Who's afraid of the Big, Bad DSP?

Implementation of Advanced Modulation for Single-Lambda Short-Reach Optical Interconnects

Paul Little
Fujitsu Semiconductor America, Inc.

Ethernet Alliance 100GbE per Lambda for Data Center Workshop – June 12th 2014

100G Single-Lambda for Datacenter

FUJITSU

- Why bother with 100G Single-Lambda?
  - Because network capacity isn't going down…
  - Reduction of optical component count needed to reduce cost
  - Cost-effective building block needed for next generation 400G and beyond

- Straw Polls from Norfolk IEEE 802.3 interim meeting indicated strong interest in single lambda solutions for 400GbE for 2km reach.

Strawpoll #5 (Chicago Rules)

For 2km duplex SMF 400GbE PMD, I believe the TF should select a proposal based on an effective bit rate per wavelength per direction of

a)	25G	1+1+3 = 5
b)	50G	16+17+18=51
c)	100G	23+24+30 = 77
d)	400G	4+4+2 = 10

Strawpoll #6 (Chicago Rules)

For 10km duplex SMF 400GbE PMD, I believe the TF should select a proposal based on an effective bit rate per wavelength per direction of

a)	25G	2+1+2 = 5
b)	50G	14+19+20 = 53
c)	100G	24+22+28=74
d)	400G	4+2+5 = 11

DSP for 100G Single-Lambda

- DSP can offer a cost effective solution to increase link capacity
  - 100G NRZ single lambda optics are a few years away …
  - QSFP28 target form factor for 100G and possibly 400G in the future
  - DSP needed to meet power, cost and TTM requirements
- DSP shifts the complexity into the silicon
  - CMOS DSP implementations are predictable, scalable and cost effective
  - Decades of experience using advanced modulation in other applications
- Significant achievements in DSP power reduction in past few years
  - Courtesy of coherent transceivers for long-haul optical transport
  - Power reduction already driving 100G Coherent Metro devices into CFP module
- Noise and thermal issues have already been solved for long haul
- "Big, scary" DSPs are already lurking in your home
  - xDSL, brought to you by DMT (the little modulation format that could…)

DSP for 100G Single-Lambda Options

- Options for Single Lambda
  - DP-QSPK
  - OFDM
  - PAM-8
  - **PAM-4**
  - **Discrete Multi-Tone (DMT)**
  - Others….?

- This presentation will focus on DSP using Discrete Multi-Tone (DMT) and PAM-4

Discrete Multi-Tone Quick Overview

FUJITSU

- DMT uses a series of uniformly separately subcarriers (e.g 256) to transmit data, where each subcarrier uses QAM-n modulation
- Level of modulation is adjusted based on the available SNR within each subcarrier to maximize the number of bits per symbol within that subcarrier
- Transmitter and Receiver communicate to determine optimal settings for transmission
- Flexible modulation on each of the subcarriers allows DMT to compensate for link impairments and achieve the best use of the available signal channel bandwidth and SNR
- Frequency domain calculations with efficient iFFT & FFT implementations
- Transmit data is assembled in the frequency domain before passing through an iFFT and high speed DAC
- Receive DMT data is captured using a high-speed ADC, converted back into the frequency domain via FFT where subcarrier amplitude and phase are mapped into received data

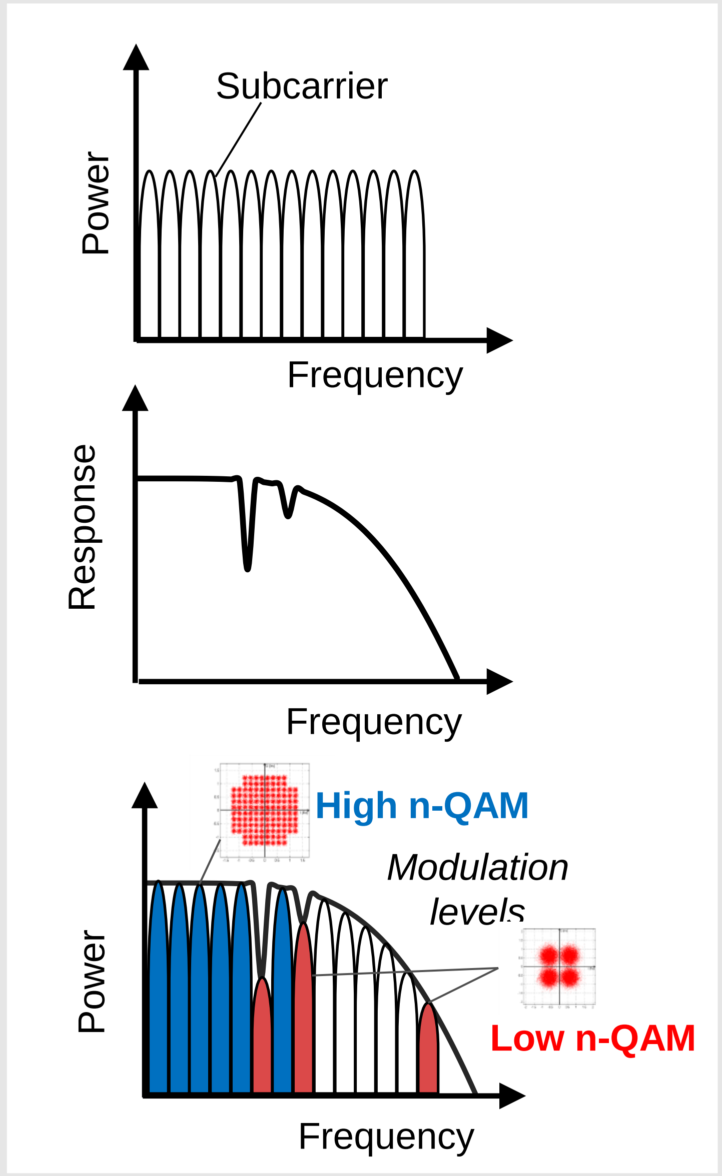

DMT Link Negotiation & Communication

FUJITSU

- 2 fixed location subcarriers are used to create a Link Communication Channel (LCC)
- DMT requires initial negotiation between the Tx and Rx on startup
- Tx sends known pilot tones (fixed constellation and bit sequence) for probing channel SNR via EVM across all 256 subcarriers
- Rx analyzes SNR and exchanges back bit allocations tables with Tx over LCC channel
- Target link negotiation time around 10ms

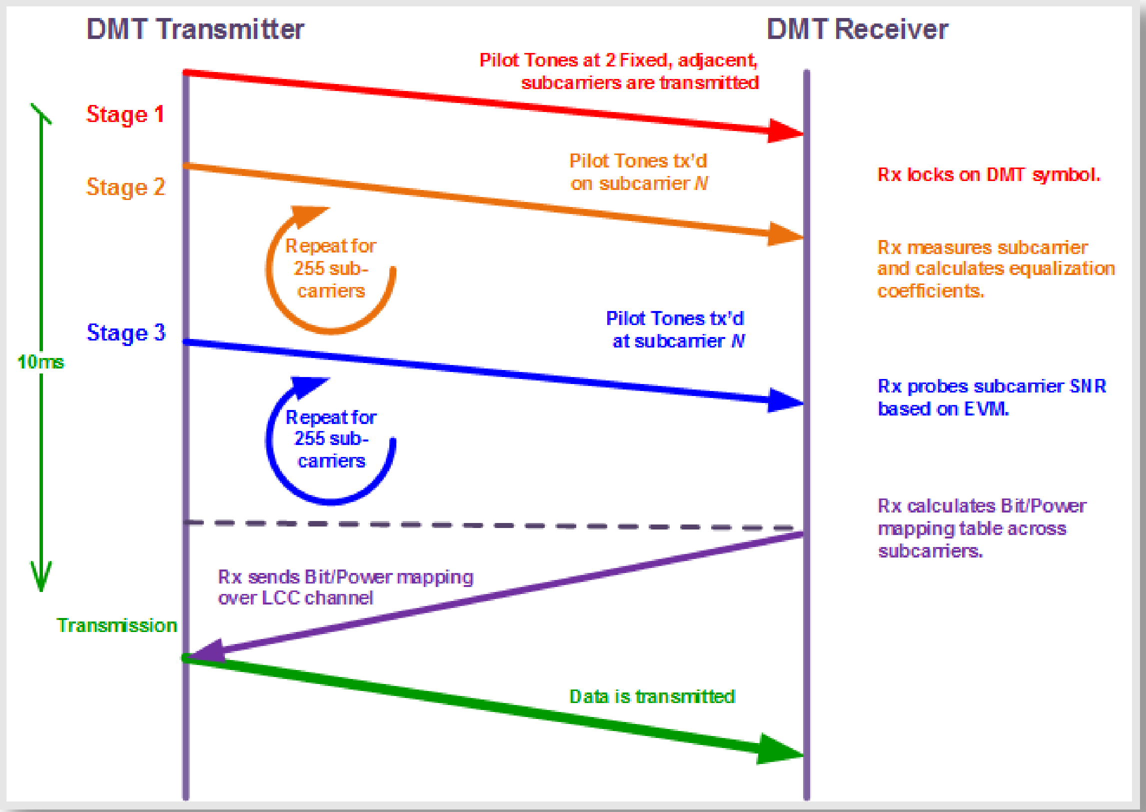

Advantages of DMT Technology for 100G

- Maximizes number of allocated bits for any given channel
  - High spectral efficiency (max bits for channel at low baud rate)
  - Reduces required bandwidth by stuffing the most bits into the lower frequencies
- Adaptive bit allocation for each subcarrier with low baud rate
  - Relaxes requirements on shape of channel frequency response
  - Doesn't need to be flat across the entire band
- Robust modulation scheme
  - More immune to channel reflections, ISI and crosstalk
- Comparable implementation complexity to that of xDSL (just a little faster!)
  - Excellent for IC development
  - Well known, efficient Digital Processing techniques (FFT, IFFT)
  - DMT Algorithms easily handled in non-realtime
  - Low power and latency

DMT Electrical B2B Measurements

FUJITSU

- Electrical B2B measurements made using existing ADC & DAC silicon
  - Channel probing shows ~6dB SNR improvement between 40nm and 28nm
  - Excess bitrate (>116Gbps) gives indication of relative margin in the system for other component losses

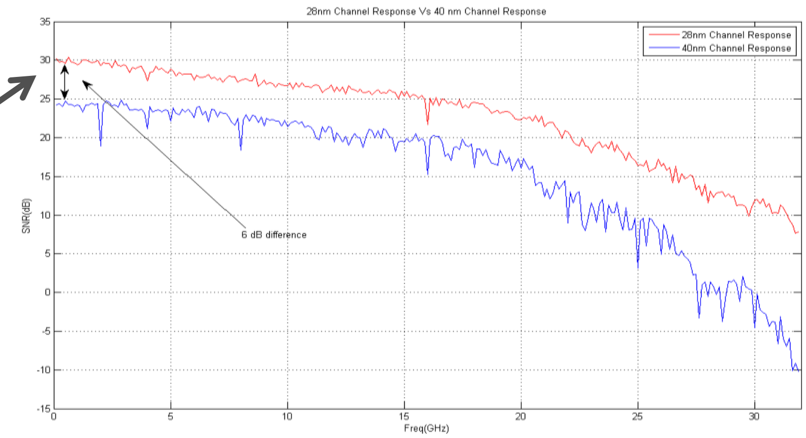

- 40nm Silicon Measurements
  - **136Gbps** throughput
  - 14% margin

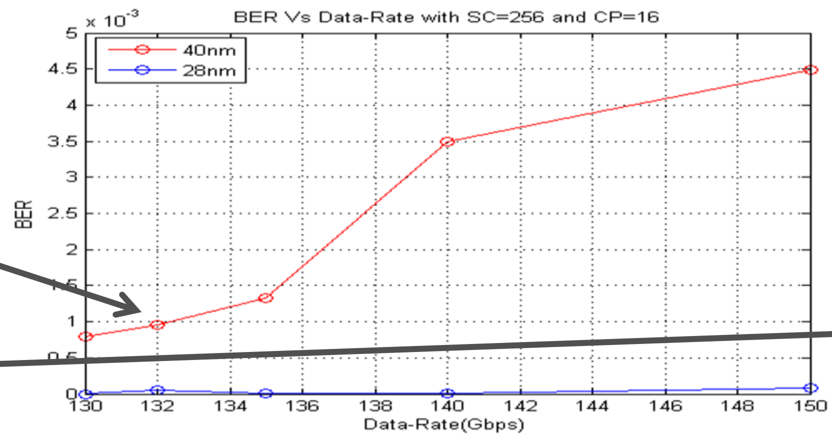

- 28nm Silicon Measurement
  - **170Gbps** throughput
  - 46% margin

- Initial optical experiments using 28nm converters presented, work in progress

Plots show 256 subcarriers, Cyclic Prefix of 16
Bit rate measured at FEC threshold of 1e-3

Optical Transmission Experiments (40nm)

FUJITSU

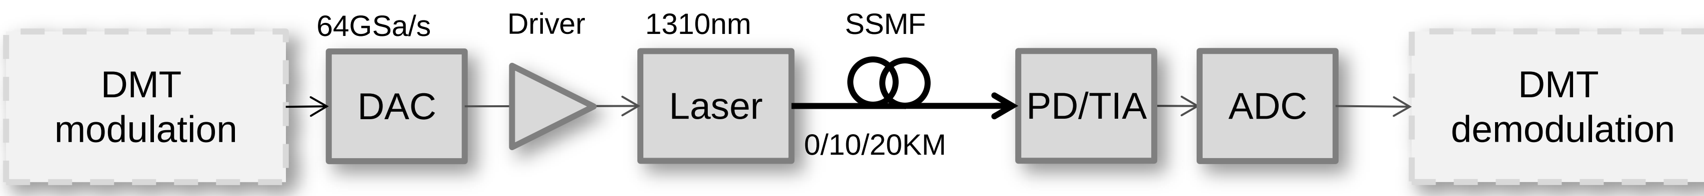

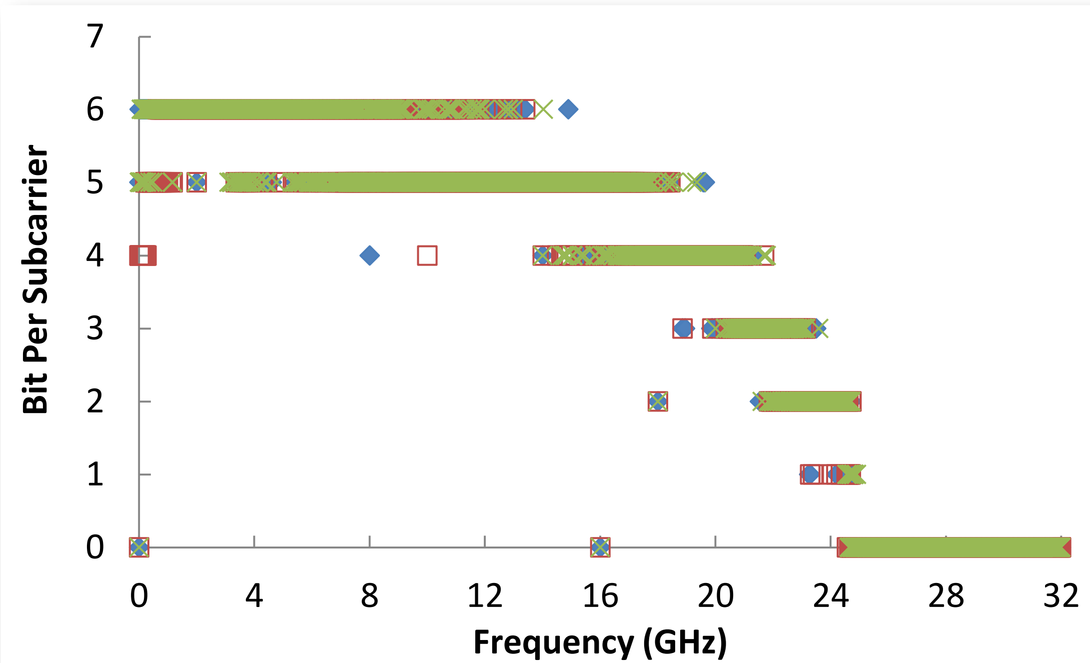

Distance	Capacity	BER
500m	116Gb/s	3.70E-04
10km	116Gb/s	3.56E-04
20km	116Gb/s	2.67E-04

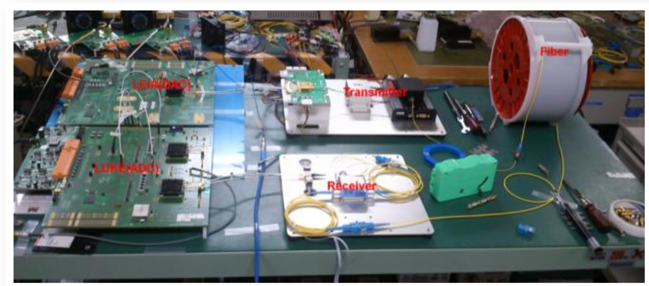

- Single DMT transceiver can cover multiple transmission distances and meet the input BER target for BCH FEC

100G DMT Transceiver IC Overview

FUJITSU

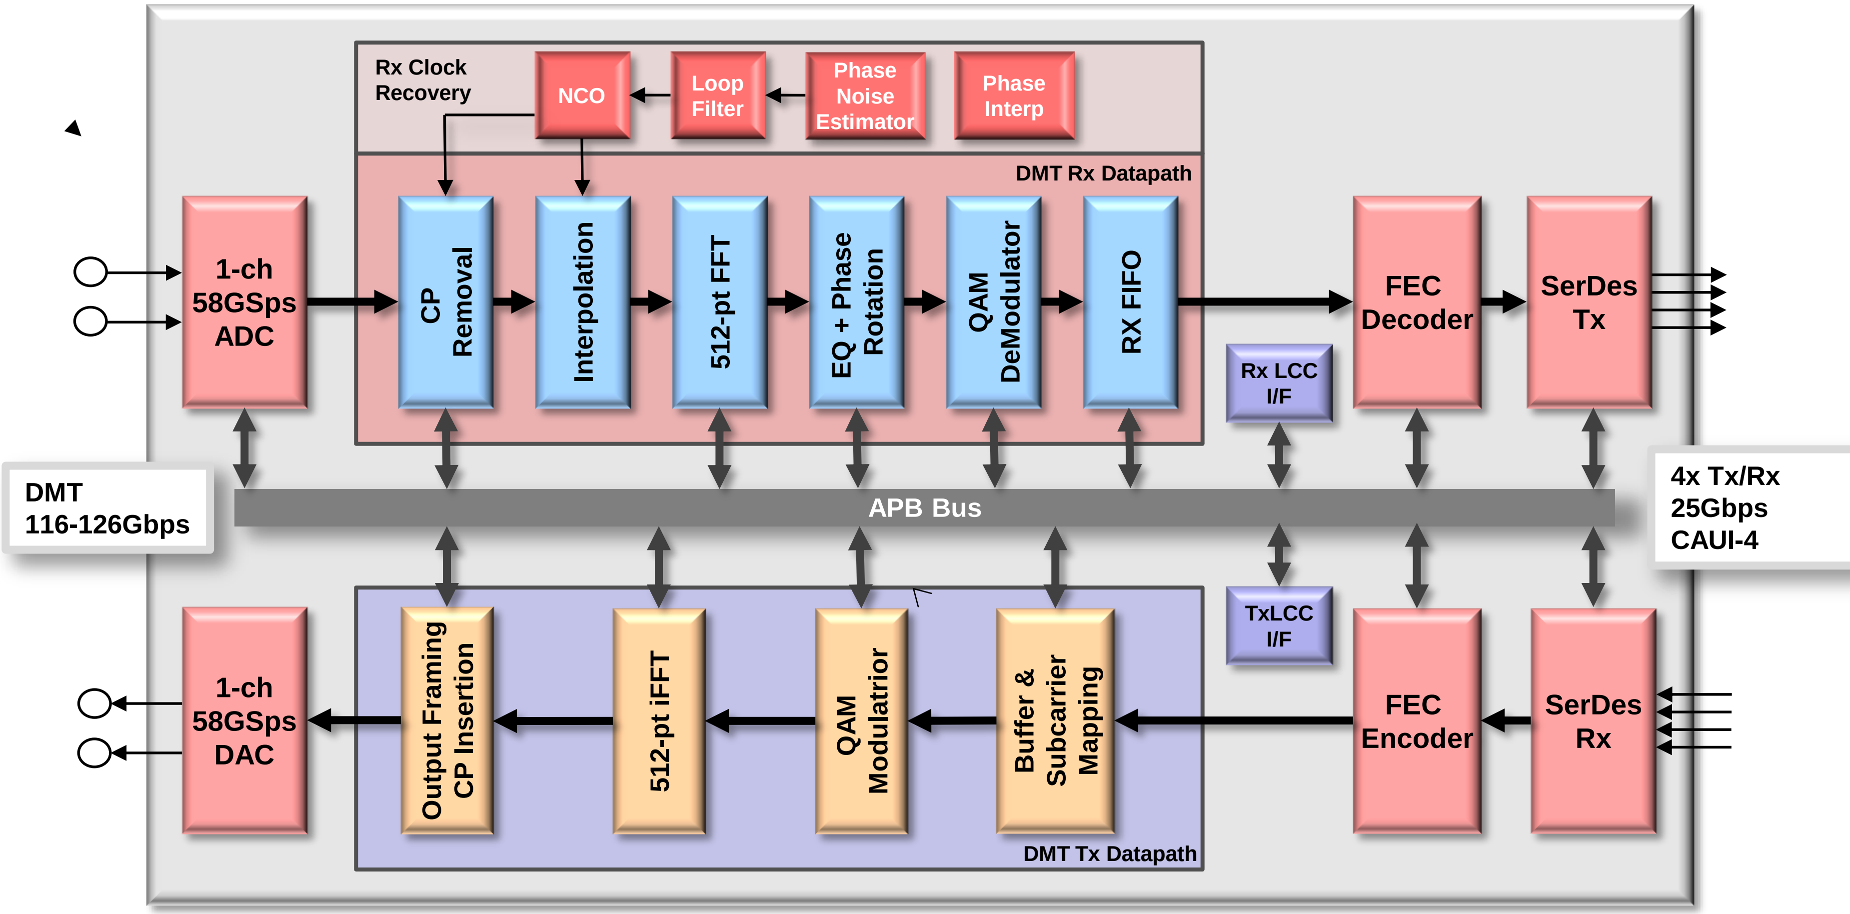

- Single chip integrated CMOS transceiver
  - Low cost, low loss 10mm x 10mm SHDBU package for QSFP28 module integration
  - 3.5W in 28nm (silicon measurement, synthesis and layout results) - CFP4 OK.
  - **1.8W in 14nm (estimated)**

DMT IC Complexity

Item		100G Coherent (DP-QPSK)	100G DMT	Notes
ADC	Channels	4	**1**	Direct detection
	Sample Rate	64GSa/s	58GSa/s	
	ENOB	>4.5	>5	ENOB/Bandwidth combination is key
DAC	Channels	4	**1**	
SerDes	4-lane (28Gbps)	CEI-28G-SR/MR	**CEI-28G-VSR**	
DSP	Logic Gate Size	~30-50M	**5M**	Optimized hard-wired FFT/iFFT. Based on 28nm design
FEC	Type	SD (CI-BCH)	BCH	
	Coding Gain	11-12dB	8.7dB	
	Line Rate OH	20%	12.5%	
	Latency	< 10 us	< 150ns	
	Gate Count	~20M	**0.75M**	

DSP for DMT is < 1/10th size of a coherent DSP

100G DMT Transceivers

- Can support multiple PMDs for 500m, 2km, 10km…..
- Can potentially support flexible line rates (25G,50G,75G,100G)

100GbE module (SR/LR/ER)

Mature CMOS process, 28nm

IC power dissipation estimate: 3.5 W

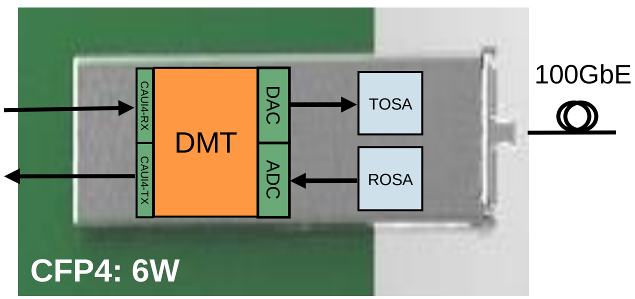

100GbE module (SR/LR/ER)

Next generation CMOS process, 14nm

IC power dissipation estimate: **<1.8W**

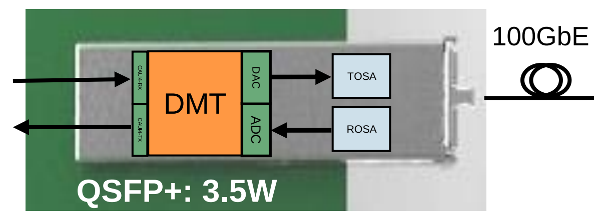

100GbE Host-based (SR/LR/ER)

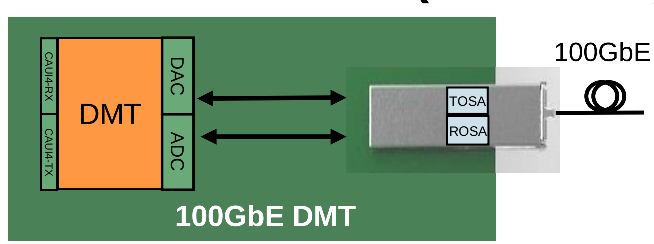

400G DMT Transceivers

FUJITSU

- Can support optical PMDs for 500m, 2km, 10km…..
- Can potentially support flexible line rates
- 4 x 1-ch 100G transceivers or dedicated 4-ch 400G transceiver
- Host-based transceiver provides opportunity for delivering 400G in a QSFP+ form factor
  - Optics only inside QSFP+ module, with 100G DMT electrical interface
  - 116Gbps electrical interface from chip to module through connector may not be possible with PAM-4, PAM-8
  - Further experiments necessary, but B2B electrical results already show good feasibility

400GbE host (LR4/ER4)

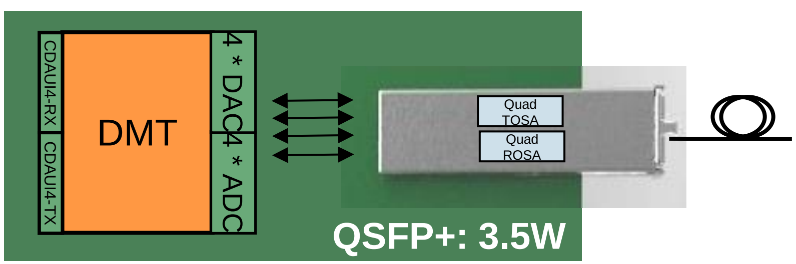

2 x 400GbE host (LR4/ER4)

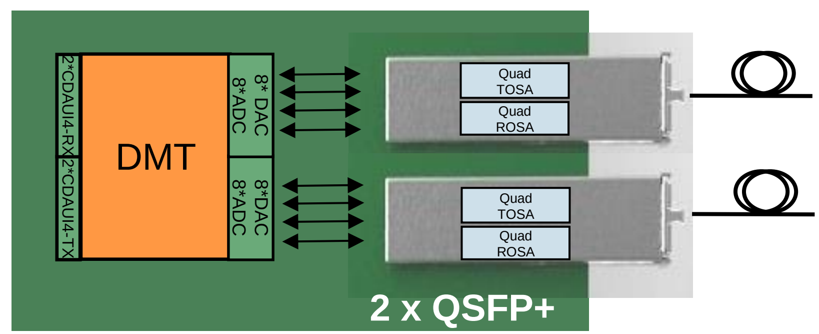

Addressing Concerns on DMT

FUJITSU

- *"DMT is overkill for optical transmission as the channel is flat"*
  - It's not just about the optics: The electrical channel is a factor @ 56Gbaud
    - ADC & DAC have non-flat responses
    - Electrical discontinuities can close a PAM-4 eye requiring further equalization
    - PAM4 is more susceptible to reflections and overshoot/ringing in the channel.
      - Channel includes the chip, package, PCB, laser driver, TIA.
      - 1 baud (18ps) is about 3mm so delayed reflections can fall outside the span of a low-power digital equalizer … more taps mean more power
    - Matching (S11) of all components up to around 35-40GHz instead of maybe 20GHz
  - PAM-4 needs roughly constant SNR across the full bandwidth
  - DMT adjusts the modulation complexity to the available SNR
    - The majority of bits are loaded into lower subcarriers, where SNR is generally higher

- *"DMT's link negotiation is too complicated"*
  - Link negotiation handshaking protocol is not complex (see Page 5) .. just new
  - Link negotiation needs minimal processing using on-chip ARM™ CPU

- *"DMT is more susceptible to non-linearities and requires use of Volterra NLC"*
  - Experiments show that NLC may not actually be necessary for <2km links
  - Experiments show NLC can be used effectively to achieve 40km

Addressing Concerns on DMT

- *"Implementation of FFT & iFFT is 10x more complex than FFE for PAM-4"*
  - FFT & iFFT can be efficiently realized
    - 6 "scaling" operations per sample - multiplication by hardcoded constants
    - Lower clock frequency (e.g F_s/528) translates to lower overall power
  - FFE requires real multipliers with programmable coefficients
    - N-tap FFE needs N-real multipliers per sample.
    - Parallelism needed to reduce clock rates (e.g. N*64 multipliers needed for F_s/64)

- *"Where's my eye diagram and how do I test this?"*
  - Optical Modulation Analyzers supporting QAM & OFDM can be software-upgraded to support DMT to provide relevant statistics and measurement data
  - Further discussion with test & measurement vendors is needed on requirements

- *"I need multiple sources for modules/transceivers"*
  - Several IC companies have announced ADC/DAC technology that could readily be adapted for DMT transceiver designs…if they choose to do so.

Raising Concerns on PAM-4

- *Baud Rate Sampling - Timing and Clock recovery loop*
  - All experiments are based on offline processing of data
  - How to reliably and cost effectively (power) do this @ 56Gbaud?

- *ADC & DAC Bandwidth Requirements*
  - Most if not all experiments with 28 & 56Gbaud PAM-4 use high performance DSOs
  - 33GHz bandwidth requirement is not easy to achieve without power penalties

- *Equalization Requirements*
  - Complexity varies over the various studies, unclear on the power implications
  - Assumptions on EQ also based on extremely high BW ADC & DAC (see above)
  - MLSE looks promising, but timing & clock recovery still needs to be addressed

- *Optical Component Bandwidth Requirements*
  - …

Bandwidth Requirements (PAM-4)

FUJITSU

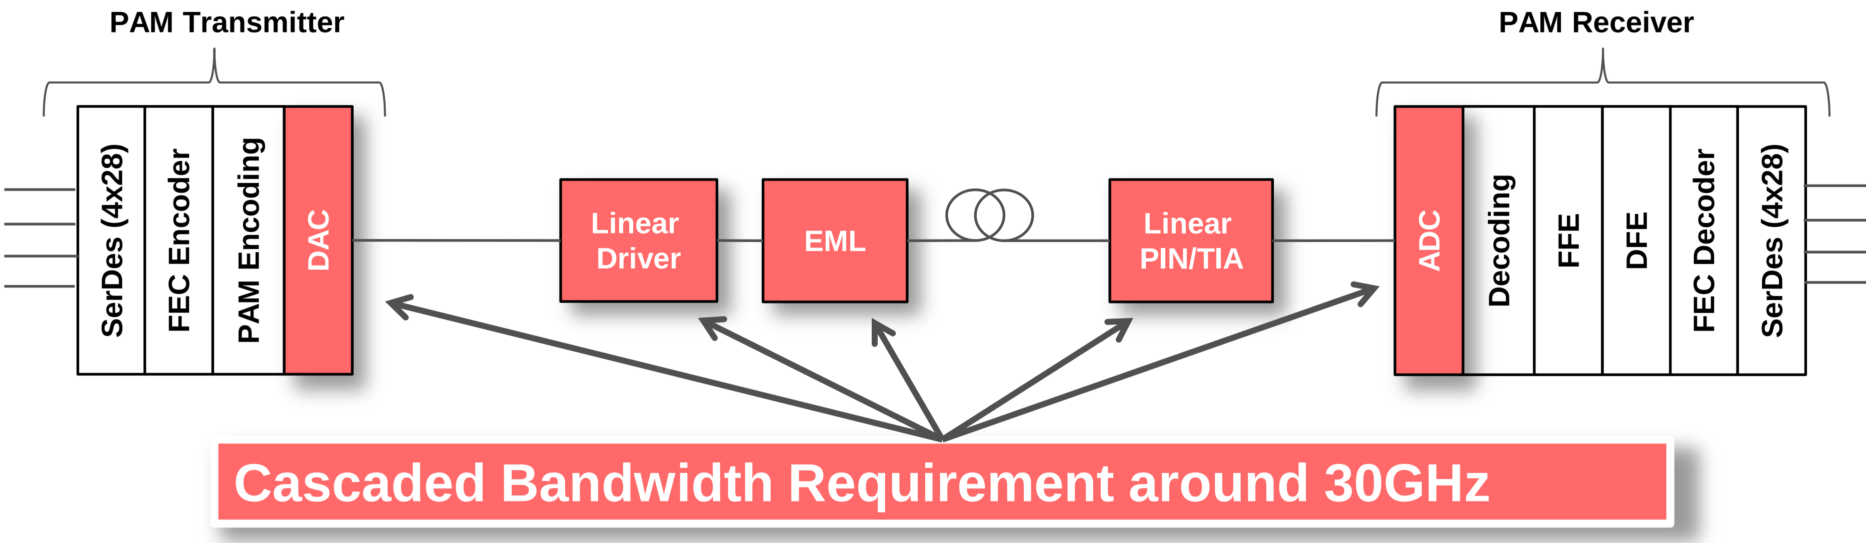

- 40/56G optics will be required
- ADC with 33GHz front-end bandwidth is very challenging
  - Should not be underestimated…even with reduced resolution
  - Increase in equalization complexity will be needed to compensate for lower bandwidth
- 33GHz Packaging & PCB should also not be neglected
  - Insertion and return loss in package double from 15GHz-30GHz
- Bandwidth is never free!
  - Need to carefully review and consider the power & cost implications of these higher bandwidth requirements

Bandwidth Requirements (DMT)

FUJITSU

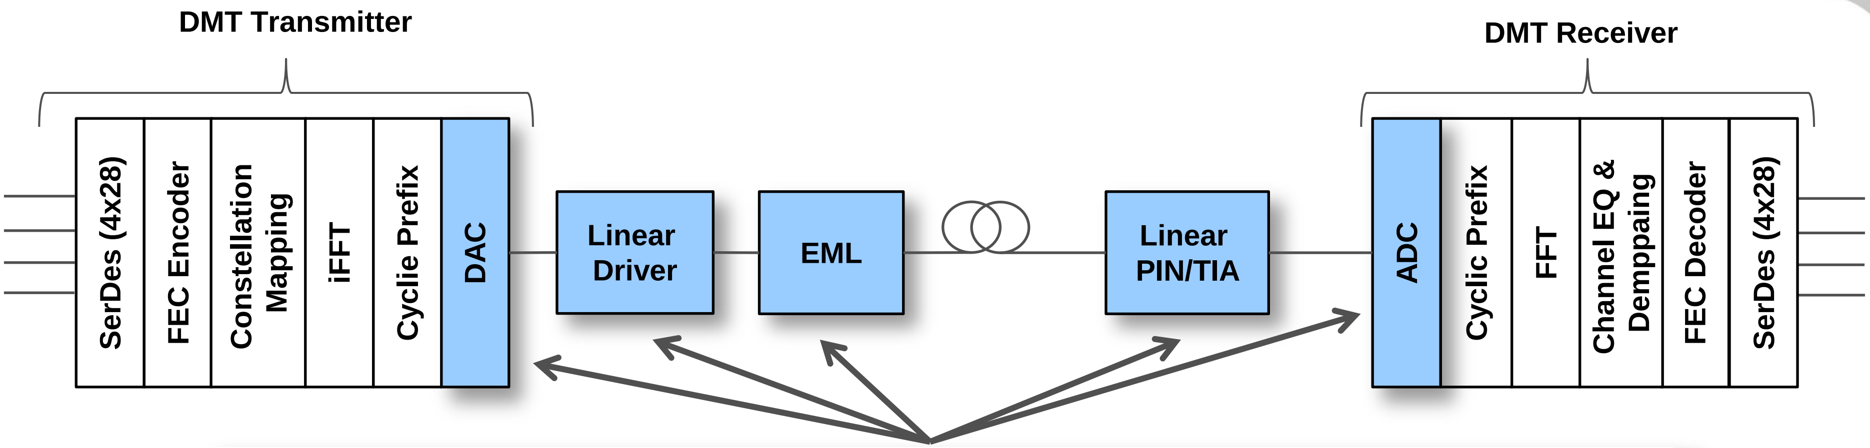

Cascaded Bandwidth Requirement around 16GHz

- 25G optics (and less) can be used
- Existing ADC & DAC (28nm) can already generate 170Gbps electrical back-to-back transmission (@ BER 1e-03) over “bandwidth limited components”
  - Increases in bandwidth due to process node migration will add more margin (excess capacity)
- Packaging & PCB losses @ 16GHz-20GHz manageable
- Bandwidth is never free!
  - Lower Bandwidth requirements should reduce overall system complexity & cost

Complexity Summary

FUJITSU

Item		PAM-4 (56Gbaud)	DMT	Notes
ADC	Sample Rate	56GSa/s *1	58GSa/s	*1 Baud rate sampling may increase complexity in timing recovery loop
	ENOB	4 *2	> 5	*2 >4-bits needed for post-ADC equalization?
	Bandwidth	**33GHz** *3	18GHz	*3 Experimental results all use high-bandwidth DSOs
DAC	Sample Rate	56GSa/s ? *4	58GSa/s	*4 Pre-emphasis can impact sample rate and resolution
	ENOB	2 ? *4	> 5	
	Bandwidth	**> 30GHz ?**	16GHz	
SerDes	4-lane 28Gbps	CEI-28G-VSR	CEI-28G-VSR	
FEC	Coding Gain	~9dB	~8dB	
DSP	Logic Gate Size	3-5M *5	5M	*5 Equalizer implementation?
Linear TIA		**40G/56G class**	< 25G class	
Linear Driver		**40G/56G class**	< 25G class	
Laser (EML)		**40G/56G class**	< 25G class	

Post 100G/lambda...

- 100G/lambda is expected for 400Gbps, but what about 1Tb?
- Higher-order modulation will be required for 400G and 1Tb transceivers
- Low-power coherent DP-QSPK single wavelength 400Gbps already being proposed for short-reach applications
  - Proven technology and modulation format, power reduction key
- Key building blocks for all will be the same
  - High speed ADC, DAC & DSP
- All of this could be realizable with the next generation of converter technology and low power CMOS process nodes

Summary

- 100G/lambda is readily achievable through use of modulation and ADC/DSP/DAC architectures
  - CMOS technology can offer performance and cost advantages
- Implementation of DSP solutions is not a concern
  - Silicon issues already solved for coherent long-haul
  - Can help decrease complexity in the optics
- Both PAM-4 and DMT can provide solutions for 100Gbps/lambda
  - DMT offers potentially the lowest cost solution due to significantly lower bandwidth requirements across the entire signal chain
  - PAM-4 increases complexity in both silicon AND optics
  - DMT increases complexity in silicon but allows usage of 10G/25G-class optics
  - DMT works with today’s technology allowing optimization with tomorrow’s technologies
- Bandwidth requirements need to be carefully considered
  - Bandwidth isn’t free, there are always tradeoffs
    - Lower bandwidth generally equals lower cost and power
  - Focus should not only be on just the DSP silicon, but all components

FUJITSU

shaping tomorrow with you