EoBP Auto-Negotiation Requirements

Howard Baumer
Ali Ghiasi

BROADCOM®

Contributors and Supporters

Howard Baumer – Broadcom

Ali Ghiasi – Broadcom

Jeff Lynch - IBM

Joe Abler - IBM

Tom Palkert - Xilinx

Brian Brunn - Xilinx

Outline

- A Few Reminders
- EoBP Phy Implementations & Connections
- Auto-Negotiation Logical Structure
- Auto-Negotiation Signaling Method
- Possible Signaling Methods
- Summary

The EoBP Phys

Per the IEEE P802.3ap Objectives
Dated July 2004

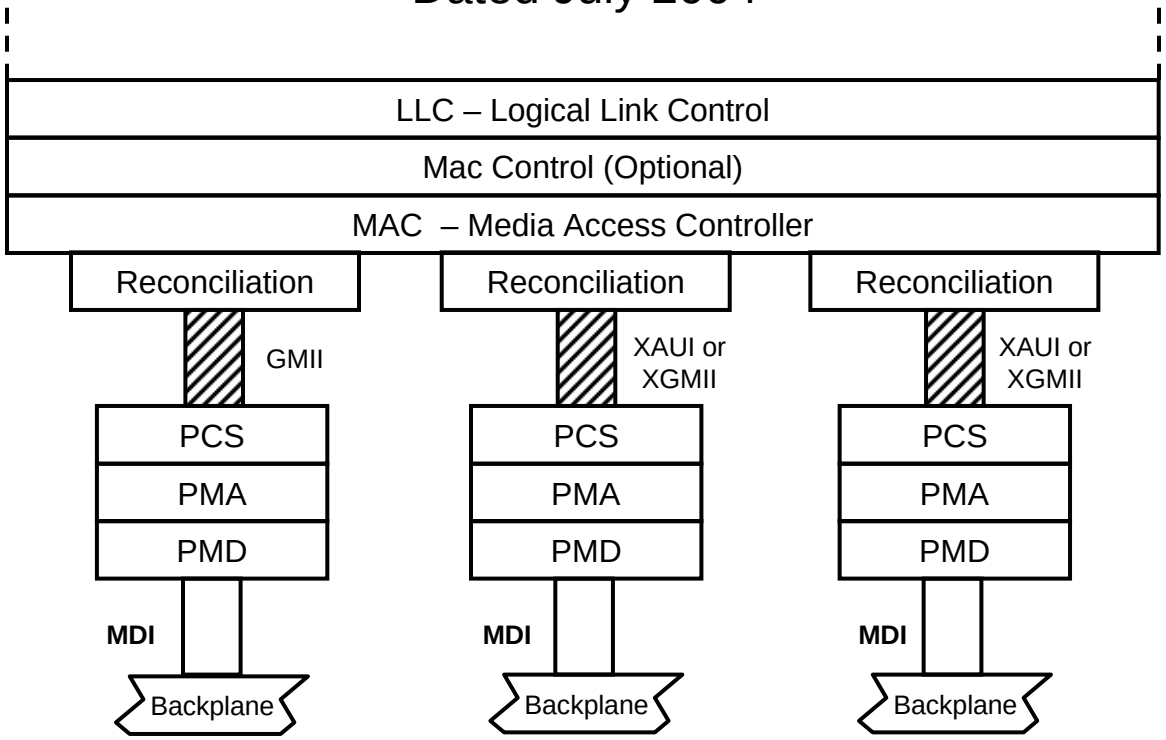

Why Auto-Negotiation

- From the IEEE P802.3ap Objectives, July 2004:
  - "Consider Auto-Negotiation"
- From the July 2004 Portland Plenary, Motion #4:
  - "Auto-Negotiation as a minimum include port-type (e.g. 1G-1-lane, 10G-4-lane, 10G-1-lane) negotiation and any parameter exchange required to select the proper PMA." Yes – 39, No – 0, Abstain - 8
- This gives a good sense that there will be an Auto-Negotiation Clause!

What Auto-Negotiation Does

- Pick the Phy type from a multi-Phy implementation
- Pass parameters to the remote link partner
  - Establish rate matching (full/half duplex, pause abilities, etc.)
- Configure the remote link partner's Transmitter
  - Currently no other 802.3 Phy does this
  - There is a good chance 10GBASE-T will need to do this
  - There is a good chance 10G-1-Lane will need to do this
  - There could be some benefit for 10G-4-Lane

Outline

- A Few Reminders
- EoBP Phy Implementations & Connections
- Auto-Negotiation Logical Structure
- Auto-Negotiation Signaling Method
- Possible Signaling Methods
- Summary

Possible Phy Implementations

- Single Phy Only
  - 1G-1-Lane (1G1): Exists in .3z as 1000BASE-CX
  - 10G-4-Lane (10G4): Exists in .3ae as XAUI and in .3ak as 10GBASE-CX4
  - 10G-1-Lane (10G1): To be defined!
- Multi Phys
  - 10G-4-Lane & 1G-1-Lane: Already implemented in many XAUI and CX4 products
  - 10G-1-Lane & 1G-1-Lane: Very probable
  - 10G-1-Lane & 10G-4-Lane: Probability low?
  - 10G-1-Lane, 10G-4-Lane & 1G-1-Lane: Probability low?

BROADCOM®

1G-1-Lane EoBP Connections

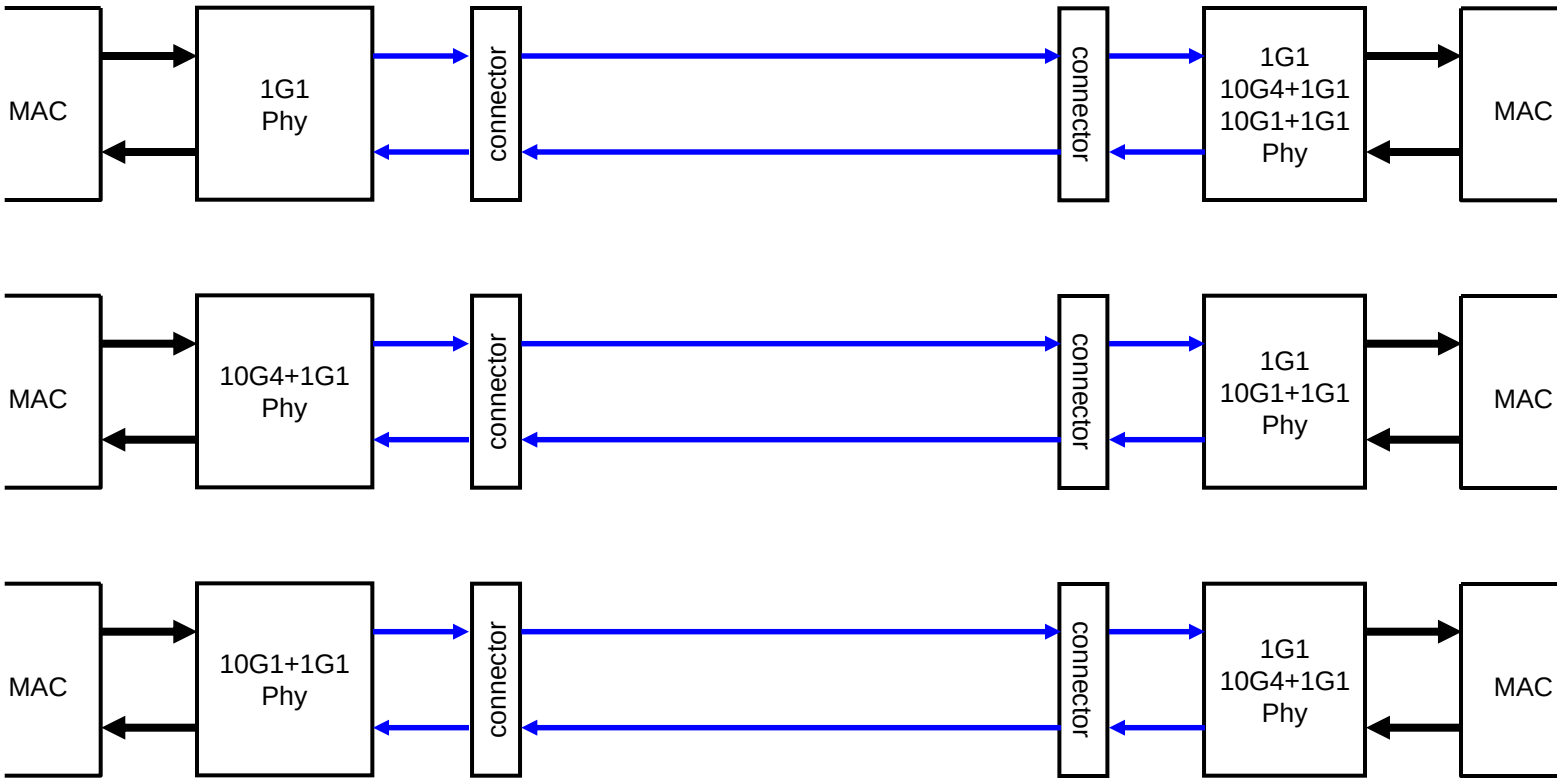

Auto-Negotiation for Backwards Compatibility

10G-4-Lane EoBP Connections

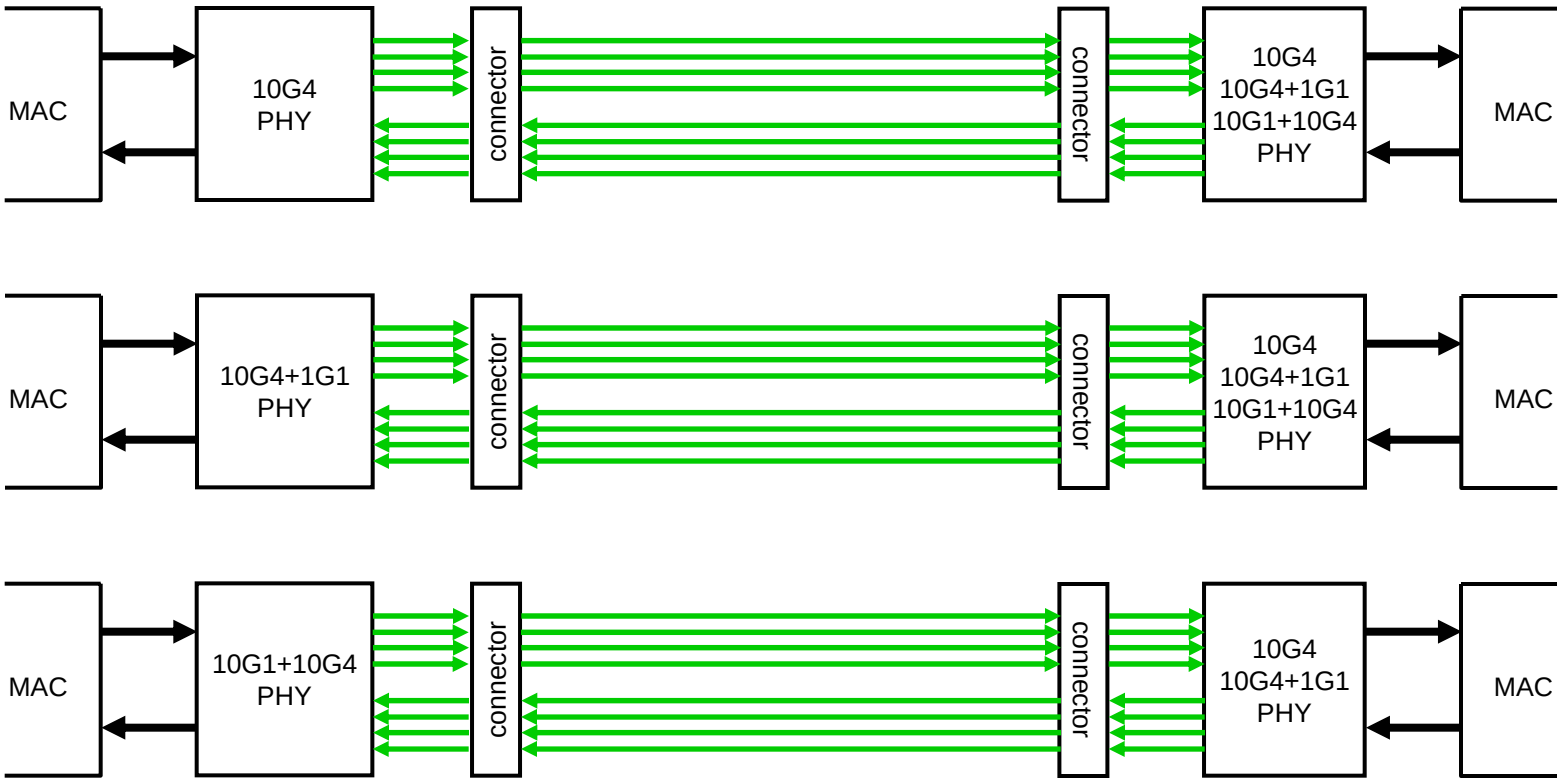

Auto-Negotiation for Backwards Compatibility and Possible Transmitter Configuration

10G-1-Lane EoBP Connections

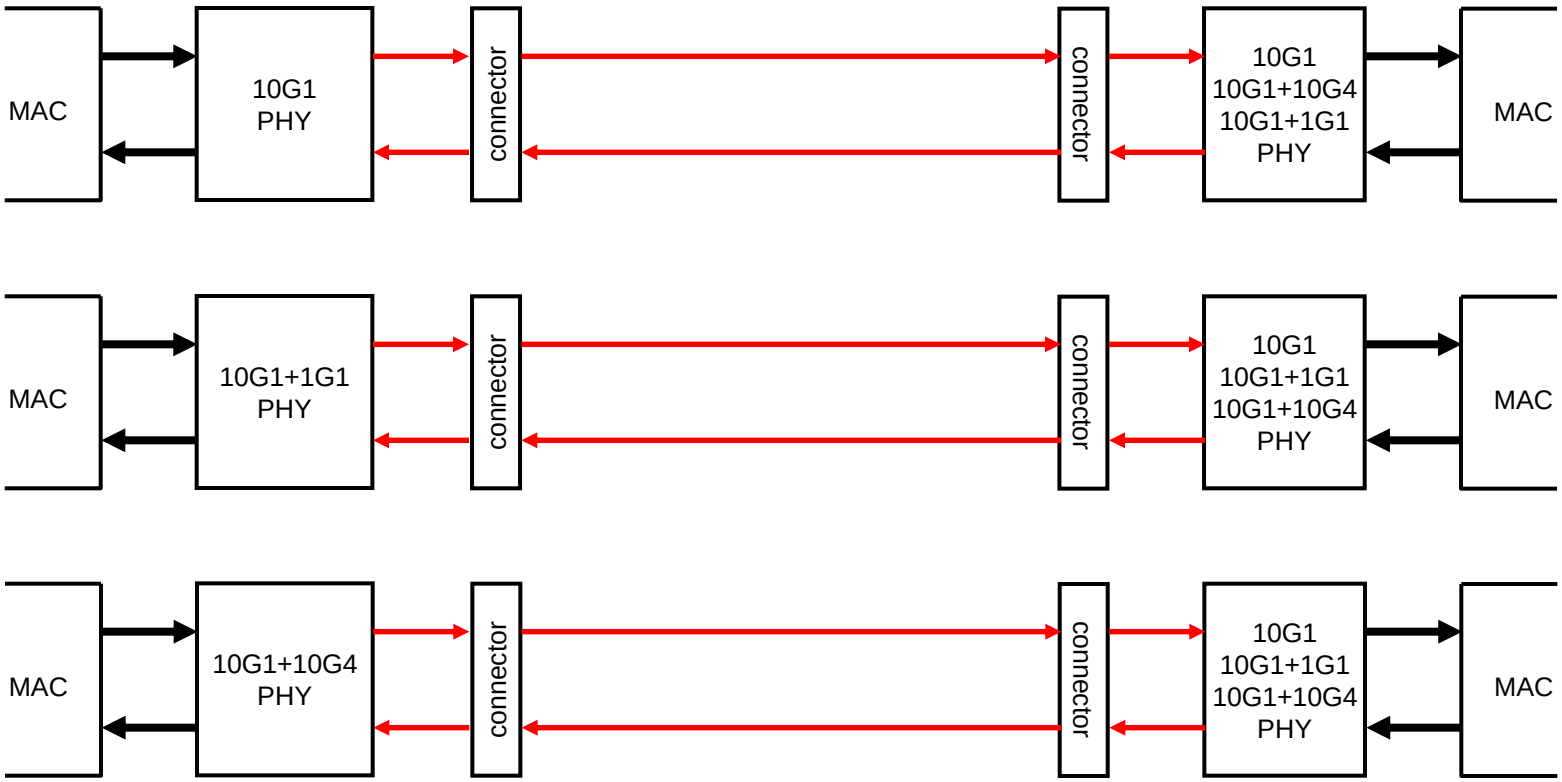

Auto-Negotiation for Possible Transmitter Configuration

Outline

- A Few Reminders
- EoBP Phy Implementations & Connections
- Auto-Negotiation Logical Structure
- Auto-Negotiation Signaling Method
- Possible Signaling Methods
- Summary

Auto-Negotiation's Logical Structure

- State diagrams
- Bit definitions
- Register definitions
- Bit organization within the signaling method
- Etc.

Logical Structure Requirements

1. Compatibility to existing 1000BASE-CX Phy
2. Select PHY/PMA type from multi-Phy implementations
3. Pass parameters between link partners
4. Train and Configure Link Partner Transmitter

Outline

- A Few Reminders
- EoBP Phy Implementations & Connections
- Auto-Negotiation Logical Structure
- Auto-Negotiation Signaling Method
- Possible Signaling Methods
- Summary

Signaling Method Requirements

1. Backwards compatibility with existing phys, 1000BASE-CX.
2. Minimize adding circuits on the raw signal path, **ideally none!**
3. Quick, unnoticeable link up time.
4. Minimize transmit parameter passing turn around time.

Why Backwards Compatibility?

- 1000BASE-X phys have been and are being used as a 1G-1-Lane serdes phys in backplanes for many years now.
- 10GBASE-X phys, using Tx pre-emphasis have been also been and are also being used as a 10G-4-Lane serdes phys in backplanes as well.
- Many backplanes in current use, that use these phys, will not meet the channel definition the Channel Ad Hoc is developing.
- The 802.3ap task force should not replace these existing products with something else.

BROADCOM®

Transmit Parameter Passing Time

- Assume 7 5-bit Transmit parameters to be exchanged
- It takes 1 Base, 1 Message & 4 Unformatted pages for a total of 6 pages
- Each page will take a minimum of 6 transmissions
- Total minimum transmissions is 36.
- Clause 28 FLP style (SSP) transmission time is 1ms / transmission
- Clause 37 @ 1G /C/ transmission time is 32ns / transmission
- Clause 37 @ 10G ||C|| transmission time is 3.2ns / transmission
- For one set of 7 parameters to be passed the total time is:

  CL 28: **36ms** CL 37 @ 1G : **192ns** CL 37 @ 10G: **19.2ns**

HOW MUCH TIME DO YOU WANT TO TAKE?

Outline

- A Few Reminders
- EoBP Phy Implementations & Connections
- Auto-Negotiation Logical Structure
- Auto-Negotiation Signaling Method
- Possible Signaling Methods
- Summary

Phy Types and Connections Revisited

- 5 implementations: 1G1, 10G4, 10G1, 10g1+1G1, 10G4+1G1
- Only one implementation doesn't have 1G and that is the 10G1 only! (Almost all 10G4 phys are 1G capable)
- The lowest common rate for virtually all connections is 1G!

- But, should the 10G only implementation be burden with 1G?

No it shouldn't be!

BROADCOM®

Auto-Negotiation Signaling Possibilities

- A new approach
- Clause 28 FLP style
- Clause 37 via 1000BASE-CX signals
- Clause 37 style using Implemented Data Rates

Common Additions

- Parallel PMA detection (for single speed implementations w/o auto-negotiation)
- 1000BASE-CX, Clause 37 as it is (for backwards compatibility)
- A mechanism for switching between parameter passing and training pattern

A New Approach

- Pros
  - Can make an optimal solution for task at hand
- Cons
  - Additional circuitry and logic since Clause 37 is needed for backwards compatibility
  - Lots more work!

Clause 28 FLP style

- Pros
  - Asynchronous signaling, doesn't need complex CDR or decoder to recover the data
- Cons
  - Adds more load to incoming 10G signal path
  - Adds extra analog circuits: integrators, etc.
  - 1G1 mode still needs Clause 37 for backwards compatibility
  - Adds extra state machines:
    - Will need Clause 28 FLP state machine
    - Need state machine to recover asynchronous data
  - Clause 28 register space is not really reusable, forces extra useless words to be transferred
  - Slowest parameter passing times

Clause 37 via 1000BASE-CX signals

- Pros
  - Already needed for backwards compatibility
  - Fast parameter passing times
- Cons
  - Requires all implementations to have 1000BASE-CX Phy

Clause 37 style using Implemented Data Rates

- Pros
  - Least amount of additional anything
  - Fast parameter passing times
- Cons
  - Needs a slower rate mechanism to handle transmitter configuration parameter passing

10G vs. 1G

- 10G-1-Lane 64B/66B NRZ, Duo-Binary
  - Baud rate = 10.31GBd
  - 1/8th of this = 1.29GBd
- 10G-1-Lane 64B/66B PAM4
  - Baud rate = 5.16GBd
  - 1/4th of this = 1.29GBd
- 10G-4-Lane 8B/10B NRZ
  - Baud Rate = 3.125GBd
  - 1/2 of this = 1.5625GBd
- 1G-1-Lane 8B/10B NRZ
  - Baud rate = 1.25GBd
  - Deviation from above = 3%, 3% & 25%

Outline

- A Few Reminders
- EoBP Phy Implementations & Connections
- Auto-Negotiation Logical Structure
- Auto-Negotiation Signaling Method
- Possible Signaling Methods
- Summary

Auto-Negotiation Requirement Summary

- Parallel detection for single rate PHY / PMA implementations that do not support Auto-Negotiation
- Clause 37 support for PHY / PMA implementations that support 1G-1-Lane: For backwards compatibility to 1000BASE-CX
- Lower speed signaling for 10G PHY / PMA implementations to pass transmitter parameters during training
- Fast parameter passing to allow many training sequence iterations

BROADCOM®

Contributors and Supporters

Multi-Lane EoBP Connections

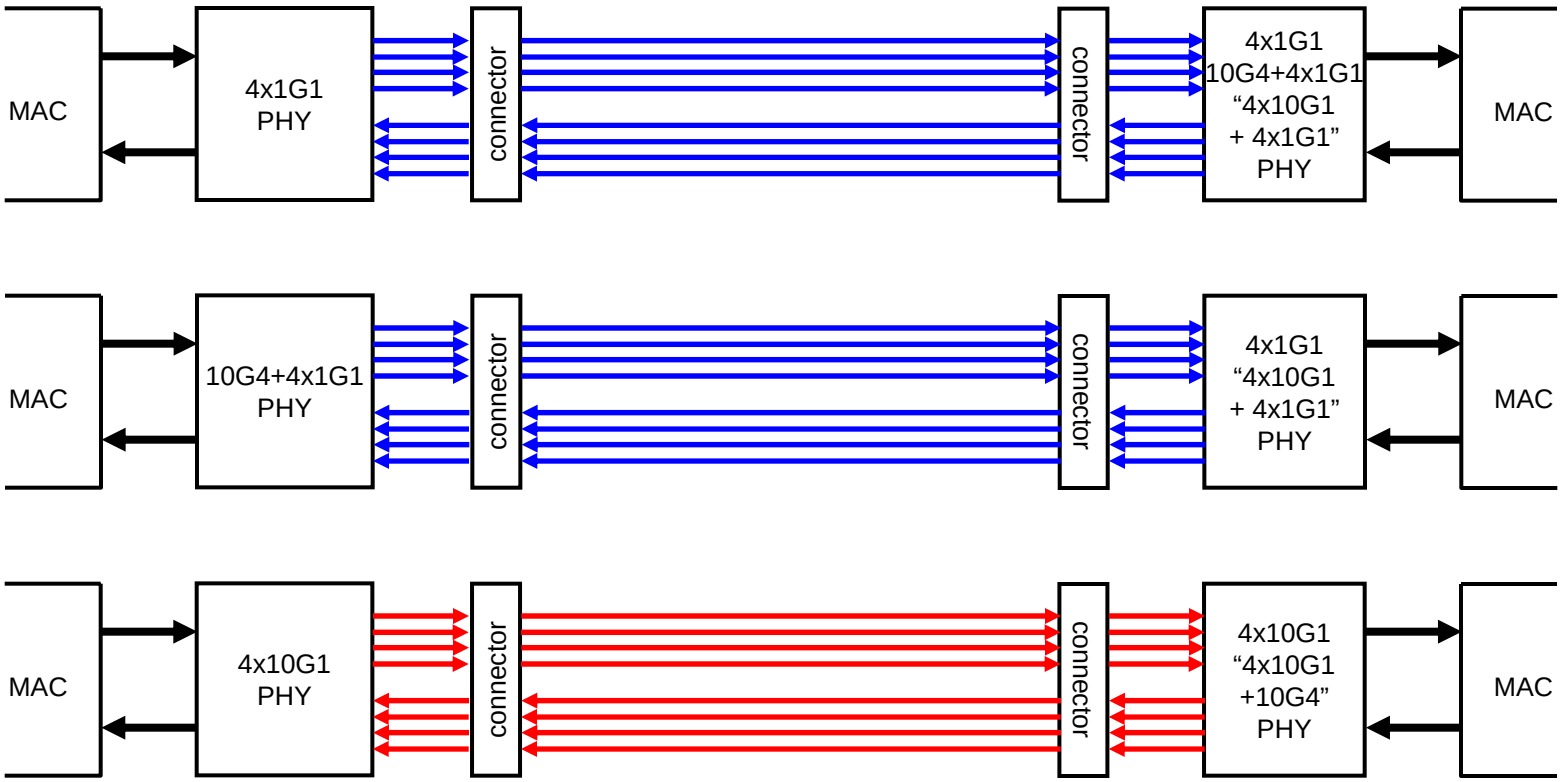

Auto-Negotiation to handle 4-Lane to quad 1-Lane resolution?