



A Gigabit “Subset PHY” Approach for 10GBASE-T Energy Efficient Ethernet

Scott Powell, Wael Diab Broadcom

Atlanta, Georgia Nov, 2007


Subset PHY Review

- 802.3 defines a set of distinct copper PHYs
 - Each PHY defined for single speed operation
 - Ability to rapidly switch between PHYs not considered
- Alternative is to define a set of line codes specifically designed to support rapid switching between data rates
 - Highest data rate “parent” PHY is a standard n BASE-T PHY
 - Lower data rate subset PHYs designed to retain information necessary for rapid switching back to highest data rate
 - Synchronization, equalizer coefficients, canceller coefficients, frame alignment
- Line code for lower data rate is a simple subset of the higher data rate (standard) mode
 - Subset PHY implemented by simply turning off elements of higher data rate standard parent PHY

Potential Subset Modes for 10GBASE-T

1. Turn off 3 of the four channels

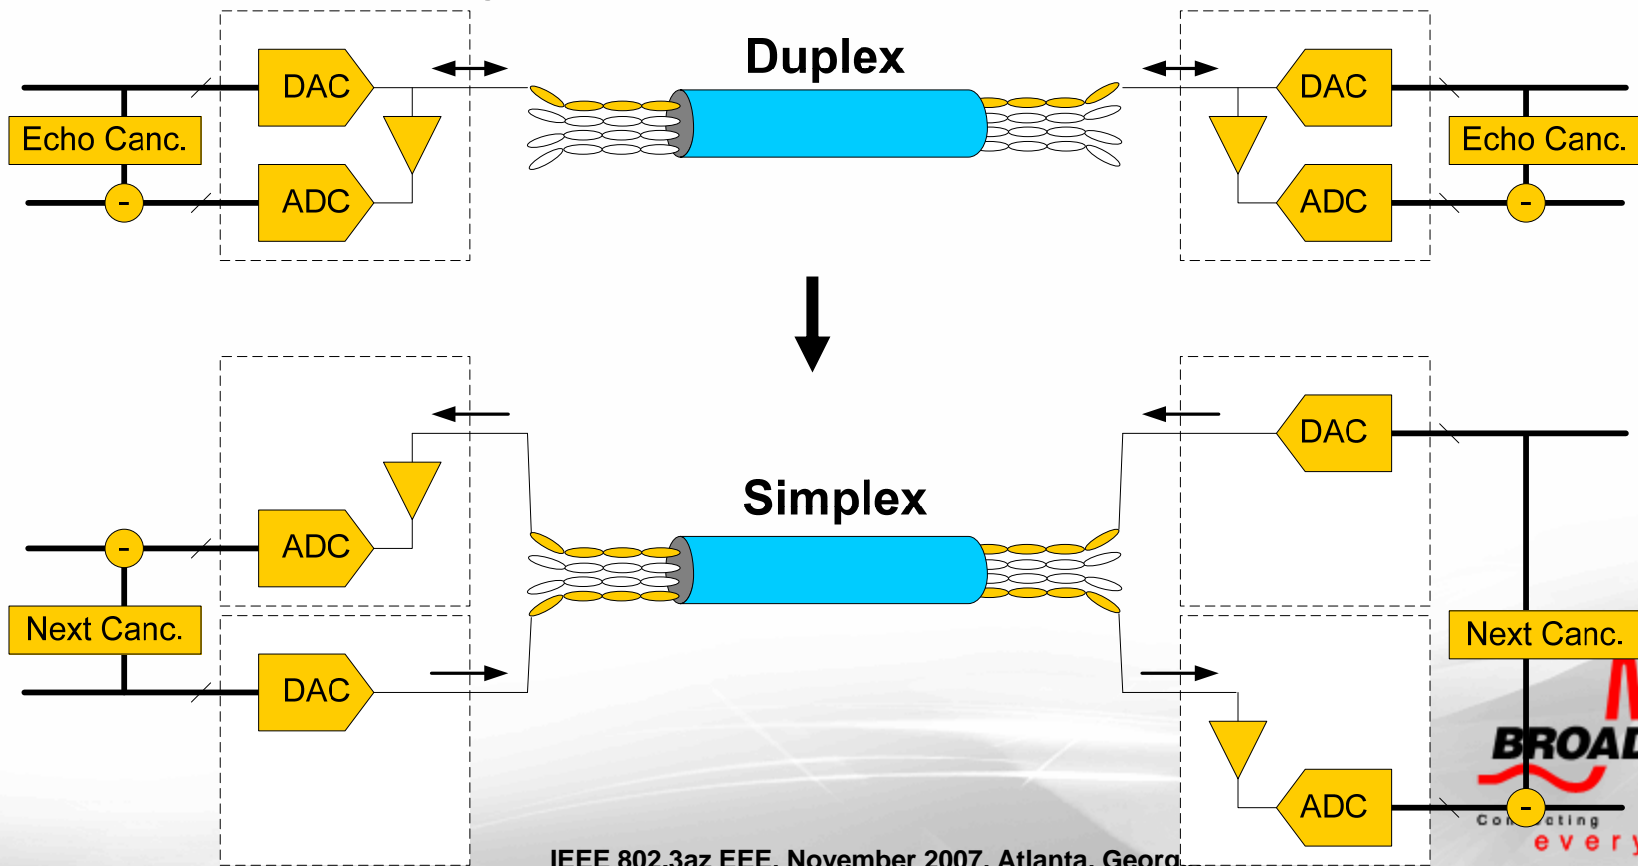
- 800Mps symbol rate remains the same for the active channel
 - Equalizer and echo canceller continue to adapt for active channel
- Synchronization and master/slave loop timing retained
- First step in reducing rate to 1Gbps



Potential Subset Modes for 10GBASE-T

2. Turn off echo cancellers and hybrids (duplex to simplex)

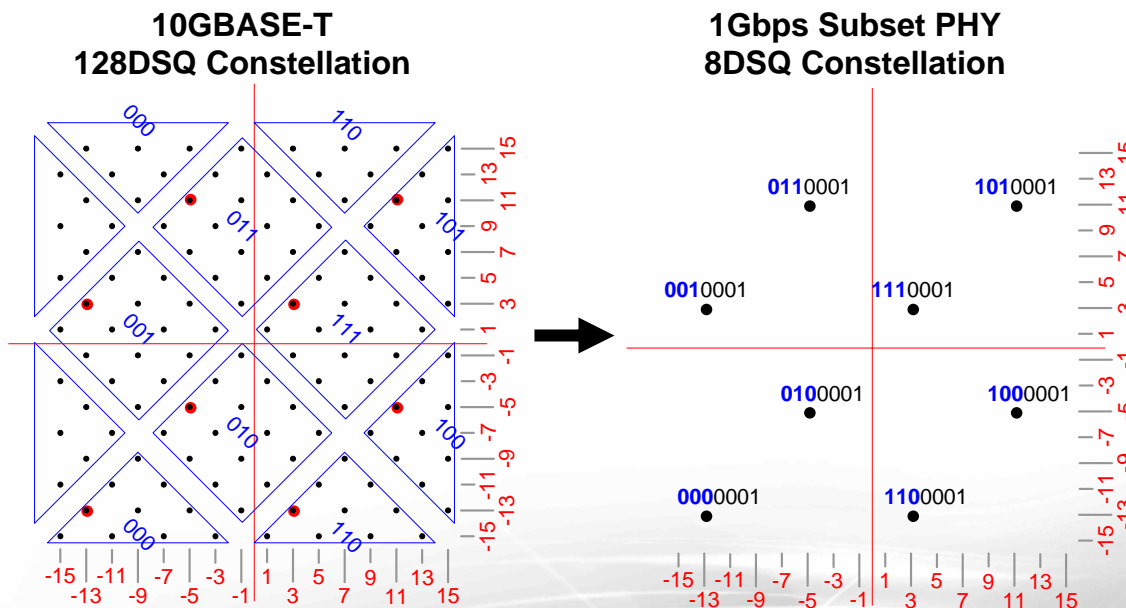
- Next canceller is usually less complex
- Better dynamic range on ADC (no echo)
- Does not change the number of active transmitters or receivers



Potential Subset Modes for 10GBASE-T

3. Turn off extra constellation points

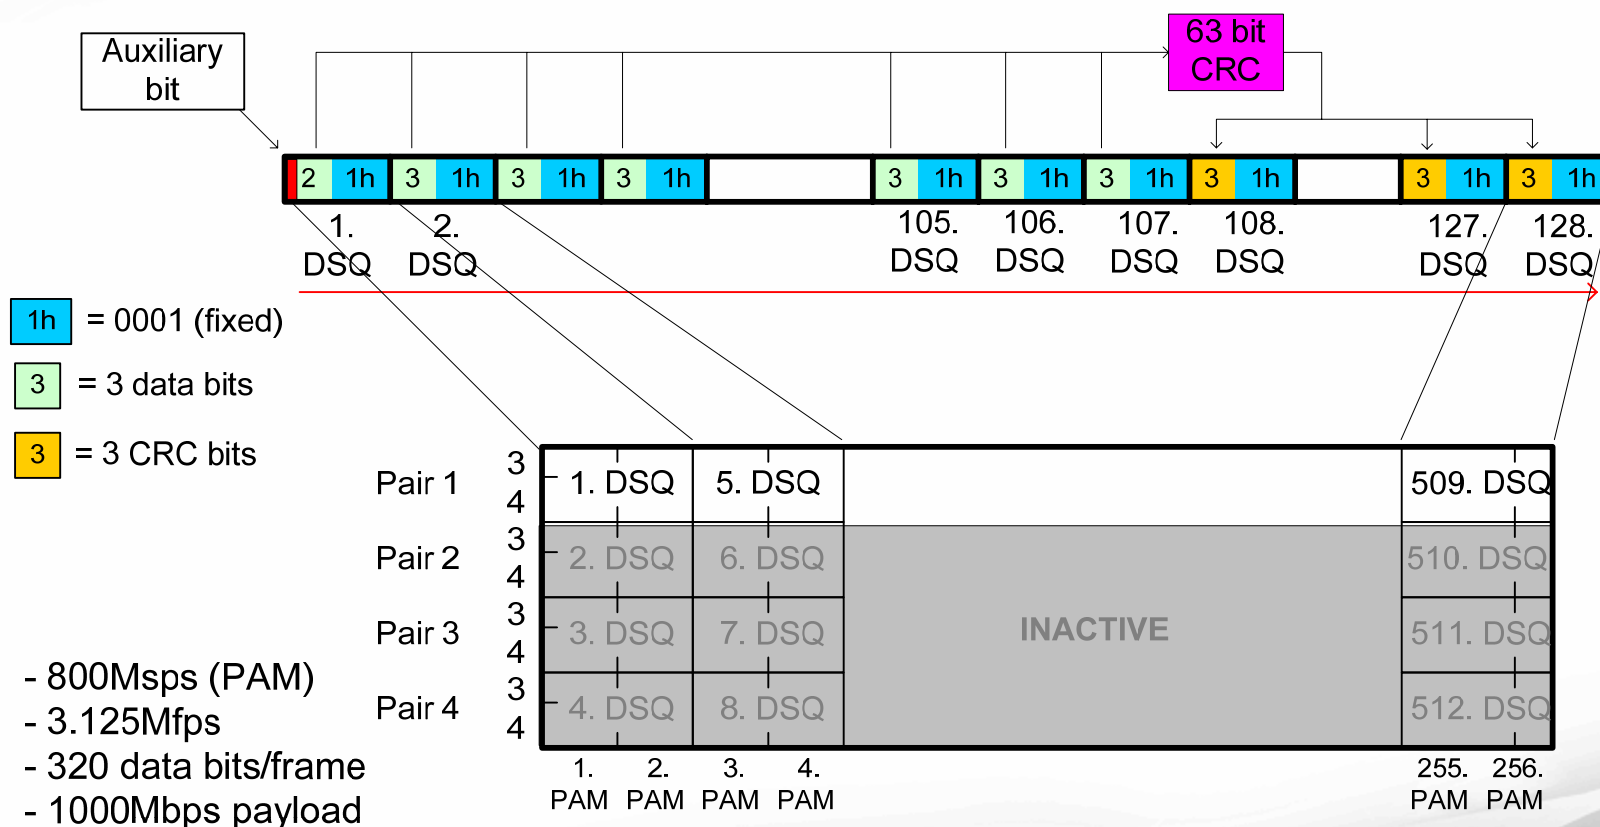
- 8DSQ symbols formed by 10GBASE-T 128DSQ subset partitions
- Turn off LDPC encoder and decoder (big power savings)
- Data protected by subset partition gain
 - Same BER as full rate 10GBASE-T
- Enough "spare" bits to add 63b CRC for error detection ($P_{\text{undet}} < \sim 10^{-19} \cdot P_{\text{err}}$)



- 8 Subsets (uncoded)
- 16 points per subset (coded)

1Gbps Subset PHY Framing/mapping

- Simple modification to 10GBASE-T Clause 55 framing:



Updating Inactive Channels

- Zimmerman_01_0307.pdf presents measurements indicating:
 - The state of a 10GBASE-T PHY need only be updated every 3~5 minutes
 - Updates should require less than 20ms
- Inactive channels and cancellers not being used for anything else
 - Available to be used for refresh operation
 - Does not impact data transfer, link maintained
 - Does not require switching back to higher speed
- Inactive channels can be updated by turning them on for 20ms every 3 minutes
 - Transmit idles
 - Very low duty cycle makes update power negligible

An Additional OOB Channel

- Subset PHY approach provides an additional OOB channel that is independent of the packet transfer
 - One use could be to synchronize updates of inactive channels
- 10GBASE-T Clause 55 reserves one bit per LDPC frame for use as an OOB communication channel
 - “The use of this bit for vendor-specific communication is outside the scope of this document. For the purposes of this standard it is ignored by the link partner.”
- The 10GBASE-T auxiliary channel allows for 3.125Mbps full duplex communication
 - Does not interfere with data traffic



1G Subset PHY Power Estimate

- Analog

- 3 of 4 ADCs and DACs powered down
- All 4 hybrids powered down
- [Estimate 85% power saving](#)

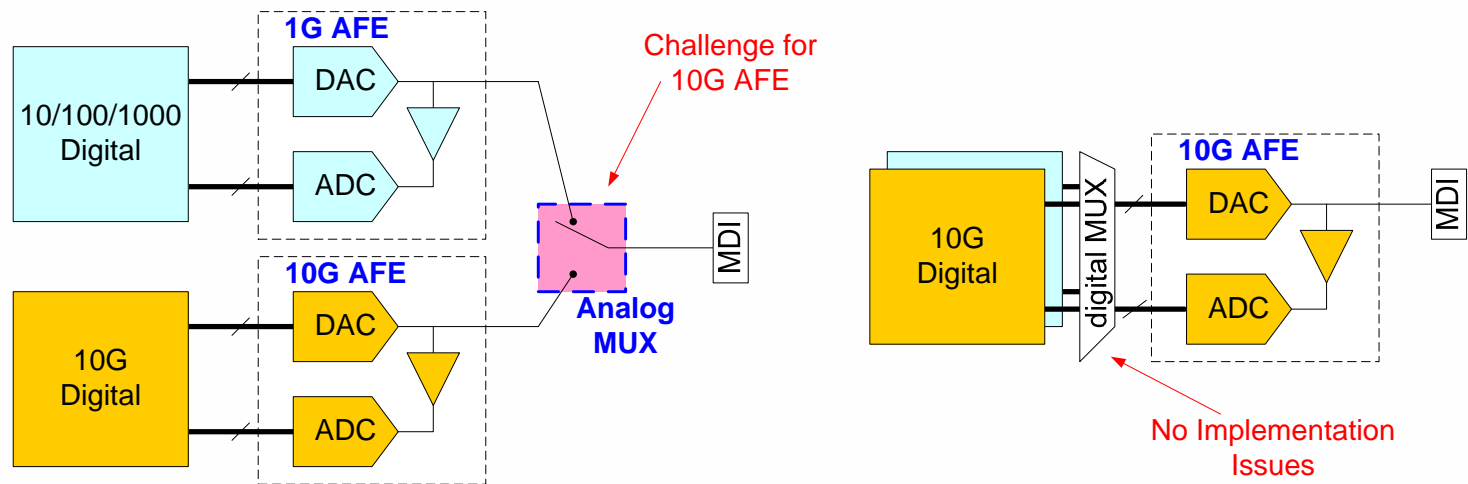
- Digital

- All 4 echo cancellers powered down
- 11 of 12 NEXT cancellers powered down
- All 12 FEXT cancellers powered down
- 3 of 4 equalization filters powered down
- 3 of 4 Precoder filters powered down
- LDPC encoder and decoder powered down
- [Estimate >85% power savings](#)

- Estimate ~85% power reduction in 1G subset PHY mode
 - About 7x power reduction

Practical Considerations for Switching Between Standard PHY Types

- 10/100/1000BASE-T and 10GBASE-T must share the same MDI
 - Likely to have one AFE for all speeds, could have separate digital



Option 1: Switch at AFE Output

Option 2: Switch at AFE Input

- AFE must be sized to support 10GBASE-T
 - Number of bits ADC/DAC, PGA, linearity, jitter, etc. set by 10G spec
- 10/100/1000BASE-T may not be the lowest power way to implement 10/100/1000 Mbps data rates within a 10GBASE-T PHY

Summary

- Line code for lower data rate is a simple subset of the higher data rate (standard) mode
 - Subset PHY implemented by simply turning off elements of higher data rate standard parent PHY
 - A subset PHY is designed specifically to permit rapid speed change
- Subset PHY approach enables PHY-level synchronization
 - Further reduces time to switch over packet based synchronization
- Subset PHY maintains state of “parent” PHY
 - Echo, NEXT, FEXT, equalizer, and sample clock are all kept current without requiring a speed change
- Micro-second order of magnitude data rate transition time technically feasible
 - ~7x or greater power reduction