Architectural Choices in 802.3dg

IEEE P802.3dg 100 Mb/s Long-Reach Single-Pair Ethernet Task Force

George Zimmerman

CME Consulting/ADI, APL Group, Cisco, Marvell, OnSemi, SenTekSe, Sony

802.3dg Interim Meeting, 16 September 2024

Choices for 802.3 PHY Architecture

- There are many choices in IEEE Std 802.3 for our PHY architecture:
  - “Old Style” 10/100 Mbps PHYs
  - “Gigabit Style” PHYs (1000BASE-T and 100BASE-T1)
  - “10 Gigabit Style” PHYs (2.5/5/10GBASE-T and BASE-T1)
  - Gigabit and MultiGigabit BASE-T1 PHYs
  - Common Industry Augmentations – including other interfaces

Signalling Control Values across the Link

- Control from the RS
  - "Old Style 10/100" - Cl 22 MII, Cl 24 PCS – limited encodings
    - SSD, ESD, IDLE, Error or data each get 4B/5B code group (Table 24-1)
  - "Gigabit Style" – Clause 36 – Ordered Sets, encoded by PCS
    - Configuration, Error propagation, Configuration, LPI (Table 36-3)
    - Encoded into special modulated symbols (Table 40-2)
  - "10 Gigabit Style" – Clause 46 – Sequence Ordered Sets, adds "sequence ordered sets" controls are block encoded

Table 46–5—Sequence ordered sets

Lane 0	Lane 1	Lane 2	Lane 3	Description
Sequence	0x00	0x00	0x00	Reserved
Sequence	0x00	0x00	0x01	Local Fault
Sequence	0x00	0x00	0x02	Remote Fault
Sequence	0x00	0x00	0x03	Link Interruption
Sequence	≥ 0x00	≥ 0x00	≥ 0x04	Reserved

NOTE—Values in Lane 1, Lane 2, and Lane 3 columns are in hexadecimal, most significant bit to least significant bit (i.e., <7:0>). The link fault signaling state diagram allows future standardization of reserved Sequence ordered sets for functions other than link fault indications

"The link fault signaling state diagram allows future standardization of reserved Sequence ordered sets for functions other than link fault indications"

A 24 bit address space with only 4 defined values…

BASE-T1 PHYs – add OAM

- From PHY to PHY, outside data path for the host to read and do something…
  - Framing structure had room for extra bits
  - Warns of impending doom
  - Triggers refresh during LPI
- Adds complexity, extra:
  - State diagrams
  - Corner cases
  - Bits to send…
- BASE-T and 100BASE-T1 don't have this, don't seem to need it

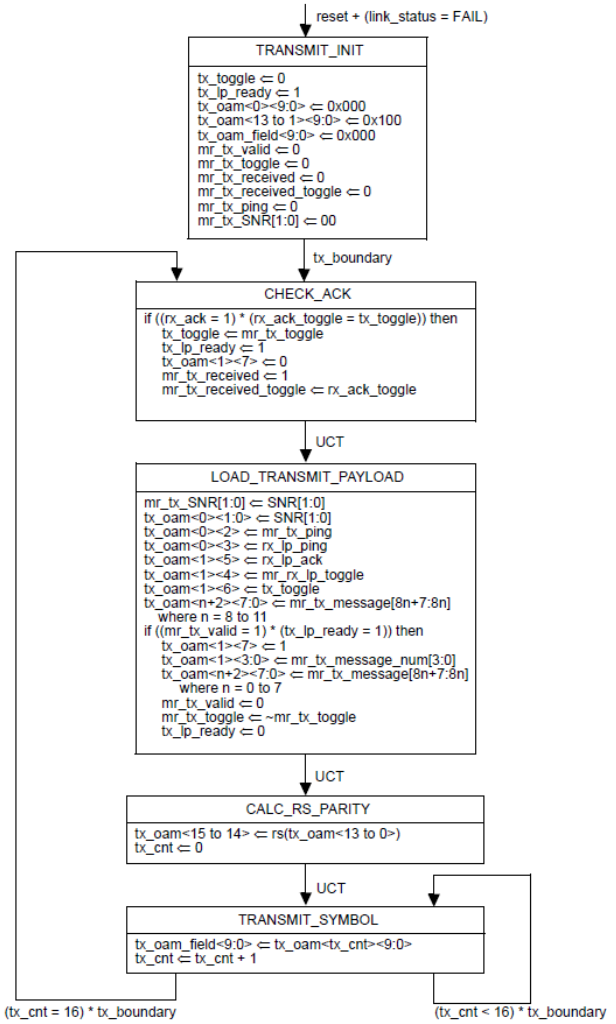

Figure 149–25—Transmit state diagram

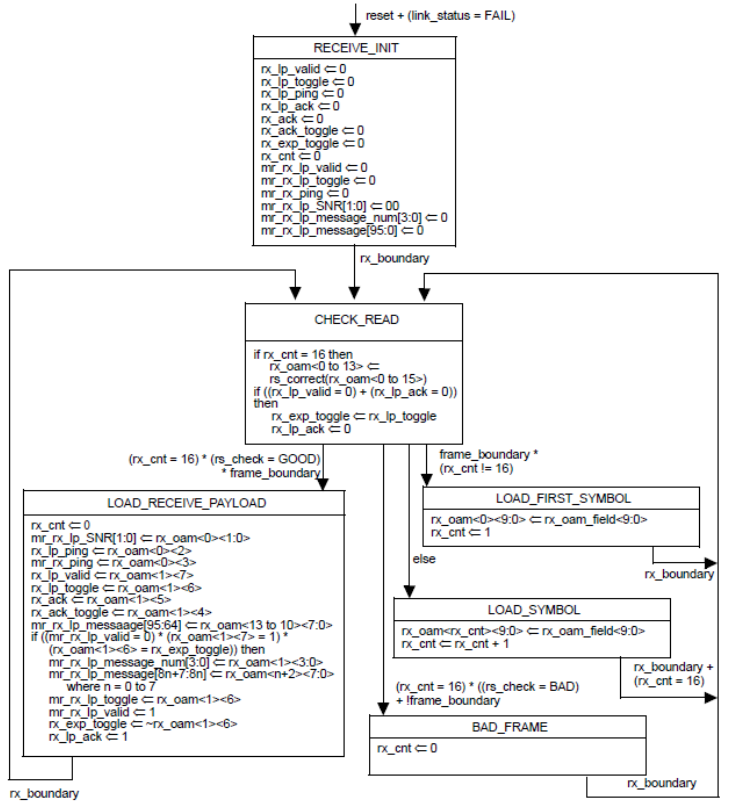

Figure 149–24—Receive state diagram

BASE-T vs. BASE-T1 PHYs - retrains

- BASE-T PHYs generally took seconds to restart
  - Clause 28 break_link_timer is 1200 to 1500 msec
  - MDI/MDI-X resolution sample timer (62 msec in clause 40)
- Cost of an auto-neg/restart cycle for a BASE-T PHY was high….
  - Link drops were avoided (worked through dropping frames)
  - OR faster retraining means were desired and specified,.
- BASE-T1 PHYs train in milliseconds – generally < 100 msec
  - Clause 98 break_link_timers are < 10 msec (300usec for HSM)
  - No MDI/MDI-X resolution
- ***Is it preferable just to drop and restart a marginal link?***
  - ***Perhaps with a 'downshift' function (see jones_3dg_01_08172022***

Common Industry Interface Features

- "link status" – is the link up?
  - Locally used
  - Most PHYs have a "Link Monitor" state machine
  - Status passed to Autoneg (on the TDI, not the MII)
  - Status may be used by other protocols (see 802.1AC MAC_Operational)
  - Familiar light on the RJ-45 port
    - Externally pinned out

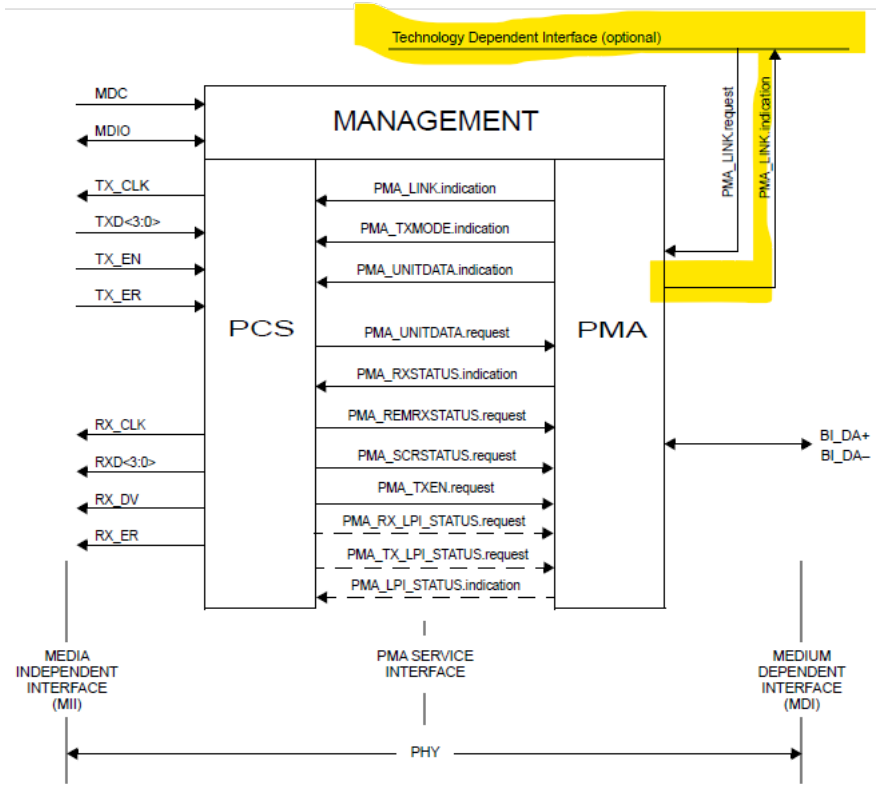

NOTE—Service interface primitives shown with dashed lines are required only for EEE capability.

Figure 146–2—10BASE-T1L PHY interfaces

Should a new interface allow for encoding Non-MII signals like this?

Relation to 100BASE-T1L

- 100BASE-T1L is a return to 100 Mbps
- We have a chance to blend architectures
- What we choose impacts the structure the editors take for the text – so we need to make general choices soon
- Things to consider:
  - Highly constrained on power, capacity, latency
  - How much functionality we add impacts usability
  - How much do we have to send across the link?

Thank You!