Improved specifications for 40GBASE-LR2 40 Gb/s PMD for 10 km duplex SMF

Piers Dawe

Avago Technologies, Fiber Optics Product Division

September 2008

Contents

1. 2 x 20G SMF proposal with improved spec tables
   - Also describes lighter spec variant
2. Annex describing how LR2 technology could provide lighter 100G SMF solution
3. Slides from dawe_01_0708 repeated for convenient reference

Introduction

- The anticipated gulf between parallel MMF and duplex SMF before dawe_01_0708 was too great
  - Size, power and cost
- dawe_01_0708 showed that 2, 20G wavelengths are much preferable to one, 40G wavelength (serial) or four, 10G wavelength (son of LX4), except for extremely short or long time horizons
  - Lower power
  - Smaller size
  - Lower cost
- 40GBASE-LR2 as in dawe_01_0708 narrowed the gap significantly for 40G, SMF vs. parallel MMF, and provided a duplex solution
- This presentation shows even better 2 x 20 specifications
- Annex describes 5-lane versions for 100G
  - Better power, size, cost than 100GBASE-LR4
  - Better reach and duplex, compared with attempts to twist 100GBASE-SR10 into a repeat of 10GBASE-SR

2 x 20G module building blocks

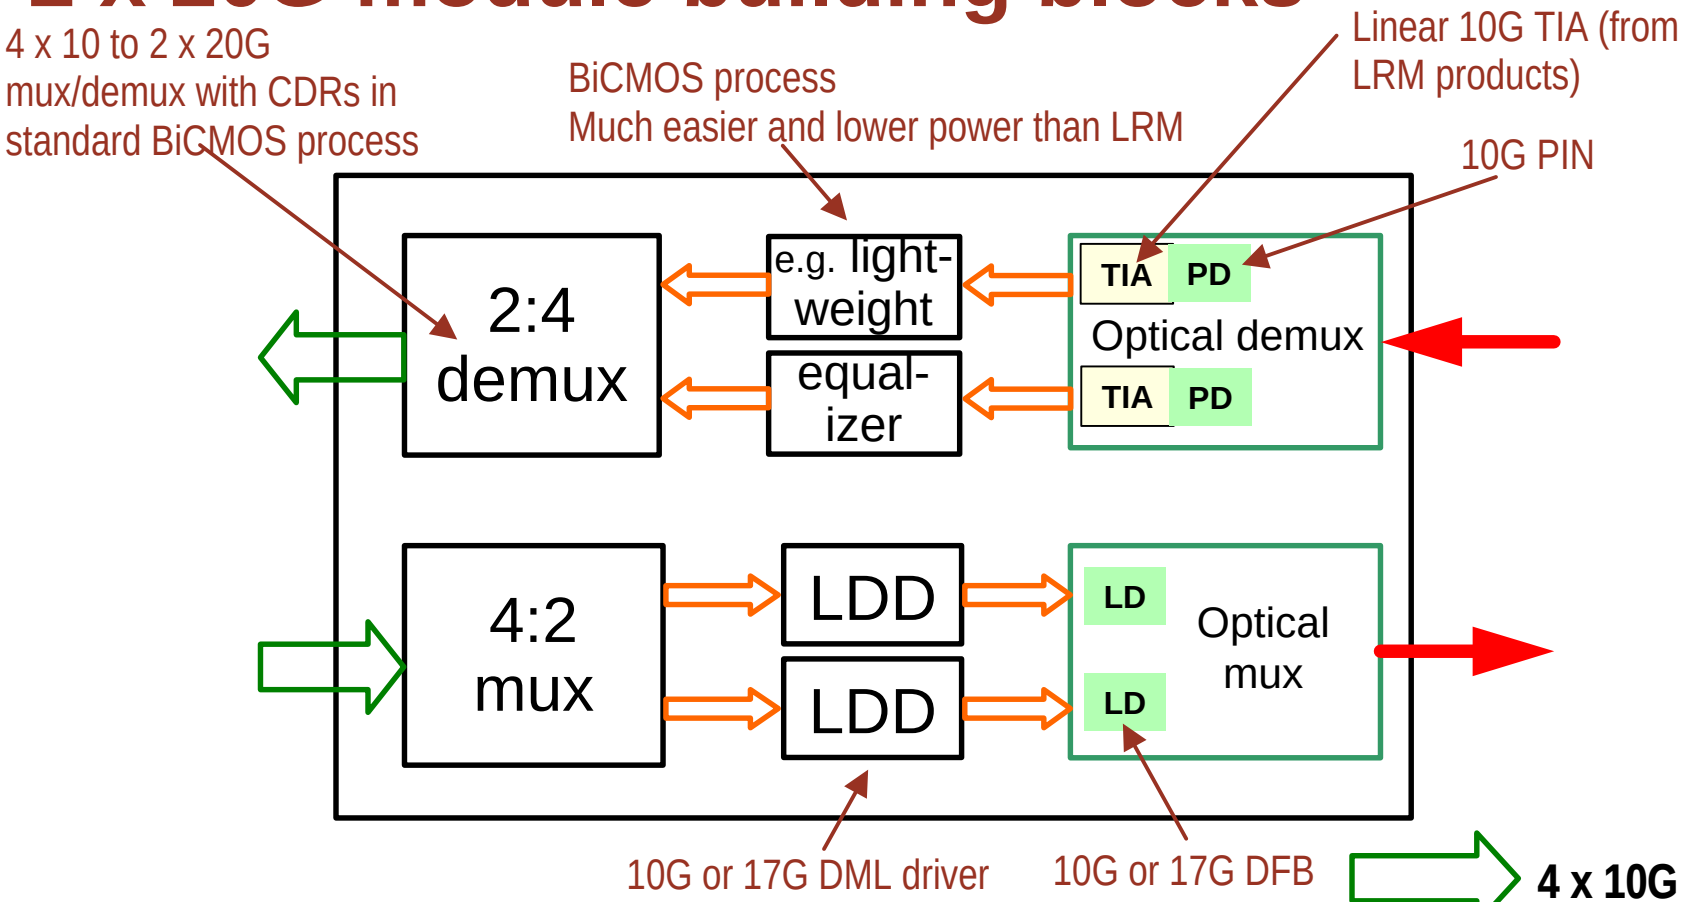

- This is one possible implementation; others are feasible
- Power consumption much the same as dawe_01_0708 (rebalanced between Tx and Rx)
- Expected cost is reduced vs. July presentation by increased use of 10G/17G grade components

Opportunities for further improvement over dawe_01_0708

- dawe_01_0708 has a tight power budget
  - Overload and Tx max-min setup and tracking range are both demanding
- dawe_01_0708 assumed 1.5 dB connector loss; 2 dB may be preferable
  - Significant interest in high loss/short distance links
- Reducing transmitter optical power would reduce module (thermal) power
- Relaxing transmitter eye, risetime and setup window requirements would reduce transmitter cost and power and ease over-temperature operation
- With 20G lanes, RIN becomes even more challenging than at 10G
  - Needs to be better per Hz because doubled noise bandwidth, but can be worse per Hz at higher modulation rate. Revised specs better optimise RIN
- Use of modern technology (equaliser and/or FEC) significantly eases these burdens and provides a self-optimising, monitorable, manageable, robust link
- dawe_01_0708 did not address 100G

Common themes for improved spec numbers

- **Two options presented which have 2 dB to 3.5 dB connector loss to allow more connectors**
- **Can fit in a QSFP in the medium term**
  - For high density data center-optimized designs. Fully retimed, unlike SFP+
- **Specs with an attractive cost point have more than 40 ps indicative max. rise time**
  - Hence high bandwidth receiver and/or receive side equaliser, or other implementation choice
  - Note the synergy between 20G lanes and 17 GBd in Fibre Channel
  - CDRs in module and no heavy equaliser on host board eases SFP+'s interop issues. One vendor is responsible.
  - Roughly speaking, can use the good end of **10G direct mod laser and driver**, use (good) **10G linear TIA** or 17G or 25G class linear TIA, with equaliser or other receiver implementation choice
- **2-wavelength plan**
  - Two wavelengths surprisingly much better than 4: see next slide (much as in dawe_01_0708)
  - Wavelengths can be chosen to balance chromatic dispersion penalty and wavelength-dependent loss
  - Wavelength plan can be tweaked so that lasers to another CWDM wavelength plan can be used – supply benefit
  - 5-wavelength plan for 100G using 5 ITU-T CWDM channels. Loses the benefit of simple diplexer
- **RefClock input? Not required, we think: easier jitter/phase noise than 40G serial**
- **Improved receiver reflectance spec to avoid reflection noise**
- **MR2 at 2 km is 1.5 dB easier than dawe_01_0708 AND allows 2 dB more connector loss!**
  - MUCH more attractive for cost and size than 40G serial or 4 wavelengths
- **We also considered Fabry-Perot variants but these have sub 1 km reach because of mode partition noise. Uncooled direct modulated DFB seems best – like 10GBASE-LR**

Are two wavelengths really that much better than four?

- In a typical 4-CWDM mux or demux component, there are 3 different dichroic mirror filters and one lane (wavelength) goes through all 3
- With 2 wavelengths there is only 1 filter. Less loss, much simpler optical path
- Wider wavelength bands allow for wider wavelength spec on lasers (yield/cost), and wider operating temperature range enabled
- Wider keep-out band between channels allows for filters with fewer dielectric layers – lower loss again
- Three fewer channels to cause crosstalk, relaxes rejection ratio spec by 5 dB
- In 2-wavelength plan, the filter steepness needed is fully 3x easier than ITU-T CWDM, allowing more technology choices
- Plus the basic benefits e.g. fewer lasers
- Summary: yes

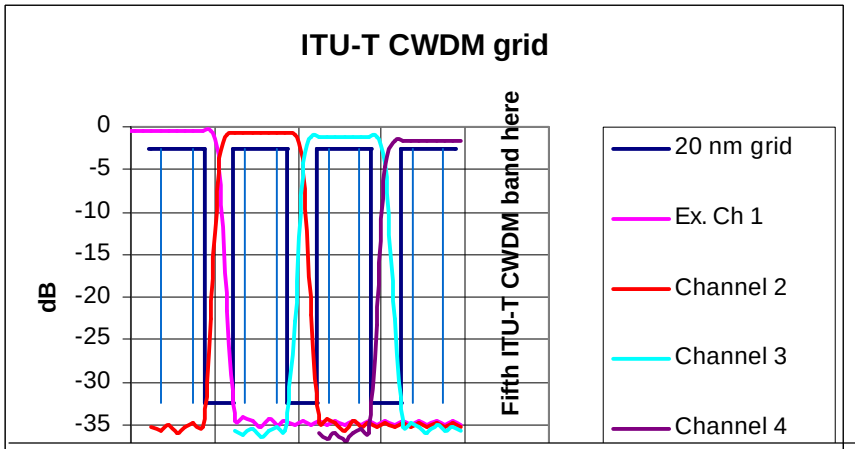

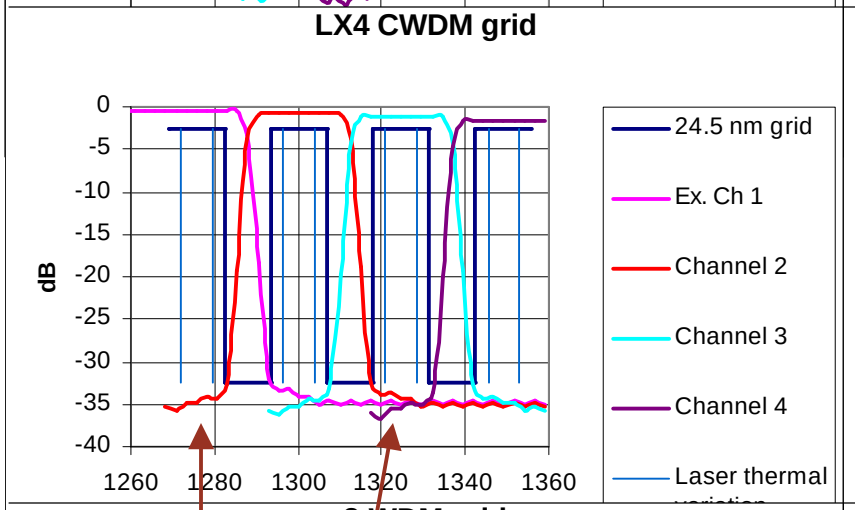

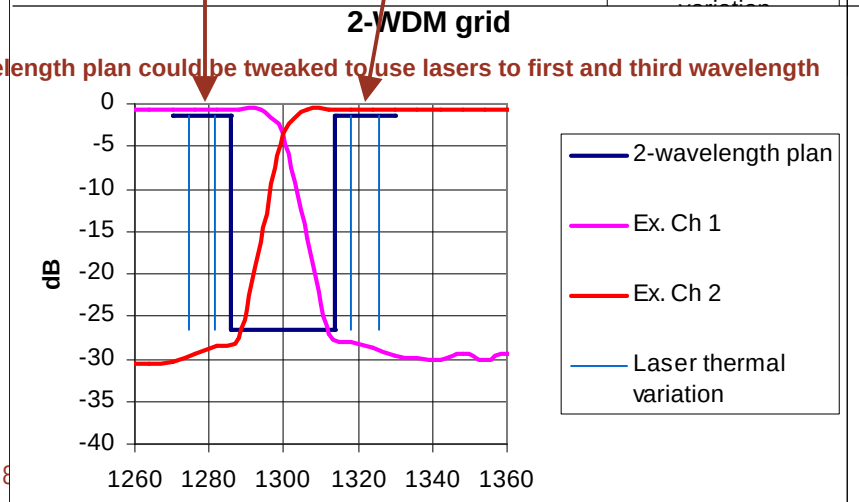

LR2 and lower spec MR2 variant:
Introduction to tables

- In the next two slides,
- Left spec column is from dawe_01_0708
- Next column is improved 10 km spec (LR2) if FEC is not acceptable
- Next three columns show a single PMD (MR2) optimized to shorter distances, showing three use scenarios
  - 10 km with FEC
  - 5 km, no FEC
  - 2 km, no FEC, extra 1.5 dB connector loss
- FEC is 10GBASE-KR FEC, giving nominally 1 dBo SNR gain, up to 1.5 dB when second order effects (e.g. RIN, P_{cross}) taken into account

LR2 and lower spec MR2 variant 1/2

Interesting rows arrowed | July specs | For adoption | For interest |

Transmit characteristics		dawe_01_0708	10 km	10 km FEC	5 km	2 km	
Description	**Type**		**40GBASE-LR2**	**40GBASE-MR2**			**Unit**
Signaling speed, each lane	nom	20.625 (same for Rx)					GBd
Signaling speed variation from nominal	max	±100 (same for Rx)					ppm
→ Lane 1: Center wavelength	range	1270 to 1286 TBC (same for Rx)					nm
Lane 2: Center wavelength	range	1314 to 1330 TBC (same for Rx)					nm
Side mode suppression ratio	min	30					dB
→ Average launch power, two lanes	max	7	6.5	6.5			dBm
Average launch power, each lane	max	4	3.5	3.5			dBm
Average launch power, each lane	min	-2	-3	-3.5			dBm
Optical modulation amplitude^^	max	3.5	3	3			dBm
→ **Optical modulation amplitude**	**min**	**-0.5**	**-1.5**	**-2**			**dBm**
→ **Example rise time**	**max**	**25**	**39**	**45**			**ps**
TWDP^^, each lane	max	3.6 TBC*	TBD^	TBD^			dB
and/or TDP	max	TBD					dB
Average launch power OFF, each lane	max	-30					dBm
→ **Extinction ratio, each lane**	**min**	**3.5**	**3**				**dB**
Peak launch power	max	5	4.5	4.5			dBm
→ **RIN_12_OMA or RIN_18_OMA**	**max**	**-135**	**-130**	**-128**			**dB/Hz**
Optical return loss tolerance	min	-12	-18				dB
→ Transmitter reflectance	max	-12?	-20				dB

All transmit and receive powers at TP2, TP3 (after mux loss and before demux loss)

^^ Likely combined in a single useful transmitted power metric

^ TWDP(4,1) 4 FFE taps, 1 DFE taps or similar metric

Old LR2 PMD | Single MR2 PMD

Receive characteristics

Description	**Type**	**dawe_01_0708**	**40GBASE-LR2**	**40GBASE-MR2**		**Unit**
Stressed receiver sensitivity in OMA, eac	max	-7.5	-8.6	-9.2	-8.4	dBm
Overload in OMA, each lane	min	3.5	3.0	3		dBm
Vertical eye closure penalty, each lane		3 TBC	TBD	TBD	TBD	dB
→ **Receiver reflectance**	**max**	**-12**	**-22**			**dB**

Next two lines for information in this presentation: not to go in the draft

Unstressed receiver sensitivity in OMA, e max	-11.0	-12.0	-11.0		dBm
→ **Unstressed sensitivity without FEC (OMA)**		**-11.0**			**dBm**

LR2 and lower spec MR2 variant 2/2

Interesting rows arrowed | July specs | For adoption | For interest |

Signal in & at the receiving end of, a compliant channel (informative)				40GBASE-MR2			
Description	**Type**	**dawe_01_0708**	**10 km**	**10 km FEC**	**5 km**	**2 km**	**Unit**
Highest power in OMA, each lane	max	3.5	3.0		3.0		dBm
Lowest power in OMA, each lane	min	-6 s/b -6.5	-8.0	-8.5	-5.7	-4.9	dBm
Highest average power, one channel	max	4.0	3.5		3.5		dBm
Lowest average power, one channel	min	-7.5 s/b -8.0	-9.5	-10.0	-7.2	-6.4	dBm
Highest average power, both channels	max	7.0	6.5		6.5		dBm
Lowest average power, both channels	min	-4.5 s/b -5.0	-6.5	-7.0	-4.2	-3.4	dBm
Peak power, one channel	max	5.0	4.5		4.5		dBm

Power budgets (informative)		40GBASE-LR2	40GBASE-MR2			
Transmit characteristics	**dawe_01_0708**	**10 km**	**10 km FEC**	**5 km**	**2 km**	**Unit**
Power budget	10.5	9.5	10.0	9.0	9.0	dB
Ex-FEC budget	**10.5**	**9.5**	**9.0**	**9.0**	**9.0**	**dB**
Operating distance	10			5	2	km
Channel insertion loss (max) at 1310 nm	**5.5**	**6**		**4**	**2.8**	**dB**
Channel insertion loss (min)	0					dB
Allocation for penalties	5.0	3.5	4.0	5.0	6.2	dB
Additional insertion loss allowed	**0**			**0**	**1.5**	**dB**

The maximum channel insertion loss is based on the cable attenuation at the target distance and nominal measurement wavelength, 1310 nm). The channel insertion loss also includes the loss for connectors, splices and other passive components.

The allocation for penalties is the difference between the available power budget and the channel insertion loss; insertion loss difference between nominal and worst-case operating wavelength is considered a penalty.

- **LR2 is better than specs in dawe_01_0708, which in turn is better than 4 wavelengths or serial**
- **MR2 goes with FEC at 10 km, or provides better optimization for lengths between 100 m and 2-5 km**
- **Could make LR2 and MR2 interoperable, not sure what the cost of harmonised overload spec would be**

Transmitter power tolerancing

- **Graphical representation of approximate region of signal compliance (informative)**
- **Optical powers at lane transmitters and receivers are higher/ lower than TP2 / TP3 powers by the WDM losses**
- **Wider setup window than dawe_01_0708 in both diagonal directions**

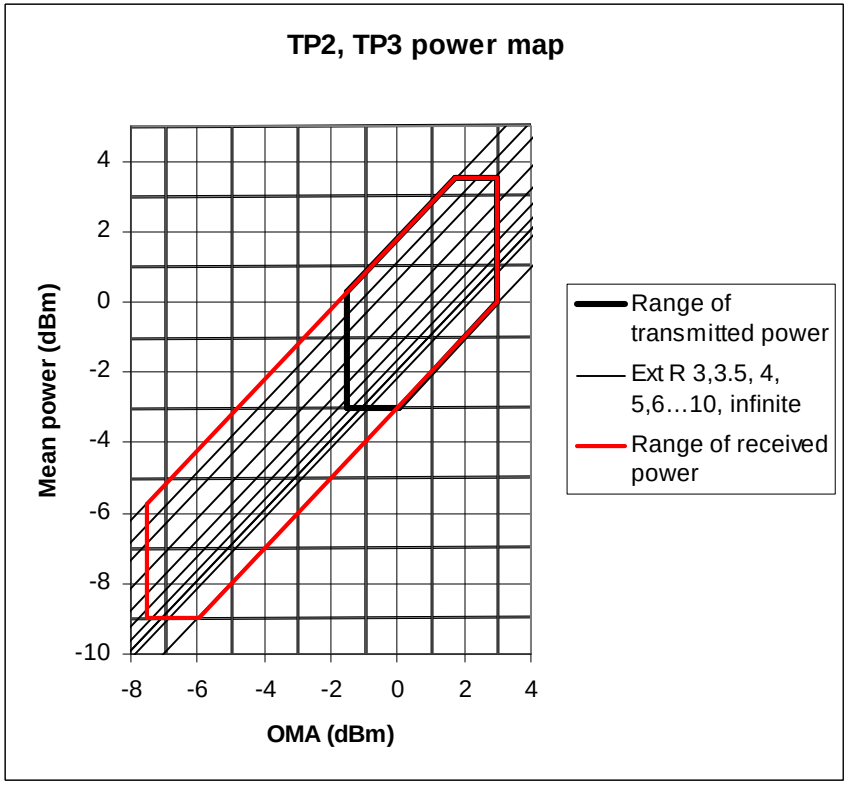

Mechanical

	Format	Power	Dimensions
40G	Shortest term: e.g. X2	4 W (power Class 1)	42* mm W x 12^ mm H x 77.2% mm D (PCI/low)
	Medium term: QSFP	3.5 W (Power Level 4)	19 mm W x 14 mm H x 52.4% mm D
100G	New form factor	10 W TBC	TBD (<50.8) mm W x 17? mm H x 77.2? mm D

* Width of body is 36 mm

^ Height of nose. 19.46 mm if midpak

% Inside the chassis

What might happen in future at 40G

	Near term	Reduced power CDRs	SFP+ style un-retimed LR4 module	Photonic integration	Low cost microwave inter-connect	20G lanes in host	40G lanes in host	40G direct mod laser
Four wave-lengths		Reduced power	Power reduced to a little more than 2 wave-lengths	High NRE, reduced size		2:4 inverse gearbox	1:4 inverse gearbox	
Two wave-length	Best power, size and cost	Reduced power	Still best power, size and cost, robust	High NRE, reduced size *		4:2 gearbox replaced by CDRs: power reduced again	1:2 inverse gearbox	
One wave-length		Reduced power		High NRE, reduced size ^	Reduced cost	4:2 gearbox replaced by 2:1: reduced power	Gearbox replaced by CDRs: reduced power	Reduced power
	2 wavelengths is best until ...						**Significant technology improvement in the future**	

* But is small enough anyway?

^ Maybe laser/EML integration is assumed anyway for serial

Conclusion

- 40GBASE-LR2 is the best of the options for power, size and cost
- Now and in the future until it's time to go serial at 40G
- Adopt "10 km 40GBASE-LR2" column of slides 9 and 10 for the 10 km 40G SMF objective

Annex: Discussion of 100G options

Can extend the 2 x 20G concept to 5 x 20G for 100GbE

- Could provide a shorter-reach duplex PMD for lengths between 100 m and 2-5 km

- **Use ITU-T CWDM wavelength plan which supports 5 lanes in the 1310 nm window**
  - 1264.5 to 1277.5, 1284.5 to 1297.5, 1304.5 to 1317.5, 1324.5 to 1337.5, 1334.5 to 1357.5
- **Have to find room for ~1 dB more chromatic dispersion penalty at worst loss because can't keep all wavelengths so near the chromatic dispersion sweet spot as with two wavelengths**
  - Lower loss of low water (G.652 C & D) fibre would compensate for this
- **Have to provide extra laser power and sensitivity, ~1/2 dB each end for increased WDM losses**
- **Therefore, sweet spot for LR2 technology at 100G is up to 2 km to 4 km**
- **Still significantly smaller and lower power than 4 x 25G with coolers and/or EML**
  - For the same laser, 4 lanes x 25.78125 GBd operation needs ~3 dB more equalisation than 20.625 GBd: not viable, would need different laser and/or cooler
  - 10:5 gearbox significantly more straightforward than 10:4
  - Exact power savings yet to be quantified
- **Synergy with 20G electrical lanes – MUCH easier than 25G electrical because of PCB losses**

Backup: most of the slides from dawe_01_0708

- Wavelength plan not repeated as shown in slide 7
- Transmitter power tolerancing not repeated
- Spec tables not repeated as numbers are shown in slides 8, 9

Comparison of options for 40 Gb/s PMD for 10 km duplex SMF and recommendations

Piers Dawe, David Cunningham and Dan Rausch

Avago Technologies, Fiber Optics Product Division

July 2008

Outline

- **Example 40 Gigabit Ethernet implementations**
- **High level one page summary for each SMF option**
- **Comparison tables**
- **Summary**
- **Proposed baseline specification tables for 2 x 20G WDM or BiDi PMD for duplex SMF**

Context

Question from May meeting, comparing 4 x 10G CWDM with 1 x 40G serial

Straw Poll #11: I believe that a baseline proposal for the 40GBE 10-km SMF PMD should be based on:

A) 4x10G CWDM (as per “cole_03_0508.pdf”)

B) 40G Serial (as per “jewell_03_0508.pdf”)

C) I need more information and presentation material before deciding.

D) I will abstain now and later.

Results:

A) 25

B) 23

C) 35

D) 6

Approximate Room Count: 108

Debate expressed in terms of either-or

But would a compromise between these be better?

Meanwhile,

- kropp_01_0508 proposed 2 x 20G for MMF
- Several presentations have considered 25 GBd direct modulated lasers – cutting edge
- Fibre Channel FC-PI-5 are working on 14.025 or 17 GBd serial

But not 2 x 20G for SMF. This presentation fills this gap

Example 40G implementation architectures

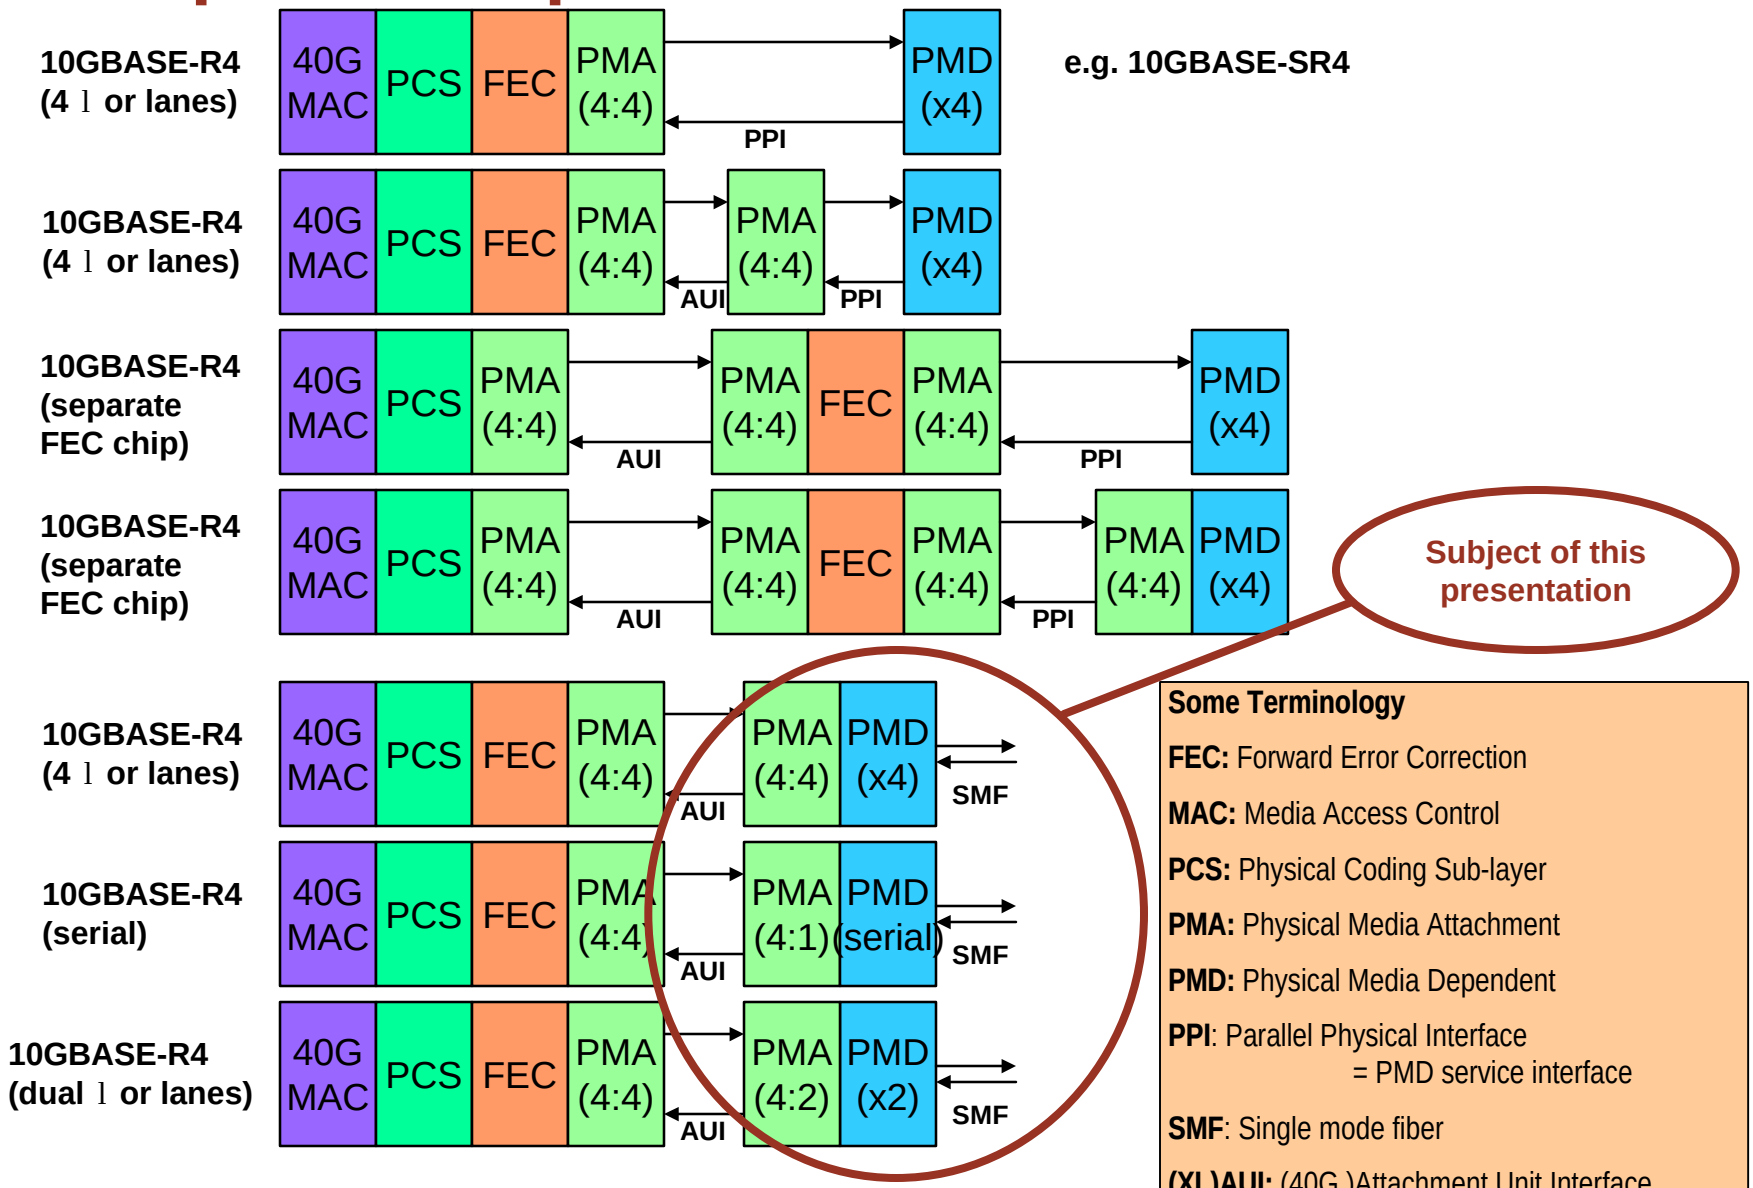

40G SMF: 4 x 10G CWDM
The past

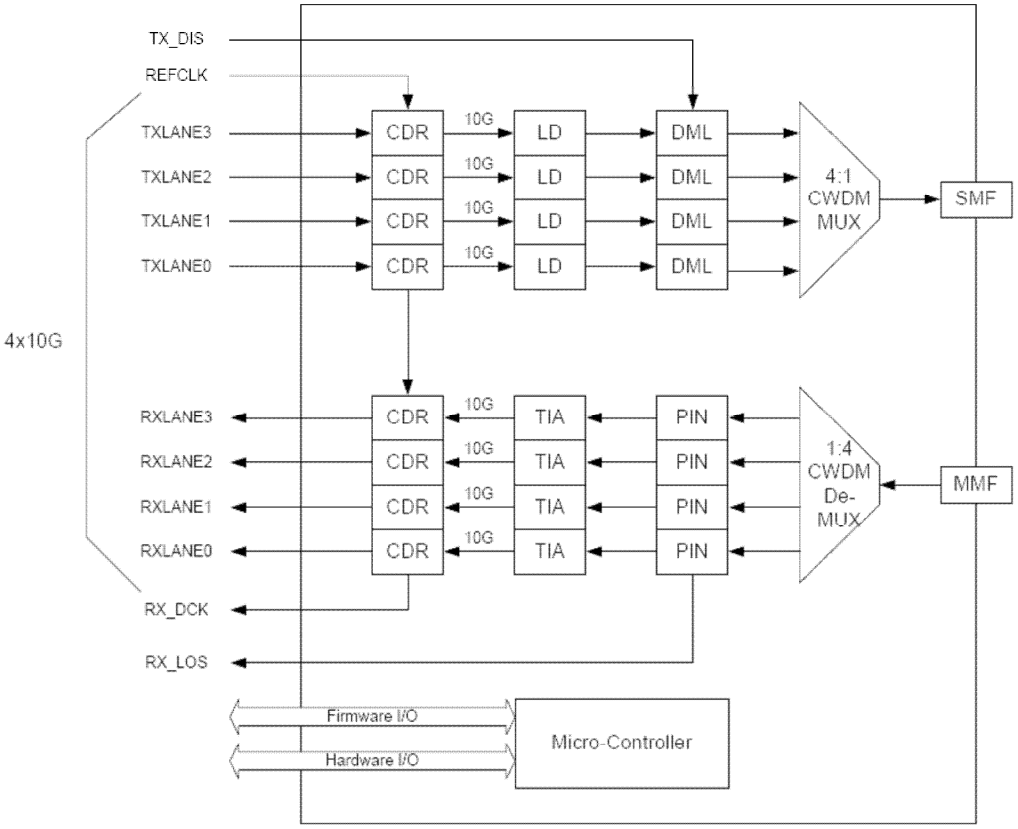

- **4-wavelength CWDM near 1310 nm**
- **Internal optical mux and de-mux**
- **SMF:** 10 km
- **OM1-3:** < 75 m without EDC. MMF compromises Rx WDM demux and PIN. But, the MMF specification is not part of the 40G standards proposal
- **10G Optics and electronics**
- **Power dissipation:** ~ 6 W
- **Size:** Large ~ XENPAK form-factor
- **Cost:** High due to complex WDM-Opto-mechanical design and 4 x 10G optical channel elements. Too much NRE for the volume/lifespan to be worth fixing.

40G SMF: 4 x 10G CWDM Alternative

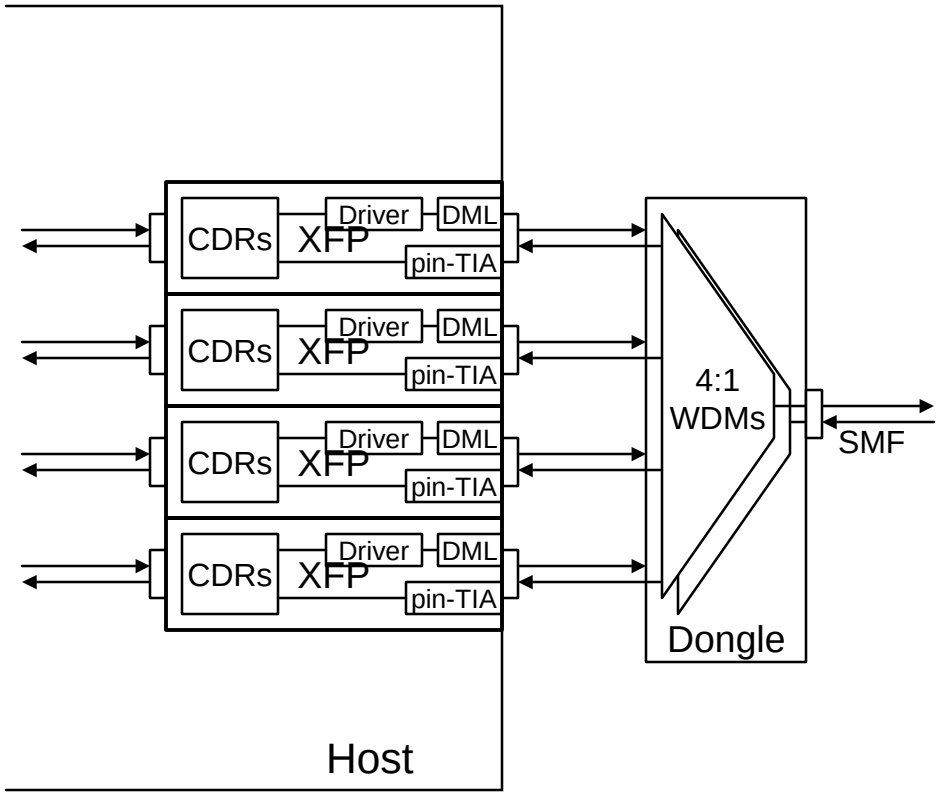

- **4-wavelength CWDM near 1310 nm**
- **External optical mux and de-mux**
- **SMF:** 10 km
- **OM1-3:** < 75 m without EDC. MMF compromises Rx WDM demux and PIN. But, the MMF specification is not part of the 40G standards proposal
- **10G Optics and electronics**
- **Power dissipation:** ~ 6 W
- **Size:** Large e.g. 4 x XFP
- **Cost:** High due to WDM, 4 x 10G optical channel elements and cases, **but** much lower NRE and can pay-as-you-grow
- **Comments:** Best short term option.Too easy? As power hungry as previous slide. If you don't like this you should like the previous slide even less

40G serial NRZ modulation
The future

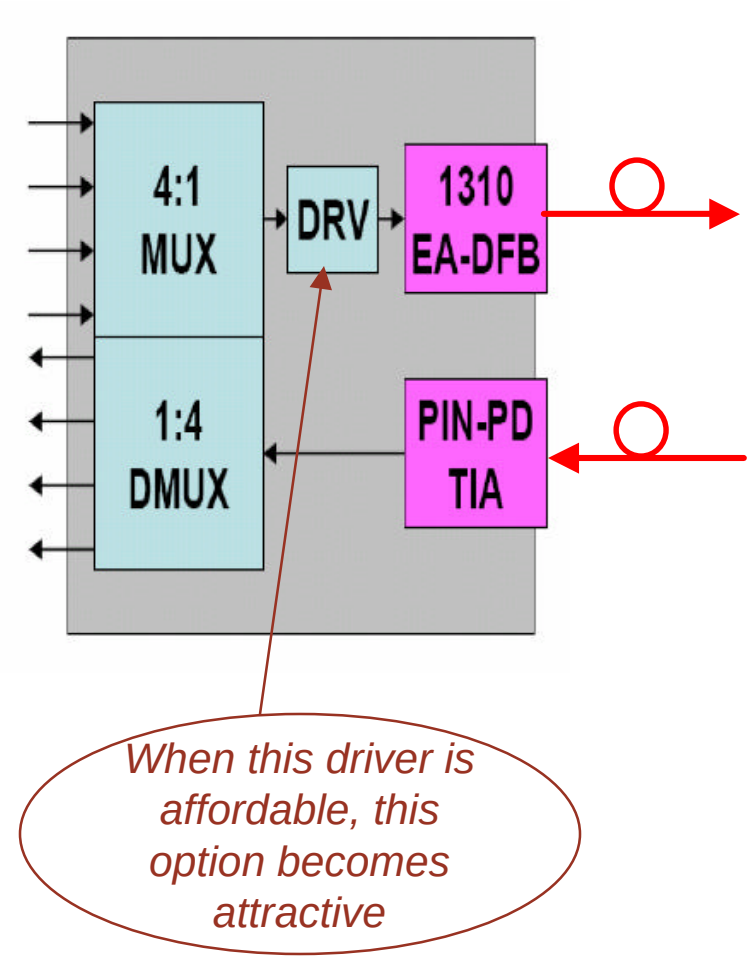

- **Single wavelength near 1310 nm**
- **SMF:** 10 km
- **OM1-3:** < 20 m without EDC. MMF badly compromises PIN. But, the MMF specification is not part of the 40G standards proposal
- **40G Optics:**
  - cooled EAM-DFB
- **40G electronics:**
  - InP laser driver and TIA
  - SiGe mux and demux
- **Power dissipation:** ~ 5 W
- **Size:** Big ~ X2 form-factor
- **Cost:** High due to 40G InP optical and electronic components

40G SMF: 2 x 20G, 1310 nm, BiDi or WDM

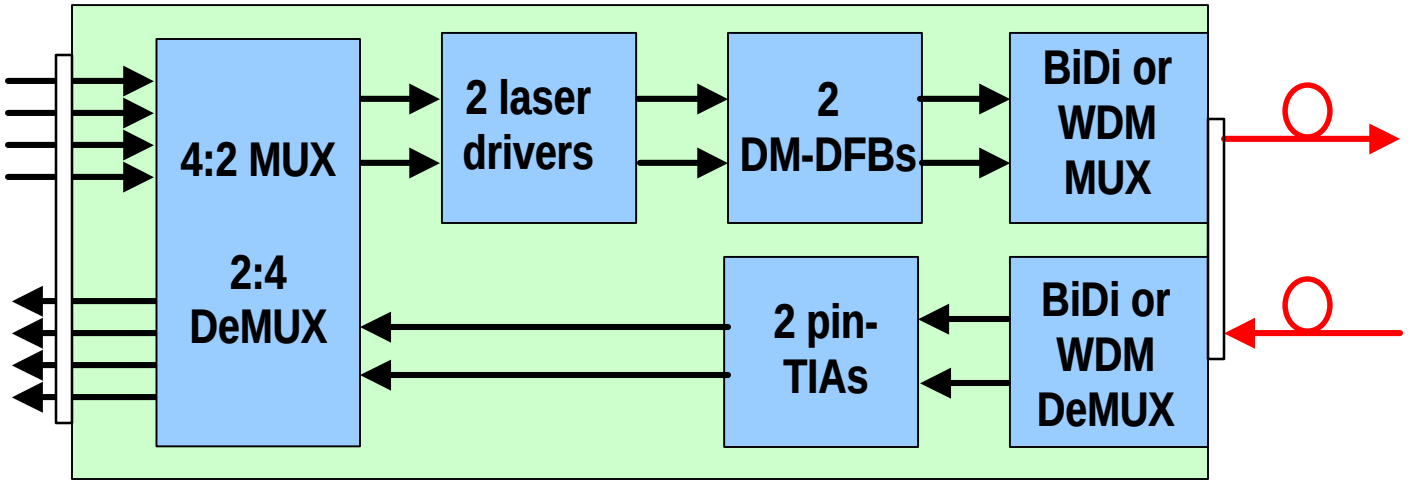

- **Wavelength plan: Two wavelengths in 1310 nm window:** up to 10 km, SMF only
  - Either both transmitters on same fiber ("regular WDM" or one on each ("BiDi")
- **Optics:** Direct-modulated DFB, 20 GBd grade TIA: very similar to FC-PI-5 (14.025 or 17 GBd). Routing for BiDi would cross over (not shown for simplicity). No cooler.
- **Electronics:** SiGe BiCMOS compatible. No exotic IC technology.
- **Power Dissipation:** ~ 3 W, much less than the other SMF 40G options being considered
- **Size:** Small, compatible with QSFP form-factor (3.5 W max)
- **Cost:** Low compared to the other 40G options being considered

Power dissipation comparison table

40G 10km 4-CWDM	2009 power (W)	2011 power (W)	40G 10 km serial	2009 power (W)	2011 power (W)	40G 10 km 2-WDM	2009 power (W)	2011 power (W)
4x DML, TOSA-WDM Mux	0	0	EML TOSA/TEC	1.5	1.0	2x DML, TOSA-Mux	0	0
4x DML driver	2.1	1.7	EML driver	0.8	0.6	2x DML driver	1.3	1.0
4x XFI Tx-Rx CDRs	1.8	1.0	4:1 & 1:4 MUX-CDR	2.0	1.5	4:2 & 2:4 MUX & deMux	1.7	0.9
4x PIN-TIA ROSA-deMux	0.7	0.5	PIN-TIA	0.4	0.3	2x PIN-TIA ROSA-deMux	0.6	0.5
Other	0.4	0.4	Other	0.1	0.1	Other	0.1	0.1
Total power	5	3.6	Total power	4.8	3.5	Total power	3.7	2.5
Ratio to CWDM	1	1	Ratio to CWDM	0.96	0.97	Ratio to CWDM	0.74	0.69

Based on jewell_03_0508

- Apples to apples comparison – aggressive
- Clearly, 2 x 20G WDM or BiDi is the lowest power option

Sept update: LR2 Tx power reduced, Rx power increased, net no change

Relative cost comparison table

40G 10km 4-CWDM	Cost factor 2009	Cost factor 2011	40G 10 km serial	Cost factor 2009	Cost factor 2011	40G 10km 2-WDM	Cost factor 2009	Cost factor 2011
4x DML, TOSA-WDM Mux	1	1	EML TOSA/TEC	2-3x	1x	2x DML, TOSA-Mux	0.7x	0.6x
4x DML driver	1	1	EML driver	4x	<1x	2x DML driver	0.6x	0.5x
4x XFI Tx-Rx CDRs	1	1	4:1 & 1:4 MUX-CDR	3x	2x	4:2 & 2:4 Mux deMux	1.3x	1.1x
4x PIN-TIA ROSA-deMux	1	1	PIN-TIA	4x	1x	2x PIN-TIA ROSA-deMux	0.7x	0.6x
Form-factor	1	1	Form-factor	1x	1x	Form-factor	0.25x	0.25x
Total Cost	1	1	Total Cost	2.8x	1x	Total Cost	0.8x	0.65x
Based on jewell_03_0508			Form factor row is our addition					

- Cost/benefit of PCB real-estate and data center air-conditioning with significantly smaller, lower-power module not included

- Clearly, 2 x 20G WDM or BiDi is the lowest relative cost option

Sept update: LR2 Tx cost reduced, Rx cost increased, net reduction

Comparison summary table

	4x10G 10 km CWDM one box	4x10G 10 km CWDM 4 boxes	40G 10 km serial	2 x 20G WDM or BiDi
Present situation	Highest power	Highest power	Highest cost now	*Lowest power, size and cost*
Form factor	Big	Big but ready now	Big now, can shrink somewhat	QSFP possible
Physical lanes	Too many	Too many	Too fast - for now	Just right
Prospects	Optical plumbing/ mechanical costs endure	Optical plumbing/ mechanical costs endure	Modulator driver technology/ cost evolution is key. Mechanical costs can shrivel, but stuck with cooler	Synergy with 16GFC and some synergy with 25G lanes in HSE
Summary	Dead-end technology	Existing technology	The future	Best in near term and more

Summary of comparison

- **For the 10 km 40G SMF objective, it has been clearly shown that the 2 x 20G WDM or BiDi PMD is optimum in terms of the relative power, cost, size and performance.**
- **As was said in kropp_01_0508, the 2 x 20G WDM or BiDi PMD can also leverage the components that will be developed for 16G Fibre Channel**
  - This will significantly decrease its relative cost
  - IC synergy with kropp_01_0708. Also some synergy with 25G lanes in P802.3ba
- **Both the 2 x 20G WDM or BiDi PMD and the 40G serial PMD will use new, efficient, and potentially compact technologies.**
  - Will 40G serial need a cooler forever?
- **But, the 40G CWDM PMD uses old, power hungry and bulky technologies.**
- **Proposal: Given the above it would seem that the wise decision would be to plan for both the 40G serial PMD and the 2 x 20G WDM or BiDi PMD:**
  - 40G serial for the future, however distant. Specs e.g. per jewell_03_0508. New CFI?
  - 2 x 20G WDM or BiDi for the near term, medium term and a bit of the long term too. Specs on following slides