

802.3ab

# A Tutorial Presentation

# 1000BASE-T Tutorial Structure

---

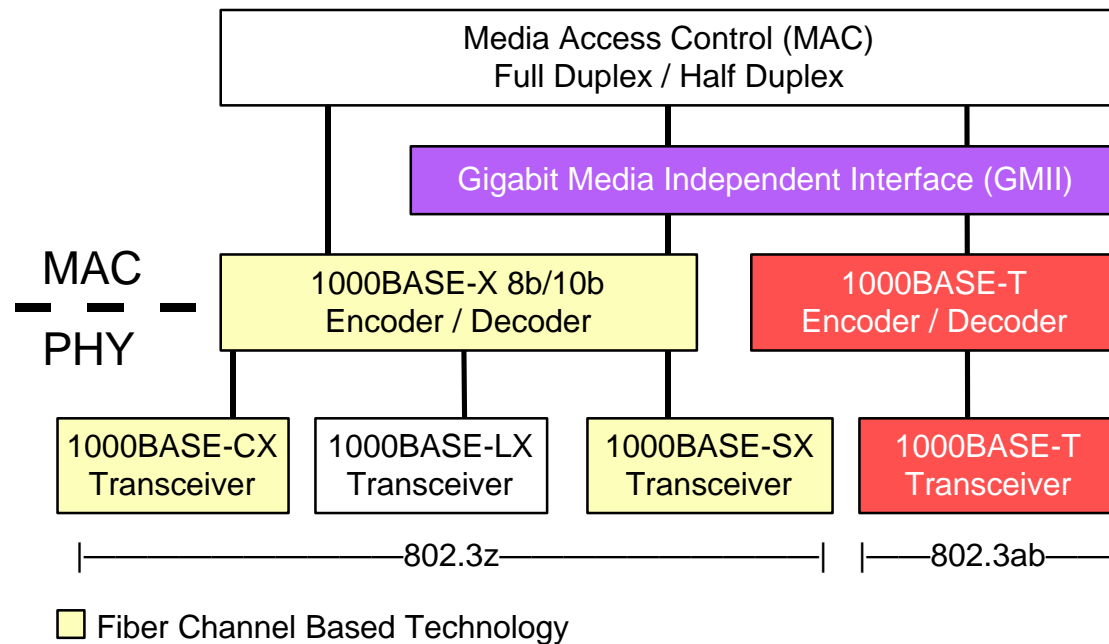
- Introduction, Market & History
  - Colin Mick, The Mick Group
- Cabling
  - Chris DiMinico, Cabletron
- Channel and Overall Architecture
  - Sreen Raghavan, ComCore Semiconductor
- Technical Details
  - Sailesh Rao, Level One Communications
- Detailed VLSI Implementation
  - Mehdi Hatamian, Broadcom Corporation

# One New Design Task for 1000BASE-T

---

- 5 Level Signaling costs 6dB in SNR
  - Get back with Forward Error Correction (FEC)
- Everything else has been done in:
  - 100Base-TX
    - 125Msps
    - DSP works
  - 100Base-T4
    - Transmit/Receive on 4 channels
  - 100Base-T2
    - ECHO and NEXT canceling


# Where Does 802.3ab Fit?



# Target Market

---

- 70% of installed UTP is CAT 5
- CAT 5 installed footage is growing 30% annually



# Market Applications

---

- Server Farms
- High performance work groups
- Graphic-based applications
- Network computers
- Shared gigabit networks

# Objectives

---

- Comply with specifications for GMII of 802.3z.
- Provide line transmission which supports full and half duplex operation.
- Provide FCC Class A/CISPR or better operation
- Support operation over 100 meters of Category 5 balanced cabling
- Achieve bit Error Rate better than to  $10^{-10}$
- Support Auto-Negotiation (Clause 28)
- Meet susceptibility requirements
- Support the objectives of 802.3z of Nov. 13, 1996

# Foundations

---

- 100BASE-TX demonstrates sending a 3-level symbol stream over Category 5 cable at 125 Mbaud is possible and practical.
  - 100BASE-TX DSP Based Phys now available
- 100BASE-T4 demonstrates techniques for sending multi-level coded symbols over four pairs.
- 100BASE-T2 demonstrates the use of digital signal processing (DSP), five-level coding, and simultaneous two-way data streams while dealing with alien signals in adjacent pairs

# Timeline

---

## **Milestone**

## **802.3z**

## **802.3ab**

Work starts

November 95

November 95

PAR received

March 96

March 97

1<sup>st</sup> draft

January 97

November 97

WG ballot

July 97

March 98\*

LMSC ballot

November 97

July 98\*

Ready for std

June 98\*

December 98\*

# Status

---

- Now reviewing D2-0
- Task force doing final tuning on D2-0 during this Plenary.
- On-track to request permission to go to working group ballot at this Plenary

# Cabling

---

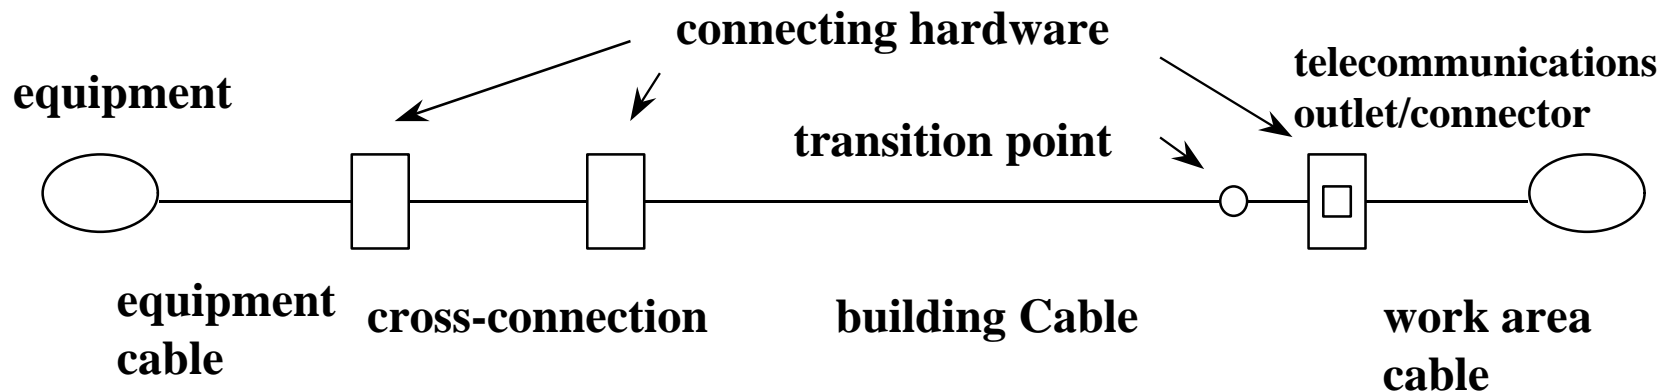
- Topology
- Cabling Specifications
- Cabling standards

# Channel Cabling Topology

---

TIA/EIA - 568-A Channel (w/o transition Point) = ISO 11801 Channel

## Cross-Connect Topology

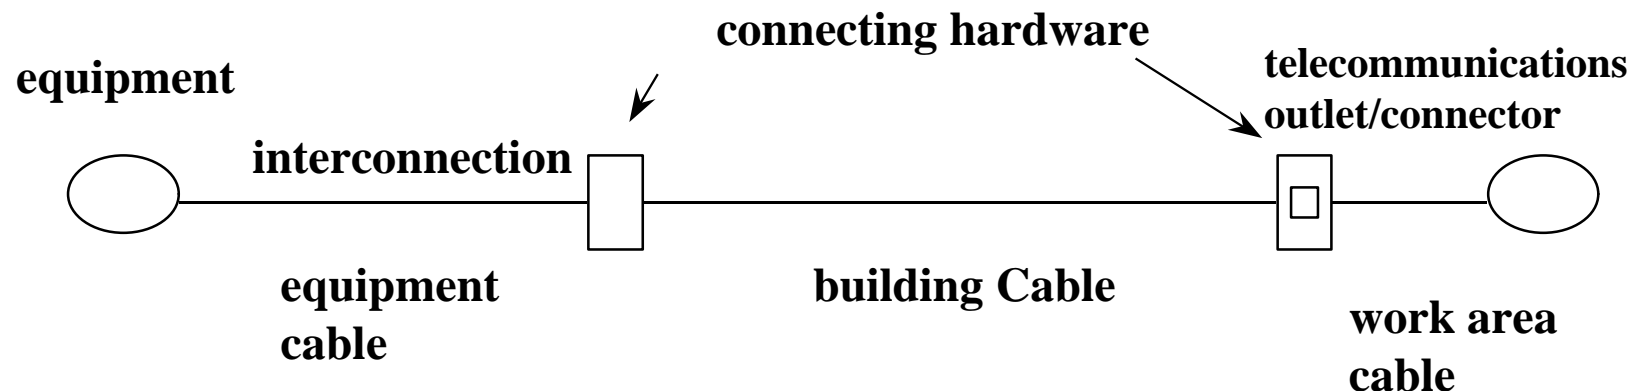


# Installed Cabling

---

## Recommended Interconnect Topology

- Minimizes connections
- Minimizes crosstalk, both near-end and far-end
- Minimizes return loss and insertion loss



# Cabling Performance Specifications

---

## **Based on Category 5 Installed Cabling**

- Project Authorization Request (PAR) for the 1000BASE-T project specifically requires operation on four pair 100 ohm Category 5 balanced copper cabling as defined by TIA/EIA-568-A, or its equivalent as built from material specified by ISO/IEC 11801: 1995 that meets the channel performance parameters specified in TIA/EIA-568-A ANNEX E.

## **Additional cabling specifications: Draft Addendum to ANSI/TIA/EIA-568-A:**

- Additional specifications for FEXT (ELFEXT) and return loss, not currently specified by TIA/EIA-568-A, are being developed developed to characterize the vast majority of the installed base of Category 5 cabling built to TIA/EIA-568-A and/or ISO/IEC 11801 clauses 6, 8 & 9.
- An addendum to ANSI/TIA/EIA-568-A will include the additional transmission performance specifications and field test parameters for FEXT (ELFEXT) and return loss.

# Field Testing

---

A verification of the installed cabling performance per the field test specification of ANSI/TIA/EIA-TSB-67- "Transmission Performance Specifications for Field Testing of Twisted Pair Cabling System" with the additional test parameters for FEXT (ELFEXT) and return loss (to be released as an addendum to TIA/EIA-568-A) will be recommended.

# Cabling Specifications

---

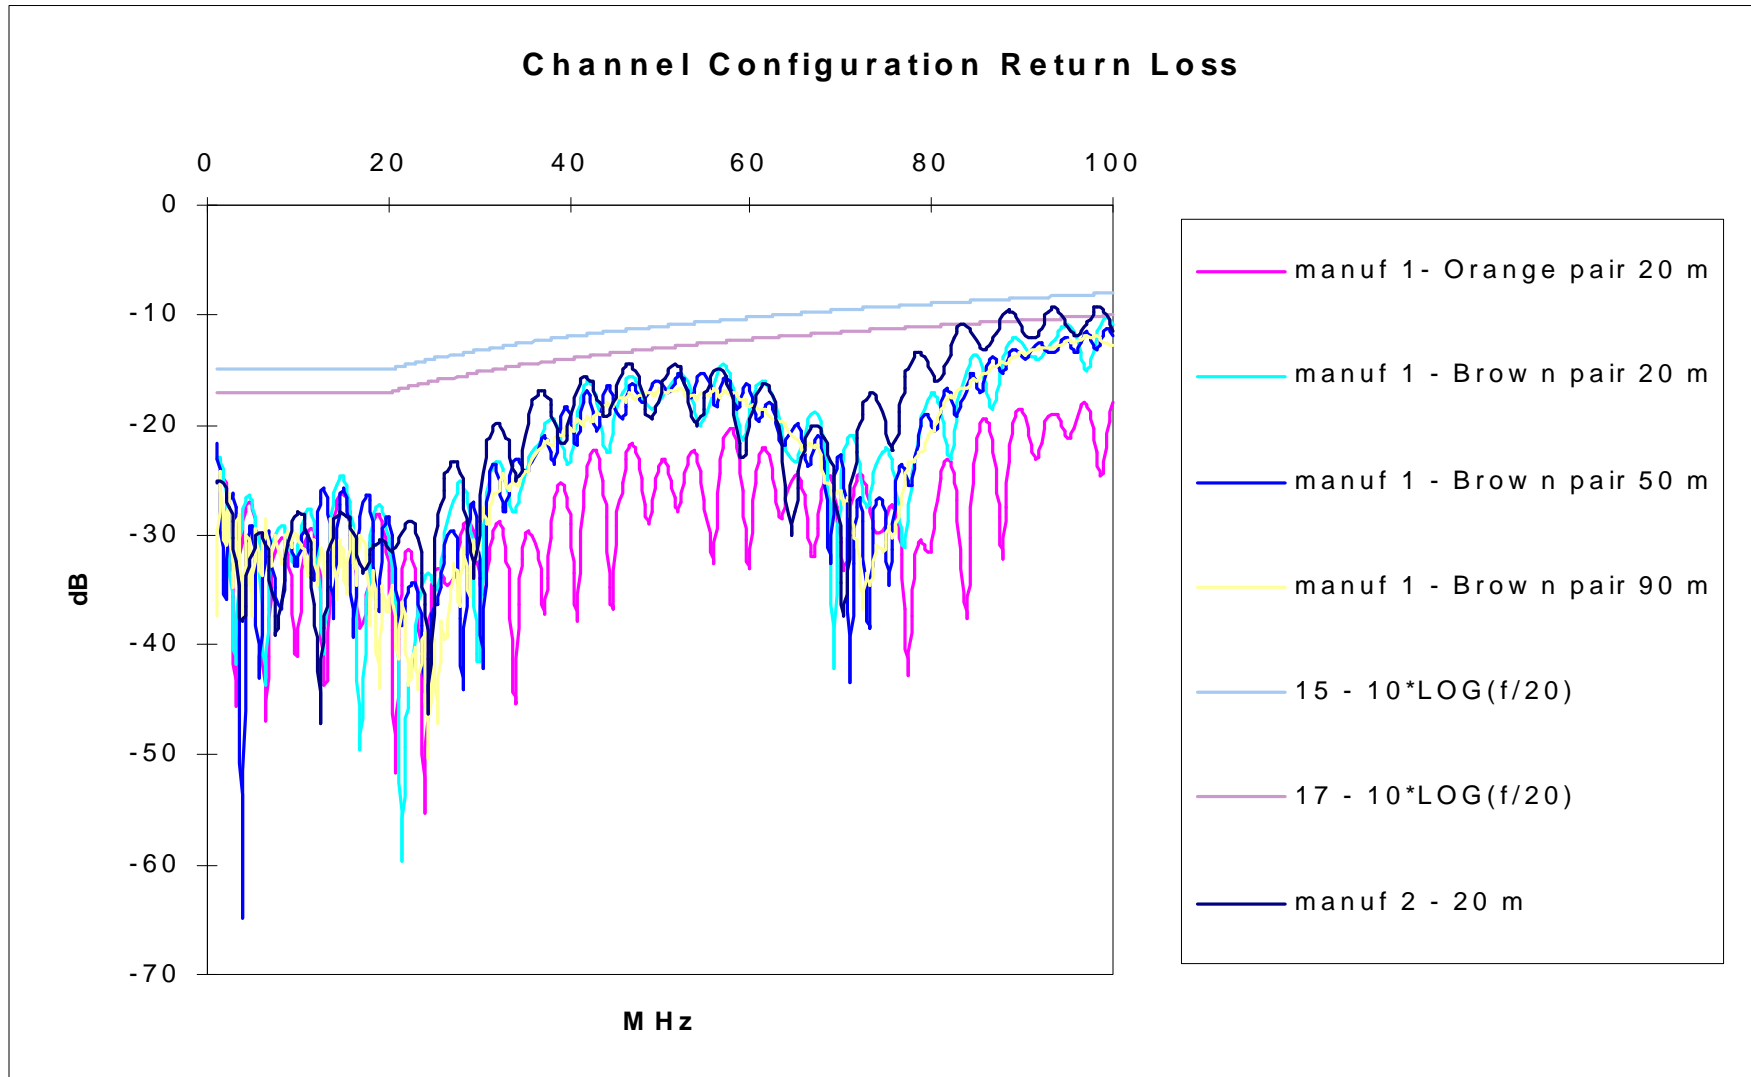
- 40.8.2.3 Return Loss

Frequency(MHz)	Return Loss(dB)
$1 \leq f < 20$	15
$20 \leq f < 100$	$15 - 10 * \text{LOG}(f/20)$

- 40.8.2.1 Insertion Loss

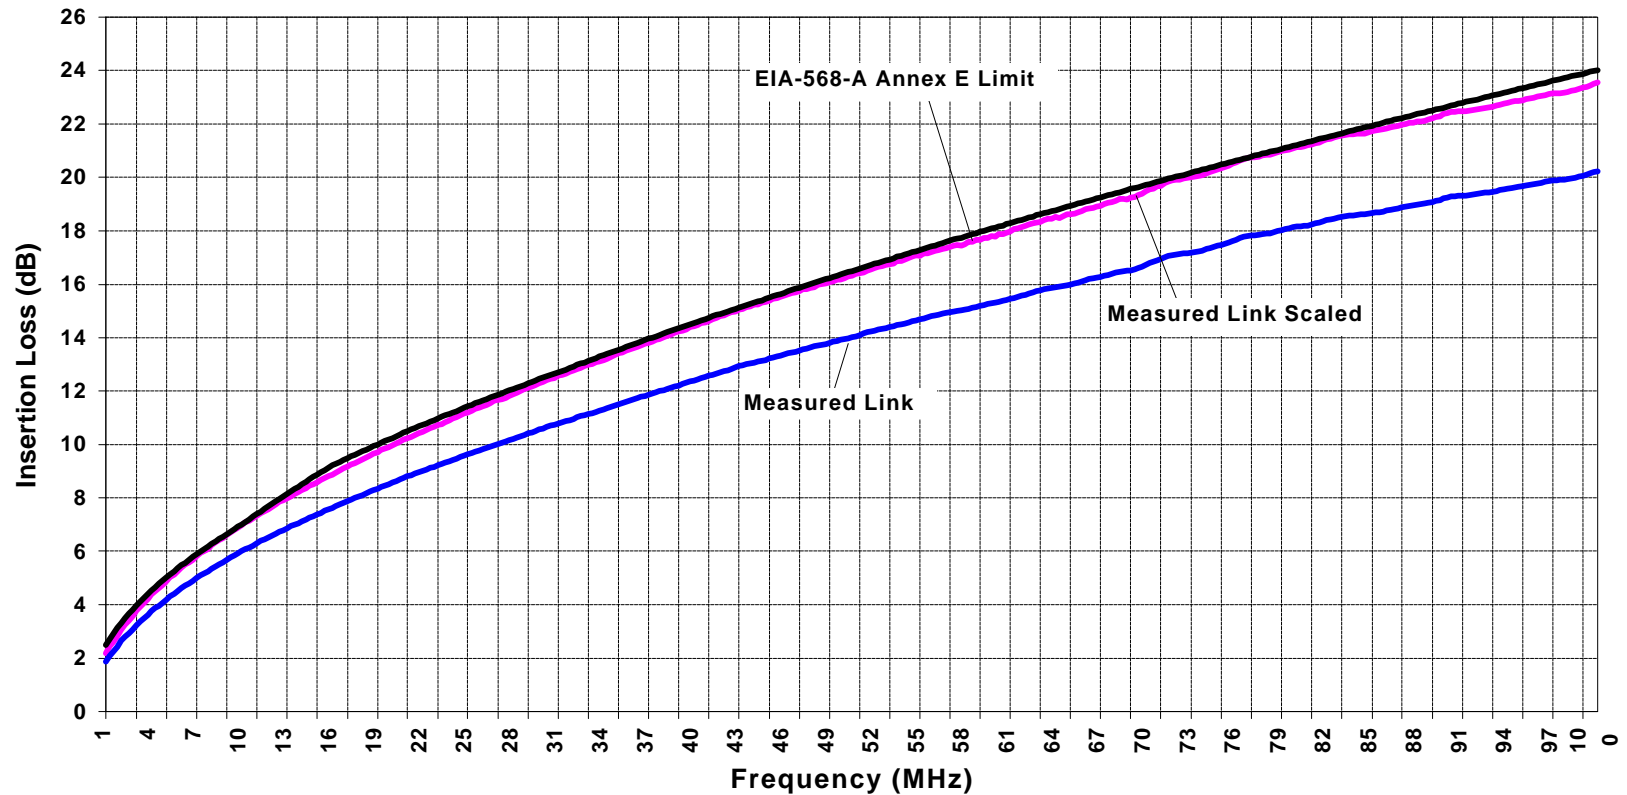
$$\text{Insertion\_loss}(f) < 2.1 * f^{0.529} + .4/f$$

# Channel Return Loss



# Channel Insertion Loss

Insertion Loss Model - Category 5 100 Meter Measured Link,  
Measured Link Scaled & EIA-568-A Annex E Insertion Loss Limit

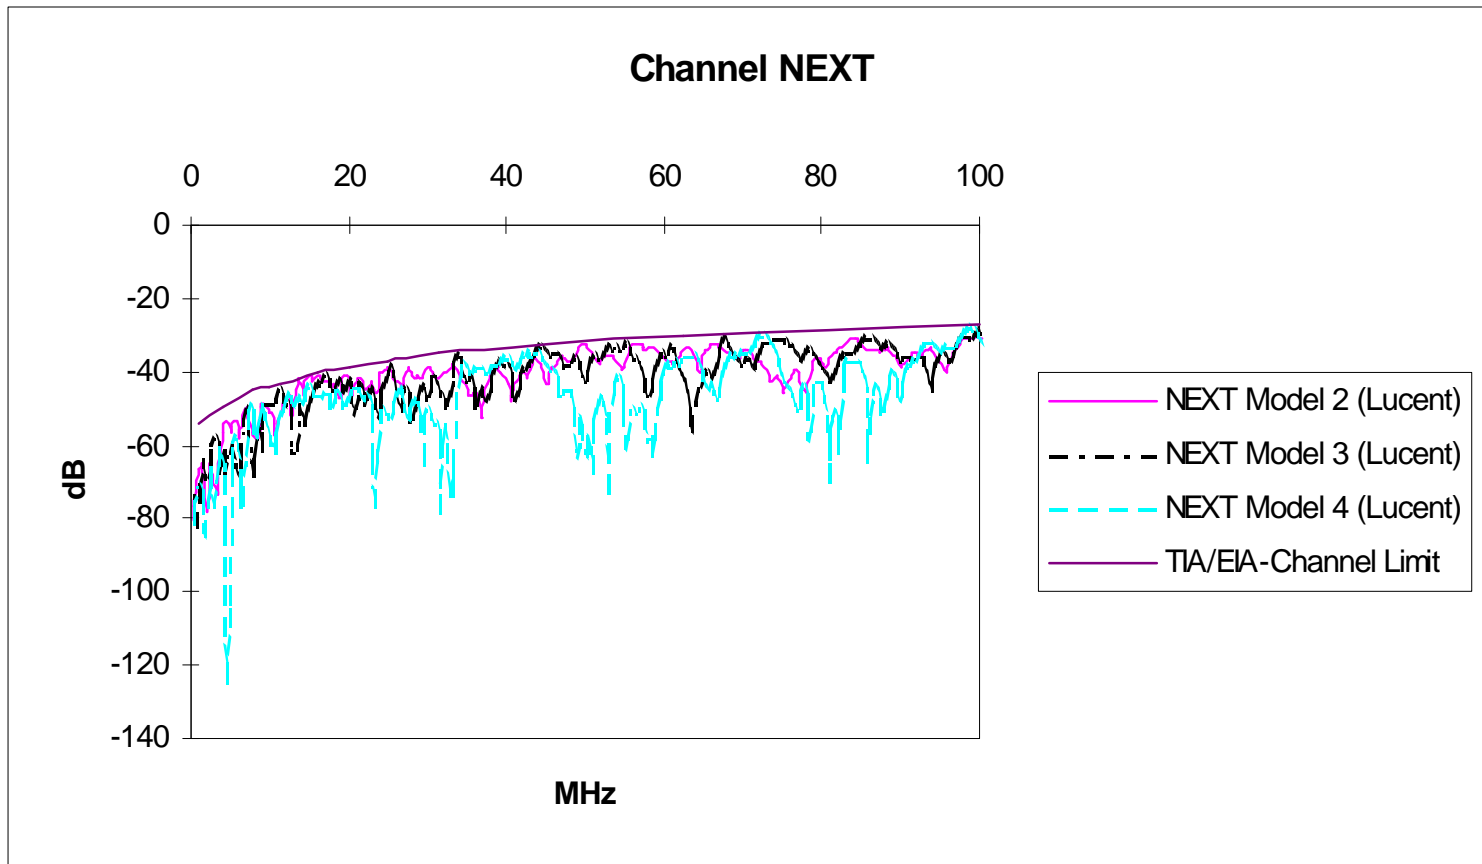


# Coupling Parameters (Crosstalk)

---

- 40.8.3.1 Differential Near-End Crosstalk (NEXT) Loss  
The NEXT loss between all duplex channels of a link segment shall be greater than  $27.1 - 16.8\text{Log}_{10}(f/100)$ .
- 40.8.3.2 Equal Level Far-End Crosstalk (ELFEXT) Loss  
The worst pair ELFEXT loss between a duplex channel shall be greater than  $17 - 20\text{Log}_{10}(f/100)$  dB
- 40.8.3.2.1 Multiple Disturber Far-end Crosstalk (ELFEXT) Loss
  1.  $17.0 - 20\text{Log}_{10}(f/100)$  dB.
  2.  $19.5 - 20\text{Log}_{10}(f/100)$  dB.
  3.  $23.0 - 20\text{Log}_{10}(f/100)$  dB.PSELFEXT loss is  $14.4 - 20 * \log(f/100)$

# Channel Next



# Channel and Overall Architecture

---

- Design Tasks
- Design Approach
- Receiver Startup

# Gigabit Over CAT-5 Copper Cable

---

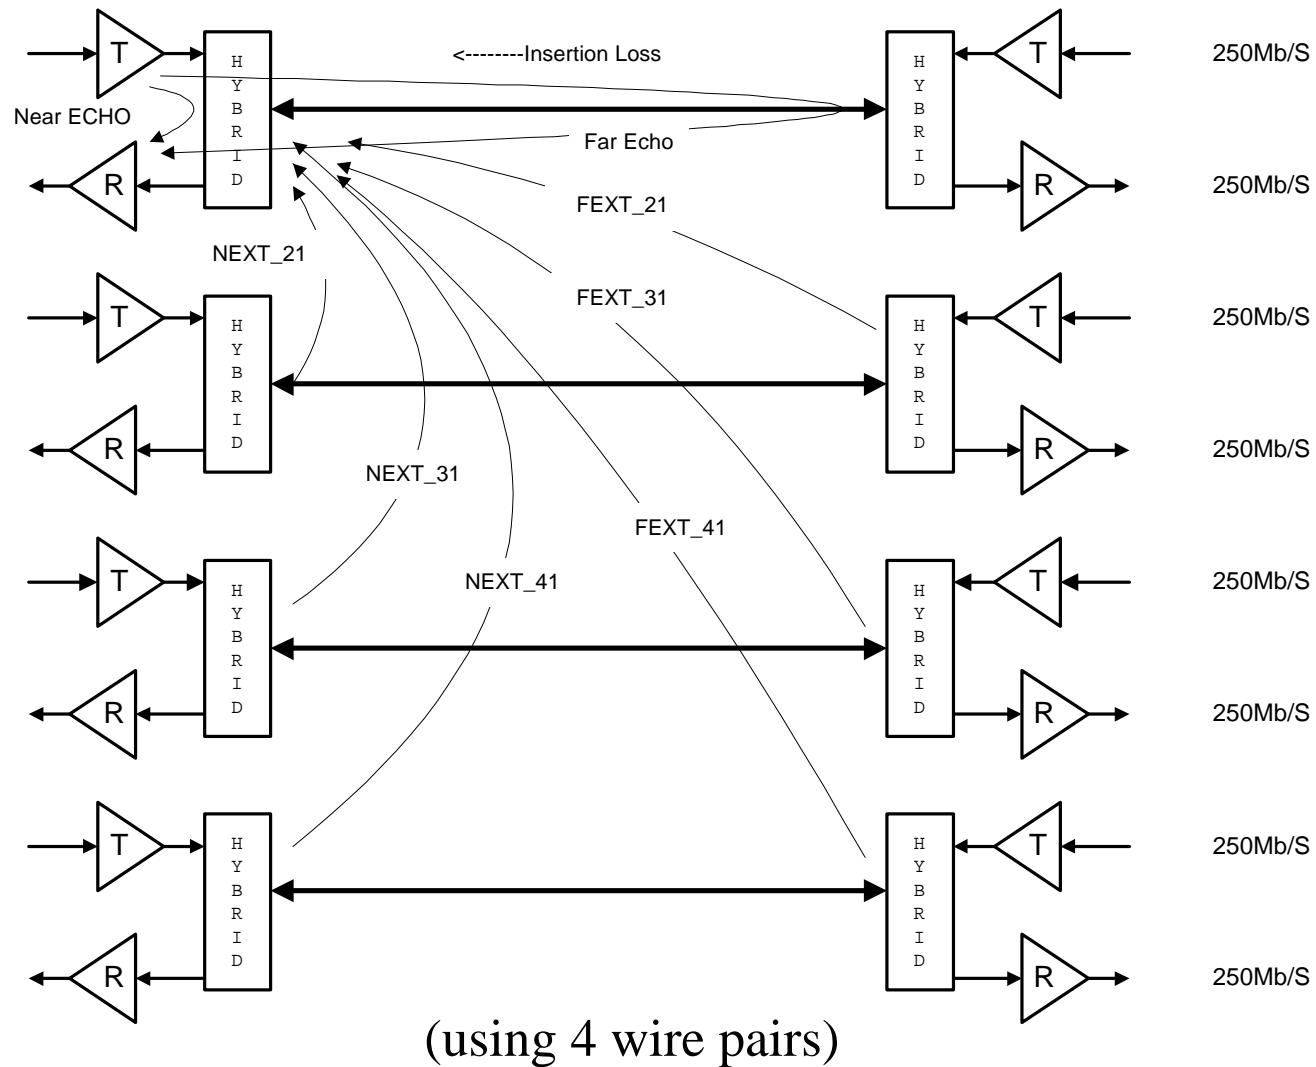
- Topology -- Link segment of a 4-Pair Cat-5 Cable. Each pair a full duplex channel supporting effective data rate of 250Mbps in both directions simultaneously
  - ECHO, SELF NEXT and FEXT
    - Hybrid to cancel most of the NEAR END ECHO
    - Adaptive cancellers to reduce remaining ECHO and SELF NEXT
- Signal to Noise Ratio (SNR) 6dB Less Than 100Base-TX due to 5-level signaling
  - Provide 6dB coding gain in the form of a 4-D Trellis code
- Channel Impairments
  - 20dB signal attenuation at 62MHz at 100 meters
    - Incorporate decision feedback channel equalization
  - External Noise
- FCC imposed limits on transmit levels
  - Limit transmit spectrum above 30MHz
    - Partial Response spectral shaping at the transmitter ( $3/4 + 1/4Z^{-1}$ )

# Major Receiver Design Tasks

---


- A/D conversion
  - 125 Mega samples per second Conversion
- Frequency Locking and Timing Recovery
  - Low Jitter Phase Locked Loop for clock recovery
- FEC – 4D-8 State Trellis Code
  - Design of critical timing path involving Trellis decoder and decision feedback equalizer

# Gigabit Ethernet over CAT-5 Cable




# MLT-3 and PAM-5 Signal Levels

Eye pattern of MLT-3 signaling

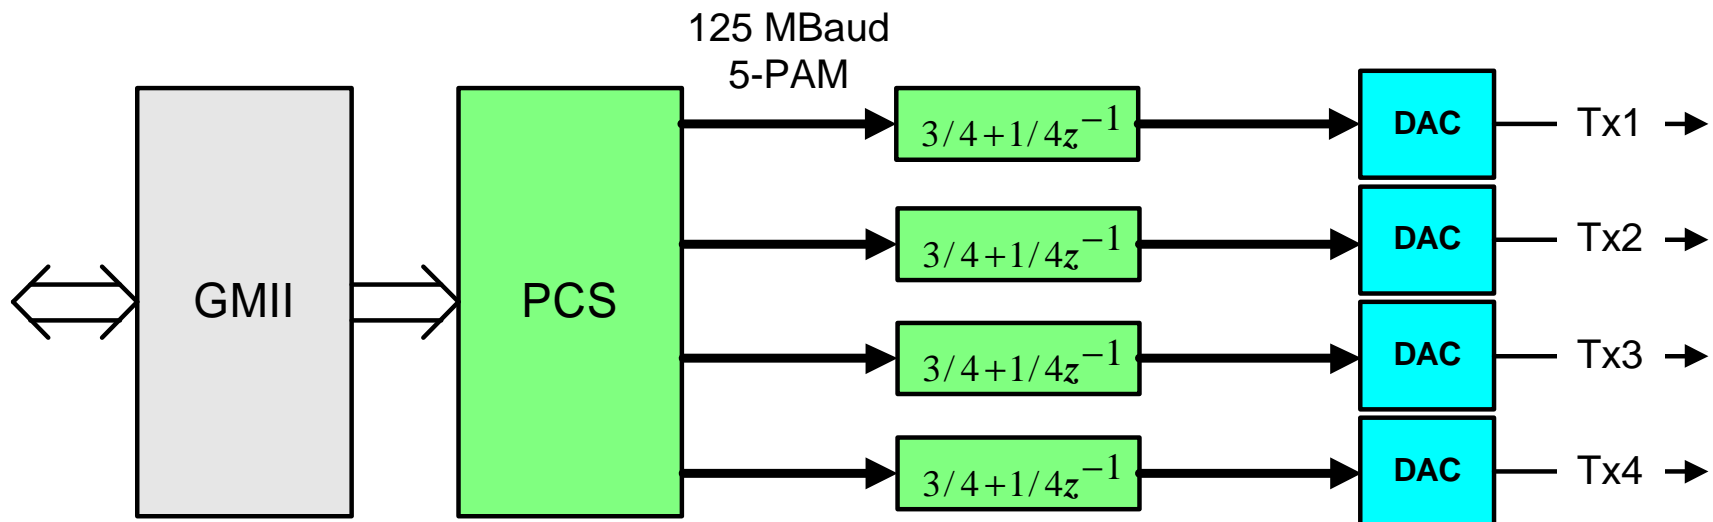


Eye pattern of PAM-5 signaling



Recover the 6dB of SNR -- Add Forward Error Correction (FEC)

# Transmitter

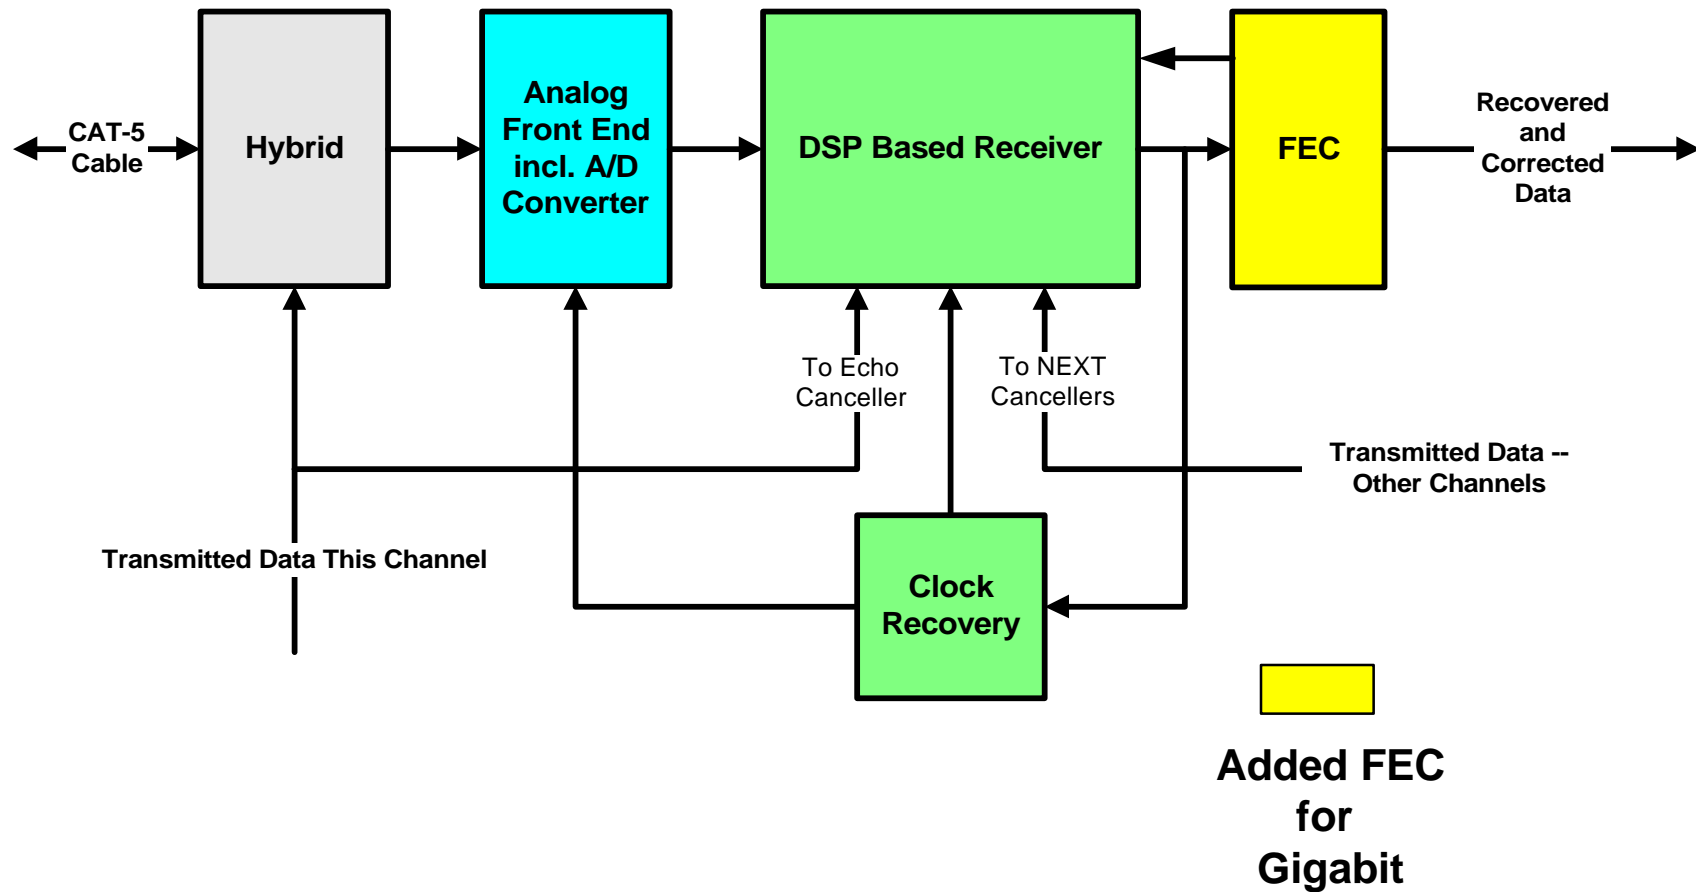


Similar to 100Base-TX

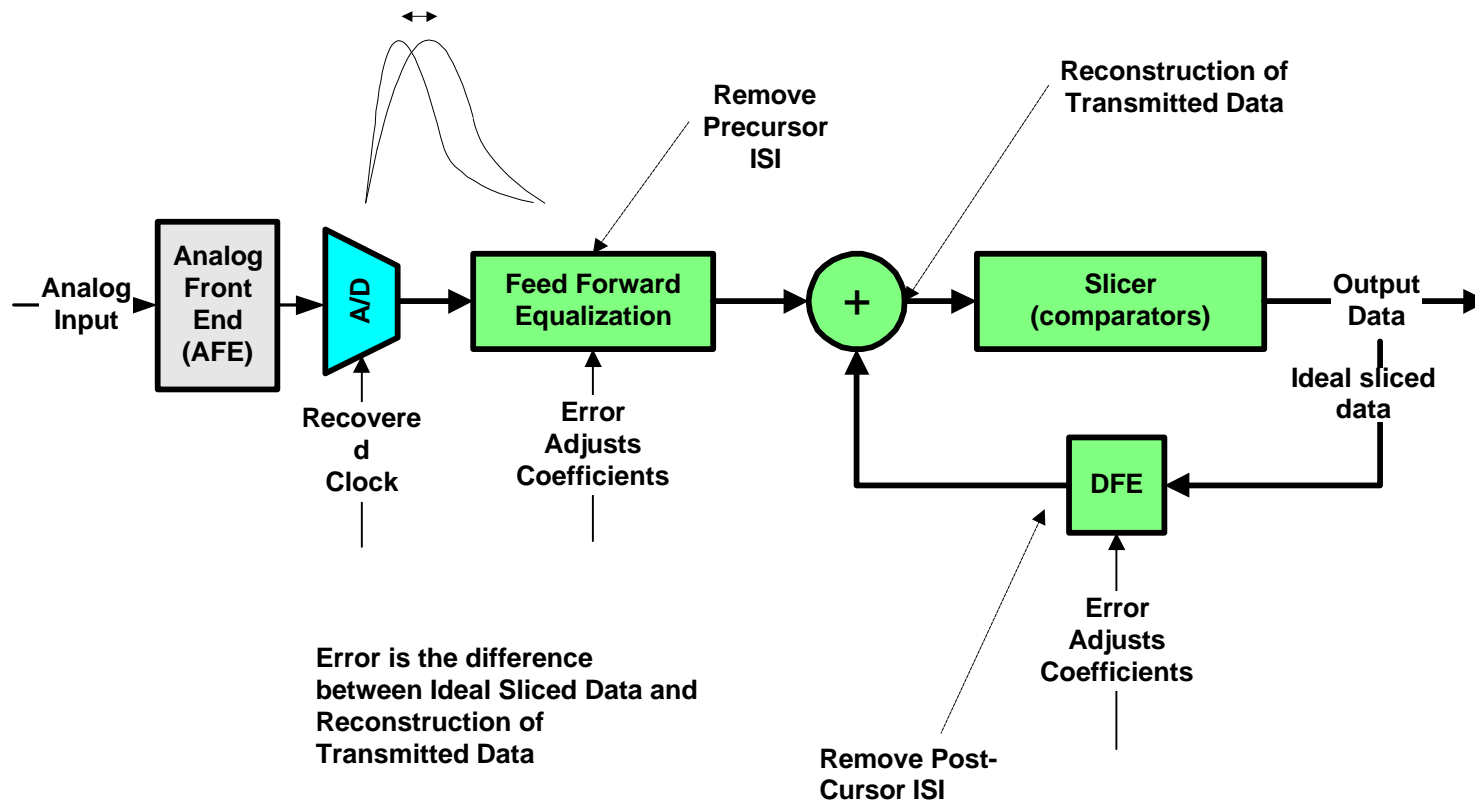
New DSP for Gigabit

New Analog for Gigabit

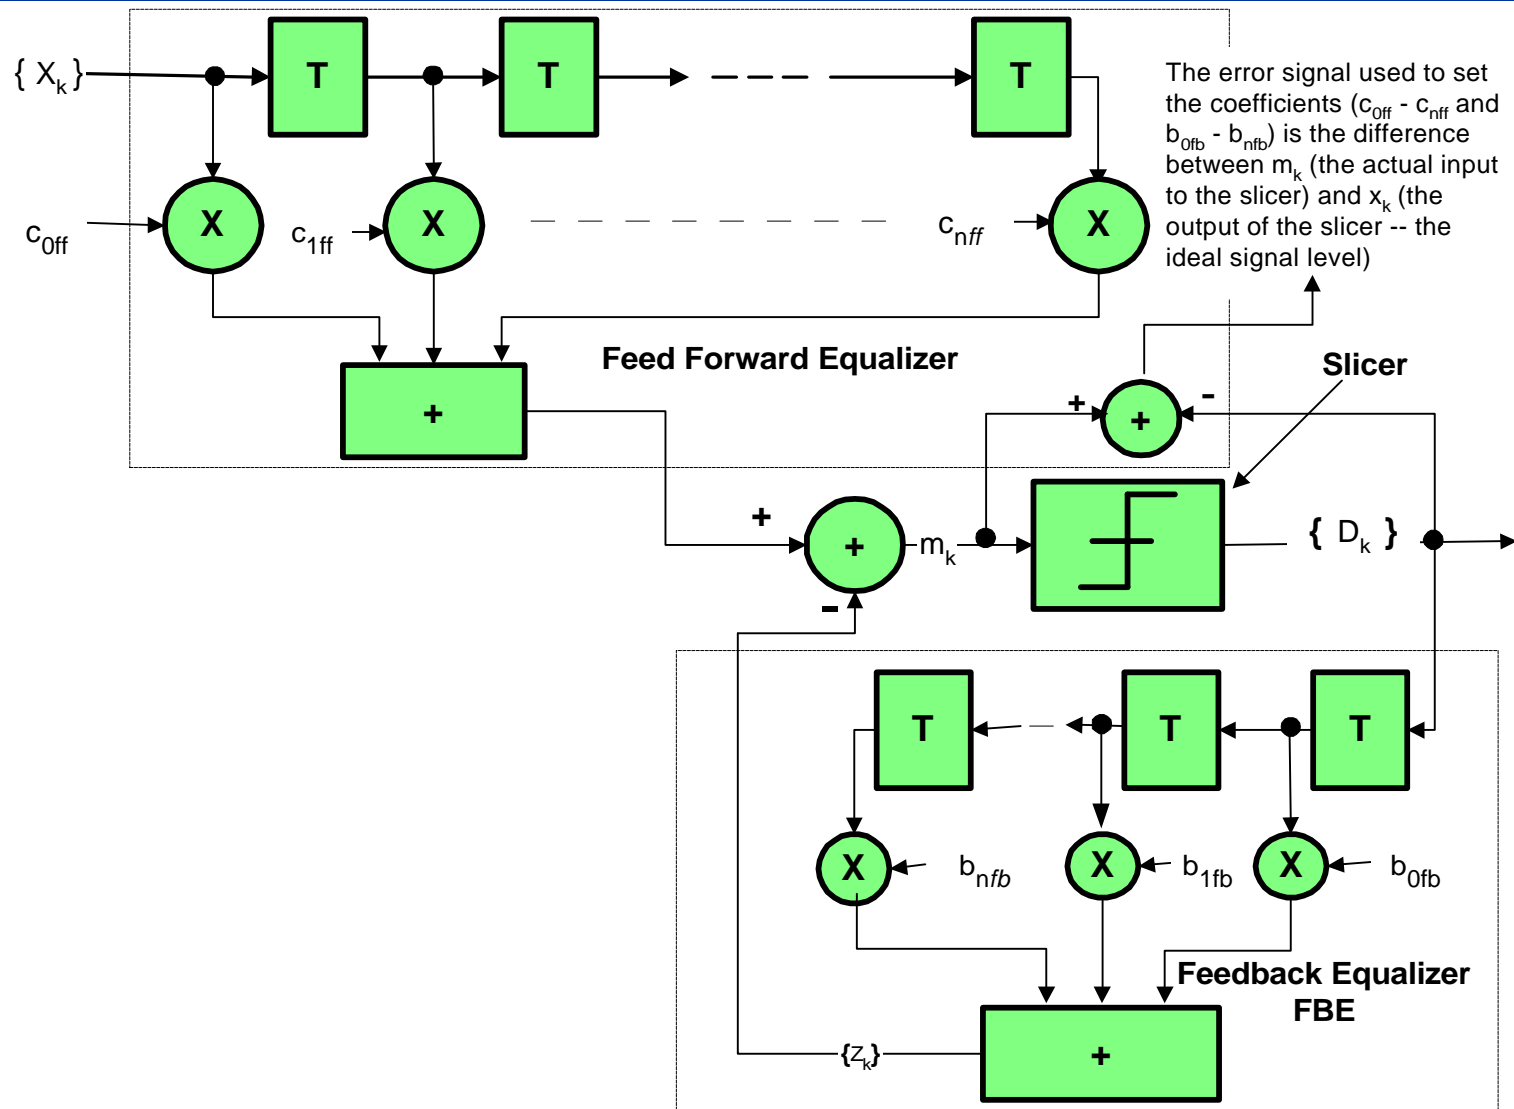
# DSP Based Receiver Block Diagram



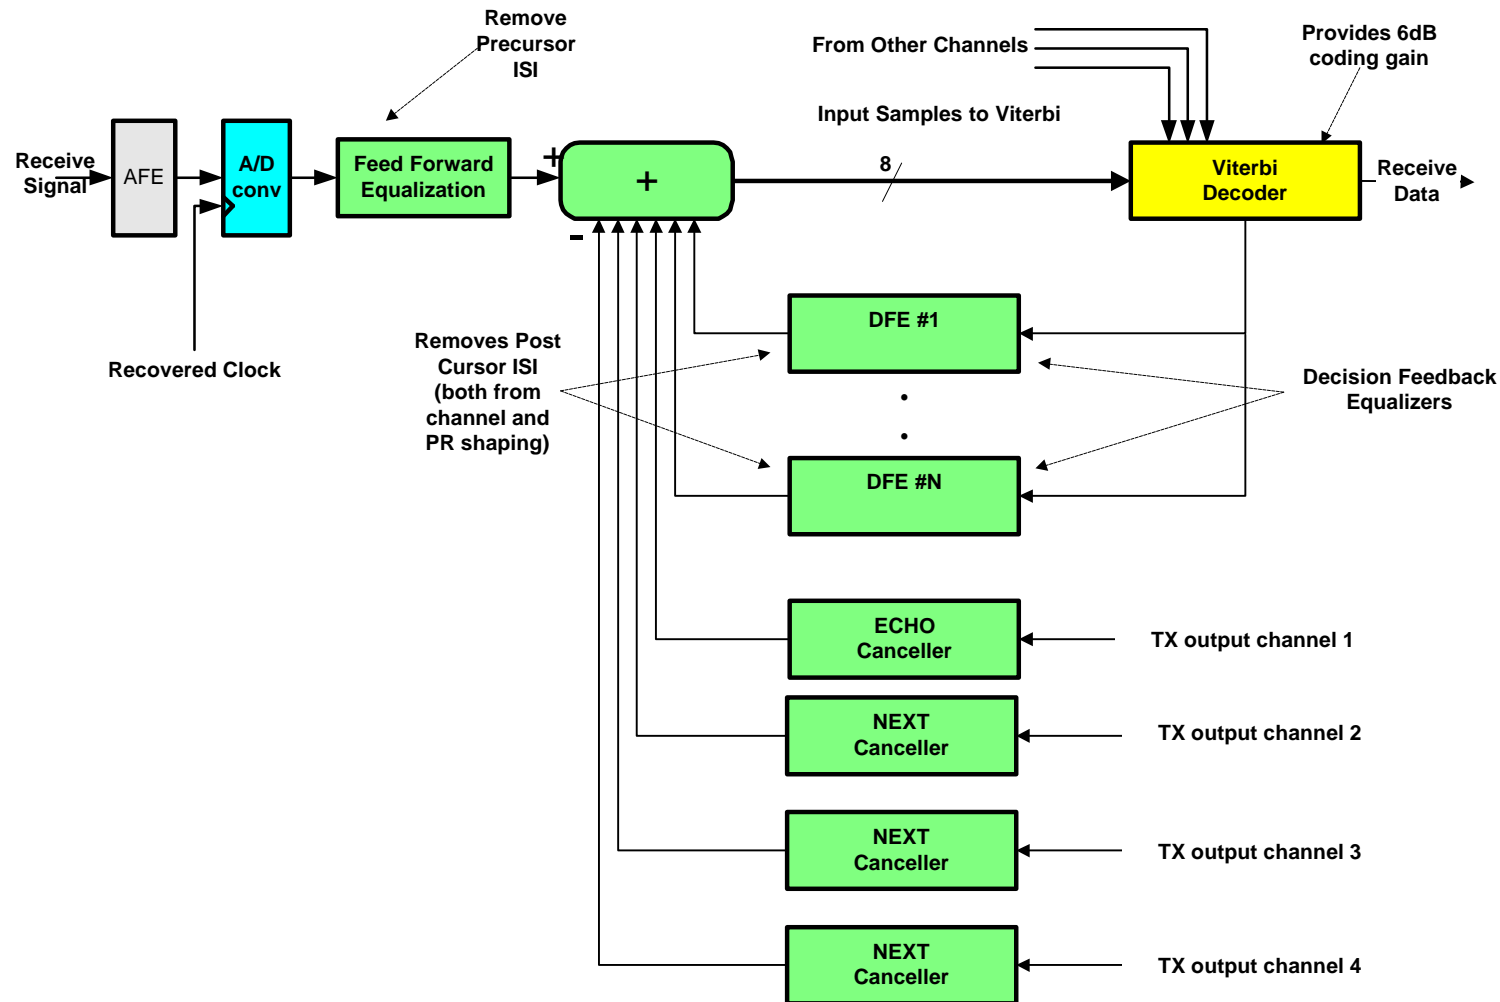
# Typical 100TX DSP Based Receiver



# Detail of DFE Based Equalizer




# Gigabit Receiver With FEC



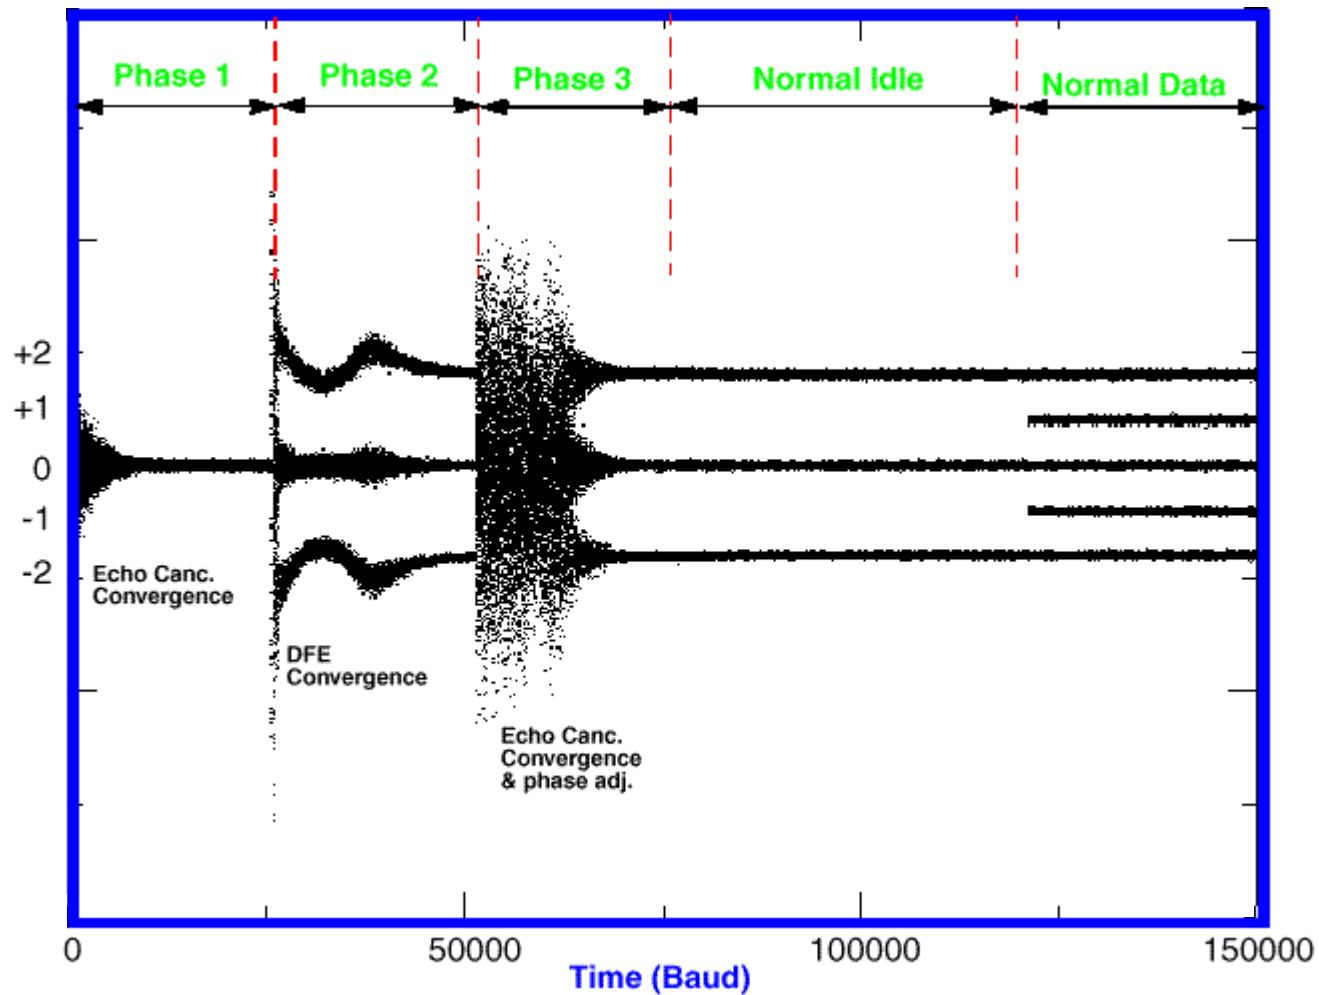
# PHY Start-up

---

- 1000BASE-T Start-up protocol provides for two different start-up procedures for maximum robustness while preserving complete interoperability:
  - *Blind Start-up*: where all adaptive filter blocks converge without sequencing. If the PHYs at the two ends of the link both implement this procedure (which is determined during the autonegotiation), the sequenced start-up is bypassed.
  - *Sequenced Start-up*: where the convergence of various adaptive blocks are separated in a 3 step sequence.



# Eye Diagrams (master sequenced start-up)



# Receiver Design Parameter Comparison

---

	<b>100BASE-TX Receiver</b>	<b>1000BASE-T Receiver</b>
A/D Conversion	5.5 bit ideal at 125MSamples/sec	7 bit ideal at 125MSamples/sec
DFE	10 Taps	14 Taps/Channel
FFE	8 Taps	12 Taps/Channel
NEXT Cancellers	0	75 Taps/Channel
ECHO Canceller	0	60 Taps
Critical Path	3 Input Add + 3 Input Select + 1 Slicer	4 Input Add-Compare Select + 3 Input Add + 5 Input Select + Branch Metric Compute
Normalized Gate Complexity	1	8

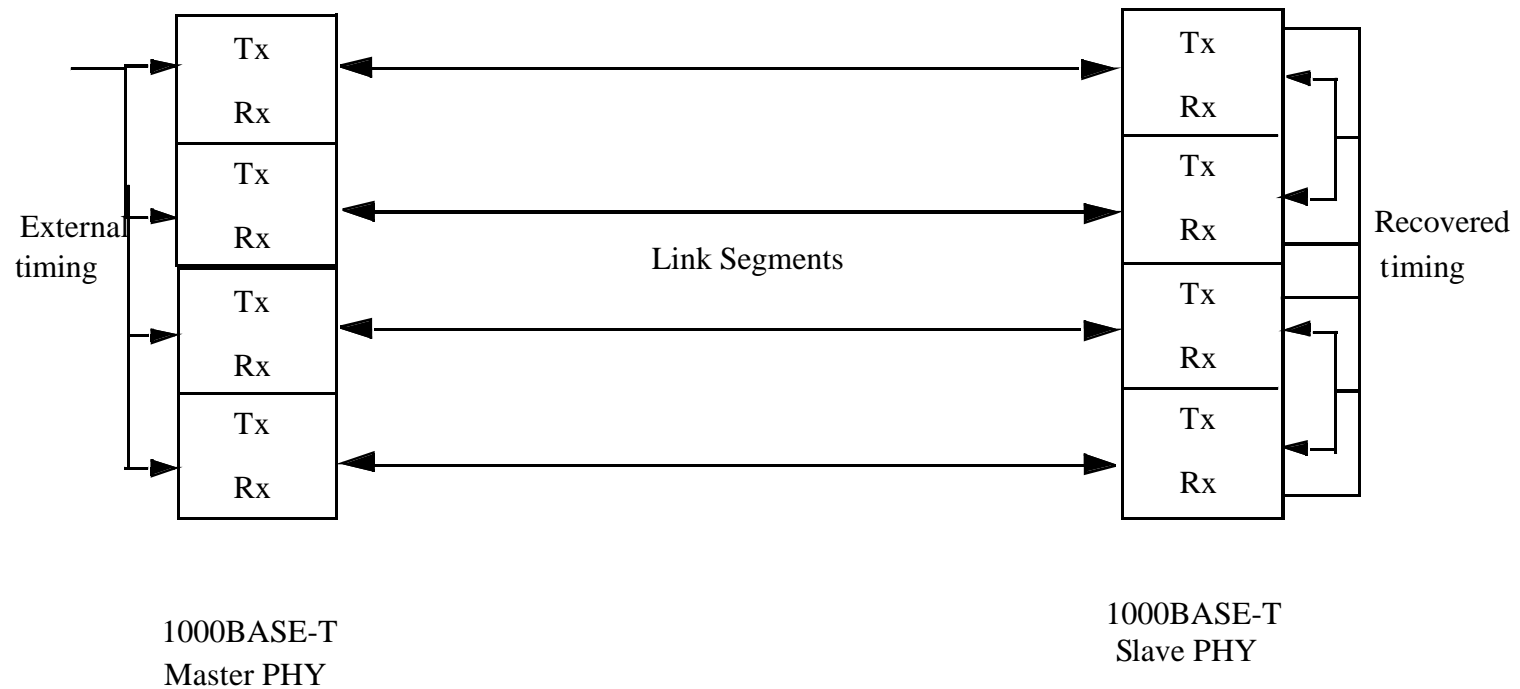
# Technical Details of 1000BASE-T

---

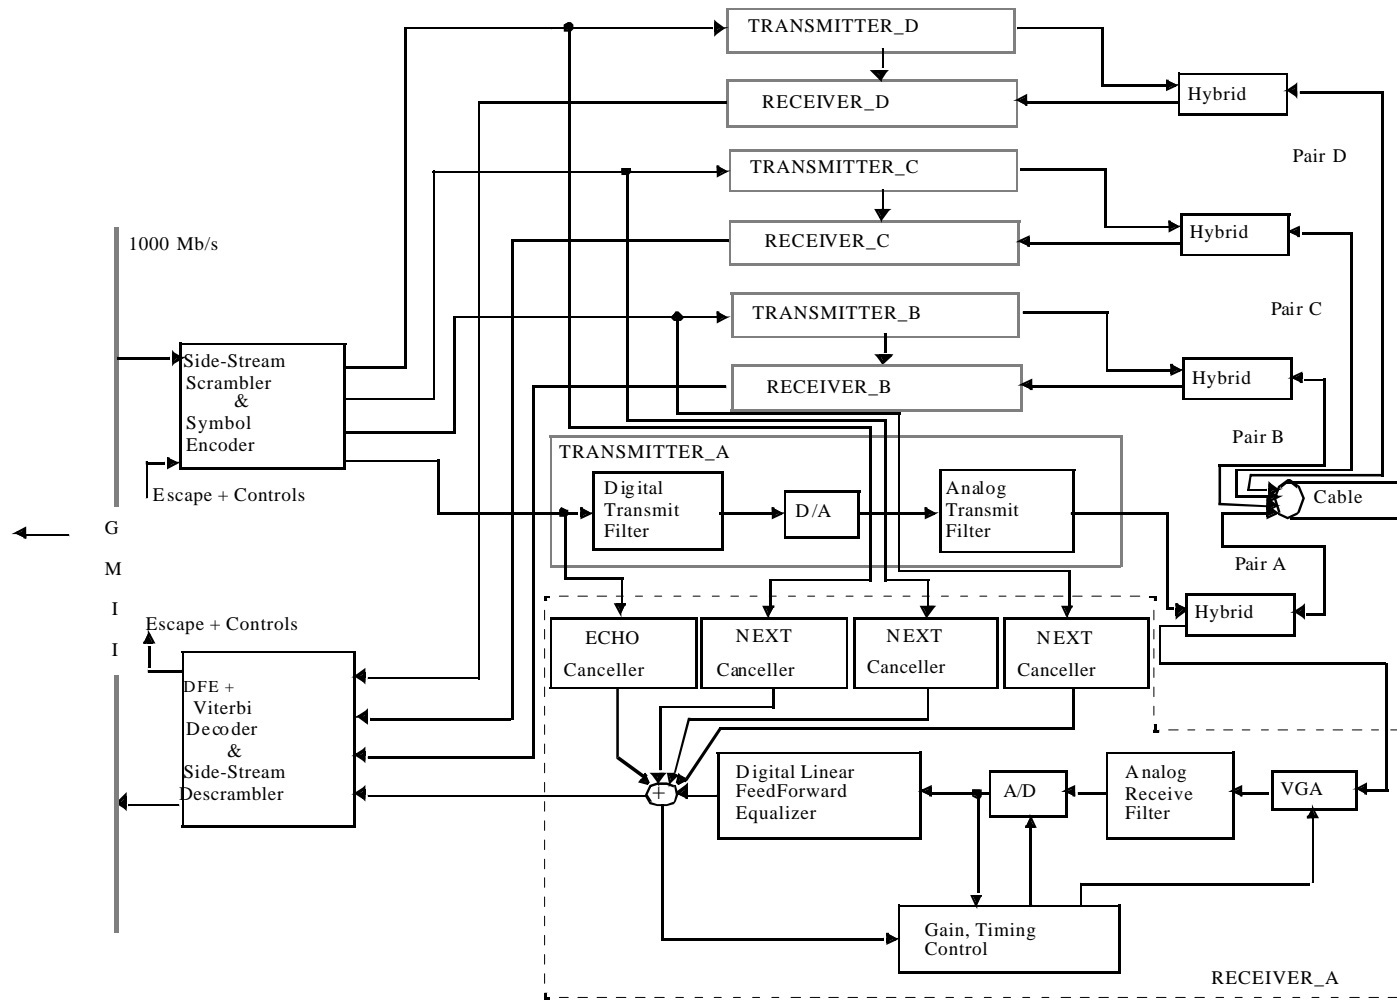
- 1000BASE-T Topology
- Operation of 1000BASE-T
- Signaling
- Trellis Coded Modulation
- Performance Evaluations

# 1000BASE-T Topology


---



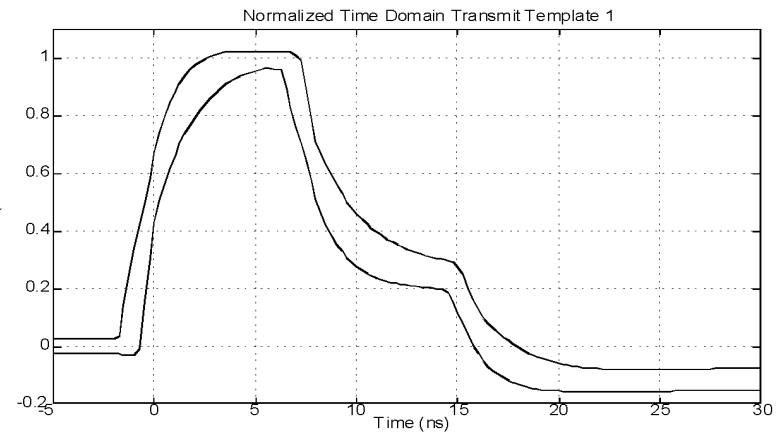
# Block Diagram of Transceiver



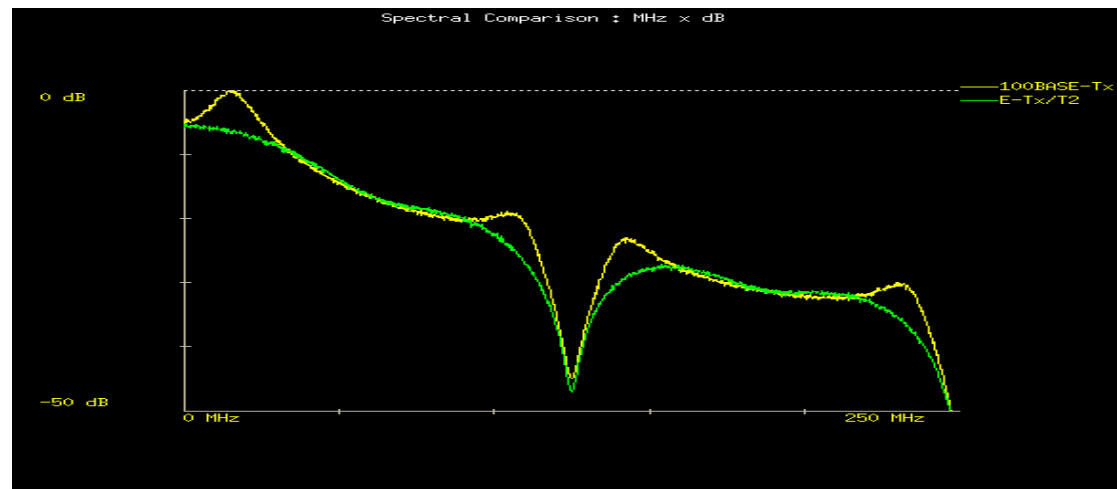
# Transmitter



Example Implementation of Transmitter



Pulse Template



Transmit Spectrum of 1000BASE-T

# Signaling Features

---

- Forward Error Correction (FEC) coded symbol mapping for data
- Algorithmic mapping and inverse mapping from octet data to a quartet of quinary symbols and back
- Uncorrelated symbols in the transmitted symbol stream
- No correlation between symbol streams traveling both directions on any pair combination
- No correlation between symbol streams on pairs A, B, C and D
- Ternary symbol mapping in idle and training modes to ease blind startup and retraining

# Signaling Features (cont.)


---

- Ability to rapidly or immediately determine if a symbol stream represents data, idle or carrier extension
- Robust delimiters for SSD, ESD, and other control signals
- Ability to signal the status of the local receiver to the remote PHY to indicate that the local receiver is not operating reliably and requires retraining
- Ability to automatically detect and correct for pair swapping and unexpected cross-over connections
- Ability to automatically detect and correct for incorrect polarity in the connections
- Ability to automatically correct for differential delay variations across the wire-pairs


# Side Stream Scrambler

---

- Side Stream Scrambler employed by Master PHY

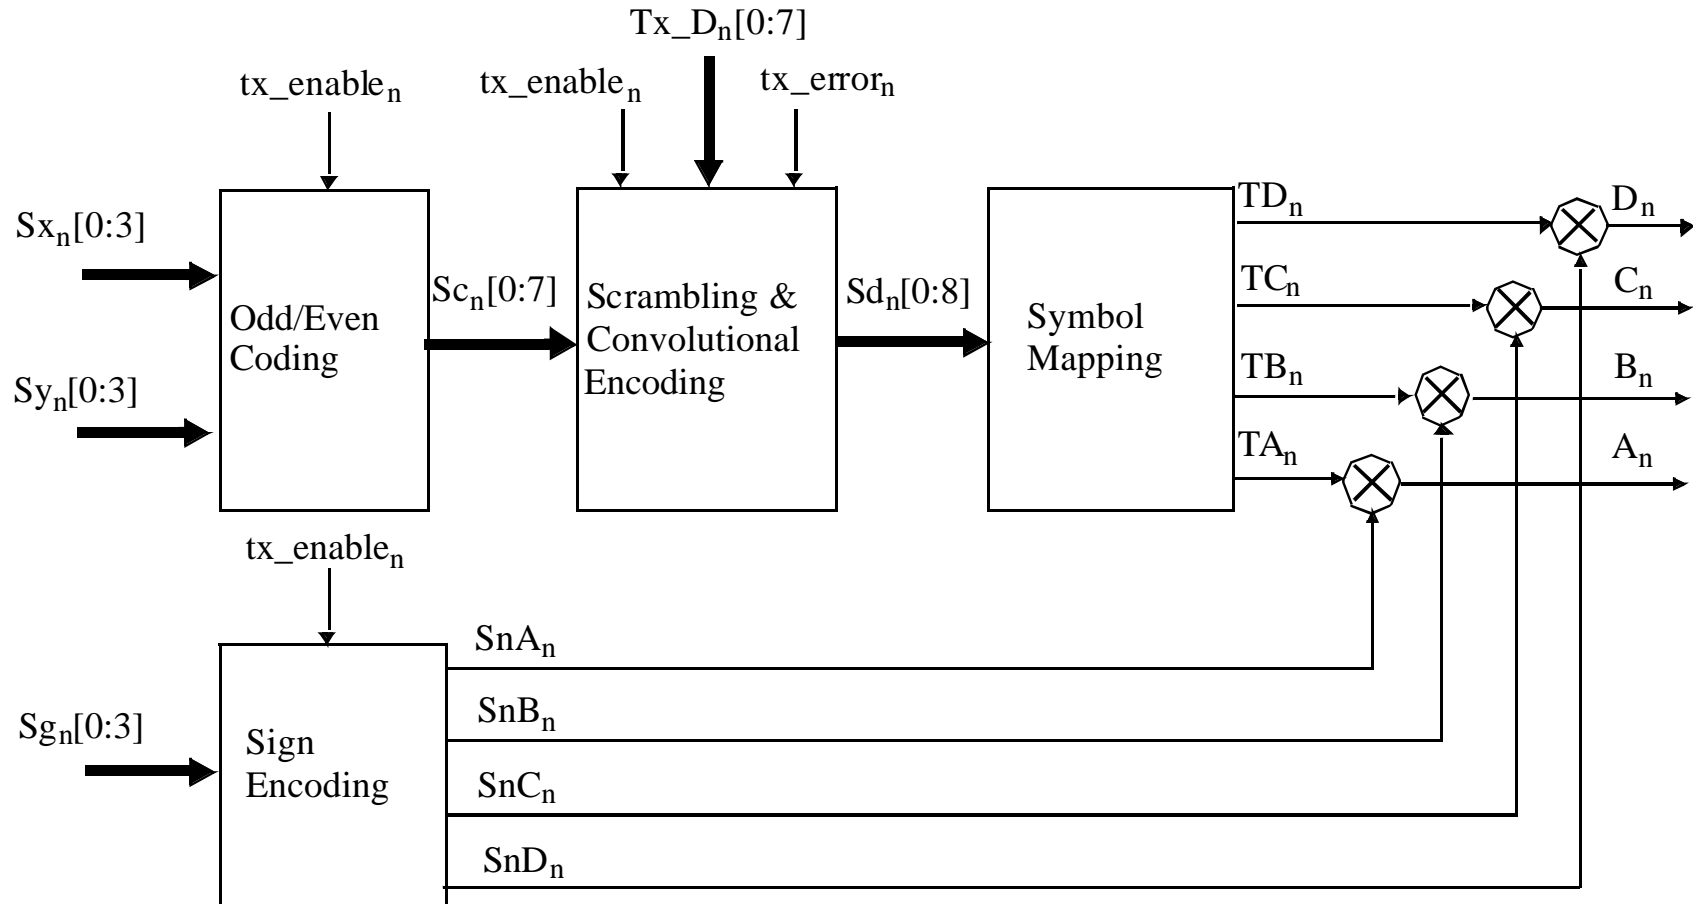


- Side Stream Scrambler employed by Slave PHY





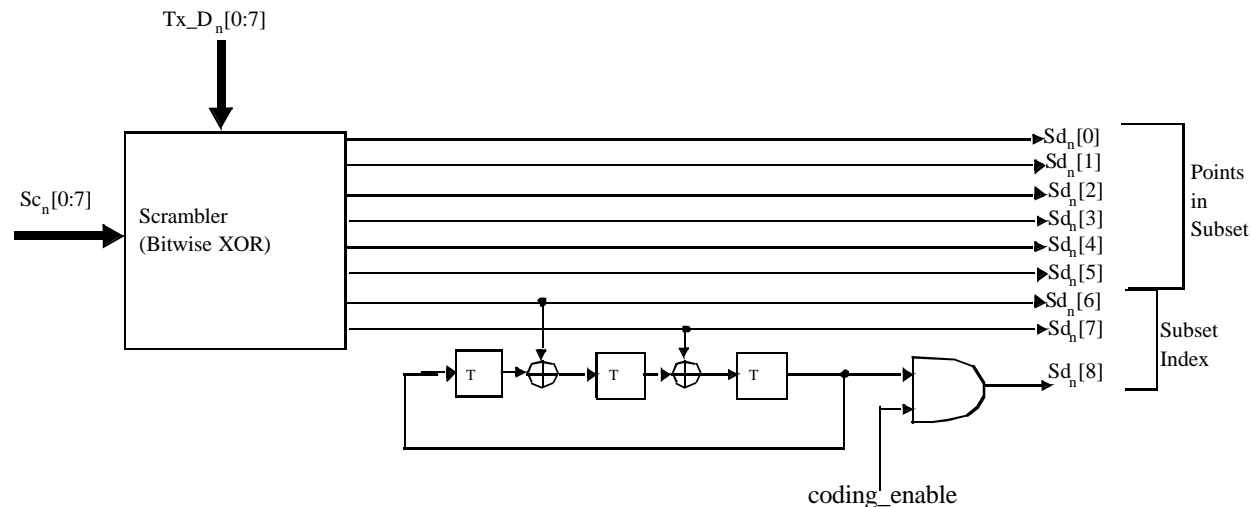
# Symbol Mapping Reference Diagram



# Trellis Coded Modulation

Purpose: To match the robustness of 3-level 100BASE-TX signaling with 5-level 1000BASE-T signaling

- Implemented as a two-step approach:
  - Convolutional Encoding to convert scrambled octet data to 9-bit word.
  - Mapping by Set Partitioning to get 6dB noise immunity gain.



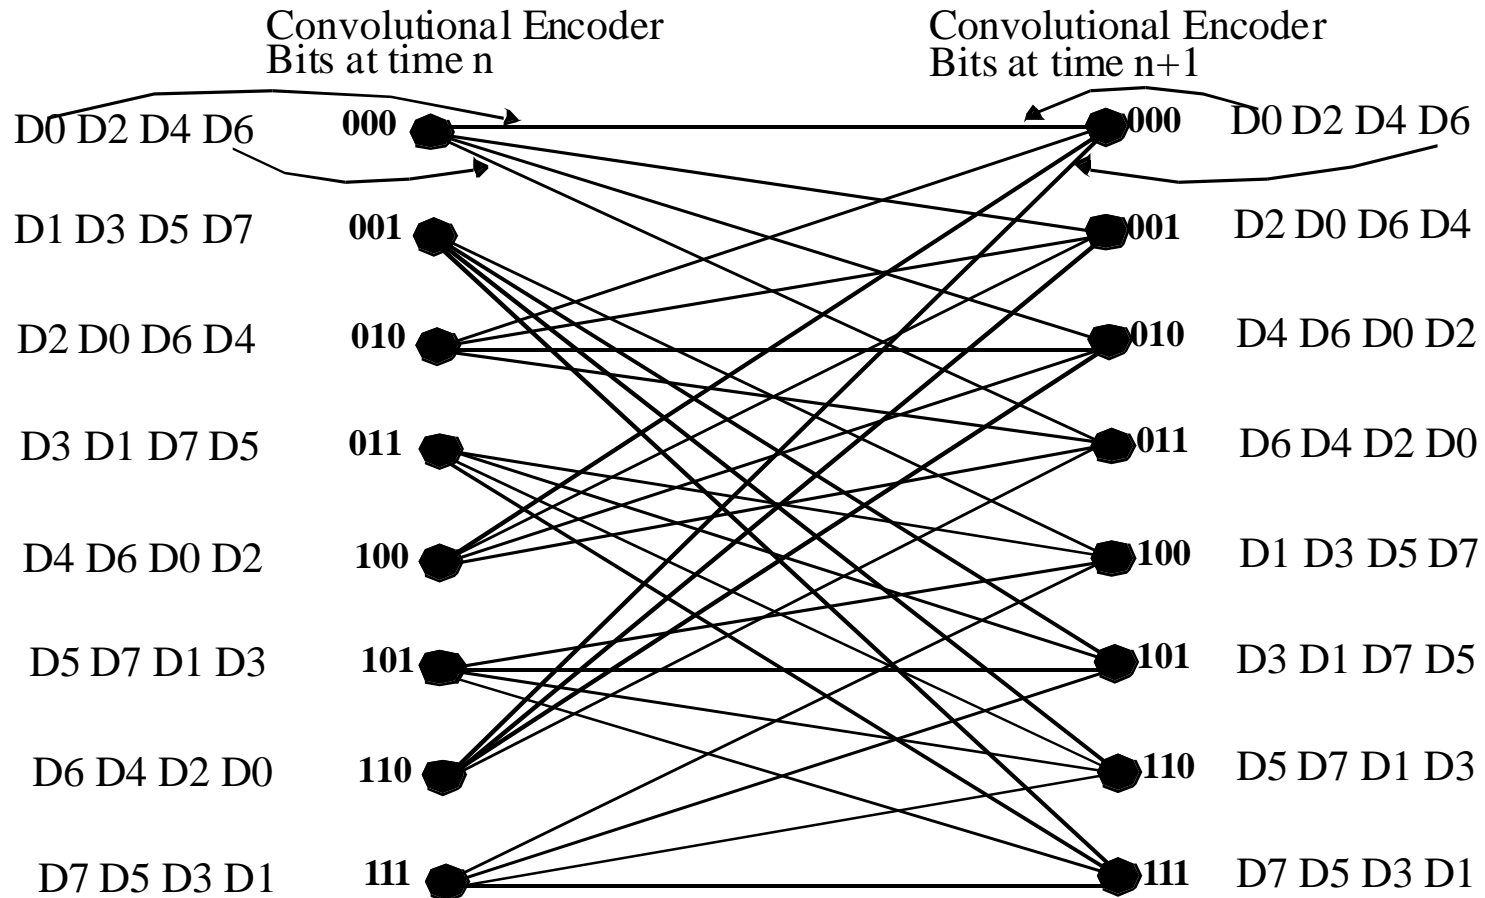
# Subset Mapping

---

- Partition 5-levels into 2 1-D subsets:  
 $X=\{-1,+1\}$ ,  $Y=\{-2,0,+2\}$ .
- Squared distance between elements in 1D subset is 4.

Subset	$Sd_n$ [6:8]	Pattern 1 $TA_n TB_n TC_n$ $TD_n$	Number of Elements in 1	Pattern 2 $TA_n TB_n TC_n$ $TD_n$	Number of Elements in 2	Total Number of Elements
D0	000	XXXX	16	YYYY	81	97
D2	010	XXYY	36	YYXX	36	72
D4	100	XYYX	36	YXXY	36	72
D6	110	XYXY	36	YXYX	36	72
D1	001	XXXY	24	YYYX	54	78
D3	011	XXYX	24	YYXY	54	78
D5	101	XYYY	54	YXXX	24	78
D7	111	XYXX	24	YXYY	54	78


# Trellis Diagram



# Sequential Decoding

---


- Squared distance between valid paths is also 4.



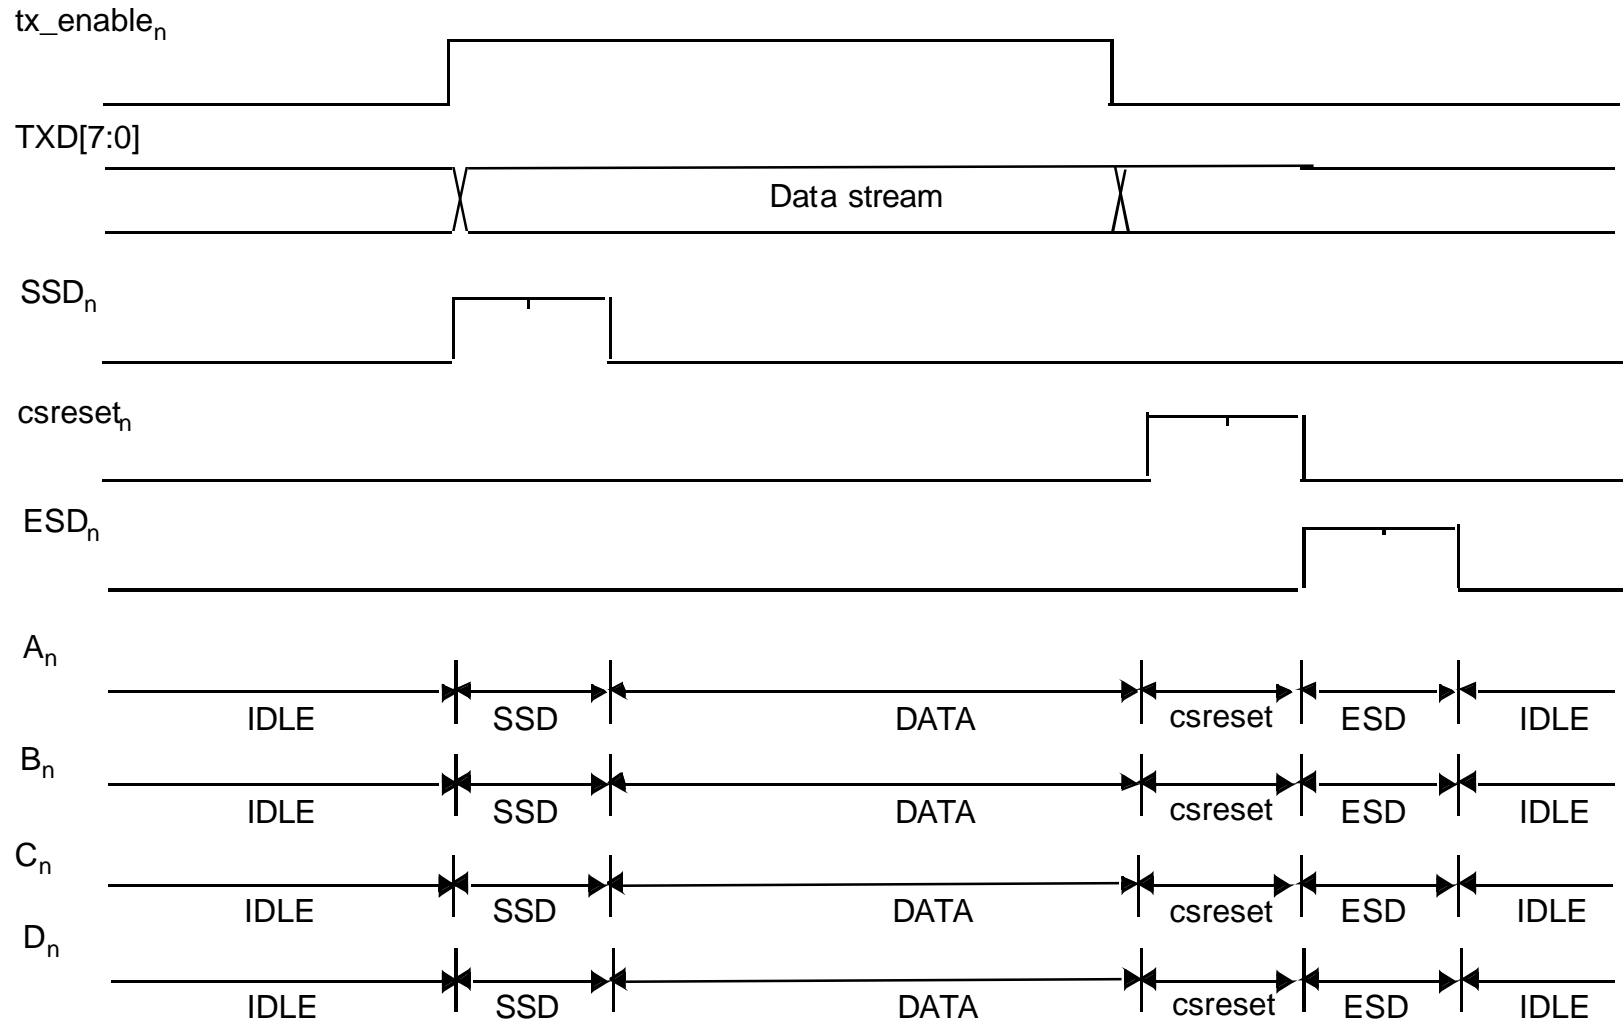
# Packetizing the Trellis Code

---

- Idle/Carrier Extension use 3-level signaling i.e. D0 subset
- Data uses Trellis Coding
- Reset convolutional encoder states to zero using two symbol periods at End of Packet



# 1000BASE-T Frame Structure

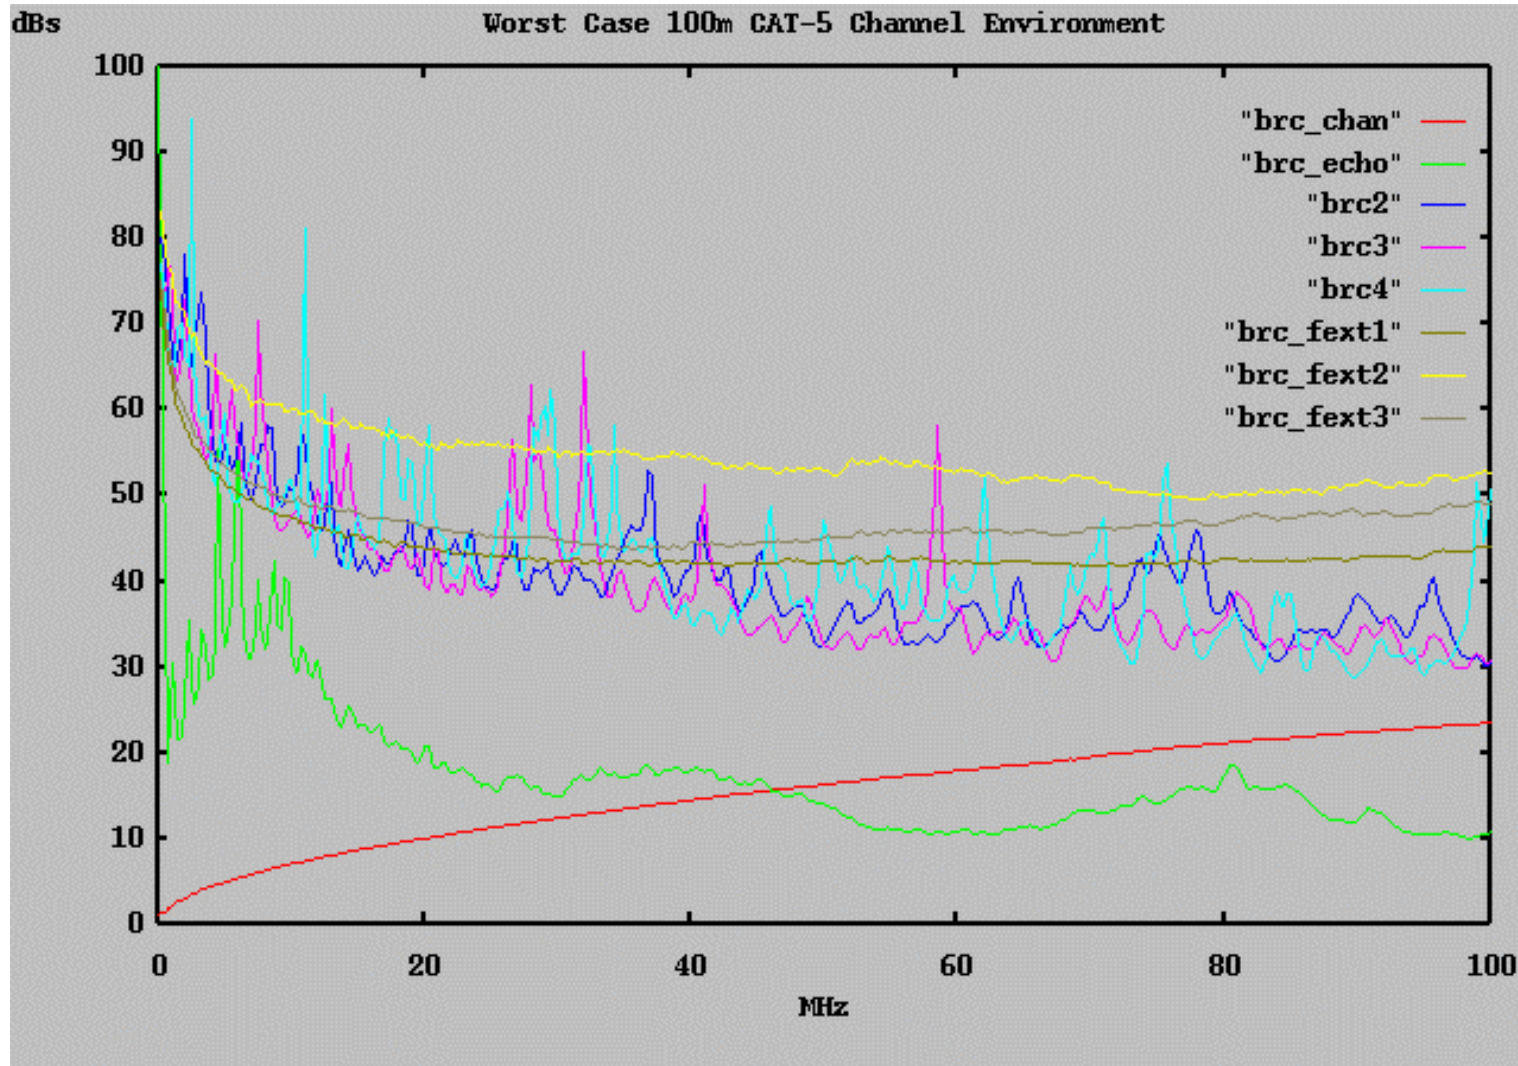


# Performance Evaluations

---

- Worst Case 100m UTP-5 Channel characteristics.
- Design Point - 10dB
- Design Point Simulations
- Matlab Code Published

# Worst Case 100m UTP-5 Channel




# Design Point - 10dB

---

- D/A: 17 levels at 125MHZ
- Launch Level: 2V P-P
- Analog Transmit Filter: Single pole RC
- Analog Receive Filter: BW2@100MHz
- A/D: 6.5bits ideal at 125MHz  
2.2 Volts p-p (prob. of clipping 1E-25)
- Baseline Wander Correction: Digital
- FFE - #taps: 16 taps at 125MHz
- DFE - #taps: 12 taps at 125MHz
- NEXT Cancellers - #taps: 72 taps at 125MHz
- Echo Canceller - #taps: 120 taps at 125MHz
- Viterbi Decoder: 12-stage
- Total worst-case latency: 31BT < 40BT
- Uniform Jitter Tolerance for 0dB margin: 1.5ns P-P [>10ns P-P Gaussian]
- Worst-Case noise immunity: Crane Test: 140mV P-P
- Est. Gate Count/Power Consumption: 330K/4W
- Margin without FEXT: 10.5dB
- Margin with Worst-Case FEXT: 5.7dB


# Design Results - 10dB

---



# Design Results - 10dB

---

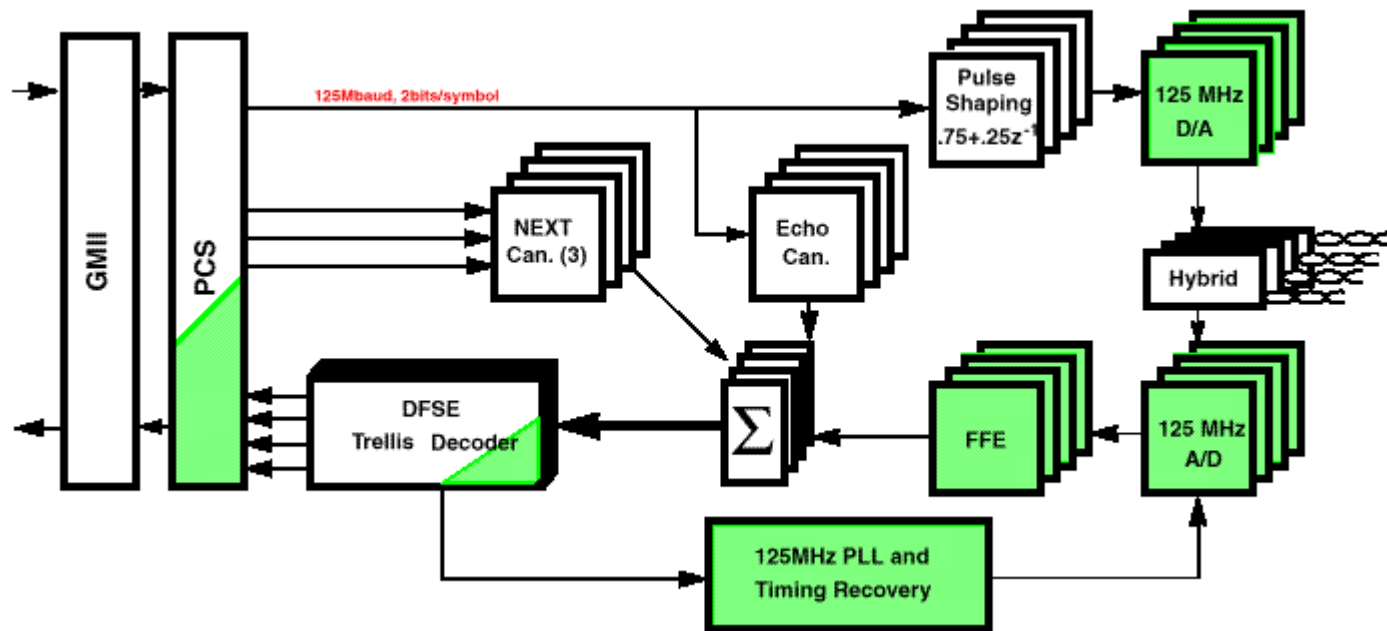



# Detailed VLSI Implementation

---

- DSP Requirements
- Technology Background
- Power requirements
- Size and Layout

# Transceiver Block Diagram



: Blocks common between a Quad 100Base-TX and a 1000BASE-T transceiver. NEXT and Echo cancellers are the major blocks contributing to the added complexity over a Fully digital Quad-TX.


# DSP Requirements

---

- 4 Equalizers (FFE/ FBE), 4 echo cancellers, 12 NEXT cancellers running at 125 MHz clock rate which is also the Baud rate
- The bulk of the computation is in the adaptive Finite Impulse Response (FIR) filters
- The regularity of the filter structures makes the design task quite manageable

# Direct Form FIR Filter

---



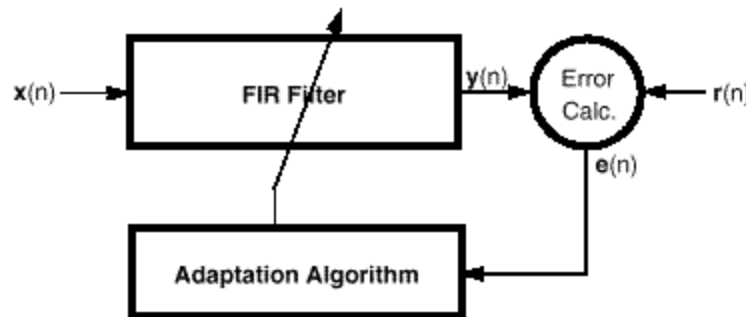
$$y(n) = h_0x(n) + h_1x(n-1) + \dots + h_{K-1}x(n-K+1)$$

$$y(n) = \sum_{k=0}^{K-1} h_kx(n-k)$$

# Adaptive FIR

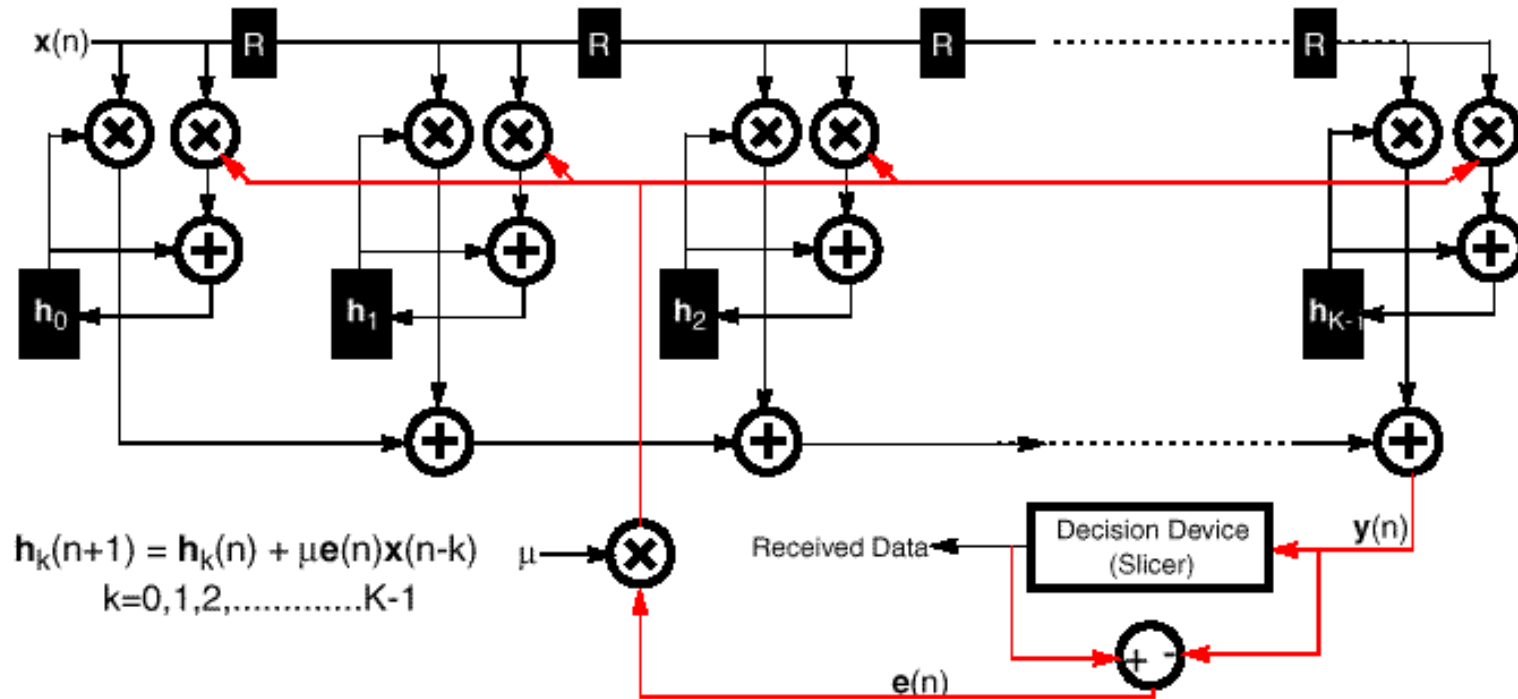
---

In an adaptive filter, in general, the output of the filter,  $y(\mathbf{n})$ , is compared against a desired reference  $r(\mathbf{n})$  which produces a *measure* to be used in an algorithm for modifying the filter coefficients (a time varying system). The filter coefficients are modified in a way to minimize the *measure* and bring the filter output close to the desired output. In an important class of adaptive filters, adaptation of the coefficients take place without any reference input. This is called blind adaptation.



The most widely used adaptation algorithm is the Least Mean Square (LMS) method where  $e(\mathbf{n}) = r(\mathbf{n}) - y(\mathbf{n})$  and the update is chosen to minimize  $e^2$ .

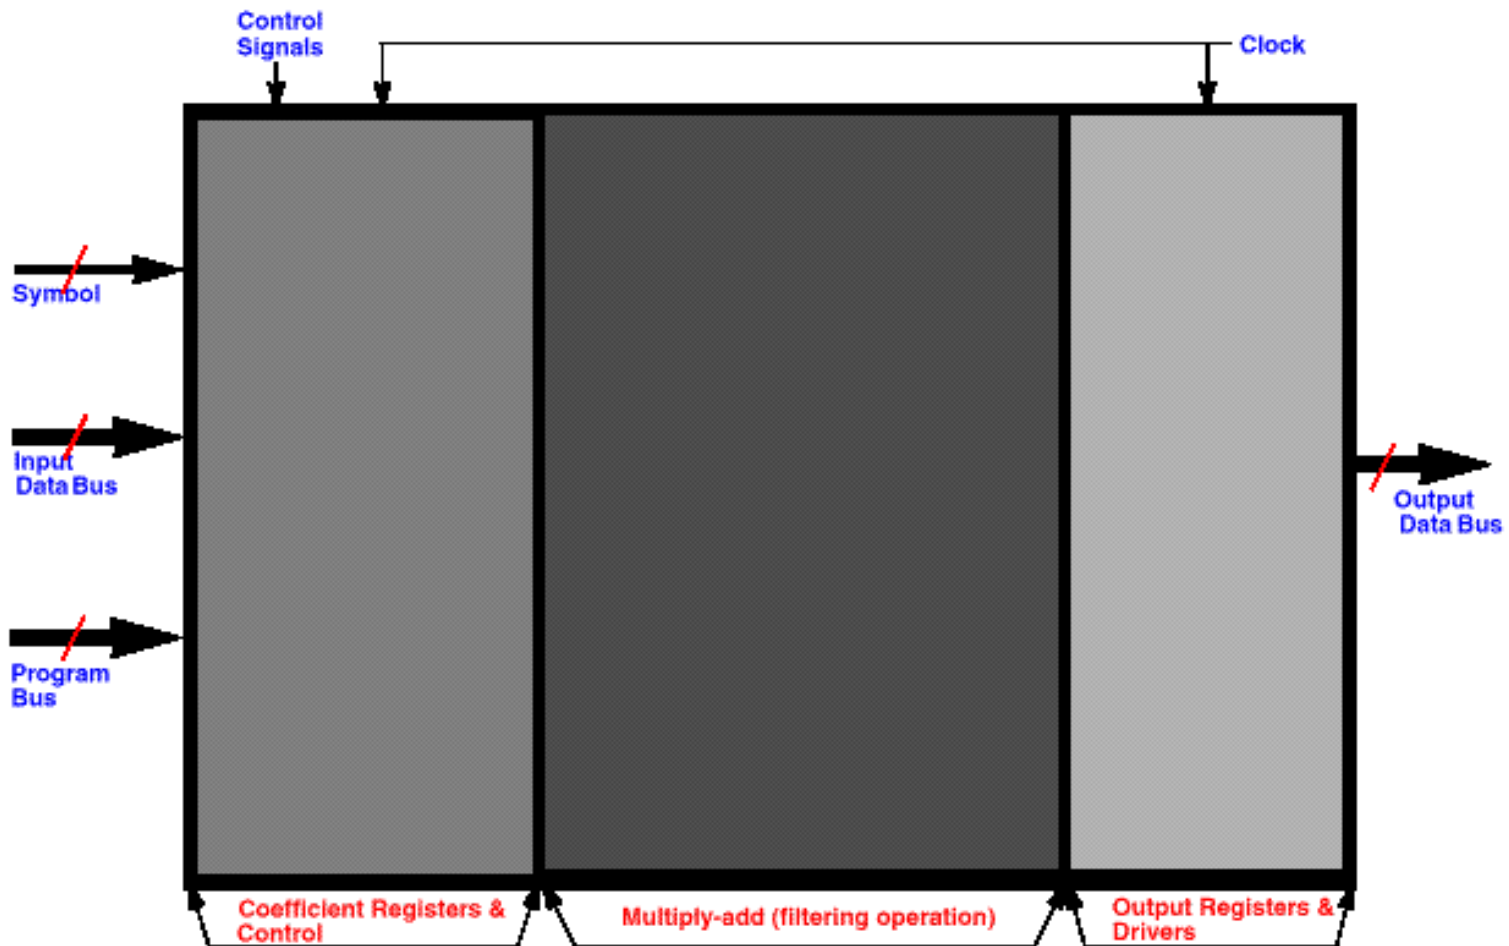
# Linear Equalizer with LMS Update



The equalizers, NEXT cancellers and echo cancellers in a 1000Base-T transceiver are all variations of this filter structure.

# Echo Canceller Tap

---



# Power Consumption

---


- A worst case DSP complexity providing a very high margin of operation:

FFE Taps	8
DFE Taps	14
NEXT Taps	80
Echo Taps	60

- Corresponds to about 400,000 gates
- Comparing to Quad-TX:
  - Fully digital Quad-TX, 3.3V, 0.50  $\mu$  CMOS 2.2W
  - 1000BASE-T, 2.5V, 0.25  $\mu$  CMOS 3.1W

# 1000BASE-T Copper Transceiver Layout

- 0.25mm CMOS, 5 layer metal
- Based on actual layout from a 0.25mm CMOS library
- 4mm X 4mm active area (65% - 75% of the area of most existing 10/100 PHYs)
- 160 pins



# 3 Steps From 100Base-TX to 1000BASE-T

---

- Start with a 100Base-TX DSP Based PHY
- Use all four channels full duplex
  - 400 MBps in both directions
  - Requires ECHO and NEXT canceling
  - Requires Master/Slave Clocking
- Remove 4B/5B Encoding
  - 500 MBps in both directions
- 5 Level Signaling - 2 Bits/Symbol
  - 1000 MBps in both directions
  - Requires FEC to get back the 6dB SNR
- The major task was adding FEC