Update on Technical Feasibility of EEE with 10GBASE-T

George Zimmerman

Shalini Rajan

Outline

- Goals for EEE Transitions
- Strawman for EEE in 10GBASE-T
- Restart Sequence
- Restart Time Considerations – revisited
  - Minimal Changes
- Laboratory Results
- Conclusions

Goals for EEE Transitions

- Be safe: do no harm
  - Base results on WORKING systems
  - No change to operational mode of existing PHYs
- Be lazy: don't invent unnecessary things
  - Transition would minimally impact existing specifications
    - Reuse of existing 802.3an PHY control as much as possible
- Be quick: get PHY transition times down
  - Need for transitions of <10msec, pref ~1msec
  - Need to minimize retraining time

Strawman for EEE in 10GBASE-T

- Freeze stored 10GBASE-T state while 1000BASE-T or other lower speed is running
  - Feasibility Question: How long before the transceiver state typically gets stale?
    - Answered in zimmerman_1_0307: ~ 5 minutes (reconfirmed)
    - Repeated experiments suggest it is limited by environmental changes and timing stability.
- Restart 10GBASE-T transmission by entering final stages of PHY-control startup sequence
  - Feasibility Question: How short might a transition be made using the existing framework with minimal modifications?
    - zimmerman_1_0307 suggested a blind use of the protocol would yield 10-20msec transitions into 10GBASE-T
    - NEW DATA suggests minimal adjustment of the timing parameters can minimize transition time to ~1-4msec

Restart Sequence

- 10GBASE-T PHY Control State Machine (Fig. 55.4.6.1)
- Entrance points for EEE state-restoral:
  - PMA_Fine_Adjust or
  - PCS_Test (1msec fixed)
    - Required to maintain quality
    - Test time limited by desire to see enough LDPC frames
  - Only PMA_Fine_Adjust entry considered here
- Full Retrain triggered if PCS_Test fails, dropping link

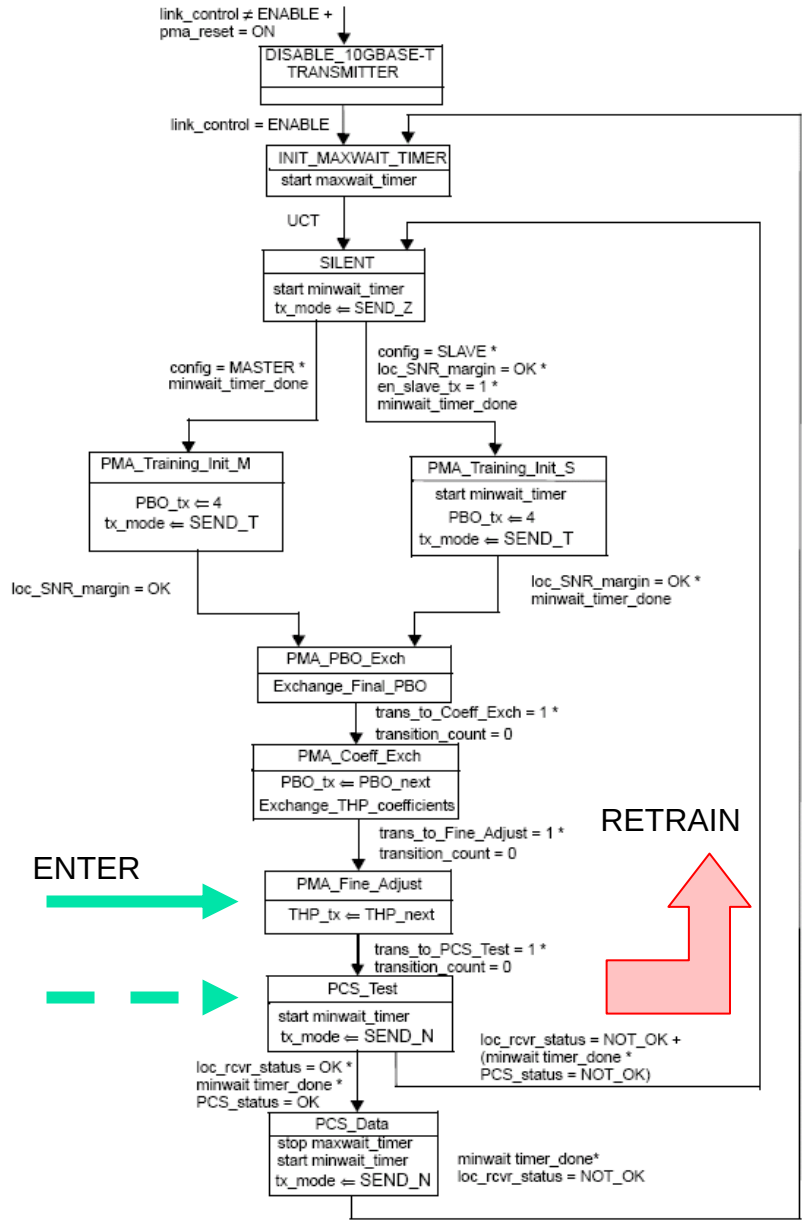

Restart Time Considerations - revisited

- PCS_Test time is 1msec
  - Any signaling or entry time is in addition
- When PMA_Fine_Adjust is required, time is controlled by Infofield protocol
  - Infofield every 16384 symbols (20.48 usec)
  - Count down 512 infofields
  - 10.48 msec state transition + retraining time + PCS_Test (1msec)

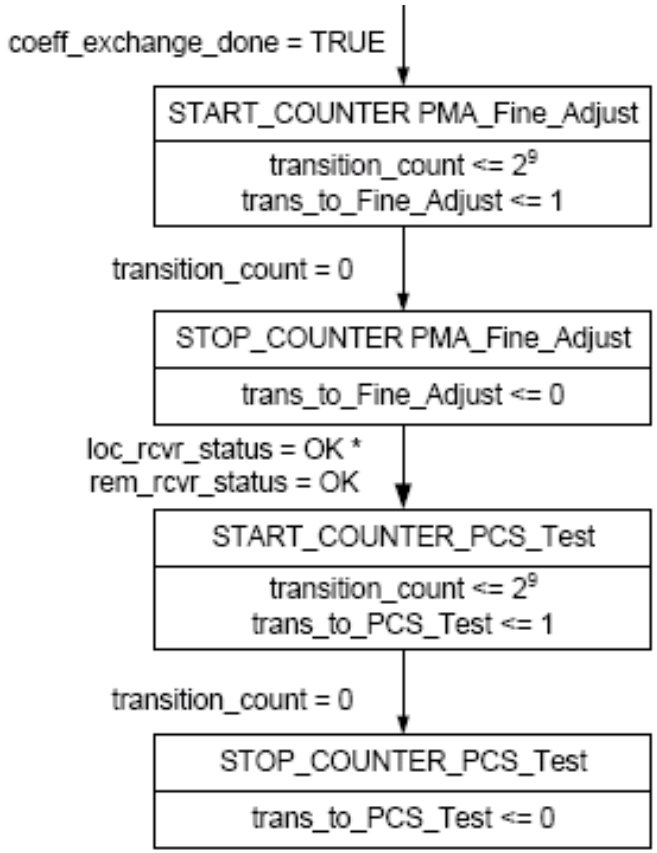

- Conclusion: Transition time is controlled by 512 infofield count
  - Rationale was design allowed for controller sync to be sloppy – not consistent with EEE assumptions of ~1msec transitions

Minimal Changes

- Change transition_count value in Fig. 55-25 (MASTER transition count) from 2^9 to 2^3
  - Corresponding change of transition count for response in Fig. 55-26 (Slave) from 2^6 to 2^2
- Minimum PMA_Fine_Adjust time reduced from 10.5 msec to 164 usec
  - Limitation should now be training time, not protocol
  - Still allows plenty of time (>50usec) for Master-Slave state change synchronization
- 1msec PCS_Test state time remains

- **Enables transitions down to ~1-2 msec by reducing unnecessary overhead with minimal standards changes**

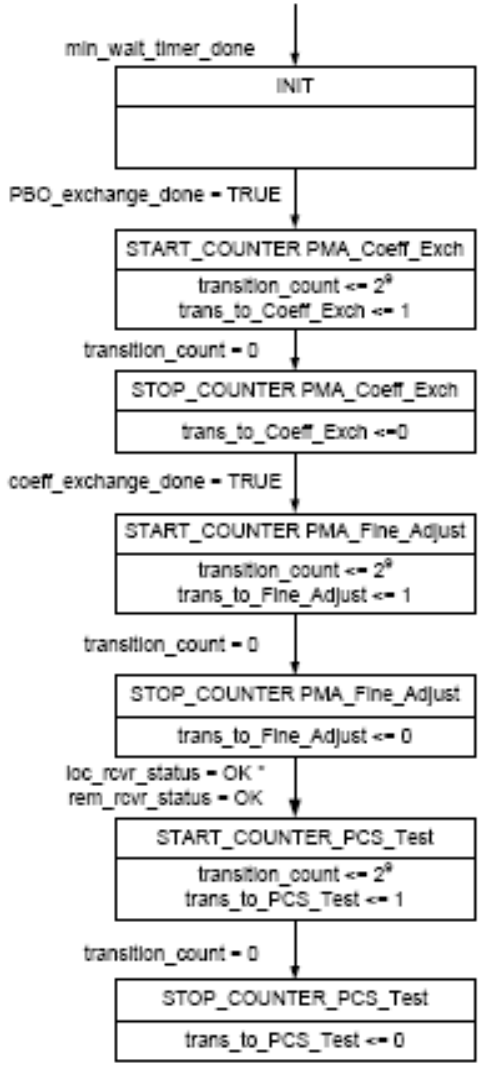

Laboratory Results

- Question: Can a transceiver be fine-trained in 1-5 msec?
  - Link setup with 4 connector, 100m channel
  - Link trained with modified counters on transition to PCS_test
    - Transition counters reduced as described previously
  - Timing and phase readjust at entry to PMA_Fine_Adjust state
    - Act as though link was being recovered from a stored state
    - Stable timing at << 1 msec
    - Similar results should be feasible for other than 10GBASE-T
  - Training time at PMA_Fine_Adjust varied to determine limitation
    - Assumes worst-case re-entry (all equalizers and cancellers need adjustment)
- Consistent demonstrations show SNR and Ethernet Frame Error Rate are uncompromised by 3-4 msec retrain time
  - Additional training improvements (vendor-specific, without standards changes) are likely to improve transition time

Conclusions

- Fast restart can reuse existing PHY control states
  - Fine Adjust, PCS_Test and retrain mechanisms exist
- Restart time is largely controlled by overly conservative transition count-down
- Simple changes to the transition counter allow development of transitions within ~1-2msec
- Laboratory results demonstrate the feasibility of 3-4msec retrains today, even on 100m 10GBASE-T links