and is to be used in conjunction with bit 9.11.	GMII register 9.12 (MASTER-SLAVE Control register)
PAGE 2 (Unformatted Next Page)		
U10	1000BASE-T MASTER-SLAVE Seed Bit 10 (SB10) (MSB)	MASTER-SLAVE Seed Value (10:0)
U9	1000BASE-T MASTER-SLAVE Seed Bit 9 (SB9)	
U8	1000BASE-T MASTER-SLAVE Seed Bit 8 (SB8)	
U7	1000BASE-T MASTER-SLAVE Seed Bit 7 (SB7)	
U6	1000BASE-T MASTER-SLAVE Seed Bit 6 (SB6)	
U5	1000BASE-T MASTER-SLAVE Seed Bit 5 (SB5)	
U4	1000BASE-T MASTER-SLAVE Seed Bit 4 (SB4)	
U3	1000BASE-T MASTER-SLAVE Seed Bit 3 (SB3)	
U2	1000BASE-T MASTER-SLAVE Seed Bit 2 (SB2)	
U1	1000BASE-T MASTER-SLAVE Seed Bit 1 (SB1)	
U0	1000BASE-T MASTER-SLAVE Seed Bit 0 (SB0)	

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 197 of 379

40.5.1.3 Sending Next Pages

Implementors who do not wish to send additional Next Pages (i.e., Next Pages in addition to those required to perform PHY configuration as defined in this clause) can use Auto-Negotiation as defined in Clause 28 and the Next Pages defined in 40.5.1.2. Implementors who wish to send additional Next Pages are advised to consult Annex 40C.

40.5.2 MASTER-SLAVE configuration resolution

Since both PHYs that share a link segment are capable of being MASTER or SLAVE, a prioritization scheme exists to ensure that the correct mode is chosen. The MASTER-SLAVE relationship shall be determined during Auto-Negotiation using Table 40–5 with the 1000BASE-T Technology Ability Next Page bit values specified in Table 40–4 and information received from the link partner. This process is conducted at the entrance to the FLP LINK GOOD CHECK state shown in the Arbitration state diagram (Figure 28–13.)

The following four equations are used to determine these relationships:

manual_MASTER = U0 * U1
manual_SLAVE = U0 * !U1
single-port device = !U0 * !U2,
multiport device = !U0 * U2

where

U0 is bit 0 of unformatted page 1,
U1 is bit 1 of unformatted page 1, and
U2 is bit 2 of unformatted page 1 (see Table 40–4).

A 1000BASE-T PHY is capable of operating either as the MASTER or SLAVE. In the scenario of a link between a single-port device and a multiport device, the preferred relationship is for the multiport device to be the MASTER PHY and the single-port device to be the SLAVE. However, other topologies may result in contention. The resolution function of Table 40–5 is defined to handle any relationship conflicts.

Table 40–5—1000BASE-T MASTER-SLAVE configuration resolution table

Local device type	Remote device type	Local device resolution	Remote device resolution
single-port device	multiport device	SLAVE	MASTER
single-port device	manual_MASTER	SLAVE	MASTER
manual_SLAVE	manual_MASTER	SLAVE	MASTER
manual_SLAVE	multiport device	SLAVE	MASTER
multiport device	manual_MASTER	SLAVE	MASTER
manual_SLAVE	single-port device	SLAVE	MASTER
multiport device	single-port device	MASTER	SLAVE
multiport device	manual_SLAVE	MASTER	SLAVE

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 198 of 379

Table 40–5—1000BASE-T MASTER-SLAVE configuration resolution table *(continued)*

Local device type	Remote device type	Local device resolution	Remote device resolution
manual_MASTER	manual_SLAVE	MASTER	SLAVE
manual_MASTER	single-port device	MASTER	SLAVE
single-port device	manual_SLAVE	MASTER	SLAVE
manual_MASTER	multiport device	MASTER	SLAVE
multiport device	multiport device	The device with the higher SEED value is configured as MASTER, otherwise SLAVE.	The device with the higher SEED value is configured as MASTER, otherwise SLAVE.
single-port device	single-port device	The device with the higher SEED value is configured as MASTER, otherwise SLAVE	The device with the higher SEED value is configured as MASTER, otherwise SLAVE.
manual_SLAVE	manual_SLAVE	MASTER-SLAVE configuration fault	MASTER-SLAVE configuration fault
manual_MASTER	manual_MASTER	MASTER-SLAVE configuration fault	MASTER-SLAVE configuration fault

The rationale for the hierarchy illustrated in Table 40–5 is straightforward. A 1000BASE-T multiport device has higher priority than a single-port device to become the MASTER. In the case where both devices are of the same type, e.g., both devices are multiport devices, the device with the higher MASTER-SLAVE seed bits (SB0...SB10), where SB10 is the MSB, shall become the MASTER and the device with the lower seed value shall become the SLAVE. In case both devices have the same seed value, both should assert link_status_1GigT=FAIL (as defined in 28.3.1) to force a new cycle through Auto-Negotiation. Successful completion of the MASTER-SLAVE resolution shall be treated as MASTER-SLAVE configuration resolution complete.

The method of generating a random or pseudorandom seed is left to the implementor. The generated random seeds should belong to a sequence of independent, identically distributed integer numbers with a uniform distribution in the range of 0 to $2^{11}-2$. The algorithm used to generate the integer should be designed to minimize the correlation between the number generated by any two devices at any given time. A seed counter shall be provided to track the number of seed attempts. The seed counter shall be set to zero at start-up and shall be incremented each time a seed is generated. When MASTER-SLAVE resolution is complete, the seed counter shall be reset to 0 and bit 10.15 shall be set to logical zero. A MASTER-SLAVE resolution fault shall be declared if resolution is not reached after the generation of seven seeds.

The MASTER-SLAVE Manual Configuration Enable bit (control register bit 9.12) and the MASTER-SLAVE Config Value bit (control register bit 9.11) are used to manually set a device to become the MASTER or the SLAVE. In case both devices are manually set to become the MASTER or the SLAVE, this condition shall be flagged as a MASTER-SLAVE Configuration fault condition, thus the MASTER-SLAVE Configuration fault bit (status register bit 10.15) shall be set to logical one. The MASTER-SLAVE Configuration fault condition shall be treated as MASTER-SLAVE configuration resolution complete and link_status_1GigT shall be set to FAIL, because the MASTER-SLAVE relationship was not resolved. This will force a new cycle through Auto-Negotiation after the link_fail_inhibit_timer has expired. Determination of MASTER-SLAVE values occur on the entrance to the FLP LINK GOOD CHECK state (Figure 28–16) when the highest common denominator (HCD) technology is 1000BASE-T. The resulting MASTER-SLAVE value is used by the 1000BASE-T PHY control (40.4.2.4).

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 199 of 379

If MASTER-SLAVE Manual Configuration is disabled (bit 9.12 is set to 0) and the local device detects that both the local device and the remote device are of the same type (either multiport device or single-port device) and that both have generated the same random seed, it generates and transmits a new random seed for MASTER-SLAVE negotiation by setting link_status to FAIL and cycling through the Auto-Negotiation process again.

The MASTER-SLAVE configuration process returns one of the three following outcomes:

a) *Successful:* Bit 10.15 of the 1000BASE-T Status Register is set to logical zero and bit 10.14 is set to logical one for MASTER resolution or for logical zero for SLAVE resolution. 1000BASE-T returns control to Auto_Negotiation (at the entrance to the FLP LINK GOOD CHECK state in Figure 28–16) and passes the value MASTER or SLAVE to PMA_CONFIG.indicate (see 40.2.4.)
b) *Unsuccessful:* link_status_1GigT is set to FAIL and Auto-Negotiation restarts (see Figure 28–16.)
c) *Fault detected:* (This happens when both end stations are set for manual configuration and both are set to MASTER or both are set to SLAVE.) Bit 10.15 of the 1000BASE-T Status Register is set to logical one to indicate that a configuration fault has been detected. This bit also is set when seven attempts to configure the MASTER SLAVE relationship via the seed method have failed. When a fault is detected, link_status_1GigT is set to FAIL, causing Auto-Negotiation to cycle through again.

NOTE—MASTER-SLAVE arbitration only occurs if 1000BASE-T is selected as the highest common denominator; otherwise, it is assumed to have passed this condition.

40.6 PMA electrical specifications

This subclause defines the electrical characteristics of the PMA.

Common-mode tests use the common-mode return point as a reference.

40.6.1 PMA-to-MDI interface tests

40.6.1.1 Isolation requirement

The PHY shall provide electrical isolation between the port device circuits, including frame ground (if any) and all MDI leads. This electrical separation shall withstand at least one of the following electrical strength tests:

a) 1500 V rms at 50 Hz to 60 Hz for 60 s, applied as specified in Section 5.3.2 of IEC 60950: 1991.
b) 2250 Vdc for 60 s, applied as specified in Section 5.3.2 of IEC 60950: 1991.
c) A sequence of ten 2400 V impulses of alternating polarity, applied at intervals of not less than 1 s. The shape of the impulses shall be 1.2/50 μs (1.2 μs virtual front time, 50 μs virtual time or half value), as defined in IEC 60060.

There shall be no insulation breakdown, as defined in Section 5.3.2 of IEC 60950: 1991, during the test. The resistance after the test shall be at least 2 MΩ, measured at 500 Vdc.

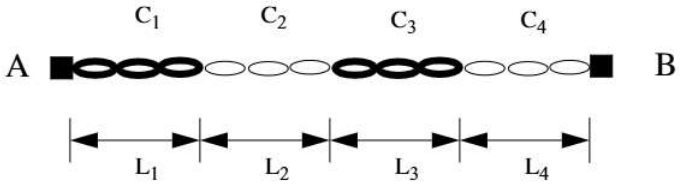
40.6.1.1.1 Test channel

To perform the transmitter MASTER-SLAVE timing jitter tests described in this clause, a test channel is required to ensure that jitter is measured under conditions of poor signal to echo ratio. This test channel shall be constructed by combining 100 and 120 Ω cable segments that both meet or exceed ISO/IEC 11801 Category 5 specifications for each pair, as shown in Figure 40–18, with the lengths and additional restrictions on parameters described in Table 40–6. The ends of the test channel shall be terminated with connectors meeting or exceeding ANSI/TIA/EIA-568-A:1995 or ISO/IEC 11801:1995 Category 5 specifications. The

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 200 of 379

return loss of the resulting test channel shall meet the return loss requirements of 40.7.2.3 and the crosstalk requirements of 40.7.3.



Identical for each of the four pairs.

Figure 40–18—Test channel topology for each cable pair

Table 40–6—Test channel cable segment specifications

Cable segment	Length (meters)	Characteristic impedance (at frequencies > 1 MHz)	Attenuation (per 100 meters at 31.25 MHz)
1	L_1=1.20	120 ± 5Ω	7.8 to 8.8 dB
2	L_2=x	100 ± 5Ω	10.8 to 11.8 dB
3	L_3=1.48	120 ± 5Ω	7.8 to 8.8 dB
4	L_4=y	100 ± 5Ω	10.8 to 11.8 dB

NOTE—x is chosen so that the total delay of segments C1, C2, and C3, averaged across all pairs, is equal to 570 ns at 31.25 MHz; however, if this would cause the total attenuation of segments C1, C2, and C3, averaged across all pairs, to exceed the worst case insertion loss specified in 40.7.2.1 then x is chosen so that the total attenuation of segments C1, C2, and C3, averaged across all pairs, does not violate 40.7.2.1 at any frequencies. The value of y is chosen so that the total attenuation of segments C1, C2, C3, and C4, averaged across all pairs, does not violate 40.7.2.1 at any frequency (y may be 0).

40.6.1.1.2 Test modes

The test modes described below shall be provided to allow for testing of the transmitter waveform, transmitter distortion, and transmitted jitter.

For a PHY with a GMII interface, these modes shall be enabled by setting bits 9.13:15 (1000BASE-T Control Register) of the GMII Management register set as shown in Table 40–7. These test modes shall only change the data symbols provided to the transmitter circuitry and shall not alter the electrical and jitter characteristics of the transmitter and receiver from those of normal (non-test mode) operation. PHYs without a GMII shall provide a means to enable these modes for conformance testing.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 201 of 379

Table 40–7—GMII management register settings for test modes

Bit 1 (9.15)	Bit 2 (9.14)	Bit 3 (9.13)	Mode
0	0	0	Normal operation
0	0	1	Test mode 1—Transmit waveform test
0	1	0	Test mode 2—Transmit jitter test in MASTER mode
0	1	1	Test mode 3—Transmit jitter test in SLAVE mode
1	0	0	Test mode 4—Transmitter distortion test
1	0	1	Reserved, operations not identified.
1	1	0	Reserved, operations not identified.
1	1	1	Reserved, operations not identified.

When test mode 1 is enabled, the PHY shall transmit the following sequence of data symbols A_n, B_n, C_n, D_n, of 40.3.1.3.6 continually from all four transmitters:

{{+2 followed by 127 0 symbols}, {–2 followed by 127 0 symbols},{+1 followed by 127 0 symbols}, {–1 followed by 127 0 symbols}, {128 +2 symbols, 128 –2 symbols, 128 +2 symbols, 128 –2 symbols}, {1024 0 symbols}}

This sequence is repeated continually without breaks between the repetitions when the test mode is enabled. A typical transmitter output is shown in Figure 40–19. The transmitter shall time the transmitted symbols from a 125.00 MHz ± 0.01% clock in the MASTER timing mode.

When test mode 2 is enabled, the PHY shall transmit the data symbol sequence {+2, –2} repeatedly on all channels. The transmitter shall time the transmitted symbols from a 125.00 MHz ± 0.01% clock in the MASTER timing mode.

When test mode 3 is enabled, the PHY shall transmit the data symbol sequence {+2, –2} repeatedly on all channels. The transmitter shall time the transmitted symbols from a 125.00 MHz ± 0.01% clock in the SLAVE timing mode. A typical transmitter output for transmitter test modes 2 and 3 is shown in Figure 40–20.

When test mode 4 is enabled, the PHY shall transmit the sequence of symbols generated by the following scrambler generator polynomial, bit generation, and level mappings:

$$g_{s1} = 1 + x^9 + x^{11}$$

The maximum-length shift register used to generate the sequences defined by this polynomial shall be updated once per symbol interval (8 ns). The bits stored in the shift register delay line at a particular time n are denoted by $Scr_n[10:0]$. At each symbol period the shift register is advanced by one bit and one new bit represented by $Scr_n[0]$ is generated. Bits $Scr_n[8]$ and $Scr_n[10]$ are exclusive OR'd together to generate the next $Scr_n[0]$ bit. The bit sequences, $x0_n$, $x1_n$, and $x2_n$, generated from combinations of the scrambler bits as shown in the following equations, shall be used to generate the quinary symbols, s_n, as shown in Table 40–8. The quinary symbol sequence shall be presented simultaneously to all transmitters. The transmitter shall time the transmitted symbols from a 125.00 MHz ± 0.01% clock in the MASTER timing mode. A typical transmitter output for transmitter test mode 4 is shown in Figure 40–21.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 202 of 379

$$x0_n = Scr_n[0]$$

$$x1_n = Scr_n[1] \wedge Scr_n[4]$$

$$x2_n = Scr_n[2] \wedge Scr_n[4]$$

Table 40–8—Transmitter test mode 4 symbol mapping

x2n	x1n	x0n	quinary symbol, s_n
0	0	0	0
0	0	1	1
0	1	0	2
0	1	1	–1
1	0	0	0
1	0	1	1
1	1	0	–2
1	1	1	–1

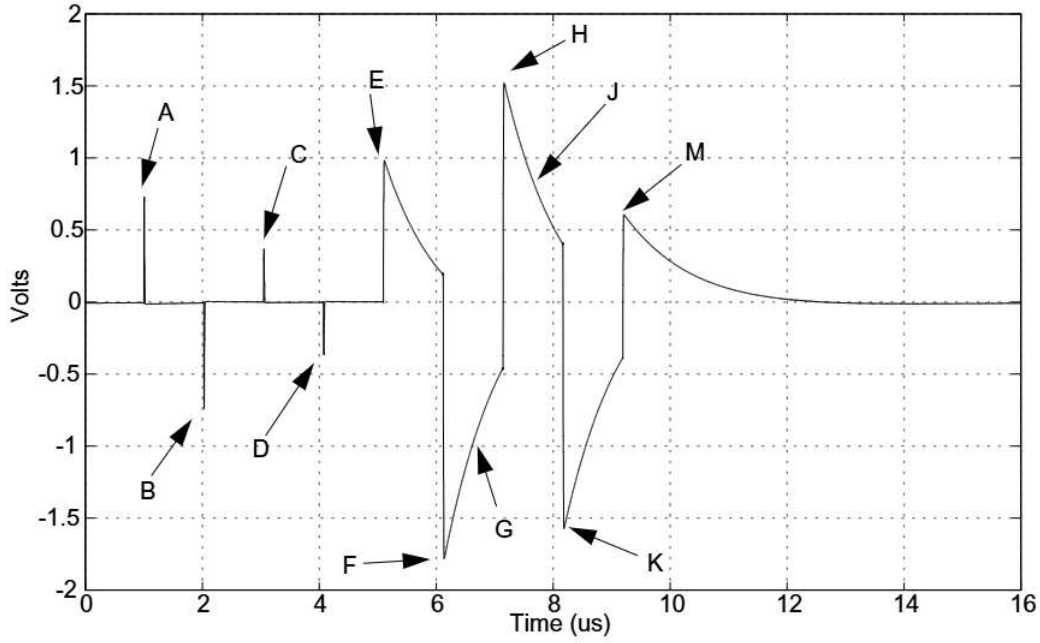


Figure 40–19—Example of transmitter test mode 1 waveform (1 cycle)

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 203 of 379

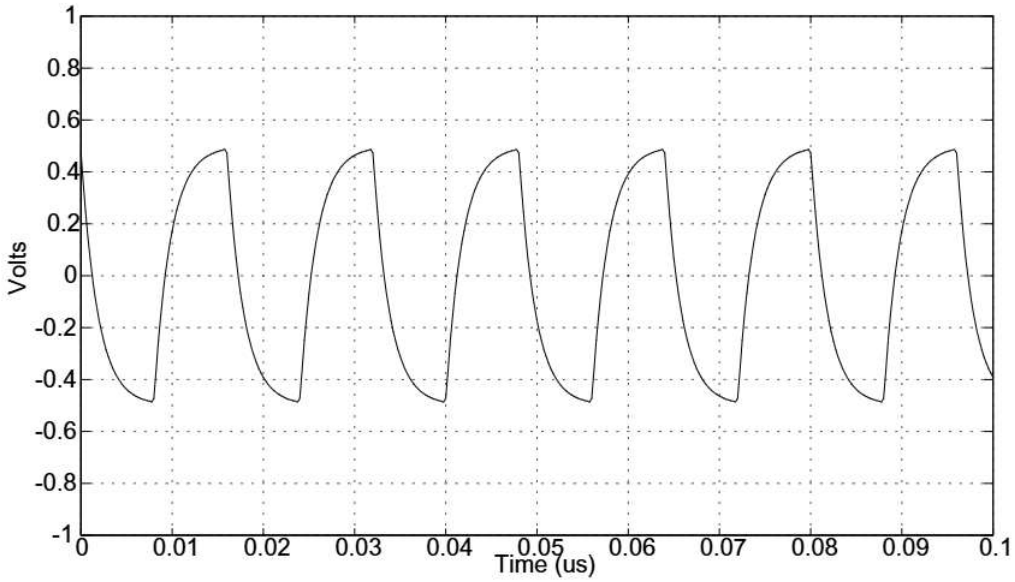


Figure 40–20—Example of transmitter test modes 2 and 3 waveform

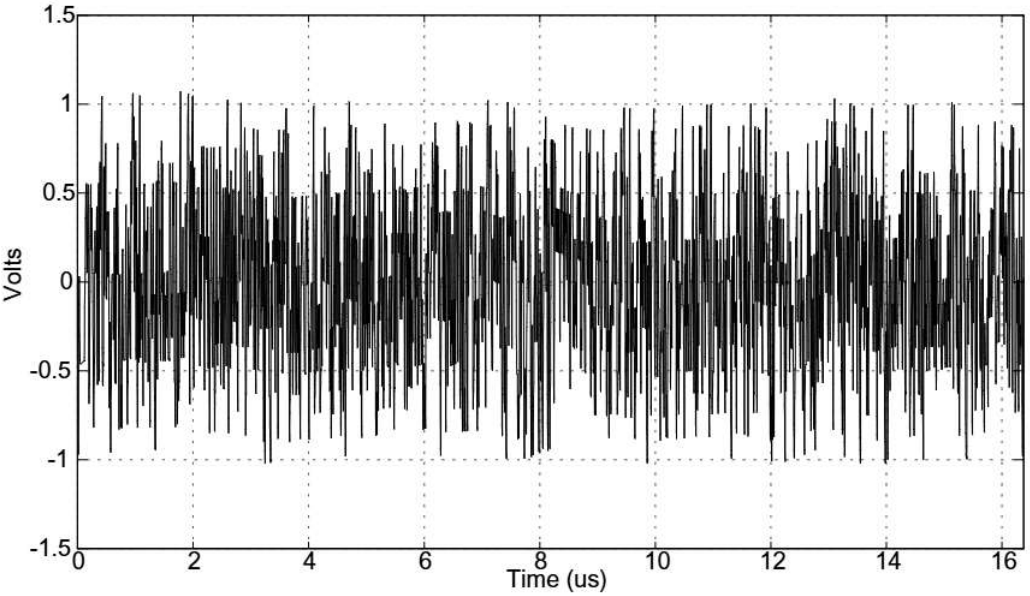


Figure 40–21—Example of Transmitter Test Mode 4 waveform (1 cycle)

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 204 of 379

40.6.1.1.3 Test Fixtures

The following fixtures (illustrated by Figure 40–22, Figure 40–23, Figure 40–24, and Figure 40–25), or their functional equivalents, shall be used for measuring the transmitter specifications described in 40.6.1.2.

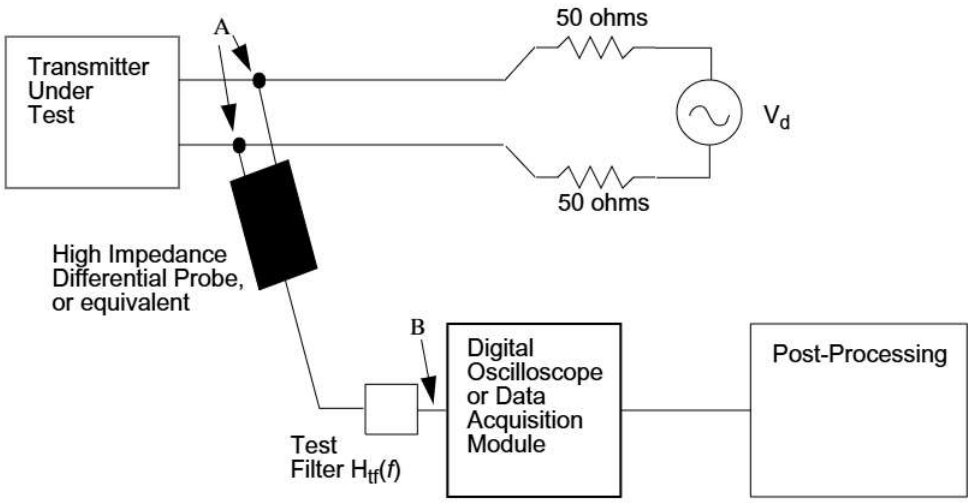


Figure 40–22—Transmitter test fixture 1 for template measurement

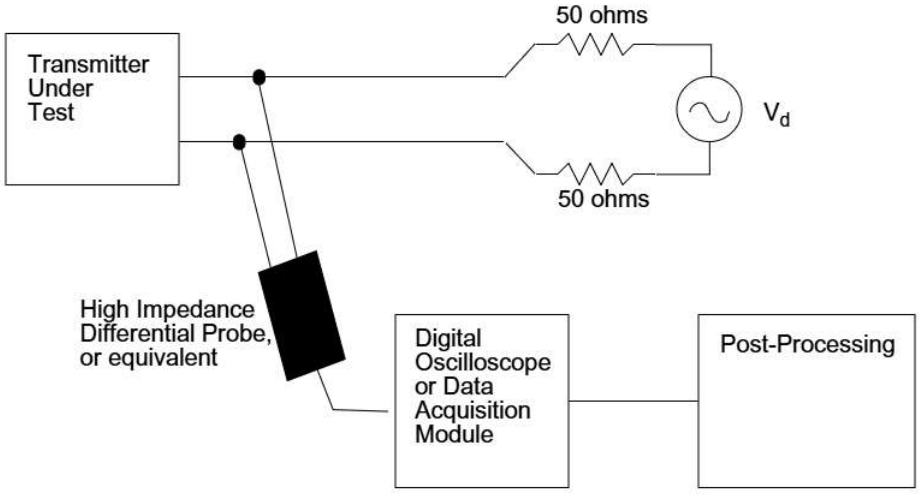


Figure 40–23—Transmitter test fixture 2 for droop measurement

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 205 of 379

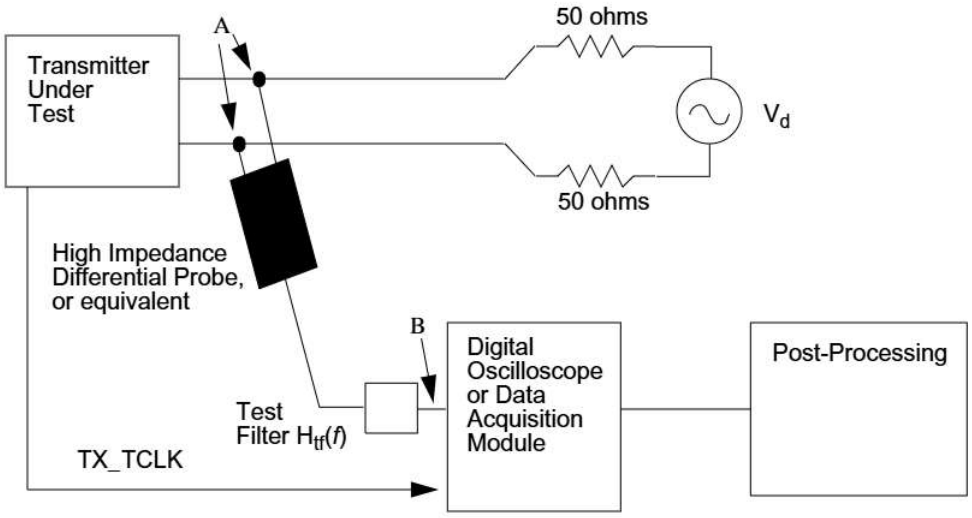


Figure 40–24—Transmitter test fixture 3 for distortion measurement

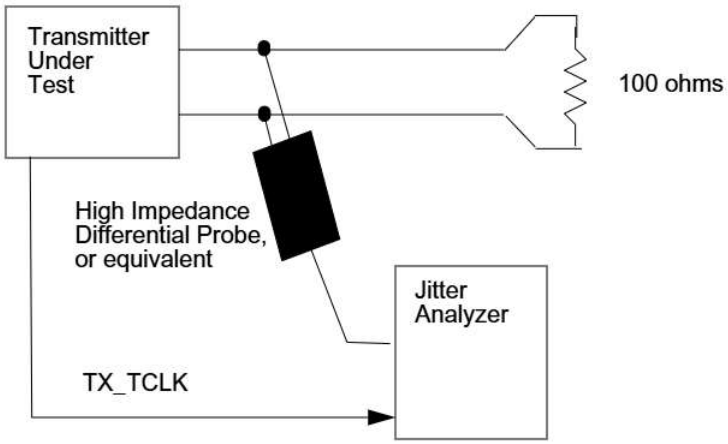


Figure 40–25—Transmitter test fixture 4 for transmitter jitter measurement

Copyright © 2002 IEEE. All rights reserved.

The test filter, $H_{tf}(f)$, used in transmitter test fixtures 1 and 3 may be located between the points A and B as long as the test filter does not significantly alter the impedance seen by the transmitter. The test filter may instead be implemented as a digital filter in the post processing block. The test filter shall have the following continuous time transfer function or its discrete time equivalent:

$$H_{tf}(f) = \frac{jf}{jf + 2\times10^{6}} \qquad f \text{ in Hz}$$

NOTE—j denotes the square root of –1.

The disturbing signal, V_d, shall have the characteristics listed in Table 40–9.

Table 40–9—V_d Characteristics

Characteristic	Transmit test fixture 1	Transmit test fixture 2	Transmit test fixture 3
Waveform	Sine wave		
Amplitude	2.8 volts peak-to-peak	2.8 volts peak-to-peak	5.4 volts peak-to-peak
Frequency	31.25 MHz	31.25 MHz	20.833 MHz (125/6 MHz)
Purity	All harmonics >40 dB below fundamental		

The post-processing block has two roles. The first is to remove the disturbing signal from the measurement. A method of removing the disturbing signal is to take a single shot acquisition of the transmitted signal plus test pattern, then remove the best fit of a sine wave at the fundamental frequency of the disturbing signal from the measurement. It will be necessary to allow the fitting algorithm to adjust the frequency, phase, and amplitude parameters of the sine wave to achieve the best fit.

The second role of the post-processing block is to compare the measured data with the templates, droop specification, or distortion specification.

Trigger averaging of the transmitter output to remove measurement noise and increase measurement resolution is acceptable provided it is done in a manner that does not average out possible distortions caused by the interaction of the transmitter and the disturbing voltage. For transmitter template and droop measurements, averaging can be done by ensuring the disturbing signal is exactly synchronous to the test pattern so that the phase of the disturbing signal at any particular point in the test pattern remains constant. Trigger averaging also requires a triggering event that is synchronous to the test pattern. A trigger pulse generated by the PHY would be ideal for this purpose; however, in practice, triggering off the waveform generated by one of the other transmitter outputs that does not have the disturbing signal present may be possible.

NOTE—The disturbing signal may be made synchronous to the test pattern by creating the disturbing signal using a source of the transmit clock for the PHY under test, dividing it down to the proper frequency for the disturbing signal, passing the result through a high Q bandpass filter to eliminate harmonics and then amplifying the result to the proper amplitude.

The generator of the disturbing signal must have sufficient linearity and range so it does not introduce any appreciable distortion when connected to the transmitter output (see Table 40–9). This may be verified by replacing the transmitter under test with another identical disturbing signal generator having a different frequency output and verifying that the resulting waveform's spectrum does not show significant distortion products.

Additionally, to allow for measurement of transmitted jitter in master and slave modes, the PHY shall provide access to the 125 MHz symbol clock, TX_TCLK, that times the transmitted symbols (see 40.4.2.2). The PHY shall provide a means to enable this clock output if it is not normally enabled.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 207 of 379

40.6.1.2 Transmitter electrical specifications

The PMA shall provide the Transmit function specified in 40.4.2.2 in accordance with the electrical specifications of this clause.

Where a load is not specified, the transmitter shall meet the requirements of this clause with a 100 Ω resistive differential load connected to each transmitter output.

The tolerance on the poles of the test filters used in this subclause shall be ±1%.

Practical considerations prevent measurement of the local transmitter performance in the presence of the remotely driven signal in this standard; however, the design of the transmitter to tolerate the presence of the remotely driven signal with acceptable distortion or other changes in performance is a critical issue and must be addressed by the implementor. To this end, a disturbing sine wave is used to simulate the presence of a remote transmitter for a number of the transmitter tests described in the following subordinate subclauses.

40.6.1.2.1 Peak differential output voltage and level accuracy

The absolute value of the peak of the waveform at points A and B, as defined in Figure 40–19, shall fall within the range of 0.67 V to 0.82 V (0.75 V ± 0.83 dB). These measurements are to be made for each pair while operating in test mode 1 and observing the differential signal output at the MDI using transmitter test fixture 1 with no intervening cable.

The absolute value of the peak of the waveforms at points A and B shall differ by less than 1%.

The absolute value of the peak of the waveform at points C and D as defined in Figure 40–19 shall differ by less than 2% from 0.5 times the average of the absolute values of the peaks of the waveform at points A and B.

40.6.1.2.2 Maximum output droop

The magnitude of the negative peak value of the waveform at point G, as defined in Figure 40–19, shall be greater than 73.1% of the magnitude of the negative peak value of the waveform at point F. These measurements are to be made for each pair while in test mode 1 and observing the differential signal output at the MDI using transmit test fixture 2 with no intervening cable. Point G is defined as the point exactly 500 ns after point F. Point F is defined as the point where the waveform reaches its minimum value at the location indicated in Figure 40–19. Additionally, the magnitude of the peak value of the waveform at point J as defined in Figure 40–19 shall be greater than 73.1% of the magnitude of the peak value of the waveform at point H. Point J is defined as the point exactly 500 ns after point H. Point H is defined as the point where the waveform reaches its maximum value at the location indicated in Figure 40–19.

40.6.1.2.3 Differential output templates

The voltage waveforms around points A, B, C, D defined in Figure 40–19, after the normalization described herein, shall lie within the time domain template 1 defined in Figure 40–26 and the piecewise linear interpolation between the points in Table 40–10. These measurements are to be made for each pair while in test mode 1 and while observing the differential signal output at the MDI using transmitter test fixture 1 with no intervening cable. The waveforms may be shifted in time as appropriate to fit within the template.

The waveform around point A is normalized by dividing by the peak value of the waveform at A.

The waveform around point B is normalized by dividing by the negative of the peak value of the waveform at A.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 208 of 379

The waveform around point C is normalized by dividing by 1/2 the peak value of the waveform at A.

The waveform around point D is normalized by dividing by the negative of 1/2 the peak value of the waveform at A.

The voltage waveforms around points F and H defined in Figure 40–19, after the normalization described herein, shall lie within the time domain template 2 defined in Figure 40–26 and the piecewise linear interpolation between the points in Table 40–11. These measurements are to be made for each pair while in test mode 1 and while observing the differential signal output at the MDI using transmitter test fixture 1 with no intervening cable. The waveforms may be shifted in time as appropriate to fit within the template.

The waveform around point F is normalized by dividing by the peak value of the waveform at F.

The waveform around point H is normalized by dividing by the peak value of the waveform at H.

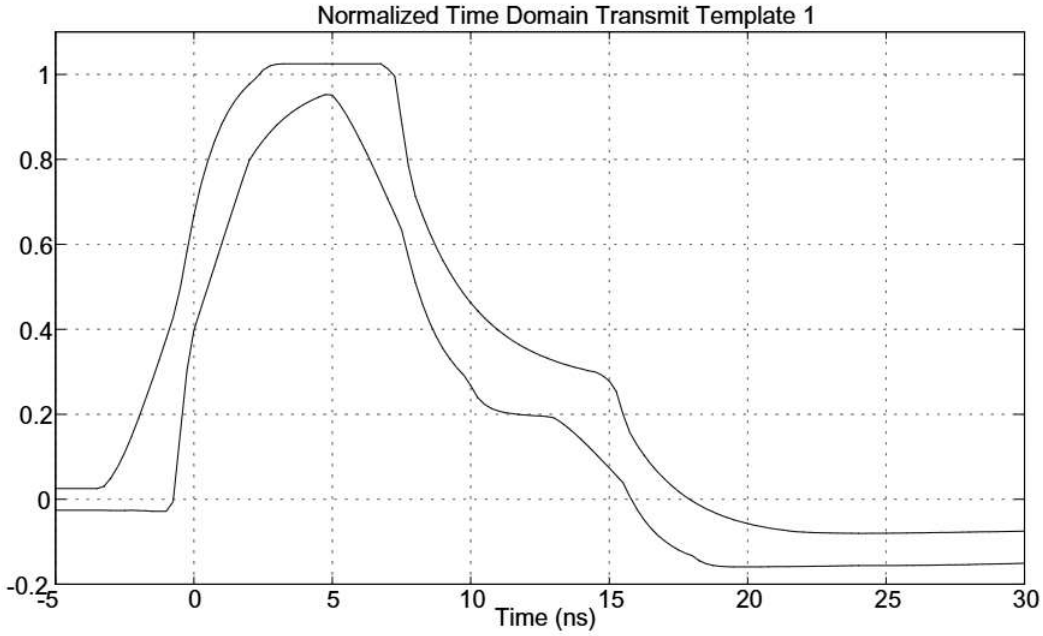
NOTE—The templates were created with the following assumptions about the elements in the transmit path:

1) Digital Filter: $0.75 + 0.25\ z^{-1}$
2) Ideal DAC
3) Single pole continuous time low pass filter with pole varying from 70.8 MHz to 117 MHz or linear rise/fall time of 5 ns.
4) Single pole continuous time high-pass filter (transformer high pass) with pole varying from 1 Hz to 100 kHz.
5) Single pole continuous time high-pass filter (test filter) with pole varying from 1.8 MHz to 2.2 MHz.
6) Additionally, +0.025 was added to the upper template and –0.025 was added to the lower template to allow for noise and measurement error.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 209 of 379

NOTE—The transmit templates are not intended to address electromagnetic radiation limits.



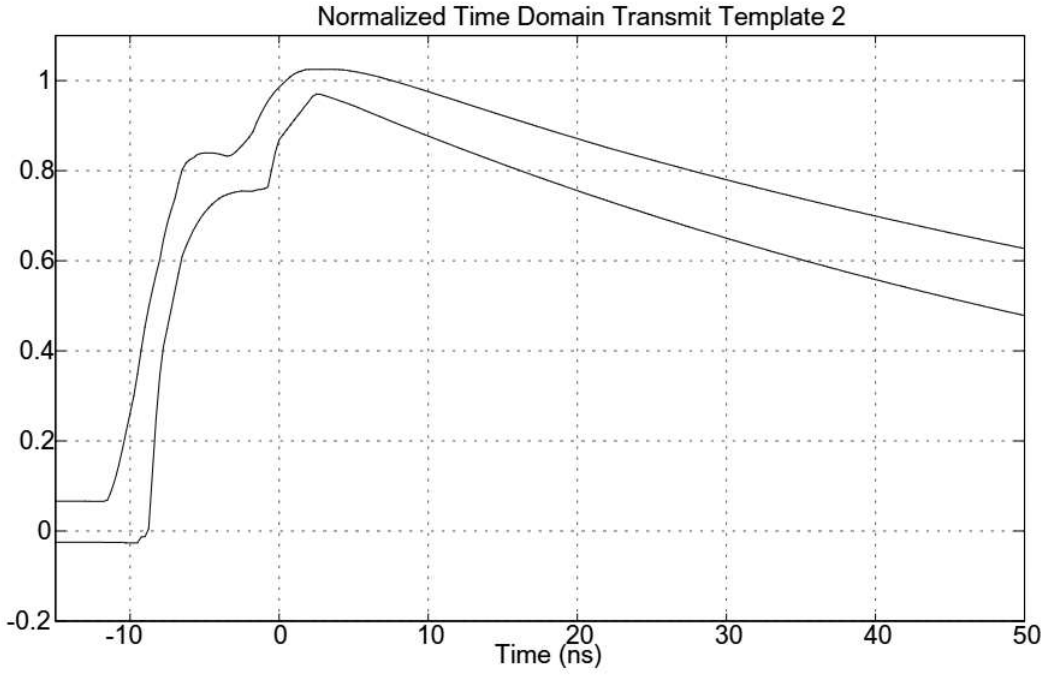


Figure 40–26—Normalized transmit templates as measured at MDI using transmit test fixture 1

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 210 of 379

NOTE—The ASCII for Tables 40–10 and 40–11 is available from http://www.ieee802.org/3/publication/index.html.[9]

Table 40–10—Normalized time domain voltage template 1

Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit	Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit
–5.00	0.025	–0.026	12.75	0.332	0.195
–4.75	0.025	–0.026	13.00	0.326	0.192
–4.50	0.025	–0.026	13.25	0.320	0.181
–4.25	0.025	–0.026	13.50	0.315	0.169
–4.00	0.025	–0.026	13.75	0.311	0.155
–3.75	0.025	–0.026	14.00	0.307	0.140
–3.50	0.025	–0.026	14.25	0.303	0.124
–3.25	0.031	–0.026	14.50	0.300	0.108
–3.00	0.050	–0.026	14.75	0.292	0.091
–2.75	0.077	–0.026	15.00	0.278	0.074
–2.50	0.110	–0.026	15.25	0.254	0.056
–2.25	0.148	–0.026	15.50	0.200	0.039
–2.00	0.190	–0.027	15.75	0.157	0.006
–1.75	0.235	–0.027	16.00	0.128	–0.023
–1.50	0.281	–0.028	16.25	0.104	–0.048
–1.25	0.329	–0.028	16.50	0.083	–0.068
–1.00	0.378	–0.028	16.75	0.064	–0.084
–0.75	0.427	–0.006	17.00	0.047	–0.098
–0.50	0.496	0.152	17.25	0.032	–0.110
–0.25	0.584	0.304	17.50	0.019	–0.119
0.00	0.669	0.398	17.75	0.007	–0.127
0.25	0.739	0.448	18.00	–0.004	–0.133
0.50	0.796	0.499	18.25	–0.014	–0.145
0.75	0.844	0.550	18.50	–0.022	–0.152
1.00	0.882	0.601	18.75	–0.030	–0.156
1.25	0.914	0.651	19.00	–0.037	–0.158
1.50	0.940	0.701	19.25	–0.043	–0.159

[9]*Copyright release for 802.3® template data:* Users of this standard may freely reproduce the template data in this subclause so it can be used for its intended purpose.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 211 of 379

Table 40–10—Normalized time domain voltage template 1 *(continued)*

Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit	Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit
1.75	0.960	0.751	19.50	–0.048	–0.159
2.00	0.977	0.797	19.75	–0.053	–0.159
2.25	0.992	0.822	20.00	–0.057	–0.159
2.50	1.010	0.845	20.25	–0.061	–0.159
2.75	1.020	0.864	20.50	–0.064	–0.159
3.00	1.024	0.881	20.75	–0.067	–0.159
3.25	1.025	0.896	21.00	–0.070	–0.159
3.50	1.025	0.909	21.25	–0.072	–0.159
3.75	1.025	0.921	21.50	–0.074	–0.158
4.00	1.025	0.931	21.75	–0.076	–0.158
4.25	1.025	0.939	22.00	–0.077	–0.158
4.50	1.025	0.946	22.25	–0.078	–0.158
4.75	1.025	0.953	22.50	–0.079	–0.158
5.00	1.025	0.951	22.75	–0.079	–0.157
5.25	1.025	0.931	23.00	–0.079	–0.157
5.50	1.025	0.905	23.25	–0.080	–0.157
5.75	1.025	0.877	23.50	–0.080	–0.157
6.00	1.025	0.846	23.75	–0.080	–0.156
6.25	1.025	0.813	24.00	–0.080	–0.156
6.50	1.025	0.779	24.25	–0.080	–0.156
6.75	1.025	0.743	24.50	–0.080	–0.156
7.00	1.014	0.707	24.75	–0.080	–0.156
7.25	0.996	0.671	25.00	–0.080	–0.156
7.50	0.888	0.634	25.25	–0.080	–0.156
7.75	0.784	0.570	25.50	–0.080	–0.156
8.00	0.714	0.510	25.75	–0.079	–0.156
8.25	0.669	0.460	26.00	–0.079	–0.156
8.50	0.629	0.418	26.25	–0.079	–0.156
8.75	0.593	0.383	26.50	–0.079	–0.155

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 212 of 379

Table 40–10—Normalized time domain voltage template 1 *(continued)*

Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit	Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit
9.00	0.561	0.354	26.75	–0.079	–0.155
9.25	0.533	0.330	27.00	–0.078	–0.155
9.50	0.507	0.309	27.25	–0.078	–0.155
9.75	0.483	0.292	27.50	–0.078	–0.154
10.00	0.462	0.268	27.75	–0.078	–0.154
10.25	0.443	0.239	28.00	–0.077	–0.154
10.50	0.427	0.223	28.25	–0.077	–0.153
10.75	0.411	0.213	28.50	–0.077	–0.153
11.00	0.398	0.208	28.75	–0.076	–0.153
11.25	0.385	0.204	29.00	–0.076	–0.152
11.50	0.374	0.201	29.25	–0.076	–0.152
11.75	0.364	0.199	29.50	–0.076	–0.152
12.00	0.355	0.198	29.75	–0.075	–0.151
12.25	0.346	0.197	30.00	–0.075	–0.151
12.50	0.339	0.196			

Table 40–11—Normalized time domain voltage template 2

Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit	Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit
–15.00	0.066	–0.025	18.00	0.891	0.779
–14.50	0.066	–0.025	18.50	0.886	0.773
–14.00	0.066	–0.025	19.00	0.881	0.767
–13.50	0.066	–0.025	19.50	0.876	0.762
–13.00	0.066	–0.025	20.00	0.871	0.756
–12.50	0.066	–0.025	20.50	0.866	0.750
–12.00	0.066	–0.025	21.00	0.861	0.745
–11.50	0.069	–0.025	21.50	0.856	0.739
–11.00	0.116	–0.025	22.00	0.852	0.734
–10.50	0.183	–0.025	22.50	0.847	0.728

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 213 of 379

Table 40–11—Normalized time domain voltage template 2 *(continued)*

Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit	Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit
–10.00	0.261	–0.027	23.00	0.842	0.723
–9.50	0.348	–0.027	23.50	0.838	0.717
–9.00	0.452	–0.013	24.00	0.833	0.712
–8.50	0.535	0.130	24.50	0.828	0.707
–8.00	0.604	0.347	25.00	0.824	0.701
–7.50	0.683	0.451	25.50	0.819	0.696
–7.00	0.737	0.531	26.00	0.815	0.691
–6.50	0.802	0.610	26.50	0.811	0.686
–6.00	0.825	0.651	27.00	0.806	0.680
–5.50	0.836	0.683	27.50	0.802	0.675
–5.00	0.839	0.707	28.00	0.797	0.670
–4.50	0.839	0.725	28.50	0.793	0.665
–4.00	0.837	0.739	29.00	0.789	0.660
–3.50	0.832	0.747	29.50	0.784	0.655
–3.00	0.839	0.752	30.00	0.780	0.650
–2.50	0.856	0.755	30.50	0.776	0.645
–2.00	0.875	0.755	31.00	0.772	0.641
–1.50	0.907	0.758	31.50	0.767	0.636
–1.00	0.941	0.760	32.00	0.763	0.631
–0.50	0.966	0.803	32.50	0.759	0.626
0.00	0.986	0.869	33.00	0.755	0.621
0.50	1.001	0.890	33.50	0.751	0.617
1.00	1.014	0.912	34.00	0.747	0.612
1.50	1.022	0.933	34.50	0.743	0.607
2.00	1.025	0.954	35.00	0.739	0.603
2.50	1.025	0.970	35.50	0.734	0.598
3.00	1.025	0.967	36.00	0.730	0.594
3.50	1.025	0.962	36.50	0.727	0.589
4.00	1.025	0.956	37.00	0.723	0.585

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 214 of 379

Table 40–11—Normalized time domain voltage template 2 *(continued)*

Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit	Time, ns	Normalized transmit time domain template, upper limit	Normalized transmit time domain template, lower limit
4.50	1.023	0.950	37.50	0.719	0.580
5.00	1.020	0.944	38.00	0.715	0.576
5.50	1.017	0.937	38.50	0.711	0.571
6.00	1.014	0.931	39.00	0.707	0.567
6.50	1.010	0.924	39.50	0.703	0.563
7.00	1.005	0.917	40.00	0.699	0.558
7.50	1.001	0.910	40.50	0.695	0.554
8.00	0.996	0.903	41.00	0.692	0.550
8.50	0.991	0.897	41.50	0.688	0.546
9.00	0.986	0.890	42.00	0.684	0.541
9.50	0.981	0.884	42.50	0.680	0.537
10.00	0.976	0.877	43.00	0.677	0.533
10.50	0.970	0.871	43.50	0.673	0.529
11.00	0.965	0.864	44.00	0.669	0.525
11.50	0.960	0.858	44.50	0.666	0.521
12.00	0.954	0.852	45.00	0.662	0.517
12.50	0.949	0.845	45.50	0.659	0.513
13.00	0.944	0.839	46.00	0.655	0.509
13.50	0.938	0.833	46.50	0.651	0.505
14.00	0.933	0.827	47.00	0.648	0.501
14.50	0.928	0.820	47.50	0.644	0.497
15.00	0.923	0.814	48.00	0.641	0.493
15.50	0.917	0.808	48.50	0.637	0.490
16.00	0.912	0.802	49.00	0.634	0.486
16.50	0.907	0.796	49.50	0.631	0.482
17.00	0.902	0.791	50.00	0.627	0.478
17.50	0.897	0.785			

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 215 of 379

40.6.1.2.4 Transmitter distortion

When in test mode 4 and observing the differential signal output at the MDI using transmitter test fixture 3, for each pair, with no intervening cable, the peak distortion as defined below shall be less than 10 mV.

The peak distortion is determined by sampling the differential signal output with the symbol rate TX_TCLK at an arbitrary phase and processing a block of any 2047 consecutive samples with the MATLAB (see 1.3) code listed below or equivalent. Note that this code assumes that the differential signal has already been filtered by the test filter.

NOTE—The ASCII for the following MATLAB code is available from http://www.ieee802.org/3/publication/index.html.[10]

MATLAB code for Distortion Post Processing is as follows:

```
%
% Distortion Specification Post Processing
%

% Initialize Variables
clear
symbolRate=125e6;                                % symbol rate
dataFile=input('Data file name: ','s')

% Generate test pattern symbol sequence

scramblerSequence=ones(1,2047);
for i=12:2047
  scramblerSequence(i)=mod(scramblerSequence(i-11) + scramblerSequence(i-9),2);
end

for i=1:2047
  temp=scramblerSequence(mod(i-1,2047)+1) + ...
       2*mod(scramblerSequence(mod(i-2,2047)+1) + scramblerSequence(mod(i-5,2047)+1),2)
+ ...
      4*mod(scramblerSequence(mod(i-3,2047)+1) + scramblerSequence(mod(i-5,2047)+1),2);
  switch temp
    case 0,
      testPattern(i)=0;
    case 1,
      testPattern(i)=1;
    case 2,
      testPattern(i)=2;
    case 3,
      testPattern(i)=-1;
    case 4,
      testPattern(i)=0;
    case 5,
      testPattern(i)=1;
    case 6,
      testPattern(i)=-2;
    case 7,
      testPattern(i)=-1;
  end
end



% Input data file
fid=fopen(dataFile,'r');
sampledData=fscanf(fid,'%f');
```

[10]*Copyright release for MATLAB code*: Users of this standard may freely reproduce the MATLAB code in this subclause so it can be used for its intended purpose.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 216 of 379

```
fclose(fid);
sampledData=sampledData.';


if (length(sampledData) < 2047)
  error('Must have 2047 consecutive samples for processing');
elseif (length(sampledData) > 2047)
  fprintf(1,'\n Warning - only using first 2047 samples in data file');
  sampledData=sampledData(1:2047);
end


% Fit a sine wave to the data and temporarily remove it to yield processed data

options=foptions;
options(1)=0;
options(2)=1e-8;
options(3)=1e-8;
options(14)=2000;
gradfun=zeros(0);
P=fmins('sinefit',[2.0 0 125/6.],options,gradfun,sampledData,symbolRate);

P

processedData=sampledData - ...
    P(1)*sin(2*pi*(P(3)*1e6*[0:2046]/symbolRate + P(2)*1e-9*symbolRate));



% LMS Canceller

numberCoeff=70;  % Number of coefficients in canceller
coefficients=zeros(1,numberCoeff);
delayLine=testPattern;

% Align data in delayLine to sampled data pattern
temp=xcorr(processedData,delayLine);
index=find(abs(temp)==max(abs(temp)));
index=mod(mod(length(processedData) - index(1),2047)+numberCoeff-10,2047);
delayLine=[delayLine((end-index):end) delayLine(1:(end-index-1))];


% Compute coefficients that minimize squared error in cyclic block

for i=1:2047
  X(i,:)=delayLine(mod([0:(numberCoeff-1)]+i-1,2047)+1);
end
coefficients=(inv(X.' * X)*(processedData*X).').';



% Canceller
for i=1:2047
   err(i)=processedData(i) - sum(delayLine(1+mod((i-1):(i+numberCoeff-2),2047)).*coef-
ficients);
end

% Add back temporarily removed sine wave

err=err+P(1)*sin(2*pi*(P(3)*1e6*[0:2046]./symbolRate + P(2)*1e-9*symbolRate));


% Re-fit sine wave and do a final removal
```

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 217 of 379

```
options=foptions;
options(1)=0;
options(2)=1e-12;
options(3)=1e-12;
options(14)=10000;
gradfun=zeros(0);
P=fmins('sinefit',[2.0 0 125/6.],options,gradfun,err,symbolRate);

P

processedData=sampledData - ...
    P(1)*sin(2*pi*(P(3)*1e6*[0:2046]/symbolRate + P(2)*1e-9*symbolRate));


% Compute coefficients that minimize squared error in cyclic block

coefficients=(inv(X.' * X)*(processedData*X).').';



% Canceller
for i=1:2047
   err(i)=processedData(i) - sum(delayLine(1+mod((i-1):(i+numberCoeff-2),2047)).*coef-
ficients);
end

% SNR Calculation
signal=0.5;
noise=mean(err.^2);

SNR=10*log10(signal./noise);

% Output Peak Distortion

peakDistortion=max(abs(err))


% Function for fitting sine wave
function err=sinefit(parameters,data,symbolRate)
err=sum((data- ...
    parameters(1)*sin(2*pi*(parameters(3)*1e6*[0:(length(data)-1)]/symbolRate + param-
eters(2)*1e-9*symbolRate))).^2);
```

40.6.1.2.5 Transmitter timing jitter

When in test mode 2 or test mode 3, the peak-to-peak jitter J_{txout} of the zero crossings of the differential signal output at the MDI relative to the corresponding edge of TX_TCLK is measured. The corresponding edge of TX_TCLK is the edge of the transmit test clock, in polarity and time, that generates the zero-crossing transition being measured.

When in the normal mode of operation as the MASTER, the peak-to-peak value of the MASTER TX_TCLK jitter relative to an unjittered reference shall be less than 1.4 ns. When the jitter waveform on TX_TCLK is filtered by a high-pass filter, $H_{jf1}(f)$, having the transfer function below, the peak-to-peak value of the resulting filtered timing jitter plus J_{txout} shall be less than 0.3 ns.

$$H_{jf1}(f) = \frac{\mathrm{j}f}{\mathrm{j}f + 5000} \qquad f \text{ in Hz}$$

When in the normal mode of operation as the SLAVE, receiving valid signals from a compliant PHY operating as the MASTER using the test channel defined in 40.6.1.1.1, with test channel port A connected to the SLAVE, the peak-to-peak value of the SLAVE TX_TCLK jitter relative to the MASTER TX_TCLK shall be less than 1.4 ns after the receiver is properly receiving the data and has set bit 10.13 of the GMII management register set to 1. When the jitter waveform on TX_TCLK is filtered by a high-pass filter, $H_{jf2}(f)$, having

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 218 of 379

the transfer function below, the peak-to-peak value of the resulting filtered timing jitter plus J_{txout} shall be no more than 0.4 ns greater than the simultaneously measured peak-to-peak value of the MASTER jitter filtered by $H_{jf1}(f)$.

$$H_{jf2}(f) = \frac{jf}{jf + 32000} \qquad f \text{ in Hz}$$

NOTE—j denotes the square root of –1.

For all high-pass filtered jitter measurements, the peak-to-peak value shall be measured over an unbiased sample of at least 10^5 clock edges. For all unfiltered jitter measurements, the peak-to-peak value shall be measured over an interval of not less than 100 ms and not more than 1 second.

40.6.1.2.6 Transmit clock frequency

The quinary symbol transmission rate on each pair of the master PHY shall be 125.00 MHz ± 0.01%.

40.6.1.3 Receiver electrical specifications

The PMA shall provide the Receive function specified in 40.4.2.3 in accordance with the electrical specifications of this clause. The patch cabling and interconnecting hardware used in test configurations shall be within the limits specified in 40.7.

40.6.1.3.1 Receiver differential input signals

Differential signals received at the MDI that were transmitted from a remote transmitter within the specifications of 40.6.1.2 and have passed through a link specified in 40.7 are translated into one of the PMA_UNITDATA.indicate messages with a 4-D symbol error rate less than 10^{-10} and sent to the PCS after link reset completion. Since the 4-D symbols are not accessible, this specification shall be satisfied by a frame error rate less than 10^{-7} for 125 octet frames.

40.6.1.3.2 Receiver frequency tolerance

The receive feature shall properly receive incoming data with a 5-level symbol rate within the range 125.00 MHz ± 0.01%.

40.6.1.3.3 Common-mode noise rejection

This specification is provided to limit the sensitivity of the PMA receiver to common-mode noise from the cabling system. Common-mode noise generally results when the cabling system is subjected to electromagnetic fields. Figure 40–27 shows the test configuration, which uses a capacitive cable clamp, that injects common-mode signals into a cabling system.

A 100-meter, 4-pair Category 5 cable that meets the specification of 40.7 is connected between two 1000BASE-T PHYs and inserted into the cable clamp. The cable should be terminated on each end with an MDI connector plug specified in 40.8.1. The clamp should be located a distance of ~20 cm from the receiver. It is recommended that the cable between the transmitter and the cable clamp be installed either in a linear run or wrapped randomly on a cable rack. The cable rack should be at least 3 m from the cable clamp. In addition, the cable clamp and 1000BASE-T receiver should be placed on a common copper ground plane and the ground of the receiver should be in contact with the ground plane. The chassis grounds of all test equipment used should be connected to the copper ground plane. No connection is required between the copper ground plane and an external reference. A description of the cable clamp, as well as the validation procedure, can be found in Annex 40B.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 219 of 379

A signal generator with a 50 Ω impedance is connected to one end of the clamp and an oscilloscope with a 50 Ω input is connected to the other end of the clamp. The signal generator shall be capable of providing a sine wave signal of 1 MHz to 250 MHz. The output of the signal generator is adjusted for a voltage of 1.0 Vrms (1.414 Vpeak) on the oscilloscope.

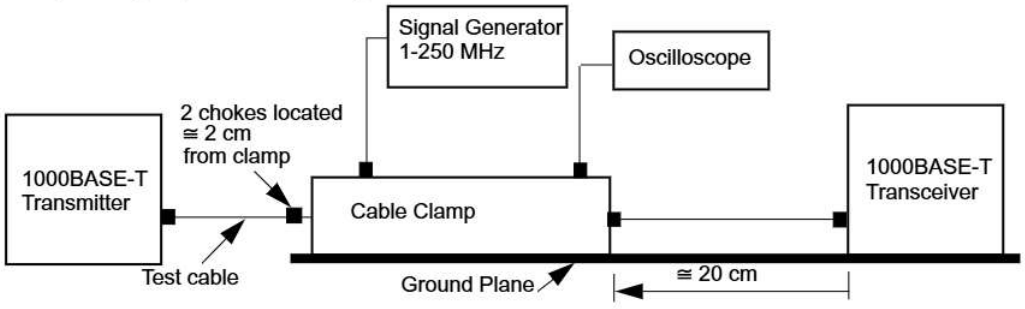


Figure 40–27—Receiver common-mode noise rejection test

While sending data from the transmitter, the receiver shall send the proper PMA_UNITDATA.indicate messages to the PCS as the signal generator frequency is varied from 1 MHz to 250 MHz.

NOTE—Although the signal specification is constrained within the 1–100 MHz band, this test is performed up to 250 MHz to ensure the receiver under test can tolerate out-of-band (100–250 MHz) noise.

40.6.1.3.4 Alien Crosstalk noise rejection

While receiving data from a transmitter specified in 40.6.1.2 through a link segment specified in 40.7 connected to all MDI duplex channels, a receiver shall send the proper PMA_UNITDATA.indicate message to the PCS when any one of the four pairs is connected to a noise source as described in Figure 40–28. Because symbol encoding is employed, this specification shall be satisfied by a frame error rate of less than 10^{-7} for 125 octet frames. The level of the noise signal at the MDI is nominally 25 mV peak-to-peak. (Measurements are to be made on each of the four pairs.) The noise source shall be connected to one of the MDI inputs using Category 5 balanced cable of a maximum length of 0.5 m.

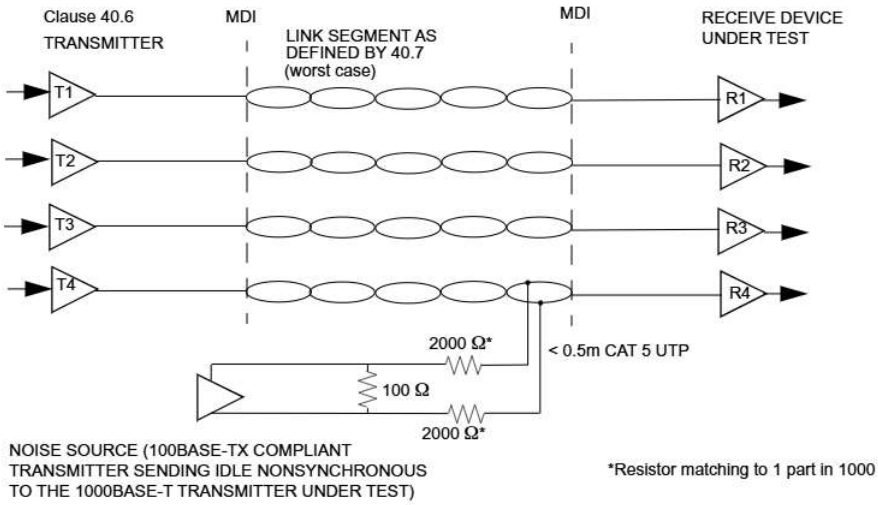


Figure 40–28—Differential mode noise rejection test

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 220 of 379

40.7 Link segment characteristics

1000BASE-T is designed to operate over a 4-pair Category 5 balanced cabling system. Each of the four pairs supports an effective data rate of 250 Mbps in each direction simultaneously. The term "link segment" used in this clause refers to four duplex channels. The term "duplex channel" will be used to refer to a single channel with full duplex capability. Specifications for a link segment apply equally to each of the four duplex channels. All implementations of the balanced cabling link shall be compatible at the MDI.

40.7.1 Cabling system characteristics

The cabling system used to support 1000BASE-T requires 4 pairs of Category 5 balanced cabling with a nominal impedance of 100 Ω. The cabling system components (cables, cords, and connectors) used to provide the link segment shall consist of Category 5 components as specified in ANSI/TIA/EIA-568-A:1995 and ISO/IEC 11801:1995. Additionally:

a) 1000BASE-T uses a star topology with Category 5 balanced cabling used to connect PHY entities.
b) 1000BASE-T is an ISO/IEC 11801 Class D application, with additional installation requirements and transmission parameters specified in Annex 40A.
c) The width of the PMD transmit signal spectrum is approximately 80 MHz.
d) The use of shielding is outside the scope of this standard.

40.7.2 Link transmission parameters

The transmission parameters contained in this subclause are specified to ensure that a Category 5 link segment of up to at least 100 m will provide a reliable medium. The transmission parameters of the link segment include insertion loss, delay parameters, characteristic impedance, NEXT loss, ELFEXT loss, and return loss.

Link segment testing shall be conducted using source and load impedances of 100 Ω. The tolerance on the poles of the test filter used in this subclause shall be no worse than 1%.

40.7.2.1 Insertion loss

The insertion loss of each duplex channel shall be less than

$$\text{Insertion\_Loss}(f) < 2.1\ f^{0.529} + 0.4/f \quad \text{(dB)}$$

at all frequencies from 1 MHz to 100 MHz. This includes the attenuation of the balanced cabling pairs, including work area and equipment cables plus connector losses within each duplex channel. The insertion loss specification shall be met when the duplex channel is terminated in 100 Ω.

NOTE—The above equation approximates the insertion loss specification at discrete frequencies for Category 5 100-meter links specified in ANSI/TIA/EIA-568-A Annex E and in TIA/EIA TSB 67.

40.7.2.2 Differential characteristic impedance

The nominal differential characteristic impedance of each link segment duplex channel, which includes cable cords and connecting hardware, is 100 Ω for all frequencies between 1 MHz and 100 MHz.

40.7.2.3 Return loss

Each link segment duplex channel shall meet or exceed the return loss specified in the following equation at all frequencies from 1 MHz to 100 MHz

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 221 of 379

$$\text{Return\_Loss}(f) \begin{cases} 15 & (1 - 20\ \text{MHz}) \\ 15 - 10\log_{10}(f/20) & (20 - 100\ \text{MHz}) \end{cases} \text{(dB)}$$

where f is the frequency in MHz. The reference impedance shall be 100 Ω.

40.7.3 Coupling parameters

In order to limit the noise coupled into a duplex channel from adjacent duplex channels, Near-End Crosstalk (NEXT) loss and Equal Level Far-End Crosstalk (ELFEXT) loss are specified for each link segment. Each duplex channel can be disturbed by more than one duplex channel. Requirements for Multiple Disturber Near-End Crosstalk (MDNEXT) are satisfied even when worst case conditions of differential pair-to-pair NEXT as specified under 40.7.3.1.1 occur. Therefore, there are no separate requirements for MDNEXT. Requirements for Multiple Disturber Equal-Level Far-End Crosstalk (MDELFEXT) loss are specified in 40.7.3.2.2.

40.7.3.1 Near-End Crosstalk (NEXT)

40.7.3.1.1 Differential Near-End Crosstalk

In order to limit the crosstalk at the near end of a link segment, the differential pair-to-pair Near-End Crosstalk (NEXT) loss between a duplex channel and the other three duplex channels is specified to meet the symbol error rate objective specified in 40.1. The NEXT loss between any two duplex channels of a link segment shall be at least

$$27.1 - 16.8\log_{10}(f/100)$$

where f is the frequency over the range of 1 MHz to 100 MHz.

NOTE—The above equation approximates the NEXT loss specification at discrete frequencies for Category 5 100-meter links specified in ANSI/TIA/EIA-568-A Annex E and in TSB-67.

40.7.3.2 Far-End Crosstalk (FEXT)

40.7.3.2.1 Equal Level Far-End Crosstalk (ELFEXT) loss

Equal Level Far-End Crosstalk (ELFEXT) loss is specified in order to limit the crosstalk at the far end of each link segment duplex channel and meet the BER objective specified in 40.6.1.3.1. Far-End Crosstalk (FEXT) is crosstalk that appears at the far end of a duplex channel (disturbed channel), which is coupled from another duplex channel (disturbing channel) with the noise source (transmitters) at the near end. FEXT loss is defined as

$$\text{FEXT\_Loss}(f) = 20\log_{10}[V\text{pds}(f)/V\text{pcn}(f)]$$

and ELFEXT_Loss is defined as

$$\text{ELFEXT\_Loss}(f) = 20\log_{10}[V\text{pds}(f)/V\text{pcn}(f)] - \text{SLS\_Loss}(f)$$

where

*V*pds	is the peak voltage of disturbing signal (near-end transmitter)
*V*pcn	is the peak crosstalk noise at far end of disturbed channel
SLS_Loss	is the insertion loss of disturbed channel in dB

The worst pair ELFEXT loss between any two duplex channels shall be greater than $17 - 20\log_{10}(f/100)$ dB where f is the frequency over the range of 1 MHz to 100 MHz.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 222 of 379

40.7.3.2.2 Multiple Disturber Equal Level Far-End Crosstalk (MDELFEXT) loss

Since four duplex channels are used to transfer data between PMDs, the FEXT that is coupled into a data carrying channel will be from the three adjacent disturbing duplex channels. This specification is consistent with three channel-to-channel disturbers—one with a ELFEXT loss of at least $17 - 20\log_{10}(f/100)$ dB, one with a ELFEXT loss of at least $19.5 - 20\log_{10}(f/100)$ dB, and one with a ELFEXT loss of at least $23 - 20\log_{10}(f/100)$ dB. To ensure the total FEXT coupled into a duplex channel is limited, multiple disturber ELFEXT loss is specified as the power sum of the individual ELFEXT losses.

The Power Sum loss between a duplex channel and the three adjacent disturbers shall be

$$\text{PSELFEXT loss} > 14.4 - 20\log_{10}(f/100)\ \text{dB}$$

where f is the frequency over the range of 1 MHz to 100 MHz.

40.7.3.2.3 Multiple-Disturber Power Sum Equal Level Far-End Crosstalk (PSELFEXT) loss

PSELFEXT loss is determined by summing the magnitude of the three individual pair-to-pair differential ELFEXT loss values over the frequency range 1 to 100 MHz as follows:

$$\text{PSELFEXT\_Loss(f)} = -10\log_{10}\sum_{i=1}^{i=3} 10^{-(NL(f)i)/10}$$

where

NL(f)i is the magnitude of ELFEXT loss at frequency f of pair combination i
i is the 1, 2, or 3 (pair-to-pair combination)

40.7.4 Delay

In order to simultaneously send data over four duplex channels in parallel, the propagation delay of each duplex channel as well as the difference in delay between any two of the four channels are specified. This ensures the 1000 Mbps data that is divided across four channels can be properly reassembled at the far-end receiver. This also ensures the round-trip delay requirement for effective collision detection is met.

40.7.4.1 Maximum link delay

The propagation delay of a link segment shall not exceed 570 ns at all frequencies between 2 MHz and 100 MHz.

40.7.4.2 Link delay skew

The difference in propagation delay, or skew, between all duplex channel pair combinations of a link segment, under all conditions, shall not exceed 50 ns at all frequencies from 2 MHz to 100 MHz. It is a further functional requirement that, once installed, the skew between any two of the four duplex channels due to environmental conditions shall not vary more than 10 ns within the above requirement.

40.7.5 Noise environment

The 1000BASE-T noise environment consists of noise from many sources. The primary noise sources that impact the objective BER are NEXT and echo interference, which are reduced to a small residual noise using cancelers. The remaining noise sources, which are secondary sources, are discussed in the following list.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 223 of 379

The 1000BASE-T noise environment consists of the following:

a) Echo from the local transmitter on the same duplex channel (cable pair). Echo is caused by the hybrid function used to achieve simultaneous bi-directional transmission of data and by impedance discontinuities in the link segment. It is impractical to achieve the objective BER without using echo cancellation. Since the symbols transmitted by the local disturbing transmitter are available to the cancellation processor, echo interference can be reduced to a small residual noise using echo cancellation methods.

b) Near-End Crosstalk (NEXT) interference from the local transmitters on the duplex channels (cable pairs) of the link segment. Each receiver will experience NEXT interference from three adjacent transmitters. NEXT cancelers are used to reduce the interference from each of the three disturbing transmitters to a small residual noise. NEXT cancellation is possible since the symbols transmitted by the three disturbing local transmitters are available to the cancellation processor. NEXT cancelers can reduce NEXT interference by at least 20 dB.

c) Far-End Crosstalk (FEXT) noise at a receiver is from three disturbing transmitters at the far end of the duplex channel (cable pairs) of the link segment. FEXT noise can be cancelled in the same way as echo and NEXT interference although the symbols from the remote transmitters are not immediately available. However, FEXT noise is much smaller than NEXT interference and can generally be tolerated.

d) Inter-Symbol Interference (ISI) noise. ISI is the extraneous energy from one signaling symbol that interferes with the reception of another symbol on the same channel.

e) Noise from non-idealities in the duplex channel, transmitters, and receivers; for example, DAC/ADC non-linearity, electrical noise (shot and thermal), and non-linear channel characteristics.

f) Noise from sources outside the cabling that couple into the link segment via electric and magnetic fields.

g) Noise from signals in adjacent cables. This noise is referred to as alien NEXT noise and is generally present when cables are bound tightly together. Since the transmitted symbols from the alien NEXT noise source are not available to the cancellation processor (they are in another cable), it is not possible to cancel the alien NEXT noise. To ensure robust operation the alien NEXT noise must meet the specification of 40.7.5.1.

40.7.6 External coupled noise

The noise coupled from external sources that is measured at the output of a filter connected to the output of the near end of a disturbed duplex channel should not exceed 40 mV peak-to-peak. The filter for this measurement is a fifth order Butterworth filter with a 3 dB cutoff at 100MHz.

40.8 MDI specification

This subclause defines the MDI. The link topology requires a crossover function in a DTE-to-DTE connection. See 40.4.4 for a description of the automatic MDI/MDI-X configuration.

40.8.1 MDI connectors

Eight-pin connectors meeting the requirements of subclause 3 and Figures 1 through 4 of IEC 60603-7: 1990 shall be used as the mechanical interface to the balanced cabling. The plug connector shall be used on the balanced cabling and the jack on the PHY. These connectors are depicted (for informational use only) in Figure 40–29 and Figure 40–30. The assignment of PMA signals to connector contacts for PHYs is shown in Table 40-12.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 224 of 379

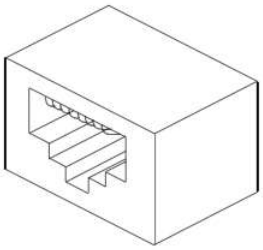


Figure 40–29—MDI connector

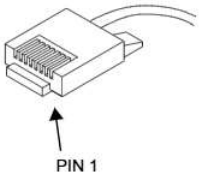


Figure 40–30—Balanced cabling connector

Table 40–12—Assignment of PMA signal to MDI and MDI-X pin-outs

Contact	MDI	MDI-X
1	BI_DA+	BI_DB+
2	BI_DA–	BI_DB–
3	BI_DB+	BI_DA+
4	BI_DC+	BI_DD+
5	BI_DC–	BI_DD–
6	BI_DB–	BI_DA–
7	BI_DD+	BI_DC+
8	BI_DD–	BI_DC–

40.8.2 Crossover function

Although the automatic MDI/MDI-X configuration (see 40.4.4) is not required for successful operation of 1000BASE-T, it is a functional requirement that a crossover function be implemented in every link segment to support the operation of Auto-Negotiation. The crossover function connects the transmitters of one PHY to the receivers of the PHY at the other end of the link segment. Crossover functions may be implemented internally to a PHY or else-where in the link segment. For a PHY that does not implement the crossover function, the MDI labels in the middle column of Table 40–12 refer to its own internal circuits. For PHYs that do implement the internal crossover, the MDI labels in the last column of Table 40–12 refer to the internal circuits of the remote PHY of the link segment. Additionally, the MDI connector for a PHY that implements the crossover function shall be marked with the graphical symbol X. The crossover function specified here is not compatible with the crossover function specified in 14.5.2 for pairs TD and RD.

When a link segment connects a single-port device to a multiport device, it is recommended that the crossover be implemented in the PHY local to the multiport device. If neither or both PHYs of a link segment contain internal crossover functions, an additional external crossover is necessary. It is recommended that the crossover be visible to an installer from one of the PHYs. When both PHYs contain internal crossovers,

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 225 of 379

it is further recommended that, in networks in which the topology identifies either a central backbone segment or a central device, the PHY furthest from the central element be assigned the external crossover to maintain consistency.

Implicit implementation of the crossover function within a twisted-pair cable or at a wiring panel, while not expressly forbidden, is beyond the scope of this standard.

40.8.3 MDI electrical specifications

The MDI connector (jack) when mated with a specified balanced cabling connector (plug) shall meet the electrical requirements for Category 5 connecting hardware for use with 100-ohm Category 5 cable as specified in ANSI/TIA/EIA-568-A:1995 and ISO/IEC 11801:1995.

The mated MDI/balanced cabling connector pair shall have a FEXT loss not less than $40 - 20\log_{10}(f/100)$ (where f is the frequency over the range 1 MHz to 100 MHz) between all contact pair combinations shown in Table 40–12.

No spurious signals shall be emitted onto the MDI when the PHY is held in power-down mode (as defined in 22.2.4.1.5) independent of the value of TX_EN, when released from power-down mode, or when external power is first applied to the PHY.

40.8.3.1 MDI return loss

The differential impedance at the MDI for each transmit/receive channel shall be such that any reflection due to differential signals incident upon the MDI from a balanced cabling having an impedance of 100 Ω ± 15% is attenuated, relative to the incident signal, at least 16 dB over the frequency range of 1.0 MHz to 40 MHz and at least $10 - 20\log_{10}(f/80)$ dB over the frequency range 40 MHz to 100 MHz (f in MHz). This return loss shall be maintained at all times when the PHY is transmitting data or control symbols.

40.8.3.2 MDI impedance balance

Impedance balance is a measurement of the impedance-to-ground difference between the two MDI contacts used by a duplex link channel and is referred to as common-mode-to-differential-mode impedance balance. Over the frequency range 1.0 MHz to 100.0 MHz, the common-mode-to-differential-mode impedance balance of each channel of the MDI shall exceed

$$34 - 19.2\log_{10}\left(\frac{f}{50}\right) \quad \text{dB}$$

where f is the frequency in MHz when the transmitter is transmitting random or pseudo random data. Test-mode 4 may be used to generate an appropriate transmitter output.

The balance is defined as

$$20\log_{10}\left(\frac{E_{cm}}{E_{dif}}\right)$$

where E_{cm} is an externally applied sine wave voltage as shown in Figure 40–31 and E_{dif} is the resulting waveform due only to the applied sine wave and not the transmitted data.

NOTES

1—Triggered averaging can be used to separate the component due to the applied common-mode sine wave from the transmitted data component.

2—The imbalance of the test equipment (such as the matching of the test resistors) must be insignificant relative to the balance requirements.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 226 of 379

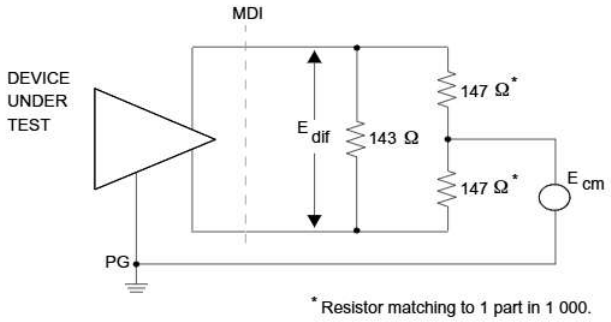


Figure 40–31—MDI impedance balance test circuit

40.8.3.3 MDI common-mode output voltage

The magnitude of the total common-mode output voltage, E_{cm_out}, on any transmit circuit, when measured as shown in Figure 40–32, shall be less than 50 mV peak-to-peak when transmitting data.

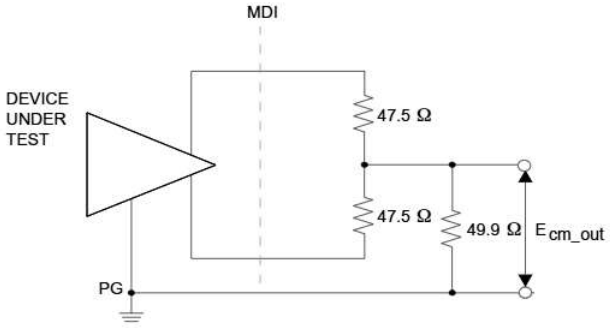


Figure 40–32—Common-mode output voltage test circuit

NOTE—The imbalance of the test equipment (such as the matching of the test resistors) must be insignificant relative to the balance requirements.

40.8.3.4 MDI fault tolerance

Each wire pair of the MDI shall, under all operating conditions, withstand without damage the application of short circuits of any wire to any other wire within the 4-pair cable for an indefinite period of time and shall resume normal operation after the short circuit(s) are removed. The magnitude of the current through such a short circuit shall not exceed 300 mA.

Each wire pair shall withstand without damage a 1000 V common-mode impulse applied at E_{cm} of either polarity (as indicated in Figure 40–33). The shape of the impulse shall be 0.3/50 μs (300 ns virtual front time, 50 μs virtual time of half value), as defined in IEC 60060.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 227 of 379

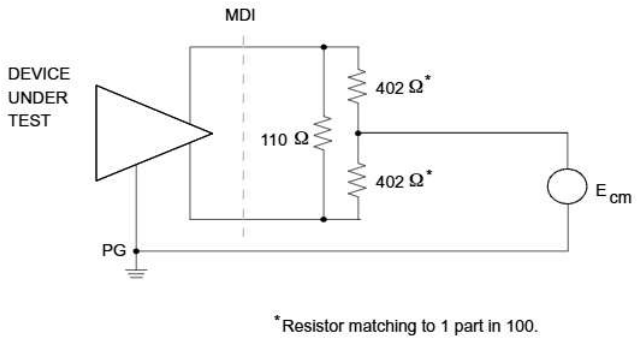


Figure 40–33—MDI fault tolerance test circuit

40.9 Environmental specifications

40.9.1 General safety

All equipment meeting this standard shall conform to IEC 60950: 1991.

40.9.2 Network safety

This subclause sets forth a number of recommendations and guidelines related to safety concerns; the list is neither complete nor does it address all possible safety issues. The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

LAN cabling systems described in this subclause are subject to at least four direct electrical safety hazards during their installation and use. These hazards are as follows:

a) Direct contact between LAN components and power, lighting, or communications circuits.
b) Static charge buildup on LAN cabling and components.
c) High-energy transients coupled onto the LAN cabling system.
d) Voltage potential differences between safety grounds to which various LAN components are connected.

Such electrical safety hazards must be avoided or appropriately protected against for proper network installation and performance. In addition to provisions for proper handling of these conditions in an operational system, special measures must be taken to ensure that the intended safety features are not negated during installation of a new network or during modification or maintenance of an existing network.

40.9.2.1 Installation

It is a mandatory requirement that sound installation practice, as defined by applicable local codes and regulations, is followed in every instance in which such practice is applicable.

40.9.2.2 Installation and maintenance guidelines

It is a mandatory requirement that, during installation and maintenance of the cabling plant, care is taken to ensure that non-insulated network cabling conductors do not make electrical contact with unintended conductors or ground.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 228 of 379

40.9.2.3 Telephony voltages

The use of building wiring brings with it the possibility of wiring errors that may connect telephony voltages to 1000BASE-T equipment. Other than voice signals (which are low voltage), the primary voltages that may be encountered are the "battery" and ringing voltages. Although there is no universal standard, the following maximums generally apply:

Battery voltage to a telephone line is generally 56 Vdc applied to the line through a balanced 400 Ω source impedance.

Ringing voltage is a composite signal consisting of an ac component and a dc component. The ac component is up to 175 V peak at 20 Hz to 60 Hz with a 100 Ω source resistance. The dc component is 56 Vdc with a 300 Ω to 600 Ω source resistance. Large reactive transients can occur at the start and end of each ring interval.

Although 1000BASE-T equipment is not required to survive such wiring hazards without damage, application of any of the above voltages shall not result in any safety hazard.

NOTE—Wiring errors may impose telephony voltages differentially across 1000BASE-T transmitters or receivers. Because the termination resistance likely to be present across a receiver's input is of substantially lower impedance than an off-hook telephone instrument, receivers will generally appear to the telephone system as off-hook telephones. Therefore, full-ring voltages will be applied for only short periods. Transmitters that are coupled using transformers will similarly appear like off-hook telephones (though perhaps a bit more slowly) due to the low resistance of the transformer coil.

40.9.3 Environment

40.9.3.1 Electromagnetic emission

A system integrating the 1000BASE-T PHY shall comply with applicable local and national codes for the limitation of electromagnetic interference.

40.9.3.2 Temperature and humidity

A system integrating the 1000BASE-T PHY is expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this standard.

It is recommended that manufacturers indicate in the literature associated with the PHY the operating environmental conditions to facilitate selection, installation, and maintenance.

40.10 PHY labeling

It is recommended that each PHY (and supporting documentation) be labeled in a manner visible to the user with at least the following parameters:

a) Data rate capability in Mb/s
b) Power level in terms of maximum current drain (for external PHYs)
c) Port type (i.e., 1000BASE-T)
d) Any applicable safety warnings

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 229 of 379

40.11 Delay constraints

In half duplex mode, proper operation of a CSMA/CD LAN demands that there be an upper bound on the propagation delays through the network. This implies that MAC, PHY, and repeater implementors must conform to certain delay minima and maxima, and that network planners and administrators conform to constraints regarding the cabling topology and concatenation of devices. MAC constraints are specified in 35.2.4. Topological constraints are contained in Clause 42.

In full duplex mode, predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) also demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementors must conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices.

The reference point for all MDI measurements is the peak point of the mid-cell transition corresponding to the reference code-bit, as measured at the MDI.

40.11.1 MDI to GMII delay constraints

Every 1000BASE-T PHY associated with a GMII shall comply with the bit delay constraints specified in Table 40–13 for half duplex operation and Table 40–14 for full duplex operation. These constraints apply for all 1000BASE-T PHYs. For any given implementation, the assertion and de-assertion delays on CRS shall be equal.

Table 40–13—MDI to GMII delay constraints (half duplex mode)

Sublayer measurement points	Event	Min (bit times)	Max (bit times)	Input timing reference	Output timing reference
GMII ⇔ MDI	TX_EN Sampled to MDI Output	—	84	GTX_CLK rising	1st symbol of SSD/CSReset/ CSExtend/ CSExtend_Err
	MDI input to CRS assert	—	244	1st symbol of SSD/CSReset	—
	MDI input to CRS de-assert	—	244	1st symbol of SSD/CSReset	—
	MDI input to COL assert	—	244	1st symbol of SSD/CSReset	—
	MDI input to COL de-assert	—	244	1st symbol of SSD/CSReset	—
	TX_EN sampled to CRS assert	—	16	GTX_CLK rising	—
	TX_EN sampled to CRS de-assert	—	16	GTX_CLK rising	—

40.11.2 DTE delay constraints (half duplex only)

Every DTE with a 1000BASE-T PHY shall comply with the bit delay constraints specified in Table 40–15 for half duplex operation.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 230 of 379

Table 40–14—MDI to GMII delay constraints (full duplex mode)

Sublayer measurement points	Event	Min (bit times)	Max (bit times)	Input timing reference	Output timing reference
GMII ⇔ MDI	TX_EN Sampled to MDI Output	—	84	GTX_CLK rising	1st symbol of SSD/CSReset/ CSExtend/ CSExtend_Err
	MDI input to RX_DV de-assert	—	244	1st symbol of CSReset	—

Table 40–15— DTE delay constraints (half duplex mode)

Sublayer measurement points	Event	Min (bit times)	Max (bit times)	Input timing reference	Output timing reference
MAC ⇔ MDI	MAC transmit start to MDI output	—	132	—	1st symbol of SSD
	MDI input to collision detect	—	292	1st symbol of SSD	—
	MDI input to MDI output (nondeferred or Jam)	—	440	1st symbol of SSD	1st symbol of SSD
	MDI Input to MDI output (worse-case non-deferred transmit)	—	440	1st symbol of SSD	1st symbol of SSD

40.11.3 Carrier de-assertion/assertion constraint (half duplex mode)

To ensure fair access to the network, each DTE operating in half duplex mode shall, additionally, satisfy the following: (MAX MDI to MAC Carrier De-assert Detect) – (MIN MDI to MAC Carrier Assert Detect) < 16 Bit Times.

40.12 Protocol implementation conformance statement (PICS) proforma for Clause 40—Physical coding sublayer (PCS), physical medium attachment (PMA) sublayer and baseband medium, type 1000BASE-T[11]

The supplier of a protocol implementation that is claimed to conform to this clause shall complete the Protocol Implementation Conformance Statement (PICS) proforma listed in the following subclauses.

Instructions for interpreting and filling out the PICS proforma may be found in Clause 21.

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

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 231 of 379

40.12.1 Identification

40.12.1.1 Implementation identification

Supplier	
Contact point for queries about the PICS	
Implementation Name(s) and Version(s)	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)	
NOTES 1—Only the first three items are required for all implementations; other information may be completed as appropriate in meeting the requirements for the identification. 2—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

40.12.1.2 Protocol summary

Identification of protocol specification	IEEE Std 802.3-2002®, Clause 40, Physical coding sub-layer (PCS), physical medium attachment (PMA) sub-layer, and baseband medium, type 1000BASE-T
Identification of amendments and corrigenda to this PICS proforma which have been completed as part of this PICS	
Have any Exceptions items been required? No [ ] Yes [ ] (See Clause 21—The answer Yes means that the implementation does not conform to the standard)	
Date of Statement	

40.12.2 Major capabilities/options

Item	Feature	Subclause	Status	Support	Value/Comment
*GMII	PHY associated with GMII	40.1	O	Yes [ ] No [ ]	
*DTE	DTE with PHY not associated with GMII	40.1	O	Yes [ ] No [ ]	
AN	Support for Auto-Negotiation (Clause 28)	40.5.1	M	Yes [ ]	Required
OMS	Operation as MASTER or SLAVE	40.5.1	M	Yes [ ]	Required
*FDX	PHY supports full duplex mode	40.1	O	Yes [ ] No [ ]	
*HDX	PHY support half duplex mode	40.1	O	Yes [ ] No [ ]	
*INS	Installation / cabling	40.7	O	Yes [ ] No [ ]	Items marked with INS include installation practices and cabling specifications not applicable to a PHY manufacturer.
*AXO	Auto-Crossover	40.4.4	O	Yes [ ] No [ ]	PHY supports auto-crossover

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 232 of 379

40.12.3 Clause conventions

Item	Feature	Subclause	Status	Support	Value/Comment
CCO1	The values of all components in test circuits shall be	40.1.6	M	Yes [ ]	Accurate to within ±1% unless otherwise stated.

40.12.4 Physical Coding Sublayer (PCS)

Item	Feature	Subclause	Status	Support	Value/Comment
PCT1	The PCS shall	40.3.1.2	M	Yes [ ]	Implement the Data Transmission Enabling process as depicted in Figure 40–8 including compliance with the associated state variables specified in 40.3.3.
PCT2	PCS Transmit function shall	40.3.1.3	M	Yes [ ]	Conform to the PCS Transmit state diagram in Figure 40–9.
PCT3	PCS Transmit shall	40.4.5.1	M	Yes [ ]	Send code-groups according to the value assumed by the tx_mode variable.
PCT4	If the parameter config provided to the PCS by the PHY Control function via the PMA_CONFIG.indicate message assumes the value MASTER, PCS Transmit shall	40.3.1.3.1	M	Yes [ ]	Employ the transmitter side-stream scrambler generator polynomial specified for use with MASTER in 40.3.1.3.1.
PCT5	If the parameter config provided to the PCS by the PHY Control function via the PMA_CONFIG.indicate message assumes the value SLAVE, PCS Transmit shall	40.3.1.3.1	M	Yes [ ]	Employ the transmitter side-stream scrambler generator polynomial specified for use with SLAVE in 40.3.1.3.1.
PCT6	In no case shall	40.3.1.3.1	M	Yes [ ]	The scrambler state be initialized to all zeros.
PCT7	If tx_error_n=1 when the condition (tx_enable_n * tx_enable_{n-2}) = 1, error indication is signaled by means of symbol substitution, wherein the values of Sd_n[5:0] are ignored during mapping and the symbols corresponding to the row denoted as "xmt_err" in Table 40–1 and Table 40–2 shall be used.	40.3.1.3.5	M	Yes [ ]	

Copyright © 2002 IEEE. All rights reserved.

Item	Feature	Subclause	Status	Support	Value/Comment
PCT8	If tx_error_n=0 when the variable $csreset_n$ = 1, the convolutional encoder reset condition is normal. This condition is indicated by means of symbol substitution, where the values of Sd_n[5:0] are ignored during mapping and the symbols corresponding to the row denoted as "CSReset" in Table 40–1 and Table 40–2 shall be used.	40.3.1.3.5	M	Yes [ ]	
PCT9	If tx_error_n=1 is asserted when the variable $csreset_n$ = 1, the convolutional encoder reset indicates carrier extension. In this condition, the values of Sd_n[5:0] are ignored during mapping and the symbols corresponding to the row denoted as "CSExtend" in Table 40–1 and Table 40–2 shall be used when TXD_n = 0x'0F, and the row denoted as "CSExtend_Err" in Table 40–1 and Table 40–2 shall be used when $TXD_n \neq$ 0x'0F.	40.3.1.3.5	M	Yes [ ]	
PCT10	In case carrier extension with error is indicated during the first octet of CSReset, the error condition shall be encoded during the second octet of CSReset, and during the subsequent two octets of the End-of-Stream delimiter.	40.3.1.3.5	M	Yes [ ]	
PCT11	The symbols corresponding to the SSD1 row in Table 40–1 shall be used when the condition (tx_enable_n) * (!tx_enable_{n-1}) = 1.	40.3.1.3.5	M	Yes [ ]	
PCT12	The symbols corresponding to the SSD2 row in Table 40–1 shall be used when the condition (tx_enable_{n-1})* (!tx_enable_{n-2}) = 1.	40.3.1.3.5	M	Yes [ ]	
PCT13	If carrier extend error is indicated during ESD, the symbols corresponding to the ESD_Ext_Err row in Table 40–1 shall be used.	40.3.1.3.5	M	Yes [ ]	
PCT14	The symbols corresponding to the ESD1 row in Table 40–1 shall be used when the condition (!tx_enable_{n-2}) * (tx_enable_{n-3}) = 1, in the absence of carrier extend error indication at time n.	40.3.1.3.5	M	Yes [ ]	

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 234 of 379

Item	Feature	Subclause	Status	Support	Value/Comment
PCT15	The symbols corresponding to the ESD2_Ext_0 row in shall be used when the condition (!tx_enable_{n-3}) * (tx_enable_{n-4}) * (!tx_error_{n}) * (!tx_error_{n-1}) = 1.	40.3.1.3.5	M	Yes [ ]	
PCT16	The symbols corresponding to the ESD2_Ext_1 row in Table 40–1 shall be used when the condition (!tx_enable_{n-3}) * (tx_enable_{n-4}) * (!tx_error_{n}) * (tx_error_{n-1}) * (tx_error_{n-2}) * (tx_error_{n-3})= 1.	40.3.1.3.5	M	Yes [ ]	
PCT17	The symbols corresponding to the ESD2_Ext_2 row in Table 40–1 shall be used when the condition (!tx_enable_{n-3}) * (tx_enable_{n-4}) * (tx_error_{n}) * (tx_error_{n-1}) * (tx_error_{n-2}) * (tx_error_{n-3}) * (TXD_{n}=0x0F)= 1, in the absence of carrier extend error indication.	40.3.1.3.5	M	Yes [ ]	

40.12.4.1 PCS receive functions

Item	Feature	Subclause	Status	Support	Value/Comment
PCR1	PCS Receive function shall	40.3.1.4	M	Yes [ ]	Conform to the PCS Receive state diagram shown in Figure 40–10a including compliance with the associated state variables as specified in 40.3.3.
PCR2	The PHY shall	40.3.1.4.2	M	Yes [ ]	Descramble the data stream and return the proper sequence of data bits RXD<7:0> to the GMII.
PCR3	For side-stream descrambling, the MASTER PHY shall employ	40.3.1.4.2	M	Yes [ ]	The receiver scrambler generator polynomial specified for MASTER operation in 40.3.1.4.2.
PCR4	For side-stream descrambling, the SLAVE PHY shall employ	40.3.1.4.2	M	Yes [ ]	The receiver scrambler generator polynomial specified for SLAVE operation in 40.3.1.4.2.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 235 of 379

40.12.4.2 Other PCS functions

Item	Feature	Subclause	Status	Support	Value/Comment
PCO1	The PCS Reset function shall	40.3.1.1	M	Yes [ ]	Be executed any time "power on" or receipt of a request for reset from the management entity occurs, including compliance with the associated state variables as specified in 40.3.3.
PCO2	The PCS shall	40.3.1.5	M	Yes [ ]	Implement the Carrier Sense process as depicted in Figure 40–11, including compliance with the associated state variables as specified in 40.3.3.
PCO3	Symb-timer shall be generated	40.3.3.3	M	Yes [ ]	Synchronously with TX_TCLK.

40.12.5 Physical Medium Attachment (PMA)

Item	Feature	Subclause	Status	Support	Value/Comment
PMF1	PMA Reset function shall be executed	40.4.2.1	M	Yes [ ]	At power on and upon receipt of a reset request from the management entity or from PHY Control.
PMF2	PMA Transmit shall	40.4.2.2	M	Yes [ ]	Continuously transmit onto the MDI pulses modulated by the quinary symbols given by tx_symb_vector[BI_DA], tx_symb_vector[BI_DB], tx_symb_vector[BI_DC], and tx_symb_vector[BI_DD], respectively.
PMF3	The four transmitters shall be driven by the same transmit clock, TX_TCLK	40.4.2.2	M	Yes [ ]	
PMF4	PMA Transmit shall	40.4.2.2	M	Yes [ ]	Follow the mathematical description given in 40.4.3.1.
PMF5	PMA Transmit shall comply with	40.4.2.2	M	Yes [ ]	The electrical specifications given in 40.6.
PMF6	When the PMA_CONFIG.indicate parameter config is MASTER, the PMA Transmit function shall	40.4.2.2	M	Yes [ ]	Source the transmit clock TX_TCLK from a local clock source while meeting the transmit jitter requirements of 40.6.1.2.5.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 236 of 379

Item	Feature	Subclause	Status	Support	Value/Comment
PMF7	When the PMA_CONFIG.indicate parameter config is SLAVE, the PMA Transmit function shall	40.4.2.2	M	Yes [ ]	Source the transmit clock TX_TCLK from the recovered clock of 40.4.2.5 while meeting the jitter requirements of 40.6.1.2.5.
PMF8	PMA Receive function shall	40.4.2.3	M	Yes [ ]	Translate the signals received on pairs BI_DA BI_DB, BI_DC and BI_DD into the PMA_UNITDATA.indicate parameter rx_symb_vector with a symbol error rate of less than one part in 10^{10}.
PMF9	PHY Control function shall	40.4.2.4	M	Yes [ ]	Comply with the state diagram descriptions given in Figure 40–15.
PMF10	The Link Monitor function shall	40.4.2.5	M	Yes [ ]	Comply with the state diagram shown in Figure 40–16.
PMF11	Clock Recovery function shall provide	40.4.2.6	M	Yes [ ]	Provide clocks suitable for signal sampling on each line so that the symbol-error rate indicated in 40.4.2.3 is achieved.
PMF12	The symbol response shall comply with	40.4.3.1	M	Yes [ ]	The electrical specifications given in 40.6.
PMF13	The four signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DD shall be processed within the PMA Receive function to yield	40.4.3.2	M	Yes [ ]	The quinary received symbols rx_symb_vector[BI_DA], rx_symb_vector[BI_DB], rx_symb_vector[BI_DC], and rx_symb_vector[BI_DD].
PMF14	If an automatic configuration method is used, it shall	40.4.4	M	Yes [ ]	Comply with the specifications of 40.4.4.
PMF15	The PMA shall	40.4.5.1	M	Yes [ ]	Generate the config variable continuously and pass it to the PCS via the PMA_CONFIG.indicate primitive.
PMF16	The variable link_det shall take the value	40.4.5.1	AXO:M	N/A [ ] Yes [ ]	TRUE or FALSE as per 40.4.4.1.
PMF17	The variable MDI_status shall take the value	40.4.5.1	AXO:M	N/A [ ] Yes [ ]	MDI or MDI-X as per Table 40–12.
PMF18	PCS Transmit shall	40.4.5.1	M	Yes [ ]	Send code-groups according to the value assumed by tx_mode.
PMF19	The A_timer shall have a period of	40.4.5.2	AXO:M	N/A [ ] Yes [ ]	1.3s ± 25%.
PMF20	The maxwait_timer timer shall expire	40.4.5.2	M	Yes [ ]	750 ± 10 ms if config = MASTER or 350 ± 5ms if config = SLAVE

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 237 of 379

Item	Feature	Subclause	Status	Support	Value/Comment
PMF21	The minwait_timer timer shall expire	40.4.5.2	M	Yes [ ]	1 ±0.1μs after being started.
PMF22	The sample_timer shall have a period of	40.4.5.2	AXO:M	N/A [ ] Yes [ ]	62 ± 2ms.
PMF23	The stabilize_timer shall expire	40.4.5.2	M	Yes [ ]	1 ± 0.1 μs after being started.

40.12.6 Management interface

Item	Feature	Subclause	Status	Support	Value/Comment
MF1	All 1000BASE-T PHYs shall provide support for Auto-Negotiation (Clause 28) and shall be capable of operating as MASTER or SLAVE.	40.5.1	M	Yes [ ]	
MF2	A 100BASE-T PHY shall	40.5.1.1	M	Yes [ ]	Use the management register definitions and values specified in Table 40–3.

40.12.6.1 1000BASE-T Specific Auto-Negotiation Requirements

Item	Feature	Subclause	Status	Support	Value/Comment
AN1	1000BASE-T PHYs shall	40.5.1.2	M	Yes [ ]	Exchange one Auto-Negotiation Base Page, a 1000BASE-T formatted Next Page, and two 1000BASE-T unformatted Next Pages in sequence, without interruption, as specified in Table 40–4.
AN2	The MASTER-SLAVE relationship shall be determined during Auto-Negotiation	40.5.2	M	Yes [ ]	Using Table 40–5 with the 1000BASE-T Technology Ability Next Page bit values specified in Table 40–4 and information received from the link partner.
AN3	Successful completion of the MASTER-SLAVE resolution shall	40.5.2	M	Yes [ ]	Be treated as MASTER-SLAVE configuration resolution complete.
AN4	A seed counter shall be provided to	40.5.2	M	Yes [ ]	Track the number of seed attempts.
AN5	At start-up, the seed counter shall be set to	40.5.2	M	Yes [ ]	Zero.
AN6	The seed counter shall be incremented	40.5.2	M	Yes [ ]	Every time a new random seed is sent.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 238 of 379

Item	Feature	Subclause	Status	Support	Value/Comment
AN7	When MASTER-SLAVE resolution is complete, the seed counter shall be reset to 0 and bit 10.15 shall be set to logical zero.	40.5.2	M	Yes [ ]	
AN8	Maximum seed attempts before declaring a MASTER_SLAVE configuration Resolution Fault	40.5.2	M	Yes [ ]	Seven.
AN9	During MASTER_SLAVE configuration, the device with the higher seed value shall	40.5.2	M	Yes [ ]	Become the MASTER.
AN10	During MASTER_SLAVE configuration, the device with the lower seed value shall	40.5.2	M	Yes [ ]	Become the SLAVE.
AN11	Both PHYs set in manual mode to be either MASTER or SLAVE shall be treated as	40.5.2	M	Yes [ ]	MASTER-SLAVE resolution fault (failure) condition.
AN12	MASTER-SLAVE resolution fault (failure) condition shall result in	40.5.2	M	Yes [ ]	MASTER-SLAVE Configuration Resolution Fault bit (10.15) to be set to logical one.
AN13	MASTER-SLAVE Configuration resolution fault condition shall be treated as	40.5.2	M	Yes [ ]	MASTER-SLAVE Configuration Resolution complete.
AN14	MASTER-SLAVE Configuration resolution fault condition shall	40.5.2	M	Yes [ ]	Cause link_status_1000BASE-T to be set to FAIL.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 239 of 379

40.12.7 PMA Electrical Specifications

Item	Feature	Subclause	Status	Support	Value/Comment
PME15	The PHY shall provide electrical isolation between	40.6.1.1	M	Yes [ ]	The port device circuits including frame ground, and all MDI leads.
PME16	PHY-provided electrical separation shall withstand at least one of three electrical strength tests	40.6.1.1	M	Yes [ ]	a) 1500 V rms at 50Hz to 60Hz for 60 s, applied as specified in Section 5.3.2 of IEC 60950:1991. b) 2250 Vdc for 60 s, applied as specified in Section 5.3.2 of IEC 60950: 1991. c) A sequence of ten 2400 V impulses of alternating polarity, applied at intervals of not less than 1 s. The shape of the impulses shall be 1.2/50 μs. (1.2 μs virtual front time, 50 μs virtual time or half value), as defined in IEC 60060.
PME17	There shall be no insulation breakdown as defined in Section 5.3.2 of IEC 60950, during the test.	40.6.1.1	M	Yes [ ]	
PME18	The resistance after the test shall be at least	40.6.1.1	M	Yes [ ]	2 MΩ, measured at 500 Vdc.
PME19	The transmitter MASTER-SLAVE timing jitter test channel shall	40.6.1.1.1	M	Yes [ ]	Be constructed by combining 100 Ω and 120 Ω cable segments that meet or exceed ISO.IEC 11801 Category 5 specifications for each pair as shown in Figure 40–18 with the lengths and additional restrictions on parameters described in Table 40–6.
PME20	The ends of the MASTER-SLAVE timing jitter test channel shall	40.6.1.1.1	M	Yes [ ]	Be connectorized with connectors meeting or exceeding ANSI/TIA/EIA-568-A:1995 or ISO/IEC 11801:1995 Category 5 specifications.
PME21	The return loss of the MASTER-SLAVE timing jitter test channel shall	40.6.1.1.1	M	Yes [ ]	Meet the return loss requirements of 40.7.2.3.
PME22	The return loss of the MASTER-SLAVE timing jitter test channel shall	40.6.1.1.1	M	Yes [ ]	Meet the crosstalk requirements of 40.7.3 on each pair.
PME23	The test modes described in 40.6.1.1.2 shall be provided for testing of the transmitted waveform, transmitter distortion and transmitted jitter.	40.6.1.1.2	M	Yes [ ]	

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 240 of 379

Item	Feature	Subclause	Status	Support	Value/Comment
PME24	For a PHY with a GMII interface the test modes shall be enabled by	40.6.1.1.2	M	Yes [ ]	Setting bits 9:13-15 (1000BASE-T Control Register) of the GMII Management register set as shown in Table 40–7.
PME25	The test modes shall only change the data symbols provided to the transmitter circuitry and shall not alter the electrical and jitter characteristics of the transmitter and receiver from those of normal operation.	40.6.1.1.2	M	Yes [ ]	
PME26	A PHY without a GMII shall provide a means to enable the test modes for conformance testing.	40.6.1.1.2	M	Yes [ ]	
PME27	When transmit test mode 1 is enabled, the PHY shall transmit	40.6.1.1.2	M	Yes [ ]	The sequence of data symbols specified in 40.6.1.1.2 continuously from all four transmitters.
PME28	When in test mode 1, the transmitter shall time the transmitted symbols	40.6.1.1.2	M	Yes [ ]	From a 125.00 MHz ± 0.01% clock in the MASTER timing mode.
PME29	When test mode 2 is enabled, the PHY shall transmit	40.6.1.1.2	M	Yes [ ]	The data symbol sequence {+2,–2} repeatedly on all four channels.
PME30	When in test mode 2, the transmitter shall time the transmitted symbols	40.6.1.1.2	M	Yes [ ]	From a 125.00 MHz ± 0.01% clock in the MASTER timing mode.
PME31	When transmit test mode 3 is enabled, the PHY shall transmit	40.6.1.1.2	M	Yes [ ]	The data symbol sequence {+2,–2} repeatedly on all four channels.
PME32	When in test mode 3, the transmitter shall time the transmitted symbols	40.6.1.1.2	M	Yes [ ]	From a 125 MHz ± 1% clock in the SLAVE timing mode.
PME33	When test mode 4 is enabled, the PHY shall transmit	40.6.1.1.2	M	Yes [ ]	The data symbols generated by the scrambler polynomial specified in 40.6.1.1.2.
PME34	When test mode 4 is enabled, the PHY shall	40.6.1.1.2	M	Yes [ ]	Use the bit sequences generated by the scrambler bits shown in 40.6.1.1.2 to generate the quinary symbols, s_n, as shown in Table 40–8.
PME35	When test mode 4 is enabled, the maximum-length shift register used to generate the sequences defined by this polynomial shall be	40.6.1.1.2	M	Yes [ ]	Updated once per symbol interval (8 ns).

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 241 of 379

Item	Feature	Subclause	Status	Support	Value/Comment
PME36	When test mode 4 is enabled, the bit sequences, $x0_n$, $x1_n$, and $x2_n$, generated from combinations of the scrambler bits shown in 40.6.1.1.2 shall be	40.6.1.1.2	M	Yes [ ]	Used to generate the quinary symbols, s_n, as shown in Table 40–8.
PME37	When test mode 4 is enabled, the quinary symbol sequence shall be	40.6.1.1.2	M	Yes [ ]	Presented simultaneously to all transmitters.
PME38	When in test mode 4, the transmitter shall time the transmitted symbols	40.6.1.1.2	M	Yes [ ]	From a 125.00 MHz ± 0.01% clock in the MASTER timing mode.
PME39	The test fixtures defined in Figure 40–22, Figure 40–23, Figure 40–24, and Figure 40–25 or their functional equivalents shall be used for measuring transmitter specifications.	40.6.1.1.3	M	Yes [ ]	
PME40	The test filter used in transmitter test fixtures 1 and 3 shall	40.6.1.1.3	M	Yes [ ]	Have the continuous time transfer function specified in 40.6.1.1.3 or its discrete time equivalent.
PME41	The disturbing signal V_d shall	40.6.1.1.3	M	Yes [ ]	Have the characteristics listed in Table 40–9.
PME42	To allow for measurement of transmitted jitter in MASTER and SLAVE modes the PHY shall provide access to the 125 MHz symbol clock, TX_TCLK that times the transmitted symbols.	40.6.1.1.3	M	Yes [ ]	
PME43	To allow for measurement of transmitted jitter in MASTER and SLAVE modes the PHY shall provide a means to enable the TX_TCLK output if it is not normally enabled.	40.6.1.1.3	M	Yes [ ]	
PME44	The PMA shall	40.6.1.2	M	Yes [ ]	Provide the Transmit function specified in 40.4.2.2 in accordance with the electrical specifications of this clause.
PME45	Where a load is not specified, the transmitter shall	40.6.1.2	M	Yes [ ]	Meet the requirements of this clause with a 100 Ω resistive differential load connected to each transmitter output.
PME46	The tolerance on the poles of the test filters used in 40.6 shall be ± 1%.	40.6.1.2	M	Yes [ ]	

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 242 of 379

Item	Feature	Subclause	Status	Support	Value/Comment
PME47	When in transmit test mode 1 and observing the differential signal output at the MDI using test fixture 1, for each pair, with no intervening cable, the absolute value of the peak of the waveform at points A and B as defined in Figure 40–19 shall fall within	40.6.1.2.1	M	Yes [ ]	The range of 0.67 V to 0.82 V (0.75 V ± 0.83 dB).
PME48	The absolute value of the peak of the waveforms at points A and B shall	40.6.1.2.1	M	Yes [ ]	Differ by less than 1%.
PME49	The absolute value of the peak of the waveform at points C and D as defined in Figure 40–19 shall differ	40.6.1.2.1	M	Yes [ ]	From 0.5 times the average of the absolute values of the peaks of the waveform at points A and B by less than 2%.
PME50	When in transmit test mode 1 and observing the differential transmitted output at the MDI, for either pair, with no intervening cabling, the peak value of the waveform at point F as defined in Figure 40–19 shall be	40.6.1.2.2	M	Yes [ ]	Greater than 73.1% of the magnitude of the negative peak value of the waveform at point F. Point G is defined as the point exactly 500 ns after point F. Point F is defined as the point where the waveform reaches it's minimum value at the location indicated in Figure 40–19.
PME51	When in transmit test mode 1 and observing the differential transmitted output at the MDI, for either pair, with no intervening cabling, the peak value of the waveform at point J as defined in Figure 40–19 shall be	40.6.1.2.2	M	Yes [ ]	Greater than 73.1% of the magnitude of the peak value of the waveform at point H. Point J is defined as the point exactly 500 ns after point H. Point H is defined as the point where the waveform reaches it's maximum value at the location indicated in Figure 40–19.
PME52	When in test mode 1 and observing the differential signal output at the MDI using transmitter test fixture 1, for each pair, with no intervening cable, the voltage waveforms at points A, B, C, D defined in Figure 40–19, after the normalization described within the referenced subclause, shall	40.6.1.2.3	M	Yes [ ]	Lie within the time domain template 1 defined in Figure 40–26 and the piecewise linear interpolation between the points in Table 40–10. The waveforms may be shifted in time as appropriate to fit within the template.
PME53	When in test mode 1 and observing the differential signal output at the MDI using transmitter test fixture 1, for each pair, with no intervening cable, the voltage waveforms at points F and H defined in Figure 40–19, after the normalization described within the referenced subclause, shall	40.6.1.2.3	M	Yes [ ]	Lie within the time domain template 2 defined in Figure 40–26 and the piecewise linear interpolation between the points in Table 40–11. The waveforms may be shifted in time as appropriate to fit within the template.

Copyright © 2002 IEEE. All rights reserved.

Item	Feature	Subclause	Status	Support	Value/Comment
PME54	When in test mode 4 and observing the differential signal output at the MDI using transmitter test fixture 3, for each pair, with no intervening cable, the peak distortion as defined below shall be	40.6.1.2.4	M	Yes [ ]	Less than 10 mV.
PME55	When in the normal mode of operation as the MASTER, the peak-to-peak value of the MASTER TX_TCLK jitter relative to an unjittered reference shall be	40.6.1.2.5	M	Yes [ ]	Less than 1.4 ns.
PME56	When the jitter waveform on TX_TCLK is filtered by a high-pass filter, $H_{jf1}(f)$ having the transfer function specified in 40.6.1.2.5, the peak-to-peak value of the resulting filtered timing jitter plus J_{txout}, shall be	40.6.1.2.5	M	Yes [ ]	Less than 0.3 ns.
PME57	When in the normal mode of operation as the SLAVE, receiving valid signals from a compliant PHY operating as the MASTER using the test channel defined in 40.6.1.1.1, with test channel port A connected to the SLAVE, the peak-to-peak value of the SLAVE TX_TCLK jitter relative to the MASTER TX_TCLK shall be	40.6.1.2.5	M	Yes [ ]	Less than 1.4 ns after the receiver is properly receiving the data and has set bit 10.13 of the GMII management register set to 1.
PME58	When the jitter waveform on TX_TCLK is filtered by a high-pass filter, $H_{jf2}(f)$, having the transfer function specified in 40.6.1.2.5, the peak-to-peak value of the resulting filtered timing jitter plus J_{txout} shall be	40.6.1.2.5	M	Yes [ ]	No more than 0.4 ns greater than the simultaneously measured peak-to-peak value of the MASTER jitter filtered by $H_{jf1}(f)$
PME59	For all jitter measurements the peak-to-peak value shall be	40.6.1.2.5	M	Yes [ ]	Measured over an unbiased sample of at least 10^5 clock edges.
PME60	For all unfiltered jitter measurements the peak-to-peak value shall be	40.6.1.2.5	M	Yes [ ]	Measured over an interval of not less than 100 ms and not more than 1 second.
PME61	The quinary symbol transmission rate on each pair of the MASTER PHY shall be	40.6.1.2.6	M	Yes [ ]	125.00 MHz ± 0.01%
PME62	The PMA shall provide the Receive function specified in 40.3.1.4 in accordance with the electrical specifications of this clause.	40.6.1.3	M	Yes [ ]	

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 244 of 379

Item	Feature	Subclause	Status	Support	Value/Comment
PME63	The patch cabling and interconnecting hardware used in test configurations shall be	40.6.1.3	M	Yes [ ]	Within the limits specified in 40.7.
PME64	Differential signals received on the receive inputs that were transmitted within the specifications given in 40.6.1.2 and have then passed through a link compatible with 40.7, shall be translated into	40.6.1.3.1	M	Yes [ ]	One of the PMA_UNITDATA.indicate messages with a 4-D symbol rate error less than 10^{-10} and sent to the PCS after link bring-up. Since the 4-D symbols are not accessible, this specification shall be satisfied by a frame error rate less than 10^{-7} for 125 octet frames.
PME65	The receive feature shall	40.6.1.3.2	M	Yes [ ]	Properly receive incoming data with a 5-level symbol rate within the range 125.00 MHz ± 0.01%.
PME66	The signal generator for the common-mode test shall be	40.6.1.3.3	M	Yes [ ]	Capable of providing a sine wave signal of 1 MHz to 250 MHz.
PME67	While sending data from the transmitter the receiver shall	40.6.1.3.3	M	Yes [ ]	Send the proper PMA_UNITDATA.indicate messages to the PCS as the signal generator frequency is varied from 1 MHz to 250 MHz.
PME68	While receiving data from a transmitter specified in 40.6.1.2 through a link segment specified in 40.7 connected to all MDI duplex channels, a receiver shall	40.6.1.3.4	M	Yes [ ]	Send the proper PMA_UNITDATA.indicate message to the PCS when any one of the four pairs is connected to a noise source as described in Figure 40–28.
PME69	The alien crosstalk test specified in 40.6.1.3.4 shall be satisfied by	40.6.1.3.4	M	Yes [ ]	A frame error rate of less than 10^{-7} for 125 octet frames
PME70	The noise source shall be	40.6.1.3.4	M	Yes [ ]	Connected to one of the MDI inputs using Category 5 balanced cable of a maximum length of 0.5 m.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 245 of 379

40.12.8 Characteristics of the link segment

Item	Feature	Subclause	Status	Support	Value/Comment
LKS1	All implementations of the balanced cabling link shall	40.7.1	M	Yes [ ]	Be compatible at the MDI.
LKS2	1000BASE-T links shall	40.7.1	M	Yes [ ]	Consist of Category 5 components as specified in ANSI/TIA/EIA-568-A:1995 and ISO/IEC 11801:1995.
LKS3	Link segment testing shall be conducted using	40.7.2	M	Yes [ ]	Source and load impedances of 100 Ω.
LKS4	The tolerance on the poles of the test filter used in this section shall be	40.7.2		Yes [ ]	± 1%.
LKS5	The insertion loss of each duplex channel shall be	40.7.2.1	M	Yes [ ]	Less than $2.1 f^{0.529} + 0.4/f$ (dB) at all frequencies from 1 MHz to 100 MHz. This includes the attenuation of the balanced cabling pairs, connector losses, and patch cord losses of the duplex channel.
LKS6	The insertion loss specification shall be met when	40.7.2.1	M	Yes [ ]	The duplex channel is terminated in 100 Ω.
LKS7	The return loss of each duplex channel shall be	40.7.2.3	M	Yes [ ]	As specified in 40.7.2.3 at all frequencies from 1 MHz to 100 MHz.
LKS8	The reference impedance for return loss measurement shall be	40.7.2.3	M	Yes [ ]	100 Ω.
LKS9	The NEXT loss between duplex channel pairs of a link segment shall be	40.7.3.1.1	M	Yes [ ]	At least $27.1 - 16.8\log_{10}(f/100)$ (where f is the frequency in MHz over the frequency range 1 MHz to 100 MHz.)
LKS10	The worst case ELFEXT loss between duplex channel pairs of a link segment shall be	40.7.3.2	M	Yes [ ]	Greater than $17 - 20\log_{10}(f/100)$ dB (where f is the frequency in MHz) over the frequency range 1 MHz to 100 MHz.
LKS11	The Power Sum loss between a duplex channel and the three adjacent disturbers shall be	40.7.3.2.2	M	Yes [ ]	Greater than $14.4 - 20\log_{10}(f/100)$ dB where f is the frequency in MHz over the frequency range of 1 MHz to 100 MHz.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 246 of 379

Item	Feature	Subclause	Status	Support	Value/Comment
LKS12	The propagation delay of a link segment shall	40.7.4.1	M	Yes [ ]	Not exceed 570 ns at all frequencies from 2 MHz to 100 MHz.
LKS13	The difference in propagation delay, or skew, between all duplex channel pair combinations of a link segment under all conditions shall not exceed	40.7.4.2	M	Yes [ ]	50 ns at all frequencies between 2 MHz and 100 MHz.
LKS14	Once installed, the skew between pairs due to environmental conditions shall not vary	40.7.4.2	M	Yes [ ]	More than ± 10 ns.

40.12.9 MDI requirements

Item	Feature	Subclause	Status	Support	Value/Comment
MDI1	MDI connector	40.8.1	M	Yes [ ]	8-Way connector as per IEC 60603-7: 1990.
MDI2	Connector used on cabling	40.8.1	M	Yes [ ]	Plug.
MDI3	Connector used on PHY	40.8.1	M	Yes [ ]	Jack (as opposed to plug).
MDI4	MDI connector	40.8.2	M	Yes [ ]	A PHY that implements the crossover function shall be marked with the graphical symbol X.
MDI5	The MDI connector (jack) when mated with a balanced cabling connector (plug) shall	40.8.3	M	Yes [ ]	Meet the electrical requirements for Category 5 connecting hardware for use with 100 Ω Category 5 cable as specified in ANSI/TIA/EIA-568-A:1995 and ISO/IEC 11801:1995.
MDI6	The mated MDI connector and balanced cabling connector shall	40.8.3	M	Yes [ ]	Not have a FEXT loss greater than $40 - 20\log_{10}(f/100)$ over the frequency range 1 MHz to 100 MHz between all contact pair combinations shown in Table 40–12.
MDI7	No spurious signals shall be emitted onto the MDI when the PHY is held in power down mode as defined in 22.2.4.1.5, independent of the value of TX_EN, when released from power down mode, or when external power is first applied to the PHY.	40.8.3	M	Yes [ ]	

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 247 of 379

Item	Feature	Subclause	Status	Support	Value/Comment
MDI8	The differential impedance as measured at the MDI for each transmit/receive channel shall be such that	40.8.3.1	M	Yes [ ]	Any reflection due to differential signals incident upon the MDI from a balanced cabling having an impedance of 100 Ω ± 15% is at least 16 dB over the frequency range of 2.0 MHz to 40 MHz and at least $10 - 20\log_{10}(f/80)$ dB over the frequency range 40 MHz to 100 MHz (f in MHz).
MDI9	This return loss shall be maintained	40.8.3.1	M	Yes [ ]	At all times when the PHY is transmitting data or control symbols.
MDI10	The common-mode to differential-mode impedance balance of each transmit output shall exceed	40.8.3.2	M	Yes [ ]	The value specified by the equations specified in 40.8.3.2. Test mode 4 may be used to generate an appropriate transmitter output.
MDI11	The magnitude of the total common-mode output voltage, E_{cm_out}, on any transmit circuit, when measured as shown in Figure 40–32, shall be	40.8.3.3	M	Yes [ ]	Less than 50 mv peak-to-peak when transmitting data.
MDI12	Each wire pair of the MDI shall	40.8.3.4	M	Yes [ ]	Withstand without damage the application of short circuits across the MDI port for an indefinite period of time without damage.
MDI13	Each wire pair of the MDI shall resume	40.8.3.4	M	Yes [ ]	Normal operation after such faults are removed.
MDI14	The magnitude of the current through the short circuit specified in PME64 shall not exceed	40.8.3.4	M	Yes [ ]	300 mA.
MDI15	Each wire pair shall withstand without damage	40.8.3.4	M	Yes [ ]	A 1000 V common-mode impulse of either polarity (E_{cm} as indicated in Figure 40–33).
MDI16	The shape of the impulse shall be	40.8.3.4	M	Yes [ ]	0.3/50 μs (300 ns virtual front time, 50 μs virtual time of half value), as defined in IEC 60060.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 248 of 379

40.12.10 General safety and environmental requirements

Item	Feature	Subclause	Status	Support	Value/Comment
ENV1	Conformance to safety specifications	40.9.1	M	Yes [ ]	IEC 60950.
ENV2	Installation practice	40.9.2.1	INS:M	N/A [ ] Yes [ ]	Sound practice, as defined by applicable local codes.
ENV3	Care taken during installation to ensure that non-insulated network cabling conductors do not make electrical contact with unintended conductors or ground.	40.9.2.2	INS:M	N/A [ ] Yes [ ]	
ENV4	1000BASE-T equipment shall be capable of withstanding a telephone battery supply from the outlet as described in 40.9.2.3.	40.9.2.3	M	Yes [ ]	
ENV5	A system integrating the 1000BASE-T PHY shall comply with applicable local and national codes for the limitation of electromagnetic interference.	40.9.3.1	INS:M	N/A [ ] Yes [ ]	

40.12.11 Timing requirements

Item	Feature	Subclause	Status	Support	Value/Comment
TR1	Every 1000BASE-T PHY associated with a GMII shall	40.11.1	M	Yes [ ]	Comply with the bit delay constraints specified in Table 40–13 for half duplex operation and Table 40–14 for full duplex operation. These constraints apply for all 1000BASE-T PHYs.
TR2	For any given implementation, the assertion delays on CRS shall	40.11.1	M	Yes [ ]	Be equal.
TR3	Every DTE with a 1000BASE-T PHY shall	40.11.2	M	Yes [ ]	Comply with the bit delay constraints specified in Table 40–15.
TR4	To ensure fair access to the network, each DTE operating in half duplex mode shall, additionally, satisfy the following:	40.11.3	M	Yes [ ]	(MAX MDI to MAC Carrier De-assert Detect) – (MIN MDI to MAC Carrier Assert Detect) < 16 Bit Times.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 249 of 379

41. Repeater for 1000 Mb/s baseband networks

41.1 Overview

41.1.1 Scope

Clause 41 defines the functional and electrical characteristics of a repeater for use with ISO/IEC 8802-3 1000 Mb/s baseband networks. A repeater for any other ISO/IEC 8802-3 network type is beyond the scope of this clause. The relationship of this standard to the entire ISO/IEC 8802-3 CSMA/CD LAN standard is shown in Figure 41–1. The purpose of the repeater is to provide a simple, inexpensive, and flexible means of coupling two or more segments.

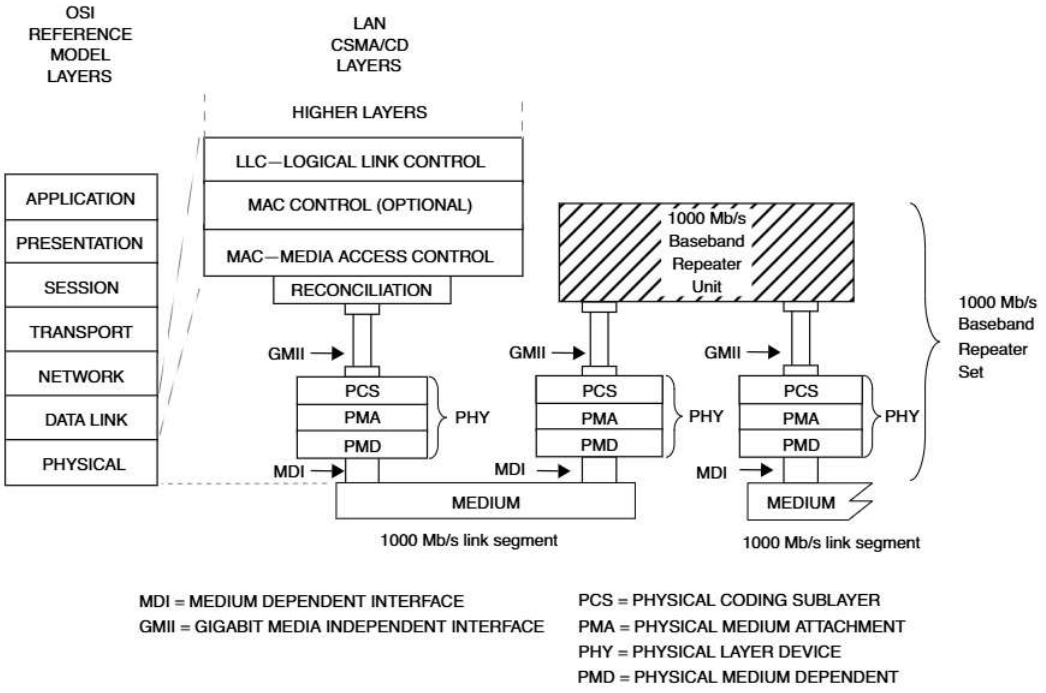


Figure 41–1—1000 Mb/s repeater set relationship to the ISO/IEC OSI reference model

41.1.1.1 Repeater set

Repeater sets are an integral part of all 1000 Mb/s baseband networks with more than two DTEs and are used to extend the physical system topology by providing a means of coupling two or more segments. A single repeater set is permitted within a single collision domain to provide the maximum connection path length. Allowable topologies contain only one operative signal path between any two points on the network. A repeater set is not a station and does not count toward the overall limit of 1024 stations on a network.

A repeater set can receive and decode data from any segment under worst-case noise, timing, and signal amplitude conditions. It retransmits the data to all other segments attached to it with timing, amplitude, and coding restored. The retransmission of data occurs simultaneously with reception. If a collision occurs, the repeater set propagates the collision event throughout the network by transmitting a Jam signal. A repeater set also provides a degree of protection to a network by isolating a faulty segment's carrier activity from propagating through the network.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 250 of 379

41.1.1.2 Repeater unit

A repeater unit is a subset of a repeater set containing all the repeater-specific components and functions, exclusive of PHY components and functions. A repeater unit connects to the PHYs using the Gigabit Media Independent Interface (GMII) defined in Clause 35.

41.1.2 Application perspective

This subclause states the broad objectives and assumptions underlying the specification defined through Clause 41.

41.1.2.1 Objectives

a) Provide physical means for coupling two or more LAN segments at the Physical Layer.
b) Support interoperability of independently developed physical, electrical, and optical interfaces.
c) Provide a communication channel with a mean bit error rate, at the physical service interface equivalent to that for the attached PHY.
d) Provide for ease of installation and service.
e) Ensure that fairness of DTE access is not compromised.
f) Provide for low-cost networks, as related to both equipment and cabling.

41.1.2.2 Compatibility considerations

All implementations of the repeater set shall be compatible at the MDI. The repeater set is defined to provide compatibility among devices designed by different manufacturers. Designers are free to implement circuitry within the repeater set in an application-dependent manner provided the appropriate PHY specifications are met.

41.1.2.2.1 Internal segment compatibility

Implementations of the repeater set that contain a MAC layer for network management or other purposes, irrespective of whether they are connected through an exposed repeater port or are internally ported, shall conform to the requirements of Clause 30 on that port if repeater management is implemented.

41.1.3 Relationship to PHY

A close relationship exists between Clause 41 and the GMII clause (Clause 35) and the PHY clauses (Clauses 36–39 for 1000BASE-X PHYs and Clause 40 for 1000BASE-T PHYs). The PHY's PMA, PCS, and MDI specification provide the actual medium attachment, including drivers, receivers, and Medium Interface Connectors for the various supported media. The repeater clause does not define a new PHY; it utilizes the existing PHYs complete and without modification. The repeater_mode variable in each PHY is set, so that the CRS signal of the GMII is asserted only in response to receive activity (see 36.2.5.1.3).

41.2 Repeater functional specifications

A repeater set provides the means whereby data from any segment can be received under worst-case noise, timing, and amplitude conditions and then retransmitted with timing and amplitude restored to all other attached segments. Retransmission of data occurs simultaneously with reception. If a collision occurs, the repeater set propagates the collision event throughout the network by transmitting a Jam signal. If an error is received by the repeater set, no attempt is made to correct it and it is propagated throughout the network by transmitting an explicit error code.

The repeater set provides the following functional capability to handle data flow between ports:

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 251 of 379

a) *Signal restoration*. Provides the ability to restore the timing and amplitude of the received signal prior to retransmission.
b) *Transmit function*. Provides the ability to output signals on the appropriate port and encoded appropriately for that port. Details of signal processing are described in the specifications for the PHYs.
c) *Receive function*. Provides the ability to receive input signals presented to the ports. Details of signal processing are described in the specifications for the PHYs.
d) *Data-Handling function*. Provides the ability to transfer code-elements between ports in the absence of a collision.
e) *Received Event-Handling requirement*. Provides the ability to derive a carrier signal from the input signals presented to the ports.
f) *Collision-Handling function*. Provides the ability to detect the simultaneous reception of frames at two or more ports and then to propagate a Jam message to all connected ports.
g) *Error-Handling function*. Provides the ability to prevent substandard links from generating streams of false carrier and interfering with other links.
h) *Partition function*. Provides the ability to prevent a malfunctioning port from generating an excessive number of consecutive collisions and indefinitely disrupting data transmission on the network.
i) *Receive Jabber function*. Provides the ability to interrupt the reception of abnormally long streams of input data.

41.2.1 Repeater functions

The repeater set shall provide the Signal Restoration, Transmit, Receive, Data Handling, Received Event Handling, Collision Handling, Error Handling, Partition, and Receive Jabber functions. The repeater is transparent to all network acquisition activity and to all DTEs. The repeater will not alter the basic fairness criterion for all DTEs to access the network or weigh it toward any DTE or group of DTEs regardless of network location.

The Transmit and Receive functional requirements are specified by the PHY clauses, Clause 40 for 1000BASE-T and Clauses 36 to 39 for 1000BASE-X.

41.2.1.1 Signal restoration functional requirements

41.2.1.1.1 Signal amplification

The repeater set (including its integral PHYs) shall ensure that the amplitude characteristics of the signals at the MDI outputs of the repeater set are within the tolerances of the specification for the appropriate PHY type. Therefore, any loss of signal-to-noise ratio due to cable loss and noise pickup is regained at the output of the repeater set as long as the incoming data is within system specification.

41.2.1.1.2 Signal wave-shape restoration

The repeater set (including its integral PHYs) shall ensure that the wave-shape characteristics of the signals at the MDI outputs of a repeater set are within the specified tolerance for the appropriate PHY type. Therefore, any loss of wave-shape due to PHYs and media distortion is restored at the output of the repeater set.

41.2.1.1.3 Signal retiming

The repeater set (including its integral PHYs) shall ensure that the timing of the encoded data output at the MDI outputs of a repeater set are within the specified tolerance for the appropriate PHY type. Therefore, any receive jitter from the media is removed at the output of the repeater set.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 252 of 379

41.2.1.2 Data-handling functional requirements

41.2.1.2.1 Data frame forwarding

The repeater set shall ensure that the data frame received on a single input port is distributed to all other output ports in a manner appropriate for the PHY type of that port. The data frame is that portion of the packet after the SFD and before the end-of-frame delimiter. The only exceptions to this rule are when contention exists among any of the ports, when the receive port is partitioned as defined in 41.2.1.6, when the receive port is in the Jabber state as defined in 41.2.1.7, or when the receive port is in the Link Unstable state as defined in 41.2.1.5.1. Between unpartitioned ports, the rules for collision handling (see 41.2.1.4) take precedence.

41.2.1.2.2 Received code violations

The repeater set shall ensure that any code violations received while forwarding a packet are propagated to all outgoing segments. These code violations shall be replaced by a code-group that provide an explicit indication that an error was received, as appropriate for the outgoing PHY type. Once a received code violation has been replaced by a code-group indicating a receive error, this substitution shall continue for the remainder of the received event regardless of its content. The only exception to this rule is when contention exists among any of the ports, where the rules for collision handling (see 41.2.1.4) then take precedence.

41.2.1.3 Received event-handling functional requirements

41.2.1.3.1 Received event handling

For all its ports, the repeater set shall detect received events by monitoring the port for any assertion of the GMII CRS signal that is the result of receive activity. The repeater_mode variable in the PHY shall be set to ensure that the CRS signal is not asserted in response to transmit activity. Received events include both the data frame and any encapsulation of the data frame such as Preamble, SFD, start and end of packet delimiters, carrier extension symbols, and error propagation symbols. A received event is exclusive of the IDLE pattern. Upon detection of a received event from one port, the repeater set shall repeat all received signals in the data frame from that port to the other ports as described in Figure 41–2.

41.2.1.3.2 Preamble regeneration

The repeater set shall output preamble as appropriate for the outgoing PHY type followed by the SFD. The duration of the output preamble shall not vary more than 8 bit times from the duration of the received preamble.

41.2.1.3.3 Start-of-packet propagation delay

The start-of-packet propagation delay for a repeater set is the time delay between the start of a received event on a repeated-from (input) port to the start of transmit on the repeated-to (output) port (or ports). This parameter is referred to as the SOP delay, and is measured at the MDI of the repeater ports. The maximum value of this delay is constrained such that the sum of the SOP delay and SOJ delay shall not exceed the value specified in 41.2.1.4.3.

41.2.1.3.4 Start-of-packet variability

The start-of-packet variability for a repeater set is defined as the total worst-case difference between start-of-packet propagation delays for successive received events separated by 112 bit times or less at the same input port. The variability shall be less than or equal to 16 bit times.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 253 of 379

41.2.1.4 Collision-handling functional requirements

41.2.1.4.1 Collision detection

The repeater performs collision detection by monitoring all its enabled input ports for received events. When the repeater detects received events on more than one input port, it shall enter a collision state and transmit the Jam message to all of its output ports.

41.2.1.4.2 Jam generation

While a collision is occurring between any of its ports, the repeater unit shall transmit the Jam message to all of the ports. The Jam message shall be transmitted in accordance with the repeater state diagram in Figure 41–2. The Jam message is signalled across the GMII using the Transmit Error Propagation encoding if the collision is detected during Normal Data Transmission, or using the Carrier Extend Error encoding if the collision is detected during Carrier Extension.

41.2.1.4.3 Start-of-collision-jam propagation delay

The start-of-collision Jam propagation delay for a repeater set is the time delay between the start of the second received event (that results in a collision) to arrive at its port and the start of Jam out on all ports. This parameter is referred to as the SOJ delay, and is measured at the MDI of the repeater ports. The sum of the SOP delay and SOJ delay shall not exceed 976 bit times (BT).

41.2.1.4.4 Cessation-of-collision Jam propagation delay

The cessation-of-collision Jam propagation delay for a repeater set is the time delay between the end of the received event that creates a state such that Jam should end at a port and the end of Jam at that port. The states of the input signals that should cause Jam to end are covered in detail in the repeater state diagram in Figure 41–2. This parameter is referred to as the EOJ delay. This delay shall not exceed the SOP delay.

41.2.1.5 Error-handling functional requirements

41.2.1.5.1 Carrier integrity functional requirements

It is desirable that the repeater set protect the network from some transient fault conditions that would disrupt network communications. Potential likely causes of such conditions are DTE and repeater power-up and power-down transients, cable disconnects, and faulty wiring.

The repeater unit shall provide a self-interrupt capability at each port, as described in Figure 41–5, to prevent a segment's spurious carrier activity from propagating through the network.

At each port the repeater shall count consecutive false carrier events signalled across the GMII. The count shall be incremented on each false carrier event and shall be reset on reception of a valid carrier event. In addition, each port shall have a false carrier timer, which is enabled at the beginning of a false carrier event and reset at the conclusion of such an event. A repeater unit shall transmit the Jam signals to all ports for the duration of the false carrier event or until the duration of the event exceeds the time specified by the false_carrier_timer (see 41.2.2.1.4), whichever is shorter. The Jam message shall be transmitted in accordance with the repeater state diagram in Figure 41–2. The LINK UNSTABLE condition shall be detected when the False Carrier Event Count equals the value FCELimit (see 41.2.2.1.1) or the duration of a false carrier event exceeds the time specified by the false_carrier_timer. In addition, the LINK UNSTABLE condition shall be detected upon power-up reset.

Upon detection of LINK UNSTABLE at a port, the repeater unit shall perform the following:

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 254 of 379

a) Inhibit sending further messages from that port to the repeater unit.
b) Inhibit sending further output messages to that port from the repeater unit.
c) Continue to monitor activity on that port.

The repeater unit shall exit the LINK UNSTABLE condition at the port when one of the following is met:

a) The repeater has detected no activity (Idle) for more than the time specified by ipg_timer plus idle_timer (see 41.2.2.1.4) on port X.
b) A valid carrier event with a duration greater than the time specified by valid_carrier_timer (see 41.2.2.1.4) has been received, preceded by no activity (Idle) for more than the time specified by ipg_timer (see 41.2.2.1.4) on port X.

The false_carrier_timer duration is longer than the maximum round-trip latency from a repeater to a DTE, but less than a slot time. This allows a properly functioning DTE to respond to the Jam message by detecting collision and terminating the transmission prior to the expiration of the timer. The upper limit on the false_carrier_timer prevents the Jam message from exceeding the maximum fragment size.

The combination of the ipg_timer, idle_timer, and valid_carrier_timer filter transient activity that can occur on a link during power cycles or mechanical connection. The duration of the ipg_timer is greater than two-thirds of the minimum IPG, and less than the minimum IPG less some shrinkage. The idle_timer is specified as approximately 320 μs based upon empirical data on such transients. The valid_carrier_timer duration is less than the duration of a minimum valid carrier event, but long enough to filter most spurious carrier events (note that there can be no valid collision fragments on an isolated link in a single repeater topology). The range of the valid_carrier_timer is specified to be the same as the false_carrier_timer range for the convenience of implementations.

41.2.1.5.2 Speed handling

If the PHY has the capability of detecting speeds other than 1000 Mb/s, then the repeater set shall have the capability of blocking the flow of non-1000 Mb/s signals. The incorporation of 1000 Mb/s and 100 Mb/s or 10 Mb/s repeater functionality within a single repeater set is beyond the scope of this standard.

41.2.1.6 Partition functional requirements

It is desirable that the repeater set protect the network from some fault conditions that would disrupt network communications. A potentially likely cause of this condition could be due to a cable fault.

The repeater unit shall provide a self-interrupt capability at each port, as described in Figure 41–4, to prevent a faulty segment's carrier activity from propagating through the network. The repeater unit shall count consecutive collision events at each port. The count shall be incremented on each transmission that suffers a collision and shall be reset on a successful transmission or reception. If this count equals or exceeds the value CELimit (see 41.2.2.1.1), the Partition condition shall be detected. In addition, the partition condition shall be detected due to a carrier event of duration in excess of jabber_timer in which a collision has occurred.

Upon detection of Partition at a port, the repeater unit shall perform the following:

a) Inhibit sending further input messages from that port to the repeater unit.
b) Continue to output messages to that port from the repeater unit.
c) Continue to monitor activity on that port.

The repeater unit shall reset the Partition function at the port when one of the following conditions is met:

— On power-up reset.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 255 of 379

— The repeater has detected activity on the port for more than the number of bits specified for no_collision_timer (see 41.2.2.1.4) without incurring a collision.
— The repeater has transmitted on the port for more than the number of bits specified for no_collision_timer (see 41.2.2.1.4) without incurring a collision.

The no_collision_timer duration is longer than the maximum round-trip latency from a repeater to a DTE (maximum time required for a repeater to detect a collision), and less than the minimum valid carrier event duration (slot time plus header_size minus preamble shrinkage).

41.2.1.7 Receive jabber functional requirements

The repeater unit shall provide a self-interrupt capability at each port, as described in Figure 41–3, to prevent an illegally long reception of data from propagating through the network. The repeater unit shall provide a window of duration jabber_timer bit times (see 41.2.2.1.4) during which the input messages from a port may be passed on to other repeater unit functions. If a reception exceeds this duration, the jabber condition shall be detected.

Upon detection of the jabber condition at a port, the repeater unit shall perform the following:

a) Inhibit sending further input messages from that port to the repeater unit.
b) Inhibit sending further output messages to that port from the repeater unit.

The repeater shall reset the Jabber function at the port, and re-enable data transmission and reception, when either one of the following conditions is met:

— On power-up reset.
— When carrier is no longer detected at that port.

The lower bound of the jabber_timer is longer than the carrier event of a maximum length burst. The upper bound is large enough to permit a wide variety of implementations.

41.2.2 Detailed repeater functions and state diagrams

A precise algorithmic definition is given in this subclause, providing a complete procedural model for the operation of a repeater, in the form of state diagrams. Note that whenever there is any apparent ambiguity concerning the definition of repeater operation, the state diagrams should be consulted for the definitive statement.

The model presented in this subclause is intended as a primary specification of the functions to be provided by any repeater unit. It is important to distinguish, however, between the model and a real implementation. The model is optimized for simplicity and clarity of presentation, while any realistic implementation should place heavier emphasis on such constraints as efficiency and suitability to a particular implementation technology.

It is the functional behavior of any repeater set implementation that is expected to match the standard, not the internal structure. The internal details of the procedural model are useful only to the extent that they help specify the external behavior clearly and precisely. For example, the model uses a separate Receive Port Jabber state diagram for each port. However, in actual implementation, the hardware may be shared.

The notation used in the state diagram follows the conventions of 1.2.1. Note that transitions shown without source states are evaluated at the completion of every state and take precedence over other transition conditions.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 256 of 379

41.2.2.1 State diagram variables

41.2.2.1.1 Constants

CELimit

The number of consecutive Collision Events that must occur before a segment is partitioned.

Values: Positive integer greater than 60.

FCELimit

The number of consecutive False Carrier Events that must occur before a segment is isolated.

Value: 2.

41.2.2.1.2 Variables

begin

The Interprocess flag controlling state diagram initialization values.

Values: true
false

CRS(X), RXD(X), RX_DV(X), RX_ER(X), TXD(X), TX_EN(X), TX_ER(X)

GMII signals received from or sent to the PHY at port X (see Clause 35). The repeater_mode variable in the PHY is set to ensure that the CRS(X) signal is asserted in response to receive activity only.

RXERROR(X)

A combination of the GMII signal encodings indicating that the PHY has detected a Data Error, Carrier Extend Error, or False Carrier Error.

Value: RXERROR(X) ⇐ ((RX_ER(X) = true) *
((RX_DV(X) = true) + (RXD(X) = FalseCarrier) + (RXD(X) = CarrierExtendError)))

TX(X)

A combination of the GMII signal encodings indicating that port X is transmitting a frame.

Value: TX(X) ⇐ ((TX_EN(X) = true) + (TX_ER(X) = true))

isolate(X)

Flag from Carrier Integrity state diagram for port X, which determines whether a port should be enabled or disabled.

Values: true; the Carrier Integrity Monitor has determined the port should be disabled.
false; the Carrier Integrity Monitor has determined the port should be enabled.

force_jam(X)

Flag from Carrier Integrity state diagram for port X, which causes the Repeater Unit to enter the Jam state.

Values: true; the port is in the False Carrier state.
false; the port is not in the False Carrier state.

jabber(X)

Flag from Receive Timer state diagram for port X which indicates that the port has received excessive length activity.

Values: true; port has exceeded the continuous activity limit.
false; port has not exceeded the continuous activity limit.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 257 of 379

link_status(X)

Indication from the Auto-Negotiation process (Clauses 28 and 37) that Auto-Negotiation has completed and the priority resolution function has determined that the link will be operated in half duplex mode.

Values: OK; the link is operational in half duplex mode.
FAIL; the link is not operational in half duplex mode.

partition(X)

Flag from Partition state diagram for port X, which determines whether a port receive path should be enabled or disabled.

Values: true; port has exceeded the consecutive collision limit.
false; port has not exceeded the consecutive collision limit.

41.2.2.1.3 Functions

port(Test)

A function that returns the designation of a port passing the test condition. For example, port(CRS = true) returns the designation: X for a port for which CRS is asserted. If multiple ports meet the test condition, the Port function will be assigned one and only one of the acceptable values.

41.2.2.1.4 Timers

All timers operate in the same fashion. A timer is reset and starts timing upon entering a state where "start x_timer" is asserted. At time "x" after the timer has been started, "x_timer_done" is asserted and remains asserted until the timer is reset. At all other times, "x_timer_not_done" is asserted.

When entering a state where "start x_timer" is asserted, the timer is reset and restarted even if the entered state is the same as the exited state.

The timers used in the repeater state diagrams are defined as follows:

false_carrier_timer

Timer for length of false carrier (41.2.1.5.1) that must be present to set isolate(X) to true. The timer is done when it reaches 3600–4000 BT.

idle_timer

Timer for length of time without carrier activity that must be present to set isolate(X) to false. The timer is done when it reaches 240 000–400 000 BT.

ipg_timer

Timer for length of time without carrier activity that must be present before carrier integrity tests (41.2.1.5.1) are re-enabled. The timer is done when it reaches 64–86 BT.

jabber_timer

Timer for length of carrier which must be present before the Jabber state is entered (41.2.1.7). The timer is done when it reaches 80 000–150 000 BT.

no_collision_timer

Timer for length of packet without collision before partition(X) is set to false (41.2.1.6). The timer is done when it reaches 3600–4144 BT.

valid_carrier_timer

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 258 of 379

Timer for length of valid carrier that must be present to cause isolate(X) to be set to false at the end of the carrier event. The timer is done when it reaches 3600–4000 BT.

41.2.2.1.5 Counters

CE(X)

Consecutive port Collision Event count for port X. Partitioning occurs on a terminal count of CELimit being reached.

Values: Non-negative integers up to a terminal count of CELimit.

FCE(X)

False Carrier Event count for port X. Isolation occurs on a terminal count of FCELimit being reached.

Values: Non-negative integers up to a terminal count of FCELimit.

41.2.2.1.6 Port designation

Ports are referred to by number. Port information is obtained by replacing the X in the desired function with the number of the port of interest. Ports are referred to in general as follows:

X

Generic port designator. When X is used in a state diagram, its value is local to that diagram and not global to the set of state diagrams.

N

Identifies the port that caused the exit from the IDLE or JAM states of Figure 41–2. The value is assigned in the term assignment statement on the transition out of these states (see 1.2.1 for State Diagram Conventions).

ALL

Indicates all repeater ports are to be considered. The test passes when all ports meet the test conditions.

ALLXJIPN

The test passes when all ports, excluding those indicated by J, I, P, or N, meet the test conditions. One or more of the J, I, P, or N indications are used to exclude from the test ports with Jabber = true, Isolate = true, Partition = true, or port N, respectively.

ANY

Indicates all ports are to be considered. The test passes when one or more ports meet the test conditions.

ANYXJIPN

The test passes when one or more ports, excluding those indicated by J, I, P, or N, meet the test conditions. One or more of the J, I, P, or N indications are used to exclude from the test ports with Jabber = true, Isolate = true, Partition = true, or port N, respectively.

ONLY1

Indicates all ports except those with Jabber = true, Isolate = true, or Partition = true are to be considered. The test passes when one and only one port meet the test conditions.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 259 of 379

41.2.2.2 State diagrams

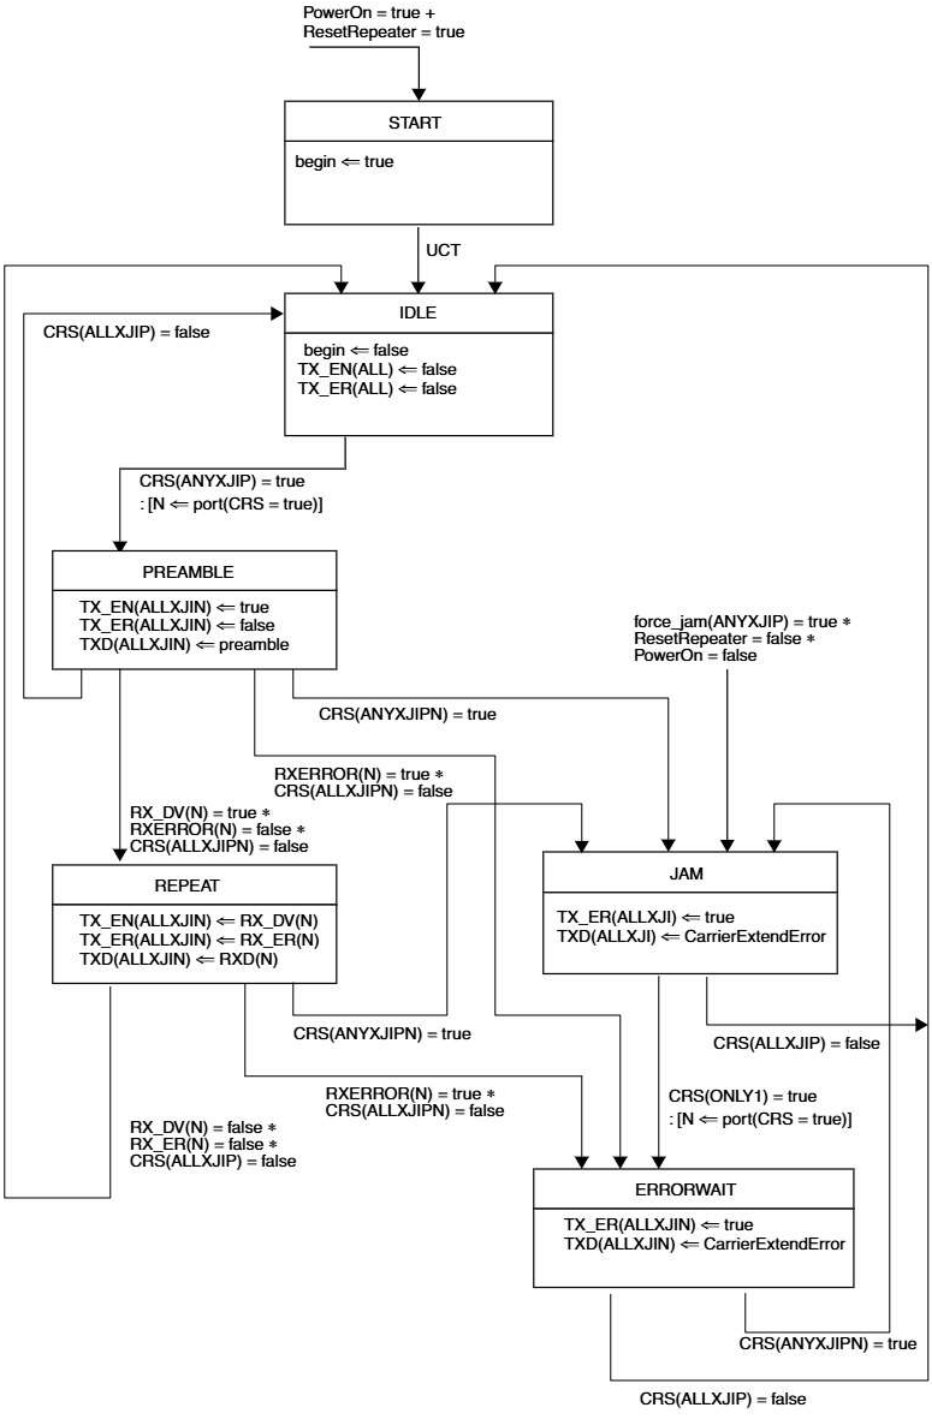


Figure 41–2—Repeater unit state diagram

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 260 of 379

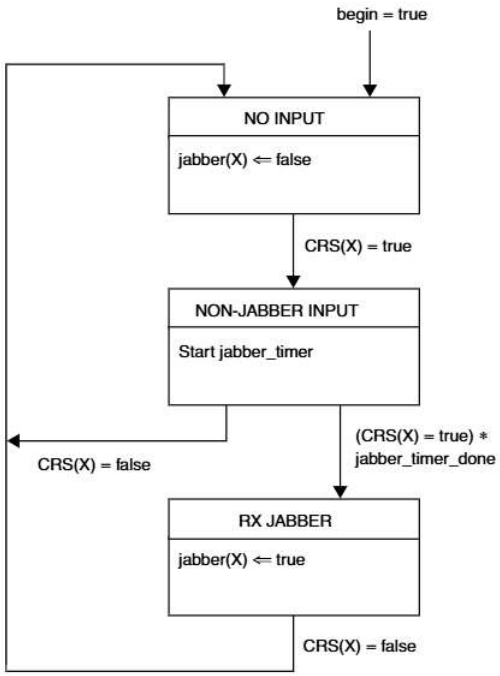


Figure 41–3—Receive timer state diagram for port X

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 261 of 379

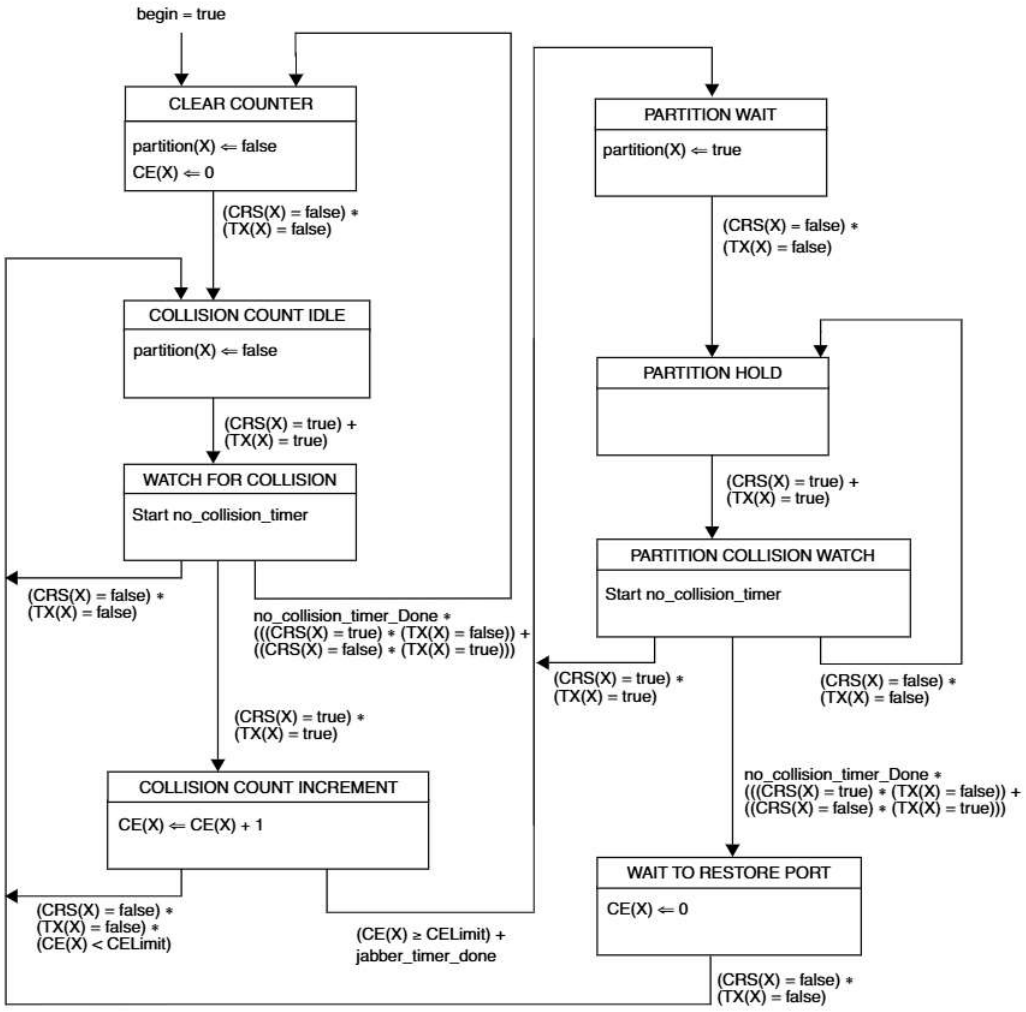


Figure 41–4—Partition state diagram for port X

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 262 of 379

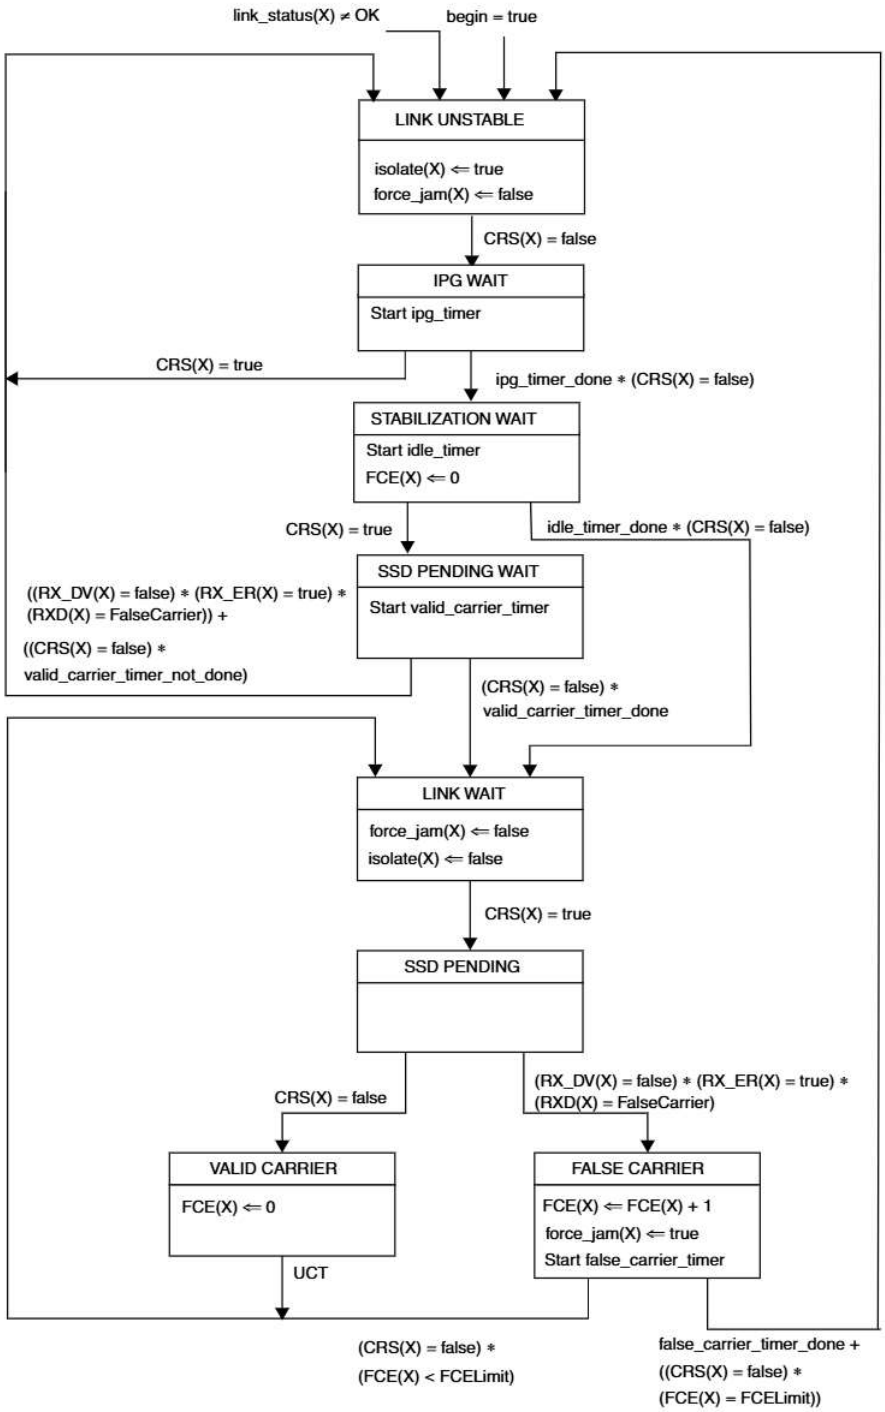


Figure 41–5—Carrier integrity monitor state diagram for port X

Copyright © 2002 IEEE. All rights reserved.

41.3 Repeater electrical specifications

41.3.1 Electrical isolation

Network segments that have different isolation and grounding requirements shall have those requirements provided by the port-to-port isolation of the repeater set.

41.4 Environmental specifications

41.4.1 General safety

All equipment meeting this standard shall conform to IEC 60950: 1991.

41.4.2 Network safety

This subclause sets forth a number of recommendations and guidelines related to safety concerns; the list is neither complete nor does it address all possible safety issues. The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

LAN cable systems described in this subclause are subject to at least four direct electrical safety hazards during their installation and use. These hazards are as follows:

a) Direct contact between LAN components and power, lighting, or communications circuits.
b) Static charge buildup on LAN cables and components.
c) High-energy transients coupled onto the LAN cable system.
d) Voltage potential differences between safety grounds to which the various LAN components are connected.

Such electrical safety hazards must be avoided or appropriately protected against for proper network installation and performance. In addition to provisions for proper handling of these conditions in an operational system, special measures must be taken to ensure that the intended safety features are not negated during installation of a new network or during modification or maintenance of an existing network. Isolation requirements are defined in 41.4.3.

41.4.2.1 Installation

Sound installation practice, as defined by applicable local codes and regulations, shall be followed in every instance in which such practice is applicable.

41.4.2.2 Grounding

The safety ground, or chassis ground for the repeater set, shall be provided through the main ac power cord via the third wire ground as defined by applicable local codes and regulations.

If the MDI connector should provide a shield connection, the shield may be connected to the repeater safety ground. A network segment connected to the repeater set through the MDI may use a shield. If both ends of the network segment have a shielded MDI connector available, then the shield may be grounded at both ends according to local regulations and ISO/IEC 11801: 1995, and as long as the ground potential difference between both ends of the network segment is less than 1 V rms.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 264 of 379

WARNING

It is assumed that the equipment to which the repeater is attached is properly grounded and not left floating nor serviced by a "doubly insulated ac power distribution system." The use of floating or insulated equipment, and the consequent implications for safety, are beyond the scope of this standard.

41.4.2.3 Installation and maintenance guidelines

During installation and maintenance of the cable plant, care should be taken to ensure that uninsulated network cable connectors do not make electrical contact with unintended conductors or ground.

41.4.3 Electrical isolation

There are two electrical power distribution environments to be considered that require different electrical isolation properties:

a) *Environment A*. When a LAN or LAN segment, with all its associated interconnected equipment, is entirely contained within a single low-voltage power distribution system and within a single building.
b) *Environment B*. When a LAN crosses the boundary between separate power distribution systems or the boundary of a single building.

41.4.3.1 Environment A requirements

Attachment of network segments via repeater sets requires electrical isolation of 500 V rms, one-minute withstand, between the segment and the protective ground of the repeater unit.

41.4.3.2 Environment B requirements

The attachment of network segments that cross environment B boundaries requires electrical isolation of 1500 Vrms, one-minute withstand, between each segment and all other attached segments and also the protective ground of the repeater unit.

The requirements for interconnected electrically conducting LAN segments that are partially or fully external to a single building environment may require additional protection against lightning strike hazards. Such requirements are beyond the scope of this standard. It is recommended that the above situation be handled by the use of nonelectrically conducting segments (e.g., fiber optic).

It is assumed that any nonelectrically conducting segments will provide sufficient isolation within that media to satisfy the isolation requirements of environment B.

41.4.4 Reliability

A two-port repeater set shall be designed to provide a mean time between failure (MTBF) of at least 50 000 hours of continuous operation without causing a communications failure among stations attached to the network medium. Repeater sets with more than two ports shall add no more than 3.46×10^{-6} failures per hour for each additional port.

The repeater set electronics should be designed to minimize the probability of component failures within the repeater electronics that prevent communications among other PHYs on the individual segments. Connectors and other passive components comprising the means of connecting the repeater to the cable should be designed to minimize the probability of total network failure.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 265 of 379

41.4.5 Environment

41.4.5.1 Electromagnetic emission

The repeater shall comply with applicable local and national codes for the limitation of electromagnetic interference.

41.4.5.2 Temperature and humidity

The repeater is expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this standard.

It is recommended that manufacturers indicate in the literature associated with the repeater the operating environmental conditions to facilitate selection, installation, and maintenance.

41.5 Repeater labeling

It is required that each repeater (and supporting documentation) shall be labeled in a manner visible to the user with these parameters:

a) Crossover ports appropriate to the respective PHY shall be marked with an X.

Additionally it is recommended that each repeater (and supporting documentation) also be labeled in a manner visible to the user with at least these parameters:

b) Data rate capability in Mb/s
c) Any applicable safety warnings
d) Port type, i.e., 1000BASE-CX, 1000BASE-SX, 1000BASE-LX, and 1000BASE-T
e) Worst-case bit time delays between any two ports appropriate for
    1) Start-of-packet propagation delay
    2) Start-of-collision Jam propagation delay
    3) Cessation-of-collision Jam propagation delay

41.6 Protocol Implementation Conformance Statement (PICS) proforma for Clause 41, Repeater for 1000 Mb/s baseband networks[12]

41.6.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 41, Repeater for 1000 Mb/s baseband networks, 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.

[12]*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.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 266 of 379

41.6.2 Identification

41.6.2.1 Implementation identification

Supplier	
Contact point for enquiries about the PICS	
Implementation Name(s) and Version(s)	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Names(s)	
NOTE 1—Only the first three items are required for all implementations; other information may be completed as appropriate in meeting the requirements for the identification. NOTE 2—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

41.6.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2002®, Clause 41, Repeater for 1000 Mb/s baseband networks
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 802.3-2002®.)	

Date of Statement	

41.6.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*SXP	Repeater supports 1000BASE-SX connections	41.1.2.2		O	Yes [ ] No [ ]
*LXP	Repeater supports 1000BASE-LX connections	41.1.2.2		O	Yes [ ] No [ ]
*CXP	Repeater supports 1000BASE-CX connections	41.1.2.2		O	Yes [ ] No [ ]
*TP	Repeater supports 1000BASE-T connections	41.1.2.2		O	Yes [ ] No [ ]
*PHYS	PHYs capable of detecting non 1000 Mb/s signals	41.2.1.5.2		O	Yes [ ] No [ ]

In addition, the following predicate name is defined for use when different implementations from the set above have common parameters:

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 267 of 379

*XP:SXP or LXP or CXP

41.6.4 PICS proforma tables for the Repeater for 1000 Mb/s baseband networks

41.6.4.1 Compatibility considerations

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	1000BASE-SX port compatible at the MDI	41.1.2.2		SXP:M	Yes [ ] N/A [ ]
CC2	1000BASE-LX port compatible at the MDI	41.1.2.2		LXP:M	Yes [ ] N/A [ ]
CC3	1000BASE-CX port compatible at the MDI	41.1.2.2		CXP:M	Yes [ ] N/A [ ]
CC4	1000BASE-T port compatible at the MDI	41.1.2.2		TP:M	Yes [ ] N/A [ ]
CC5	Internal segment compatibility	41.1.2.2.1	Internal port meets Clause 30 when repeater management implemented	M	Yes [ ]

41.6.4.2 Repeater functions

Item	Feature	Subclause	Value/Comment	Status	Support
RF1	Signal Restoration	41.2.1		M	Yes [ ]
RF2	Data Handling	41.2.1		M	Yes [ ]
RF3	Received Event Handling	41.2.1		M	Yes [ ]
RF4	Collision Handling	41.2.1		M	Yes [ ]
RF5	Error Handling	41.2.1		M	Yes [ ]
RF6	Partition	41.2.1		M	Yes [ ]
RF7	Received Jabber	41.2.1		M	Yes [ ]
RF8	Transmit	41.2.1		M	Yes [ ]
RF9	Receive	41.2.1		M	Yes [ ]

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 268 of 379

41.6.4.3 Signal restoration function

Item	Feature	Subclause	Value/Comment	Status	Support
SR1	Output amplitude as required by 1000BASE-SX	41.2.1.1.1		SXP:M	Yes [ ] N/A [ ]
SR2	Output amplitude as required by 1000BASE-LX	41.2.1.1.1		LXP:M	Yes [ ] N/A [ ]
SR3	Output amplitude as required by 1000BASE-CX	41.2.1.1.1		CXP:M	Yes [ ] N/A [ ]
SR4	Output amplitude as required by 1000BASE-T	41.2.1.1.1		TP:M	Yes [ ] N/A [ ]
SR5	Output signal wave-shape as required by 1000BASE-SX	41.2.1.1.2		SXP:M	Yes [ ] N/A [ ]
SR6	Output signal wave-shape as required by 1000BASE-LX	41.2.1.1.2		LXP:M	Yes [ ] N/A [ ]
SR7	Output signal wave-shape as required by 1000BASE-CX	41.2.1.1.2		CXP:M	Yes [ ] N/A [ ]
SR8	Output signal wave-shape as required by 1000BASE-T	41.2.1.1.2		TP:M	Yes [ ] N/A [ ]
SR9	Output data timing as required by 1000BASE-SX	41.2.1.1.3		SXP:M	Yes [ ] N/A [ ]
SR10	Output data timing as required by 1000BASE-LX	41.2.1.1.3		LXP:M	Yes [ ] N/A [ ]
SR11	Output data timing as required by 1000BASE-CX	41.2.1.1.3		CXP:M	Yes [ ] N/A [ ]
SR12	Output data timing as required by 1000BASE-T	41.2.1.1.3		TP:M	Yes [ ] N/A [ ]

41.6.4.4 Data-Handling function

Item	Feature	Subclause	Value/Comment	Status	Support
DH1	Data frames forwarded to all ports except receiving port	41.2.1.2.1		M	Yes [ ]
DH2	Code Violations forwarded to all transmitting ports	41.2.1.2.2		M	Yes [ ]
DH3	Received Code Violation forwarded as code-group explicitly indicating received error	41.2.1.2.2		M	Yes [ ]
DH4	Code element substitution for remainder of packet after received Code Violation	41.2.1.2.2		M	Yes [ ]

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 269 of 379

41.6.4.5 Receive Event-Handling function

Item	Feature	Subclause	Value/Comment	Status	Support
RE1	Detect all received events	41.2.1.3.1		M	Yes [ ]
RE2	Repeat all received signals	41.2.1.3.1		M	Yes [ ]
RE3	Preamble repeated as required	41.2.1.3.2		M	Yes [ ]
RE4	Start-of-packet propagation delay	41.2.1.3.3	SOP + SOJ ≤ 976 BT	M	Yes [ ]
RE5	Start-of-packet variability	41.2.1.3.4	SOP variation ≤ 16 BT	M	Yes [ ]
RE6	PHY repeater_mode variable	41.2.1.3.1	Shall be set to ensure CRS signal not asserted in response to transmit activity	M	Yes [ ]
RE7	Output preamble variation	41.2.1.3.2	Variation between received and transmitted preamble ≤ 8 BT	M	Yes [ ]

41.6.4.6 Collision-Handling function

Item	Feature	Subclause	Value/Comment	Status	Support
CO1	Collision Detection	41.2.1.4.1	Receive event on more than one port	M	Yes [ ]
CO2	Jam Generation	41.2.1.4.2	Transmit Jam message while collision is detected	M	Yes [ ]
CO3	Collision-Jam Propagation delay	41.2.1.4.3	SOP + SOJ ≤ 976 BT	M	Yes [ ]
CO4	Cessation of Collision Propagation delay	41.2.1.4.4	EOJ ≤ SOP	M	Yes [ ]

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 270 of 379

41.6.4.7 Error-Handling function

Item	Feature	Subclause	Value/Comment	Status	Support
EH1	Carrier Integrity function implementation	41.2.1.5.1	Self-interrupt of data reception	M	Yes [ ]
EH2	False Carrier Event count for Link Unstable detection	41.2.1.5.1	False Carrier Event count equals FCELimit	M	Yes [ ]
EH3	False carrier count reset	41.2.1.5.1	Count reset on valid carrier	M	Yes [ ]
EH4	False carrier timer for Link Unstable detection	41.2.1.5.1	False carrier of length in excess of false_carrier_timer	M	Yes [ ]
EH5	Jam message duration	41.2.1.5.1	Equals duration of false carrier event, but not greater than duration of false_carrier_timer	M	Yes [ ]
EH6	Link Unstable detection	41.2.1.5.1	False Carrier Event count equals FCELimit or False carrier exceeds the false_carrier_timer or power-up reset	M	Yes [ ]
EH7	Messages sent to repeater unit in Link Unstable state	41.2.1.5.1	Inhibited sending messages to repeater unit	M	Yes [ ]
EH8	Messages sent from repeater unit in Link Unstable state	41.2.1.5.1	Inhibited sending output messages	M	Yes [ ]
EH9	Monitoring activity on a port in Link Unstable state	41.2.1.5.1	Continue monitoring activity at that port	M	Yes [ ]
EH10	Reset of Link Unstable state	41.2.1.5.1	No activity for more than ipg_timer plus idle_timer or Valid carrier event of duration greater than valid_carrier_timer preceded by Idle of duration greater than ipg_timer	M	Yes [ ]
EH11	Block flow of non-1000 Mb/s signals	41.2.1.5.2		M	Yes [ ]

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 271 of 379

41.6.4.8 Partition function

Item	Feature	Subclause	Value/Comment	Status	Support
PA1	Partition function implementation	41.2.1.6	Self-interrupt of data reception	M	Yes [ ]
PA2	Consecutive Collision Event count for entry into partition state	41.2.1.6	Consecutive Collision Event count equals or exceeds CELimit	M	Yes [ ]
PA3	Excessive receive duration with collision for entry into partition state.	41.2.1.6	Reception duration in excess of jabber_timer with collision	M	Yes [ ]
PA4	Consecutive Collision Event counter incrementing	41.2.1.6	Count incremented on each transmission that suffers a collision	M	Yes [ ]
PA5	Consecutive Collision Event counter reset	41.2.1.6	Count reset on successful transmission or reception	M	Yes [ ]
PA6	Messages sent to repeater unit in Partition state	41.2.1.6	Inhibited sending messages to repeater unit	M	Yes [ ]
PA7	Messages sent from repeater unit in Partition state	41.2.1.6	Continue sending output messages	M	Yes [ ]
PA8	Monitoring activity on a port in Partition state	41.2.1.6	Continue monitoring activity at that port	M	Yes [ ]
PA9	Reset of Partition state	41.2.1.6	Power-up reset or transmitting or detecting activity for greater than duration no_collision_timer without a collision	M	Yes [ ]

41.6.4.9 Receive Jabber function

Item	Feature	Subclause	Value/Comment	Status	Support
RJ1	Receive Jabber function implementation	41.2.1.7	Self-interrupt of data reception	M	Yes [ ]
RJ2	Excessive receive duration timer for Receive Jabber detection	41.2.1.7	Reception duration in excess of jabber_timer	M	Yes [ ]
RJ3	Messages sent to repeater unit in Receive Jabber state	41.2.1.7	Inhibit sending input messages to repeater unit	M	Yes [ ]
RJ4	Messages sent from repeater unit in Receive Jabber state	41.2.1.7	Inhibit sending output messages	M	Yes [ ]
RJ5	Reset of Receive Jabber state	41.2.1.7	Power-up reset or Carrier no longer detected	M	Yes [ ]

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 272 of 379

41.6.4.10 Repeater state diagrams

Item	Feature	Subclause	Value/Comment	Status	Support
SD1	Repeater unit state diagram	41.2.2.2	Meets the requirements of Figure 41–2	M	Yes [ ]
SD2	Receive timer for port X state diagram	41.2.2.2	Meets the requirements of Figure 41–3	M	Yes [ ]
SD3	Repeater partition state diagram for port X	41.2.2.2	Meets the requirements of Figure 41–4	M	Yes [ ]
SD4	Carrier integrity monitor for port X state diagram	41.2.2.2	Meets the requirements of Figure 41–5	M	Yes [ ]

41.6.4.11 Repeater electrical

Item	Feature	Subclause	Value/Comment	Status	Support
EL1	Port-to-port isolation	41.3.1	Satisfies isolation and grounding requirements for attached network segments	M	Yes [ ]
EL2	Safety	41.4.1	IEC 60950: 1991	M	Yes [ ]
EL3	Installation practices	41.4.2.1	Sound, as defined by local code and regulations	M	Yes [ ]
EL4	Grounding	41.4.2.2	Chassis ground provided through ac mains cord	M	Yes [ ]
EL5	Two-port repeater set MTBF	41.4.4	At least 50 000 hours	M	Yes [ ]
EL6	Additional port effect on MTBF	41.4.4	No more than 3.46×10^{-6} increase in failures per hour	M	Yes [ ]
EL7	Electromagnetic interference	41.4.5.1	Comply with local or national codes	M	Yes [ ]

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 273 of 379

41.6.4.12 Repeater labeling

Item	Feature	Subclause	Value/Comment	Status	Support
LB1	Crossover ports	41.5	Marked with an X	M	Yes [ ]
LB2	Data Rate	41.5	1000 Mb/s	O	Yes [ ] No [ ]
LB3	Safety warnings	41.5	Any applicable	O	Yes [ ] No [ ]
LB4	Port Types	41.5	1000BASE-SX, 1000BASE-LX, 1000BASE-CX, or 1000BASE-T	O	Yes [ ] No [ ]
LB5	Worse-case start-of-packet propagation delay	41.5	Value in bit times (BT)	O	Yes [ ] No [ ]
LB6	Worse-case start-of-collision-Jam propagation delay	41.5	Value in BT	O	Yes [ ] No [ ]
LB7	Worse-case Cessation-of-Collision Jam propagation delay	41.5	Value in BT	O	Yes [ ] No [ ]

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 274 of 379

42. System considerations for multisegment 1000 Mb/s networks

42.1 Overview

This clause provides information on building 1000 Mb/s networks. The 1000 Mb/s technology is designed to be deployed in both homogenous 1000 Mb/s networks and 10/100/1000 Mb/s mixed networks using bridges and/or routers. Network topologies can be developed within a single 1000 Mb/s collision domain, but maximum flexibility is achieved by designing multiple collision domain networks that are joined by bridges and/or routers configured to provide a range of service levels to DTEs. For example, a combined 1000BASE-T/100BASE-T/10BASE-T system built with repeaters and bridges can deliver dedicated or shared service to DTEs at 1000 Mb/s, 100 Mb/s, or 10 Mb/s.

Linking multiple collision domains with bridges maximizes flexibility. Bridged topology designs can provide single data rate (Figure 42–1) or multiple data rate (Figure 42–2) services.

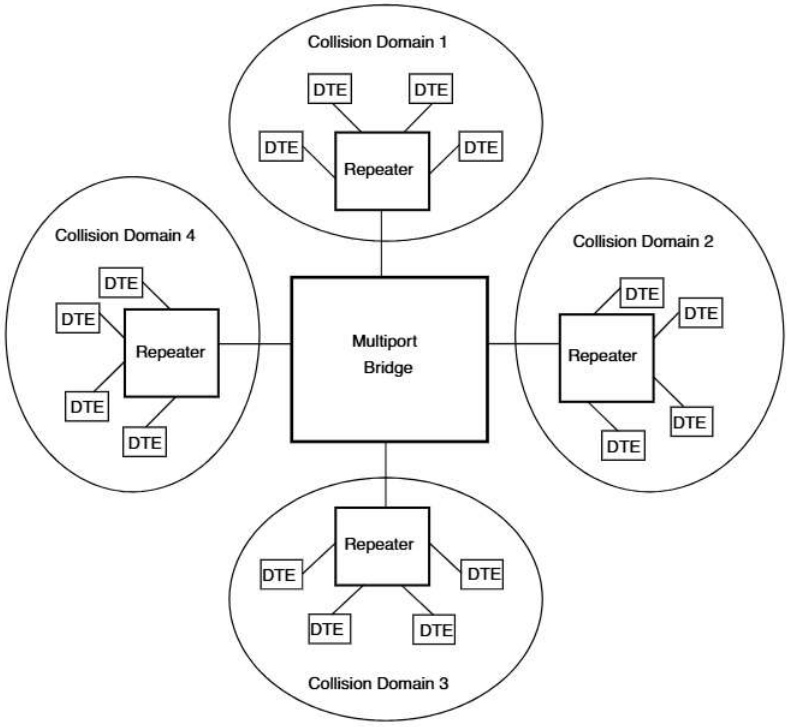


Figure 42–1—1000 Mb/s multiple collision domain topology using multiport bridge

Individual collision domains can be linked by single devices (as shown in Figure 42–1 and Figure 42–2) or by multiple devices from any of several transmission systems. The design of multiple-collision-domain networks is governed by the rules defining each of the transmission systems incorporated into the design.

The design of shared bandwidth 10 Mb/s collision domains is defined in Clause 13; the design of shared bandwidth 100 Mb/s CSMA/CD collision domains is defined in Clause 29; the design of shared bandwidth 1000 Mb/s CSMA/CD collision domains is defined in the following subclauses.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 275 of 379

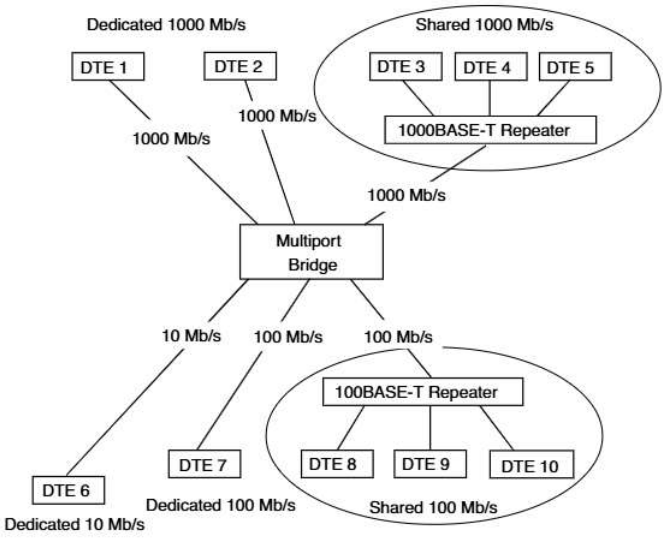


Figure 42–2—Multiple data rate, multiple collision domain topology using multiport bridge

42.1.1 Single collision domain multisegment networks

This clause provides information on building 1000 Mb/s CSMA/CD multisegment networks within a single collision domain. The proper operation of a CSMA/CD network requires the physical size of the collision domain to be limited in order to meet the round-trip propagation delay requirements of 4.2.3.2.3 and 4.4.2.1, and requires the number of repeaters to be limited to one so as not to exceed the InterFrameGap shrinkage noted in 4.4.2.4.

This clause provides two network models. Transmission System Model 1 is a set of configurations that have been validated under conservative rules and have been qualified as meeting the requirements set forth above. Transmission System Model 2 is a set of calculation aids that allow those configuring a network to test a proposed configuration against a simple set of criteria that allows it to be qualified. Transmission System Model 2 validates an additional broad set of topologies that are fully functional and do not fit within the simpler, but more restrictive rules of Model 1.

The physical size of a CSMA/CD network is limited by the characteristics of individual network components. These characteristics include the following:

a) Media lengths and their associated propagation time delay.
b) Delay of repeater units (start-up, steady-state, and end of event).
c) Delay of MAUs and PHYs (start-up, steady-state, and end of event).
d) Interpacket gap shrinkage due to repeater units.
e) Delays within the DTE associated with the CSMA/CD access method.
f) Collision detect and deassertion times associated with the MAUs and PHYs.

Table 42–1 summarizes the delays, measured in Bit Times (BTs), for 1000 Mb/s media segments.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 276 of 379

Table 42–1—Delays for network media segments Model 1

Media type	Maximum number of PHYs per segment	Maximum segment length (m)	Maximum medium round-trip delay per segment (BT)
Category 5 UTP Link Segment (1000BASE-T)	2	100	1112
Shielded Jumper Cable Link Segment (1000BASE-CX)	2	25	253
Optical Fiber Link Segment (1000BASE-SX, 1000BASE-LX)	2	316[a]	3192

[a]May be limited by the maximum transmission distance of the link.

42.1.2 Repeater usage

Repeaters are the means used to connect segments of a network medium together into a single collision domain. Different physical signaling systems (e.g., 1000BASE-CX, 1000BASE-SX, 1000BASE-LX, 1000BASE-T) can be joined into a common collision domain using a repeater. Bridges can also be used to connect different signaling systems; however, if a bridge is so used, each LAN connected to the bridge will comprise a separate collision domain.

42.2 Transmission System Model 1

The following network topology constraints apply to networks using Transmission System Model 1.

a) Single repeater topology maximum.
b) Link distances not to exceed the lesser of 316 m or the segment lengths as shown in Table 42–1.

42.3 Transmission System Model 2

Transmission System Model 2 is a single repeater topology with the physical size limited primarily by round-trip collision delay. A network configuration must be validated against collision delay using a network model for a 1000 Mb/s collision domain. The modeling process is quite straightforward and can easily be done either manually or with a spreadsheet.

The model proposed here is derived from the one presented in 13.4. Modifications have been made to accommodate adjustments for DTE, repeater, and cable speeds.

For a network consisting of two 1000BASE-T DTEs as shown in Figure 42–3, a crossover connection may be required if the auto-crossover function is not implemented. See 40.4 and 40.8.

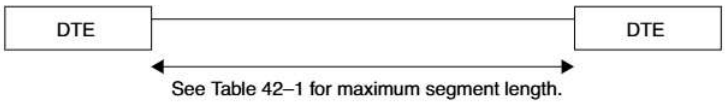


Figure 42–3—Model 1: Two DTEs, no repeater

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 277 of 379

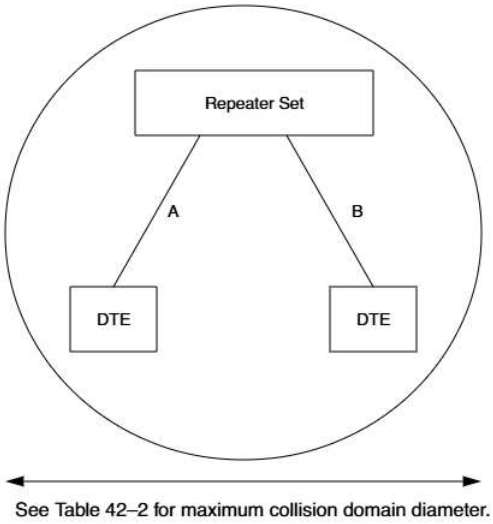


Figure 42–4—Model 1: Single repeater

Table 42–2—Maximum Model 1 collision domain diameter[a]

Configuration	Category 5 UTP (T)	Shielded Jumper Cable (CX)	Optical Fiber (SX/LX)	Mixed Category 5 and Fiber (T and SX/LX)	Mixed Shielded Jumper and Fiber (CX and SX/LX)
DTE-DTE (see Figure 42–3)	100	25	316[b]	NA	NA
One repeater (see Figure 42–4)	200	50	220	210[c]	220[d]

[a]In meters.
[b]May be limited by the maximum transmission distance of the link.
[c]Assumes 100 m of Category 5 UTP and one Optical Fiber link of 110 m.
[d]Assumes 25 m of Shielded Jumper Cable and one Optical Fiber link of 195 m.

42.3.1 Round-trip collision delay

For a network configuration to be valid, it must be possible for any two DTEs on the network to properly arbitrate for the network. When two or more stations attempt to transmit within a slot time interval, each station must be notified of the contention by the returned "collision" signal within the "collision window" (see 4.1.2.2). Additionally, the maximum length fragment created on a 1000 Mb/s network must contain less than 512 bytes after the Start Frame Delimiter (SFD). These requirements limit the physical diameter (maximum distance between DTEs) of a network. The maximum round-trip delay must be qualified between all pairs of DTEs in the network. In practice this means that the qualification must be done between those that, by inspection of the topology, are candidates for the longest delay. The following network modeling methodology is provided to assist that calculation.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 278 of 379

42.3.1.1 Worst-case path delay value (PDV) selection

The worst-case path through a network to be validated is identified by examination of aggregate DTE delays, cable delays, and repeater delay. The worst case consists of the path between the two DTEs at opposite ends of the network that have the longest round-trip time. Figure 42–5 shows a schematic representation of a one-repeater path.

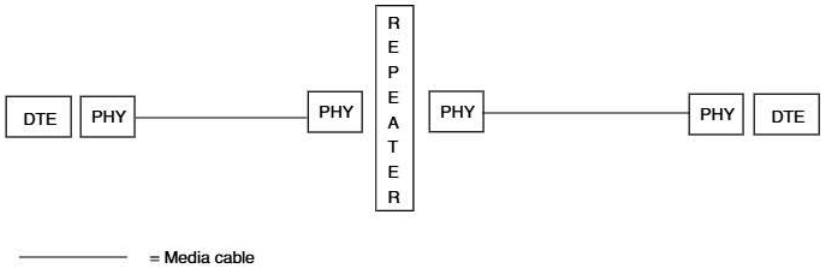


Figure 42–5——System Model 2: Single repeater

42.3.1.2 Worst-case PDV calculation

Once a set of paths is chosen for calculation, each is checked for validity against the following formula:

PDV = ∑link delays (LSDV) + repeater delay + DTE delays + safety margin

Values for the formula variables are determined by the following method:

a) Determine the delay for each link segment (Link Segment Delay Value, or LSDV), using the formula

LSDV=2 (for round-trip delay) × segment length × cable delay for this segment

NOTE 1—Length is the sum of the cable lengths between the PHY interfaces at the repeater and PHY interfaces at the farthest DTE. All measurements are in meters.

NOTE 2—Cable delay is the delay specified by the manufacturer or the maximum value for the type of cable used as shown in Table 42–3. For this calculation, cable delay must be specified in bit times per meter (BT/m). Table 42–4 can be used to convert values specified relative to the speed of light (%c) or nanoseconds per meter (ns/m).

NOTE 3—When actual cable lengths or propagation delays are not known, use the Max delay in bit times as specified in Table 42–3 for copper cables. Delays for fiber should be calculated, as the value found in Table 42–3 will be too large for most applications.

NOTE 4—The value found in Table 42–3 for Shielded Jumper Cable is the maximum delay for cable with solid dielectric. Cables with foam dielectric may have a significantly smaller delay.

b) Sum together the LSDVs for all segments in the path.

c) Determine the delay for the repeater. If model-specific data is not available from the manufacturer, enter the appropriate default value from Table 42–3.

d) Use the DTE delay value shown in Table 42–3 unless your equipment manufacturer defines a different value. If the manufacturer's supplied values are used, the DTE delays of both ends of the worst-case path should be summed together.

e) Decide on appropriate safety margin—0 to 40 bit times—for the PDV calculation. Safety margin is used to provide additional margin to accommodate unanticipated delay elements, such as extra-long connecting cable runs between wall jacks and DTEs. (A safety margin of 32 BT is recommended.)

f) Insert the values obtained through the calculations above into the following formula to calculate the PDV. (Some configurations may not use all the elements of the formula.)

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 279 of 379

PDV = ∑link delays (LSDV) + repeater delay + DTE delay + safety margin

g) If the PDV is less than 4096, the path is qualified in terms of worst-case delay.
h) Late collisions and/or CRC errors may be indications that path delays exceed 4096 BT.

Table 42–3—Network component delays, Transmission System Model 2

Component	Round-trip delay in bit times per meter (BT/m)	Maximum round-trip delay in bit times (BT)
Two DTEs		864
Category 5 UTP Cable segment	11.12	1112 (100 m)
Shielded Jumper Cable segment	10.10	253 (25 m)
Optical Fiber Cable segment	10.10	1111 (110 m)
Repeater		976

Table 42–4—Conversion table for cable delays

Speed relative to c	ns/m	BT/m
0.4	8.34	8.34
0.5	6.67	6.67
0.51	6.54	6.54
0.52	6.41	6.41
0.53	6.29	6.29
0.54	6.18	6.18
0.55	6.06	6.06
0.56	5.96	5.96
0.57	5.85	5.85
0.58	5.75	5.75
0.5852	5.70	5.70
0.59	5.65	5.65
0.6	5.56	5.56
0.61	5.47	5.47
0.62	5.38	5.38
0.63	5.29	5.29

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 280 of 379

Table 42–4—Conversion table for cable delays *(continued)*

Speed relative to c	ns/m	BT/m
0.64	5.21	5.21
0.65	5.13	5.13
0.654	5.10	5.10
0.66	5.05	5.05
0.666	5.01	5.01
0.67	4.98	4.98
0.68	4.91	4.91
0.69	4.83	4.83
0.7	4.77	4.77
0.8	4.17	4.17
0.9	3.71	3.71

42.4 Full duplex 1000 Mb/s topology limitations

Unlike half duplex CSMA/CD networks, the physical size of full duplex 1000 Mb/s networks is not limited by the round-trip collision propagation delay. Instead, the maximum link length between DTEs is limited only by the signal transmission characteristics of the specific link.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 281 of 379

43. Link Aggregation

43.1 Overview

This clause defines an optional Link Aggregation sublayer for use with CSMA/CD MACs. Link Aggregation allows one or more links to be aggregated together to form a Link Aggregation Group, such that a MAC Client can treat the Link Aggregation Group as if it were a single link. To this end, it specifies the establishment of DTE to DTE logical links, consisting of N parallel instances of full duplex point-to-point links operating at the same data rate.

43.1.1 Terminology

In this clause, unless otherwise noted, the term *link* refers to an *Aggregation Link* and the term *port* refers to *Aggregation Port,* as defined in 1.4. This allows for better readability of this clause while avoiding conflicting use of these terms in other clauses of this standard. Similarly, the term *Key* when used in this clause is synonymous with *Aggregation Key,* and the term *System* is synonymous with *Aggregation System*.

43.1.2 Goals and objectives

Link Aggregation, as specified in this clause, provides the following:

a) **Increased bandwidth**—The capacity of multiple links is combined into one logical link.
b) **Linearly incremental bandwidth**—Bandwidth can be increased in unit multiples as opposed to the order-of-magnitude increase available through Physical Layer technology options (10 Mb/s, 100 Mb/s, 1000 Mb/s, etc.).
c) **Increased availability**—The failure or replacement of a single link within a Link Aggregation Group need not cause failure from the perspective of a MAC Client.
d) **Load sharing**—MAC Client traffic may be distributed across multiple links.
e) **Automatic configuration**—In the absence of manual overrides, an appropriate set of Link Aggregation Groups is automatically configured, and individual links are allocated to those groups. If a set of links can aggregate, they will aggregate.
f) **Rapid configuration and reconfiguration**—In the event of changes in physical connectivity, Link Aggregation will quickly converge to a new configuration, typically on the order of 1 second or less.
g) **Deterministic behavior**—Depending on the selection algorithm chosen, the configuration can be made to resolve deterministically; i.e., the resulting aggregation can be made independent of the order in which events occur, and be completely determined by the capabilities of the individual links and their physical connectivity.
h) **Low risk of duplication or mis-ordering**—During both steady-state operation and link (re)configuration, there is a high probability that frames are neither duplicated nor mis-ordered.
i) **Support of existing IEEE 802.3® MAC Clients**—No change is required to existing higher-layer protocols or applications to use Link Aggregation.
j) **Backwards compatibility with aggregation-unaware devices**—Links that cannot take part in Link Aggregation—either because of their inherent capabilities, management configuration, or the capabilities of the devices to which they attach—operate as normal, individual IEEE 802.3® links.
k) **Accommodation of differing capabilities and constraints**—Devices with differing hardware and software constraints on Link Aggregation are, to the extent possible, accommodated.
l) **No change to the IEEE 802.3® frame format**—Link aggregation neither adds to, nor changes the contents of frames exchanged between MAC Clients.
m) **Network management support**—The standard specifies appropriate management objects for configuration, monitoring, and control of Link Aggregation.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 282 of 379

Link Aggregation, as specified in this clause, does not support the following:

n) **Multipoint Aggregations**—The mechanisms specified in this clause do not support aggregations among more than two Systems.
o) **Dissimilar MACs**—Link Aggregation is supported only on links using the IEEE 802.3® MAC.
p) **Half-duplex operation**—Link Aggregation is supported only on point-to-point links with MACs operating in full duplex mode.
q) **Operation across multiple data rates**—All links in a Link Aggregation Group operate at the same data rate (e.g., 10 Mb/s, 100 Mb/s, or 1000 Mb/s).

43.1.3 Positioning of Link Aggregation within the IEEE 802.3® architecture

Link Aggregation comprises an optional sublayer between a MAC Client and the MAC (or optional MAC Control sublayer). Figure 43–1 depicts the positioning of the Link Aggregation sublayer in the CSMA/CD layer architecture, and the relationship of that architecture to the Data Link and Physical Layers of the OSI Reference Model. The figure also shows the ability of the Link Aggregation sublayer to combine a number of individual links in order to present a single MAC interface to the MAC Client.

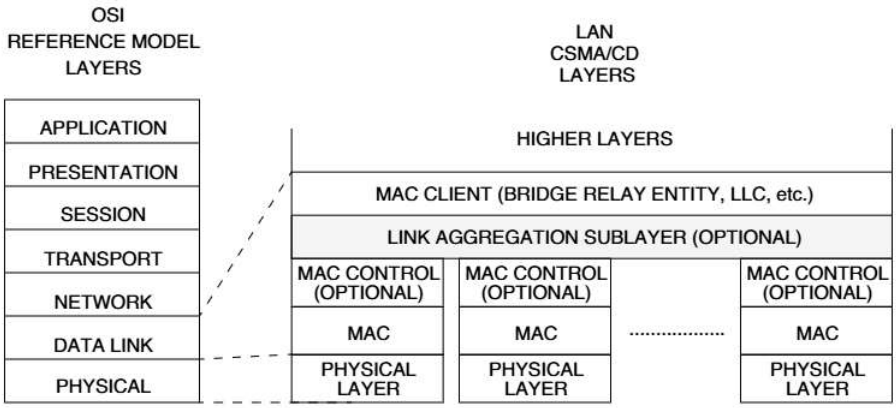


Figure 43–1—Architectural positioning of Link Aggregation sublayer

Figure 43–2 depicts the major blocks that form the Link Aggregation sublayer, and their interrelationships.

It is possible to implement the optional Link Aggregation sublayer for some ports within a System while not implementing it for other ports; i.e., it is not necessary for all ports in a System to be subject to Link Aggregation. A conformant implementation is not required to be able to apply the Link Aggregation sublayer to every port.

43.1.4 State diagram conventions

Many of the functions specified in this clause are presented in state diagram notation. All state diagrams contained in this clause use the notation and conventions defined in 21.5. In the event of a discrepancy between the text description and the state diagram formalization of a function, the state diagrams take precedence.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 283 of 379

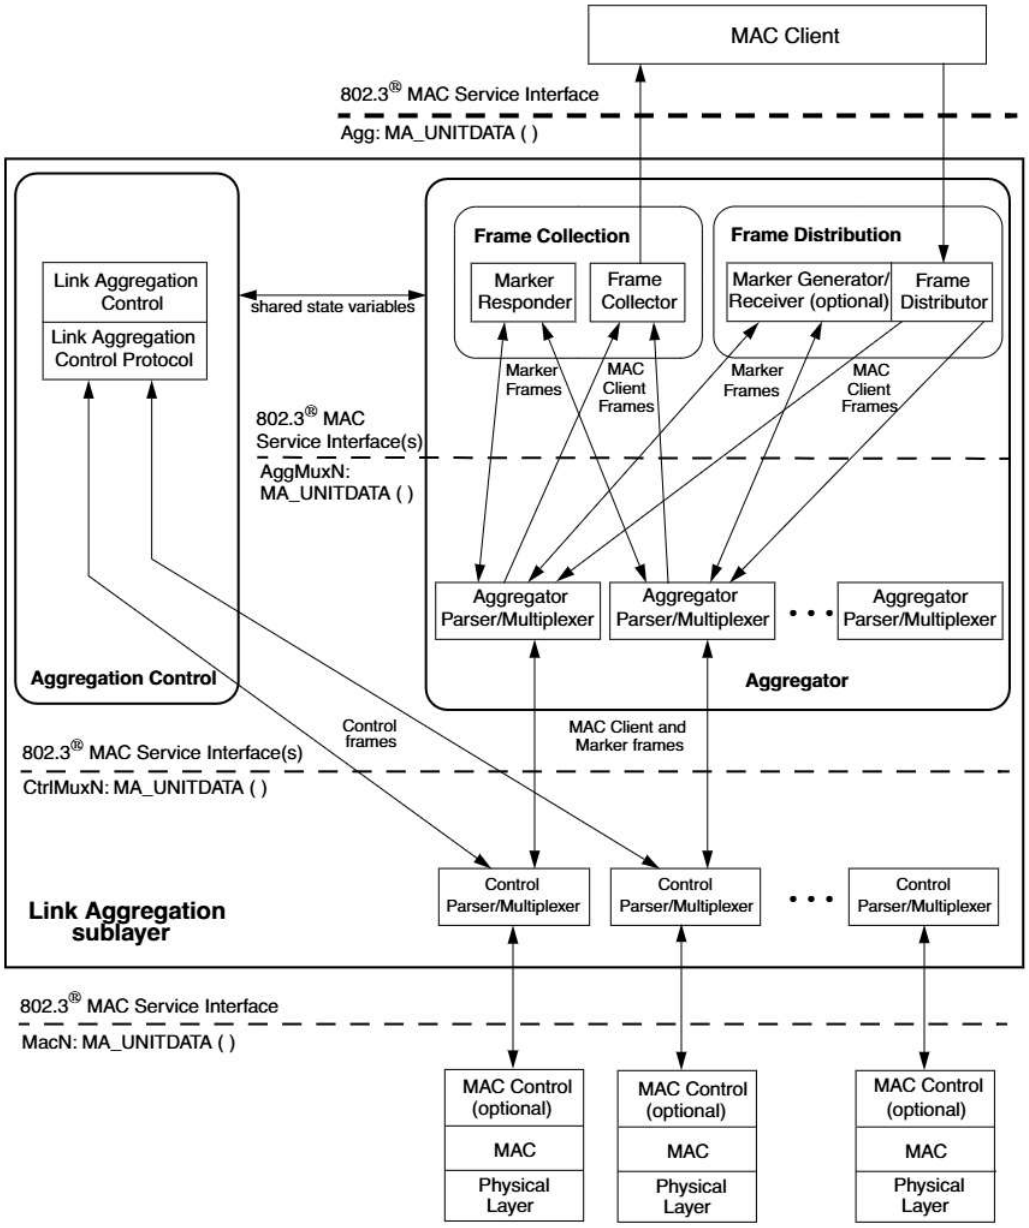


Figure 43–2—Link Aggregation sublayer block diagram

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 284 of 379

43.2 Link Aggregation operation

As depicted in Figure 43–2, the Link Aggregation sublayer comprises the following functions:

a) *Frame Distribution*. This block is responsible for taking frames submitted by the MAC Client and submitting them for transmission on the appropriate port, based on a frame distribution algorithm employed by the Frame Distributor. Frame Distribution also includes an optional *Marker Generator/Receiver* used for the Marker protocol. (See 43.2.4, 43.2.5, and 43.5.)

b) *Frame Collection*. This block is responsible for passing frames received from the various ports to the MAC Client. Frame Collection also includes a *Marker Responder,* used for the Marker protocol. (See 43.2.3 and 43.5.)

c) *Aggregator Parser/Multiplexers*. On transmit, these blocks simply pass frame transmission requests from the Distributor, Marker Generator, and/or Marker Responder to the appropriate port. On receive, these blocks distinguish among Marker Request, Marker Response, and MAC Client PDUs, and pass each to the appropriate entity (Marker Responder, Marker Receiver, and Collector, respectively).

d) *Aggregator.* The combination of Frame Distribution and Collection, along with the Aggregator Parser/Multiplexers, is referred to as the Aggregator.

e) *Aggregation Control*. This block is responsible for the configuration and control of Link Aggregation. It incorporates a *Link Aggregation Control Protocol* (LACP) that can be used for automatic communication of aggregation capabilities between Systems and automatic configuration of Link Aggregation.

f) *Control Parser/Multiplexers*. On transmit, these blocks simply pass frame transmission requests from the Aggregator and Control entities to the appropriate port. On receive, these blocks distinguish Link Aggregation Control PDUs from other frames, passing the LACPDUs to the appropriate sublayer entity, and all other frames to the Aggregator.

43.2.1 Principles of Link Aggregation

Link Aggregation allows a MAC Client to treat a set of one or more ports as if it were a single port. In doing so, it employs the following principles and concepts:

a) A MAC Client communicates with a set of ports through an Aggregator, which presents a standard IEEE 802.3® service interface to the MAC Client. The Aggregator binds to one or more ports within a System.

b) It is the responsibility of the Aggregator to distribute frame transmissions from the MAC Client to the various ports, and to collect received frames from the ports and pass them to the MAC Client transparently.

c) A System may contain multiple Aggregators, serving multiple MAC Clients. A given port will bind to (at most) a single Aggregator at any time. A MAC Client is served by a single Aggregator at a time.

d) The binding of ports to Aggregators within a System is managed by the Link Aggregation Control function for that System, which is responsible for determining which links may be aggregated, aggregating them, binding the ports within the System to an appropriate Aggregator, and monitoring conditions to determine when a change in aggregation is needed.

e) Such determination and binding may be under manual control through direct manipulation of the state variables of Link Aggregation (e.g., Keys) by a network manager. In addition, automatic determination, configuration, binding, and monitoring may occur through the use of a Link Aggregation Control Protocol (LACP). The LACP uses peer exchanges across the links to determine, on an ongoing basis, the aggregation capability of the various links, and continuously provides the maximum level of aggregation capability achievable between a given pair of Systems.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 285 of 379

f) Frame ordering must be maintained for certain sequences of frame exchanges between MAC Clients (known as conversations, see 1.4). The Distributor ensures that all frames of a given conversation are passed to a single port. For any given port, the Collector is required to pass frames to the MAC Client in the order that they are received from that port. The Collector is otherwise free to select frames received from the aggregated ports in any order. Since there are no means for frames to be misordered on a single link, this guarantees that frame ordering is maintained for any conversation.

g) Conversations may be moved among ports within an aggregation, both for load balancing and to maintain availability in the event of link failures.

h) This standard does not impose any particular distribution algorithm on the Distributor. Whatever algorithm is used should be appropriate for the MAC Client being supported.

i) Each port is assigned a unique, globally administered MAC address. This MAC address is used as the source address for frame exchanges that are initiated by entities within the Link Aggregation sublayer itself (i.e., LACP and Marker protocol exchanges).

NOTE—The LACP and Marker protocols use a multicast destination address for all exchanges, and do not impose any requirement for a port to recognize more than one unicast address on received frames.

j) Each Aggregator is assigned a unique, globally administered MAC address; this address is used as the MAC address of the aggregation from the perspective of the MAC Client, both as a source address for transmitted frames and as the destination address for received frames. The MAC address of the Aggregator may be one of the MAC addresses of a port in the associated Link Aggregation Group (see 43.2.10).

43.2.2 Service interfaces

The MAC Client communicates with the Aggregator using the standard service interface specified in Clause 2. Similarly, Link Aggregation communicates internally (between Frame Collection/Distribution, the Aggregator Parser/Multiplexers, the Control Parser/Multiplexers, and Link Aggregation Control) and with its bound ports using the same, standard service interface. No new interlayer service interfaces are defined for Link Aggregation.

Since Link Aggregation uses four instances of the MAC Service Interface, it is necessary to introduce a notation convention so that the reader can be clear as to which interface is being referred to at any given time. A prefix is therefore assigned to each service primitive, indicating which of the four interfaces is being invoked, as depicted in Figure 43–2. The prefixes are as follows:

a) *Agg:*, for primitives issued on the interface between the MAC Client and the Link Aggregation sublayer.

b) *AggMuxN:*, for primitives issued on the interface between Aggregator Parser/Multiplexer N and its internal clients (where N is the port number associated with the Aggregator Parser/Multiplexer).

c) *CtrlMuxN:*, for primitives issued on the interface between Control Parser/Multiplexer N and its internal clients (where N is the port number associated with the Control Parser/Multiplexer).

d) *MacN:*, for primitives issued on the interface between underlying MAC N and its Control Parser/Multiplexer (where N is the port number associated with the underlying MAC).

MAC Clients may generate Agg:MA_DATA.request primitives for transmission on an aggregated link. These are passed by the Frame Distributor to a port selected by the distribution algorithm. MacN:MA_DATA.indication primitives signifying received frames are passed unchanged from a port to the MAC Client by the Frame Collector.

MAC Clients that generate MA_CONTROL.request primitives (and which expect MA_CONTROL.indication primitives in response) cannot communicate through a Link Aggregation sublayer. They must communicate directly with the MAC Control entity through which these control primitives are to be sent and received.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 286 of 379

The multiplexing of such MAC Control clients with a Link Aggregation sublayer for simultaneous use of a single port is outside the scope of this standard.

The multiplexing of MAC Clients with a Link Aggregation sublayer for simultaneous use of an individual MAC that is also part of a Link Aggregation Group is outside the scope of this standard.

43.2.3 Frame Collector

A Frame Collector is responsible for receiving incoming frames (i.e., AggMuxN:MA_DATA.indications) from the set of individual links that form the Link Aggregation Group (through each link's associated Aggregator Parser/Multiplexer) and delivering them to the MAC Client. Frames received from a given port are delivered to the MAC Client in the order that they are received by the Frame Collector. Since the Frame Distributor is responsible for maintaining any frame ordering constraints, there is no requirement for the Frame Collector to perform any reordering of frames received from multiple links.

The Frame Collector shall implement the function specified in the state diagram shown in Figure 43–3 and the associated definitions contained in 43.2.3.1.

43.2.3.1 Frame Collector state diagram

43.2.3.1.1 Constants

CollectorMaxDelay
: In tens of microseconds, the maximum time that the Frame Collector may delay the delivery of a frame received from an Aggregator Parser to its MAC Client. Value is assigned by management or administration policy.
Value: Integer

43.2.3.1.2 Variables

DA
SA
m_sdu
status
: The parameters of the MA_DATA.indication primitive, as defined in Clause 2.

BEGIN
: A Boolean variable that is set to TRUE when the System is initialized or reinitialized, and is set to FALSE when (re-)initialization has completed.
Value: Boolean

43.2.3.1.3 Messages

Agg:MA_DATA.indication
AggMuxN:MA_DATA.indication
: The service primitives used to pass a received frame to a client with the specified parameters.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 287 of 379

43.2.3.1.4 State diagram

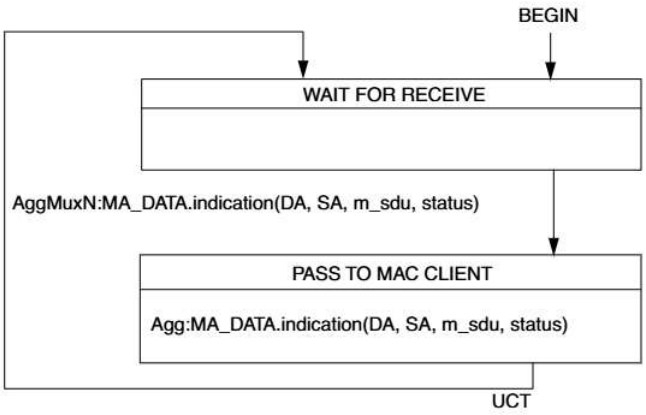


Figure 43–3—Frame Collector state diagram

The architectural models of the IEEE 802.3® MAC and Link Aggregation do not make any provision for queueing of frames between the link and the MAC Client. Furthermore, the state machine conventions used in this standard assume that actions within a state occur instantaneously (see 21.5). However, practical implementations of Link Aggregation will typically incur both queueing and delay in the Frame Collector. In order to ensure that frame delivery is not delayed indefinitely (which could cause a frame ordering problem when moving conversations from one link to another), the Frame Collector shall, upon receiving a frame from an Aggregator Parser, either deliver the frame to its MAC Client, or discard the frame within a CollectorMaxDelay time. The Frame Distributor (within the Partner System at the other end of the link) can assume that all frames transmitted on a given link have been either received by its Partner's MAC Client or discarded after a CollectorMaxDelay plus the propagation delay of the link. The use of CollectorMaxDelay is further discussed in 43A.3.

43.2.4 Frame Distributor

The Frame Distributor is responsible for taking outgoing frames from the MAC Client and transmitting them through the set of links that form the Link Aggregation Group. The Frame Distributor implements a distribution function (algorithm) responsible for choosing the link to be used for the transmission of any given frame or set of frames.

This standard does not mandate any particular distribution algorithm(s); however, any distribution algorithm shall ensure that, when frames are received by a Frame Collector as specified in 43.2.3, the algorithm shall not cause

a) Mis-ordering of frames that are part of any given conversation, or
b) Duplication of frames.

The above requirement to maintain frame ordering is met by ensuring that all frames that compose a given conversation are transmitted on a single link in the order that they are generated by the MAC Client; hence, this requirement does not involve the addition (or modification) of any information to the MAC frame, nor any buffering or processing on the part of the corresponding Frame Collector in order to re-order frames. This approach to the operation of the distribution function permits a wide variety of distribution and load balancing algorithms to be used, while also ensuring interoperability between devices that adopt differing algorithms.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 288 of 379

NOTE—The subject of distribution algorithms and maintenance of frame ordering is discussed in Annex 43A.

The Frame Distributor shall implement the function specified in the state diagram shown in Figure 43–4 and the associated definitions contained in 43.2.4.1.

43.2.4.1 Frame Distributor state diagram

43.2.4.1.1 Variables

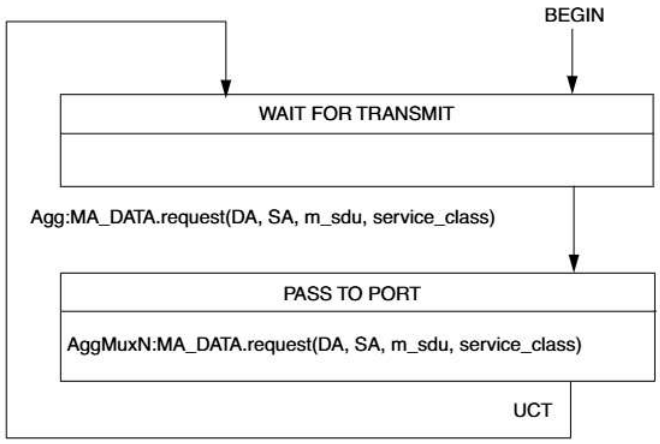
DA
SA
m_sdu
service_class
: The parameters of the MA_DATA.request primitive, as defined in Clause 2.

BEGIN
: A Boolean variable that is set to TRUE when the System is initialized or reinitialized, and is set to FALSE when (re-)initialization has completed.
Value: Boolean

43.2.4.1.2 Messages

Agg:MA_DATA.request
AggMuxN:MA_DATA.request
: The service primitives used to transmit a frame with the specified parameters.

43.2.4.1.3 State diagram



If a client issues an Agg:MA_DATA.request primitive that contains no SA parameter, the AggMuxN:MA_DATA.request primitive generated shall use the Aggregator's MAC address for the SA.

NOTE—The algorithm that the Frame Distributor uses to select the value of N in AggMuxN:MA_DATA.request for a given frame is unspecified.

Figure 43–4—Frame Distributor state diagram

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 289 of 379

43.2.5 Marker Generator/Receiver (optional)

The optional Marker Generator is used by the Marker protocol, as specified in 43.5. When implemented and so requested by the Distribution algorithm, the Marker Generator shall issue an AggMuxN:MA_DATA.request primitive, with an m_sdu containing a Marker PDU as defined in 43.5.3, to the port associated with the conversation being marked, subject to the timing restrictions for Slow Protocols specified in Annex 43B.

The optional Marker Receiver is used by the Marker protocol, as specified in 43.5. It receives Marker Response PDUs from the Aggregator Parser.

43.2.6 Marker Responder

The Marker Responder is used by the Marker protocol, as specified in 43.5. The Marker Responder receives Marker PDUs (generated by a Partner System's Marker Generator), and transmits a Marker Response PDU through the same port from which the Marker PDU was received. While implementation of the Marker Generator/Receiver is optional, the ability to respond to a Marker PDU (the Marker Responder) is mandatory. An implementation conformant to this clause shall implement the Marker Responder as specified in 43.5.4.2, thus ensuring that implementations that need to make use of the protocol can do so.

43.2.7 Aggregator Parser/Multiplexer

On transmission, the Aggregator Multiplexer shall provide transparent pass-through of frames submitted by the Marker Responder and optional Marker Generator to the port specified in the transmission request. The Aggregator Multiplexer shall provide transparent pass-through of frames submitted by the Frame Distributor to the port specified in the transmission request only when the port state is Distributing (see 43.4.15); otherwise, such frames shall be discarded.

On receipt, the Aggregator Parser decodes frames received from the Control Parser, passes those frames destined for the Marker Responder or Marker Receiver to the selected entity, and discards frames with invalid Slow Protocol subtype values (see Table 43B–2). The Aggregator Parser shall pass all other frames to the Frame Collector for passage to the MAC Client only when the port state is Collecting (see 43.4.15); otherwise, such frames shall be discarded. The Aggregator Parser shall implement the function specified in the state diagram shown in Figure 43–5 and the associated definitions contained in 43.2.7.1.

43.2.7.1 Aggregator Parser state diagram

43.2.7.1.1 Constants

Slow_Protocols_Multicast
: The value of the Slow Protocols Multicast address. (See Table 43B–1.)

Slow_Protocols_Type
: The value of the Slow Protocols Length/Type field. (See Annex 43B.)

Marker_subtype
: The value of the Subtype field for the Marker protocol. (See 43.5.3.)
Value: Integer
2

Marker_Information
: The encoding of the Marker Information TLV_type field. (See 43.5.3.)
Value: Integer
1

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 290 of 379

Marker_Response_Information
The encoding of the Marker Response Information TLV_type field. (See 43.5.3.)
Value: Integer
2

43.2.7.1.2 Variables

DA
SA
m_sdu
status
The parameters of the MA_DATA.indication primitive as defined in Clause 2.

Length/Type
The value of the Length/Type field in a received frame.
Value: Integer

Subtype
The value of the octet following the Length/Type field in a Slow Protocol frame. (See Annex 43B.)
Value: Integer

TLV_type
The value contained in the octet following the Version Number in a received Marker or Marker Response frame. This identifies the "type" for the Type/Length/Value (TLV) tuple. (See 43.5.3.)
Value: Integer

BEGIN
A Boolean variable that is set to TRUE when the System is initialized or reinitialized, and is set to FALSE when (re-)initialization has completed.
Value: Boolean

43.2.7.1.3 Messages

CtrlMuxN:MA_DATA.indication
AggMuxN:MA_DATA.indication
The service primitives used to pass a received frame to a client with the specified parameters.

43.2.8 Aggregator

An *Aggregator* comprises an instance of a Frame Collection function, an instance of a Frame Distribution function and one or more instances of the Aggregator Parser/Multiplexer function for a Link Aggregation Group. A single Aggregator is associated with each Link Aggregation Group. An Aggregator offers a standard IEEE 802.3® MAC service interface to its associated MAC Client; access to the MAC service by a MAC Client is always achieved via an Aggregator. An Aggregator can therefore be considered to be a *logical MAC*, bound to one or more ports, through which the MAC client is provided access to the MAC service.

A single, individual MAC address is associated with each Aggregator (see 43.2.10).

An Aggregator is available for use by the MAC Client if the following are all true:

a) It has one or more attached ports.
b) The Aggregator has not been set to a disabled state by administrative action (see 30.7.1.1.13).

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 291 of 379

c) The collection and/or distribution function associated with one or more of the attached ports is enabled (see 30.7.1.1.14).

NOTE—To simplify the modeling and description of the operation of Link Aggregation, it is assumed that there are as many Aggregators as there are ports in a given System; however, this is not a requirement of this standard. Aggregation of two or more ports consists of changing the bindings between ports and Aggregators such that more than one port is bound to a single Aggregator. The creation of any aggregations of two or more links will therefore result in one or more Aggregators that are bound to more than one port, and one or more Aggregators that are not bound to any port. An Aggregator that is not bound to any port appears to a MAC Client as a MAC interface to an inactive port. During times when the bindings between ports and Aggregators are changing, or as a consequence of particular configuration choices, there may be occasions when one or more ports are not bound to any Aggregator.

43.2.8.1 State diagram

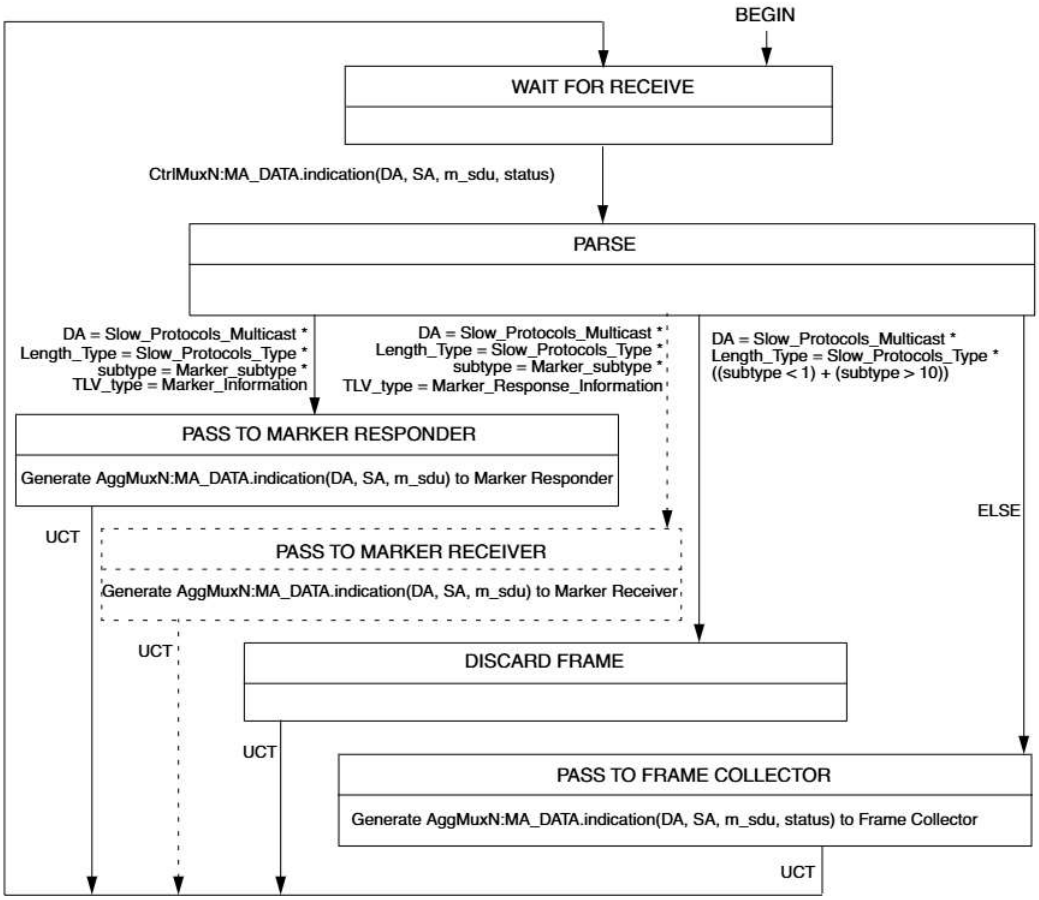


Figure 43–5—Aggregator Parser state diagram

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 292 of 379

43.2.9 Control Parser/Multiplexer

On transmission, the Control Multiplexer shall provide transparent pass-through of frames submitted by the Aggregator and Link Aggregation Control Protocol to the port specified in the transmission request.

On receipt, the Control Parser decodes frames received from the various ports in the Link Aggregation Group, passes those frames destined for the Link Aggregation Control Protocol to the appropriate entity, and passes all other frames to the Aggregator Parser. The Control Parser shall implement the function specified by the state diagram shown in Figure 43–6 and the associated definitions contained in 43.2.9.1.

43.2.9.1 Control Parser state diagram

43.2.9.1.1 Constants

Slow_Protocols_Multicast
: The value of the Slow Protocols Multicast address. (See Table 43B–1.)

Slow_Protocols_Type
: The value of the Slow Protocols Length/Type field. (See Table 43B–2.)

LACP_subtype
: The value of the Subtype field for the Link Aggregation Control Protocol. (See Table 43B–3.)
Value: Integer
1

43.2.9.1.2 Variables

DA
SA
m_sdu
status
: The parameters of the MA_DATA.indication primitive, as defined in Clause 2.

Length/Type
: The value of the Length/Type field in a received frame.
Value: Integer

Subtype
: The value of the octet following the Length/Type field in a Slow Protocol frame.
(See Annex 43B.)
Value: Integer

BEGIN
: A Boolean variable that is set to TRUE when the System is initialized or reinitialized, and is set to FALSE when (re-)initialization has completed.
Value: Boolean

43.2.9.1.3 Messages

MacN:MA_DATA.indication
CtrlMuxN:MA_DATA.indication
: The service primitives used to pass a received frame to a client with the specified parameters.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 293 of 379

43.2.9.1.4 State diagram

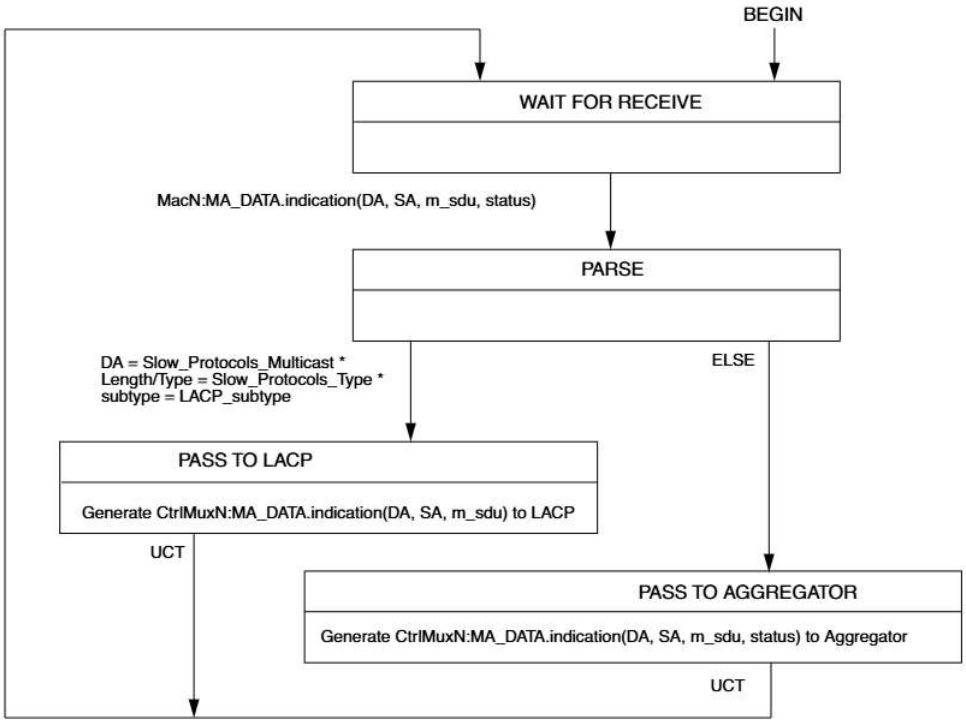


Figure 43–6—Control Parser state diagram

43.2.10 Addressing

Each IEEE 802.3® MAC has an associated globally-unique individual MAC address, whether that MAC is used for Link Aggregation or not (see Clause 4).

Each Aggregator to which one or more ports are attached has an associated globally-unique individual MAC address (see 43.3.3). The MAC address of the Aggregator may be the globally-unique individual MAC addresses of one of the MACs in the associated Link Aggregation Group, or it may be a distinct MAC address. The manner in which such addresses are chosen is not otherwise constrained by this standard.

Protocol entities sourcing frames from within the Link Aggregation sublayer (e.g., LACP and the Marker protocol) use the MAC address of the MAC within an underlying port as the source address in frames transmitted through that port. The MAC Client sees only the Aggregator and not the underlying MACs, and therefore uses the Aggregator's MAC address as the source address in transmitted frames. If a MAC Client submits a frame to the Aggregator for transmission without specifying a source address, the Aggregator inserts its own MAC address as the source address for transmitted frames.

NOTE—This behavior causes the Aggregator to behave the same way as a standard MAC with regard to frames submitted by its client.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 294 of 379

43.3 Link Aggregation Control

Link Aggregation Control configures and controls the Link Aggregation sublayer using static information local to the control function and dynamic information exchanged by means of the Link Aggregation Control Protocol.

For each aggregatable port in the System, Link Aggregation Control

a) Maintains configuration information (reflecting the inherent properties of the individual links as well as those established by management) to control aggregation.
b) Exchanges configuration information with other Systems to allocate the link to a Link Aggregation Group.

NOTE—A given link is allocated to, at most, one Link Aggregation Group at a time. The allocation mechanism attempts to maximize aggregation, subject to management controls.

c) Attaches the port to the Aggregator used by the Link Aggregation Group, and detaches the port from the Aggregator when it is no longer used by the Group.
d) Uses information from the Partner System's Link Aggregation Control entity to enable or disable the Aggregator's Collector and Distributor.

The operation of Link Aggregation Control involves the following activities, which are described in detail in the subclauses that follow:

e) Checking that candidate links can actually be aggregated.
f) Controlling the addition of a link to a Link Aggregation Group, and the creation of the group if necessary.
g) Monitoring the status of aggregated links to ensure that the aggregation is still valid.
h) Removing a link from a Link Aggregation Group if its membership is no longer valid, and removing the group if it no longer has any member links.

In order to allow Link Aggregation Control to determine whether a set of links connect to the same System, and to determine whether those links are compatible from the point of view of aggregation, it is necessary to be able to establish

i) A globally unique identifier for each System that participates in Link Aggregation (see 43.3.2).
j) A means of identifying the set of capabilities associated with each port and with each Aggregator, as understood by a given System.
k) A means of identifying a Link Aggregation Group and its associated Aggregator.

System identification allows the detection of links that are connected in a loopback configuration (i.e., both ends of the same link are connected to the same System).

43.3.1 Characteristics of Link Aggregation Control

Link Aggregation Control provides a configuration capability that is

a) **Automatic.** In the absence of manual override controls, an appropriate set of Link Aggregation Groups is automatically configured, and individual links are allocated to those groups. If a set of links can aggregate, they do aggregate.
b) **Continuous.** Manual intervention or initialization events are not a requirement for correct operation. The configuration mechanism continuously monitors for changes in state that require reconfiguration. The configuration functions detect and correct misconfigurations by performing reconfiguration and/or by taking misconfigured links out of service.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 295 of 379

c) **Deterministic.** The configuration can resolve deterministically; i.e., the configuration achieved can be made independent of the order in which events occur, and be completely determined by the combination of the capabilities of the individual links and their physical connectivity.
d) **Controllable.** The configuration capabilities accommodate devices with differing hardware and software constraints on Link Aggregation.
e) **Compatible.** Links that cannot take part in Link Aggregation, either because of their inherent capabilities or of the capabilities of the devices to which they attach, operate as normal IEEE 802.3® links. The introduction of Link Aggregation capability at one or both ends of a link should not result in a degradation of the perceived performance of the link.
f) **Rapid.** The configuration resolves rapidly to a stable configuration. Convergence can be achieved by the exchange of three LACPDUs, without dependence on timer values.

with

g) **Low risk of misdelivery.** The operation of the (re-)configuration functions minimizes the risk of frames being delivered to the wrong Aggregator.
h) **Low risk of duplication or mis-ordering.** The operation of the (re-)configuration functions minimizes the risk of frame duplication and frame mis-ordering.
i) **Low protocol overhead.** The overhead involved in external communication of configuration information between devices is small.

43.3.2 System identification

The globally unique identifier used to identify a System shall be the concatenation of a globally administered individual MAC address and the System Priority. The MAC address chosen may be the individual MAC address associated with one of the ports of the System.

Where it is necessary to perform numerical comparisons between System Identifiers, each System Identifier is considered to be an eight octet unsigned binary number, constructed as follows:

a) The two most significant octets of the System Identifier comprise the System Priority. The System Priority value is taken to be an unsigned binary number; the most significant octet of the System Priority forms the most significant octet of the System Identifier.
b) The third most significant octet of the System Identifier is derived from the initial octet of the MAC address; the least significant bit of the octet is assigned the value of the first bit of the MAC address, the next most significant bit of the octet is assigned the value of the next bit of the MAC address, and so on. The fourth through eighth octets are similarly assigned the second through sixth octets of the MAC address.

43.3.3 Aggregator identification

Each Aggregator to which one or more ports are attached shall be assigned a unique, globally administered individual MAC address. The MAC address assigned to the Aggregator may be the same as the MAC address assigned to one of its bound ports. No Aggregator shall be assigned a MAC address that is the same as that of a port bound to a different Aggregator within the System. When receiving frames, a port is never required to recognize more than one unicast address, i.e., the Aggregator's MAC address.

NOTE—This ensures that Aggregators can be uniquely addressed, and allows (but does not require) the unique address to be allocated from the same set of addresses as are assigned to the ports. It also acknowledges the fact that locally administered addresses may be used in particular implementations or environments. The stated restriction on the allocation of MAC addresses to Aggregators may have implications with regard to the choice of selection algorithm.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 296 of 379

An aggregator also shall be assigned an integer identifier that is used by Link Aggregation Control to uniquely identify the aggregator within the System. This value will typically be the same as the interface identifier (ifIndex) used for management purposes.

43.3.4 Port identification

Link Aggregation Control uses a Port Identifier, comprising the concatenation of a Port Priority and a Port Number, to identify the port. Port Numbers (and hence, Port Identifiers) shall be uniquely assigned within a System. Port Number 0 shall not be assigned to any port.

When it is necessary to perform numerical comparisons between Port Identifiers, each Port Identifier is considered to be a four octet unsigned binary number constructed as follows:

a) The most significant and second most significant octets are the first and second most significant octets of the Port Priority, respectively.
b) The third and fourth most significant octets are the first and second most significant octets of the Port Number, respectively.

43.3.5 Capability identification

The ability of one port to aggregate with another is summarized by a simple integer parameter, known as a Key. This facilitates communication and comparison of aggregation capabilities, which may be determined by a number of factors, including

a) The port's physical characteristics, such as data rate, duplexity, point-to-point or shared medium.
b) Configuration constraints established by the network administrator.
c) Use of the port by higher layer protocols (e.g. assignment of Network Layer addresses).
d) Characteristics or limitations of the port implementation itself.

Two Keys shall be associated with each port: an operational Key and an administrative Key. The operational Key is the Key that is currently in active use for the purposes of forming aggregations. The administrative Key allows manipulation of Key values by management. The administrative and operational Keys assigned to a port may differ

e) If the operation of the implementation is such that an administrative change to a Key value cannot be immediately reflected in the operational state of the port.
f) If the System supports the dynamic manipulation of Keys, as discussed in 43.6.2, either to accurately reflect changes in operational capabilities of the port (for example, as a result of Auto-Negotiation), or to provide a means of handling constraints on aggregation capability.

A given Key value is meaningful only in the context of the System that allocates it; there is no global significance to Key values. Similarly, the relationship between administrative and operational Key values is meaningful only in the context of the System that allocates it. When a System assigns an administrative Key value to a set of ports, it signifies that the set of ports have the potential to aggregate together, subject to the considerations discussed in 43.6.2. When a System assigns an operational Key value to a set of ports, it signifies that, in the absence of other constraints, the current operational state of the set of ports allows any subset of that set of ports (including the entire set) to be aggregated together from the perspective of the System making the assignment. The set of such ports that will actually be aggregated will be those that terminate at a common Partner System, and for which that Partner System has assigned a common operational Key value, local to that Partner. The set of ports in a given System that share the same operational Key value are said to be members of the same Key Group.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 297 of 379

A System may determine that a given link is not able to be aggregated with other links. Such links are referred to as Individual links (as opposed to Aggregatable links). A System may declare a link to be Individual if the inherent properties of the link allow its use as part of an aggregation, but the system is aware of no other links that are capable of aggregating with this link (e.g., the System has allocated a unique operational Key value to the link).

The capability information communicated between Systems, therefore, includes this local knowledge of the aggregation capability of the link in addition to the operational Key value; i.e., whether the System considers the link to be Aggregatable or Individual.

An administrative Key value and an operational Key value shall also be associated with each Aggregator. The operational Key is the Key that is currently in active use for the purposes of forming aggregations. The administrative Key allows manipulation of Key values by management. The values of administrative and operational Key for an Aggregator may differ in the same manner as that of port Keys, per item e) and item f), in this subclause. Ports that are members of a given Key Group can only be bound to Aggregators that share the same operational Key value.

All Keys are 16-bit identifiers. All values except the null value (all zeroes) are available for local use.

NOTE—This model allows for two convenient initial configurations. The first is achieved by assigning each port an initial administrative and operational Key value identical to its port number, and assigning the same port numbers as Keys to the corresponding Aggregators for each port. A device with this initial configuration will bring up all links as individual, non-aggregated links. The second is achieved by assigning the same administrative and operational Key values to all ports with a common set of capabilities, and also to all Aggregators. A device with this initial configuration will attempt to aggregate together any set of links that have the same Partner System ID and operational Key, and for which both Systems are prepared to allow aggregation.

43.3.6 Link Aggregation Group identification

A Link Aggregation Group consists of either

a) One or more Aggregatable links that terminate in the same pair of Systems and whose ports belong to the same Key Group in each System, or
b) An Individual link.

43.3.6.1 Construction of the Link Aggregation Group Identifier

A unique Link Aggregation Group Identifier (LAG ID) is constructed from the following parameters for each of the communicating Systems:

a) The System Identifier
b) The operational Key assigned to the ports in the LAG
c) The Port Identifier, if the link is identified as an Individual link

The local System's values for these parameters shall be non-zero. In cases where the local System is unable to determine the remote System's values for these parameters by exchange of protocol information, administrative values are used in the construction of the LAG ID. The value of these administrative parameters for the remote System may be configured as zero, provided that the port(s) concerned are also configured to be Individual.

A compound identifier formed from the System Identifiers and Key values alone is sufficient to identify a LAG comprising Aggregatable links. However, such an identifier is not sufficient for a LAG comprising a single Individual link where the Partner System Identifier and operational Key may be zero. Even if these are non-zero there may be multiple Individual Links with the same System Identifier and operational Key combinations, and it becomes necessary to include Port Identifiers to provide unique LAG IDs.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 298 of 379

Given

d) S and T are System Identifiers,
e) K and L are the operational Keys assigned to a LAG by S and T respectively, and
f) P and Q are the Port Identifiers of the ports being attached if the LAG comprises a single Individual Link and zero if the LAG comprises one or more Aggregatable links,

then the general form of the unique LAG ID is [(SKP), (TLQ)].

To simplify comparison of LAG IDs it is conventional to order these such that S is the numerically smaller of S and T.

43.3.6.2 Representation of the Link Aggregation Group Identifier

In order to allow for convenient transcription and interpretation by human network personnel, this standard provides a convention for representing compound LAG IDs. Using this format

a) All fields are written as hexadecimal numbers, 2 digits per octet, in canonical format.
b) Octets are presented in order, from left to right. Within fields carrying numerical significance (e.g., priority values), the most significant octet is presented first, and the least significant octet last.
c) Within fields that carry MAC addresses, successive octets are separated by dashes (-), in accordance with the hexadecimal representation for MAC addresses defined in IEEE Std 802-1990™ .
d) Parameters of the LAG ID are separated by commas.

For example, consider the parameters for the two Partners in a Link Aggregation Group shown in Table 43–1.

Table 43–1—Example Partner Parameters

	Partner SKP	Partner TLQ
System Parameters (S, T)	System Priority = 0x8000 (see 43.4.2.2) System Identifier = AC-DE-48-03-67-80	System Priority = 0x8000 (see 43.4.2.2) System Identifier = AC-DE-48-03-FF-FF
Key Parameter (K, L)	Key = 0x0001	Key = 0x00AA
Port Parameters (P, Q)	Port Priority = 0x80 (see 43.4.2.2) Port Number = 0x0002	Port Priority = 0x80 (see 43.4.2.2) Port Number = 0x0002

The complete LAG ID derived from this information is represented as follows, for an Individual link:

[(SKP), (TLQ)] = [(8000,AC-DE-48-03-67-80,0001,80,0002), (8000,AC-DE-48-03-FF-FF,00AA,80,0002)]

The corresponding LAG ID for a set of Aggregatable links is represented as follows:

[(SKP), (TLQ)] = [(8000,AC-DE-48-03-67-80,0001,00,0000), (8000,AC-DE-48-03-FF-FF,00AA,00,0000)]

NOTE—The difference between the two representations is that, for an Aggregatable link, the port identifier components are zero.

It is recommended that this format be used whenever displaying LAG ID information for use by network personnel.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 299 of 379

43.3.7 Selecting a Link Aggregation Group

Each port is selected for membership in the Link Aggregation Group uniquely identified by the LAG ID (composed of operational information, both derived from local administrative parameters and received through the Link Aggregation Control Protocol). Initial determination of the LAG ID is delayed to allow receipt of such information from a peer Link Aggregation Control entity; in the event such information is not received, locally configured administrative defaults are assumed for the remote port's operational parameters.

Where a particular link is known to be Individual, the complete LAG ID is not required to select the Link Aggregation Group since the link will not be aggregated with any other.

43.3.8 Agreeing on a Link Aggregation Group

Before frames are distributed and collected from a link, both the local Link Aggregation Control entity and its remote peer (if present) need to agree on the Link Aggregation Group. The Link Aggregation Control Protocol allows each of the communicating entities to check their peer's current understanding of the LAG ID, and facilitates rapid exchange of operational parameters while that understanding differs from their own. The protocol entities monitor their operation and, if agreement is not reached (perhaps due to an implementation failure), management is alerted.

The ability of LACP to signal that a particular link is Individual can accelerate the use of the link since, if both Link Aggregation Control entities know that the link is Individual, full agreement on the LAG ID is not necessary.

43.3.9 Attaching a link to an Aggregator

Once a link has selected a Link Aggregation Group, Link Aggregation Control can attach that link to a compatible Aggregator. An Aggregator is compatible if

a) The Aggregator's operational Key matches the port's operational Key, and
b) All other links currently attached to the Aggregator have selected the same Link Aggregation Group.

If several compatible Aggregators exist, Link Aggregation Control may employ a locally determined algorithm, either to ensure deterministic behavior (i.e., independence from the order in which Aggregators become available) or to maximize availability of the aggregation to a MAC Client. If no compatible Aggregator exists, then it is not possible to enable the link until such a time as a compatible Aggregator becomes available.

NOTE—In a properly configured System, there should always be a suitable Aggregator available with the proper Key assigned to serve a newly created Link Aggregation Group, so the unavailability of a compatible Aggregator is normally a temporary state encountered while links are moved between Aggregators. However, given the flexibility of the Key scheme, and given that in some implementations there may not be enough Aggregators to service a given configuration of links, it is possible to create configurations in which there is no Aggregator available to serve a newly identified LAG, in which case the links that are members of that LAG cannot become active until such a time as the configuration is changed to free up an appropriate Aggregator.

Links that are not successful candidates for aggregation (e.g., links that are attached to other devices that cannot perform aggregation or links that have been manually configured to be non-aggregatable) are enabled to operate as individual IEEE 802.3® links. For consistency of modeling, such a link is regarded as being attached to a compatible Aggregator that can only be associated with a single link. That is, from the perspective of Link Aggregation, non-aggregated links are not a special case; they compose an aggregation with a maximum membership of one link.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 300 of 379

More than one link can select the same Link Aggregation Group within a short period of time and, as these links detach from their prior Aggregators, additional compatible Aggregators can become available. In order to avoid such events causing repeated configuration changes, Link Aggregation Control applies hysteresis to the attachment process and allows multiple links to be attached to an Aggregator at the same time.

43.3.10 Signaling readiness to transfer user data

Once a link has been attached to an Aggregator (43.3.9) compatible with the agreed-upon Link Aggregation Group (43.3.8), each Link Aggregation Control entity signals to its peer its readiness to transfer user data to and from the Aggregator's MAC Client. In addition to allowing time for the organization of local Aggregator resources, including the possibility that a compatible Aggregator may not exist, explicit signaling of readiness to transfer user data can be delayed to ensure preservation of frame ordering and prevention of frame duplication. Link Aggregation Control will not signal readiness until it is certain that there are no frames in transit on the link that were transmitted while the link was a member of a previous Link Aggregation Group. This may involve the use of an explicit Marker protocol that ensures that no frames remain to be received at either end of the link before reconfiguration takes place. The operation of the Marker protocol is described in 43.5. The decision as to when, or if, the Marker protocol is used is entirely dependent upon the nature of the distribution algorithm that is employed.

43.3.11 Enabling Collection and Distribution

Every Aggregator can enable or disable Collection and Distribution of frames for each port that is attached to the Aggregator. Initially, both Collection and Distribution are disabled. Once the Link Aggregation Control entity is ready to transfer user data using the link and its peer entity has also signaled readiness, the process of enabling the link can proceed. The Collector is enabled (thus preparing it to receive frames sent over the link by the remote Aggregator's Distributor) and that fact is communicated to the Partner. Once the received information indicates that the remote Aggregator's Collector is enabled, the Distributor is also enabled.

NOTE—This description assumes that the implementation is capable of controlling the state of the transmit and receive functions of the MAC independently. In an implementation where this is not possible, the transmit and receive functions are enabled or disabled together. The manner in which this is achieved is detailed in the description of the Mux machine (see 43.4.15).

If at least one port's Mux in the Link Aggregation Group is Collecting, then the Receive state of the corresponding Aggregator will be Enabled. If at least one port's Mux in the Link Aggregation Group is Distributing, then the Transmit state of the corresponding Aggregator will be Enabled.

43.3.12 Monitoring the membership of a Link Aggregation Group

Each link is monitored in order to confirm that the Link Aggregation Control functions at each end of the link still agree on the configuration information for that link. If the monitoring process detects a change in configuration that materially affects the link's membership in its current LAG, then it may be necessary to remove the link from its current LAG and to move it to a new LAG.

43.3.13 Detaching a link from an Aggregator

A port may be detached from the Aggregator used by its Link Aggregation Group as a result of protocol (e.g., Key) changes, or because of System constraints (e.g., exceeding a maximum allowable number of aggregated links, or device failures) at either end of the link. Both classes of events will cause the LAG ID information for the link to change, and it will be necessary for Link Aggregation Control to detach the link from its current Aggregator and move it to a new LAG (if possible). At the point where the change is detected, the Collecting and Distributing states for the port are set to FALSE. The Frame Distribution function is informed that the link is no longer part of the group, the changed configuration information is

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 301 of 379

communicated to the corresponding Link Aggregation Partner, then the Frame Collection function is informed that the link is no longer part of the group.

Once a link has been removed from its Aggregator, the link can select its new Link Aggregation Group and then attach to a compatible Aggregator, as described in 43.3.7 and 43.3.9.

Any conversation that is reallocated to a different link as a result of detaching a link from an Aggregator shall have its frame ordering preserved. This may involve the use of the Marker protocol to ensure that no frames that form part of that conversation remain to be received at either end of the old link before the conversation can proceed on the new link.

43.3.14 Configuration and administrative control of Link Aggregation

Administrative configuration facilities allow a degree of control to be exerted over the way that links may be aggregated. In particular, administrative configuration allows

a) The Key values associated with a port to be identified or modified.
b) The Key values associated with an Aggregator to be identified or modified.
c) Links to be identified as being incapable of aggregation.
d) Link Aggregation Control Protocol parameters to be identified or modified.

43.3.15 Link Aggregation Control state information

The Link Aggregation Control function maintains the following information with respect to each link:

a) The identifier of the Link Aggregation Group to which it currently belongs.
b) The identifier of the Aggregator associated with that Link Aggregation Group.
c) The status of interaction between the Frame Collection function of the Aggregator and the link (Collecting TRUE or Collecting FALSE). Collecting TRUE indicates that the receive function of this link is enabled with respect to its participation in an aggregation; i.e., received frames will be passed up to the Aggregator for collection.
d) The status of interaction between the Frame Distribution function of the Aggregator and the link (Distributing TRUE or Distributing FALSE). Distributing TRUE indicates that the transmit function of this link is enabled with respect to its participation in an aggregation; i.e., frames may be passed down from the Aggregator's distribution function for transmission.

This state information is communicated directly between Link Aggregation Control and the Aggregator through shared state variables without the use of a formal service interface.

The Link Aggregation Control function maintains the following information with respect to each Aggregator:

e) The status of the Frame Collection function (Receive Enabled or Receive Disabled).
f) The status of the Frame Distribution function (Transmit Enabled or Transmit Disabled).

These status values are exactly the logical OR of the Collection and Distribution status of the individual links associated with that Aggregator; i.e., if one or more links in the Link Aggregation Group are Collecting, then the Aggregator is Receive Enabled, and if one or more links are Distributing, then the Aggregator is Transmit Enabled.

The Transmit and Receive status of the Aggregator effectively govern the point at which the Aggregator becomes available for use by the MAC Client, or conversely, the point at which it ceases to be available.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 302 of 379

43.4 Link Aggregation Control Protocol (LACP)

The Link Aggregation Control Protocol (LACP) provides a standardized means for exchanging information between Partner Systems on a link to allow their Link Aggregation Control instances to reach agreement on the identity of the Link Aggregation Group to which the link belongs, move the link to that Link Aggregation Group, and enable its transmission and reception functions in an orderly manner.

43.4.1 LACP design elements

The following considerations were taken into account during the development of the protocol described in this subclause:

a) The protocol depends upon the transmission of information and state, rather than the transmission of commands. LACPDUs sent by the first party (the Actor) convey to the second party (the Actor's protocol Partner) what the Actor knows, both about its own state and that of the Partner.
b) The information conveyed in the protocol is sufficient to allow the Partner to determine what action to take next.
c) Active or passive participation in LACP is controlled by LACP_Activity, an administrative control associated with each port, that can take the value Active LACP or Passive LACP. Passive LACP indicates the port's preference for not transmitting LACPDUs unless its Partner's control value is Active LACP (i.e., a preference not to speak unless spoken to). Active LACP indicates the port's preference to participate in the protocol regardless of the Partner's control value (i.e., a preference to speak regardless).
d) Periodic transmission of LACPDUs occurs if the LACP_Activity control of either the Actor or the Partner is Active LACP. These periodic transmissions will occur at either a slow or fast transmission rate depending upon the expressed LACP_Timeout preference (Long Timeout or Short Timeout) of the Partner System.
e) In addition to periodic LACPDU transmissions, the protocol transmits LACPDUs when there is a Need To Transmit (NTT) something to the Partner; i.e., when the Actor's state changes or when it is apparent from the Partner's LACPDUs that the Partner does not know the Actor's current state.
f) The protocol assumes that the rate of LACPDU loss is very low.

There is no explicit frame loss detection/retry mechanism employed by the LACP; however, if information is received from the Partner indicating that it does not have up-to-date information on the Actor's state, or if the next periodic transmission is due, then the Actor will transmit a LACPDU that will correctly update the Partner.

43.4.2 LACPDU structure and encoding

43.4.2.1 Transmission and representation of octets

All LACPDUs comprise an integral number of octets. The bits in each octet are numbered from 0 to 7, where 0 is the low-order bit. When consecutive octets are used to represent a numerical value, the most significant octet is transmitted first, followed by successively less significant octets.

When the encoding of (an element of) a LACPDU is depicted in a diagram

a) Octets are transmitted from top to bottom.
b) Within an octet, bits are shown with bit 0 to the left and bit 7 to the right, and are transmitted from left to right.
c) When consecutive octets are used to represent a binary number, the octet transmitted first has the more significant value.

Copyright © 2002 IEEE. All rights reserved.

LINKPLAY EXHIBIT 1017
Page 303 of 379