VITESSE®

Making Next-Generation Networks a Reality.

Speed Switching without Communication Interruption

Jim Barnette
Mandeep Chadha
Worayot Lertniphonphun

Prepared for the IEEE 802.3 Study Group

Overview

- Motivation
- Speed switching without Communication Interruption
- Speed transitioning
  - Intermediate states
- Proposed new PHY names
- 1000BASE-T-100
  - Sketch
  - Benefits and Challenges
- Special case for 100BASE-TX Energy-Efficient Ethernet
- 10GBASE-T-1G sketch
- Summary

Motivation

- DSP-based Physical-layer transceivers require significant training time when first bringing up the link at a new link speed
  - Timing acquisition
  - Equalization
  - Noise cancellation
- When transitioning between speeds at which the link was previously up, an uncertain amount of time will be required to retrain adaptive filter coefficients
  - Timing re-acquisition
  - Temperature-induced equalizer and noise canceller coefficient drift
- The IEEE 802.1 A/V bridging task force indicated that a link interruption of >1ms will lead to A/V stream interruption.
- Thus, we need a way to prepare the link at the new speed while keeping the link active at the current link speed:

Speed Switching without Communication Interruption

Speed Switching without Communication Interruption

- Use all twisted-pairs for highest speed; 1 or 2 pairs for lower speeds
  - Requires new subset PHY specifications for PHYs that run over all 4 pairs.
  - New modes are very similar to the existing one to enable transparent speed downshift.
- Tx and Rx may independently switch their speed.

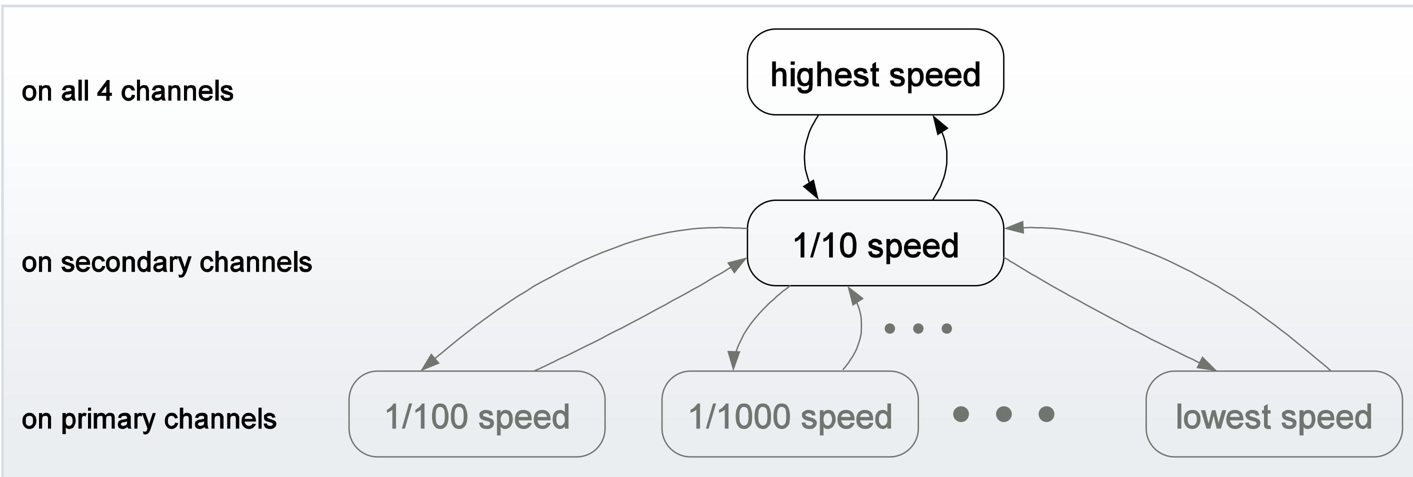

Speed Transitioning

- Downshift from all 4 twisted-pairs to 1 or 2 twisted-pairs
  - Tx indicates that it is about to downshift and then switches speed
- Other speed shift
  - Tx indicates desire to upshift and begins transmitting on “new” pairs
  - When Rx side ready, it sends a response through in-band channel
    - Possibly by encoding the response in the idle pattern during inter-packet interval
  - Tx indicates that it is about to upshift and then switches speed

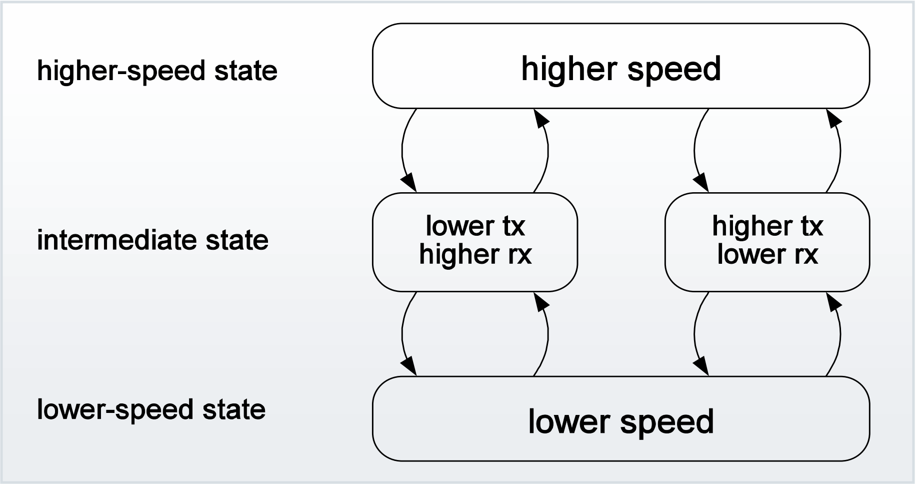

Switching speed from all 4 twisted-pairs to 1 or 2 twisted-pairs

- Using 1 pair to Tx or Rx for lower speed
- Can be implemented on duplex or simplex, prefer simplex
- One specific slave receiving channel is used for lower speed mode
  - Reason for choosing channel C/D → enables legacy 10/100 for downshift

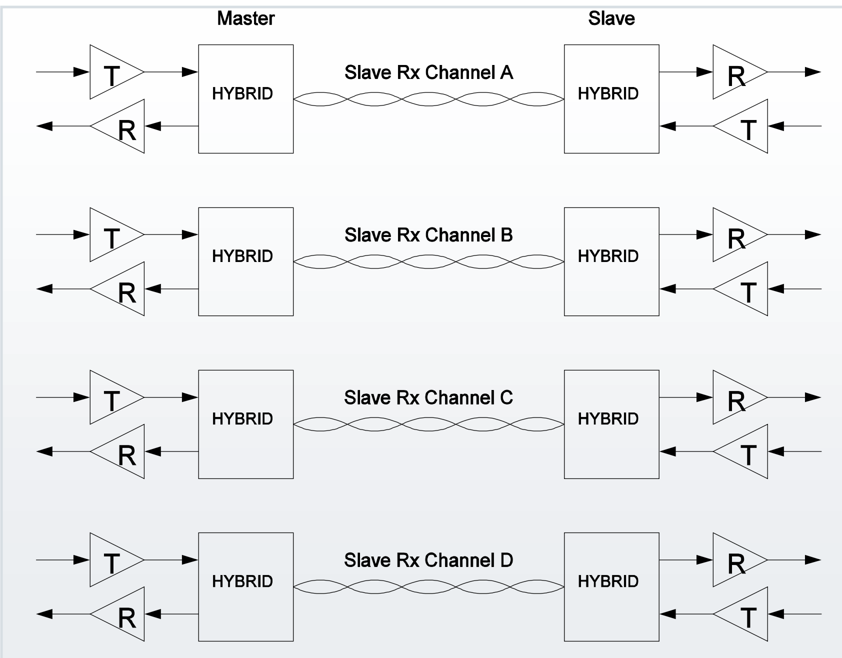

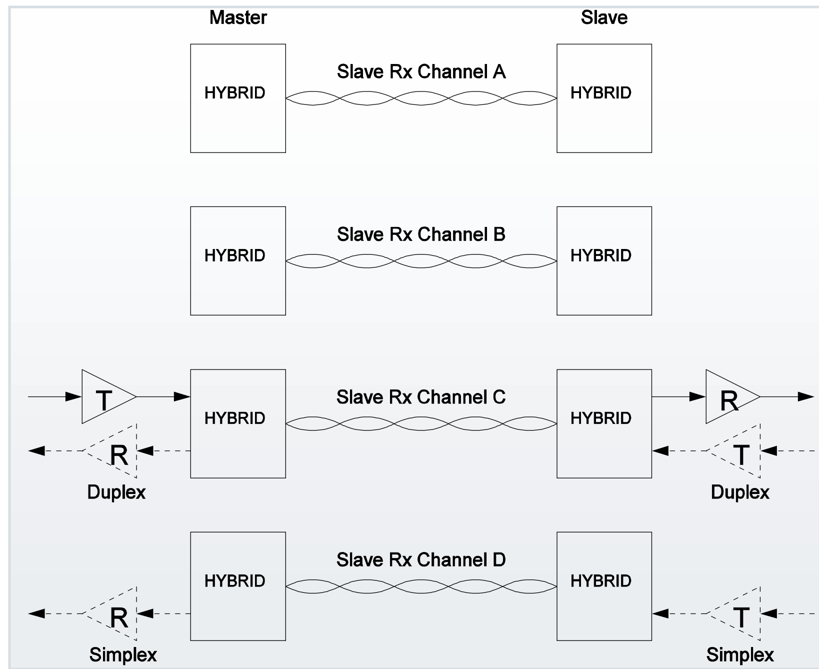

Intermediate States

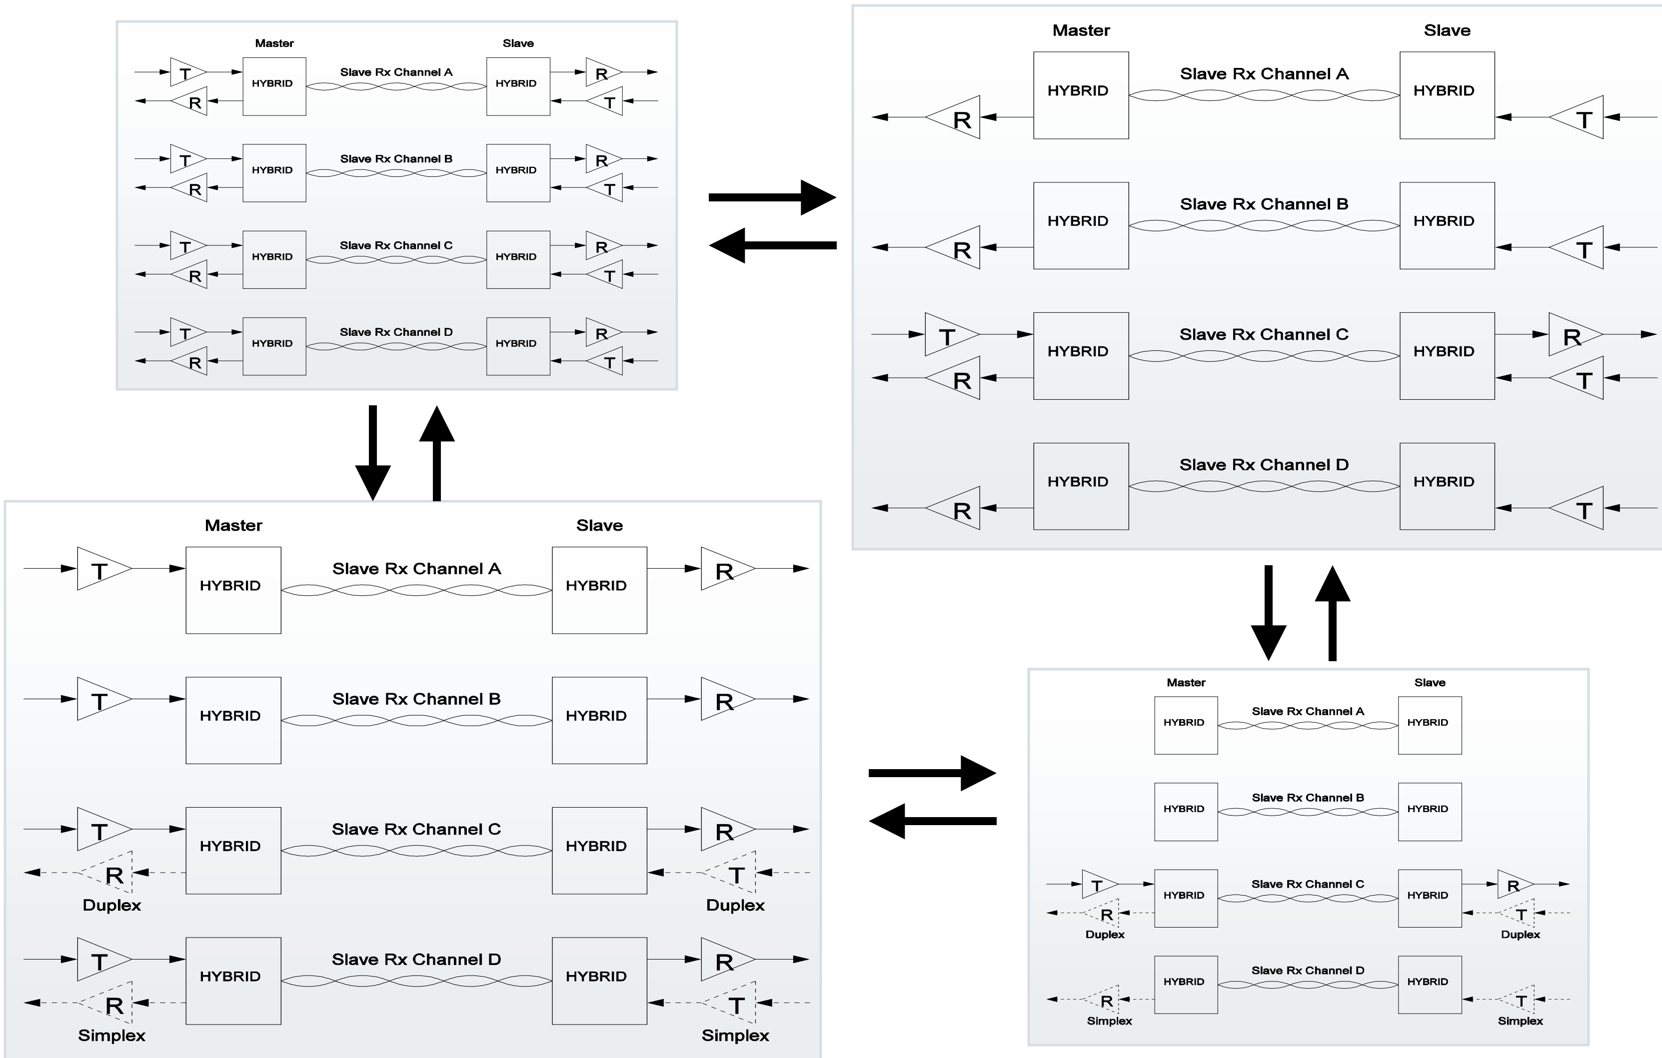

Proposed new PHY names

- 10BASE-TEEE

  EEE 10Mb PHY that is fully interoperable with 10BASE-T over 100m Cat-5

- 1000BASE-T-100

  EEE 100Mb subset PHY for 1000BASE-T

- 10GBASE-T-1G

  EEE 1000Mb subset PHY for 10GBASE-T

1000BASE-T-100: 100Mb subset PHY for 1000BASE-T

- BAUD rate remains at 125 MHz
- Reduce from 4D transmission to 1 twisted-pair
- Reduce from PAM5 {±2,±1,0} to PAM2 {±1} keeping shaping filter
- Transceiver remains locked across transition from 1000 to 100Mb by using 1000BASE-T scrambler
- Transition after defined packet or special idle code
- Must choose either simplex on two pairs or duplex on one twisted-pair: EC training drift issue
- Uses pairs C&D for 100Mb thus enabling glitch-free 2nd decade downshift to 10BASE-TEEE

1000BASE-T-100 Benefits

- Downshift may occur asymmetrically
  - Enables power savings when only one link direction requires high speed
  - After both sides at 100Mb, 2nd decade downshift may also be done asymmetrically to 10BASE-TEEE
- Lower-speed link kept until higher speed ready
  - When upshift desired, transmission begins on otherwise unused pairs enabling link-partner to re-acquire phase lock and refresh possibly drifted coefficients while still communicating at current operating speed
- Adapted receiver coefficients maintain continuity

1000BASE-T-100 Challenges (1)

- Two-pair simplex vs. one-pair duplex
  - Simplex offers simplest, 100BASE-TX-like, receiver
  - Duplex keeps EC up-to-date simplifying upshift design
- Downshift
  - Slave may downshift-at-will
  - Master must await slave downshift unless we require slave to transmit on all 4 pairs using timing recovered from one pair

1000BASE-T-100 Challenges (2)

- Upshift
  - If simplex is chosen, upshift will expose PHY to risk of drifted EC coefficients while link speed was reduced on pair still being relied upon for 100Mb communication
  - Slave must await Master upshift unless we require slave to transmit on all 4 pairs using timing recovered from one pair
- Asymmetric downshift
  - Requires support from higher layer protocols
  - Other-direction transition time may be bounded

Special Case for 100BASE-TX

- No elegant solution for 100BASE-TX EEE

1. 100Mb on pairs A&B to 10Mb also on pairs **A&B**
   - – Communication interruption (May be acceptable at this data rate)
   - \+ Requires only 2-pair connection

2. 100Mb on pairs A&B to 10Mb on pairs **C&D**
   - \+ No communication interruption
   - – Requires 4-pair connection
   - – C/D pairswap unknown

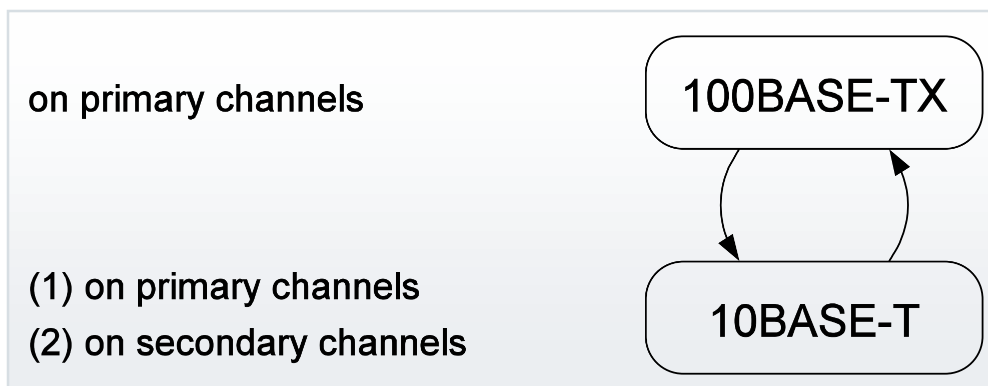

10GBASE-T-1G: 1000Mb subset PHY for 10GBASE-T

- BAUD rate remains at 800 MHz
- Reduce from 4D transmission to 1 twisted-pair
- Reduce from PAM16/DSQ128{±1,±3,…,±15} to PAM4/DSQ8{±4,±12}
- Transceiver remains locked across transition from 10G to 1000Mb by using 10GBASE-T scrambler (clock gated down to lower rate)
- Transition after defined packet or special idle code (transition at LDPC frame boundary!)
- Must choose either simplex on two pairs or duplex on one twisted-pair: EC training drift issue
- Uses pairs C&D to enable glitch-free downshifts

Summary

- The proposed technique enables the communications link to remain up at the current operating speed while training the new receiver for the new operating speed
- A feasible implementation for a subset PHY for 1000BASE-T is sketched and challenges are discussed
- Options are identified for 100BASE-TX Energy-Efficient Ethernet for subsequent decision
- A similar subset PHY approach may be applied to 10GBASE-T for EEE with the possibility of different tradeoffs than those that will be made for 1000BASE-T (e.g. simplex/duplex, master downshift)