

10GBASE-T start up: new concepts and proposed text

Gottfried Ungerboeck, Broadcom Corporation

55.4.2.5 PHY Control function

[10GBASE-T] PHY Control performs the control actions that are needed to bring a 10GBASE-T PHY into a mode of operation during which PCS data frames are exchanged with the link partner. PHY Control shall comply with the state diagram description given in Figure 55-cc. The figure also depicts the Auto Negotiation function defined in Clause 28. Auto Negotiation and PHY Control are envisaged as processes, which operate in parallel and communicate with each other through the exchange of variables.

55.4.2.5.1 Interaction between PHY Control and Auto Negotiation

Auto Negotiation controls PHY Control through variable `link_control` $\in \{\text{DISABLE}, \text{ENABLE}\}$ and specifies PHY operation further through additional variables, e.g., `config` $\in \{\text{MASTER}, \text{SLAVE}\}$. PHY Control provides variable `link_status` $\in \{\text{FAIL}, \text{OK}\}$. The interaction between Auto Negotiation and PHY Control is illustrated in Figure 55-aa.

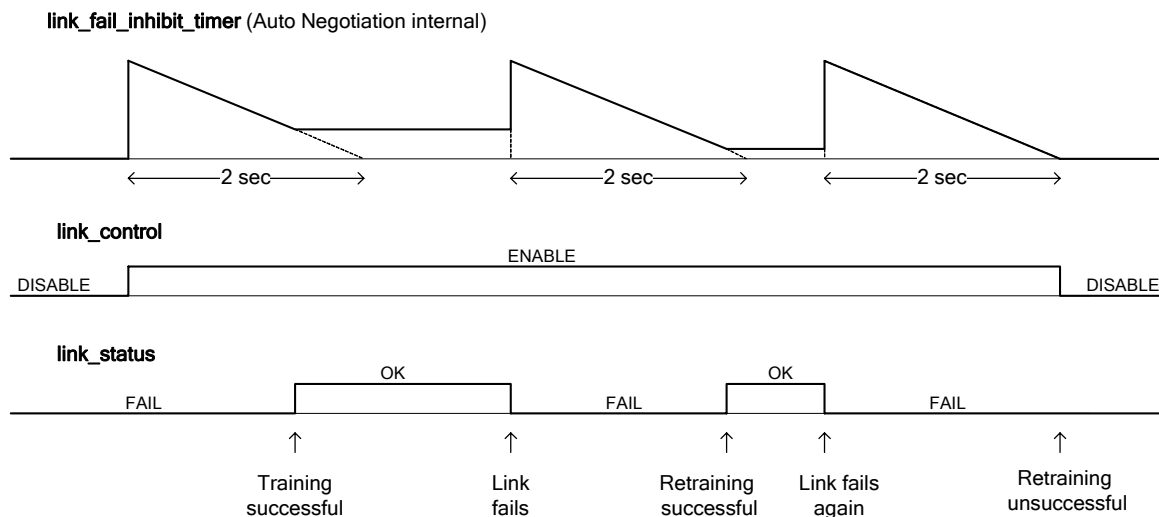


Figure 55-aa—Interaction between Auto Negotiation and PHY Control

Assertion of `link_control` = **DISABLE** by Auto Negotiation forces PHY Control unconditionally into state **PHY_Disabled**. In state **PHY_Disabled**, the PHY is silent, signals `link_status` = **FAIL**, and waits for being enabled by Auto Negotiation.

When Auto Negotiation sets link_control = ENABLE, PHY Control begins the PMA training process. AutoNegotiation also starts the internal link_fail_inhibit_timer at a nominal value of 2 sec. PHY Control is given this maximum time to reach state PCS_Data, where it sets link_status = OK. *[When Auto Negotiation observes link_status = OK, it may stop the link_fail_inhibit_timer or let it continue to count down to zero and stop there.]*

During normal PHY operation in state PCS_Data, a failure situation may occur causing PHY Control to return to state Silent, set link_status = FAIL, and then begin PMA retraining. When Auto Negotiation observes the transition of link_status from OK to FAIL while link_control = ENABLE, the link_fail_inhibit_timer is restarted at a nominal value of 2 sec. *[or a smaller value for retraining?]* This gives PHY Control another opportunity to reach state PCS_Data within this maximum time.

When the link_fail_inhibit_timer has counted down to zero and PHY Control has not yet signaled link_status = OK, Auto Negotiation sets link_control = DISABLE and thus forces PHY Control to return to state PHY_Disabled.

[By observing link_status, Auto Negotiation can determine the length of initial training periods and the length and frequency of retraining events. Several possible uses of this information are conceivable.]

55.4.2.5.2 PHY Control information exchange during PMA training

During transmission of PMA training frames (PHY Control states PMA_Train1_M, PMA_Train2_M, PMA_Train2_S, PMA_Coeff_Exch, and PMA_Fine_Adj) information is sent to the link partner in an 8-octet InfoField that is inserted in the 64 last bit positions associated with pair A of each PMA training frame (see Figure 55-13 *[to be updated]*). The link partner is not required to decode every InfoField, but InfoFields must be decoded frequently enough to ensure that PHY Control performs correct actions in response to information exchanged between the link partners.

55.4.2.5.3 InfoField format

[In some implementations, the resources required to decode an InfoField will be proportional to the length of the InfoField. This provides one motivation to shorten the InfoField from 16 octets as in Draft 2.2 to 8 octets (or even smaller length). Avoiding wasteful use of bit positions and redundancies as in the InfoFields of Draft 2.2 is another reason.]

The InfoField and its state dependent payloads are depicted in Figure 55-bb. The InfoField comprises 8 octets. The first two octets provide a start_of_InfoField_delimiter (Oct1:2), which contains the hexadecimal value 0xAB70, i.e., the binary sequence 1010101101110000 *[or whatever]*. The next four octets convey the InfoField payload (Oct3:6), which is checked by a 2-octet CRC-16 (Oct7:8). Unused bit positions in the payload may be set in the transmitter to arbitrary values and are ignored upon reception.

The CRC-16 is defined as follows. Let polynomial

$$d(x) = d_{31}x^{31} + d_{30}x^{30} + \dots d_1x + d_0$$

represent the payload bits, where $d_{31} = \text{Oct3}\langle 7 \rangle$, $d_{30} = \text{Oct3}\langle 6 \rangle$, ... $d_1 = \text{Oct6}\langle 1 \rangle$, $d_0 = \text{Oct6}\langle 0 \rangle$. Let polynomial

$$r(x) = r_{15}x^{15} + r_{14}x^{14} + \dots + r_1x + r_0$$

represent the check bits, where $r_{15} = \text{Oct7}\langle 7 \rangle$, $r_{14} = \text{Oct7}\langle 6 \rangle$, ... $r_1 = \text{Oct8}\langle 1 \rangle$, $r_0 = \text{Oct8}\langle 0 \rangle$. The polynomial $r(x)$ is the remainder of the division of $d(x)x^{16}$ by the generator polynomial

$$g(x) = (x + 1)(x^{15} + x + 1) = x^{16} + x^{15} + x^2 + 1.$$

[...This definition of CRC-16 is more compact but otherwise equivalent to the pictorial definition given in Draft 2.2, Figure 55-22. ... Protecting 4 payload octets by a CRC-16 may be considered as overkill; CRC-8 may be sufficient. Using a shortened RS or BCH code instead of a CRC may be another option.]

55.4.2.5.4 InfoField payload: state indicator

The 2-bit state_indicator (SI, Oct 3<7:6>) reflects the PHY Control state of the transmitting PHY: SI = 00 for PMA_Train1_M [the SLAVE does not transmit in state PMA_Train1_S], SI = 01 for PMA_Train2_M or PMA_Train2_S, SI = 10 for PMA_Coeff_Exch, and SI = 11 for PMA_Fine_Adj.

55.4.2.5.5 InfoField payload: power back-off (PBO)

The 3-bit field current_PBO (Oct3<5:3>) indicates the current PBO setting of the transmitting PHY.

The 3-bit field next_PBO (Oct3<2:0>) is employed to announce in connection with a non-zero value of transition_count a change in the PBO setting of the transmitting PHY. The transition takes effect immediately after transmission of the InfoField, in which the transition count reaches zero. When transition_count = 0, the field next_PBO is ignored.

The 3-bit field requested_PBO (Oct4<7:5>) is used to request that the link partner changes its PBO setting to the requested value. If the link partner receives a requested_PBO field that differs from its current PBO setting, it must announce a transition to the requested PBO setting.

The fields current_PBO, next_PBO, and requested_PBO are only used in states PMA_Train1_M, PMA_Train2_M and PMA_Train2_S.

[PBO is defined as in Draft 2.2]

[Notice that there are no fields $THP \in \{THP_{Byt}, THP_{Short}, THP_{Med}, THP_{Long}\}$. The specification of these fields in the InfoField format of Draft 2.2 is redundant and hence not needed. If fixed precoding is requested for the transmission of PMA training frames in states PMA_Train1_M and PMA_Train2_M (corresponding in Draft 2.2 to state PMA_Training_Init_M) and state PMA_Train2_S (corresponding in Draft 2.2 to state PMA_Training_Init_S), the value of THP is implied by the value of PBO; see Draft 2.2, section 55.4.2.5.11: PBO = 7 implies $THP = THP_{Short}$, PBO = 5 implies $THP = THP_{Med}$, PBO = 3 implies $THP = THP_{Long}$].

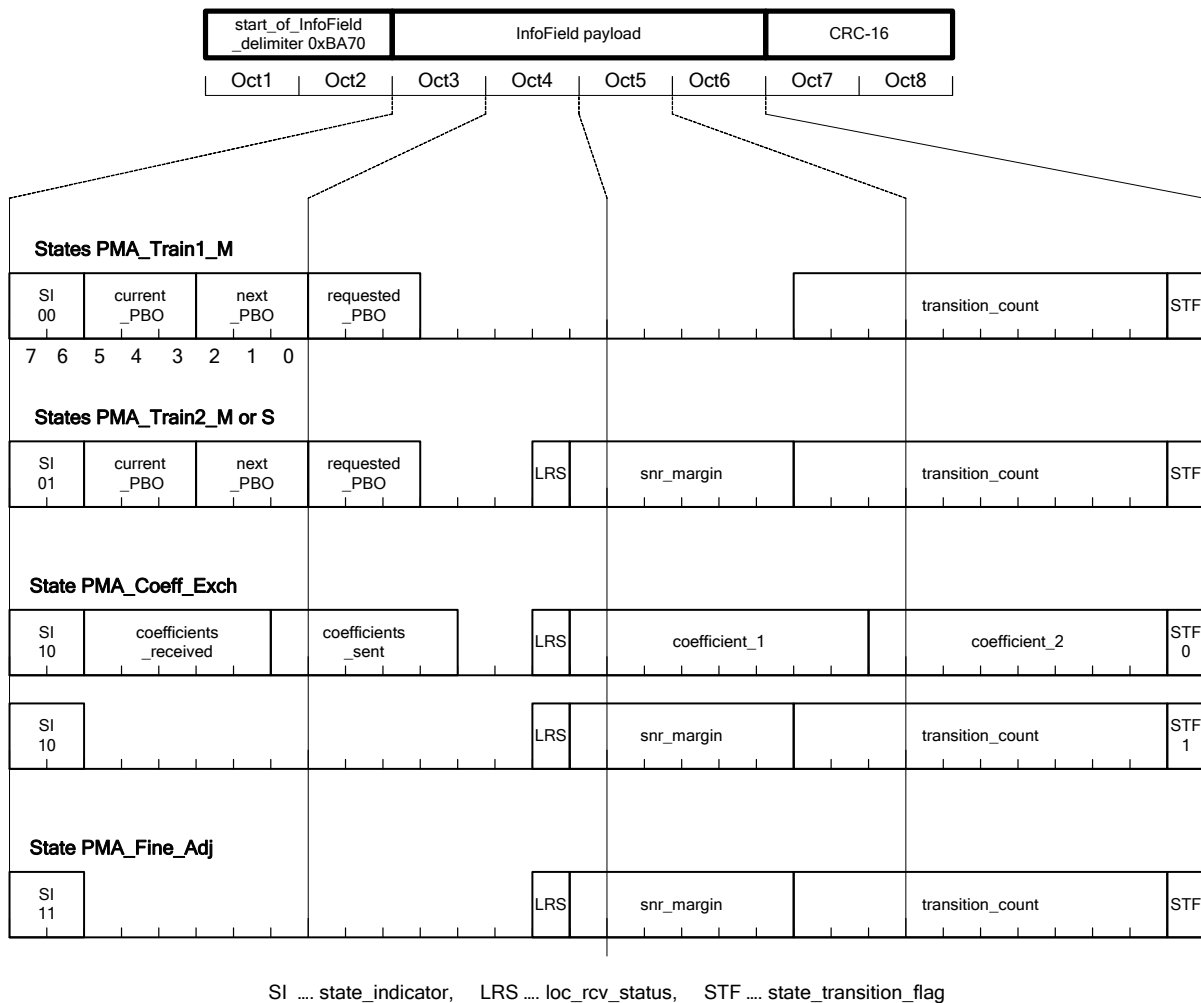


Figure 55-bb—InfoField format and payloads

55.4.2.5.6 InfoField payload: local receiver status and SNR margin

The loc_rcvr_status (LRS, Oct4<1>) bit is transmitted during all PMA training states except PMA_Train1_M. *[In state PMA_Train1_M, the MASTER has to detect the beginning of SLAVE transmission, but may not yet decode InfoFields].* loc_rcvr_status reflects a judgment of the local receiver on the projected decision-point SNR margin for reliable decoding of LDPC-coded 128 DSQ signals, i.e., decoding of PCS frames, with the current setting of PBO by the link partner and the adjustments of the local echo and next cancellers and the equalizer. loc_rcvr_status = OK indicates sufficient SNR margin and that a different PBO setting of the link partner and/or further training of the local cancellation and equalization filters is not warranted.

The information provided by local_rcvr_status is further qualified by the 6-bit field snr_margin (Oct4<0>, Oct5<7:3>). The field conveys the anticipated SNR margin for decoding PCS frames with the current PBO setting of the link partner and the local adaptive filter adjustments. The field is included

for monitoring and diagnostic purposes. The snr_margin values 000000, 000001, ... 111110, 111111 shall indicate decision-point SNR margins of ≤ -8.00 , -7.75 , ... 7.50 , ≥ 7.75 dB, respectively.

55.4.2.5.7 InfoField payload: transition counter and state transition flag

By inserting a non-zero value into the 10-bit field transition_count (Oct5<2:0>, Oct6<7:1>), a PHY announces to the link partner a transition. The transition may be a change in the PBO setting of the transmitting PHY or a state transition to the next PHY Control state. In every subsequently transmitted InfoField, transition_count will automatically be decremented by one until transition_count reaches the value zero. The transition_count then stays at zero until another transition is announced. The transition takes effect immediately after the InfoField, in which the transition_count reaches the value zero. PHY Control operation shall not depend on decoding this particular last InfoField. Announced transitions shall never be revoked.

The state_transition_flag (STF, Oct6<0>) bit is used to indicate state transitions. Setting STF = 1 and transition_count to a non-zero value announces a transition to the next state in the PHY Control state diagram.

[Add "The minimum value, to which transition_count may be set and from where transition_count is then counting down to zero, is given by variable min_transition_count." min_transition_count may be determined during Auto Negotiation as a parameter requested by the link partner. Alternatively, a fixed value may be chosen, e.g., min_transition_count = 128, corresponding to 128 PMA frames = $128 \times 20.48 \mu\text{s} = 2.62 \text{ ms}$]

55.4.2.5.8 InfoField payload: coefficient exchange and coefficient format

The InfoField payload transmitted in state PMA_Coeff_Exch depends on the value of STF. With STF = 0, two coefficients are transmitted and reception of the two last received coefficients is acknowledged. With STF = 1, a state transition to state PMA_Fine_Adj is announced.

Coefficients are transmitted in the order A/1:2; A/3-4; ... A/15:16; B/1:2; ... B/15:16, C/1:2, ... D/15:16, where A,B,C,D stands for pair A,B,C,D, respectively, and the numbers represent the indices of the coefficients transmitted in the 8-bit fields coefficient_1 (Oct4<0>, Oct5<7:1>) and coefficient_2 (Oct5<0>, Oct6<7:1>). Each pair of coefficients is repetitively sent until the link partner acknowledges successful reception.

The 8-bit coefficients are represented in two's complement format *[sx.xxxxxx]* and can take values in the range -2.0 to 1.984375 in steps of 0.015625 ($=1/64$) *[In calculations of precoding responses for worst case channel conditions this author has found optimum coefficients slightly larger than 2. The performance penalty due to constraining coefficients to the above range may be acceptable. Alternatively, notice that in the proposed payload format offers room for 9-bit coefficients.]*

The 5-bit field coefficients_sent (Oct3<0>, Oct4<7:4>) indicates which two coefficients are currently transmitted: coefficients_sent = 00000, 00001, ... 11111 designates coefficients A/1:2, A/3:4, ... D/15:16, respectively.

The 5-bit field `coefficients_received` (Oct3<0>,Oct4<7:4>) indicates which two coefficients have last been received: `coefficients_received` = 00000, 00001, 11111 designates coefficients A/1:2, A/3:4, ... D/15:16, respectively. When no coefficients have yet been received, `coefficients_received` = 11111 is sent.

When all coefficients of the link partner have been received, and the link partner acknowledges reception of all coefficients, a transition to state `PMA_Fine_Adj` is announced. The condition is met when received and transmitted `InfoFields` convey `coefficients_received` = 11111.

[This coefficient exchange has been adapted from the 4-coefficient version of Draft 2.2. The simple handshake procedure appears to be adequate for the intended use. Let two coefficients be effectively transferred to the link partner in every 128 PMA frames. Then the transfer of the $4 \times 16 = 64$ coefficients is accomplished in $64/2 \times 128 \times 20.48 \text{ us} = 84 \text{ ms}$, which is a small portion of the allowed maximum link training time of 2 sec. --- A more efficient procedure could be as follows. All coefficients are cyclically inserted into consecutive `InfoFields`. The link_partner counts received `PMA_training` frames since the local PHY entered state `PMA_Coeff_Exch`. The link partner would know which coefficients are conveyed in which `InfoField`, and can focus decoding resources on `InfoFields` containing coefficients yet to be recovered. When all coefficients have been received, the link partner sets a bit "all_coefs_received" in the `InfoFields`. ...]

55.4.2.9 Startup sequence

When Auto Negotiation sets `link_control` from DISABLE to ENABLE, PHY Control proceeds from state PHY Disabled to state Silent.

In state Silent, the PHY remains silent (`tx_mode` = SEND_Z) for a fixed period of 1 ms. Then, if `config` = MASTER, PHY Control enters state `PMA_Train1_M`. Otherwise, `config` = SLAVE and PHY Control enters state `PMA_Train1_S`.

State `PMA_Train1_M`

In state `PMA_Train1_M`, the MASTER sends PMA training frames (`tx_mode` = SEND_T), which includes transmission of `InfoFields`. The MASTER increases transmit power in up to three steps, adjusts echo and next cancellers, and sends invitations to the SLAVE to start transmission at specified times until the start of SLAVE transmission is detected.

Upon entering state `PMA_Train1_M`, the MASTER sets `master_step` = 1. This corresponds to sending PMA training frames at minimum transmit power level by using `PBO` = 7 [*nominal power -14 dB*]. If the MASTER is required to use fixed TH precoding, `PBO` = 7 implies the fixed THP setting `THPShort`. In the `InfoField` the MASTER sends `current_PBO` = `requested_PBO` = 7. After the MASTER has sufficiently adjusted its echo and next cancellers, it sets `max_incr_timer` to $168 \pm 5 \text{ ms}$ [*≈ 8200 PMA training frames*] and begins to send invitations to the SLAVE to start transmission at specified times. The invitations are expressed by sending in the `InfoField` `next_PBO` = `current_PBO` together with a non-zero value of `transition_count`, i.e., by announcing a "zero power change". If the SLAVE is already able to decode `InfoFields` with an SNR margin of at least 6 dB for binary symbol decisions, SLAVE transmis-

sion with a PBO setting equal to the value in the received current_PBO = requested_PBO fields is started when the transition_count becomes zero. The start of transmission shall be timed such that at the MDI of the SLAVE [*the AFE of the SLAVE*] the first transmitted PMA training frame and the PMA training frame received following the InfoField in which the transition_count becomes zero are time aligned within ± 500 ns. This alignment ensures that the MASTER has to check for reception of the SLAVE signal only within a time window, whose width is determined by the alignment tolerance of ± 500 ns and the maximum round-trip propagation delay of the link. If the MASTER detects the beginning of SLAVE transmission in this window, it announces a transition to state PMA_Train2_M by inserting into the InfoField a non-zero transition_count together with STF = 1. Otherwise, the MASTER sends further invitations to the SLAVE at appropriate time intervals and checks for arrival of the SLAVE signal until the max_incr_timer expires.

When the max_incr_timer expires and the MASTER has not yet detected the SLAVE signal, the MASTER sets master_step = 2. This corresponds to sending PMA training frames at minimum transmit power level + 4 dB by using PBO = 5 [*nominal power -10 dB*]. If the MASTER is required to use fixed TH precoding, PBO = 5 implies the fixed THP setting THP_{Med}. The MASTER announces the increase in transmit power by sending in the InfoField current_PBO = 7 and next_PBO = 5 together with a non-zero value of transition_count. The change in transmit power occurs immediately after the InfoField, in which transition_count becomes zero. In the InfoField the MASTER then sends current_PBO = requested_PBO = 5. After the MASTER has sufficiently readjusted its echo and next cancellers, it sets max_incr_timer to 100 ± 5 ms [*≈ 4880 PMA training frames*] and begins to send invitations to the SLAVE to start transmission at specified times. Further operations are as described above for master_step = 1.

When max_incr_timer expires again and the MASTER has still not detected the SLAVE signal, the MASTER proceeds to master_step = 3. This corresponds to sending PMA training frames at minimum transmit power level +8 dB by using PBO = 3 [*nominal power -6 dB*]. If the MASTER is required to use fixed TH precoding, PBO = 3 implies the fixed THP setting THP_{Long}. The MASTER announces the increase in transmit power by sending in the InfoField current_PBO = 5 and next_PBO = 3 together with a non-zero value of transition_count. The change in transmit power occurs immediately after the InfoField, in which transition_count becomes zero. In the InfoField the MASTER then sends current_PBO = requested_PBO = 3. At this maximum transmit power level for state PMA_Train1_M, the MASTER continues operations as described above for master_step = 1 and 2, but without setting a time limit for the detection of the SLAVE signal. The dwell time is only limited by the expiration of link_fail_inhibit_timer in Auto Negotiation. If link_fail_inhibit_timer expires, PMA training has failed. Auto Negotiation sets link_enable = DISABLE and thus forces PHY Control into state PHY_Disabled.

State PMA_Train2_M

In state PMA_Train2_M, the MASTER sends PMA training frames (tx_mode = SEND_T), maintains echo and next cancellation, and trains its receiver section. When the MASTER has acquired the ability to decode InfoFields and assess SNR margin, it may send InfoFields to the SLAVE requesting a change in the SLAVE's PBO setting to the value in the requested_PBO field. Likewise, the MASTER may receive such requests from the SLAVE and respond by announcing a change to the requested transmit

power. The fixed THP setting shall not be affected by such power changes. The training of the receiver section includes the adaptation of feedback filter coefficients, which are sent to the SLAVE in state PMA_Coeff_Exch.

The objectives of transmit-power adjustments and filter adaptation are to achieve local and remote receiver operations with sufficient SNR margins for PCS operation with lowest possible transmit powers. This condition is met when `loc_rcv_status = OK` and `rem_rcv_status = OK`, where `rem_rcv_status` is the value of `loc_rcv_status` received from the SLAVE. *[Recall from 55.4.2.5.6 that `loc_rcvr_status` reflects a judgment of the local receiver on the projected decision-point SNR margin for reliable decoding of LDPC-coded 128 DSQ signals. Setting this bit may be delayed if PHY Control observes continuing improvements of the SNR margin by further filter adaptation.]* The MASTER then announces a transition to state PMA_Coeff_Exch.

State PMA_Train1_S

In state PMA_Train1_S, the SLAVE remains silent (`tx_mode = SEND_Z`). The SLAVE acquires timing from the received MASTER signal and trains its receiver section until it can decode InfoFields with an SNR margin of at least 6 dB for binary symbol decisions. The SLAVE checks the received InfoFields for an invitation to start SLAVE transmission. When an InfoField containing `current_PBO = next_PBO` together with a non-zero `transition_count` is decoded, i.e., an announcement of a "zero power change" is received, the SLAVE advances to state PMA_Train2_S.

State PMA_Train2_S

In state PMA_Train2_S, the SLAVE sends PMA training frames (`tx_mode = SEND_T`) initially using the PBO setting recovered from the `current_PBO (= next_PBO = requested_PBO)` field in the received InfoFields. If the SLAVE is required to use fixed TH precoding, the value of PBO implies the fixed THP setting to be used. The start of transmission shall be timed such that at the MDI of the SLAVE *[the AFE of the SLAVE]* the first transmitted PMA training frame and the PMA training frame received following the InfoField in which the `transition_count` becomes zero are time aligned within ± 500 ns. This alignment ensures that the MASTER has to check for reception of the SLAVE signal only within a time window, whose width is determined by the alignment tolerance of ± 500 ns and the maximum round-trip propagation delay of the link. The SLAVE then adjusts its echo and next cancellers *[in the presence the MASTER signal]* and reestablishes receiver operation. When the SLAVE has regained the ability to decode InfoFields and assess SNR margin, it may send InfoFields to the MASTER requesting a change in the MASTER's PBO setting to the value in the `requested_PBO` field. Likewise, the SLAVE may receive such requests from the MASTER and respond by announcing a change to the requested transmit power. The fixed THP setting shall not be affected by such power changes. The training of the receiver section includes the adaptation of feedback filter coefficients, which are sent to the MASTER in state PMA_Coeff_Exch.

The objectives of these adjustments are to achieve local receiver operation and remote receiver operation with sufficient SNR margins for later PCS operation with lowest possible transmit powers. This condition is met when `loc_rcv_status = OK` and `rem_rcv_status = OK`, where `rem_rcv_status` is the value of `loc_rcv_status` received from the SLAVE. If the MASTER then announces a transition to state

PMA_Coeff_Exch, the SLAVE also announces a transition to state PMA_Coeff_Exch. The state transition of the SLAVE shall not precede the corresponding transition of the MASTER. The SLAVE's transition should be scheduled to occur simultaneously with or shortly after the transition of the MASTER.

[From a logical viewpoint, it may be acceptable that the SLAVE enters PMA_Coeff_Exch before the MASTER. Notice that state transitions are primarily a matter of the transmitters. A receiver simply tracks the state of the remote transmitter.]

State PMA_Coeff_Exch

In state PMA_Coeff_Exch, MASTER and SLAVE operations are symmetric. The PHY sends PMA training frames (tx_mode = SEND_T). The PBO and fixed THP settings remain unchanged. The exchange of the 4 x 16 coefficients is explained in 55.4.2.8.

When all coefficients of the link partner have been received and the link partner acknowledges reception of all coefficients, a transition to state PMA_Fine_Adj is announced. The state transition of the SLAVE shall not precede the corresponding transition of the MASTER. The SLAVE's transition should be scheduled to occur simultaneously with or shortly after the transition of the MASTER.

State PMA_Fine_Adj

In state PMA_Fine_Adj, MASTER and SLAVE operations are symmetric. The PHY sends PMA training frames (tx_mode = SEND_T). The PBO settings remain unchanged. TH precoding is enabled with the coefficients received from the link partner.

The objectives of operation in state PMA_Fine_Adj are to ensure proper reception of TH precoded PAM-2 symbols on both sides of the link and achieve final convergence of all adaptive filters. *[There should be a lower limit on the time spent in state PMA_Fine_Adj. In Draft 2.2, 500 ms is mentioned; it is not clear whether this is "allowed" minimum time or mandated minimum time.]* When loc_rcv_status = OK and rem_rcv_status = OK, a transition to state PCS_Test is announced. The state transition of the SLAVE shall not precede the corresponding transition of the MASTER. The SLAVE's transition should be scheduled to occur simultaneously with or shortly after the transition of the MASTER.

State PCS_Test

In state PCS_Test, MASTER and SLAVE operations are symmetric. The PHY sends a fixed number of 3125 PCS frames (tx_mode = SEND_N) for link testing purposes *[3125 x 320 ns = 1 ms]*. The link partners can no longer exchange transmission parameters or state information. After transmission of 3125 PCS frames, PHY Control transitions to state PCS_Data.

State PCS_Data

In state PCS_Data, MASTER and SLAVE operations are symmetric. The PHY transmits PCS data frames (tx_mode = SEND_N). If more than 3125 PCS frames have been received and PCS_status = OK, PHY Control asserts link_status = OK.

Return to state Silent

If in state PCS_Test or PCS_Data an error situation occurs indicated either by loc_rcvr_status = NOK or after reception of at least 3125 PCS frames PCS_status = NOK, then PHY Control returns to State Silent where link_status is set to NOK. After a silent period of 1ms, PMA retraining is started.

Returns to state Silent upon detecting error situations in the PMA training states are left to implementers.

[Section 55.3 in Draft 2.2 should define the effect of link_status on the PCS RECEIVE and PCS TRANSMIT functions. This author has not found this definition. Where is it?.....]

55.4.2.10 PHY Control variables, timers, and functions**config**

This variable is defined in 55.2.2.2.
Values MASTER or SLAVE.

link_control

This variable is defined in 28.2.6.2.
Values DISABLE or ENABLE.

link_status

This variable is defined in 45.2.7.2.6.
Values FAIL or OK

(this section is to be completed)

55.4.2.8.11 PHY Control state diagram

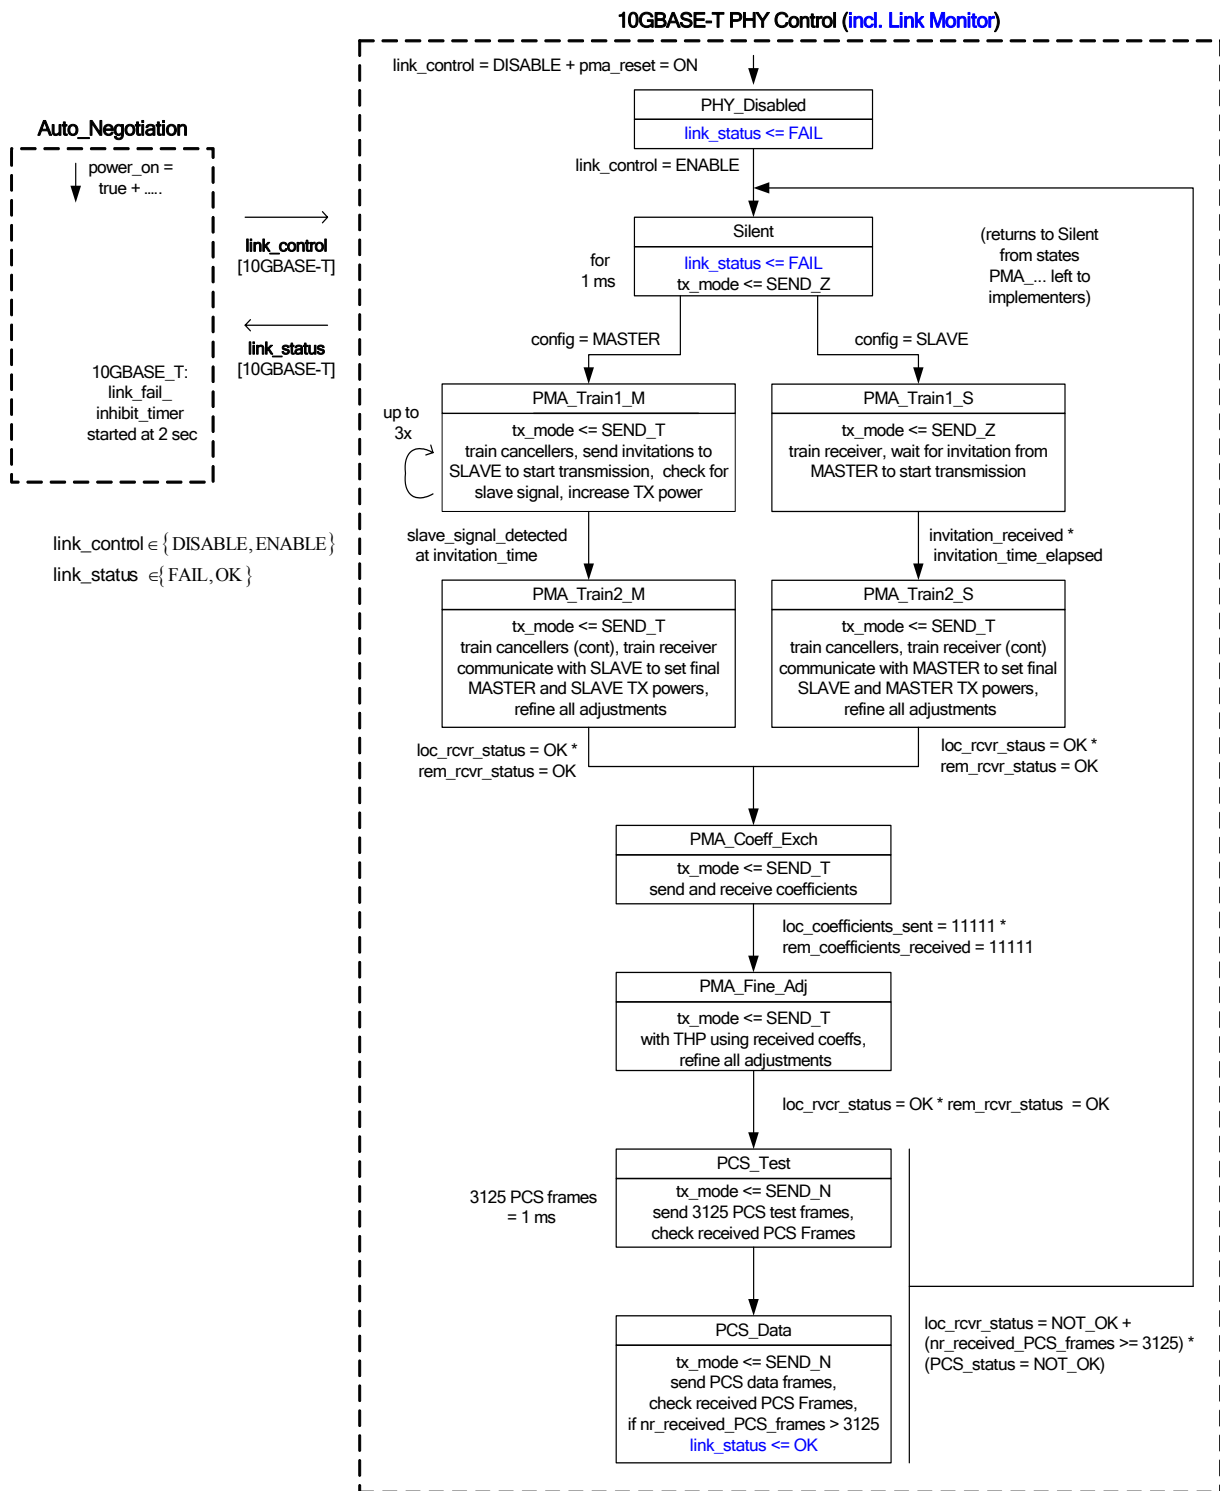


Figure 55-cc— PHY Control state diagram and relation with Auto Negotiation