The bridge to possible

Ethernet MII Interfaces

Past, Present, and Proposals for Future

Jason Potterf
Affiliated with Cisco
2024-01-25

Agenda

- Review Past MII Solutions
  - Parallel Buses - MII, RMII, GMII, RGMII
  - Command Space in Parallel Buses
  - Serial Buses – SMII, SGMII
  - Multi-Port Serial Buses – QSGMII, USGMII
  - 10G/mgig MII – XGMII, USXGMII, MP-USXGMII
- Path Forward Proposal
  - Leveraged solution for multi-port and single-port 10/100 HD&FD SPE

Motivation

Motivation

- 802.3dg wants a new MII
  - Needs to provide a modern single-port solution for 100 mbit/s data rates
  - Also need to solve multi-port applications to enable switch density
- 802.3da also wants a new MII
  - Needs to solve PLCA control challenges
  - Even legacy parallel buses are deficient for PLCA!
- Why this presentation?
  - 802.3dg would benefit from understanding the 802.3da PLCA problem
  - Most MII specifications are industry specs outside of IEEE, so an overview would be beneficial to the group

MII Evolution
The Parallel Busses

IEEE 802.3
Clause 22. Reconciliation Sublayer (RS) and Media Independent Interface (MII)

MII Score Card	
Max Data Rate	100 Mbit/s
Signal Count	16 Data + 2 Mgmt
Bus Max Clock Rate	25 MHz
Clock Scheme	PHY Synchronous
Command Space	4-bit
Commands Assigned	Tx - 4/16, Rx – 5/16
PLCA Support	Yes, Beacon and Commit

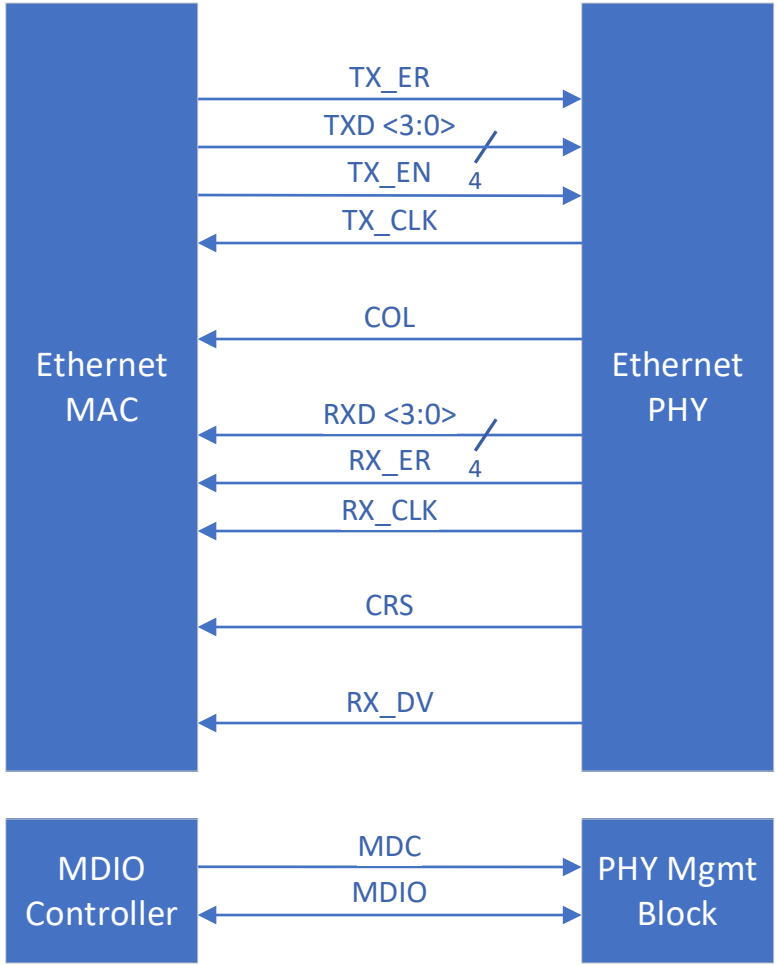

Industry Specification
Reduced Media Independent Interface

RMII Score Card	
Max Data Rate	100 Mbit/s
Signal Count	**7 Data + 1 Clk** + 2 Mgmt
Bus Max Clock Rate	**50 MHz**
Clock Scheme	System Synchronous
Command Space	2-bit
Commands Assigned	Tx - 2/4, Rx – 2/4 (Note 1)
PLCA Support	**No**

Note 1 – Some vendors have assigned values for EEE outside of the RMII spec.

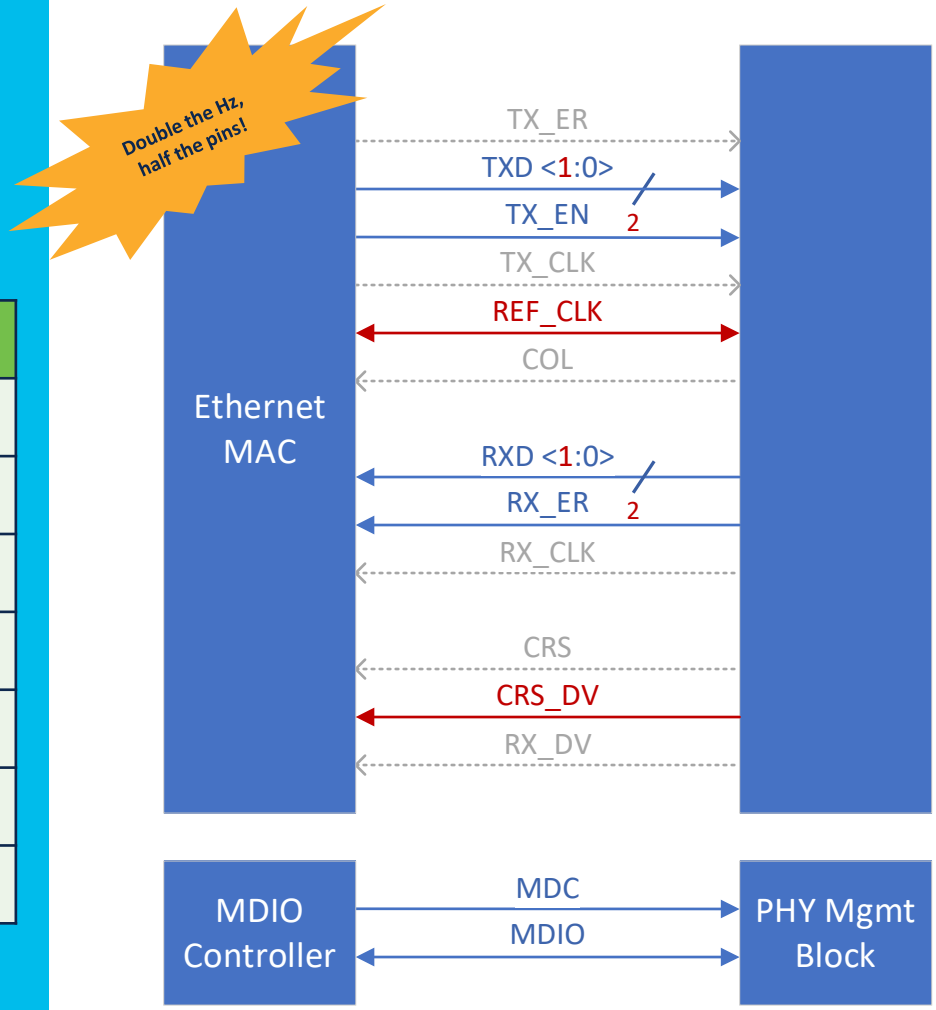

IEEE 802.3
Clause 35. Reconciliation Sublayer (RS) and Gigabit Media Independent Interface (GMII)

MII Score Card	
Max Data Rate	1000 Mbit/s
Signal Count	25 Data + 2 Mgmt
Bus Max Clock Rate	125 MHz
Clock Scheme	Source Synchronous
Command Space	8-bit
Commands Assigned	Tx - 4/256, Rx – 5/256
PLCA Support	**No**

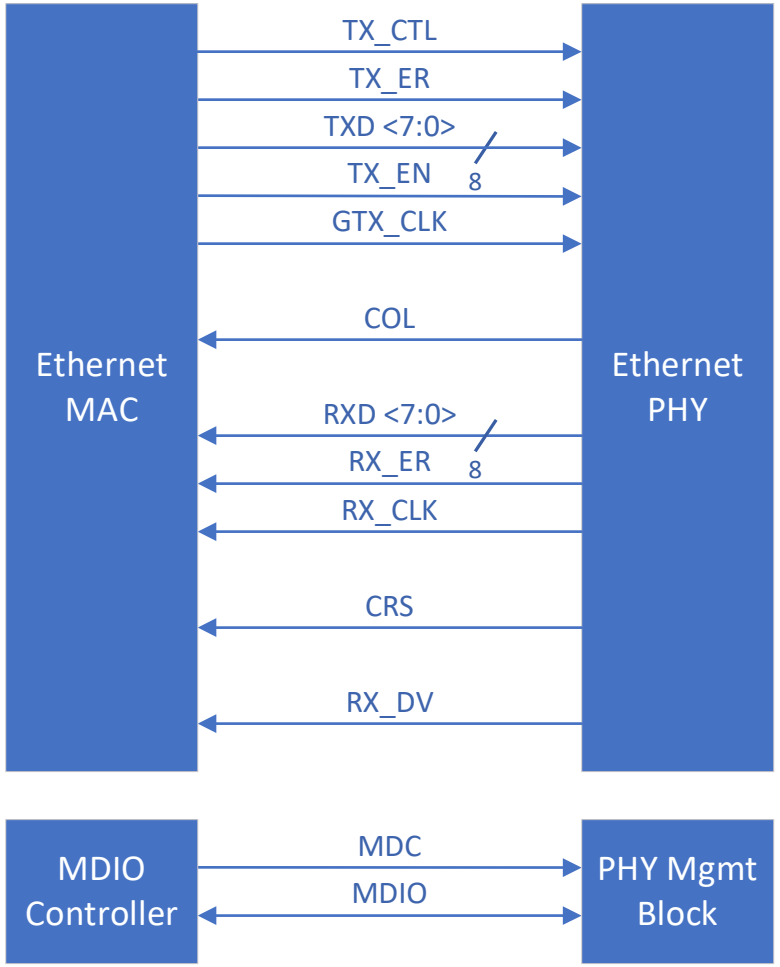

Industry Specification
Reduced Gigabit Media Independent Interface

RGMII Score Card	
Max Data Rate	1000 Mbit/s
Signal Count	**12 Data** + 2 Mgmt
Bus Max Clock Rate	125 MHz
Clock Scheme	Source Synchronous DDR
Command Space	8-bit
Commands Assigned	Tx - 2/256, Rx – 4/256
PLCA Support	**No**

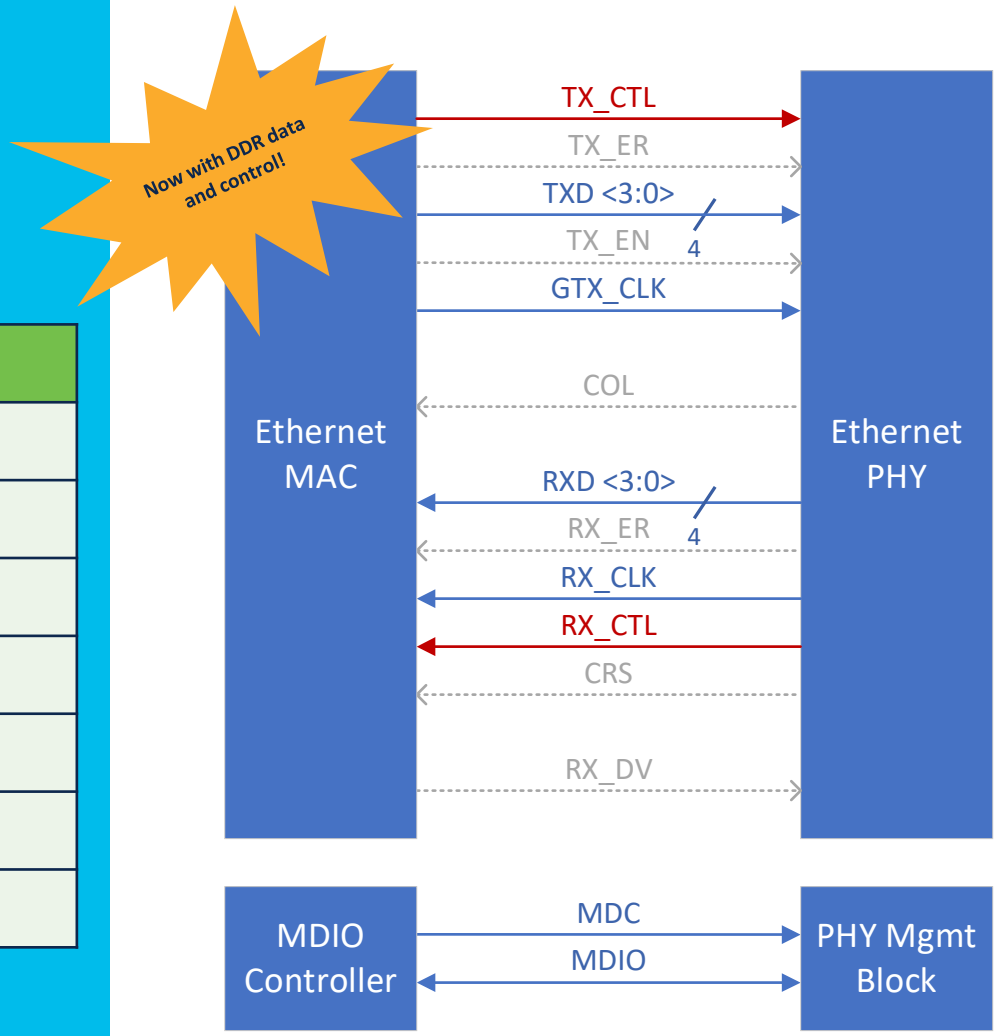

MII Evolution
Parallel Bus Command Space

MII Commands

- Cause Transmission or Indicate Reception of something other than valid data bytes on the wire
- Uses TX_EN / TX_ER / RX_En / RX_ER to create address space
- Currently Defined in 802.3 Clause 22
  - Assert LPI
  - PLCA Beacon Request / Indication
  - PLCA Commit Request / Indication
  - False Carrier Indication

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 through 1111	Reserved
1	0	0000 through 1111	Normal data transmission
1	1	0000 through 1111	Transmit error propagation

Table 22–2—Permissible encoding 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 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

GMII Commands

- GMII increases address space with wider bus
- Uses same control signals as MII, TX_EN / TX_ER / RX_En / RX_ER, to create address space
- Currently Defined in 802.3 Clause 22
  - Assert LPI
  - False Carrier indication
  - Carrier Extend/ Carrier Extend Error
- **PLCA not defined here, but RGMII would benefit when operating at 10 Mb/s data rates**

Table 35–1—Permissible encodings of TXD<7:0>, TX_EN, and TX_ER

TX_EN	TX_ER	TXD<7:0>	Description	PLS_DATA.request parameter
0	0	00 through FF	Normal inter-frame	DATA_COMPLETE
0	1	00	Reserved	—
0	1	01	Assert LPI	—
0	1	02 through 0E	Reserved	—
0	1	0F	Carrier Extend	EXTEND (eight bits)
0	1	10 through 1E	Reserved	—
0	1	1F	Carrier Extend Error	EXTEND_ERROR (eight bits)
0	1	20 through FF	Reserved	—
1	0	00 through FF	Normal data transmission	ZERO, ONE (eight bits)
1	1	00 through FF	Transmit error propagation	No applicable parameter

NOTE—Values in TXD<7:0> column are in hexadecimal.

Table 35–2—Permissible encoding of RXD<7:0>, RX_ER, and RX_DV

RX_DV	RX_ER	RXD<7:0>	Description	PLS_DATA.indication parameter
0	0	00 through FF	Normal inter-frame	No applicable parameter
0	1	00	Normal inter-frame	No applicable parameter
0	1	01	Assert LPI	No applicable parameter
0	1	02 through 0D	Reserved	—
0	1	0E	False Carrier indication	No applicable parameter
0	1	0F	Carrier Extend	EXTEND (eight bits)
0	1	10 through 1E	Reserved	—
0	1	1F	Carrier Extend Error	ZERO, ONE (eight bits)
0	1	20 through FF	Reserved	—
1	0	00 through FF	Normal data reception	ZERO, ONE (eight bits)
1	1	00 through FF	Data reception error	ZERO, ONE (eight bits)

NOTE—Values in RXD<7:0> column are in hexadecimal.

RMII / RGMII Compromises

- RMII assumed the following
  - COL can be derived
  - CRS and RX_DV can merge
  - TX_ER's purpose can be accomplished by intentional data corruption and will be obsoleted by switches
  - Two-bit data bus values during TX_EN / CRS_DV low can be use for control codes
- RGMII made similar assumptions about derived signals, but due to DDR clocking **did not compromise on control codes**

RMII

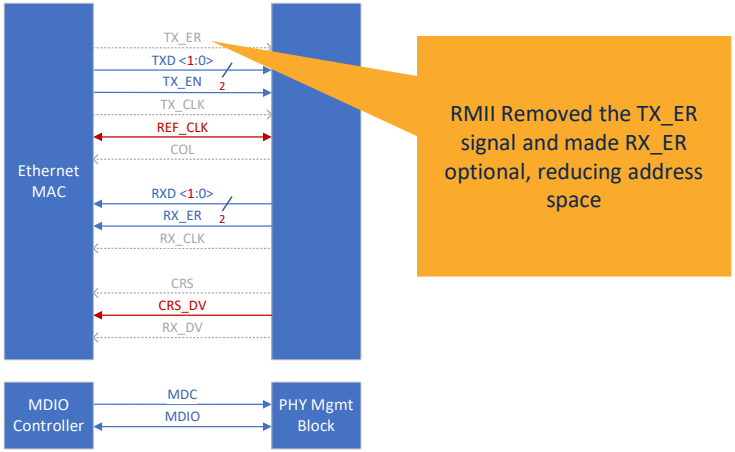

RGMII

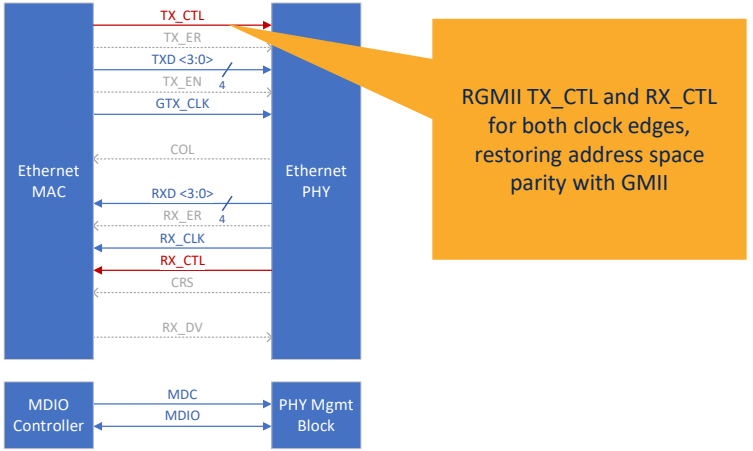

MII Evolution
The Serial Buses

Cisco Specification
Serial-MII (aka SMII)

MII Score Card	
Max Data Rate	100 Mbit/s
Signal Count	**2/port** + 2-5/PHY + 2 Mgmt
Bus Max Clock Rate	125 MHz
Clock Scheme	System or PHY Synchronous
Command Space	**8-bit bitfield**
Commands Assigned	Tx - 5/8, Rx – 7/8 bits
PLCA Support	**No**

System Diagrams

One Clock!

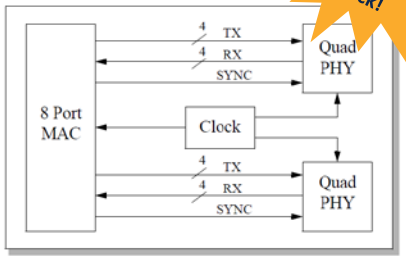

Typical SMII System

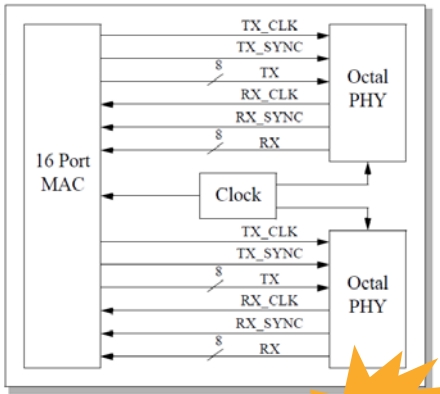

Source Synchronous SMII System

Many Clocks!

Timing Diagrams

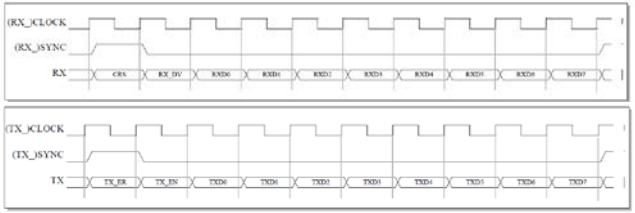

Command Bit Assignments

CRS	RX_DV	RXD0	RXD1	RXD2	RXD3	RXD4	RXD5	RXD6	RXD7
x	0	RX_ER from previous frame	Speed 0 = 10MBit 1 = 100MBit	Duplex 0 = Half 1 = Full	Link 0 = Down 1 = Up	Jabber 0 = OK 1 = Error	Upper Nibble 0 = invalid 1 = valid	False Carrier Detected	1
x	1	One Data Byte (Two MII Data Nibbles)							

RXD7-0 Encoding

TX_ER	TX_EN	TXD0	TXD1	TXD2	TXD3	TXD4	TXD7-5
x	0	Use to force an error in a direct MAC to MAC connection	1 100MBit	1 Full Duplex	1 Link Up	0 No Jabber	1
x	1	One Data Byte (Two MII Data Nibbles)					

TXD7-0 Encoding

Cisco Specification
Serial-GMII (aka SGMII)

MII Score Card	
Max Data Rate	1000 Mbit/s
Signal Count	**4/port** + 2 Mgmt
SerDes Max Rate	1.25 Gb/s
Clock Scheme	SerDes or Source Sync
Command Space	Large – Ordered Sets
Commands Assigned	8 Non-Idle Sets Defined
PLCA Support	**No**

System Diagrams (SGMII Spec)

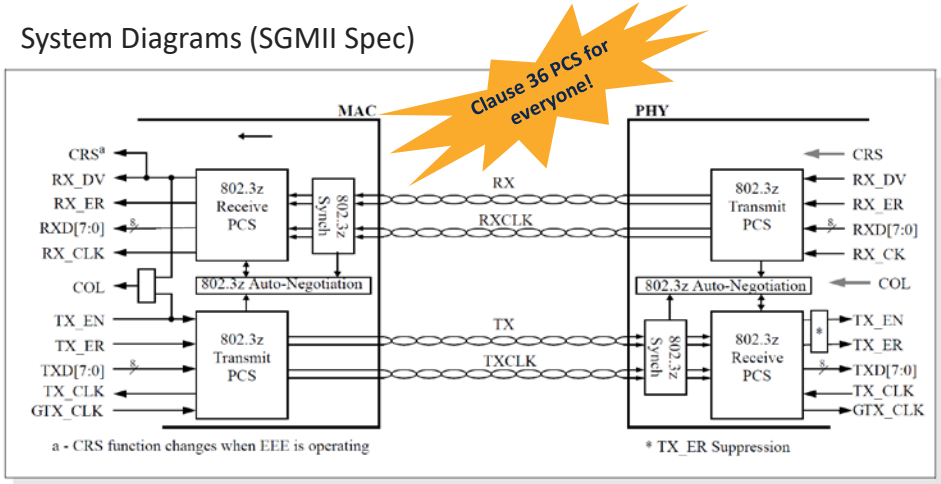

PCS Scheme (Clause 36)

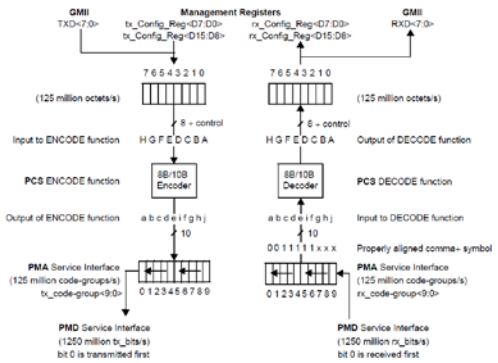

Figure 36–3—PCS reference diagram

Commands (Clause 36)

Table 36–3—Defined ordered sets

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/
	Encapsulation		
/R/	Carrier_Extend	1	/K23.7/
/S/	Start_of_Packet	1	/K27.7/
/T/	End_of_Packet	1	/K29.7/
/V/	Error_Propagation	1	/K30.7/
/LI/	**LPI**		Correcting /LI1/, Preserving /LI2/
/LI1/	LPI 1	2	/K28.5/D6.5/
/LI2/	LPI 2	2	/K28.5/D26.4/

[a]Two data code-groups representing the Config_Reg value.

SGMII Choices

- Source Synchronous vs SerDes
  - Left to implementers, but all used SerDes
- SGMII assumed the following
  - CRS and COL can be derived
  - Speeds < 1Gb/s achieved by “elongation” where bytes are duplicated 10x or 100x to allow constant SerDes data rates
- SGMII’s use of the Clause 36 PCS gives us a large control code address space via 8b/10b code group ordered sets
  - EEE modified Clause 36 to enable LPI
  - **PLCA did not define ordered sets in Clause 36 for COMMIT and BEACON as the PCS is only defined for 1 Gb/s data rates**

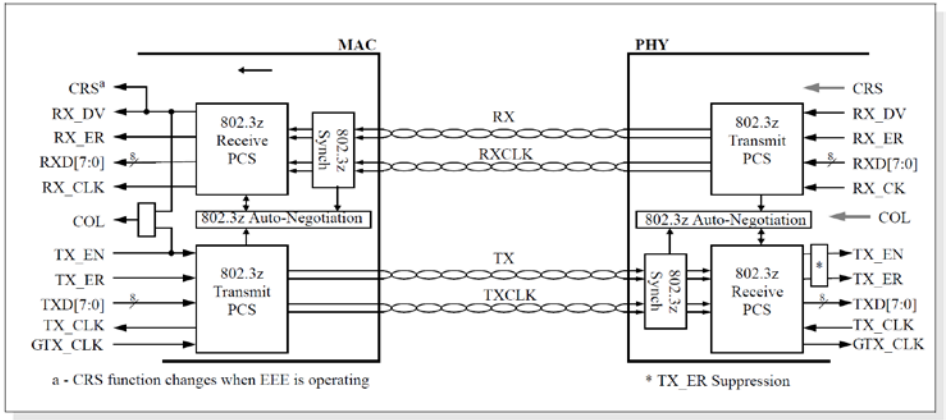

MII Evolution
The Multi-Port Serial Busses

Cisco Specification
Quad Serial GMII (aka QSGMII)

MII Score Card	
Max Data Rate	1000 Mbit/s, **4 Ports**
Signal Count	**4/Quad Phy** + 2 Mgmt
SerDes Max Rate	5.0 Gb/s
Clock Scheme	SerDes
Command Space	Large – Ordered Sets
Commands Assigned	K28.1 Reserved 8 Non-Idle Sets Defined
PLCA Support	**No**

QSGMII Overview

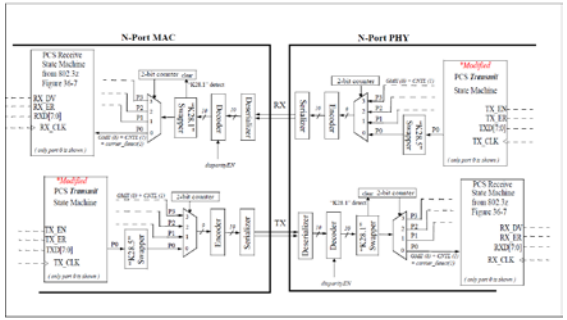

QSGMII PCS TX Scheme

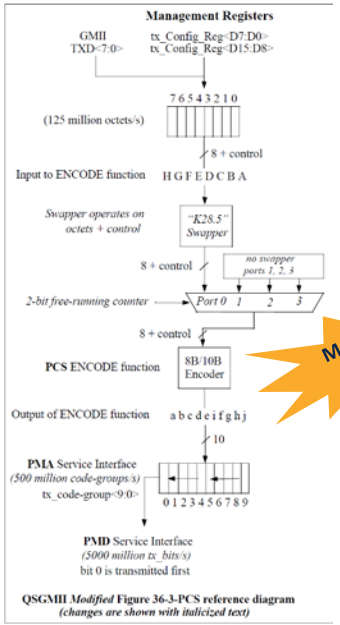

QSGMII PCS RX Scheme

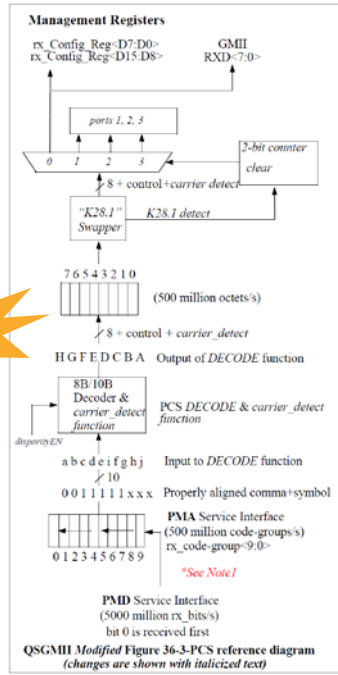

Multiplex faster!

One to Rule Them All!

We're tired of doing these!

Cisco Specification
Universal Serial GMII (aka USGMII)

MII Score Card	
Max Data Rate	1000 Mbit/s, with 8**, 4, and 1** Port Modes
Signal Count	**4/Octal PHY** + 2 Mgmt
SerDes Max Rate	10.0, 5.0, 1.25 Gb/s
Clock Scheme	SerDes
Command Space	Large – Ordered Sets Packet Control Header
Commands Assigned	K28.1 Reserved 8 Non-Idle Sets Defined **PCH Packet Types**
PLCA Support	**No**

USGMII Overview

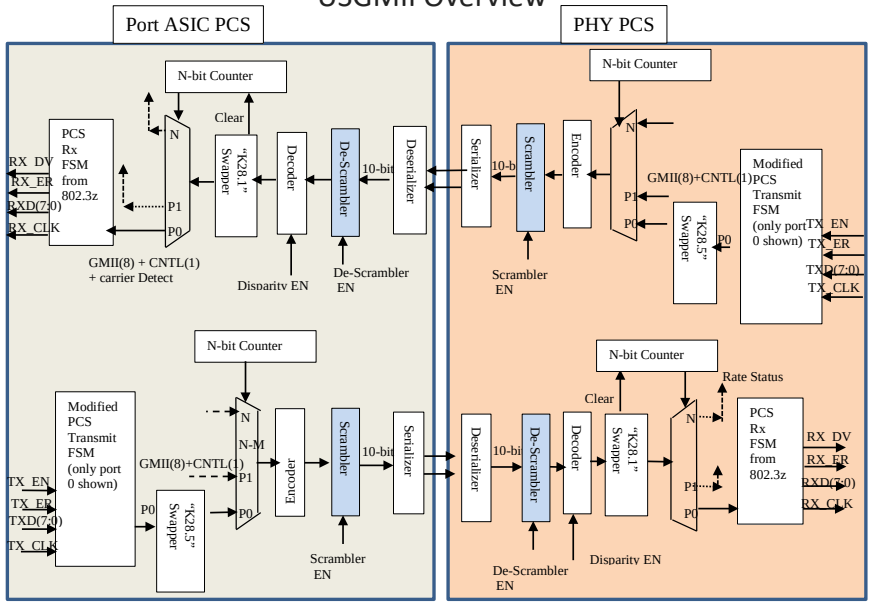

Packet Control Header

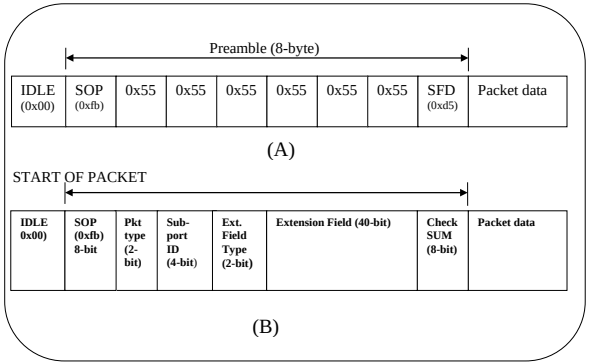

PCH Packet Types

00: Ethernet Packet with PCH
01: Ethernet packet, without PCH (packet information)
10: Idle Packet – Contains status data for a port – no packet data
11: Preemption Frame, aka Interspersing Express Traffic (IET) frame

In-Band PTP Timestamps via Extension Field

MII Evolution
10 Gig / mGig MII Interfaces

XGMII

- XGMII Defined in 802.3 Clause 46. Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII)
  - High complexity 32-bit wide bus
  - Employs lane striping
  - **Overkill for SPE applications**

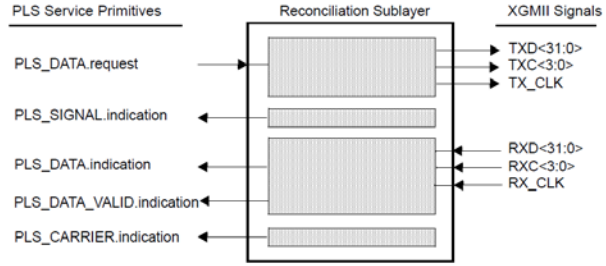

Figure 46–2—Reconciliation Sublayer (RS) inputs and outputs

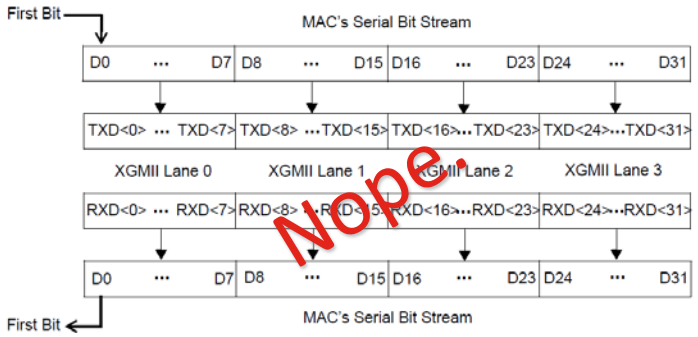

Figure 46–4—Relationship of data lanes to MAC serial bit stream

USXGMII Family

- Two Cisco Specs for mGig
  - Universal SXGMII Interface for a Single MultiGigabit Copper Network Port
  - Universal SXGMII PHY-MAC Interface for Multiple Network Ports
- All the features
  - Clause 49 PCS
  - 1-8 ports per SerDes pair
  - SerDes Speeds from 5Gb/s to 20 Gb/s
  - Packet Control Header
  - Clause 46 ordered sets
- **Again, overkill for SPE**

USXGMII Modes

USXGMII Type	Num Prts	Network Port Types	Replications – Lowest to Highest data speed	PCS	SERDES Speed (Gbps)
10G-USXGMII	1	100M/1G/2.5G/5G/10G	100, 10, 4, 2, 1	Clause 49	10.3125
5G-USXGMII	1	100M/1G/2.5/5G	50, 5, 2, 1	Clause 49	5.15625

USXGMII Modes

MP-USXGMII Type	Num Prts	Network Port Types	Replications – Lowest to Highest data speed	PCS	SERDES Speed (Gbps)
10G-SXGMII	1	10M/100M/1G/2.5G/5G/10G	1000/100/10/4/2/1	Clause 49	10.3125
5G-SXGMII	1	10M/100M/1G/2.5/5G	500/50/5/2/1	Clause 49	5.15625
10G-DXGMII	2	10M100M/1G/2.5G/5G	500/50/5/2/1	Clause 49	10.3125
5G-DXGMII	2	10M/100M/1G/2.5G	250/25/2.5/1	Clause 49	5.156
20G-QXGMII	4	10M/100M/1G/2.5G/5G	500/50/5/2/1	Clause 49	20.625
20G-DXGMII	2	10M/100M/1G/2.5G/5G/10G	1000/100/10/4/2/1	Clause 49	20.625
2.5G-SXGMII	1	10M/100M/1G/2.5G	250/25/2.5/1	Clause 49	2.578125
10G-QXGMII	4	10M/100M/1G/2.5G	250/25/2.5/1	Clause 49	10.3125
20G-OXGMII	8	10M/100M/1G/2.5G	250/25/2.5/1	Clause 49	20.625

Nope.

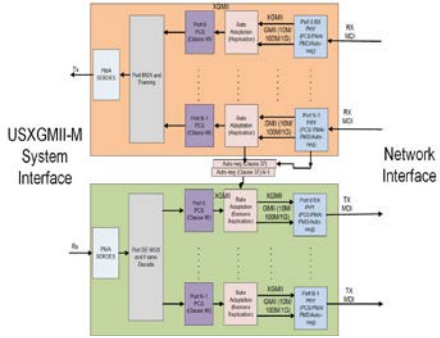

MII Evolution
Path Forward for SPE

Step 1 – PLCA

Amend Clause 36 to add ordered sets for PLCA BEACON and COMMIT functionality

- Enables SGMII to support PLCA
- Also enable USGMII for PLCA
- *Small change* to Table 36-3, add supporting text

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 through 1111	Reserved
1	0	0000 through 1111	Normal data transmission
1	1	0000 through 1111	Transmit error propagation

Table 22–2—Permissible encoding 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 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

Table 36–3—Defined ordered sets

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/
	Encapsulation		
/R/	Carrier_Extend	1	/K23.7/
/S/	Start_of_Packet	1	/K27.7/
/T/	End_of_Packet	1	/K29.7/
/V/	Error_Propagation	1	/K30.7/
/LI/	**LPI**		Correcting /LI1/, Preserving /LI2/
/LI1/	LPI 1	2	/K28.5/D6.5/
/LI2/	LPI 2	2	/K28.5/D26.4/

[a] Two data code-groups representing the Config_Reg value.

Step 2 – Interfaces

Define new logical interfaces

- Single-port 10/100 SPE interface
- Multi-port 10/100 SPE interface

- Criteria
  - Run at appropriate speeds
    - Single 10/100 Port, < 125 MHz
    - Octal 10/100 Port, < 1.25 Ghz
  - Clocking Scheme
    - Avoid SerDes for Single Port
    - Avoid complex clock trees for Multi-Port
  - Supported Features
    - PTP, PLCA, EEE

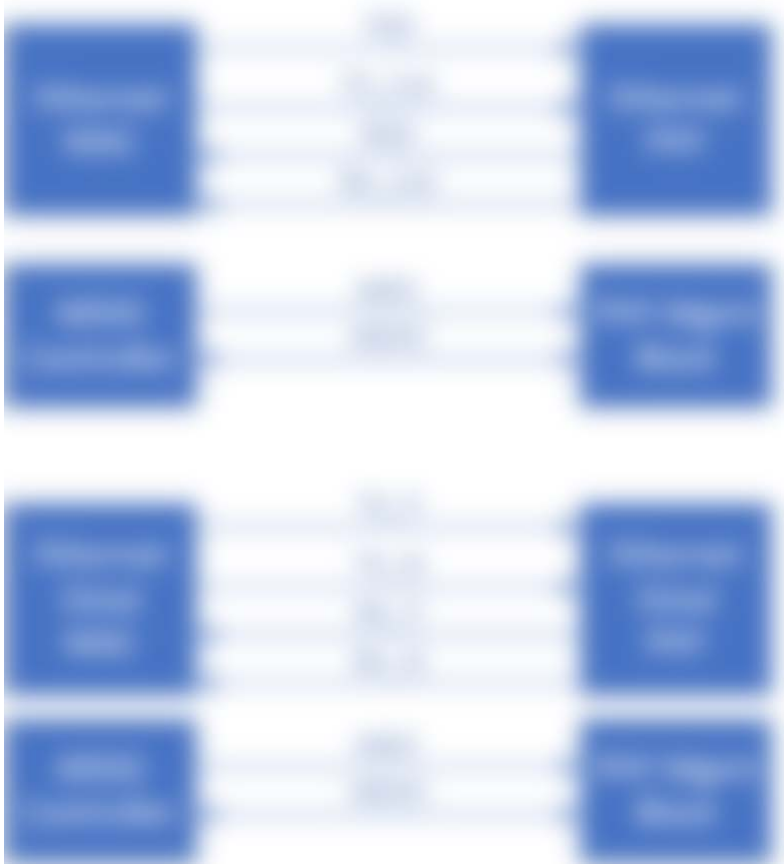

Not another one!

Make it stop!

Proposal
SPE-SMII and SPE-MP-SMII

MII Score Card	
Max Data Rate	100 Mbit/s, with 8 and 1 Port Modes
Signal Count	**4 per PHY** + 2 Mgmt
Clock Scheme	SP -125 MHz Source Sync Clock MP - 1.0 Gb/s SerDes
Command Space	Amended Clause 36 Ordered Sets, Packet Control Header
Features	**PTP Timestamp via PCH** **Preemption via PCH** **EEE via Ordered Sets** **PLCA via Ordered Sets**

Single-Port Configuration

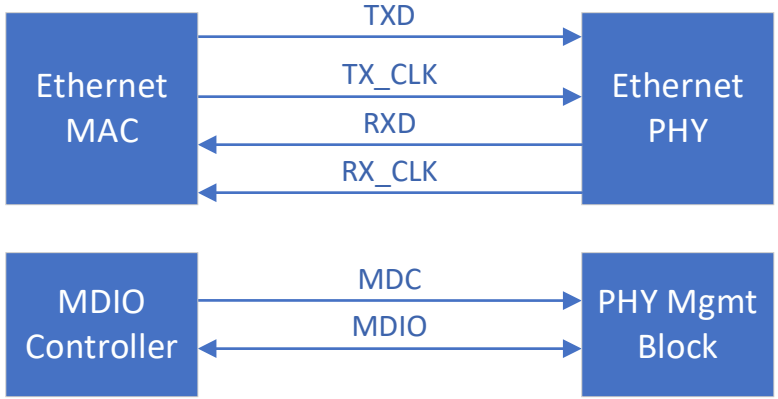

Multi-Port Configuration

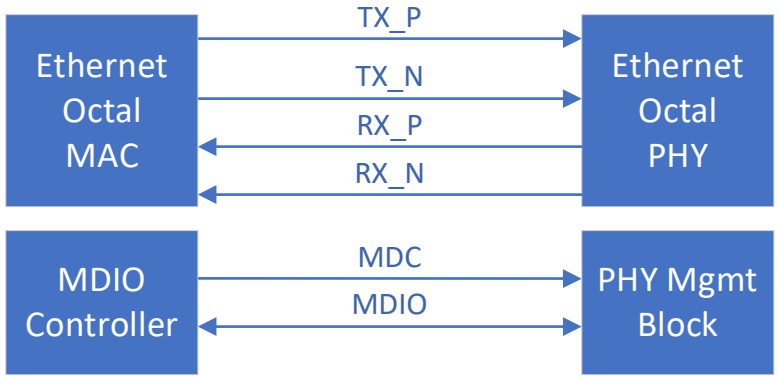

Implementation Rough Sketch

- Start with Cisco USGMII Specification
  - https://developer.cisco.com/site/usgmii-usxgmii/
- Use amended Clause 36 ordered sets with added PLCA support
- Update config register definitions
- Change rates to 1/10th USGMII data rates
  - Consider omitting scrambler
- Adopt SGMII-style source-synchronous clocking for lower complexity options:
  - Single-ended SDR data + 125 MHz Clock
  - Single-ended DDR data + 62.5 MHz Clock
  - Differential Data + Clock (at cost of 2x pins)
  - Differential Data + Clock Recovery (SerDes)

Questions and Discussion

Specific Feedback Sought

- Clause 36 PLCA Ordered Set fix should be done in IEEE 802.3, but should it be part of .3dg or a one-off quick effort?
- Should MII Spec be done as an industry spec to accelerate the existing SPE market?
- Do we want to bump the multi-port SerDes to 1.25 Gb/s to enable in-band MDIO?
- Same question, but for single-port, source-synchronous interface by faster clock and/or DDR data transfer methods to enable in-band MDIO?
- Is using a common Clause 36 PCS coding scheme for both single and multi-port interfaces beneficial?
- Are we deluding ourselves in thinking the MAC will be in charge of PLCA?
  - If so, should we pivot to a model where PHY-based PLCA is explicitly supported?

The bridge to possible