ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

Clause 22 MII Changes for 802.3dg Required by Sequence Ordered Sets

Philip Curran
Brian Murray
Jacobo Riesco

Introduction

- This presentation discusses some of the issues associated with the changes required to clause 22 to add Sequence ordered sets
- In the task force we have had quite a number of presentations on the 8N/(8N+1) block encoding and how we should add Sequence ordered sets
  - Sequence ordered sets were originally added into XGMII which aligns on 4-byte boundaries
  - There is significant complexity required to work with existing MII MAC devices to resolve the issues of odd nibble alignment
- **However, there has been very little discussion on the changes required to clause 22 to support Sequence ordered sets**
- **And very little discussion on why we are adding Sequence ordered sets**
  - **What new features / functions do Sequence ordered sets support?**

Sequence Ordered Sets

- Sequence ordered sets were first introduced in clause 46 Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII)
- At the May 2024 Interim meeting the 802.3dg task force adopted a motion add the capability to support Sequence ordered sets to the MII signaling
- At the July 2024 Plenary meeting the 802.3dg task force adopted a motion to add Sequence ordered set of the MII according to Lo_3dg_01a_0724 page 3
  - Add Sequence code for MII encoding of Sequence ordered sets
  - Eight nibble Sequence ordered set similar to the 4-byte sequence used in XGMII

TX_EN	TX_ER	TXD<3:0>	Indication
0	0	0000 to 1111	Normal inter-frame
0	1	0000	Reserved
0	1	0001	Assert LPI
0	1	0010	PLCA BEACON request (not supported)
0	1	0011	PLCA COMMIT request (not supported)
0	1	0100	Sequence (New)
0	1	0101 to 1111	Reserved
1	0	0000 to 1111	Normal data transmission
1	1	0000 to 1111	Transmit error propagation

RX_DV	RX_ER	RXD<3:0>	Indication
0	0	0000 to 1111	Normal inter-frame
0	1	0000	Normal inter-frame
0	1	0001	Assert LPI
0	1	0010	PLCA BEACON request (not supported)
0	1	0011	PLCA COMMIT request (not supported)
0	1	0100	Sequence (New)
0	1	0101 to 1101	Reserved
0	1	1110	False Carrier indication (not supported)
0	1	1111	Reserved
1	0	0000 to 1111	Normal data transmission
1	1	0000 to 1111	Transmit error propagation

IEEE 802.3dg Task Force 3 15 July 2024

Backup: Slide 3 from Lo_3dg_01a_0724

- May 2024 - **Motion #2:**
  - Move that 802.3dg add the capability to support Sequence ordered sets to the MII signaling.
- July 2024 - **Motion #2 Adopt:**
  - Sequence ordered set of the MII according to Lo_3dg_01a_0724.pdf, page 3.

Supported Codes on MII

- Eight nibble sequence ordered set similar to the 4-byte sequence used in XGMII.
  Sequence, Sequence, Data1[3:0], Data1[7:4], Data2[3:0], Data2[7:4], Data3[3:0], Data3[7:4]
- Data1, Data2, Data3 mapping and meaning same as for XGMII lane1, lane2, lane3
- Half-Duplex not supported

TX_EN	TX_ER	TXD<3:0>	Indication
0	0	0000 to 1111	Normal inter-frame
0	1	0000	Reserved
0	1	0001	Assert LPI
0	1	0010	PLCA BEACON request (not supported)
0	1	0011	PLCA COMMIT request (not supported)
0	1	0100	Sequence (New)
0	1	0101 to 1111	Reserved
1	0	0000 to 1111	Normal data transmission
1	1	0000 to 1111	Transmit error propagation

RX_DV	RX_ER	RXD<3:0>	Indication
0	0	0000 to 1111	Normal inter-frame
0	1	0000	Normal inter-frame
0	1	0001	Assert LPI
0	1	0010	PLCA BEACON request (not supported)
0	1	0011	PLCA COMMIT request (not supported)
0	1	0100	Sequence (New)
0	1	0101 to 1101	Reserved
0	1	1110	False Carrier indication (not supported)
0	1	1111	Reserved
1	0	0000 to 1111	Normal data transmission
1	1	0000 to 1111	Transmit error propagation

IEEE 802.3dg Task Force 3 15 July 2024 AXONNE

Sequence Ordered Sets

- They support link fault signaling between the remote RS and the local RS
  - Three fault signals are supported: Local Fault, Remote Fault and Link Interruption
    - Upon recognition of a fault condition a PHY sublayer indicates Local Fault status on the data path
    - When this Local Fault status reaches an RS, the RS stops sending MAC data or LPI and continuously generates a Remote Fault status on the transmit data path
    - When Remote Fault or Link Interruption status is received by an RS, the RS stops sending MAC data or LPI
      - And continuously generates Idle control characters
    - Link Interruption is indicated by the PHY receive function by continuously sending the Link Interruption ordered set
  - Local Fault indicates a fault detected on the receive data path between the remote RS and the local RS
  - Remote Fault indicates a fault on the transmit path between the local RS and the remote RS
  - Status is signaled in a four byte Sequence ordered set (as there are 4 lanes in XGMII)

Link Fault Signaling between Remote RS and Local RS

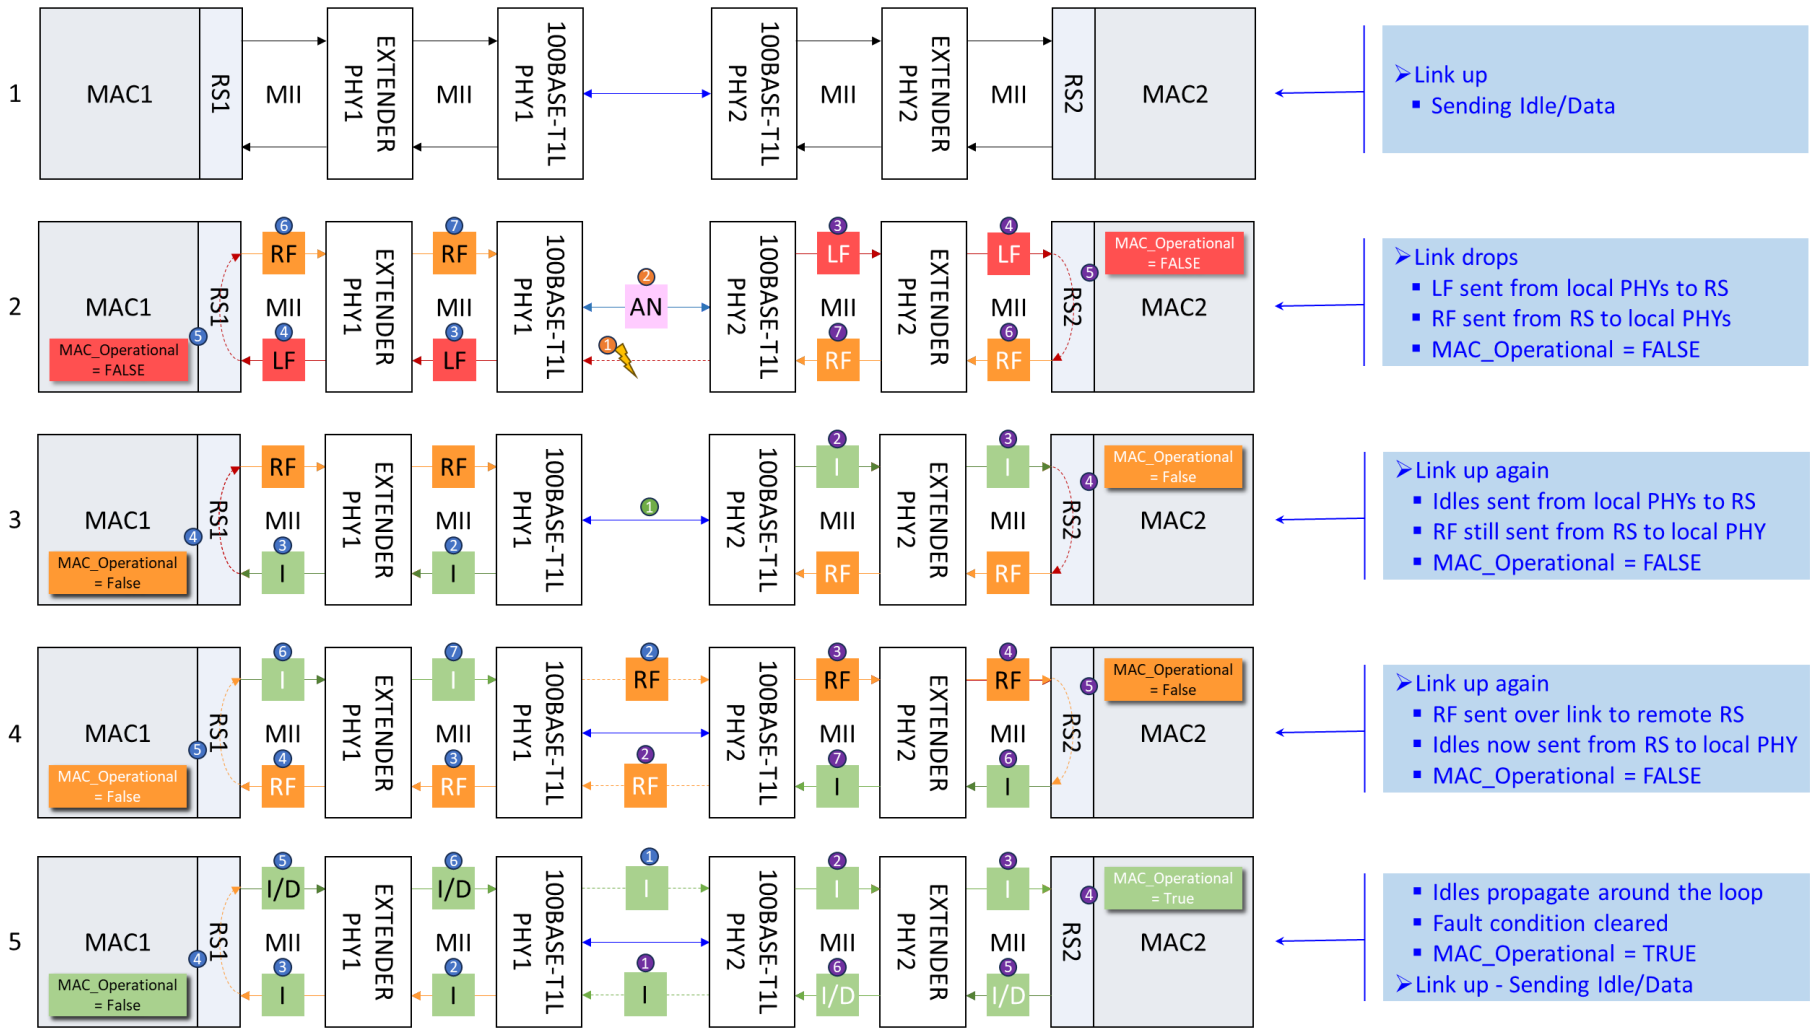

Link Status Synchronization / Link Fault

- Link fault signaling in MultiG PHYs T1(55, 126, 149, 165) is used for Link Status Synchronization (MAC_Operational Synchronization)
  - This is when the RS is ready to transmit and receive data
  - 100BASE-T1L does **not** require Local or Remote Fault signaling for this function
- The 802dg task force has adopted a means to synchronize Link Status for 100BASE-T1L
  - The PHY Control state diagram adopted in January achieves link status synchronization using loc_rcvr_status = OK AND rem_phy_ready = OK
- There is no need to send a fault indication in 100BASE-T1L
  - When loc_rcvr_status = NOT_OK we drop the link and restart Auto-Negotiation
    - Link Status signals that the link has dropped
    - Previously 'Industry Defined MAC Interfaces' have carried the Link Status information across the interface using non-standard signalling

Changes needed to Clause 22 MII - Add Clause 46.3.4

- Link fault signaling is defined in clause 46.3.4
  - To support Sequence ordered sets add this clause, the table & state diagram
  - Local Fault, Remote Fault & Link Interruption are continuous not transitory signals

46.3.4 Link fault signaling

Link fault signaling operates between the remote RS and the local RS. Faults detected between the remote RS and the local RS are received by the local RS as Local Fault. Only an RS originates Remote Fault signals.

Sublayers within the PHY are capable of detecting faults that render a link unreliable for communication. Upon recognition of a fault condition a PHY sublayer indicates Local Fault status on the data path. When this Local Fault status reaches an RS, the RS stops sending MAC data or LPI, and continuously generates a Remote Fault status on the transmit data path (possibly truncating a MAC frame being transmitted). When Remote Fault or Link Interruption status is received by an RS, the RS stops sending MAC data or LPI, and continuously generates Idle control characters. When the RS no longer receives fault status messages, it returns to normal operation, sending MAC data or LPI.

Status is signaled in a four byte Sequence ordered set as shown in Table 46–5. The PHY indicates Local Fault with a Sequence control character in lane 0 and data characters of 0x00 in lanes 1 and 2 plus a data character of 0x01 in lane 3. The RS indicates a Remote Fault with a Sequence control character in lane 0 and data characters of 0x00 in lanes 1 and 2 plus a data character of 0x02 in lane 3. Though most fault detection is on the receive data path of a PHY, in some specific sublayers, faults can be detected on the transmit side of the PHY. This is also indicated by the PHY with a Local Fault status.

For operation with links that may be temporarily interrupted, optional detection of a third fault condition, Link Interruption, is provided. Link Interruption is indicated by the PHY receive function by continuously sending the Link Interruption ordered set as defined in Table 46–5.

The RS reports the fault status of the link. Local Fault indicates a fault detected on the receive data path between the remote RS and the local RS. Remote Fault indicates a fault on the transmit path between the local RS and the remote RS. The RS shall implement the link fault signaling state diagram (see Figure 46–11).

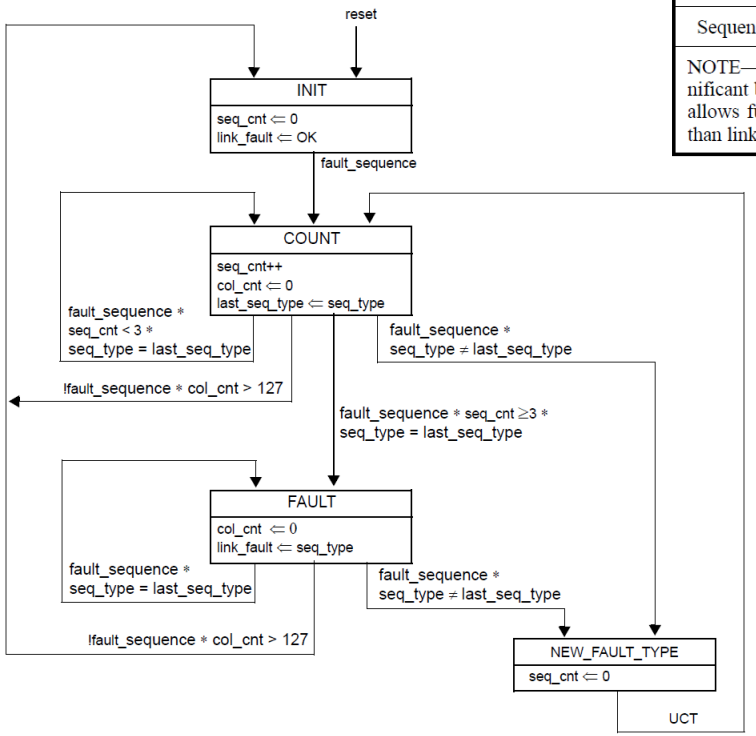

Figure 46–11—Link fault signaling state diagram

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 RS output onto TXC<3:0> and TXD<31:0> is controlled by the variable link_fault.

a) **link_fault = OK**
The RS shall send MAC frames as requested through the PLS service interface.
In the absence of MAC frames, the RS shall generate Idle control characters.

b) **link_fault = Local Fault**
The RS shall continuously generate Remote Fault Sequence ordered sets.

c) **link_fault = Remote Fault or link_fault = Link Interruption**
The RS shall continuously generate Idle control characters.

Need to Signal Sequence Ordered Sets PHY to PHY

- A means has been proposed to modify the 8N/(8N+1) block encoding so that Sequence ordered sets can be communicated across the link from PHY to PHY
  - This is to communicate Local Fault, Remote Fault and Link Interruption from PHY to RS and Remote Fault from RS to RS
  - This is proposed using the XGMII four byte sequence, supporting 24-bits of information for 3 messages
  - In XGMII it was natural to keep the four bytes as XGMII is a 4-lane interface
- Only Remote Fault is signalled over the link for 100BASE-T1L
- We have added significant complexity to the block encoding just to signal Remote Fault across the link
  - Adding Sequence ordered sets to an MII interface is more complicated than for XGMII as with XGMII you are guaranteed to align on a 4-byte boundary
- There are much simpler ways to signal Remote Fault across the link using the 8N/(8N+1) block encoding proposed for 100BASE-T1L
  - The Q code added to the block encoding could be used to signal Remote Fault

Add Assert Local / Remote Fault to Clause 22 MI

- We need to communicate Local Fault and Remote Fault between the PHY and RS
  - Add Assert Remote Fault to clause 22 MII on the Tx and Rx side
  - Add Assert Local Fault to clause 22 MII on the Rx side
  - And add a link fault signaling state diagram

Table 22–1—Permissible encodings of TXD<3:0>, TX_EN, and TX_ER

TX_EN	TX_ER	TXD<3:0>	Indication
0	0	0000 through 1111	Normal inter-frame
0	1	0000	Reserved
0	1	0001	Assert LPI
0	1	0010	PLCA BEACON request
0	1	0011	PLCA COMMIT request
0	1	0100	Assert remote fault
0	1	0101 through 1111	Reserved
1	0	0000 through 1111	Normal data transmission
1	1	0000 through 1111	Transmit error propagation

Table 22–2—Permissible encodings of RXD<3:0>, RX_ER, and RX_DV

RX_DV	RX_ER	RXD<3:0>	Indication
0	0	0000 through 1111	Normal inter-frame
0	1	0000	Normal inter-frame
0	1	0001	Assert LPI
0	1	0010	PLCA BEACON indication
0	1	0011	PLCA COMMIT indication
0	1	0100	Assert remote fault
0	1	0101	Assert local fault
0	1	0110 through 1101	Reserved
0	1	1110	False Carrier indication
0	1	1111	Reserved
1	0	0000 through 1111	Normal data reception
1	1	0000 through 1111	Data reception with errors

Assert Remote Fault with Fault Signalling State Diagram

- When the link drops the PHY signals Assert Local Fault to the RS
  - The RS signals Assert Remote Fault which is sent to the remote RS over the link
  - A link fault signaling state diagram controls the behavior of the RS

The RS output onto TX_EN, TX_ER and TXD<3:0> is controlled by the variable link_fault.

a) **link_fault = OK**
The RS shall send MAC frames as requested through the PLS service interface.
In the absence of MAC frames, the RS shall generate Idle control characters.

b) **link_fault = Local Fault**
The RS shall continuously generate Assert Remote Fault.

c) **link_fault = Remote Fault**
The RS shall continuously generate Idle control characters.

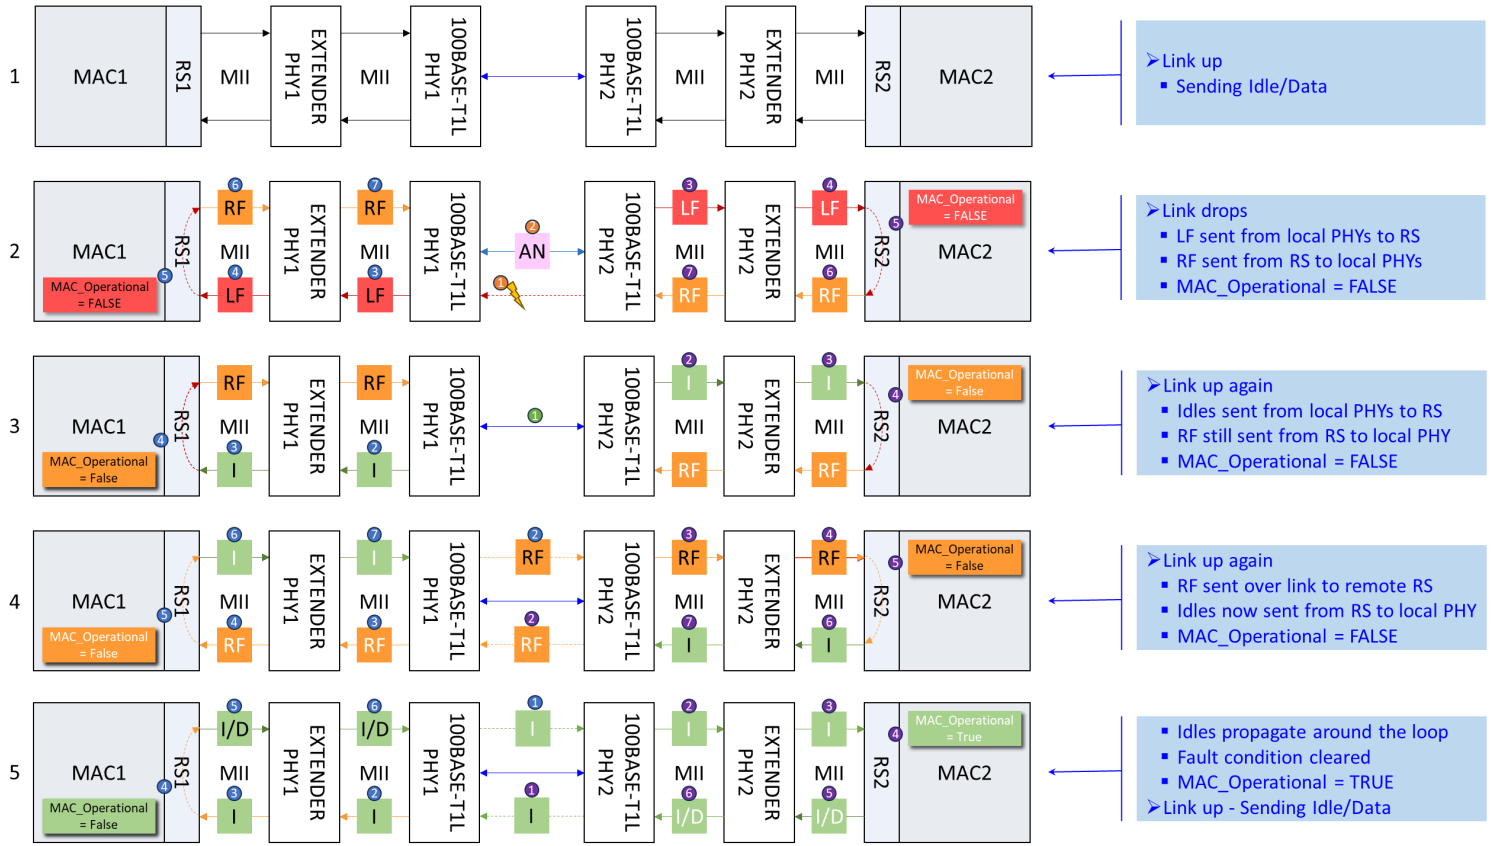

Options for Clause 22 MII Changes for 100BASE-T1L

- A/ Change the Clause 22 MII to add Clause 46 Sequence ordered sets
  - Specify clause 199 to the new clause 22 MII
  - The new MII supports Sequence ordered sets between the RS and the PHY
  - Add the encoding for Sequence ordered sets to the 8N/(8N+1) block encoding
  - Add a link fault signaling state diagram as in clause 46 to clause 22
  - Changes to clause 22 are optional
- B/ Change the Clause 22 MII to add Assert Local/Remote Fault over the MII
  - Specify clause 199 to the new clause 22 MII
  - The new MII supports Local Fault and Remote Fault signaling between the RS and the PHY
  - Don't add Sequence ordered sets
  - Use the Q code to signal Remote Fault across the link from PHY to PHY
  - Add a link fault signaling state diagram to clause 22
  - Changes to clause 22 are optional
- C/ Don't change the Clause 22 MII
  - Specify clause 199 to the existing clause 22 MII
  - Don't add Sequence ordered sets
  - Don't support link fault signaling

Conclusions

- We believe the best approach is Option B
  - Minimum change to clause 22 MII and use the Q code to signal Remote Fault across the link
  - Supports the spirit of fault signaling without adding a lot of complexity
- No existing MAC devices (MII, RMII, RGMII, SGMII, USGMII) support Local or Remote Fault signaling
- If Sequence ordered sets are to be fully supported:
  - They must be added to the 8N/(8N+1) block encoding
  - Work is required to simulate and verify that the changes to the 8N/(8N+1) block encoding
  - New text has to be devised with tables and or state diagrams
  - Significant text has to be written to add the link fault signaling and the state diagram to clause 22
- Complicated text for a featured not supported today and creates a risk of interoperability issues found after deployment

Questions ?