



PHY Proposal for 10GBASE-T: Encoding, Mapping & Framing

George Zimmerman & Brett McClellan - SolarFlare Communications

Supporters:

Rick Rabinovich, Spirent

Michael Laudon, Force 10 Networks

Mike Bennett, Lawrence Berkeley National Laboratory

Paul Kish - Nordx/CDT

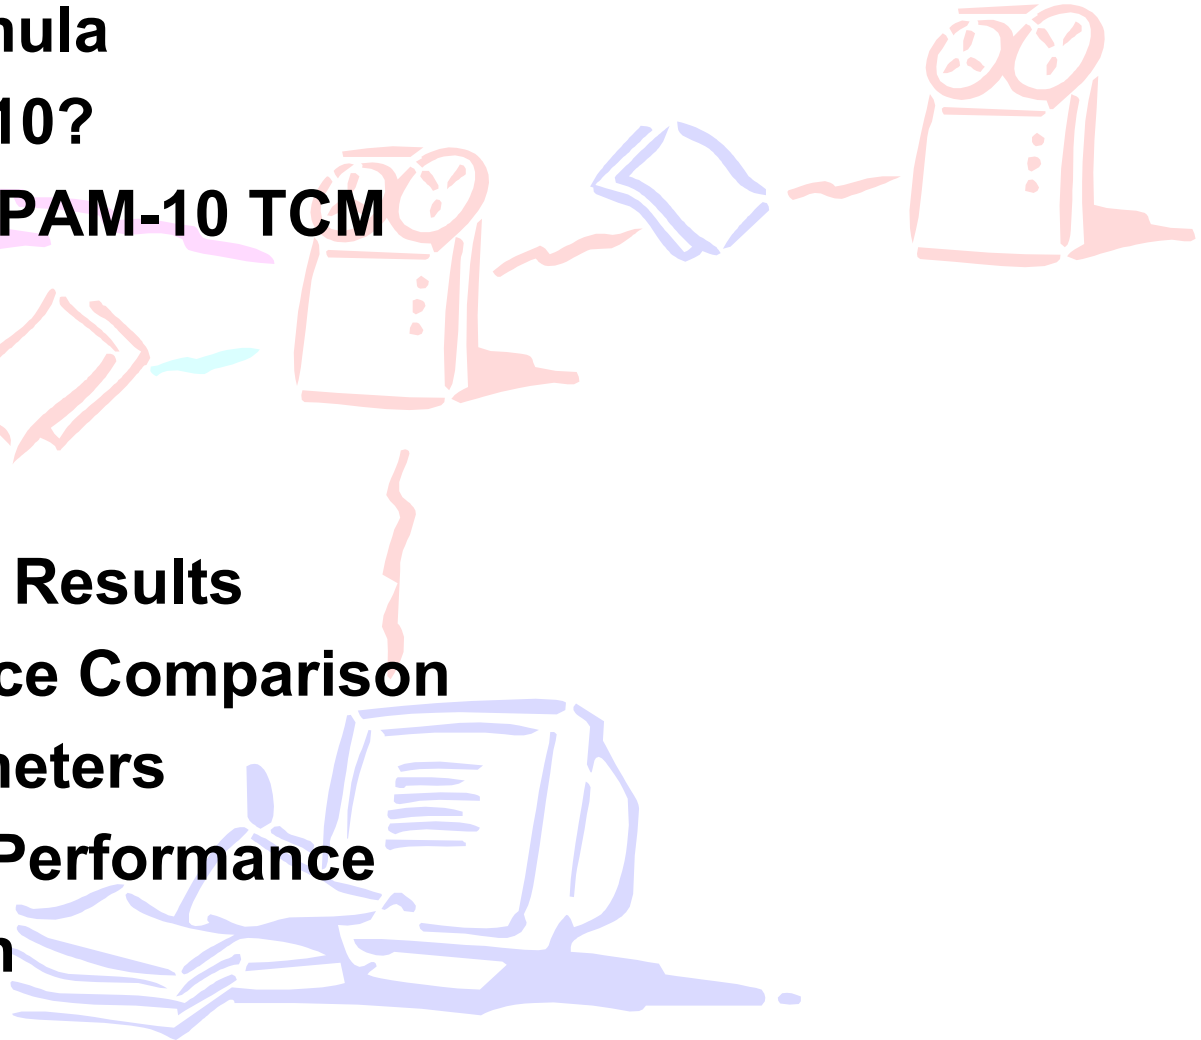
Shadi AbuGhazaleh-Hubbell Premise Wiring

Sterling Vaden- Superior Modular Products

Bernie Hammond - Krone

Outline

- **Basic Formula**
- **Why PAM-10?**
- **4D 8 State PAM-10 TCM**
- **Mapping**
- **Scrambler**
- **Framing**
- **Simulation Results**
- **Performance Comparison**
- **PHY Parameters**
- **Measured Performance**
- **Conclusion**



Line Coding: Basic Formula

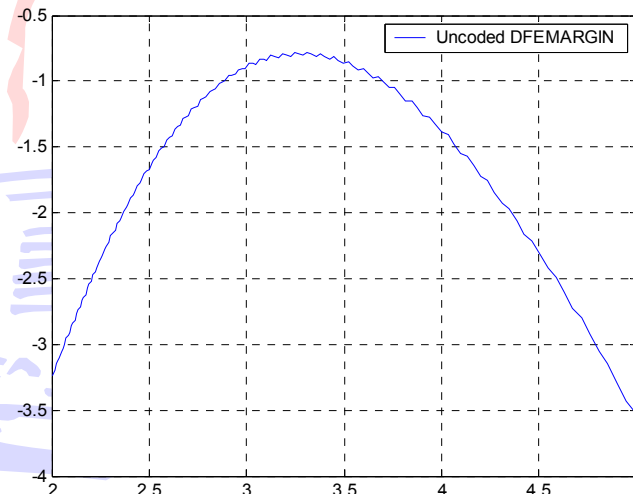
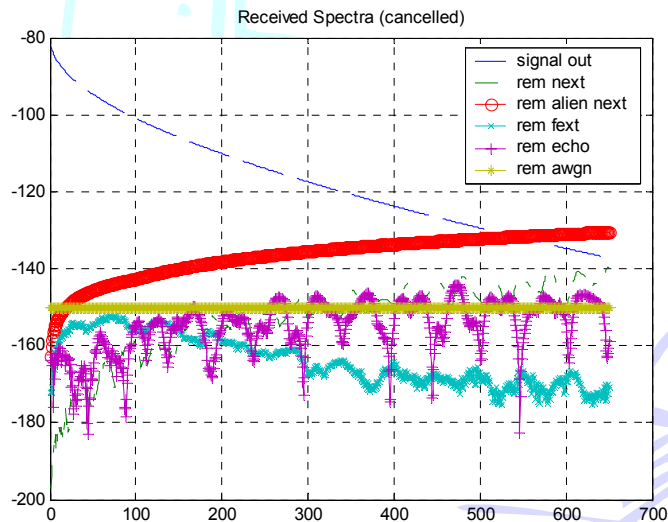
- **Start with 1000BASE-T**
- **Take away half duplex / carrier sense**
- **Support 10Gbps Data Rate: Optimize baud rate for line**
 - 8 bits / baud -> 12-bits / baud
 - PAM5 @ 125Mbaud -> PAM10 @ 833.33 Mbaud
- **4D 8-state TCM**
 - Same scrambler as 1000BASE-T
 - Same encoder as 1000BASE-T
 - Take away the “packetized” TCM
 - No convolutional reset or mode switching
 - Same 4D mapping technique as 1000BASE-T
- **Support 802.3 Clause 46**
 - XGMII Frame Structure and Control

Why PAM10?

- **Based on calculations from the channel impairments, the maximum channel capacity is achieved with a baud rate of around 800MHz.**
- **Minimal data framing complexity with byte oriented modulation**
 - 2 byte/baud -> 625 Mbaud (1 dB equalization loss)
 - 1 bytes/baud -> 1.25 Gbaud (2.5-2.9 dB loss)
- **1.5 byte/baud -> 833.33 Mbaud**
 - Hits the channel capacity sweet spot
 - Minimal additional framing complexity
- **4D PAM10 for Trellis Coded Modulation**
 - Additional 2 levels per pair provide redundancy for FEC

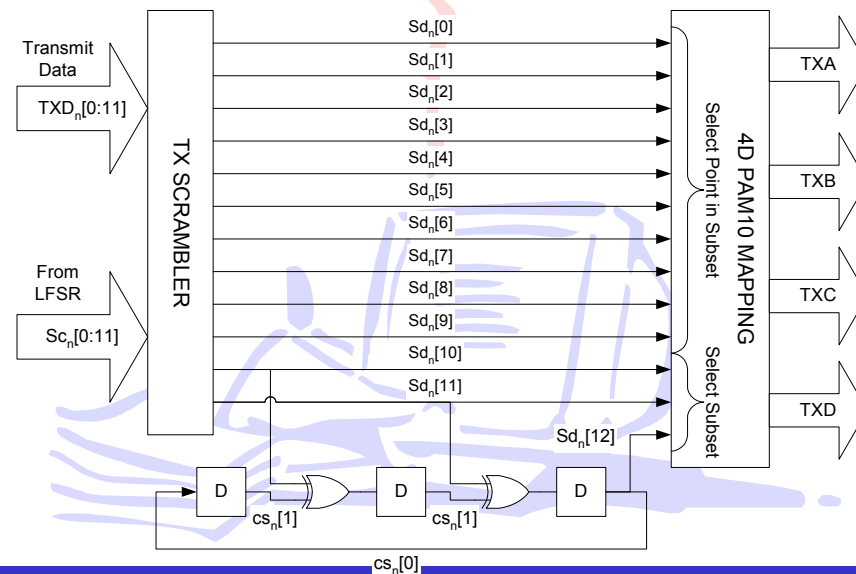
Why PAM-10?: Optimal Baud Rate

- **Solarsep7a_varlen simulations: Model 1**
 - No ANEXT Canc., 100m Cat7 IL, Baseline ANEXT (62.5 dB at 100 MHz, split slope), Default Cancellation parameters
- **833 Mbaud = 3 bits/ baud SNR = 29.4 dB**
- **Margin peak = 770 Mbaud, 0.12 dB better margin**



4D 8-State PAM10 TCM

- 4D 8-state Ungerboeck code from 1000BASE-T
 - 12-bits \rightarrow 13 bits per baud
- Continuous Trellis Coding
 - Encoder runs on a continuous basis
 - No need to switch modes between Data and Idle
 - Idle and Control symbols are covered by TCM



4D PAM10 Mapping

- **PAM10**
 - { -9 -7 -5 -3 -1 +1 +3 +5 +7 +9 }
- **13-bits / baud**
 - $2^{13} = 8192$ data symbols
 - 4D PAM10 has 10,000 constellation points
 - Remaining 1808 points are used for constellation shaping and control symbol mapping
- **Constellation shaping**
 - Eliminate any 4D points with two or more +/- 9's
 - 0.64 dB of shaping gain

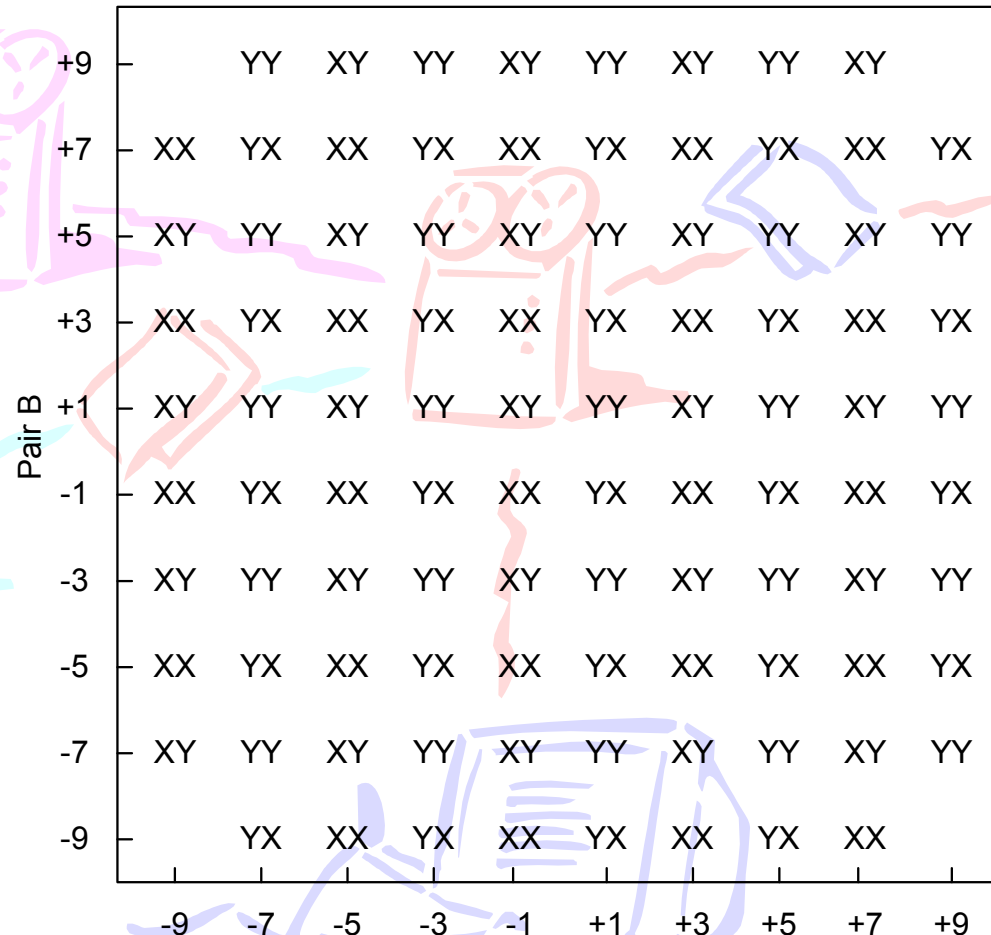
1D PAM Level Rate of Occurrence in the 4D Mapping (8192 points)

-9	-7	-5	-3	-1	+1	+3	+5	+7	+9
512	896	896	896	896	896	896	896	896	512

4D PAM10 Data Symbol Mapping

- Parity + 2 MSB's ($Sd_n12:10$) select 1 of 8 subsets
- 10 LSB's select a point within the subset
 - Sd_n9 selects normal mapping or ESC code (+/- 9)
 - Sd_n8 selects X or Y primary subset
 - $Sd_n7:0$ selects a point within the subset
 - when $Sd_n9 == 1$, $Sd_n7:6$ selects on which pair the ESC code is transmitted
- **Balanced constellation**
 - No polarity scrambler is required
 - Remaining 4 LFSR outputs are used to scramble the 4 additional data bits

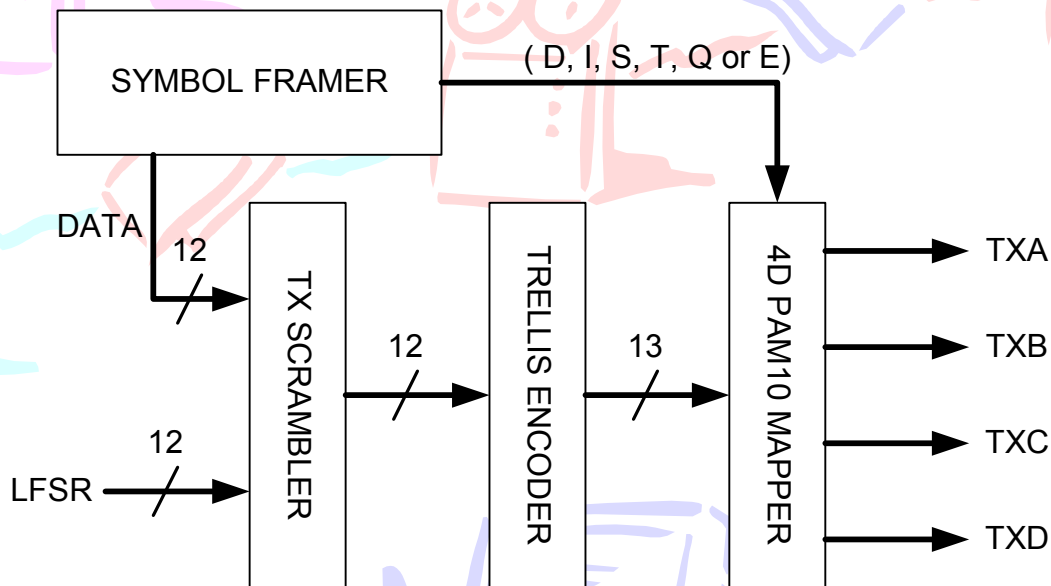
4D PAM10 Partitioning



2D Constellation Points of 4D PAM10 Mapping

Control Symbol Mapping cont.

- **Symbol Framer determines control symbol mapping**



Control Symbol Mapping

- **Map control characters into unused constellation**
 - Map to 4D PAM10 symbols containing two ± 9 's
 - 4D Symbols fall within the subset selected by encoder
 - Symbols are DC balanced
 - Examples:

	TXA	TXB	TXC	TXD
• START	± 5	± 9	± 5	± 9
• TERMINATE	± 9	± 9	± 7	± 7
• ERROR	± 1	± 1	± 9	± 9
• Q	± 1	± 9	± 1	± 9
• IDLE*	± 9	± 9	± 3	± 3
- **Infrequent nature of control symbols allows shaping gain to be maintained**

Side Stream Scrambler

- From 1000BASE-T

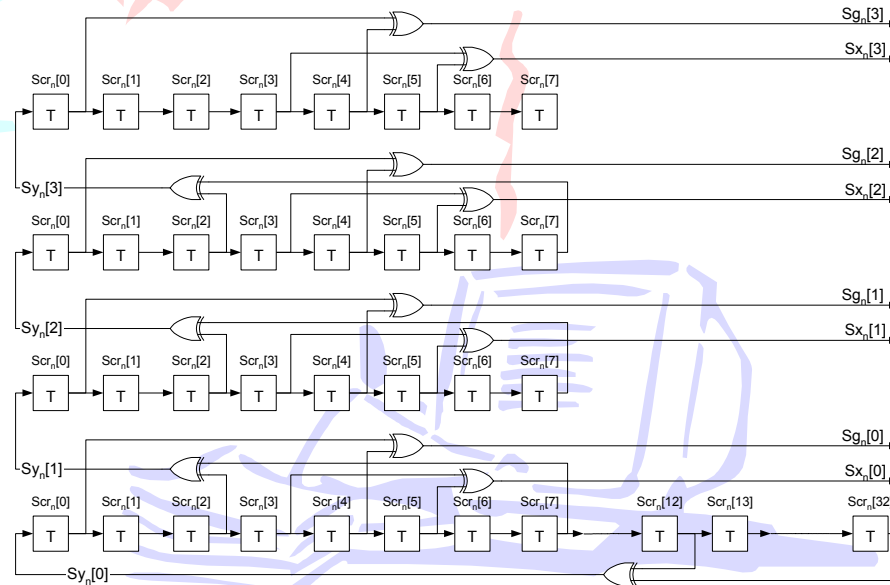
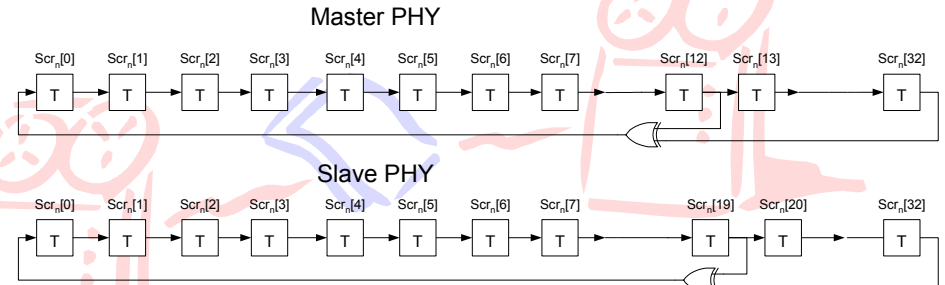
- Master TX polynomial

- $G_M(x) = 1 + x^{13} + x^{33}$

- Slave TX polynomial

- $G_M(x) = 1 + x^{20} + x^{33}$

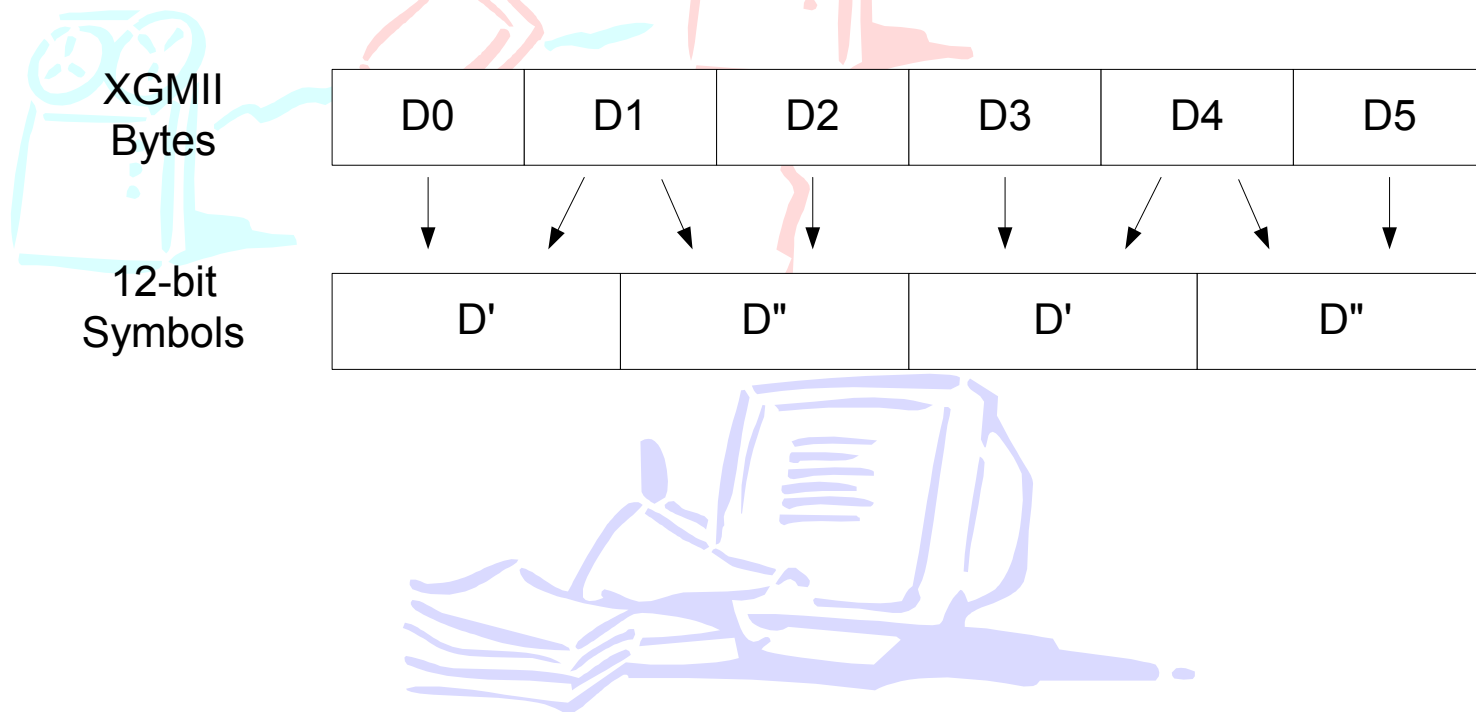
- 12 scrambler bit streams: $Sy_n[3:0]$, $Sx_n[3:0]$, $Sg_n[3:0]$



1000BASE-T Scrambler

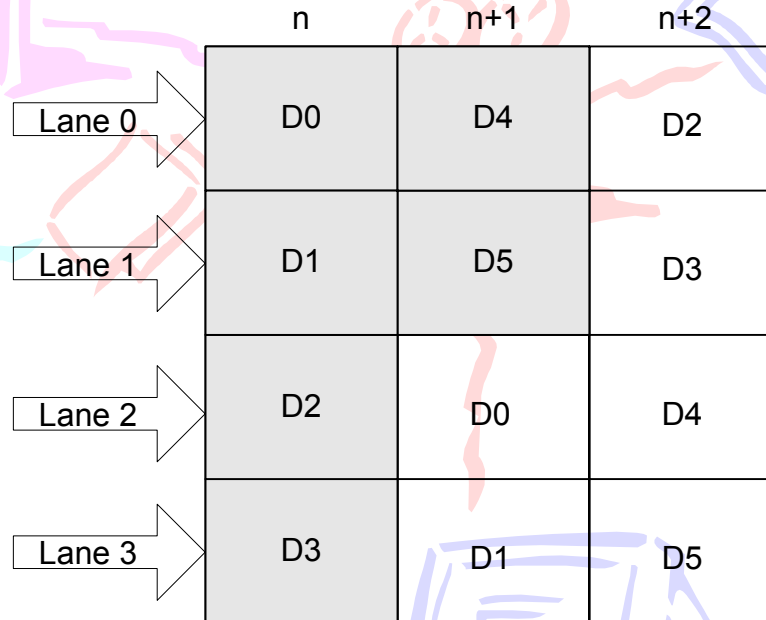
Data Framing

- **Data bytes are taken in groups of six and parsed into 4 12-bit symbols**
 - **Every third byte is split in half**



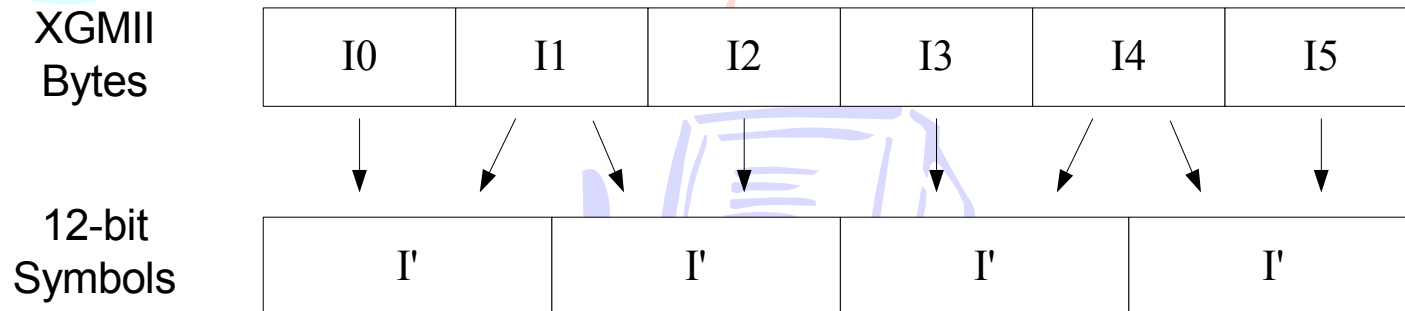
Data Framing cont.

- Three XGMII transfers per two framing blocks



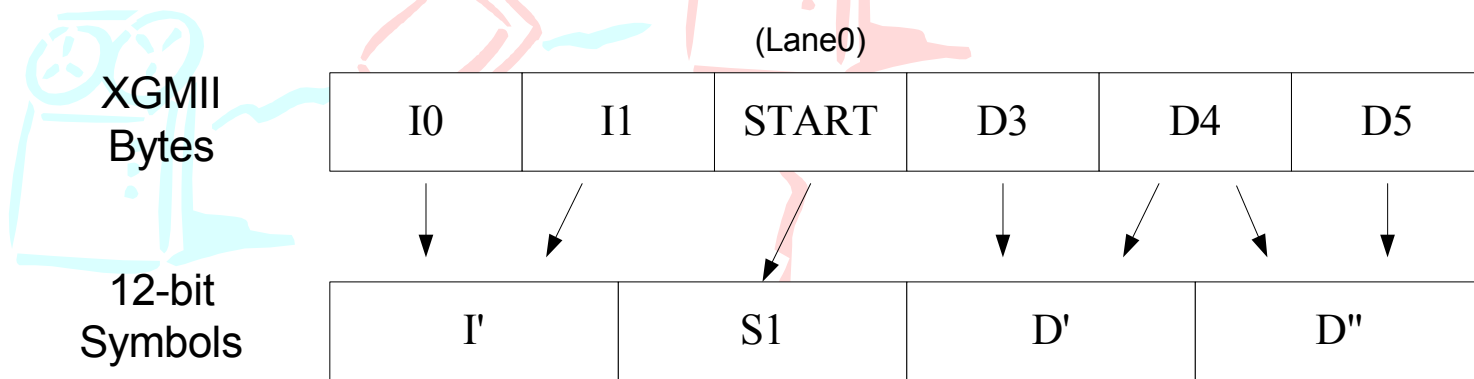
IDLE

- **XGMII IDLE bytes are replaced with 4D PAM10 trellis coded Idle symbols**
 - Idle is differentiated from data by using double ESC symbols
 - *Average TX power is maintained by mixing in non-ESC symbols



START

- **START (Start of Frame) moves into the prior/current symbol in the IPG**
 - Data bytes are left unaffected

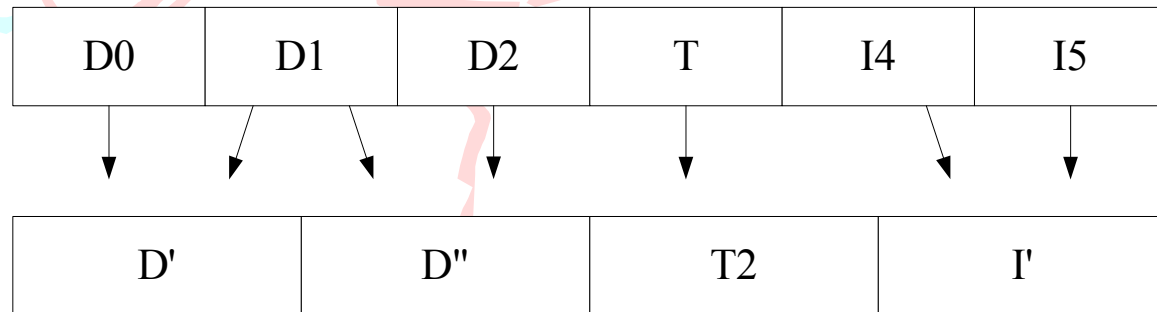


TERMINATE

- **TERMINATE (End of Frame) moves into the next/current symbol in the IPG**
 - Data bytes are left unaffected

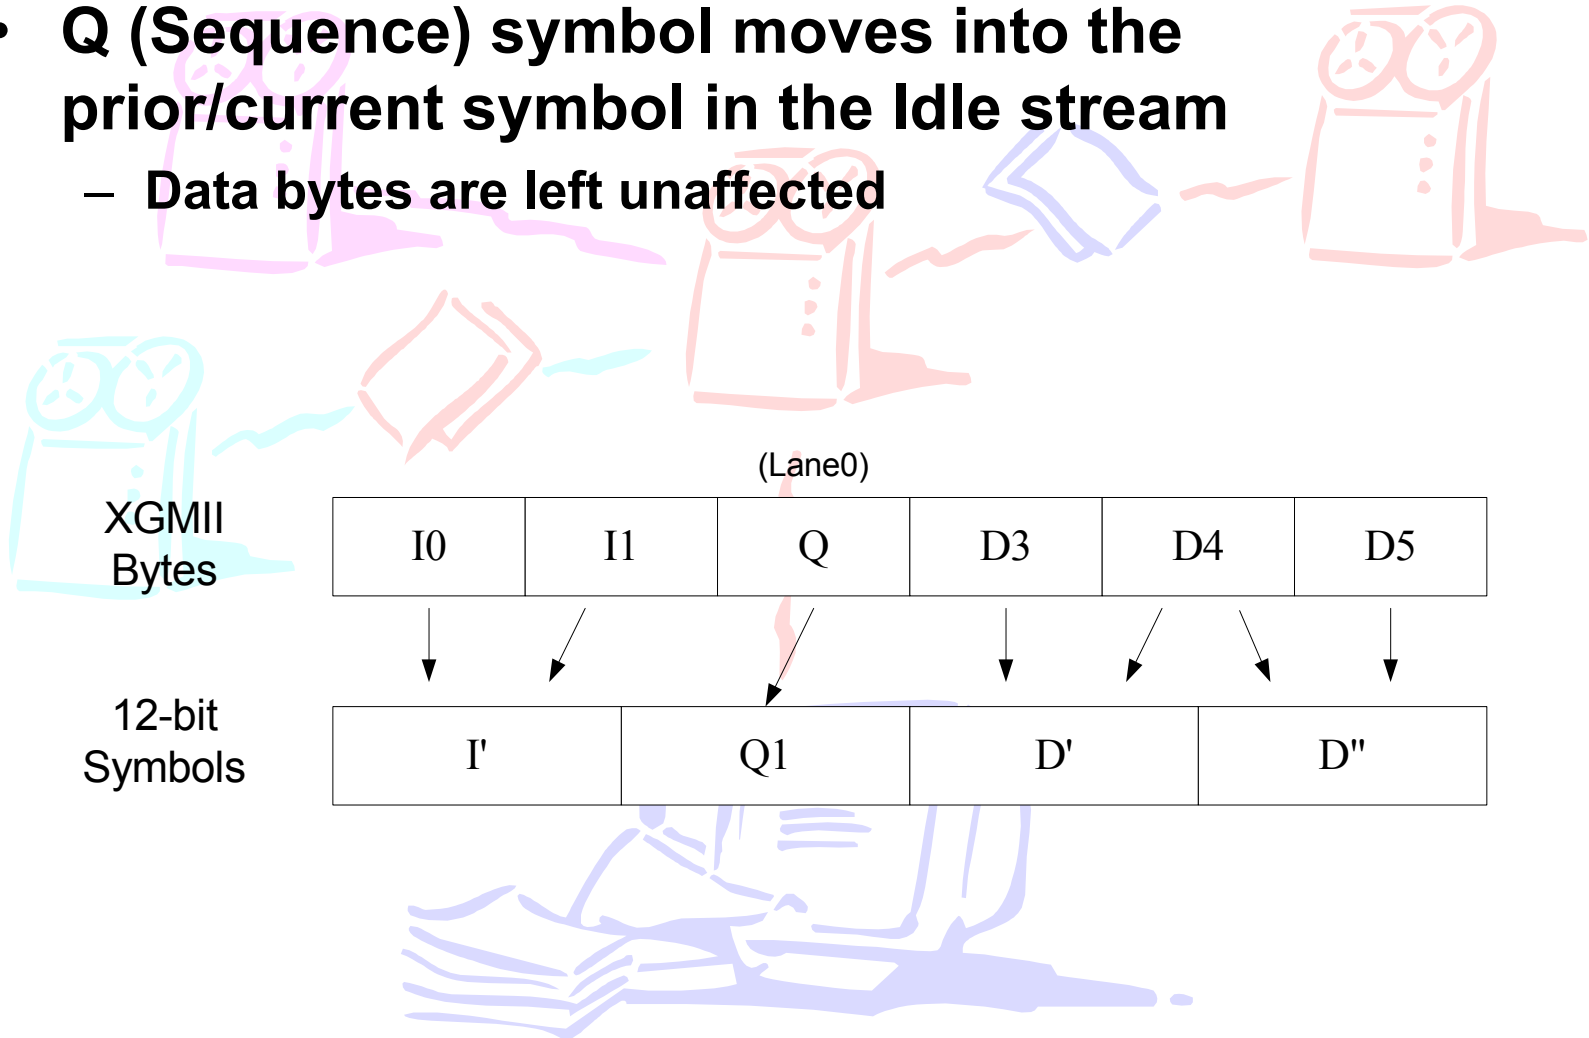
XGMII
Bytes

12-bit
Symbols



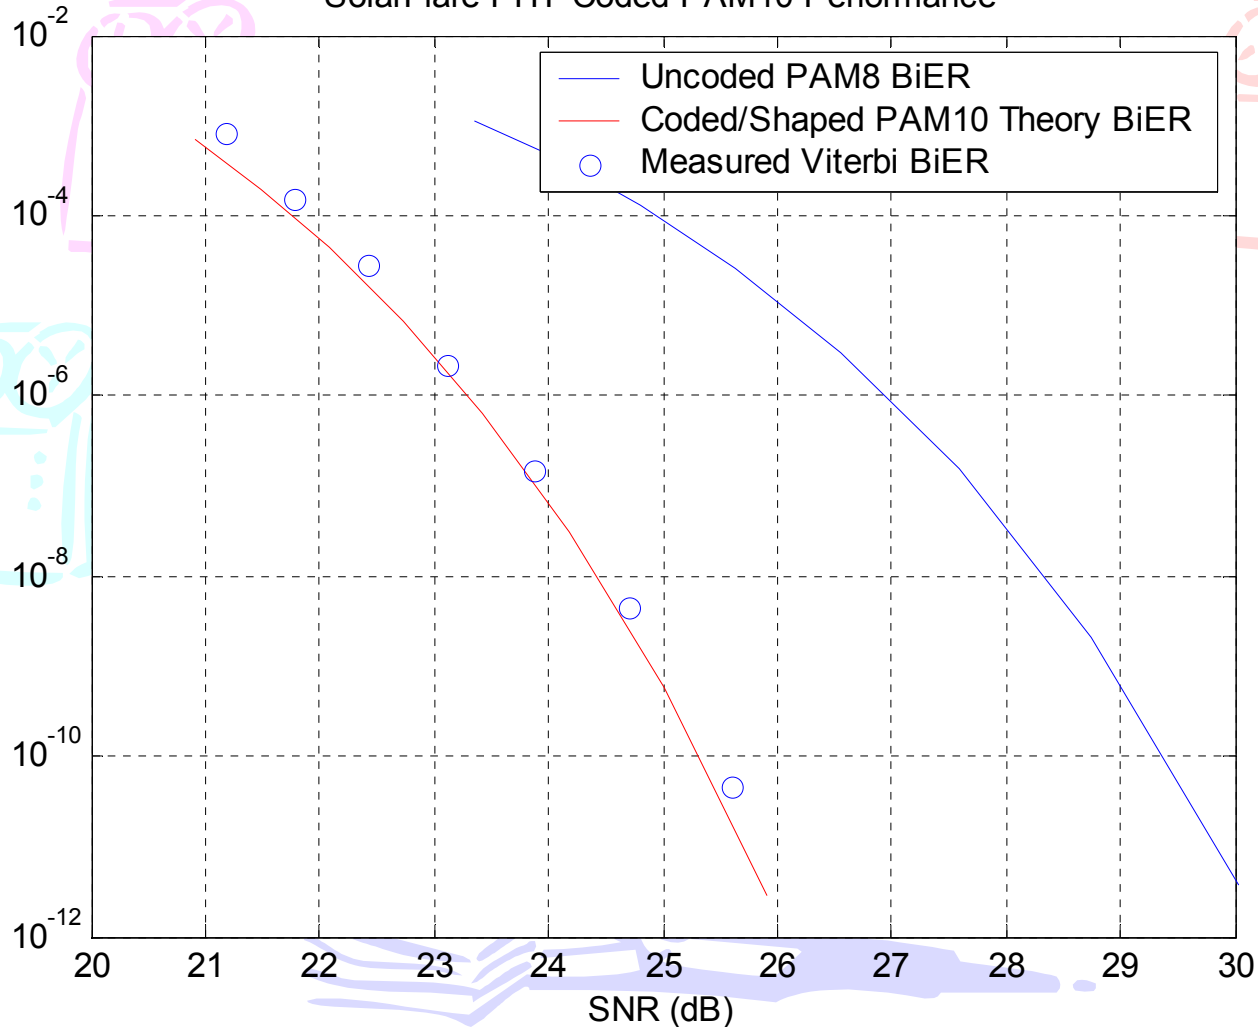
Ordered Sets

- **Q (Sequence) symbol moves into the prior/current symbol in the Idle stream**
 - Data bytes are left unaffected



4D 8-State PAM10 Performance (Simulations: AWGN)

SolarFlare PHY Coded PAM10 Performance



Relative Complexity Estimates

- **Baseline Complexity: 4.7M gate equivalents 90nm**
 - Based on existing (.13u), proven hardware design for 10GBASE-T on existing Cat6 for this proposal
 - Includes complexity for impairments from cabling configurations in the real world, but not in standards environments
 - Includes several complexity reductions for signal processing elements (estimated 6X reduction)
- **Competitive Complexity:**
 - **Cancellers & equalizers:**
 - Scale linearly with baud rate (optimistic), pessimistic would scale as the square of baud rate
 - Scale linearly with data word length – optimistic “effective symbol” scaling (+2 bits) assumed for TH precoding, proposal would scale as output symbol (+6 bits)
 - **Coding: LDPC estimates from seki_1_0304, slide 4 (2-3M gates); RS complexity from Ungerboeck verbal (500k gates)**

Performance Comparisons, Model 1

- No ANEXT Cancellation, 100m Cat7 IL, Baseline ANEXT (62.5 dB at 100 MHz, split slope), Default Cancellation parameters

Uncoded DFE Margin Peak	-0.78	dB	770	MHz					
Proposal	1e-12 SNR reqd	Bits/ sym/ pair	Ovrhd for Ctl	Baud Rate (MHz)	Opt DFE SNR at rate	Margin	dB from max	dB from DFE peak	Relative Complexity
SF 4D-4W TCM PAM10	26.2	3	0.00%	833	29.36	3.16	-3.31	-0.12	1.000
Keyeye PAM-4	24	2	3.13%	1289	20.15	-3.85	-10.3	-3.1	0.773
Rao_Nov2003	19.9	2.683	7.30%	1000	25.52	5.62	-0.85	-2.01	1.875
Rao_Mar2004	19.9	2.683	3.13%	961	26.37	6.47	0	-1.16	1.817
Ter. 12-PAM (2048,1723)	23.8	3.182	5.00%	825	29.58	5.78	-0.69	-1.01	1.860
Ter. 12-PAM (1024,781)	22.8	3.029	3.00%	850	28.95	6.15	-0.32	-0.71	1.905
Powell 4DTCM+RS(2.66b/syn	21.4	2.66	3.13%	969	26.2	4.8	-1.67	-1.19	1.807
Ung. 4DTCM+RS(2.5775b/s)	19.9	2.578	3.13%	1000	25.52	5.62	-0.85	-1.36	1.863
Plato PAM-5	20.5	2	0.00%	1250	20.8	0.3	-6.17	-2.45	1.125

Performance Comparisons, Model 3

- No ANEXT Cancellation, 100m Cat6 IL, Baseline ANEXT (64.5 dB at 100 MHz, split slope), Default Cancellation parameters

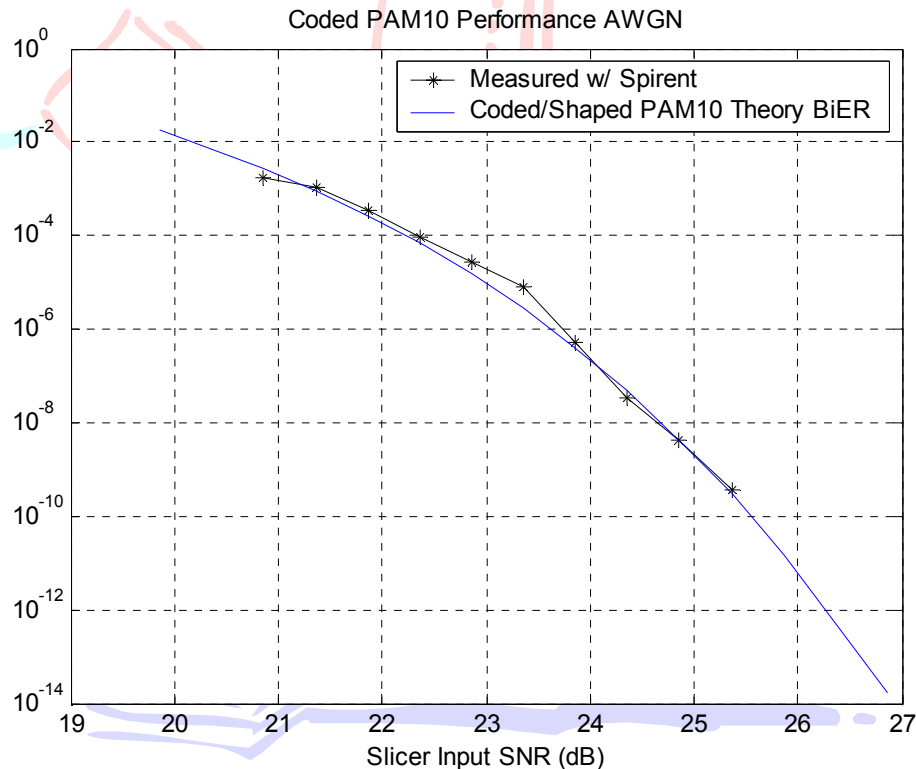
Standards Code Performance Summary									
Max Margin	6.07 dB								
Uncoded DFE Margin Peak	-0.87 dB		740 MHz						
Proposal	1e-12 SNR _{req}	Bits/sym/pair	Ovrhd for Ctl	Baud Rate (MHz)	Opt DFE SNR at rate	dB Margin	dB from max	dB from DFE peak	Relative Complexity
SF 4D-4W TCM PAM10	26.2	3	0.00%	833	29.13	2.96	-3.11	-0.2	1.000
Keyeye PAM-4	24.0	2	3.13%	1289	19.67	-4.33	-10.4	-3.3	0.770
Rao_Nov2003	19.9	2.68	7.30%	1000	25.07	5.17	-0.9	-1.22	1.870
Rao_Mar2004	19.9	2.68	3.13%	961	25.97	6.07	0	-0.92	1.820
Ter. 12-PAM (2048,1723)	23.8	3.18	5.00%	825	29.36	5.56	-0.51	-0.18	1.860
Ter. 12-PAM (1024,781)	22.8	3.03	3.00%	850	28.7	5.9	-0.17	-0.3	1.910
Powell 4DTCM+RS(2.66b/sym)	21.4	2.66	3.13%	969	25.79	4.39	-1.68	-1	1.810

PHY Parameter Highlights

- **ADC: 10 bits, 9 ENOB**
- **DAC: 10 bits, 9 ENOB**
- **Jitter: 4 psec RMS**
- **TX Power: 6 dBm (2.5-3Vpp for PAM-10, depends on measurement point)**
- **TX PSD: -80 dBm/Hz peak, with PAM spectrum**
- **Intrinsic latency: < 20 nsec**
- **TX equalization: none**
- **Margins w/noise, impairments & limits**
 - Models 1,2,3 2.6, 2.9, 2.1 dB
- **Power estimate (90nm): 8.1W**

Measured Performance

- Laboratory IC transceiver results on Cat5e cabling
- Includes Jitter, A/D, EMI Ingress, and additive noise source



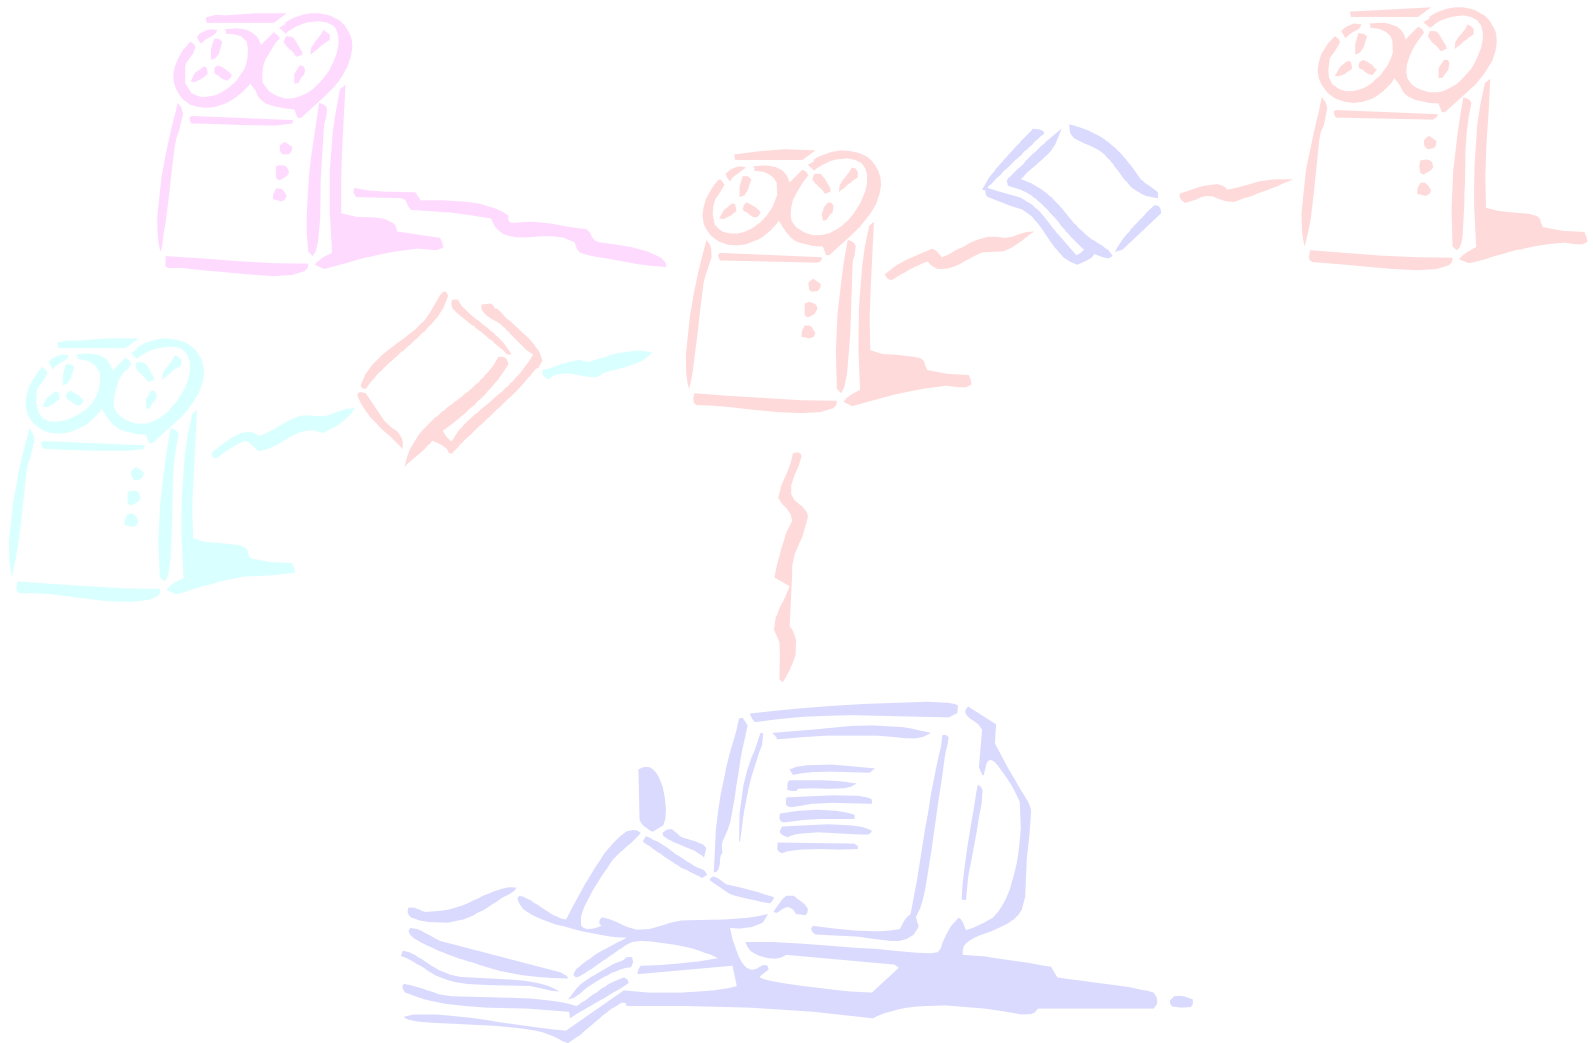
Conclusion

- **PAM-10 833 Mbaud line coding proposal is based on proven 1000BASE-T technology with extensions to meet performance goals**
- **PAM-10 833 Mbaud 3 bits/baud proposal will meet the reach requirements**
- **PAM-10 833 Mbaud 3 bits/baud proposal balances complexity with performance**
- **PAM-10 833 Mbaud 3 bits/baud proposal has intrinsic latency to scale to meet future needs**
- **PAM-10 833 Mbaud 3 bits/baud proposal is the only complete solution to date with complexity proven in hardware**

Motions

- **Motion #1:** Move that the 802.3an Task Force adopt a 10-level pulse amplitude modulation (PAM10) with a symbol rate of 833 Mbaud for D1.0 of Clause 55 as presented in (mcclellan_1_0504.pdf).
 - Moved By:
 - Seconded By:
 - Yes: No: Abstain:
- **Motion #2:** Move that the 802.3an Task Force adopt the 4D 8-State PAM10 TCM and 4D PAM10 Mapping for D1.0 of Clause 55 as presented in (mcclellan_1_0504.pdf).
 - Moved By:
 - Seconded By:
 - Yes: No: Abstain:
- **Motion #3:** Move that the 802.3an Task Force adopt the XGMII Frame Structure and Control for D1.0 of Clause 55 as presented in (mcclellan_1_0504.pdf).
 - Moved By:
 - Seconded By:
 - Yes: No: Abstain:

Backup



4D PAM10 Partitioning

- 1D values: { -9, -7, -5, -3, -1, +1, +3, +5, +7, +9 }
- This set can be split into two subsets:
 $X = \{ -9, -5, -1, +3, +7 \}$
 $Y = \{ -7, -3, +1, +5, +9 \}$
- 2D 10X10 constellation:
 { (-9,-9), (-9,-7),, (+9,+7), (+9,+9) }
- We can divide this constellation into 4 subsets:
 $XX = \{ (-9,-9), (-9,-5),, (+7,+3), (+7,+7) \}$
 $XY = \{ (-9,-7), (-9,-3),, (+7,+5), (+7,+9) \}$
 $YX = \{ (-7,-9), (-7,-5),, (+9,+3), (+9,+7) \}$
 $YY = \{ (-7,-7), (-7,-3),, (+9,+5), (+9,+9) \}$

4D PAM10 Partitioning cont.

- **4D 10X10X10X10 constellation:**
 $\{ (-9,-9,-9,-9), (-9,-9,-9,-7), \dots, (+9,+9,+9,+7), (+9,+9,+9,+9) \}$
- **We can divide this constellation into 16 subsets:**
 $XXXX = \{ (-9,-9,-9,-9), (-9,-9,-9,-5), \dots, (+7,+7,+7,+3), (+7,+7,+7,+7) \}$
 $XXXY = \{ (-9,-9,-9,-7), (-9,-9,-9,-3), \dots, (+7,+7,+7,+5), (+7,+7,+7,+9) \}$
.....
 $YYYX = \{ (-7,-7,-7,-9), (-7,-7,-7,-5), \dots, (+9,+9,+9,+3), (+9,+9,+9,+7) \}$
 $YYYY = \{ (-7,-7,-7,-7), (-7,-7,-7,-3), \dots, (+9,+9,+9,+5), (+9,+9,+9,+9) \}$
- **While the minimum distance between the points in the original 10X10X10X10 constellation was $d_{\min}=2$, the distance between points within any of the sixteen subsets is $d_{\min} = 4$.**



4D PAM10 Partitioning cont.

- For the 4D 8-state trellis code we need only 8 sublattices. Therefore, we combine pairs of the subsets shown above to get 8 subsets, D0 to D7:

$$D0 = XXXX + YYYY$$

$$D4 = XYYX + YXXY$$

$$D1 = XXXY + YYYY$$

$$D5 = XYYY + YXXX$$

$$D2 = XXYY + YYXX$$

$$D6 = XYXY + YXYX$$

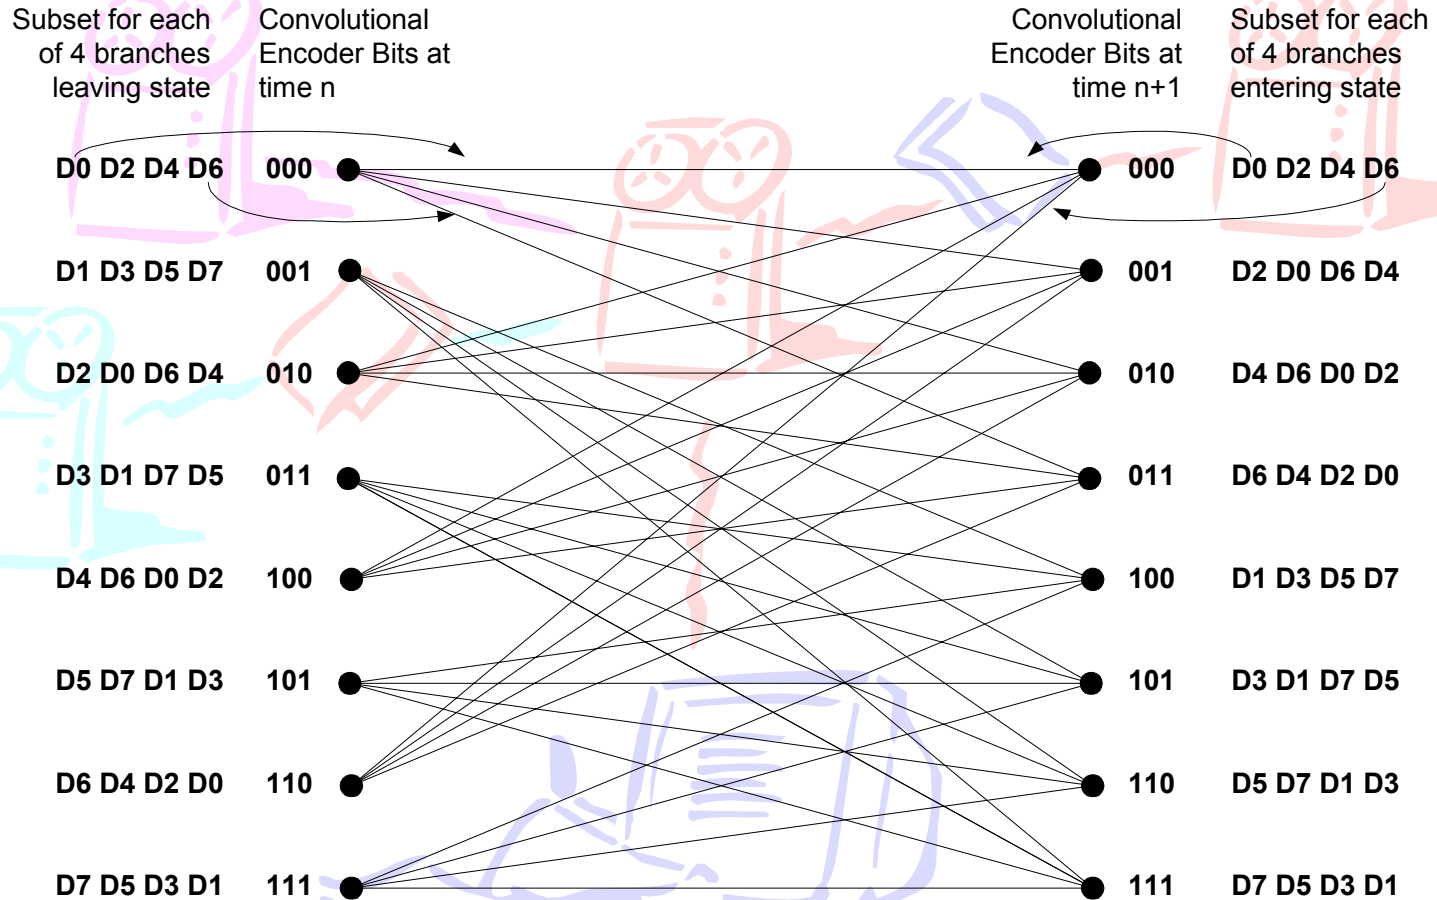
$$D3 = XXYX + YYXY$$

$$D7 = XYXX + YXYY$$

- Each subset contains 1250 points and with this grouping the minimum distance between points within any subset is still $d_{\min} = 4$.



4D 8-State TCM w/ Partitioning



Trellis Diagram for 4D 8-State Trellis Code