VITESSE

1000BT 10GBT Technology Review and Application to SAS

Vitesse - Ethernet Products Division

Houston SAS-2 Meeting

Doc # 05-215r0

Nick van Bavel
VP of Technology

nvb@vitesse.com

YOUR PARTNER FOR SUCCESS

Outline

VITESSE

- Review of 100B-T
- Review of 1000B-T
- Review of 10GB-T
- Application of 10GB-T Technology to SAS

100B-T System Topology

VITESSE

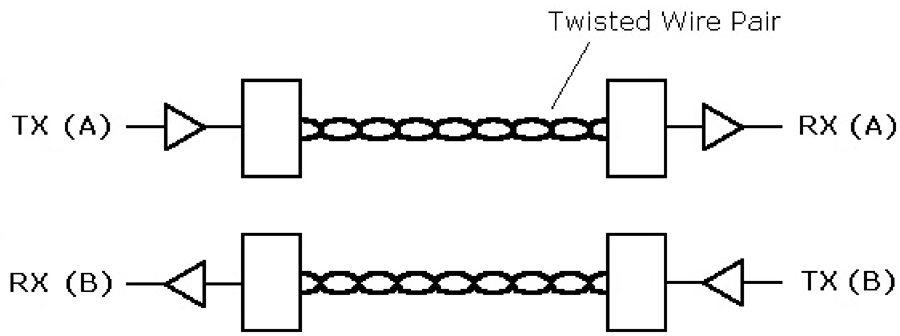

- Target Media: 100m of Cat5
- Pair Usage: 2 Pairs Half-Duplex
- Line Code: MLT3
  - Tri-Level (-1,0,+1) with Redundancy for Spectrum Control
  - 1 Bit per Symbol
- Line Rate: 125 MHz
- Encoding: 4B/5B
  - Run-length Limiting to Aid in Timing Recovery
  - Adds 25% Time Redundancy

1000B-T System Topology

VITESSE

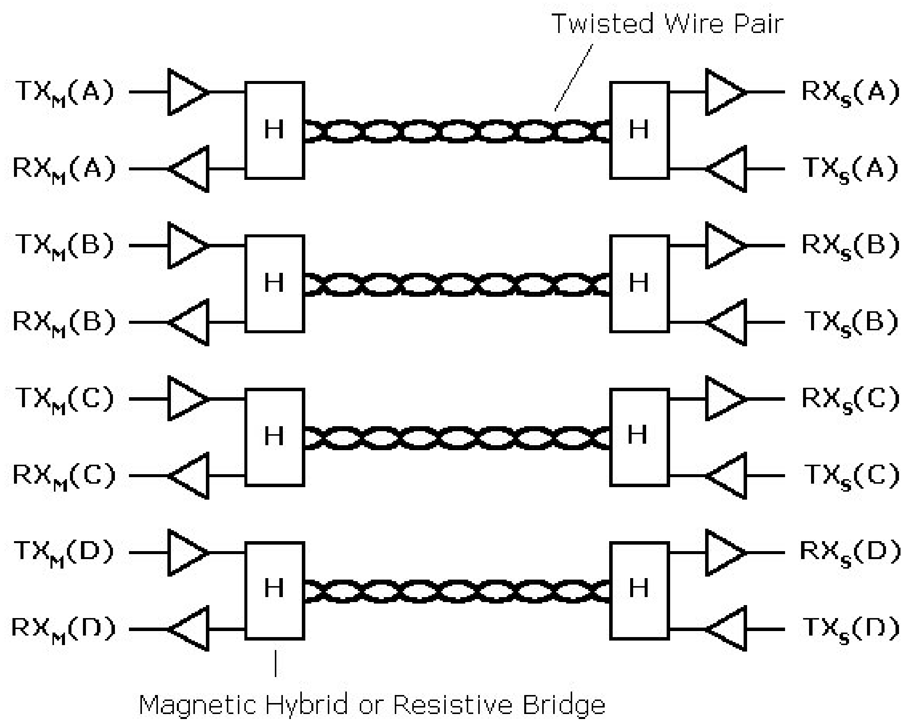

1000B-T System Topology

VITESSE

Target Media:	100m of Cat5e	(Cat5e specs FEXT)
Pair Usage:	4 Pairs Full-Duplex	(2x Pairs, 2x Duplex)
Line Code:	PAM-5	(2x Bits/Symbol)

- Five-Level (-2,-1,0,+1,+2)
- 2 Bits per Symbol
- Requires Higher SNR!

Line Rate:	125 MHz	(Same)
Encoding:	4-D Viterbi Code	(Gains 1.25x)

- No Time Redundancy = > Low Bandwidth PLL for Timing Recovery
- Takes Advantage of Amplitude Redundancy over 4 Pairs
- (5^4) > 2 x (2^2^4) gives Coset Separation
- 6 dB of Coding Gain

1000B-T Transceiver Specifications

VITESSE

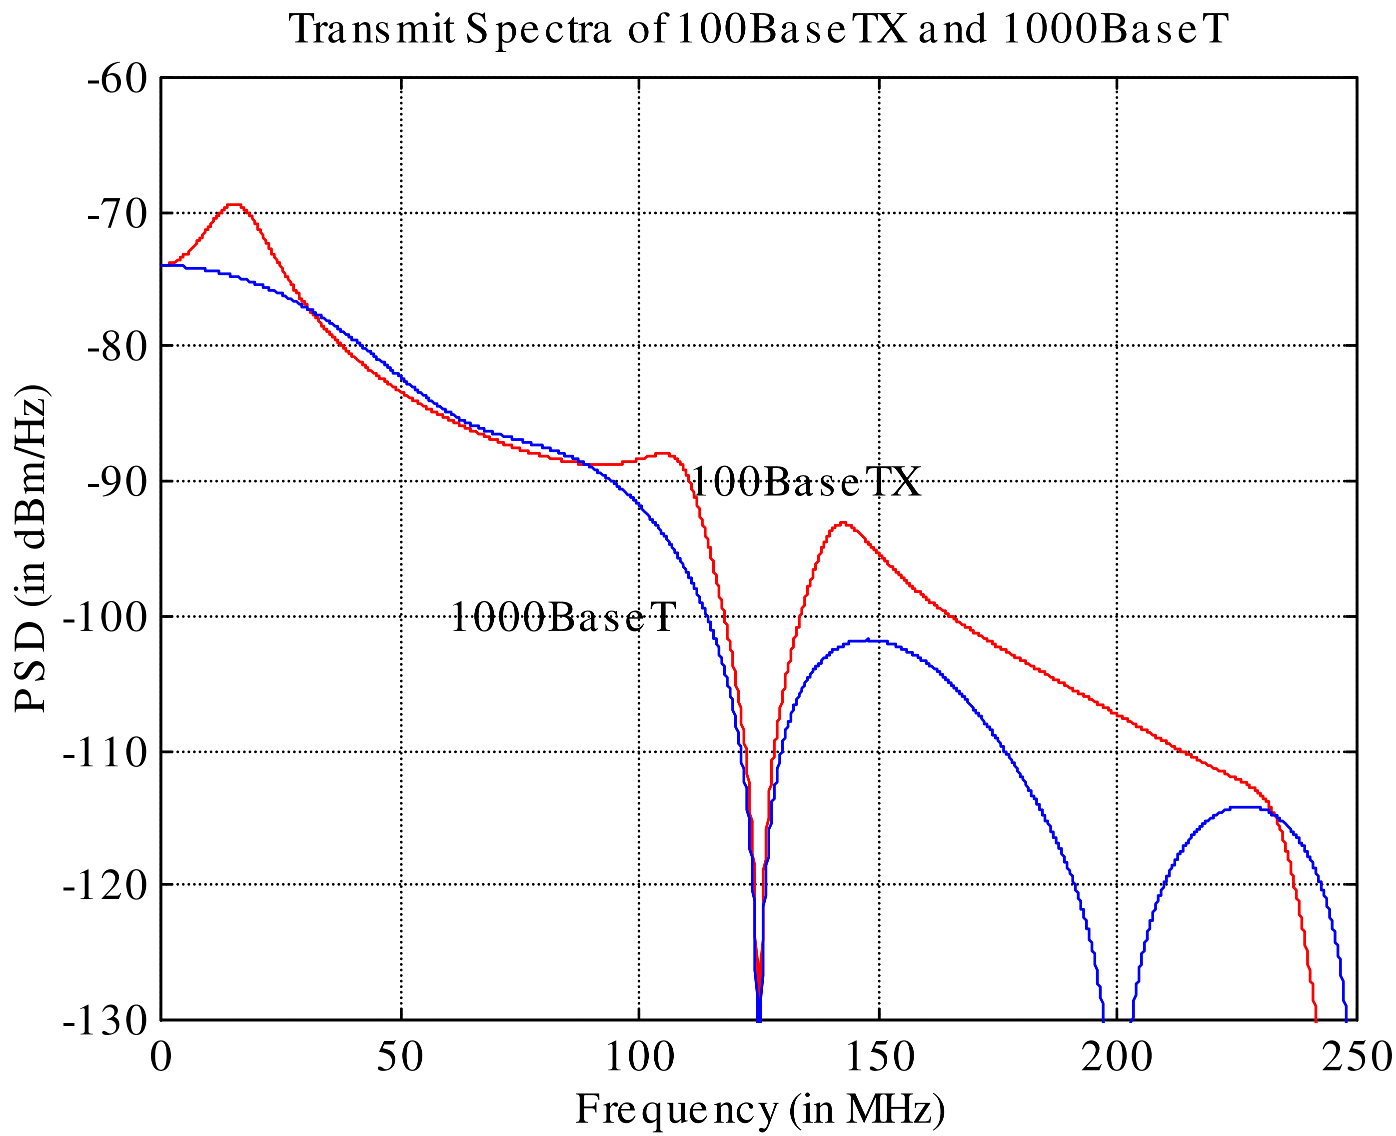

1000B-T System Impairments

VITESSE

- Cable Attenuation: 24dB @ 100MHz
- Reflections due to Impedance Mismatch (Echo)
- Near-End CrossTalk (NEXT)
- Far-End CrossTalk (FEXT)
- Noise Ingress from adjoining cabling (Alien NEXT/FEXT)

1000B-T Noise Environment

VITESSE

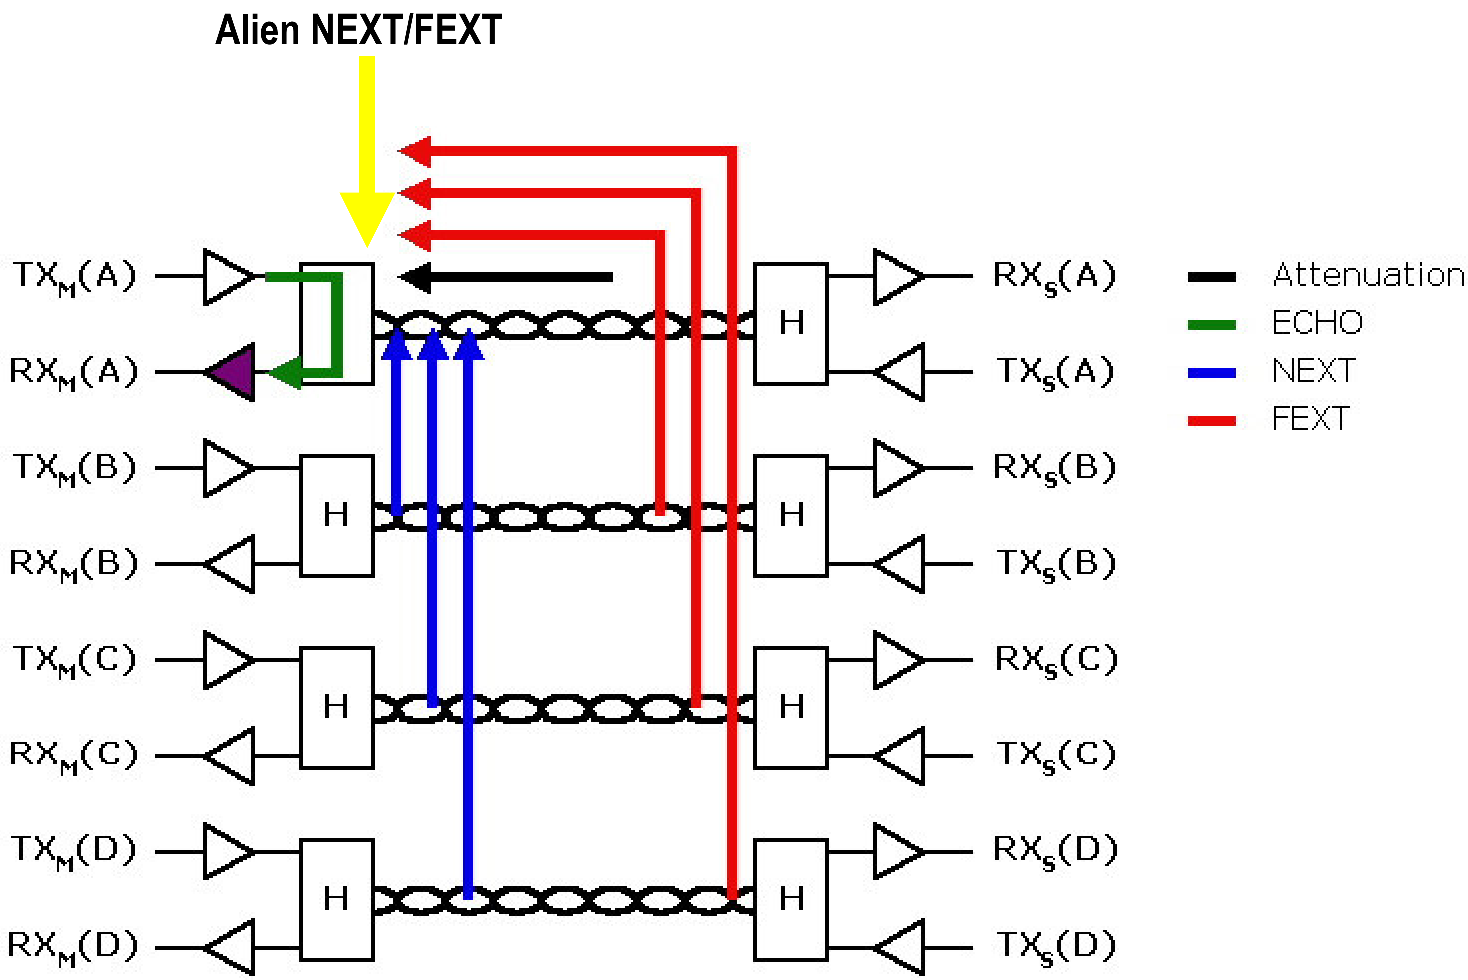

1000B-T Data Pump

VITESSE

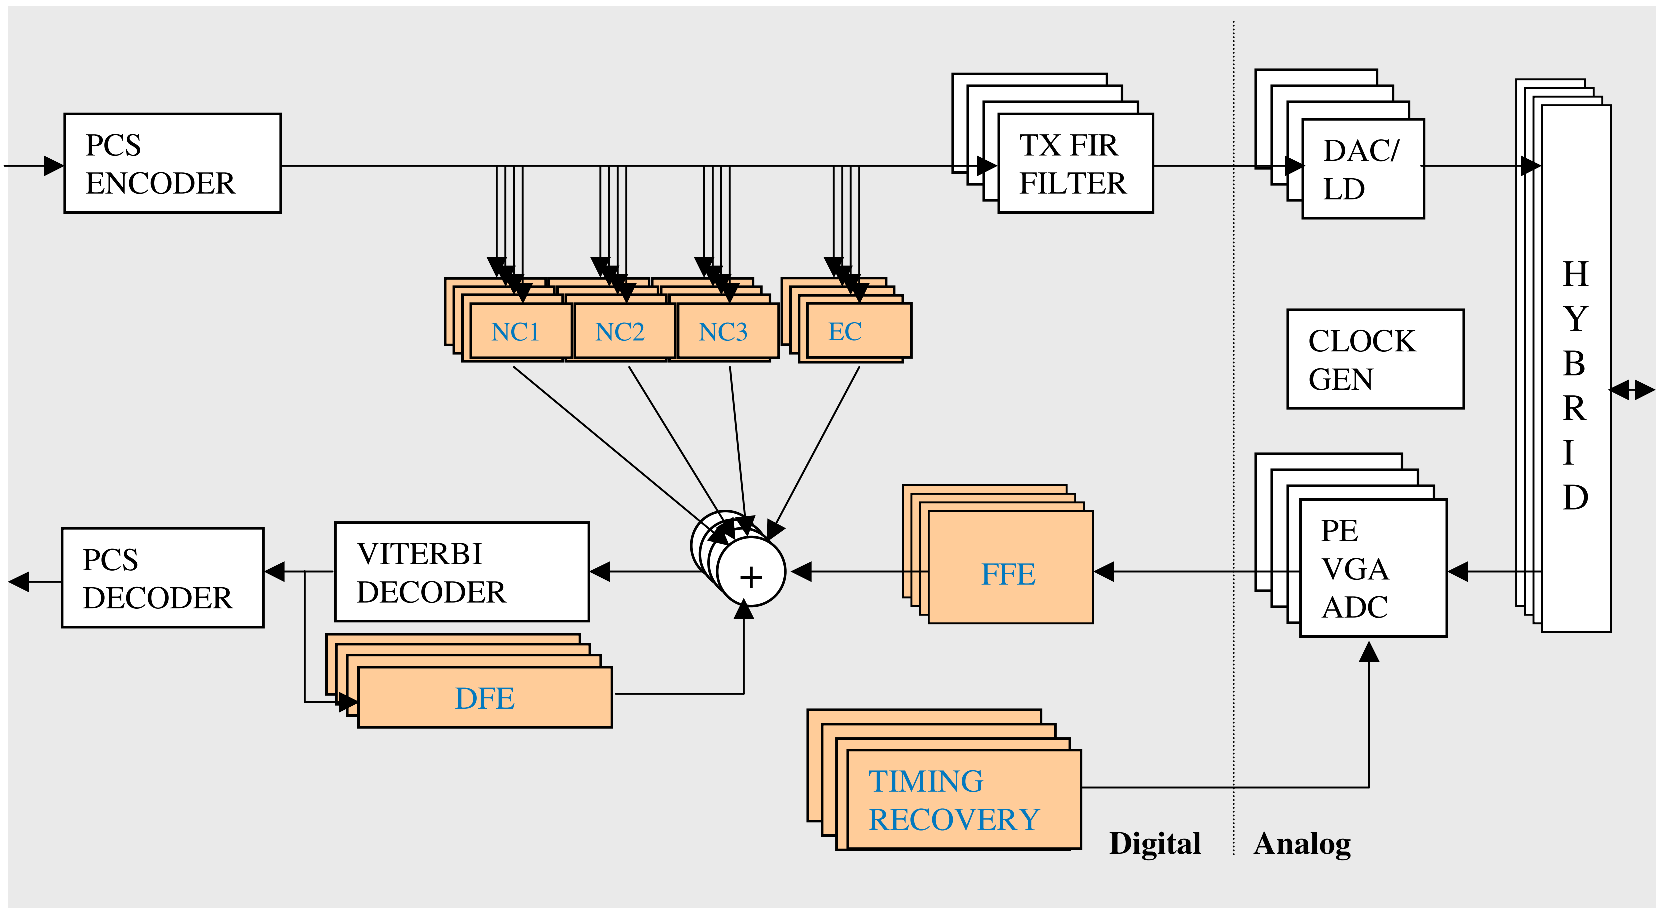

= Adaptive Filter

10GB-T System Topology

VITESSE

- Target Media: 100m of Cat7
  55m – 100m of Cat6 Augmented
- Pair Usage: 4 Pairs Full-Duplex (Same)
- Line Code: DSQ-128 + LDPC (1.5625x)
  - Two-Dimensional 128 Points (Equiv ~11 PAM)
  - 3.125 Bits per Symbol per Pair
  - Requires Higher SNR!
- Line Rate: 800 MHz (6.4x)
  - Increases Signal Attenuation by >30 dB
- Encoding: (2048,1723) Block LDPC (Same)
  - No Time Redundancy = > Low Bandwidth PLL for Timing Recovery
  - Requires THP (Tomlinson-Harashima Precoding)
  - Provides 9 dB of Coding Gain

10GB-T vs 1000B-T Implementation Complexity

VITESSE

- DSP
  - 7-10x Number of Adaptive Filter Taps
  - Must Have FEXT Cancellation
  - Running at Equivalent of 6.4x Speed
  - Necessitates Block Signal Processing
    - FFT's vs Trivial Time Domain Filters
    - Latency of LDPC Decoder plus Block Processing
    - Latency Spec'd at max 10usec
- AFE
  - 10B ENOB ADC at 800MHz vs. 8B at 125MHz
  - 6B DAC at 800MHz vs. 2B at 125MHz
  - Bandwidth of Hybrid, Analog Equalizer, AGC increased by 6.4x
  - Required TX and RX Jitter is 3-5ps rms vs 100-200ps

10GB-T vs 1000B-T Implementation Cost

VITESSE

	100B-T	1000B-T	10GB-T
Digital/DSP Gate Count	50k (1x)	500k (10x)	10M (200x)
Digital Power (90nm)	20 mW	150 mW	7.0 W
AFE Area	1x	2x	5x
AFE Power	150 mW	350 mW	7.5 W
Total Power (90nm)	170 mW	500 mW	14.5 W *

* Note: Power numbers are typical only, and are optimistic: other vendors estimates run as high as 20W in 90nm

SAS on UTP?

VITESSE

4x Infiniband Cables and Connectors are:

- Expensive (e.g., 8m is > $200 retail)
- Bulky & Cumbersome
- Limited Reach

  (10m max without integrated equalization)

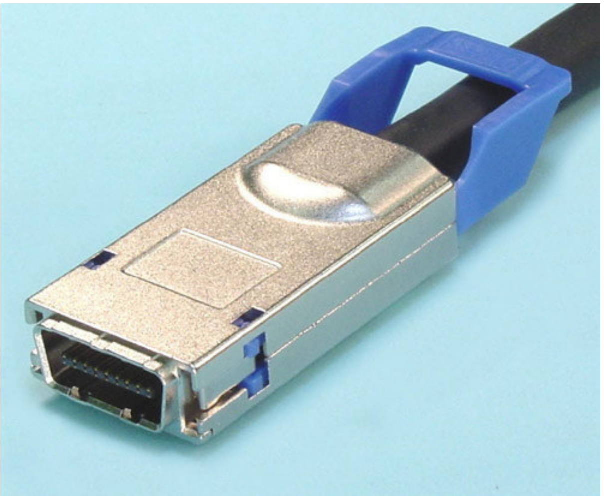

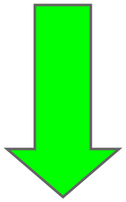

Solution:

Define a new PHY for external SAS connections

- Leverage Cat6 Ethernet Cables and Connectors:
- Cheap (e.g., 8m is <$10.00 retail)
- Flexible & Lightweight
- Reach of 25m at 10G

Problems with 10GB-T Technology for SAS

VITESSE

- Problems for SAS
  - 15w power dissipation probably unacceptable
  - 10usec latency probably unacceptable
  - Cost of 10M gates probably too high
- However 100m Reach is Overkill
- Can We Leverage 10GB-T Technology for Shorter Reach?
- Target New Specs
  - Less Than 4W Power Dissipation
  - 25m reach
  - Sub 500 nsec latency
  - Assume still have to be spectrally compatible with 10GB-T

25m CAT6 UTP results

VITESSE

- Performance gain
  - SNR only 4dB better than 55m CAT6 or 100m CAT6 Screened UTP due to
    - Power Back off to coexist in the same bundle as 10GB-T long reach
    - Cross talk from higher power 10GB-T 55m links
- Complexity and Cost savings
  - 4dB SNR can be used for
    - 7-bit ADC and full power LDPC decoder (from 9-bit))
    - 8 bit ADC and power/area optimized LDPC decoder
  - 10dB Power Back Off saves transmitter power
  - Overall AFE power reduced by factor of 2
  - Reduced DSP filter taps by factor of 2
- Latency
  - Latency gets reduced to 1.5us – 2usec

Conclusion

VITESSE

- Long Reach Power = ~14.5W (Typical, not worst-case)
- Short Reach Power = ~ 7.0W (Typical, not worst-case)
- Did not meet power target of 4W
- Latency Reduced but still above target
- Cost is still prohibitive (10x 1000B-T)

Other Ideas?

VITESSE

- What if targeted 25m of screened Cat6?
  - Reduces Alien NEXT/FEXT, SNR margin improves by another 6.5 dB
  - Still can't get rid of the LDPC decoder (b/c of the noise averaging on four pairs)
  - But could further reduce power by ~2W; still not there
- What if the PHY architecture was optimized for 25m on UTP Cat6?
  - Constrained to be spectrally compatible with 10GB-T
  - May be able to transmit at higher baud rate
  - Definitely could further reduce power/latency; merits investigation
- What if the PHY architecture was optimized, but also targets 25m of screened Cat6?
  - Definitely opens design space to very low-power, low-latency solutions

Conclusion

VITESSE

- Conclusion:
- 10GB-T is not a PHY readily adapted for SAS applications
- The architecture is not optimized for short reach
- Other architectures optimized for short reach on UTP/STP may be able to hit cost/power/latency specifications for SAS

Comments / Questions ?