An Overview of 10GBASE-T

Paul Langner – *Director Systems and Validation*
Bill Woodruff – *VP Marketing*

February 8, 2006

Overview

- 10GBASE-T Overview
- Cabling
- 10GBASE-T Standard
- Block Overview
- Typical PHY
- Questions

AQUANTIA

10GBASE-T Challenge

- The challenge is to run 10Gb/s Ethernet over 100m of CAT-6a UTP (Unshielded Twisted Pair) cable, or 55m of CAT-6e UTP cable
  - Uses regular RJ-45 type connectors

Wait a minute… CAT-6a??

- It turned out that running 10Gb/s over 100m of CAT-6 was impossible, so a new cable standard, called CAT-6a was created
- There are now 3 kinds of CAT-6 cable:
  - CAT-6: original CAT-6 cable specified to 250 MHz
  - CAT-6e: extended characterization of CAT-6 cable up to 500 MHz
  - CAT-6a: new CAT-6 cable for 10GBASE-T
    - Defined up to 625 MHz

AQUANTIA

Overview

- 10GBASE-T Overview
- Cabling
- 10GBASE-T Standard
- Block Overview
- Typical PHY
- Questions

AQUANTIA

Impairments Inside The Cable

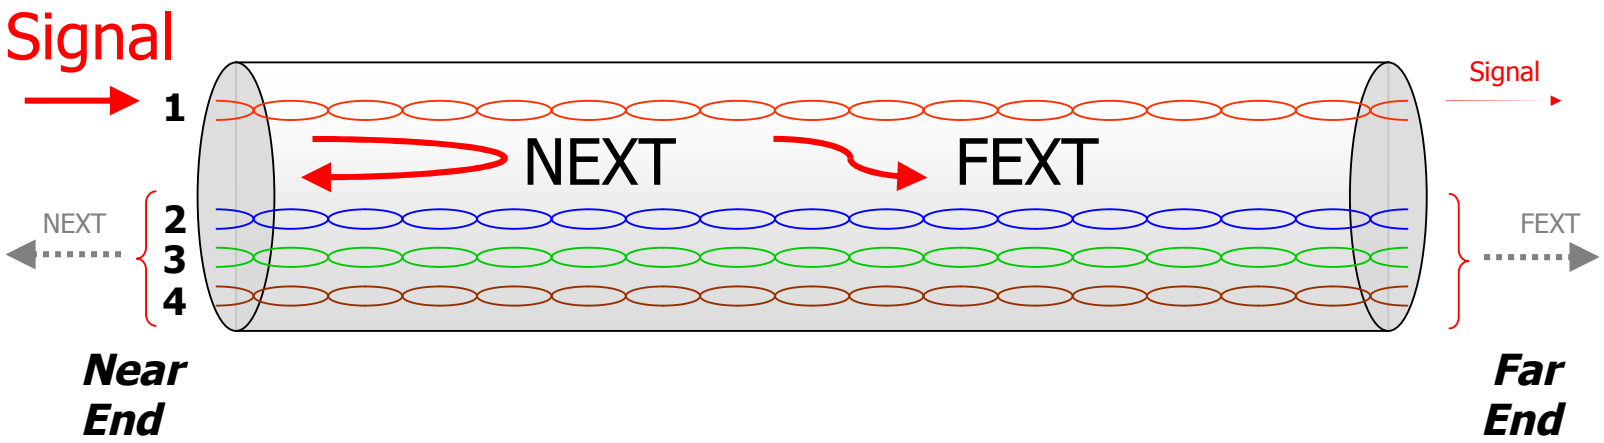

- Insertion Loss (IL) is the loss the signal experiences traveling down a pair in the cable
- Near End Cross Talk (NEXT) is what comes out on the near end of the other pairs, when you inject a signal on #1
- Far End Cross Talk (FEXT) is what comes out on the far end of the other pairs, when you inject a signal on #1

Impairments Outside The Cable

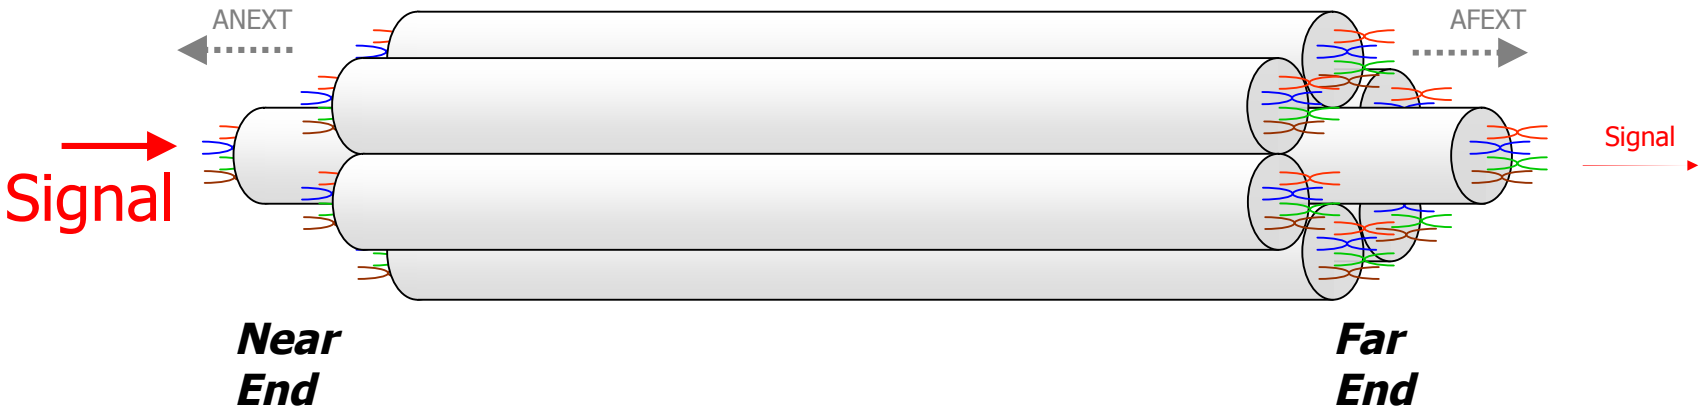

- Alien NEXT (ANEXT) is what comes out on the near end of the other pairs, when you inject a signal on a pair in the center cable
- Alien FEXT (AFEXT) is what comes out on the far end of the other lines, when you inject a signal on a pair in the center cable

Power Sum Interference

- All of these impairments add up in a power sum fashion (versus a voltage sum) to create a very challenging, noisy environment
  - 6-around-1 is "worst case" configuration
- While the A-NEXT and A-FEXT signals may resemble the NEXT and FEXT signals, they are more difficult to deal with, as the receiver doesn't have access to a clean version of the interference signal to use for equalization

CAT-6a

- To create a UTP channel capable of providing greater than 10 Gb/s of information BW, cable manufacturers did the following for CAT-6a:
  - Increased the twist
  - Varied the twist rates between the four pairs, so that coupling is controlled
  - Increased the diameter of the cable (0.31” vs. .22”)
  - Installed a separator for controlling the pair positions within the cable
  - Specs up to 625 MHz for all impairments

* Systimax Cable Cross-section

Overview

- 10GBASE-T Overview
- Cabling
- 10GBASE-T Standard
- Block Overview
- Typical PHY
- Questions

AQUANTIA

10GBASE-T Standard

- Despite these dramatic improvements in the UTP cable, every "trick in the book" of signal processing needs to be employed to achieve 10 Gb/s over 100m of CAT-6a cable
  - Still very close to the Shannon limit
- To meet this challenge, the 10GBASE-T standard (IEEE 802.3an) utilizes:
  - Self-synchronizing scrambler
  - 128-DSQ (Double Square) coset-partitioned constellation
  - LDPC (Low Density Parity Check) block codes
  - Tomlinson-Harashima Precoding (THP)

AQUANTIA

10GBASE-T Standard (continued)

- For the datapath, the 10GBASE-T standard meticulously defines what the transmitter has to do

- The receiver implementation is left almost entirely up to the individual chip manufacturer

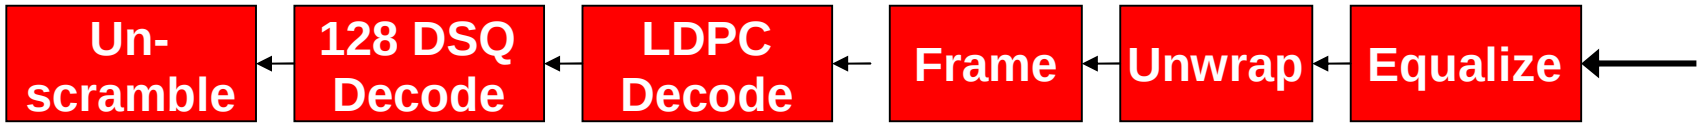

  - A typical line up for the receiver is shown above

Overview

- 10GBASE-T Overview
- Cabling
- 10GBASE-T Standard
- **Block Overview**
- Typical PHY
- Questions

AQUANTIA

Self-Sync Scramblers

- Provides clock transitions, and a statistically random power spectrum for EMI control, equalizer convergence, etc.

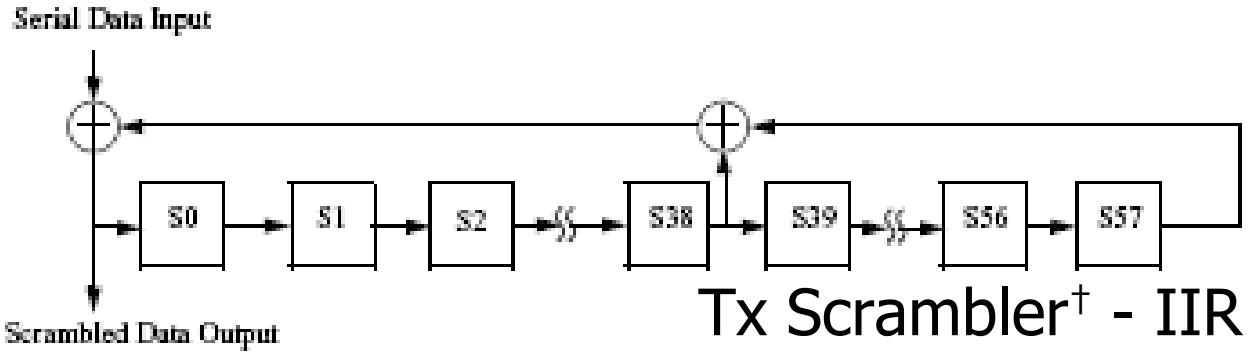

Tx Scrambler† - IIR

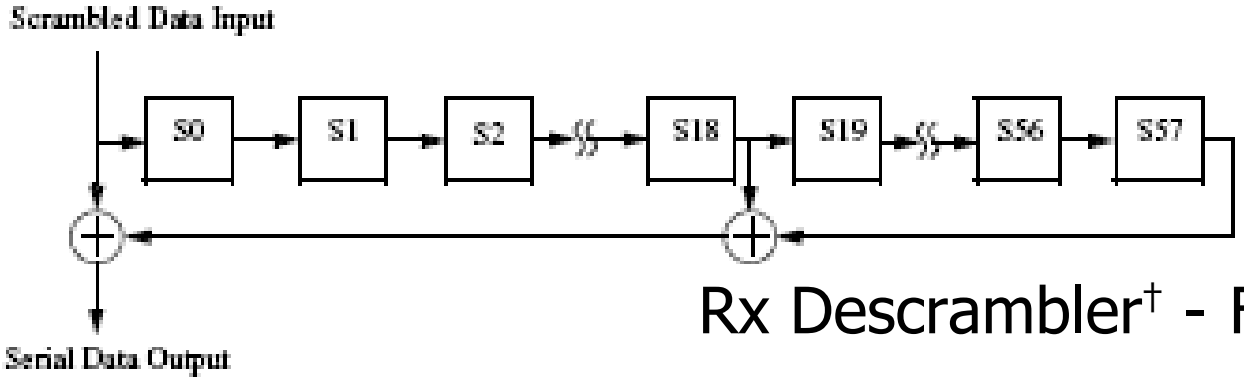

Rx Descrambler† - FIR

†Master shown

Self-Sync Scramblers (continued)

- Self-sync scramblers are the simplest to implement, as they "self-synchronize" and don't require any frame locking or out-of-band sync state communication
  - Equivalent to finite-field pole/zero cancellation
- Downside is that they suffer from:
  - error multiplication (x number of taps)
  - you have to flush out the descrambler on start-up

AQUANTIA

LDPC

- Low Density Parity Check (LDPC) block codes are one of two kinds of error correcting codes that can be used to approach the Shannon capacity of a channel
  - The other code is a Turbo code, but this is heavily patented
- Invented in 1963 by Gallager, but abandoned because the error correcting capability of these codes is not "predictable"
- 10GBASE-T standard calls for a block size of 2048 bits, with 325 check bits: a (2048, 1723) code

AQUANTIA

LDPC Decoder

- LDPC decoding is done in the receiver using a soft-decision, message passing algorithm
  - "Soft" means that the bits are treated as continuous variables which represent the probability of the value being a 0 or a 1
  - Represented in the decoding algorithm as Log Likelihood Ratio (LLRs) – ln (pr1 / pr 0)

AQUANTIA

LDPC Decoder (continued)

- Message passing algorithm:
  - sets the variable nodes to the soft receive LLRs
  - calculating the syndrome LLRs, given the variable LLRs
  - updating the variable LLRs with the new syndrome information
  - Iterate until convergence is reached

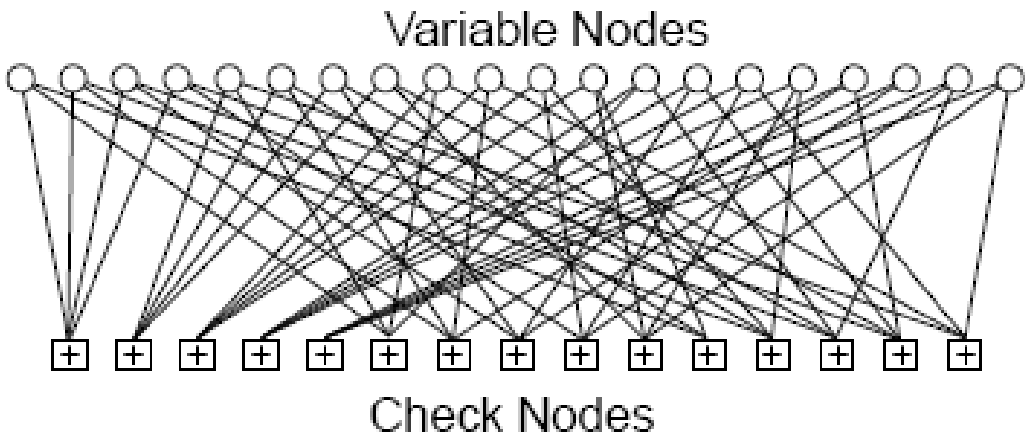

AQUANTIA

128-DSQ

- The 10GBASE-T standard uses a synthetic 2-dimensional 128 DSQ (Double SQuare) constellation, which conveys 7 bits per symbol
  - ($2^7 = 128$)
- 128 DSQ resembles a 16x16 checker-board with the white squares missing

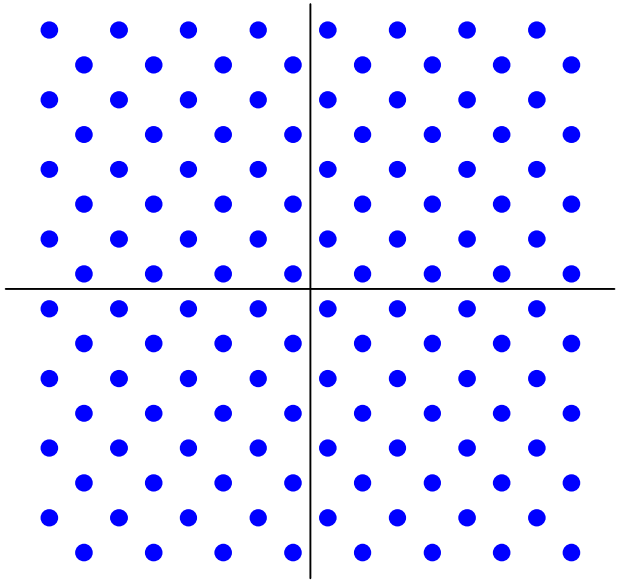

128-DSQ (continued)

- Constellation is partitioned into 8 cosets (diamond regions containing 16 points)
  - 3 bits uncoded bits are conveyed by coset, and 4 coded bits by the points in the coset
  - Effectively spaces the cosets far apart, relative to the points within them

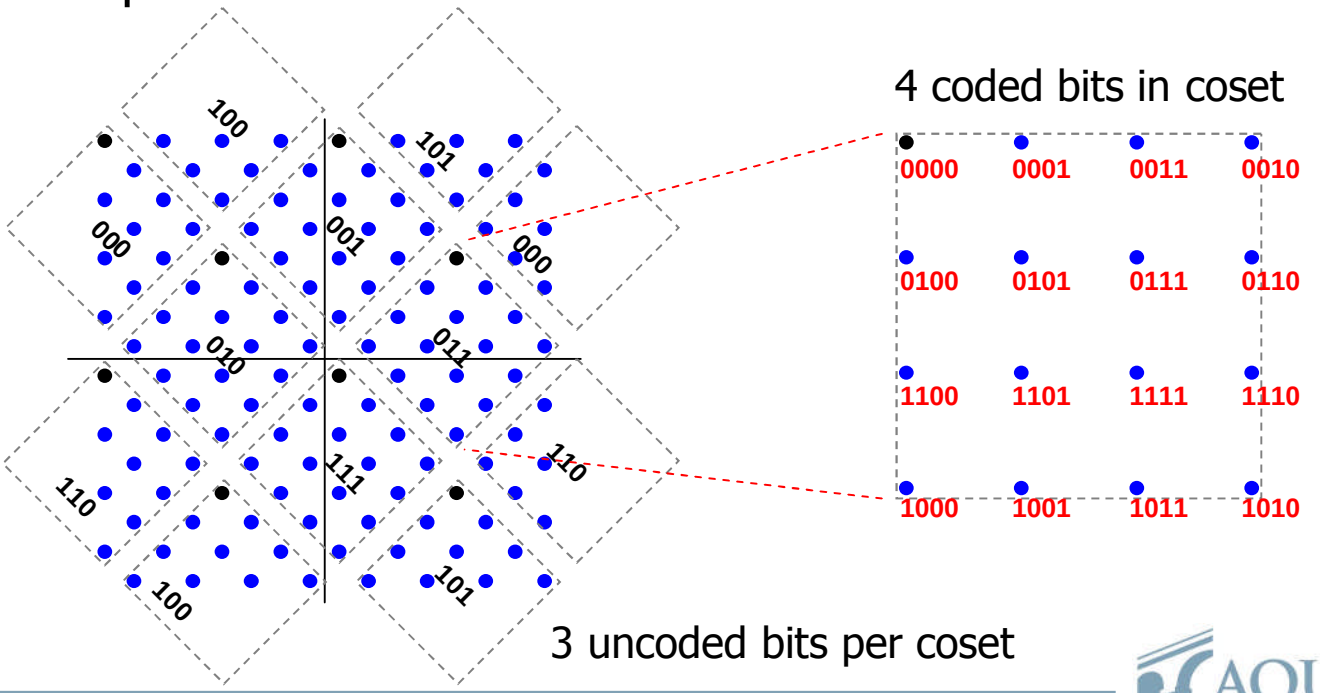

128-DSQ (continued)

- As you may have guessed the 4 coded bits are protected using the LDPC block code
- Each 128-DSQ symbol is transmitted using two, back-to-back PAM-16 (16 level Pulse Amplitude Modulation) signals
  - running at 800 MS/s (400 DSQ MS/s)
  - Spread over all 4 pairs
- Raw bit rate is 11.2 Gb/s

AQUANTIA

THP

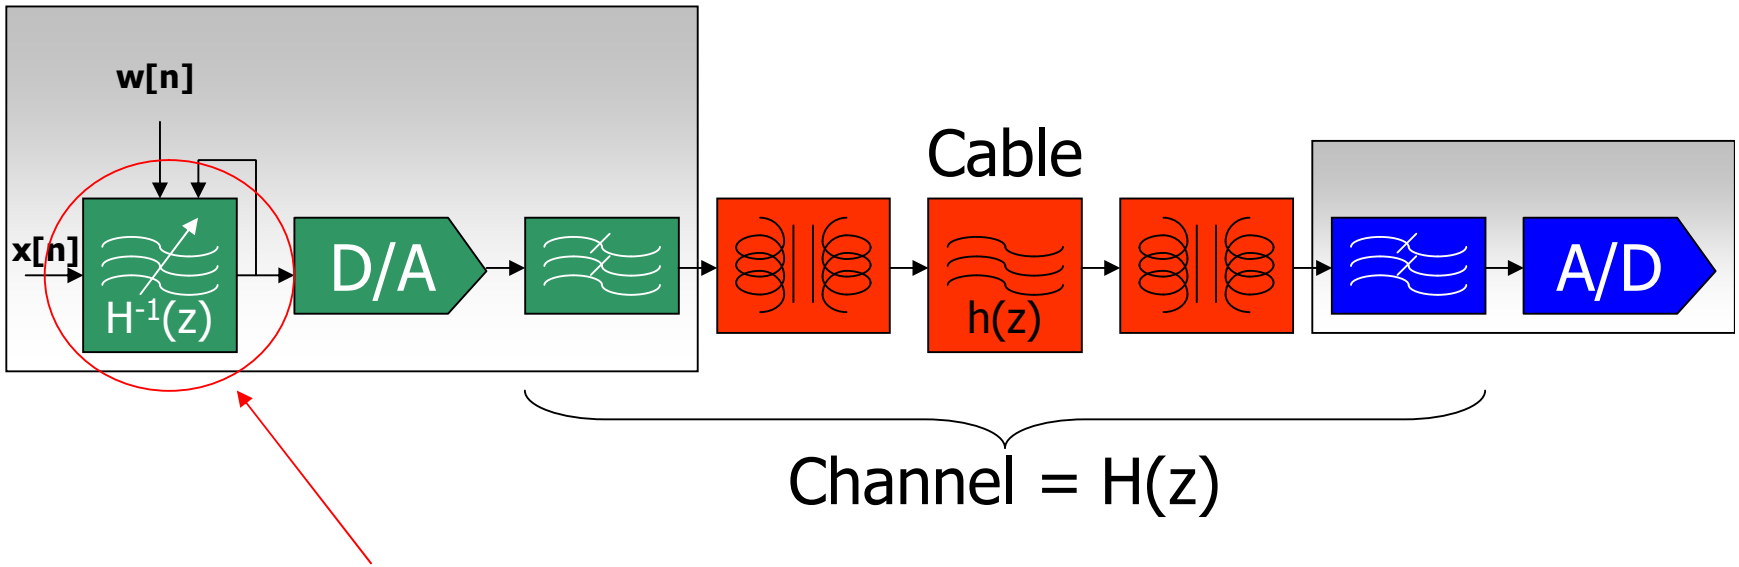

- The 10GBASE-T standard calls for the use of Tomlinson-Harashima Precoding (THP), which is a scheme in which the equalizer for the channel is placed in the transmitter
  - theoretically allows the receiver to see "perfect" symbols
- Training is accomplished during the initialization phase of the link, and then is fixed

THP (continued)

- In order to control the transmit power levels, a "dither" sequence is added into the precoder
  - THP checks what the next output is going to be
  - If it is too large, it picks an input level that is "modulo" the D/A rails in an attempt to reduce the dynamic range
    - This is called "wrapping" and is like modulo addition
- Net effect of this is that the dynamic range on the receiver is now larger, by multiples of the rail voltage
  - Typically 2-3 is the maximum over-range used, as a balance has to be struck between the receiver dynamic range, and the transmitter
- At the receiver, the received level is unwrapped via a modulo operation

10GBASE-T Frame

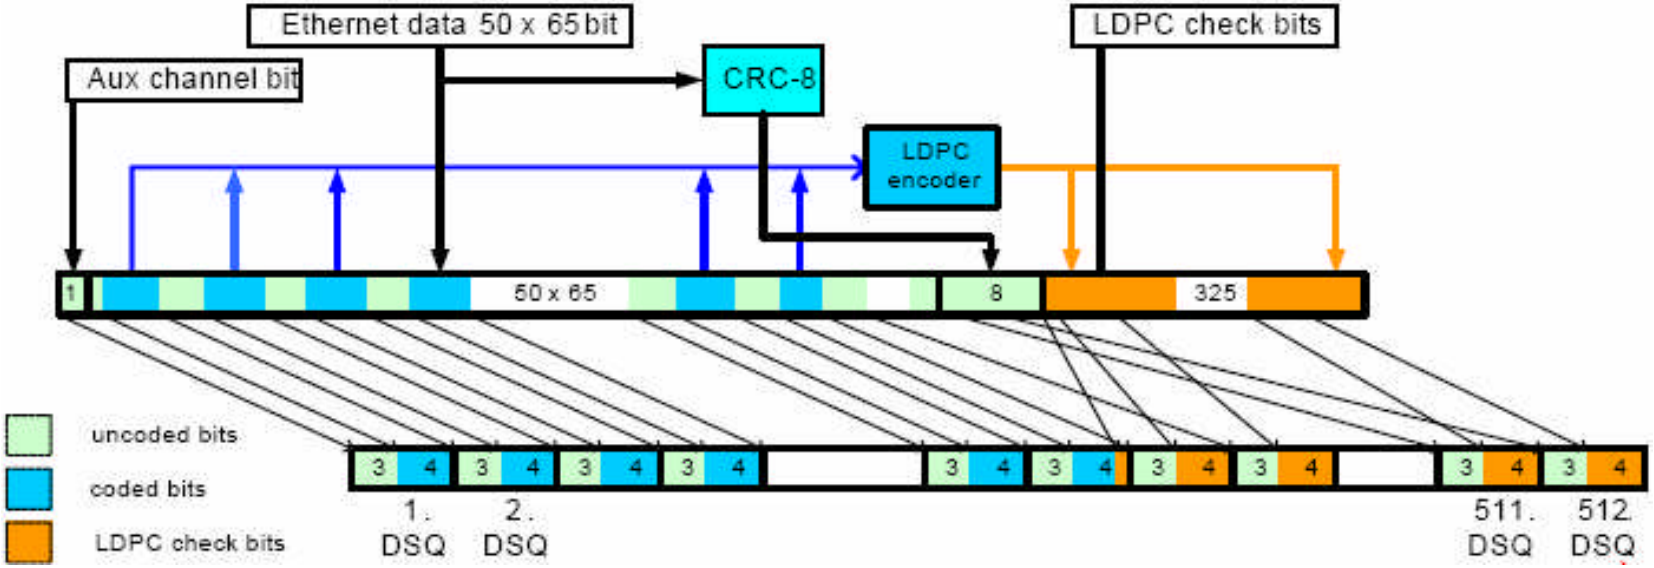

- 10GBASE-T frame collects 50x 65-bit blocks from the XGMII interface and scrambles them, adds an auxillary channel bit, and a CRC-8 for a total of 3259 bits
- Payload is divided up into 1723 bits for the LDPC coder, and 3x 512 uncoded bits (1536 bits)
  - LDPC block size after coding is 4x 512 bits (2048)
- These 512 symbols are then sent as 1024 PAM-16 symbols over the 4 pairs (3584 bits)

Overview

- 10GBASE-T Overview
- 10GBASE-T Standard
- Block Overview
- Typical PHY
- Questions

AQUANTIA

Simplified PHY Block Diagram

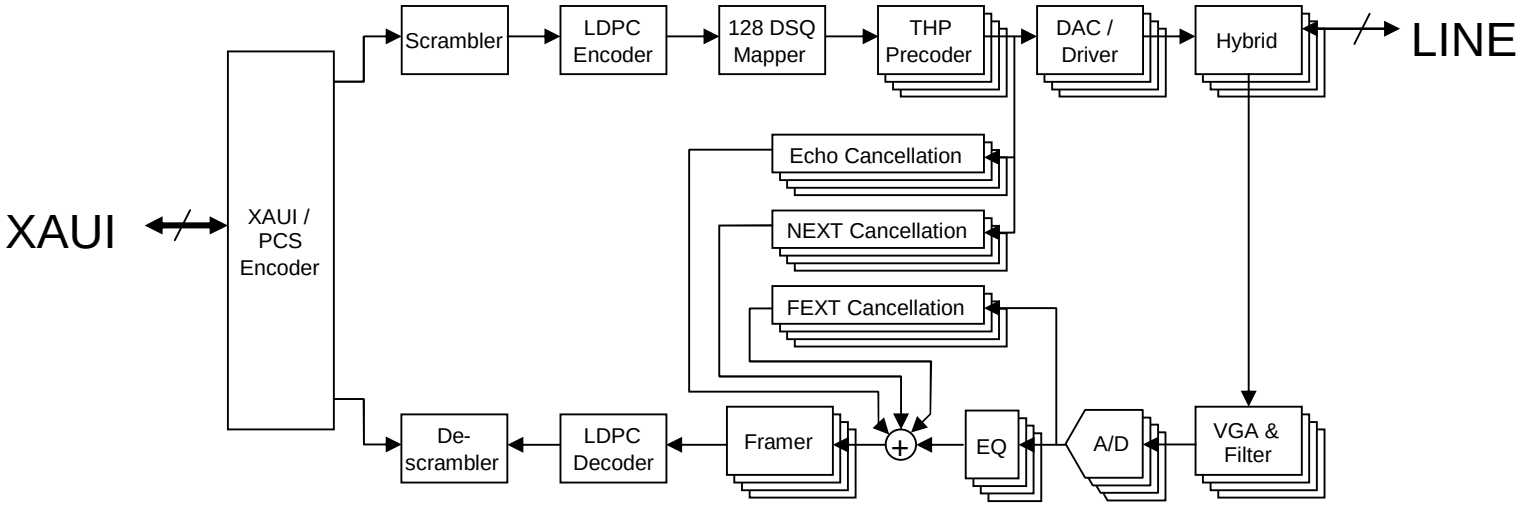

Overview

- 10GBASE-T Overview
- 10GBASE-T Standard
- Block Overview
- Typical PHY
- Questions

AQUANTIA