ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

Link Status Synchronization for 100BASE-T1L

Philip Curran
Brian Murray
Jacobo Riesco

Introduction

- This is a proposal for link status synchronization for 100BASE-T1L
  - In Industrial Ethernet applications fast link-up time is important
  - Link Status is commonly signaled on a H/W pin to MAC devices and its exact timing is much more important than for applications that just rely on register reads to determine link status
- In some early PHY standards, such as 100BASE-TX there is very poor synchronization between the time that the link partners assert *link_status* OK
- The link status synchronization is much better in later single pair Ethernet standards
  - The link status synchronization is these standards meets the requirements of their target applications – but we believe 100BASE-T1L can and should have better synchronization
  - This presentation reviews the approach for link status synchronization used in some of these standards
  - This is useful as a guide and background to the proposal for 100BASE-T1L

Link Status

- Link status is available via the *link_status* signal
  - *link_status* = FAIL: A valid 100BASE-T1L link is not established.
  - *link_status* = OK: A valid 100BASE-T1L link is established
- Link Status is available in bit 2 of the MII Status register (address 1)

22.2.4.2.13 Link Status

When read as a logic one, bit 1.2 indicates that the PHY has determined that a valid link has been established. When read as a logic zero, bit 1.2 indicates that the link is not valid. The criteria for determining link validity is PHY specific. The Link Status bit shall be implemented with a latching function, such that the occurrence of a link failure condition will cause the Link Status bit to become cleared and remain cleared until it is read via the management interface. This status indication is intended to support the management attribute defined in 30.5.1.1.4, aMediaAvailable.

  - Many PHY devices make this signal available on a H/W pin
- This signal is embedded in a number of Industry Standard MAC interfaces
  - RGMII
  - SGMII
  - USGMII

Link Status Synchronization – 100BASE-TX

- The transition to LINK_UP (*link_status* = OK) **only** depends on local variables, e.g. *signal_status* = ON
  - Some PHYs can bring up a 100BASE-TX link in as little as 10 ms
  - But the link partner can also take up to the time the *link_fail_inhibit_timer* expires to establish its link which could be as long as 1 s

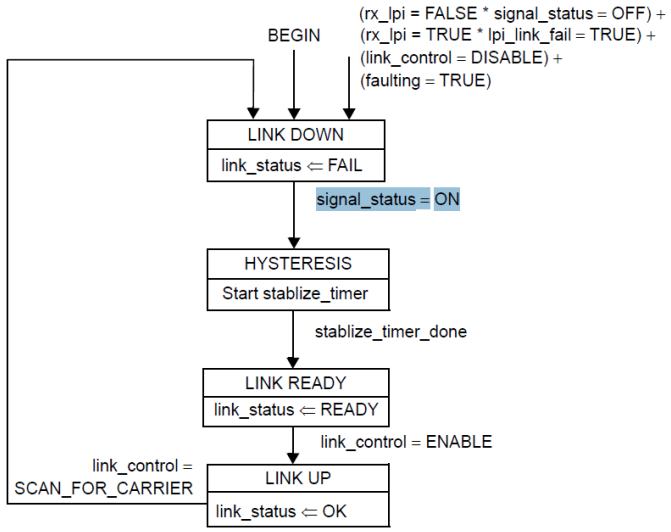

NOTE 1—The variables link_control and link_status are designated as link_control_[TX] and link_status_[TX], respectively, by the Auto-Negotiation Arbitration state diagram (Figure 28–18).

NOTE 2—The variables rx_lpi and lpi_link_fail are only required for the EEE capability and should be treated as if the value of these two variables is FALSE otherwise.

Figure 24–15—Link Monitor state diagram

signal_status

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

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

Link Status Synchronization – 1000BASE-T

- The transition between SEND_IDLE and SEND_IDLE_OR_DATA is a combination of *loc_rcvr_status* = OK **AND** *rem_rcvr_status* = OK, which is good
  - And the *stabilize_timer* has a tight bound of ± 0.1 µs, which is good
  - But in the Link Monitor the transition to LINK_UP (*link_status* = OK) **only** depends on *loc_rcvr_status* = OK

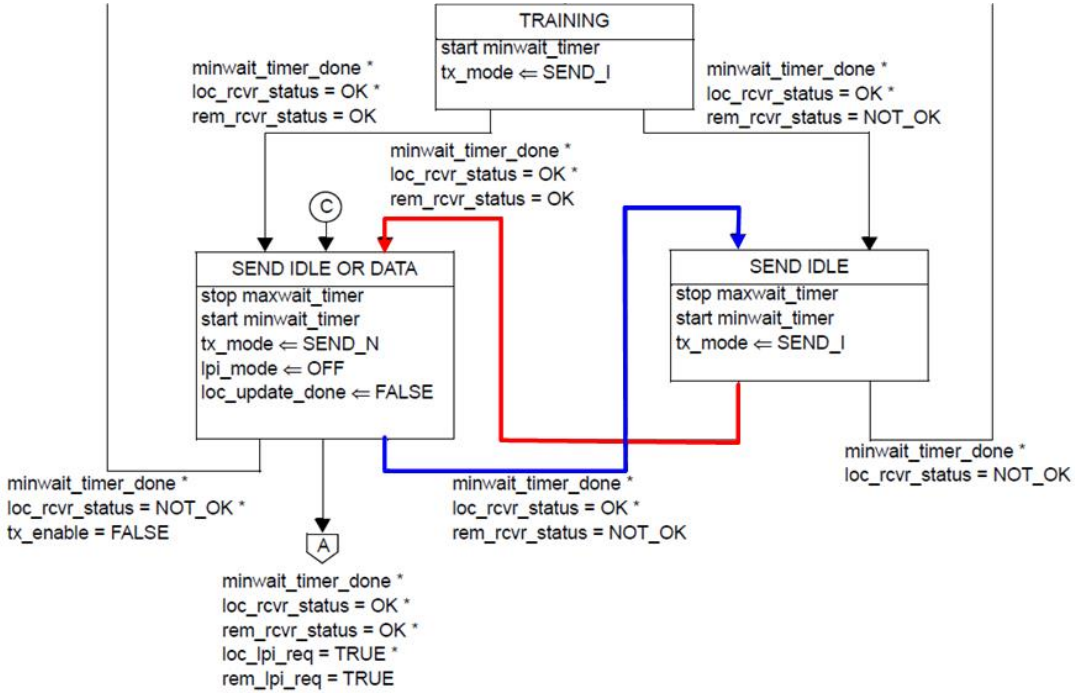

Figure 40–16a—PHY Control state diagram, part a

40.4.6.2 Link Monitor state diagram

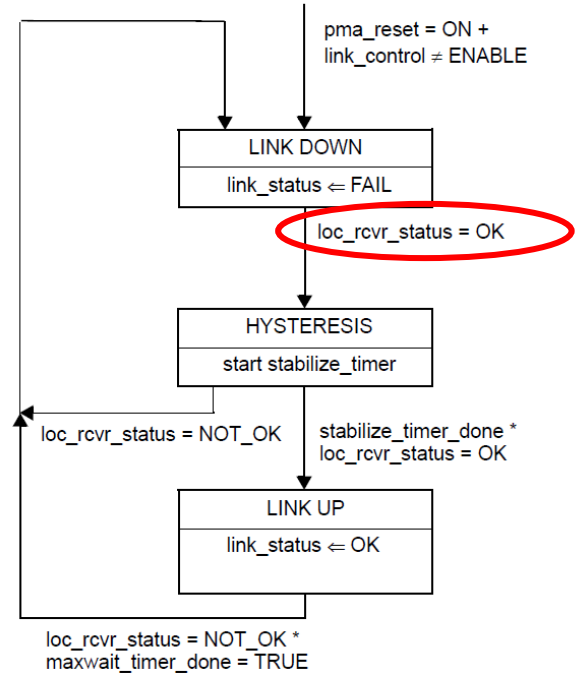

stabilize_timer
A timer used to control the minimum time that loc_rcvr_status has to be OK before a transition to Link Up can occur. The timer shall expire 1 µs ± 0.1 µs after being started.

Link Status Synchronization – 1000BASE-T1

- The transition to SEND_DATA is a combination of *loc_phy_ready* = OK **AND** *rem_phy_ready* = OK , which is good
  - But it is also combined with the *minwait_timer, which* has an uncertainty of ± 50 µs
  - And the *minwait_timer, is used again in the* Link Monitor *with a 2nd* uncertainty of ± 50 µs – a total of ± 100 µs

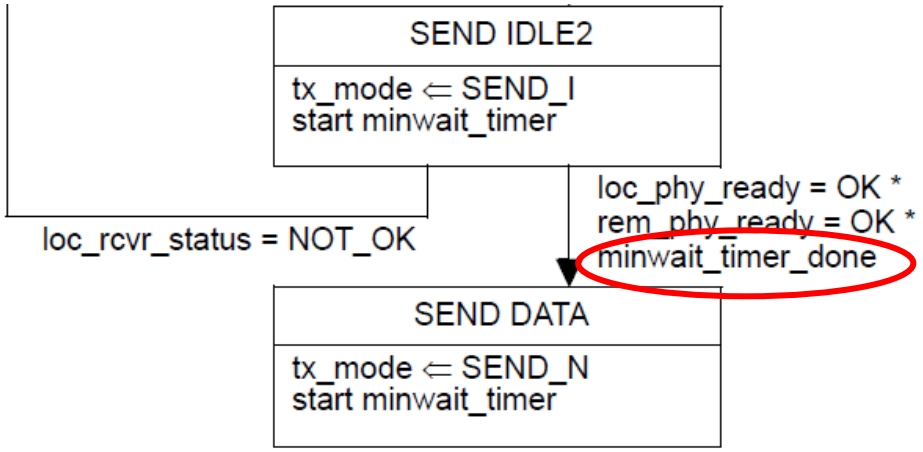

Figure 97–26—PHY Control state diagram

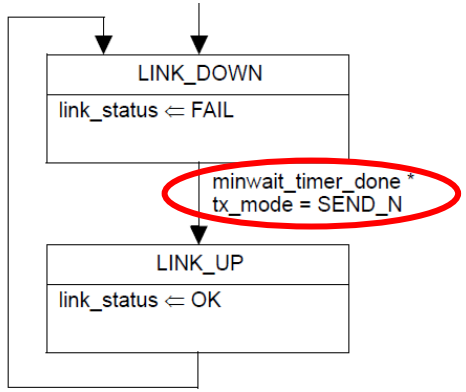

NOTE 1—maxwait_timer is started in PHY Control state diagram (see Figure 97–26).

NOTE 2—The variables link_control and link_status are designated as link_control_1GigT1 and link_status_1GigT1, respectively, by the Auto-Negotiation Arbitration state diagram (Figure 98–7) if the optional Auto-Negotiation function is implemented.

Figure 97–27—Link Monitor state diagram

minwait_timer

A timer used to determine the minimum amount of time the PHY Control stays in the SILENT, TRAINING, SEND IDLE2, and SEND DATA states. The timer shall expire 975 µs ± 50 µs after being started.

Link Status Synchronization – 2G5/ 5G/10GBASE-T1

- Clause 149 and 165 PHYs use an XGMII based MAC Interface
- In XGMII based PHYs Local / Remote Fault signalling is used to ensure link status synchronization
  - Local / Remote Fault signalling is not available today in the clause 22 MII
  - Local / Remote Fault signalling is dependent on the RS which is external to the PHY

Link Status Synchronization – 10BASE-T1L

- The transition between SEND_IDLE and SEND_IDLE_OR_DATA is a combination of *loc_rcvr_status* = OK **AND** *rem_rcvr_status* = OK, which is good
  - And the the *minwait_timer* has a tight bound of ± 1 µs, which is good

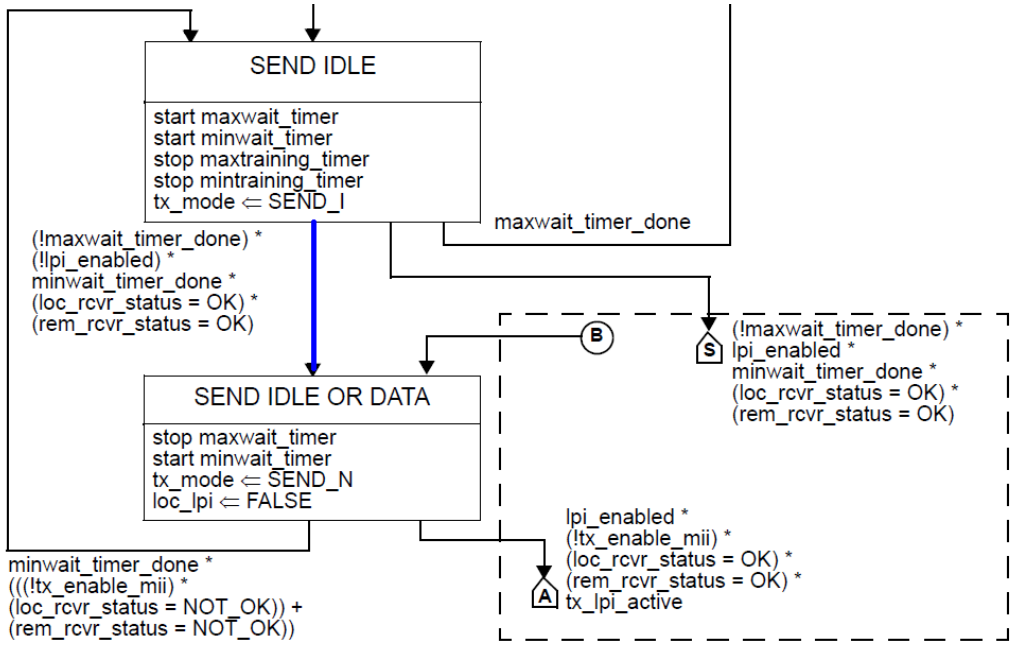

NOTE—Transitions inside dashed boxes are required only for the EEE capability.

Figure 146–15—PHY Control state diagram, part a

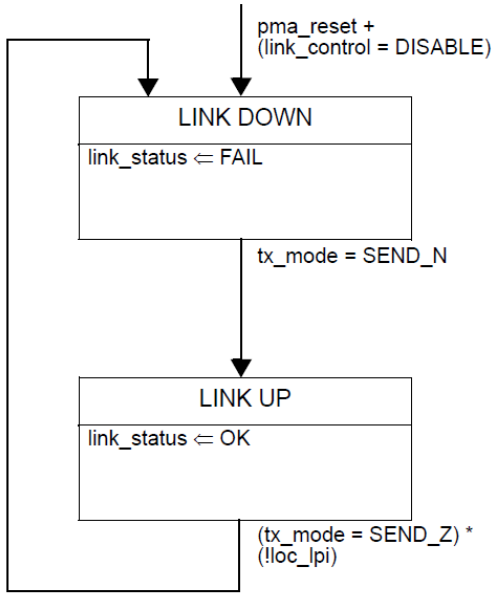

Figure 146–18—PHY Link Monitor state diagram

minwait_timer
A timer used to determine the minimum amount of time the PHY Control stays in the SEND IDLE or DATA states. The timer shall expire 20 µs ± 1 µs after being started.

Link Status Synchronization and Retrain

- Retrain was introduced in 1000BASE-T as link start-up was long (2-4s)
  - If the link partner indicates that its receiver status is NOT_OK the local PHY goes to SEND_IDLE – **but *link_status* remains OK**
  - In this scenario **packets transmitted by the local MAC go nowhere**
- We have the same issue in clause 146
- This has been resolved in the automotive PHYs by not supporting retrain
- 10GBASE-T uses retrain and it is necessary as the link up is more complicated and is much longer (~2s) in terms of bit times
- Retrain should not be supported in 100BASE-T1L
  - Link start-up in 100BASE-T1L has a target of < 100 ms
  - Most robust and fastest means to recover from a link disturbance is to drop the link and restart Auto-Negotiation

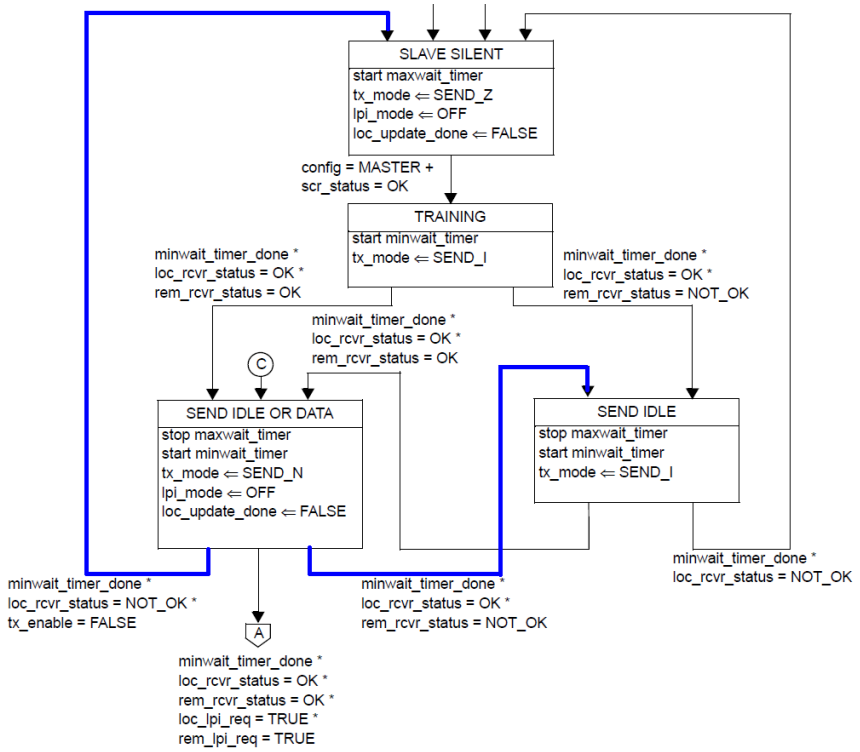

Figure 40–16a—PHY Control state diagram, part a

Proposed 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, which is good
  - There are no timers involved so no uncertainty added
  - The decoding of *rem_phy_ready* is very fast with low latency

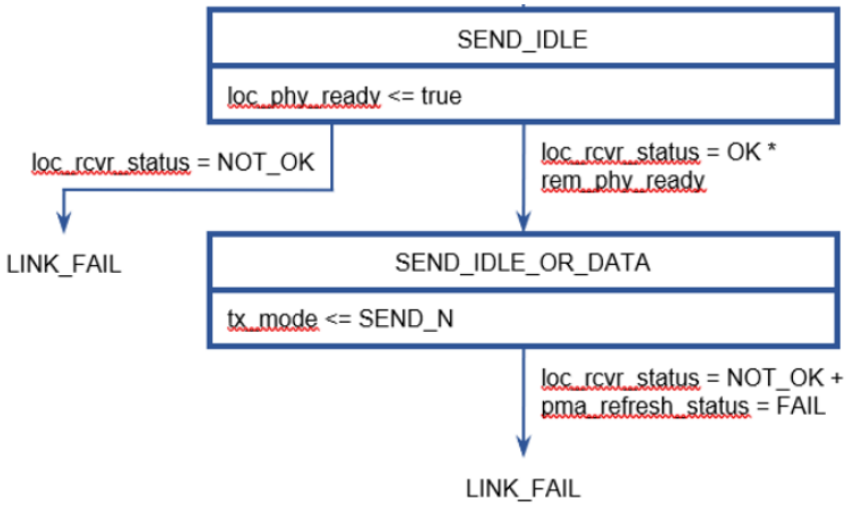

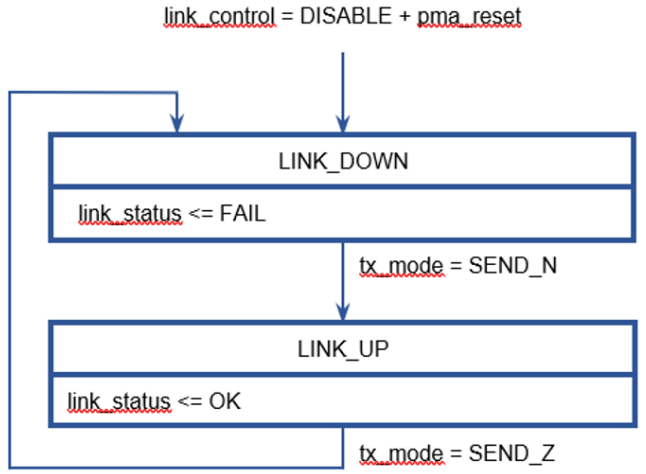

The rem_phy_ready parameter is set to FALSE when the pcs_rx_mode parameter is TRAIN. A sequence of symbols representing idle mode with PHY ready encoding is detected if 4 consecutive control characters are received, each of which is /I/.

Summary

- A review of link status synchronization for a number of lower speed Ethernet PHY standards has been provided
- Each generation of these Ethernet PHY standards has improved the synchronization of link status between local and remote PHYs
- Fast link-up time and good link status synchronization is important for Industrial Ethernet applications
- We have proposed link status synchronization for 100BASE-T1L with good synchronization between the times that the link partners assert *link_status* OK

Questions ?

Link Status Synchronization and Retrain – 10BASE-T1L

- Retrain in 10BASE-T1L is similar to 1000BASE-T
  - If the local or remote PHY indicates that its receiver status is NOT_OK the PHY goes to SEND_IDLE – **but *link_status* remains OK**
  - In this scenario **packets transmitted by the local MAC go nowhere**

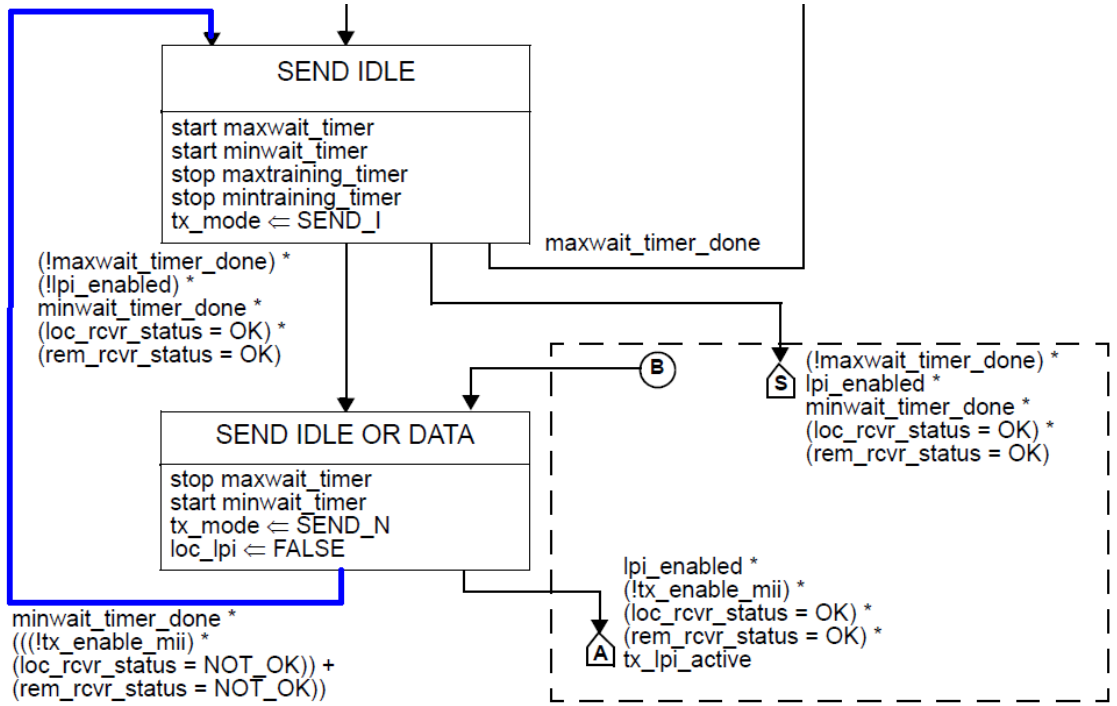

NOTE—Transitions inside dashed boxes are required only for the EEE capability.

Figure 146–15—PHY Control state diagram, part a