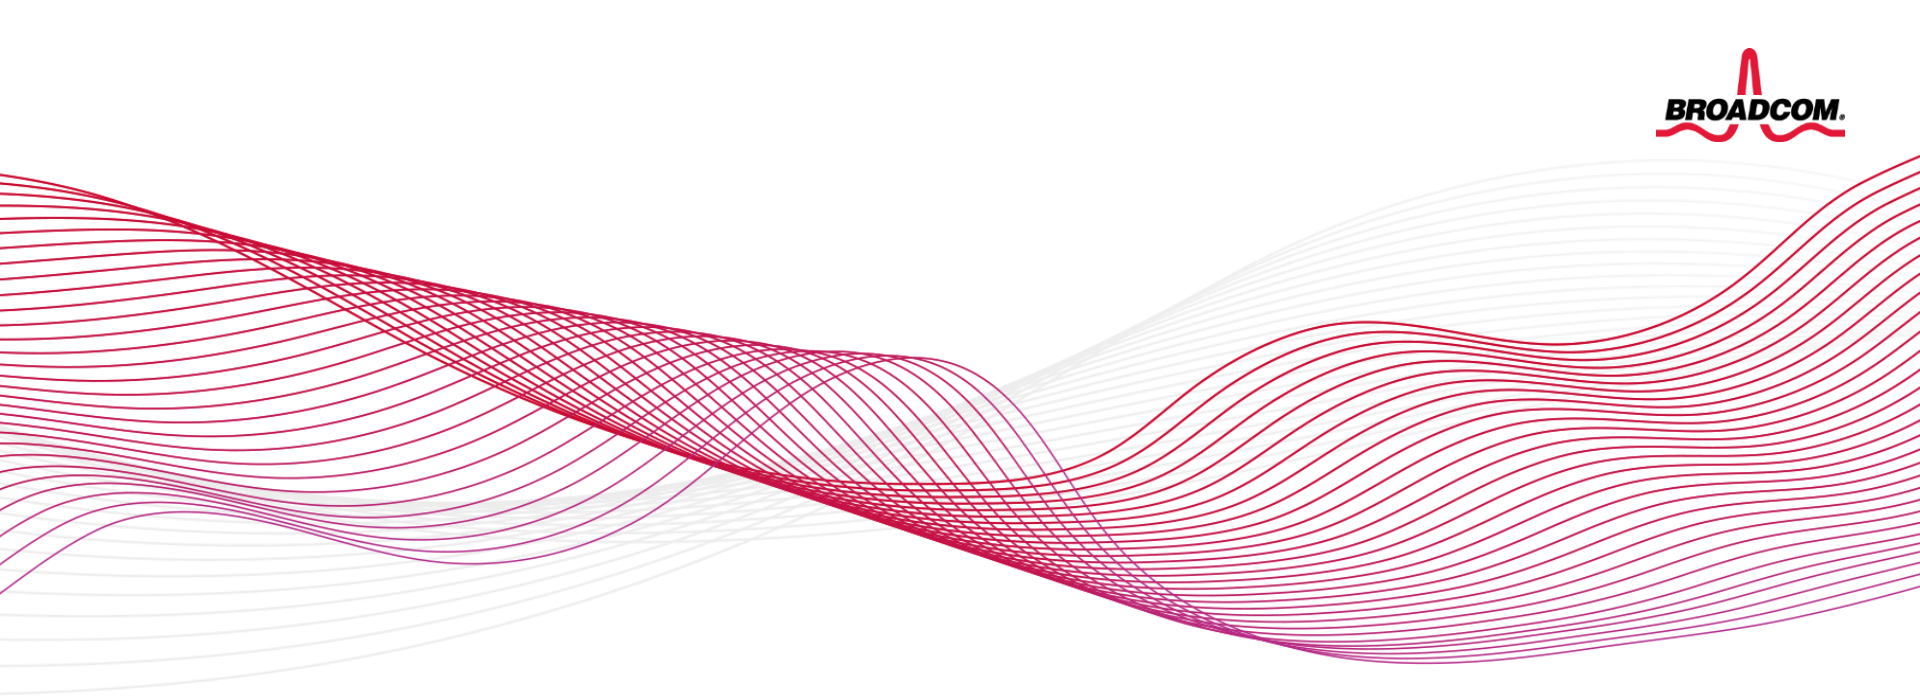

SCALING THE XGMII FOR 2.5/5GBASE-T

Howard Frazier

Architecture ad hoc committee conference call

14-April-2015

OUTLINE

- A blast from the past
- Motivation
- Layering
- XGMII features and benefits
- Scaling
- Alternatives
- Conclusions

A BLAST FROM THE PAST

Interface Locations

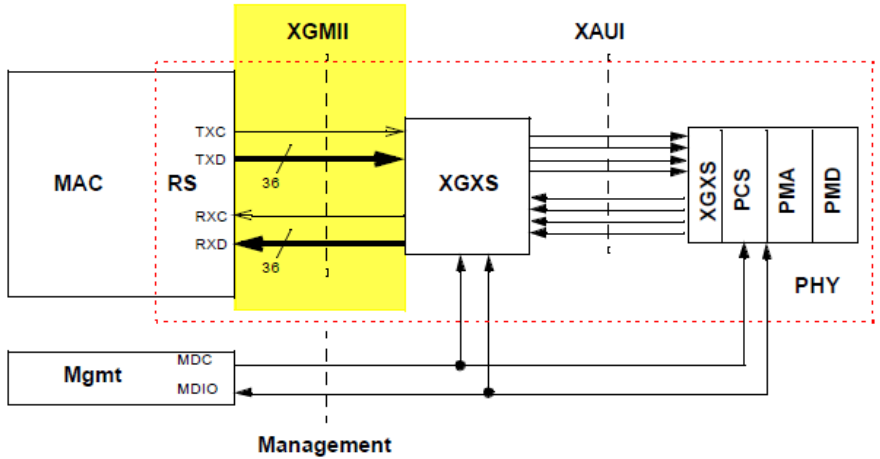

http://www.ieee802.org/3/ae/public/jul00/frazier_1_0700.pdf

IEEE 802.3ae
10 Gigabit Ethernet

MOTIVATION

- **PHYs defined in IEEE Std 802.3 are described in terms of two primary interfaces**
  - Medium dependent interface (MDI)
  - Media independent interface (MII)
- **While the MDI is visible to the user, and usually tightly specified, the MII is often used simply as a convenient way to partition the physical layer specifications from the datalink layer specifications**
- **The MII provides a reference model for conveying data and management information**
- **The MII (which ever version) might not be implemented, but the PHY is described as though it were implemented**
- **Implementations may use an industry standard derivative of the MII (e.g. SGMII, XFI)**
  - The IEEE 802.3 Ethernet Working Group has resisted writing a standard for such interfaces

LAYERING

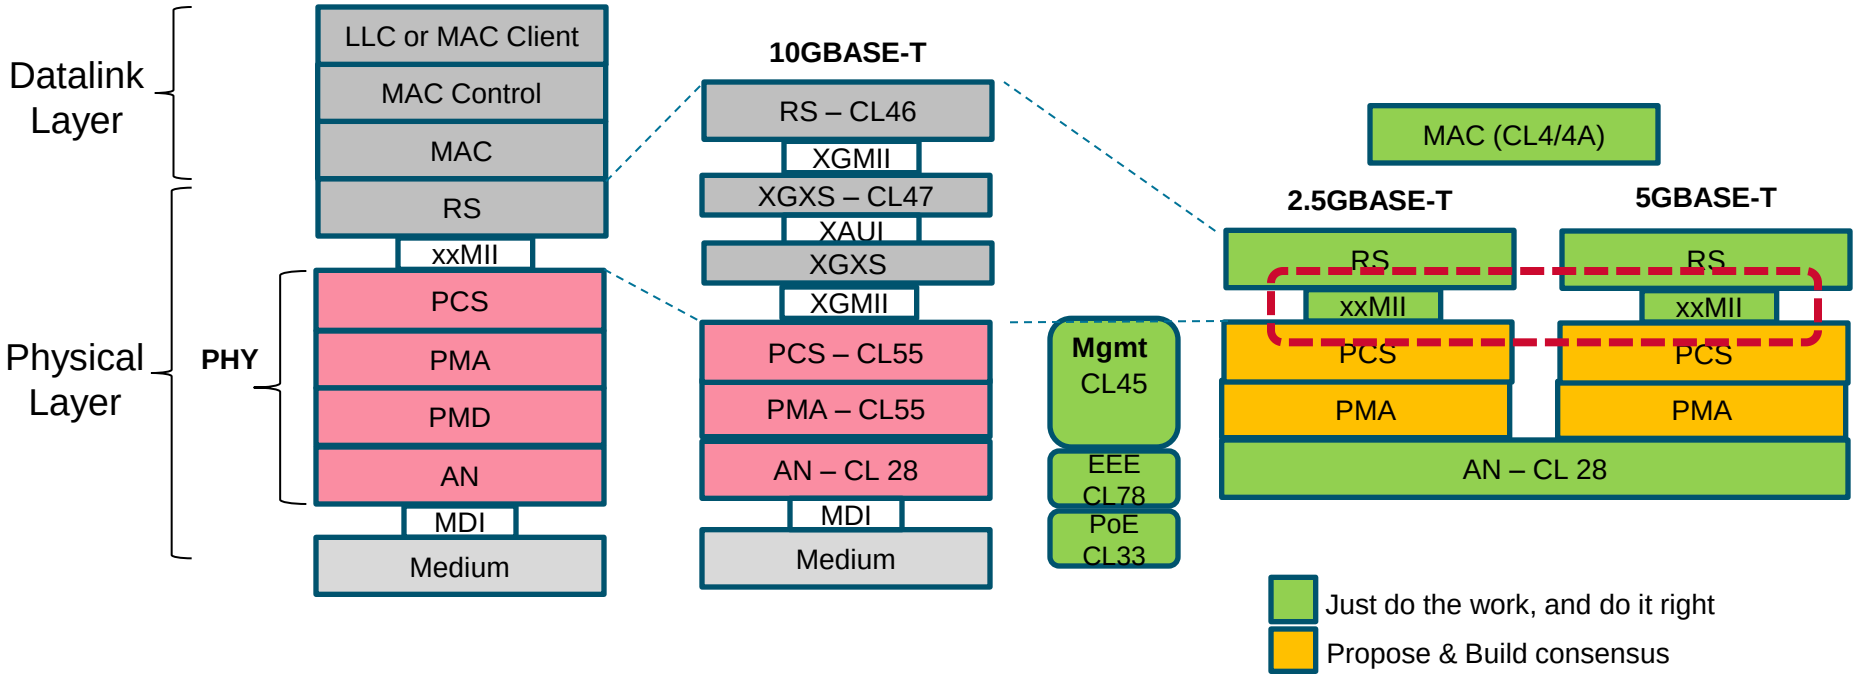
Datalink Layer
Physical Layer
PHY
LLC or MAC Client
MAC Control
MAC
RS
xxMII
PCS
PMA
PMD
AN
MDI
Medium
10GBASE-T
RS – CL46
XGMII
XGXS – CL47
XAUI
XGXS
XGMII
PCS – CL55
PMA – CL55
AN – CL 28
MDI
Medium
Mgmt CL45
EEE CL 78
PoE CL33
MAC (CL4/4A)
2.5GBASE-T
5GBASE-T
RS
xxMII
PCS
PMA
AN – CL 28
Just do the work, and do it right
Propose & Build consensus

XGMII FEATURES AND BENEFITS

- **XGMII was defined in Clause 46 as part of IEEE Std 802.3ae-2002**
  - Baseline proposal adopted in July, 2000
  - Stable, mature, well understood
- **XGMII represented a significant evolution from the MII and GMII**
  - DDR signaling
  - 32 bit data paths
  - Embedded delimiters
  - Clause 45 MDIO/MDC electricals and register addressing
- **XGMII uses source synchronous timing**
  - Amenable to frequency scaling
- **XGMII supports EEE and 802.1AS/802.3bf time synchronization**

SCALING

- **The XGMII specification is largely speed independent**
- **The changes needed to accommodate 2.5G and 5G data rates are very modest**
- **Examples follow**

SCALING (CONTINUED)

46.3.1.1 TX_CLK (~~10 Gb/s~~ transmit clock)

TX_CLK is a continuous clock ~~used for operation at 10 Gb/s~~. TX_CLK provides the timing reference for the transfer of the TXC<3:0> and TXD<31:0> signals from the RS to the PHY. The values of TXC<3:0> and TXD<31:0> shall be sampled by the PHY on both the rising edge and falling edge of TX_CLK. TX_CLK is sourced by the RS.

The TX_CLK frequency shall be ~~156.25 MHz ±0.01%,~~ one sixty-fourth of the MAC transmit data rate ±0.01% .

NOTE—For EEE capability, TX_CLK may be halted according to 46.3.1.5.

- **Make a similar change in 46.3.2.1 RX_CLK**

SCALING (CONTINUED)

46.1.3 Rate(s) of operation

The XGMII supports ~~only the 10 Gb/s~~ MAC data rates of 10 Gb/s, 5 Gb/s and 2.5Gb/s as defined within this clause.

Operation at 10 Mb/s and 100 Mb/s is supported by the MII defined in Clause 22 and operation at 1000 Mb/s by the GMII defined in Clause 35.

~~PHYs that provide an XGMII shall support the 10 Gb/s MAC data rate. 10GBASE-X and 10GBASE-R PHYs operate at a 10 Gb/s data rate.~~

SCALING (CONTINUED)

- **Suggest leaving the setup and hold time parameters in Figure 46-16 as is**
- **The t_{pwmin} parameter might be restated as 40% of the clock period**

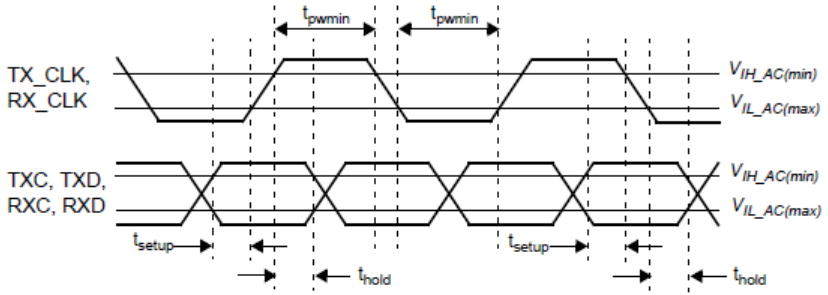

Symbol	Driver	Receiver	Units
t_{setup}	960	480	ps
t_{hold}	960	480	ps
t_{pwmin}	2.5	—	ns

Figure 46–16—TX_CLK and RX_CLK timing parameters

ALTERNATIVES

- **Scale the GMII from 1 G to 2.5G and 5G**
  - Anybody worried about specifying a 625 MHz single-ended interface?
  - Implies re-using Clause 22 MDIO/MDC electricals and register addressing
- **Adopt and scale a de facto industry standard interface such as SGMII, QSGMII, XFI**
  - Perhaps this is best left outside of IEEE 802.3
- **Create a new interface**
  - Volunteers?
- **No logical/electrical interface specification**
  - Use service primitives instead

CONCLUSIONS

- **The XGMII specification is well understood and stable**
- **The industry knows how to create serial variants**
- **The XGMII specification can be scaled for 2.5G and 5G operation with modest changes to Clause 46**
- **The Clause 45 MDIO/MDC register addressing scheme is much preferred over the Clause 22 scheme**

Gracias!