25 Gb/s and 50 Gb/s Electrical Interface Consensus Building

IEEE 802.3bs 400Gb/s Ethernet
September 2014 version r4

Overview

Quick Note

A Process - Simple Steps Forward

Looking at Implementations

Working towards possible adoptions

Steps to November

Quick Note: This is a Consensus Building Meeting

This was reviewed at each meeting:

Not an official IEEE meeting

Build Consensus towards a presentation or adoption covering interface type, rate, modulation, and loss

Offer lots of constructive comments to converge quickly

Please mute your phone

Requested Presentation Time for September

Requested Evening Time for Further Discussion

Supporters Appreciated / Attendees will be listed

There are some things discussed that are outside the scope of the project, but are necessary to complete the system level discussion and drive the best Electrical Interface possible.

Goals for this Presentation

Reviewing 3 meetings worth of thoughts and conversations towards building consensus

Is there anything we can propose / adopt at the interim?

The Process Steps

The following simplified process steps are used to build consensus:

- Defining / Discussing Reach
- System Architecture
- Channel Loss – the first crack at it because one needs an egg to get to the chicken
- Modulation
- Equalization
- Error Correction
- Power

Attendees


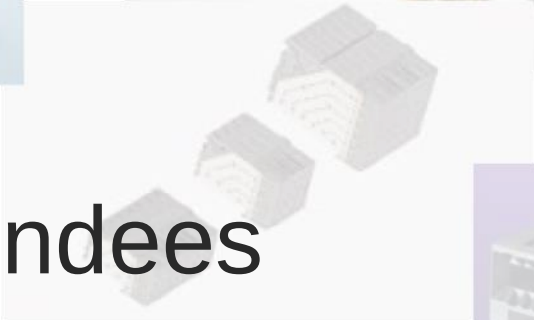

40 inches
35 dB at 12.8 GHz
Clean eye
Closed eye
Backplane

XILINX

People Who Requested to Attend

Peter	Anslow	Ciena
Vittal	Balasubramani	Dell
Arash	Behziz	TE
Liav	Ben Artsi	Marvell
Bill	Brennan	Credo
David	Brown	Semtech
Matt	Brown	APM
derek	cassidy	BT
Chris	Cole	Finisar
John	DAmbrosia	Dell
Piers	Dawe	Mellanox
Chris	Diminico	
Yuval	Domb	BroadCom
Mike	Dudek	Qlogic
John	Ewen	IBM
Edward	Frlan	Semtech
Mike	Furlong	Clariphy
Ali	Ghiasi	
Mark	Gustlin	Xilinx
Adam	Healey	Avagotech
David	Helster	TE
Yasuo	Hidaka	Fujitsu
Sammy	Hindi	Juniper
Kiyohisa	Hiramoto	Oclaro
Scott	Irwin	
Elizabeth	Kochuparambil	Cisco
Paul	Kolesar	Commscope
Ryan	Latchman	Macom
David	Lewis	JDSU
Jeffery	Maki	Juniper
Yonatan	Malkiman	Mellanox
Rich	Meltz	Intel
Andy	Moorwood	Infinera
Gary	Nicholl	Cisco
Takeshi	Nishimura	YEU
David	Ofelt	Juniper
Vasudevan	Parthasarathy	BroadCom
Scott	Powell	Clariphy
Haoli	Qian	Credo
Rick	Rabinovich	ALU
Vineet	Salunke	Cisco
Omer	Sella	Mellanox
Megha	Shanbhag	TE
Ravi	Tangirala	Ericsson
Vivek	Telang	BroadCom
Alan	Tipper	Semtech
Pirooz	Tooyserkani	Cisco
Nathan	Tracy	TE
Francois	Tremblay	Semtech
Xinyuan	Wang	Huawei
		Neophotoni
Winston	Way	cs
Helen	Xuyu	Huawei
Andrew	Zambell	FCI
Tony	Zortea	PMCS

Attendees By Meeting

Attendance on 8/27/2014
Adam Healey
Ali Ghiasi
Andy Moorwood
Beth Kochuparambil
Chris Cole
Dave Brown
David Ofelt
Gary Nicholl
Jeff Twombly
Jeff Maki
Joel Goergen
Liav Ben-Artsi
Mark Gustlin
Matt Brown
Megha Shanbhag
Mike Dudek
Nathan Tracy
Piers Dawe
Pirooz Tooyserkani
Rich Melitz
Rick Rabinovich
Scott Irwin
Tony Zortea
Vasu Parthasarathy
Vittal Balasubramanian
Vivek Telang
Xinyuan Wang
Yasuo Hidaka
Yuval Domb

Attendees on 8/20/2014
Alan Tipper
Ali Ghiasi
Andy Zambell
Beth Kochuparambil
Chris Cole
Dave Brown
David Ofelt
Derek Cassidy
Jeff Maki
Joel Goergen
Liav Ben-Artsi
Matt Brown
Megha Shanbhag
Michael Furlong
Mike Dudek
Oded Wertheim
Omer Sella
Peter Anslow
Piers Dawe
Pirooz Tooysekani
Rich Melitz
Rick Rabinovich
Scott Irwin
Takeshi Nish
Tony Zortea
Vasu Parthasarathy
Vineet Salunke
Vittal Balasubramanian
Vivek Telang
Xinyuan Wang
Yuval Domb
Arash Behziz

Attendees on 8/13/2014
Adam Healey
Ali Ghiasi
Andy Moorwood
Andy Zambell
Chris Cole
Dave Brown
David Ofelt
Derek Cassidy
Ed Frlan
Gary Nicholl
Jeff Maki
Joel Goergen
John Ewen
Liav Ben-Artsi
Mark Gustlin
Matt Brown
Megha Shanbhag
Mike Dudek
Nathan Tracy
Oded Wertheim
Omer Sella
Paul Kolesar
Peter Anslow
Piers Dawe
Pirooz Tooyserkani
Rich Melitz
Rick Rabinovich
Sammy Hindi
Scott Irwin
Tony Zortea
Vasu Parthasarthy
Vittal Bala
Vivek Telang
Xinyuan Wang
Arash Behziz

Thank You

3 meetings – over 5 hours of courteous technical conversation.

54 people volunteering what time they could free up.

Supporters Page

Vittal	Balasubramani	Dell
David	Brown	Semtech
Matt	Brown	APM
Chris	Cole	Finisar
Edward	Frlan	Semtech
Ali	Ghiasi	Ghiasi Quantum
Paul	Kolesar	Commscope
Ryan	Latchman	Macom
Jeffery	Maki	Juniper
Andy	Moorwood	Infinera
Gary	Nicholl	Cisco
Rick	Rabinovich	ALU
Megha	Shanbhag	TE
Nathan	Tracy	TE
Andrew	Zambell	FCI

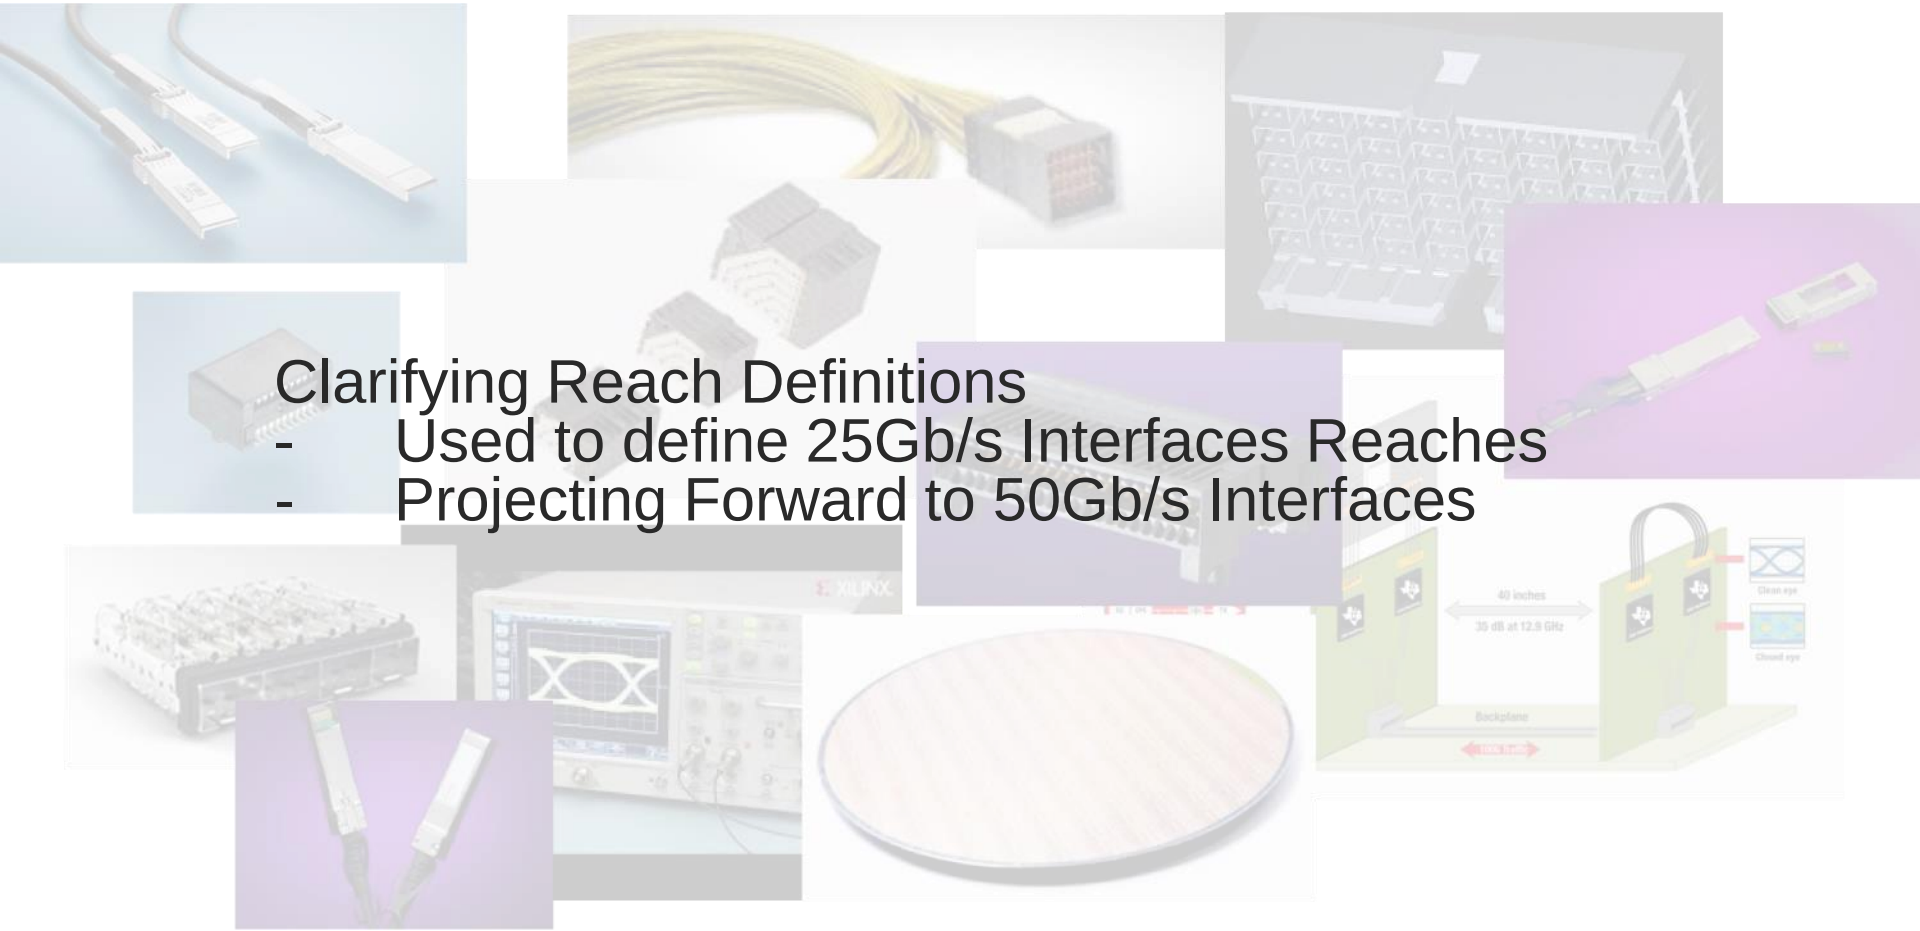

Clarifying Reach Definitions

- Used to define 25Gb/s Interfaces Reaches
- Projecting Forward to 50Gb/s Interfaces

Length, Loss & Applications

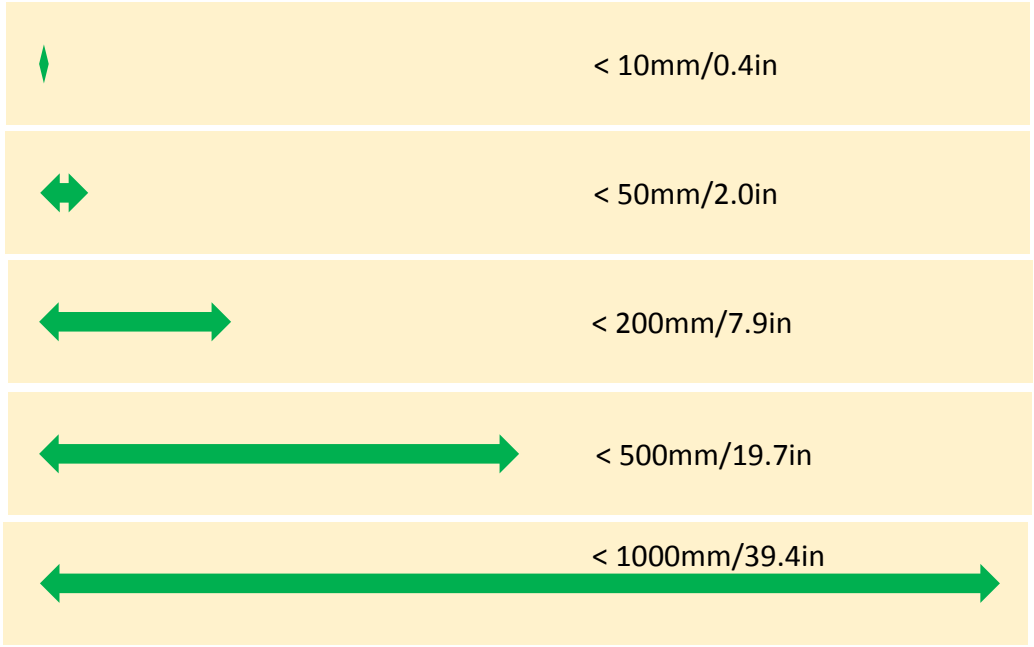

	IL	
USR	1.5dB@14GHz 3dB@28GHz	Bump-to-bump Inside MCM or 3D Stack
XSR **C2EO**	4dB@14GHz 8dB@28GHz	Ball-to-ball Across PCB
VSR **C2M**	10dB@14GHz 20dB@28GHz	Ball-to-ball
MR **C2C**	20dB@14GHz 40dB@28GHz	Ball-to-ball
LR **C2F**	35dB@14GHz	Ball-to-ball

Length, Loss & Applications:

Today a single 25GB/s SERDES core can cover all these ranges

These ranges can be easily covered with one SERDES core today for 25Gb/s.

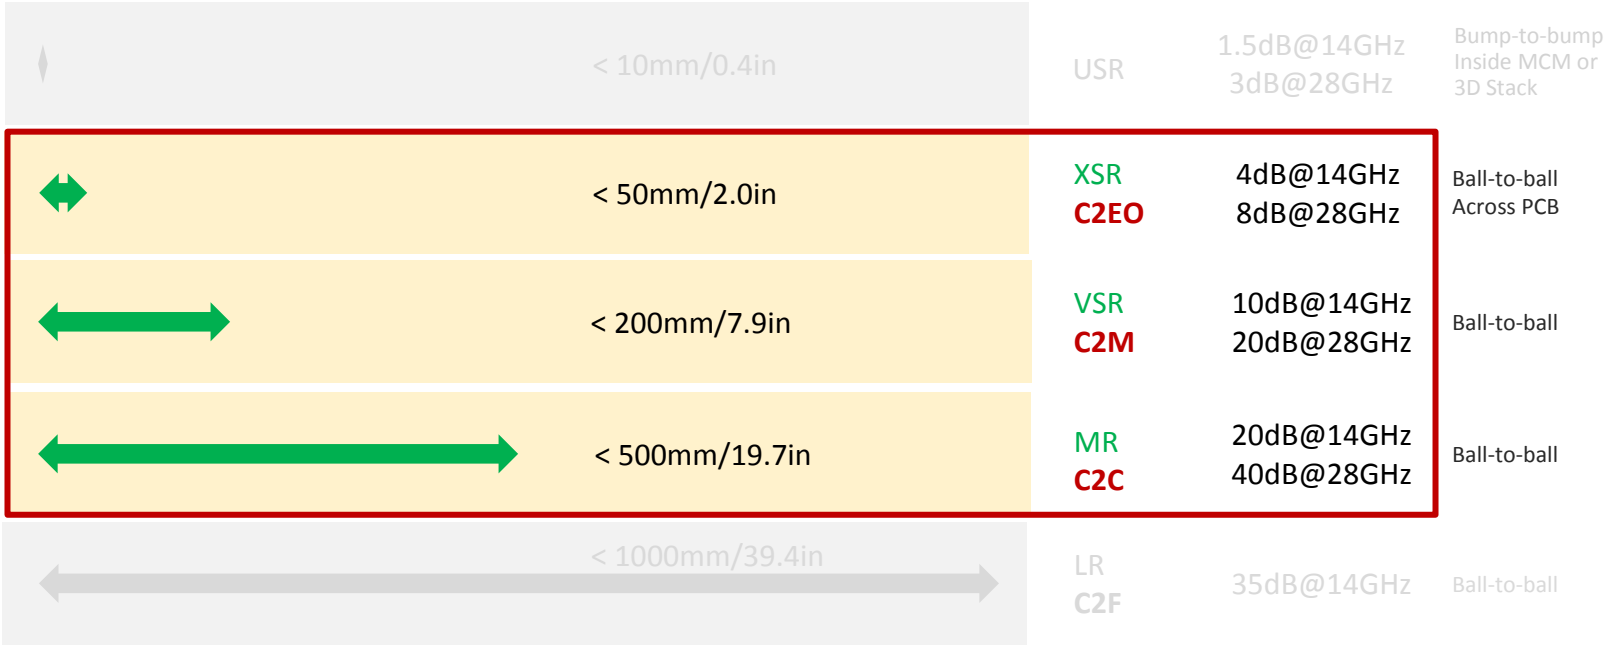

Length, Loss & Applications: Grouping (Perspective 1)

How does a 50Gb/s SERDES core cover these ranges optimally?

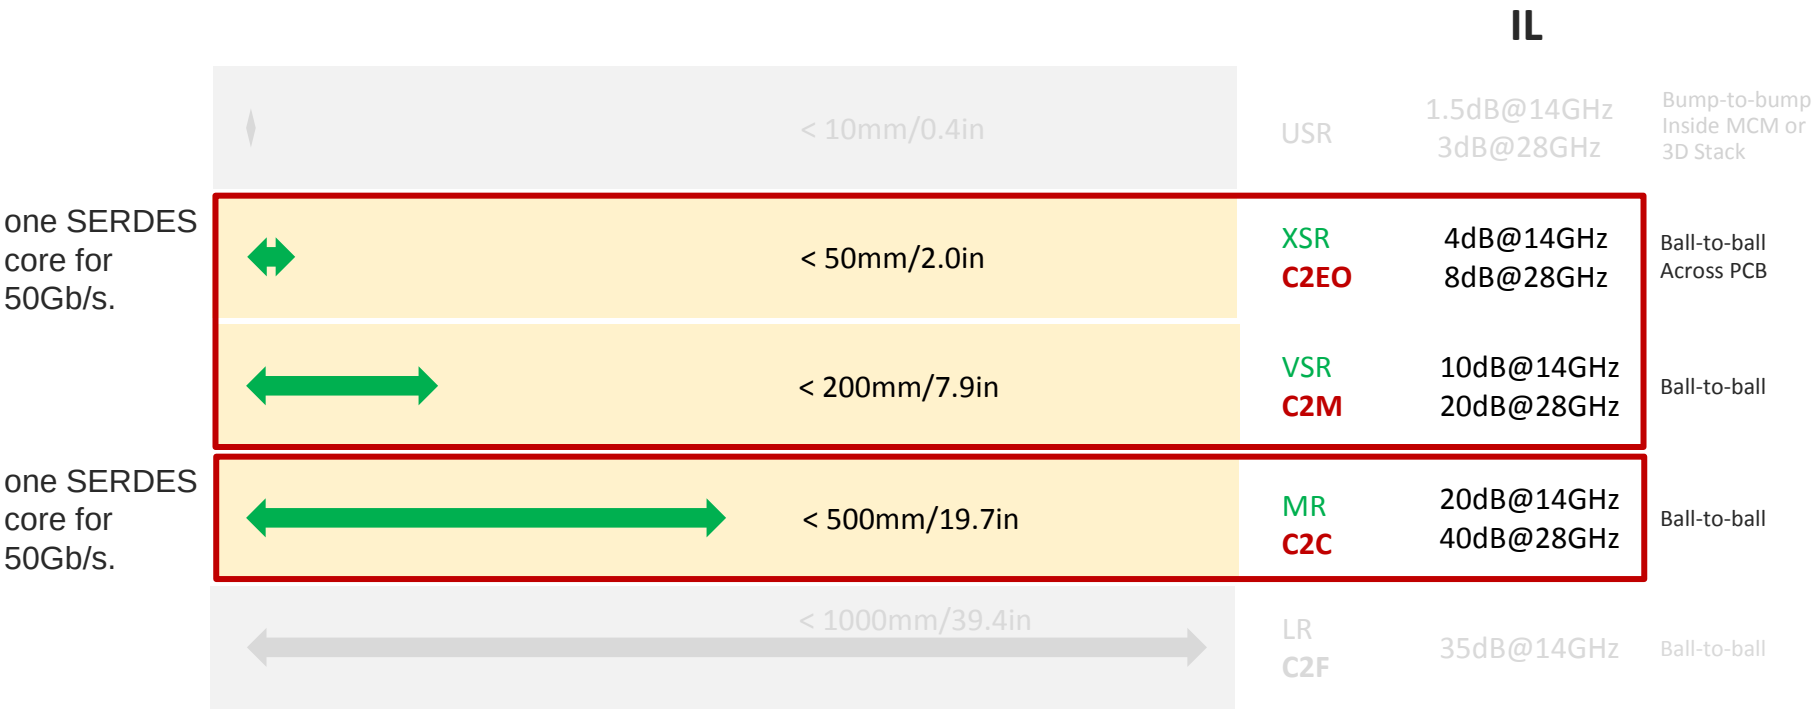

Unconfirmed but … Under the impression from the group the cores should be compatible

Length, Loss & Applications: Grouping (Perspective 2)

How does a 50Gb/s SERDES core cover these ranges optimally?

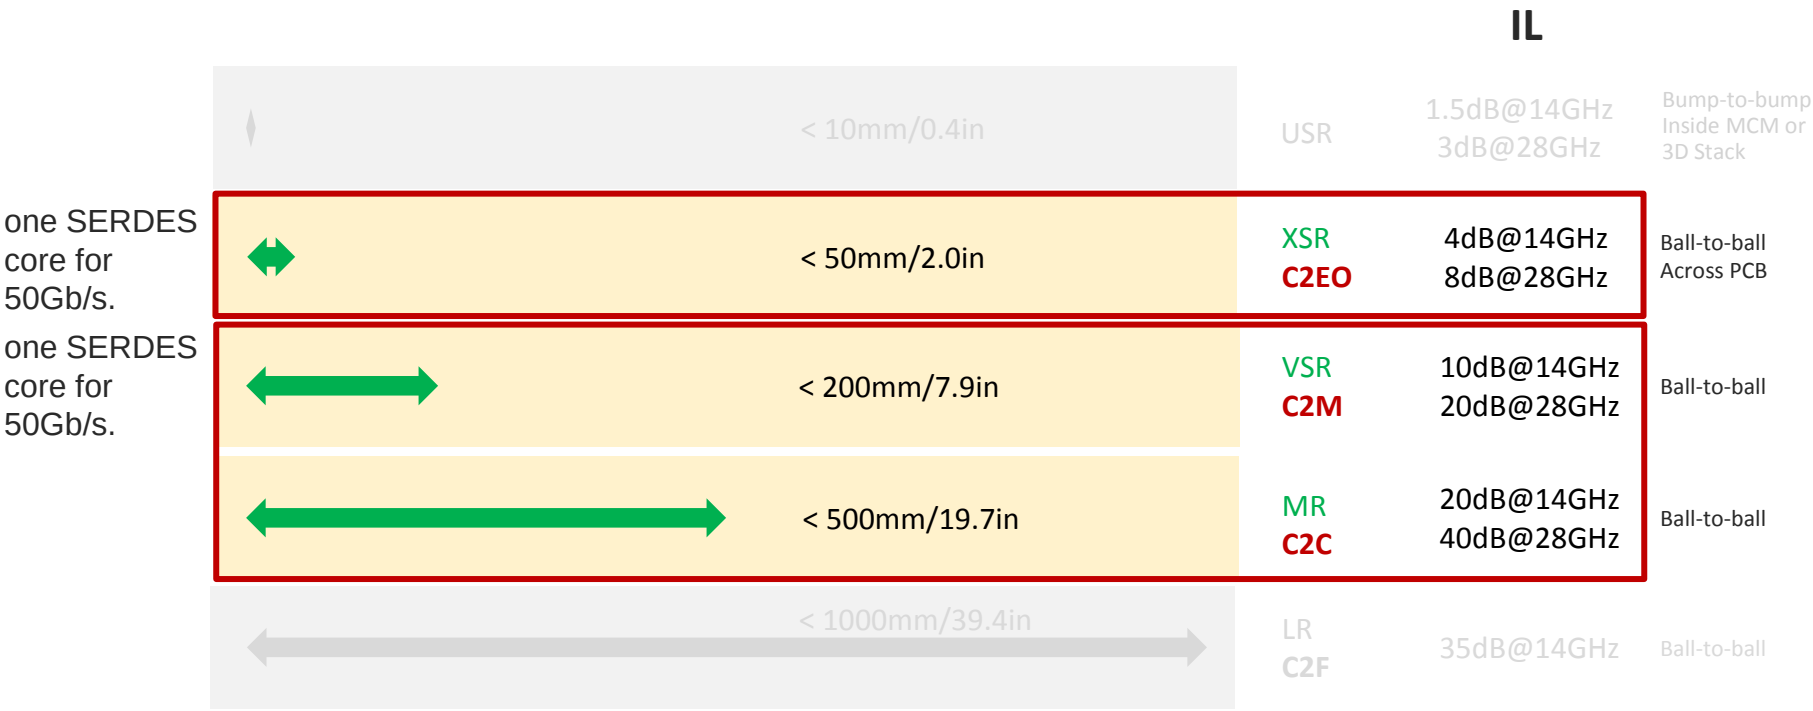

Unconfirmed but … Under the impression from the group the cores should be compatible

Length, Loss & Applications:

Future: a single 50GB/s SERDES core to cover all these ranges

These ranges should be covered with one SERDES core in the future for 50Gb/s.

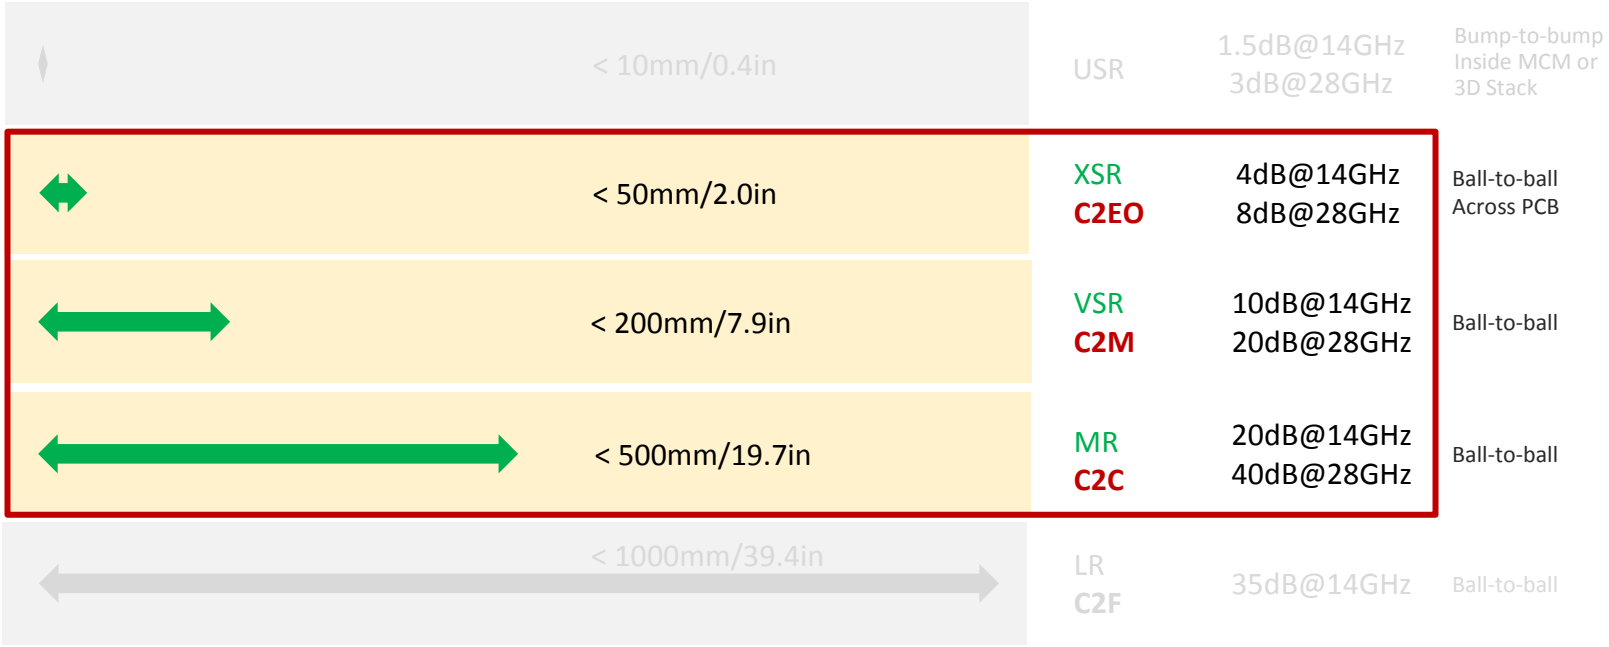

Length, Loss & Application: Technologies for 50Gb/s

These Values are under discussion

Application	Length	Loss	Modulation	pJ/bit	DFE?	FEC?
C2EO (XSR)	< 2in	<4dB@14GHz	PAM-4	TBD	TBD	TBD
		<8dB@28GHz	NRZ	TBD	TBD	TBD
C2M (VSR)	2-8in	4-10dB@14GHz	PAM-4	TBD	TBD	TBD
		8-20dB@28GHz	NRZ	TBD	TBD	TBD
C2C (MR)	8-20in	10-20dB@14GHz	PAM-4	TBD	TBD	TBD
		20-40dB@28GHz	NRZ	TBD	TBD	TBD

Knowing the reach definition allows us to begin understanding the next steps in the consensus building process

- System Architecture
- Channel Loss
- Modulation
- Equalization
- Error Correction
- Power

Looking At System Architecture


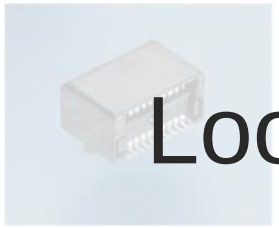

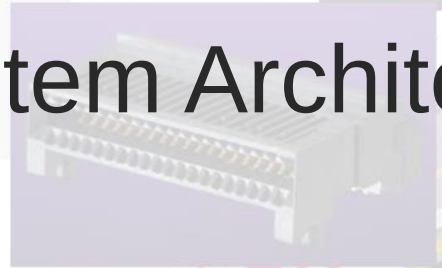


The Early Adopter Option

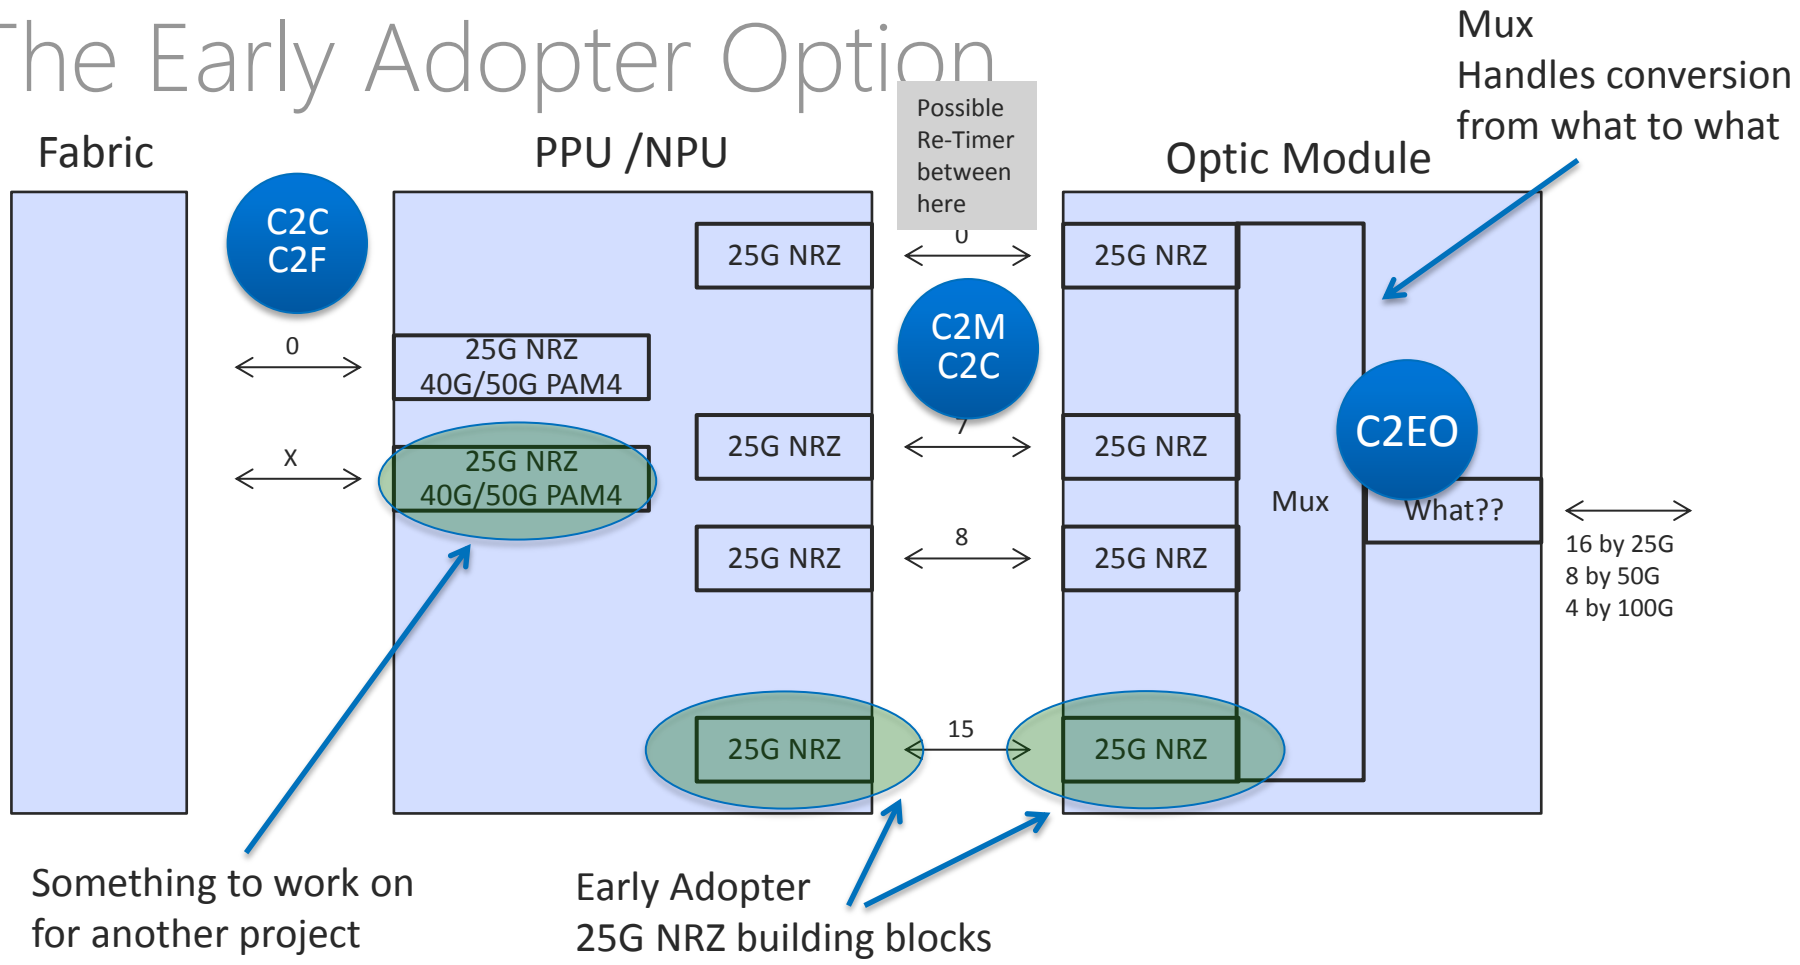

The Advanced Early Adopter Option

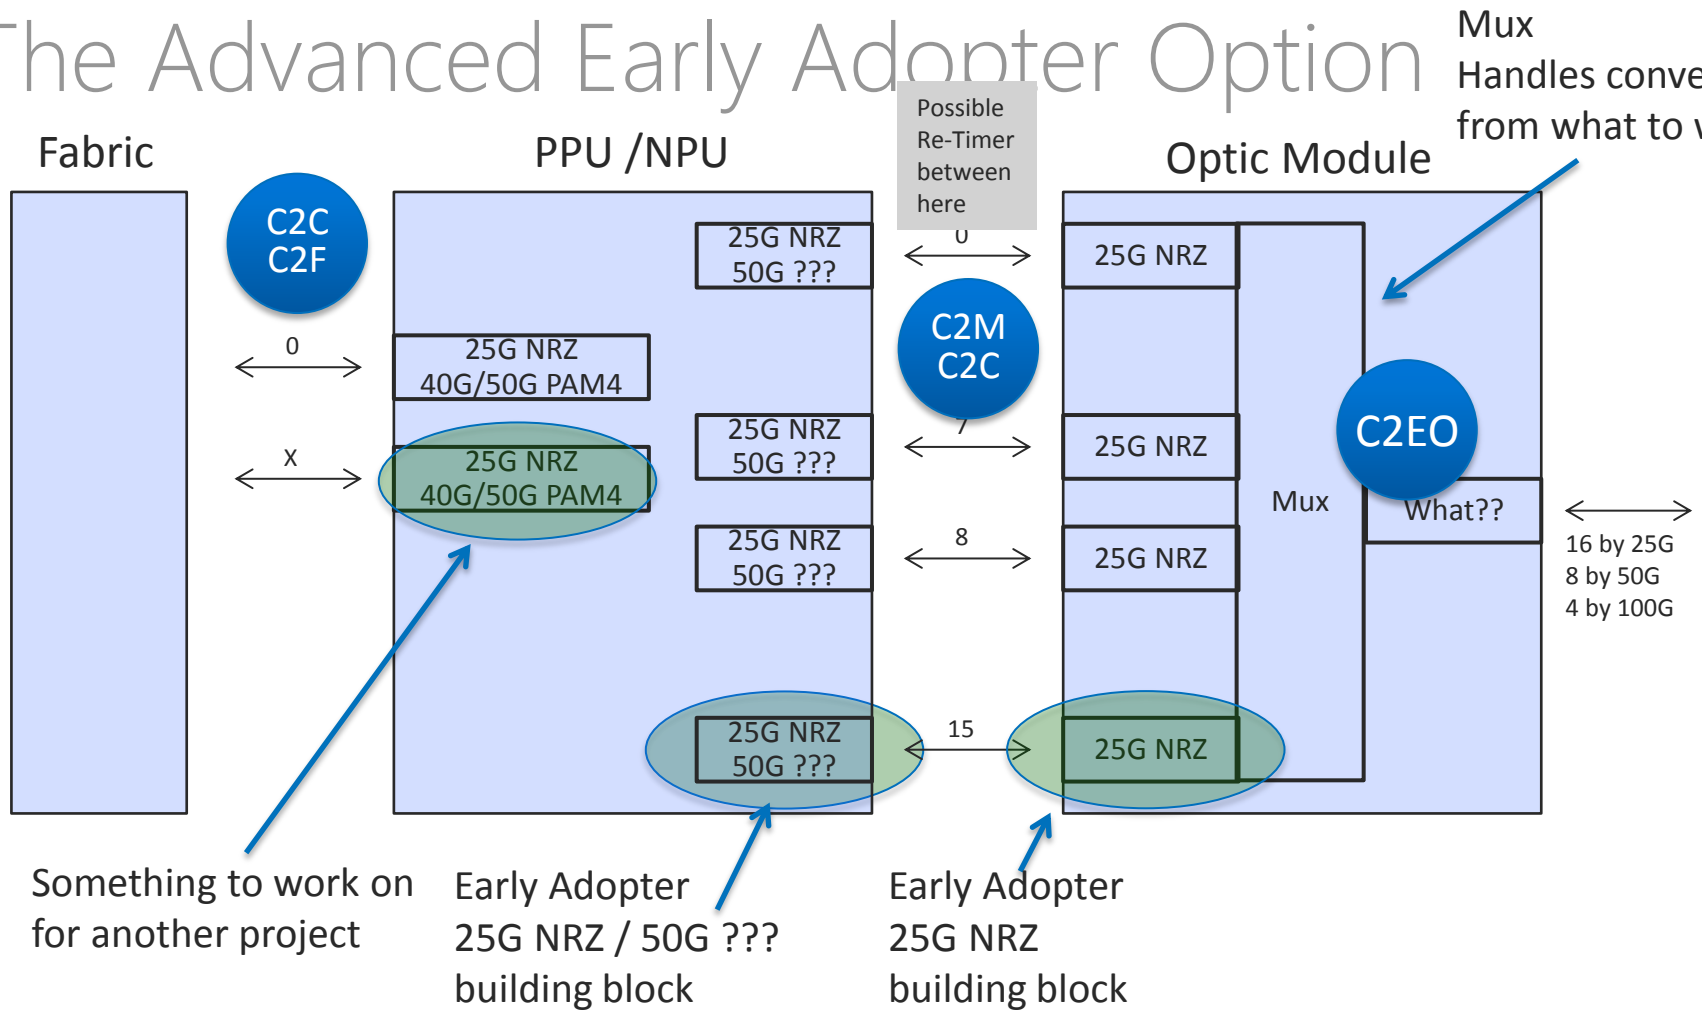

The Evolved Option

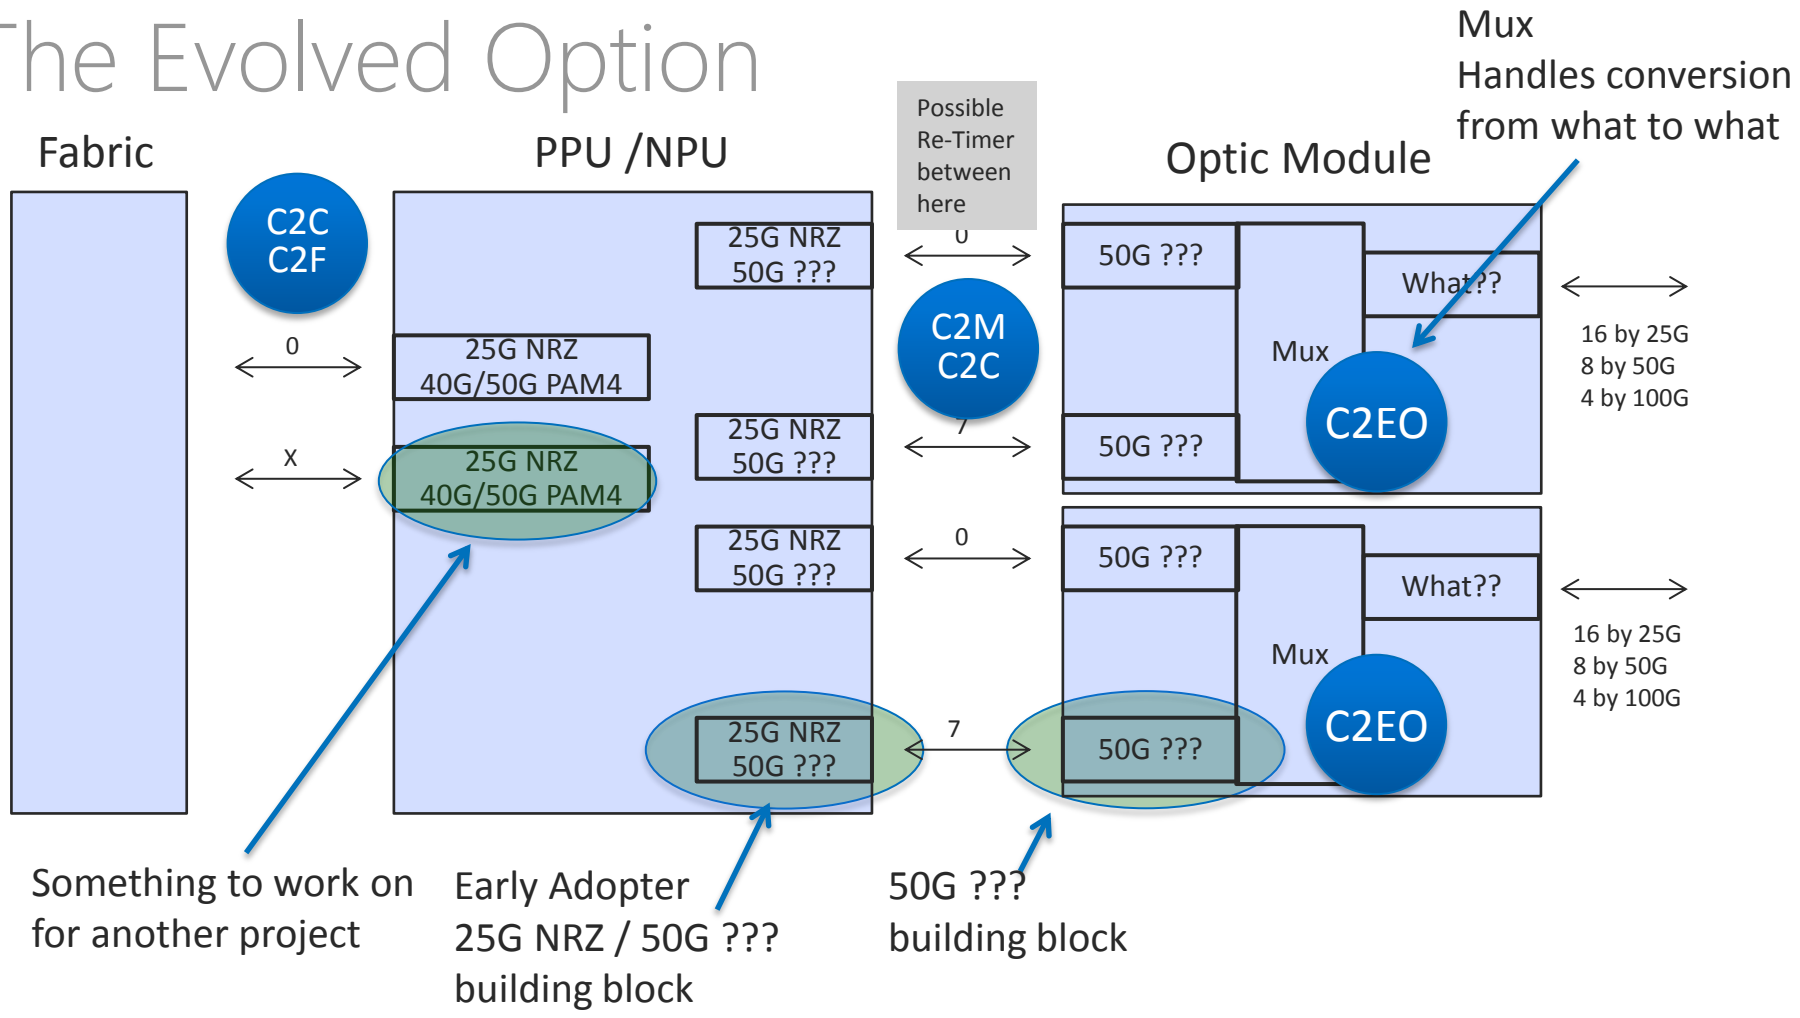

The Maybe-Some-day Option

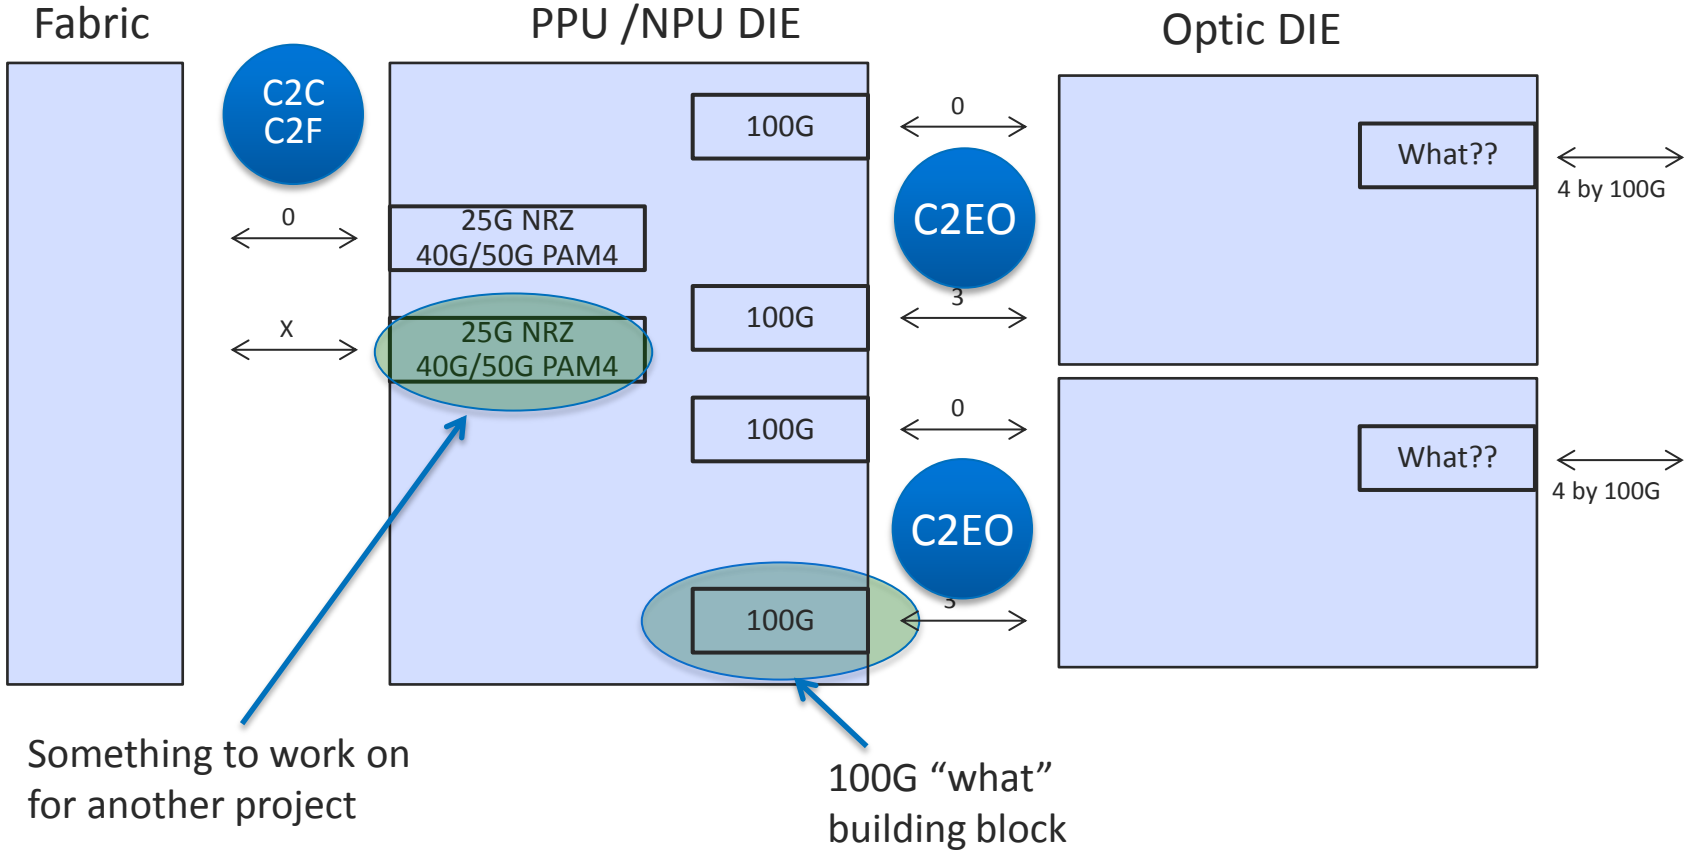

Some Thoughts

Need to continue to converge these options with work Mark Gustlin and team are addressing.

Looking at Channel Implementations For 50Gb/s




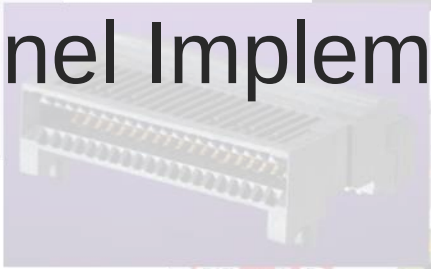

Chip-to-Embedded Optics (C2EO)

Possible Implementations

- System in package (SiP)
- 2.5D/3D Silicon interposer
- Stacked die
- Multi-chip module
- Package-on-package

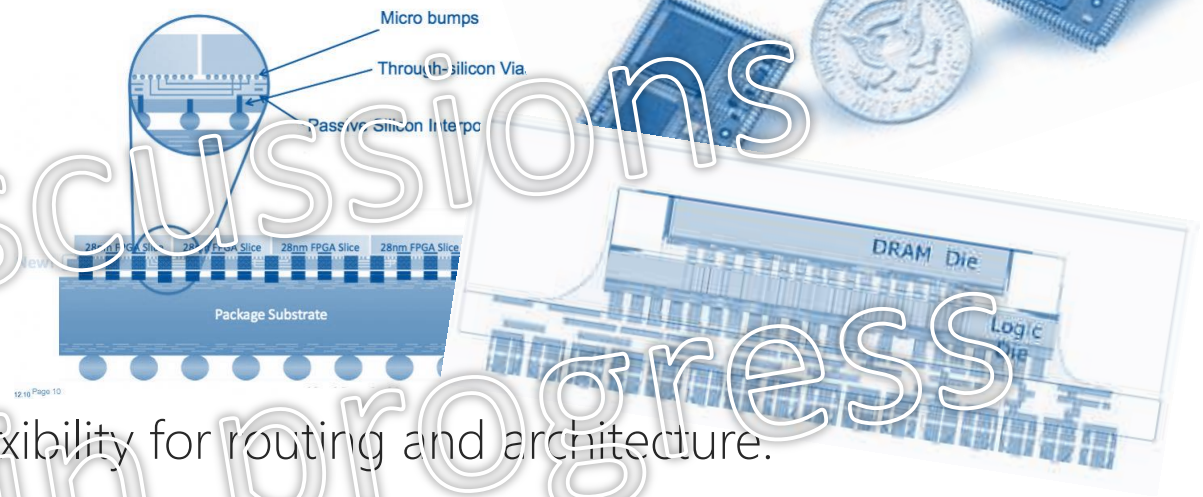

Offers system advantage/flexibility for routing and architecture.

Discussions still needed to have:

Reasonable to see C2EO interfaces commonly in industry by ____

Die-to-Die definition (instead of ball to ball)?

Can this be defined in the standards or are these proprietary links?

Discussions are in progress

Chip-to-Module (C2M)

No leap changes in this market from 100G Ethernet; albeit incremental changes have are seen (materials, connectors, quality, system constraints, etc.)

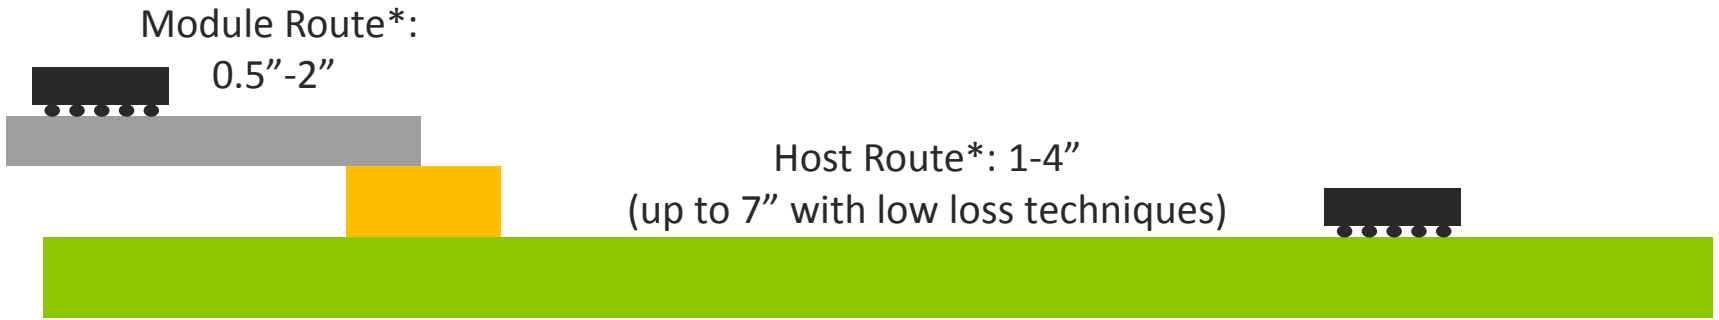

Ball-to-ball definition (ball meaning BGA on the outside of the package)?

* Looking across the industry, across multiple platforms… typical channel length ranges shown.

Chip-to-Module (C2M) – What does that mean for loss?

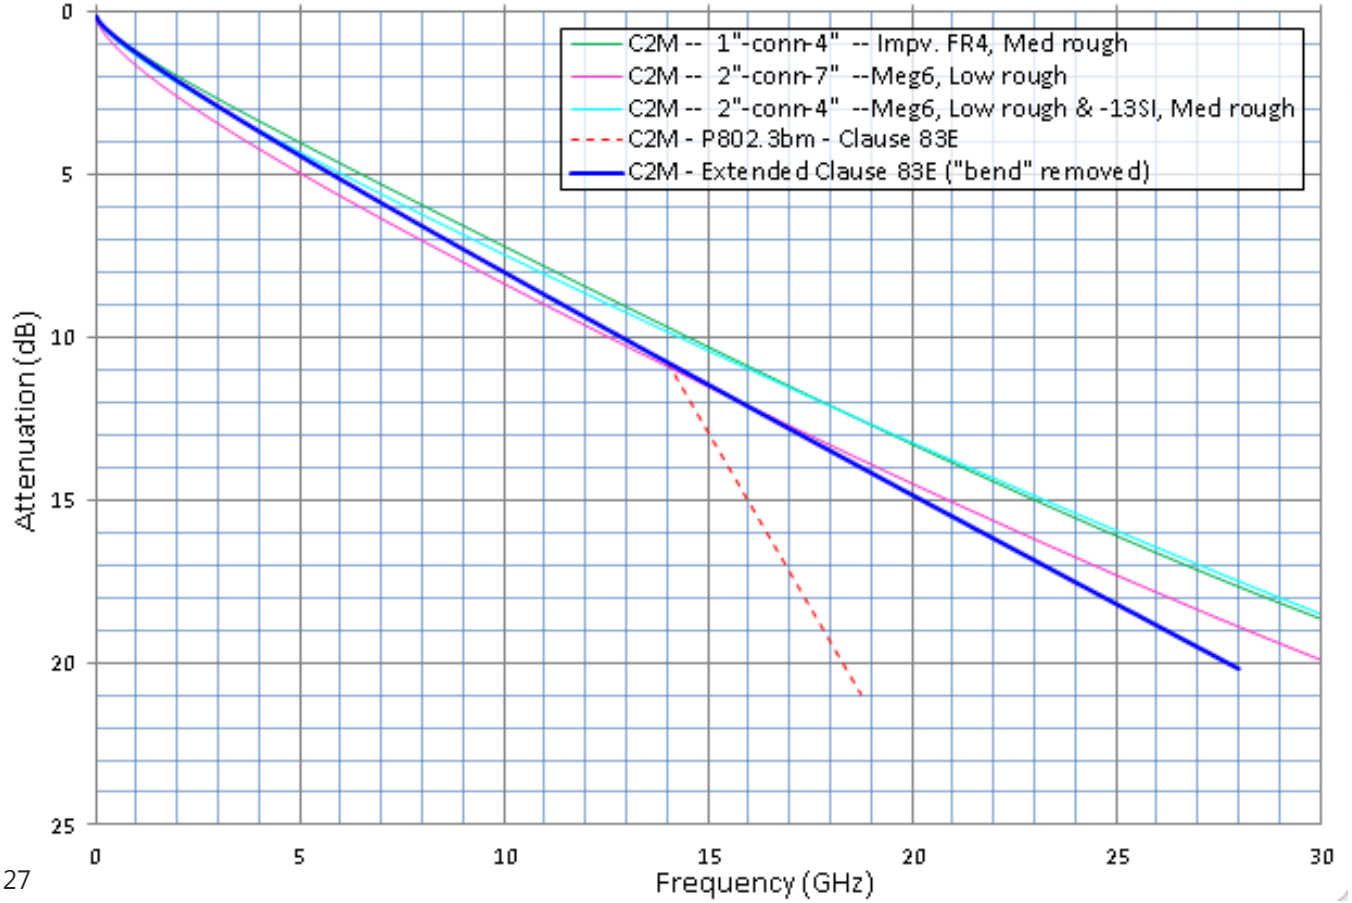

Notes:

- Uses DkDf_AlgebraicModel_v2.05. Where loss is calculated using an algebraic model for dielectric, conductor, and connector loss.
- The 4" host trace is and should be supported at mid-loss materials by 100G standards… 7" is certainly seen in designs, but industries seem to recognize and adjust with material and design tradeoffs.

Chip-to-Module (C2M) – Comparing channel data

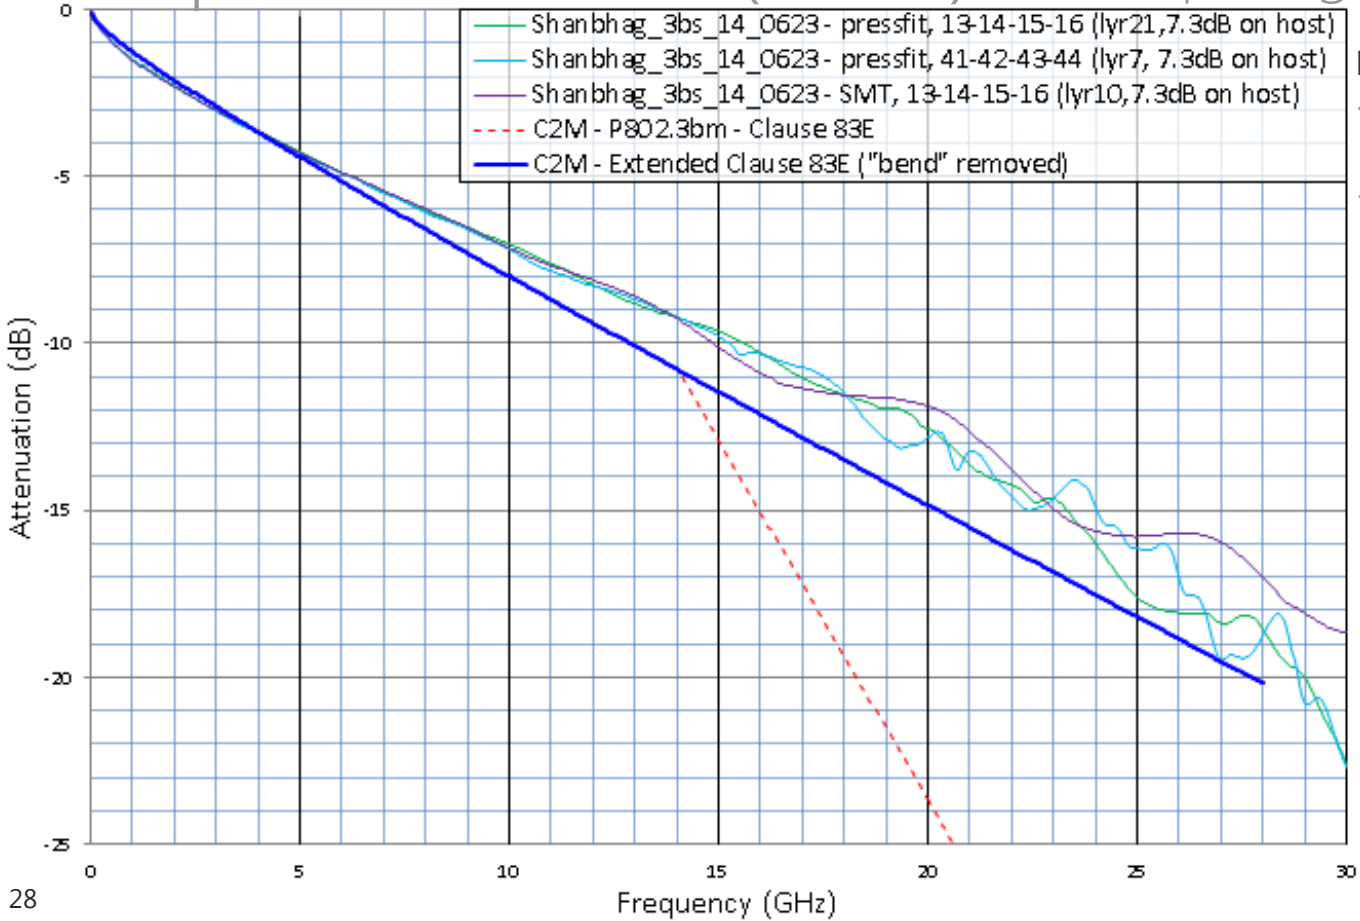

Notes:

- TE Channels developed using HFSS/ADS modeling tools.
- The 4″ host trace is and should be supported at mid-loss materials by 100G standards… 7″ is certainly seen in designs, but industries seem to recognize and adjust with material and design tradeoffs.

Propose: Use extended CAUI4 C2M–like channels for Modulation discussion/comparison

Chip-to-Chip (C2C)

- No leap changes in this market from 100G Ethernet; albeit incremental changes have are seen (materials, connectors, quality, system constraints, etc.)
  - However, there is a push in industry to longer links… while making design tradeoffs
  - How far should C2C cover? Let's look at loss!

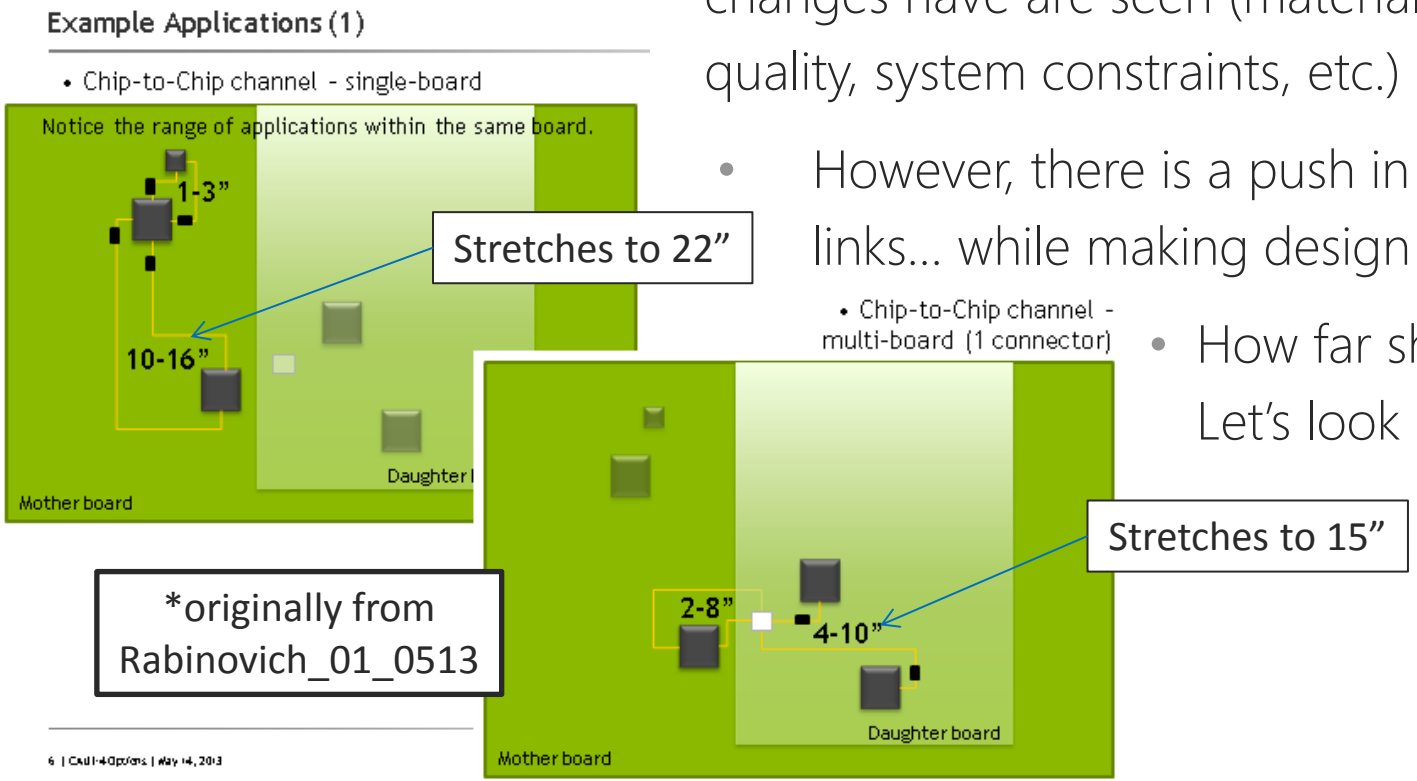

Chip-to-Chip (C2C) – What does that mean for loss?

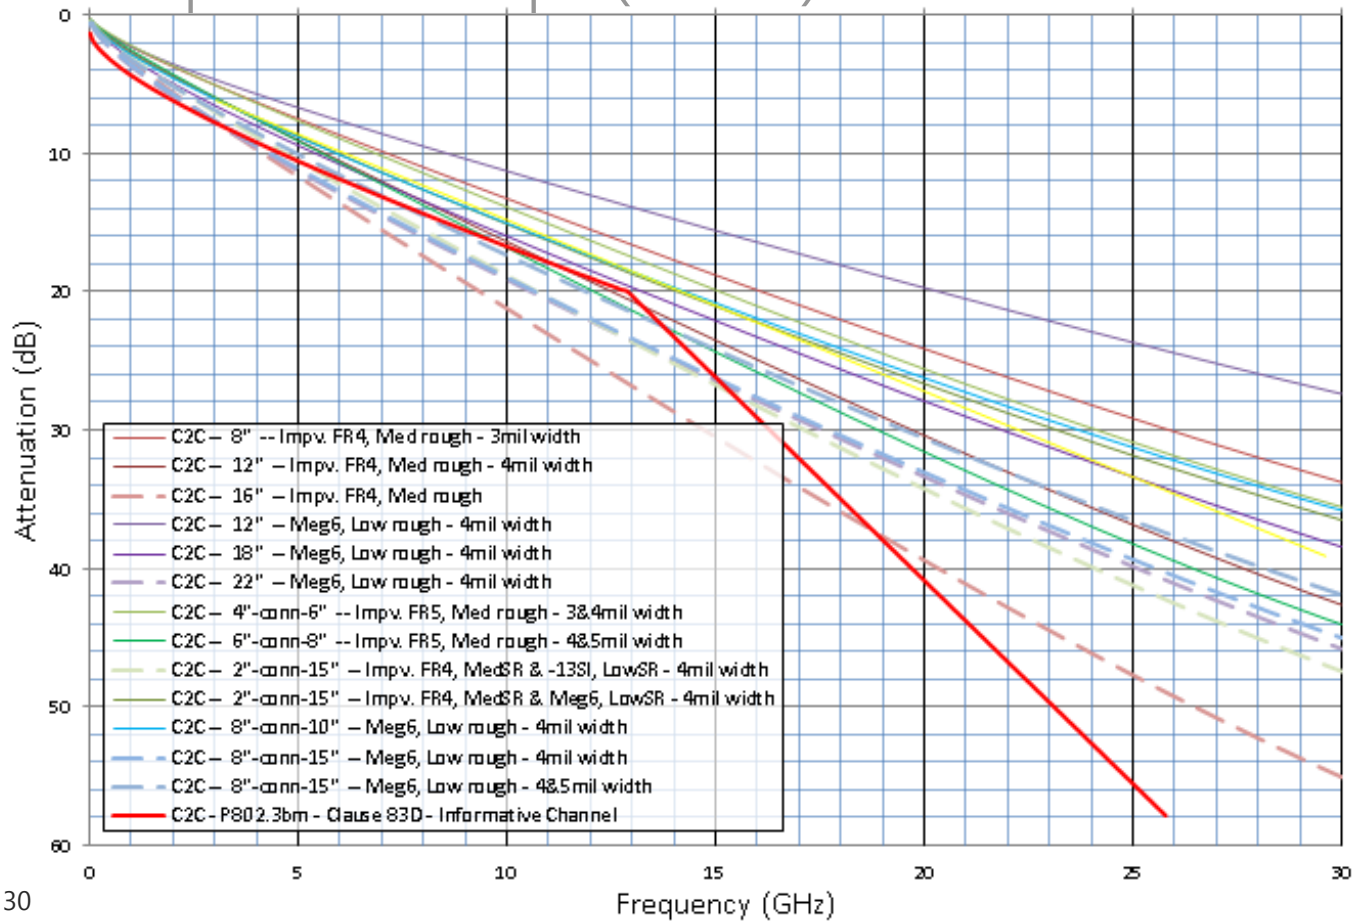

Notes:

- Uses DkDf_AlgebraicModel_v2.05. Where loss is calculated using an algebraic model for dielectric, conductor, and connector loss.
- The typical routes seem to be covered by 20dB at 12.87G… The longer links that industry stretches to reach are not.

Chip-to-Chip (C2C) – What does that mean for loss?

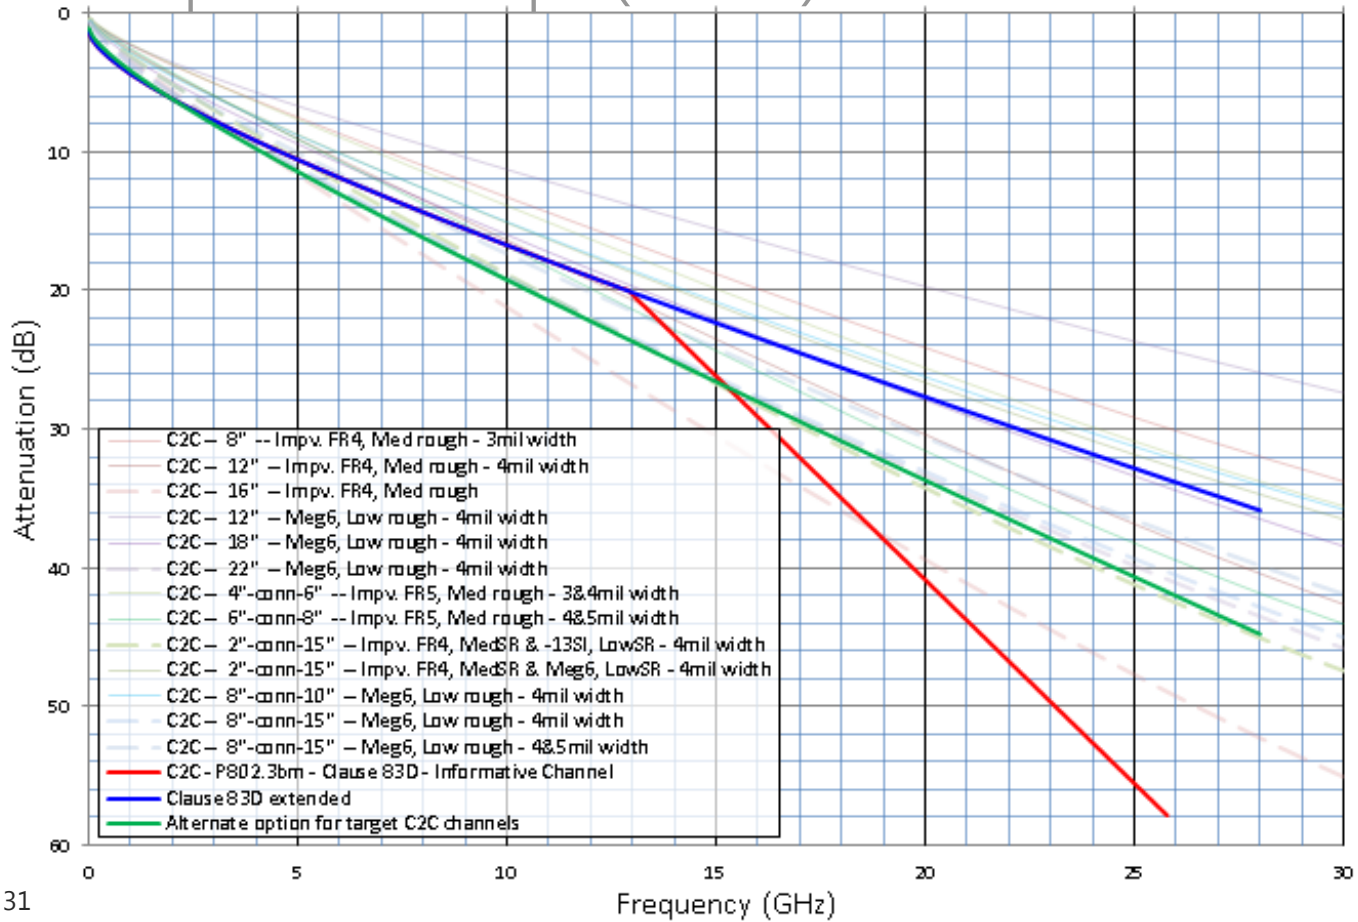

Notes:

- Uses DkDf_AlgebraicModel_v2.05. Where loss is calculated using an algebraic model for dielectric, conductor, and connector loss.
- The typical routes seem to be covered by 20dB at 12.87G… The longer links that industry stretches to reach are not.
- Do we stay with the current CAUI4 coverage or should we cover more channels now seeing industry's range of implementation??

Chip-to-Chip (C2C) – Let's compare to channel data.

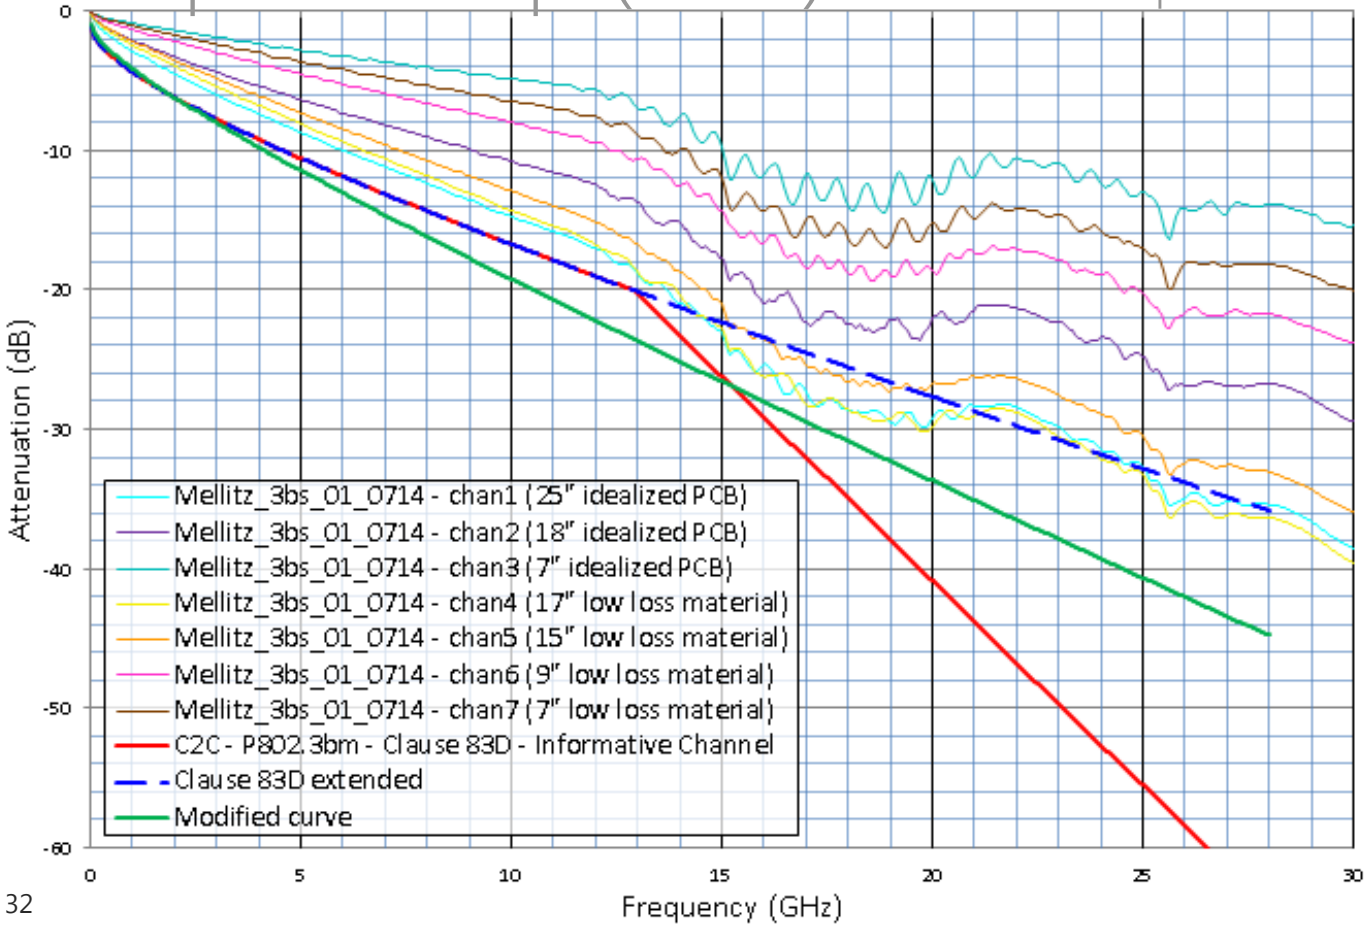

Notes:

← Channels are public on the .3bs webpage… all channels include connector and an 8% impedance variation from motherboard to daughtercard.

←Are these channels right to use for modulation discussion?

←Is ILD pessimistic for educated 50G channel design?

Chip-to-Fabric (C2F – previously Backplane)

Used for channels too long for C2C, but comes with a power/complexity penalty.

Pure loss is becoming a constricting factor for C2F designs (25-35dB of 100G standards generation). Backplanes are becoming more and more diverse

- Typical FR4 backplane
- Planned repeaters
- Cable backplane
- Mezzanine connectors
- Optics backplane

Out of scope for current project, but good to note as it is part of electrical infrastructure.

Equations plotted for C2C and C2M

802.3bm draft – C2M equation (red curve)

$$Insertion\_loss(f) \le \begin{cases} 1.076(0.075 + 0.537\sqrt{f} + 0.566f) & 0.01 \le f < 14 \\ 1.076(-18 + 2f) & 14 \le f < 18.75 \end{cases} \text{ (dB)} \qquad \text{(83E–1)}$$

802.3bm draft – c2C equation (red curve) *although, standard uses COM as the normative spec, the following is offered as the informative "limit" line

$$Insertion\_loss(f) \le \begin{cases} 1.083 + 2.543\sqrt{f} + 0.761f & 0.01 \le f < 12.89 \\ -17.851 + 2.936f & 12.89 \le f < 25.78 \end{cases} \text{ (dB)} \qquad \text{(83D–1)}$$

C2M and C2C "extended curve" (blue curves)

Simply extend first portion of above curves for full frequency range; omit 14-18.75G and 12.89-25.78G equations, respectively.

Suggested/modified curve drawn in C2C (green curve)

$$0.9 + 2.1\sqrt{f} + 1.17f$$

Starting the Modulation Discussion For 50Gb/s

Driving Modulation for 50Gb/s

Application	Length	Loss	Modulation	pJ/bit	DFE?	FEC?
C2EO (XSR)	< 2in	<4dB@14GHz	PAM-4	TBD	TBD	TBD
		<8dB@28GHz	NRZ	TBD	TBD	TBD
C2M (VSR)	2-8in	4-10dB@14GHz	PAM-4	TBD	TBD	TBD
		8-20dB@28GHz	NRZ	TBD	TBD	TBD
C2C (MR)	8-20in	10-20dB@14GHz	PAM-4	TBD	TBD	TBD
		20-40dB@28GHz	NRZ	TBD	TBD	TBD

Does filling out this chart help us close modulation discussions? Yes
Do you want to see companies provide details? Yes
Should we add a column on silicon technology?

Next Steps to Consensus


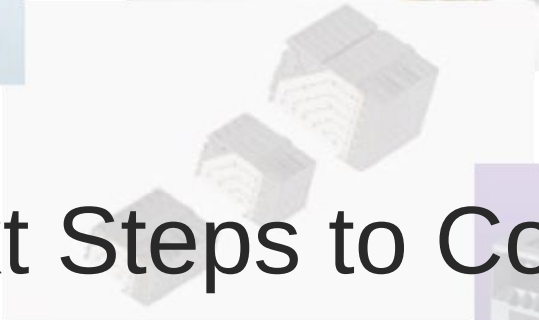
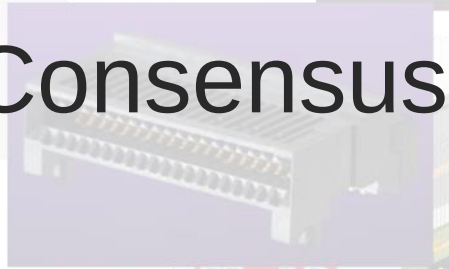


Between Now and The November Plenary

Determining the following:

Power per bit / Best fit for Reach

Modulation

Channel Loss

FEC (Y/N) and Equalization requirements

Thoughts on Past Discussions

For a 25Gb/s Electrical Interface

Adopt a 25Gb/s by 16 lane electrical C2C and C2M interface defined by .3bm specifications, using current values as starting baseline text.

For a 50Gb/s Electrical Interface

Adopt a 50Gb/s by 8 lane electrical C2C and C2M interface defined by specifications yet to be provided. Use loss/length definitions and algebraic base line channel discussions to begin baseline text and specifications.

For a 100Gb/s Electrical Interface

Should we continue C2EO discussions for this?

Thank you!

From:
Vasu Parthasarathy
Beth Kochuparambil
Vivek Telang
Joel Goergen