ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

PHY Control for 100BASE-T1L

Philip Curran
Brian Murray
Jacobo Riesco

Introduction

- The 802.3dg task force has already had a number of general presentations on training, based on the approaches used in clauses 97, 149 and 165
  - Including descriptions of training frames consistent with the adopted proposals for 8b6T encoding using 8N/(8N+1) block codes and an RS-FEC
    - Low latency mode using a 16B/17B (N=2) block code
    - Burst error protection mode using a 64B/65B (N=8) block code and RS FEC
  - And there has been discussion and general agreement on using PAM2 during training and switching to PAM3 for idle and data
- This presentation is a proposal for 100BASE-T1L link establishment and PHY Control

Purpose of Training

- The purpose of training in 100BASE-T1L is to achieve the following
  - Sequence the bring-up of the Leader and Follower PHYs using PAM2 signaling
    - PAM2 has the advantage of robustness during start-up and blind acquisition as it is easier to open the eye
  - Exchange the PHY capabilities between the local and remote PHYs
  - Align the 8N/(8N+1) and RS-FEC frames
  - Stage the transition from PAM2 to PAM3 signaling
    - Allow a short period of time to allow for any adjustment to the adaptive filter coefficients when the signaling transitions from PAM2 to PAM3
  - Synchronize the transition to idle/data
- The approach taken in this proposal is to support the necessary requirements of training, drawing from previous clauses but without carrying unnecessary elements added over the years

PHY Start-up for Echo Cancelling PHYs

- The Leader/Follower architecture has been used in high-speed echo cancelling Ethernet PHYs for many generations
  - The general approach is quite similar across the different PHY generations
  - Sequence when the Leader and Follower start to transmit to allow defined times for the Leader and Follower PHYs to acquire timing, to train their echo cancellers and equalizers
    - The transition between the various states is controlled by timers or by signaling between the Leader and Follower
- In Clause 55, 10GBASE-T introduced the concept of the PMA Training frame and InfoField
  - In 10GBASE-T Tomlinson-Harashima Precoding (THP) was used and it was necessary to exchange a set of coefficients before the receiver was fully up and achieving its full receiver gain
  - The inclusion of an InfoField allowed the exchange of data before the receiver is fully up by protecting the contents of the InfoField with a CRC and allowing it to be sent multiple times
  - This approach was continued in clauses 97, 149 and 165 even though the information to be exchanged between the PHYs was not required for PMA training
    - This information could be exchanged after both sides have asserted *loc_rcvr_status* = OK

PHY Start-up for Long Reach Echo Cancelling PHYs

- There are some differences for long reach echo cancelling PHYs like 100BASE-T1L compared to short reach PHYs that are worth highlighting
  - 100BASE-T1L operates over 500m of reach at 100 Mb/s compared to less than 40m for the more recent 1G and Multi-G single pair Ethernet PHYs
  - At 500m the insertion loss at Nyquist (40 MHz) is almost 37 dB and is much greater than for the short reach BASE-T1 PHYs
    - This results in a much greater ratio of the local echo signal power to the received signal power
    - It is expected that when the Follower starts transmitting, there may be a period of time when the transmit timing is not reliable
    - This requires additional safeguards around the behaviour of the PHYs while the echo cancellers are adapting
- This is addressed for 100BASE-T1L in this presentation
  - We follow the same approach used in 10BASE-T1L to ensure the Leader knows the time when the Follower Transmit timing is guaranteed to be reliable

Proposed Leader / Follower PHY Start-up

- The Leader / Follower PHY start-up is as follows
  - Training starts with the Leader PHY transmitting PAM2 and the Follower PHY receiving; the Leader trains its echo canceller and the Follower acquires timing lock and trains its equalizer
  - Then the Follower starts transmitting PAM2, trains its echo canceller and reacquires timing lock
  - After allowing time for the Follower to lock timing, the Leader trains its equalizer
  - When the Follower has trained its receiver, it signals *loc_rcvr_status* OK to the Leader
  - When the Leader has trained its receiver and has detected *rem_rcvr_status* OK, it transitions to transmitting PAM2 training frames
  - The Follower synchronizes its frame structure to that of the Leader and then transitions to transmitting PAM2 training frames
    - Thus, the Leader and Follower 8N/(8N+1) and RS-FEC frames are synchronized
    - And PHY capabilities are exchanged
  - The Leader and Follower then transition from PAM2 to PAM3 idle signaling
    - An allowance for PAM3 fine tuning is included here
  - And then the Leader and Follower transition to transmitting normal idle / data
  - Assertion of *link_status* OK is synchronized between the Leader and Follower

Leader / Follower PHY Start-up Flow Diagram

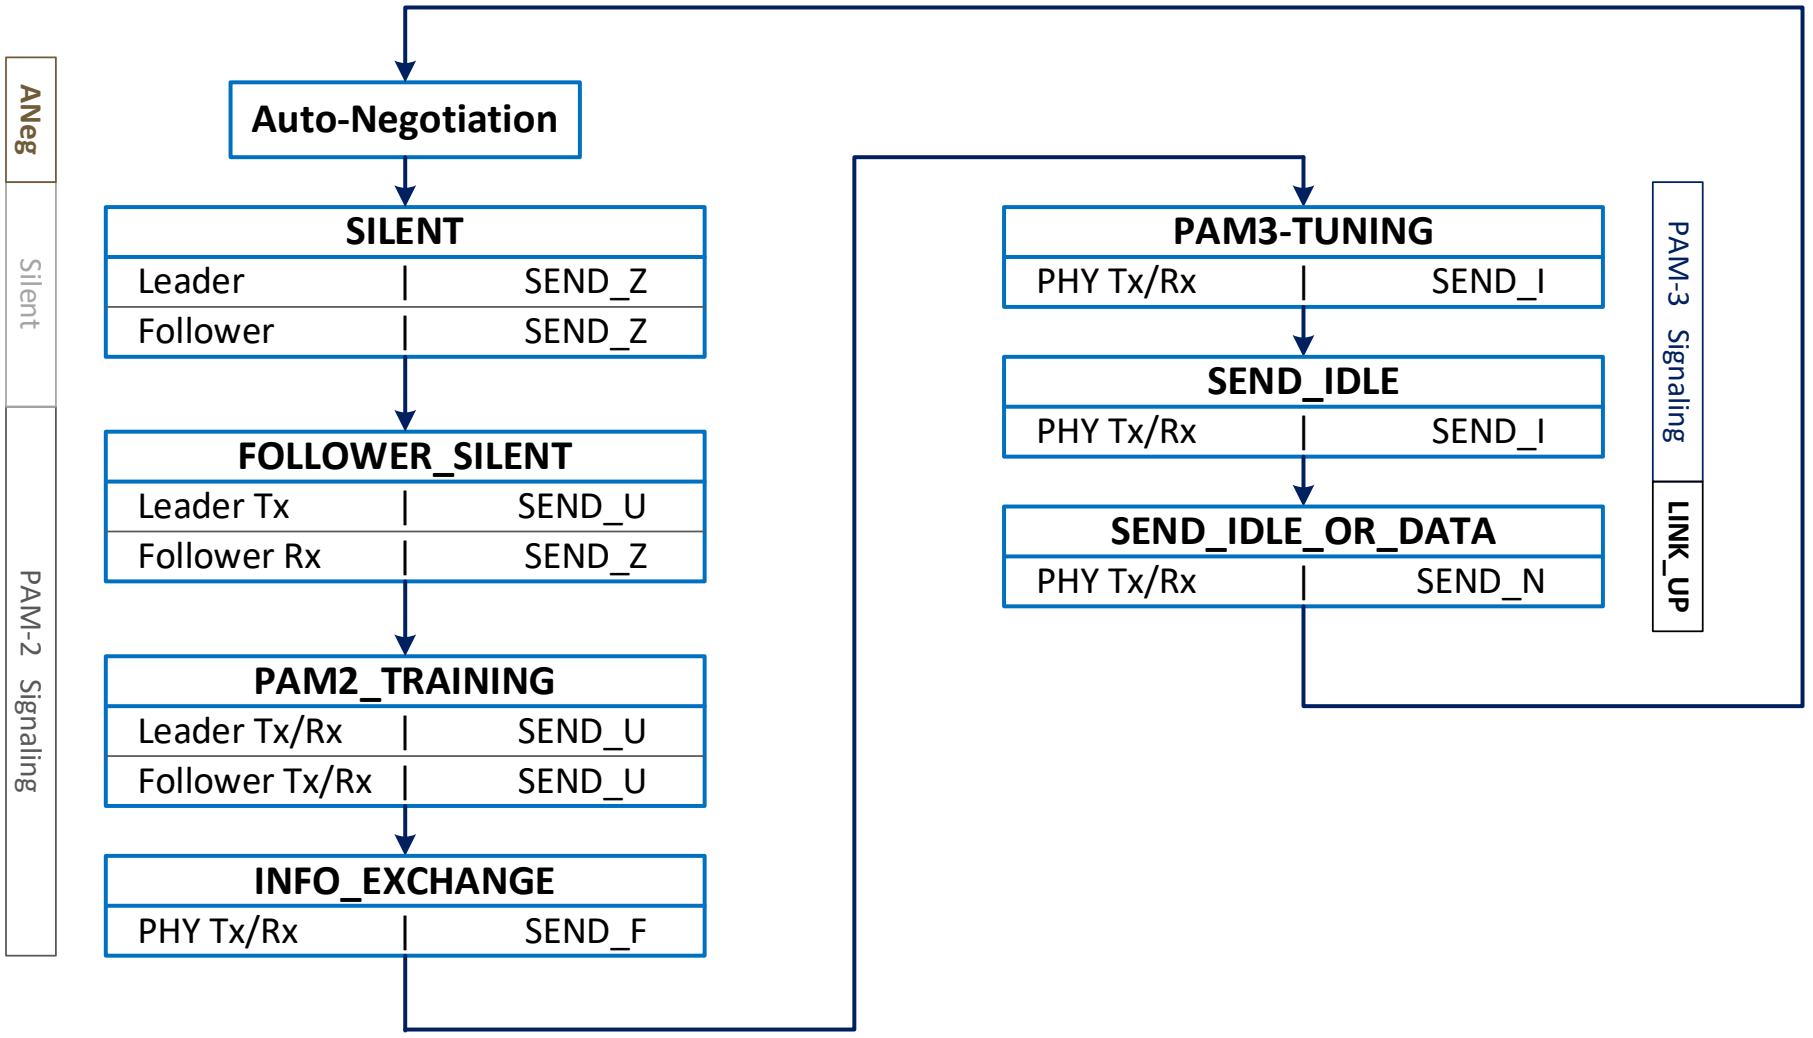

Timers for 100BASE-T1L PHY Control State Diagram

- The following timers are proposed for the 100BASE-T1L PHY Control state diagram
  - All timers are specified to ± 1000 parts per million
- **silent_timer** - 1 ms ± 1 µs
  - This timer ensures a guaranteed silent time after Auto-Negotiation completes for the Leader and Follower prior to technology dependent signaling
  - After the *silent_timer* expires the Leader starts transmitting unformatted sequences of PAM2 symbols and the Follower stays silent

Timers for 100BASE-T1L PHY Control State Diagram

- **min_follower_silent_timer** - 15 ms ± 15 µs
  - This is the minimum time that the Follower must stay silent
  - This guarantees the Leader a minimum time that the Follower is silent so it can train its echo canceller
  - The Follower may decide to stay silent for a longer time as may be required to train its receiver
  - When the Follower starts transmitting, it transmits unformatted sequences of PAM2 symbols
    - The Follower encodes *loc_rcvr_status* into this unformatted sequence
- This is a simple scheme to guarantee a silent time for the Leader
  - It is not necessary for the Leader to signal to the Follower when it is ok to start transmitting
    - That method requires that the Follower fully bring up its receiver so that it can reliably receive data from the Leader and in practice the signaling from the Leader is redundant
  - It is much simpler to just use a timer

Timers for 100BASE-T1L PHY Control State Diagram

- **follower_init_timer** - 40 ms ± 40 µs
  - This timer is started at the same time as the *min_follower_silent_timer*
  - This is the maximum time that the Follower is allowed before its transmit timing is guaranteed to be reliable
    - It implicitly imposes a limit on how long the Follower can be silent
  - At the end of the *follower_init_timer* the Leader is guaranteed that the Follower has locked its timing and that the Leader can train its receiver
    - Note, the Leader can start to train its receiver before this timer has expired, but it is not guaranteed that the Follower timing is locked
  - The Follower signals *loc_rcvr_status* OK to the Leader when it has trained its receiver
  - When the Leader has trained its receiver and has received *rem_rcvr_status* OK from the Follower it transitions to transmitting PAM2 formatted frames
- Previous standards have omitted such a timer
  - In 100BASE-T1L with the insertion loss at 500m and consequently the large ratio of echo to received signal power it is possible that there will be a period of time after the Follower starts transmitting when the transmit timing is unreliable
  - We also need to impose a limit on how long the Follower can remain silent

Timers for 100BASE-T1L PHY Control State Diagram

- After the Leader and Follower exchange PAM2 formatted frames, they switch to transmitting PAM3 idle control codes
- **min_pam3_tuning_timer** - 5 ms ± 1 µs
  - In the PAM3_TUNING state the adaptive structures are tuned for PAM3 signaling
    - The adaptation of the coefficients is dependent on the auto-correlation of the data, so it is prudent to allow some time for this adjustment
  - We exit the PAM3_TUNING state when *loc_rcvr_status* = OK
    - The implementer may have a more stringent test for *loc_rcvr_status* at this point
  - The *min_pam3_tuning_timer* is to prevent us dropping straight through this state

Transmit Signaling – SEND_Z and SEND_U

- At different stages of the training different transmit signaling is used
- **SEND_Z** – transmit zero
  - SEND_Z is used in the SILENT state, the FOLLOWER_SILENT state and the LINK_FAIL state
- **SEND_U** – transmit unformatted sequence of PAM2 symbols
  - The first stage of training transmits PAM2 with no format or frame structure
  - The PAM2 6-tuple is selected from scrambler bits $Sy_n[3:0]$ using the table of 4B6B NND PAM2 6 tuples provided in Tingting_3dg_01_18_12_2024
  - When the Follower has brought up its receiver, it signals *loc_rcvr_status* OK by inverting bit 3
    - This is a very simple scheme and is very easy to detect as *loc_rcvr_status* is encoded every 6-tuple
  - Running disparity is controlled as described in Tingting_3dg_01_18_12_2024

Transmit Signaling – SEND_F

- **SEND_F** – transmit formatted training frames
  - The second stage of training transmits PAM2 formatted training frames
  - Each PCS partial frame is 32 6-tuples long (128 bits during training)
  - Each training frame is composed of 16 PCS partial frames
- All the bits in the training frame are zero except for:
  - The second bit in all PCS partial frames is set to 1, as shown in red
  - The InfoField, as shown in blue

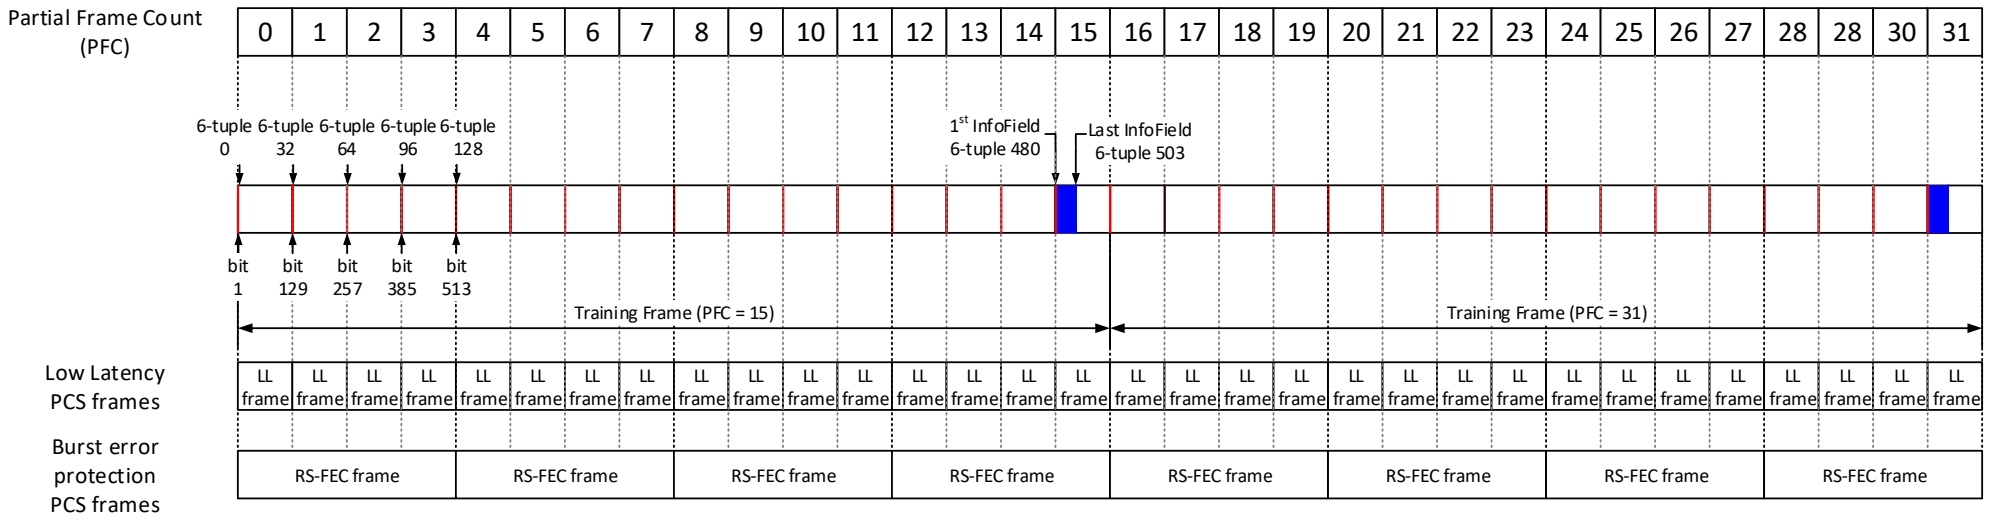

Transmit Signaling – InfoField Format

- The definition of the SEND_F formatted frames is such that:
  - The second bit of each PCS partial frame is set to 1 to facilitate alignment
  - The 3rd bit of the 16th PCS partial frame is set to 1 to facilitate PCS frame alignment
  - The first 12 octets (96 bits) in the 16th PCS partial frame are XORed with the contents of the InfoField[0:95]
    - The InfoField allows the PHY Capabilities to be exchanged between Leader and Follower
- Format of the InfoField

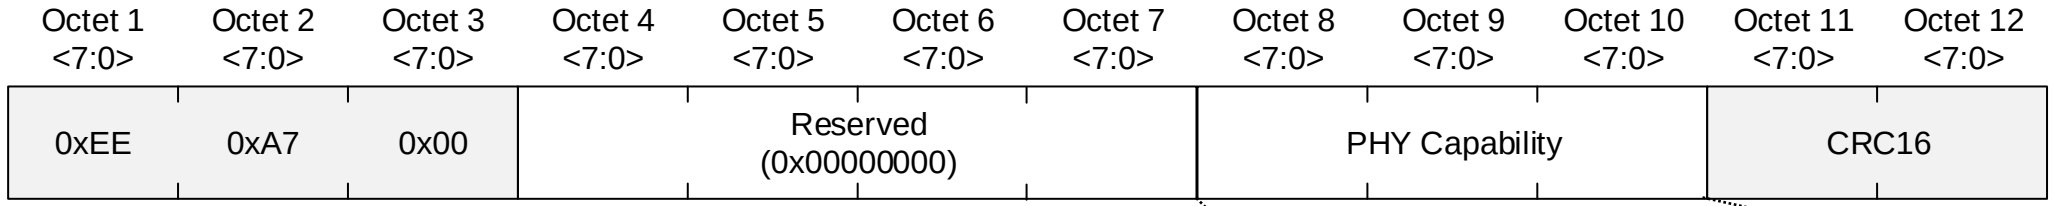

  - The InfoField has a Header, a PHY Capability field and a CRC
    - *rs_adv* advertises support for RS-FEC
    - *eee_adv* advertises support for Energy Efficient Ethernet
    - *lpi_adv* advertises support for Low Power Idle
    - *seq_adv* advertises support for sequence ordered sets

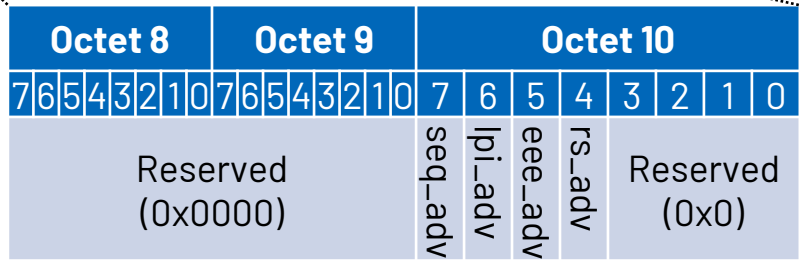

Frame Alignment

- When the Leader has trained its receiver and has received *loc_rcvr_status* = OK from the Follower it starts to transmit PAM2 formatted frames
  - The formatted frames allow the Follower to synchronize its PCS frame
  - When the Follower starts receiving formatted frames and detects a valid InfoField it can synchronize its training frame to the Leader
  - The start of the Follower training frame is delayed by not more than one PCS partial frame with reference to the start of the Leader training frame as seen at the Follower MDI
    - This aligns the Follower frame to the Leader frame
    - This ensures that the Energy Efficient Ethernet LPI quiet / refresh cycles are offset
  - This occurs when the Follower is in the PAM2_TRAINING state
- The Follower only starts sending formatted frames to the Leader after it has synchronized its frame structure to that of the Leader

PHY Capability Exchange

- PHY capabilities are exchanged using the InfoField during training
  - Each PHY advertises its capabilities to the link partner in the InfoField
  - Each PHY determines the capabilities of the link partner when it receives a valid InfoField
    - An InfoField is considered valid if the CRC is correct
  - The PHY Control function resolves the abilities
    - An ability is only enabled if advertised by both sides
    - The resolved abilities are communicated by PHY Control to the PCS
  - This occurs when the PHY is in the INFO_EXCHANGE state
- Once the PHY receives a valid InfoField, it sends 3 more formatted frames
  - This is to ensure that the link partner has received a valid InfoField

Transmit signaling – SEND_I

- **SEND_I** – transmit idle control codes
  - Transmit PAM3 symbols using 8b6T encoding and an 8N/(8N+1) block code representing control information corresponding to the idle mode
    - These can be Normal Inter-Frame (Idle) with *loc_phy_ready* = TRUE or = FALSE
    - This is the transmission of **I** or **Ix** control codes
- PAM3_TUNING and SEND_IDLE states
  - After the PHY has sent the last formatted training frame it transitions to transmitting **Ix** control codes
  - And the PHY receiver starts looking to detect PAM3 signalling
    - This is the detection of a signal is that is compatible with PAM3 signaling from the remote PHY and is not compatible with PAM2 signaling from the remote PHY
  - When the PHY has detected PAM3 signalling it waits in the PAM3_TUNING state for a minimum of the *min_pam3_tuning_timer* time
  - If *loc_rcvr_status* = OK it starts looking for idle signalling from the link partner
  - And if *loc_rcvr_status* = OK and the remote PHY is sending idle it transitions to transmitting **I** control codes

Transmit Signaling –SEND_N

- **SEND_N** – transmit normal data / idle
  - Transmit PAM3 symbols using 8b6T encoding and an 8N/(8N+1) block code representing an MII stream of data and control
- SEND_IDLE_OR_DATA state
  - When the PHY detects **I** or **Ix** control codes it moves to the SEND_IDLE state and transitions to transmitting **I** control codes
  - When the PHY detects **I** control codes it moves to the SEND_IDLE_OR_DATA state and sets *tx_mode* to SEND_N

Link Monitor and LINK_FAIL State

- LINK_FAIL state
  - If at any point the operation of the receive link for the local PHY is unreliable the *loc_rcvr_status* is set to NOT_OK
  - The PHY moves to the LINK_FAIL state and sets *tx_mode* to SEND_Z
  - There is no retrain
- Link Monitor
  - When *tx_mode* is set to SEND_N the Link Monitor changes from LINK_DOWN to LINK_UP state and *link_status* is set to OK
  - When *tx_mode* is set to SEND_Z the Link Monitor changes from LINK_UP to LINK_DOWN state and *link_status* is set to FAIL

Link Status Synchronization for 100BASE-T1L

- The transition between SEND_IDLE and SEND_IDLE_OR_DATA is a combination of *loc_rcvr_status* = OK **AND** *rem_phy_ready* = OK
  - There are no timers involved so no uncertainty added
  - The decoding of *rem_phy_ready* is very fast with low latency

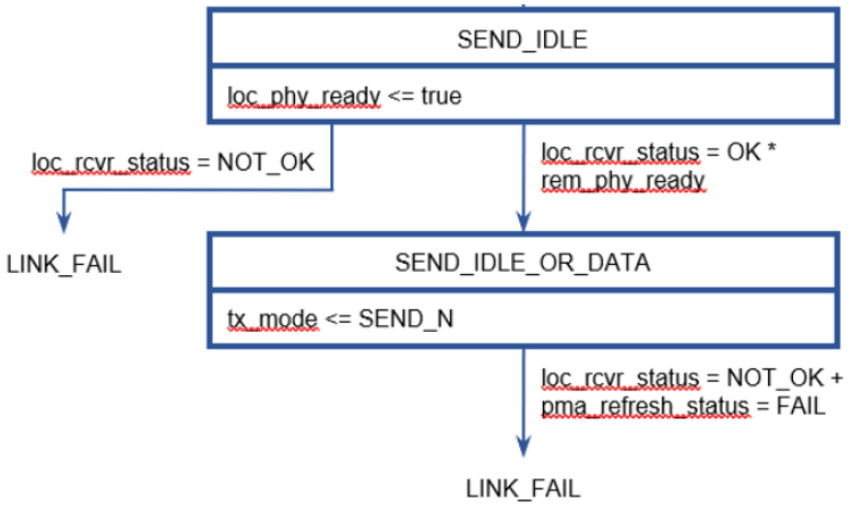

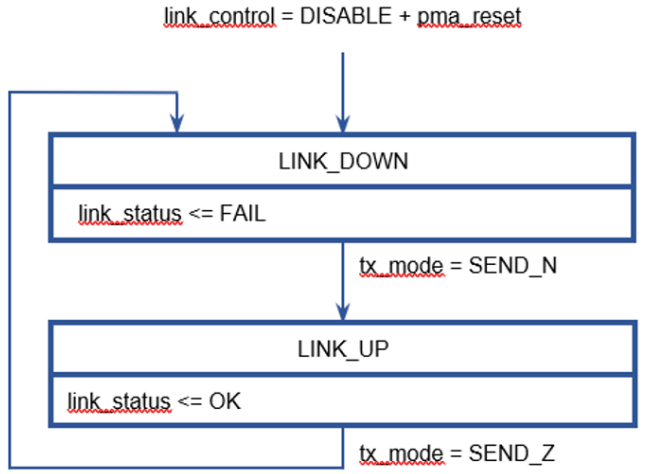

The rem_phy_ready parameter is set to FALSE when the pcs_rx_mode parameter is TRAIN. A sequence of symbols representing idle mode with PHY ready encoding is detected if 4 consecutive control characters are received, each of which is /I/.

Summary

- We have explicitly defined a means for the Leader to know the Follower timing is locked
  - Clauses 97, 149 and 165 all use a scheme (*timing_lock_OK*) that does not work in general
- In clauses 97, 146, 149 and 165 there is no limit on how long the Follower can stay silent
  - There is an implicit assumption that the available time to bring up the link (<100 ms) is shared equitably between the Leader and Follower
  - We have a defined limit on how long the Follower can stay silent
- There are unnecessary complications in clauses 97, 149 and 165 in the InfoField
  - There is no requirement for the countdown and there is no need to transmit a partial frame count
  - There is no need to do the PHY Capability exchange before the receiver is fully up
  - We have not carried unnecessary complications from these clauses into 100BASE-T1L
- We have addressed the issue of Link Status synchronization
  - We have removed the unnecessary time variability in clause 97
  - The approach used in clauses 149 and 165 are not applicable as we have to be able to work with existing clause 22 MAC devices

Conclusion

- We have presented a proposal for PHY Control for 100BASE-T1L
- We have addressed issues in previous clauses as follows
  - Have a simple and reliable scheme to determine the start of Follower transmission
  - Ensure the Leader knows exactly when the Follower transmit timing is reliable
  - Removed complications in the InfoField and frame synchronization in previous clauses that are not necessary for 100BASE-T1L
  - Ensure good link status synchronization
- We support the necessary requirements of training, drawing from previous clauses but without carrying unnecessary elements added over the years

Questions ?