ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

Legacy MAC Interfaces for 10/100/1000 PHYs and Link Fault Conditions

Brian Murray, Jacobo Riesco, Philip Curran

Existing Silicon and Legacy MAC Interfaces

- The IEEE standards define MII and GMII MAC interface for PHYs operating at 10, 100 and 1000 Mb/s speeds
  - MII is defined in Clause 22 and supports 10 Mb/s and 100 Mb/s operation
  - GMII is defined in Clause 35 and supports 1000 Mb/s operation
  - These clauses define the logical operation of the signals at the interface
- In theory this is the interface between a PHY chip and a MAC chip
  - And in early implementations of PHYs this was also the actual chip interface
  - The first/second generation of Gigabit PHYs were in 128-pin TQFP 14 x 20 mm packages supporting a 25-pin GMII interface
- In practice with actual silicon today and for the last 25 years, industry standard, lower pin count MAC / PHY interfaces are used
  - Some 100M Industrial Ethernet PHYs do continue to support MII at the pins as this interface has the lowest latency which is important for some Industrial Ethernet protocols
    - Even then COL / CRS / RX_ER are not always brought out or they are shared pins
  - But these PHYs also support the industry standard, lower pin count MAC / PHY interfaces

Industry Standard Lower Pin Count MAC / PHY Interfaces

- The high pin count of MII / GMII and the drive for smaller packages resulted in the introduction of lower pin count MAC interfaces, e.g. RMII, RGMII and SGMII
  - These interface trade bus speed for pin count
  - Multi-port PHYs drove higher SerDes speeds so multiple ports could be combined on a single 4-pin SerDes interface, for example QSGMII, USGMII and USXGMII
- Industry Standard MAC / PHY Interfaces
  - Potterf_MII_Proposals_Expansion_2024-03-12 is an excellent detailed review of these MAC interfaces

MAC Interface	Data/Control Pins	Bus Clock Rate	Clocking	Mgt Pins (MDIO)	Ports
RMII	8	50 MHz	System Synchronous	2	1
RGMII	12	125 MHz	Source Synchronous	2	1
SGMII	4	1.25 GHz	SerDes	2	1
QSGMII	4	5 GHz	SerDes	2	4
USGMII	4	10 GHz	SerDes	2	1 / 4 / 8
UXSGMII	4	10 GHz	SerDes	2	1 / 4 / 8
MII	16	25 MHz	PHY Synchronous	2	1
GMII	25	125 MHz	Source Synchronous	2	1

Most common MAC/PHY chip interface for Single 100M PHYs

Most common MAC/PHY chip interfaces for Single 1G PHYs

Most common MAC/PHY chip interface for Quad 1G PHYs

Most common MAC/PHY chip interface for Octal 1G PHYs

Rarely available at chip pins

A Link Status pin from PHY to MAC is commonly provided at the chip interface for single PHYs.

Industry Standard MAC Interfaces and In-Band Signalling

- Industry standard MAC interfaces can use in-band signalling to transfer information about the link to the MAC
  - RGMII defined optional in-band signalling for Link Status, Speed/Duplex
  - SGMII uses clause 36 (1000BASE-X) to support Link Status, Speed/Duplex signalling
  - However, many existing silicon single PHY to MAC interfaces do not uses these signals

MAC Interface	In-band / Clause 36		
	Link Status	ANEG (Speed/Duplex)	EEE - LPI
RMII	No	No	No
RGMII [1]	Yes	Yes	Yes
SGMII	Yes	Yes	Yes
QSGMII	Yes	Yes	Yes
USGMII	Yes	Yes	Yes
UXSGMII	Yes	Yes	Yes
MII	No	No	Yes

[1] RGMII uses optional In band signalling for Link Status, Speed/Duplex.

Typically single PHYs have a dedicated Link Status Pin.

Config_Reg encoding for SGMII/USGMII/USXGMII

LSB															MSB
D0	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11	D12	D13	D14	D15
1	rsvd	rsvd	rsvd	rsvd	rsvd	rsvd	EEEcs	EEEen	SPEED			Duplex	rsvd	rsvd	Link

Code	Ordered Set	Number of Code-Groups	Encoding
/C/	**Configuration**		Alternating /C1/ and /C2/
/C1/	Configuration 1	4	/K28.5/D21.5/Config_Reg[a]
/C2/	Configuration 2	4	/K28.5/D2.2/Config_Reg[a]
/I/	**IDLE**		Correcting /I1/, Preserving /I2/
/I1/	IDLE 1	2	/K28.5/D5.6/
/I2/	IDLE 2	2	/K28.5/D16.2/

Table 36–3—Defined ordered sets

IEEE Std 802.3-2022, IEEE Standard for Ethernet

Reference: Cisco USGMII Functional Specification: EDCS-1155168, Rev 4.2

Industry Standard MAC Interfaces and In-Band Signalling

- Industry standard MAC interfaces can use in-band signalling to transfer information about the link to the MAC
  - RGMII defined optional in-band signalling for Link Status, Speed/Duplex
  - SGMII uses clause 36 (1000BASE-X) to support Link Status, Speed/Duplex signalling
  - However, many existing silicon single PHY to MAC interfaces do not uses these signals

MAC Interface	In-band / Clause 36			Sequenced Ordered Sets		PCH
	Link Status	ANEG (Speed/Duplex)	EEE - LPI	Local / Remote Fault	Link Interruption	Time Synchronization
RMII	No	No	No	No	No	No
RGMII [1]	Yes	Yes	Yes	Possible [2]	Possible [2]	No
SGMII	Yes	Yes	Yes	No	No	No
QSGMII	Yes	Yes	Yes	No	No	No
USGMII	Yes	Yes	Yes	No	No	Yes [3]
UXSGMII	Yes	Yes	Yes	Yes	No	Yes [3]
MII	No	No	Yes	Possible [4]	Possible [4]	No

None of the commonly used MAC interfaces support sequence ordered sets and cannot transfer Local / Remote Fault or Link Interruption

[1] RGMII uses optional In band signalling for Link Status, Speed/Duplex.
Typically single PHYs have a dedicated Link Status Pin.

[2] RGMII can disable in-band signalling and could then use sequence ordered sets to indicate Remote Fault and Link Interruption, but cannot then indicate Link Status or ANeg information. *RMII, SGMII, QSGMII, USGMII cannot support sequenced ordered sets.*

[3] USGMII/USXGMII use a Packet Control Header during Idle to send PTP Timestamps. *Industrial Ethernet single PHYs use Tx/Rx SOP pins.*

[4] MII could use sequenced ordered sets to indicate Remote Fault and Link Interruption, but no existing silicon supports this.

Existing Silicon and Remote Fault

- There are some MAC interfaces (USXGMII) that support the transfer of Local / Remote Fault over the MAC interface
  - However, it appears that existing silicon does not use this signalling
  - And USXGMII does not support the transfer of Link Interruption
- The 1000BASE-T standard had *loc_rcvr_status* / *rem_rcvr_status* variables and the standard requires that the PHY does a retrain if these are *NOT_OK*
  - However, there is no means to communicate this back to the MAC over the interface and consequently any data sent from the MAC to the PHY during retrain is dropped
  - There is no reliable means to communicate loss of receiver performance to the link partner
  - In our experience retrain just delays the inevitable dropping of the link and restart of Auto-Neg and in practice just adds to the time taken to re-establish reliable communication
- We believe the original intention of retrain was to avoid the 1.3s break link
  - Clause 98 Auto-Negotiation is much, much faster
    - The *break_link_timer_[LSM]* is 8 ms and the *break_link_timer_[HSM]* is 350 µs

Proposed Handling of Link Fault Conditions in 100BASE-T1L

- For 100BASE-T1L we propose to drop the link and restart Auto-Negotiation as soon as the local PHY detects that the receiver status is bad
  - Therefore, we don't need to communicate link fault to the RS because it is **communicated via link_status to the MAC**
    - link_status is a parameter passed to the technology dependent interface (see 98.4.1.1) but should be communicated alongside MII
    - In practice, Clause 36 is used for SGMII/USGMII and a Link Status pin for Industrial Ethernet single PHYs
  - The link partner detects that the link has dropped as our signal disappears
  - We don't need to and cannot communicate sequence ordered sets to the link partner because we have restarted Auto-neg
  - When the link degrades there is no reliable mechanism to communicate link status information to the link partner
- We propose to have a fast link time
  - Time to establish link including Auto-Neg should be less than 100 ms
  - A fast link time removes the need for a retrain
- We believe this method will be the most robust and fastest means to recover from a severe link disturbance

Sequence Ordered Sets and Remote Fault

- At the July meeting the 802.3dg group approved a motion to adopt adding sequence ordered sets to the MII interface
  - This is to support signalling from the PHY to the MAC
  - But none of the existing 100M silicon or commonly used MAC interfaces support Sequence Ordered Sets
- We could add sequence ordered sets to a new MII interface
  - This would be to communicate from the PHY to the MAC
  - A new MII should communicate link_status and speed / duplex to the MAC
  - We do not want to add sequence ordered sets to support undefined capability for future definition

Questions ?