400GE Electrical Interface Thoughts

Joel Goergen – Distinguished Engineer / Cisco Systems

05May2014 ver 02

Abstract: Describe the four basic electrical interconnect building blocks and potential guidelines for each (C2M, C2C, C2F, and C2ISP).

Overview

- Channel Manufacturing Basics
- Defining Serial Interconnects
- Interconnect Power Efficiency
- Interconnect Basic Requirements

Technology Contribution: Trace Widths / Loss Variation

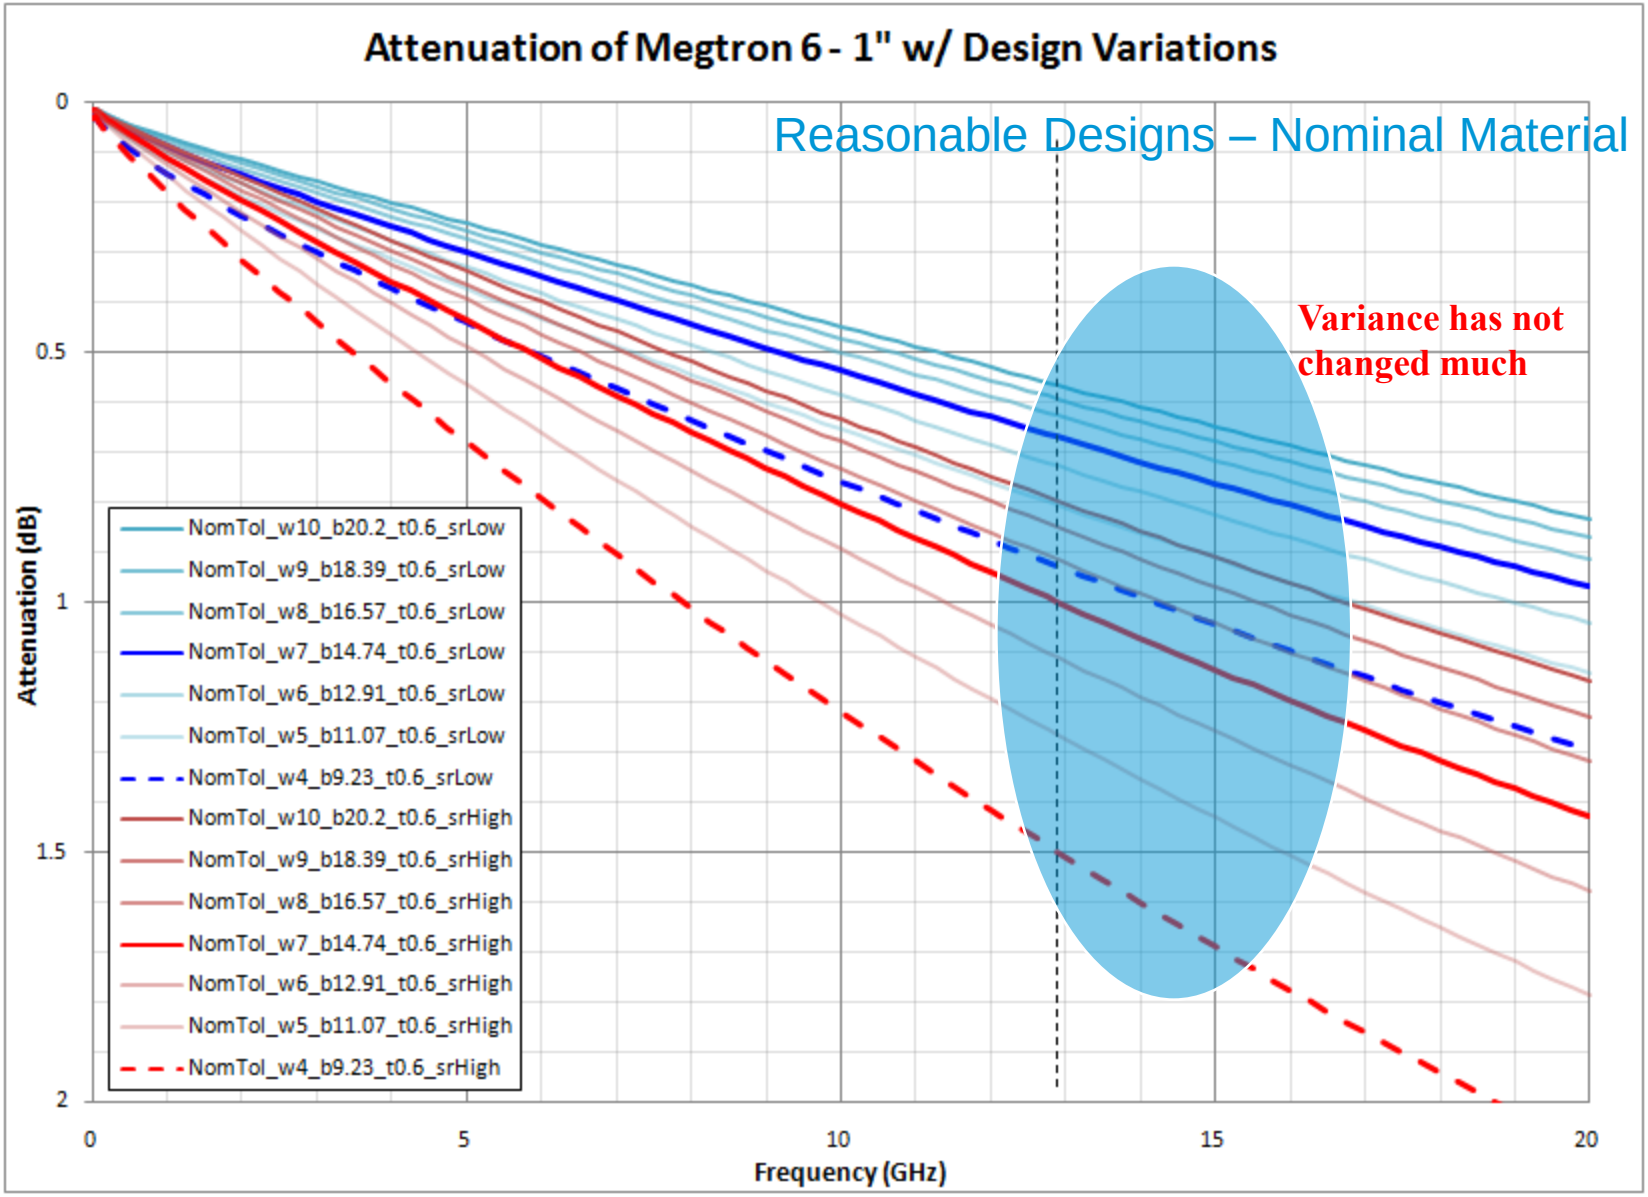

Technology Contribution: Drill Usage

Drill Bit Life Expectancy
Relative Drilled Holes on a Drill Bit

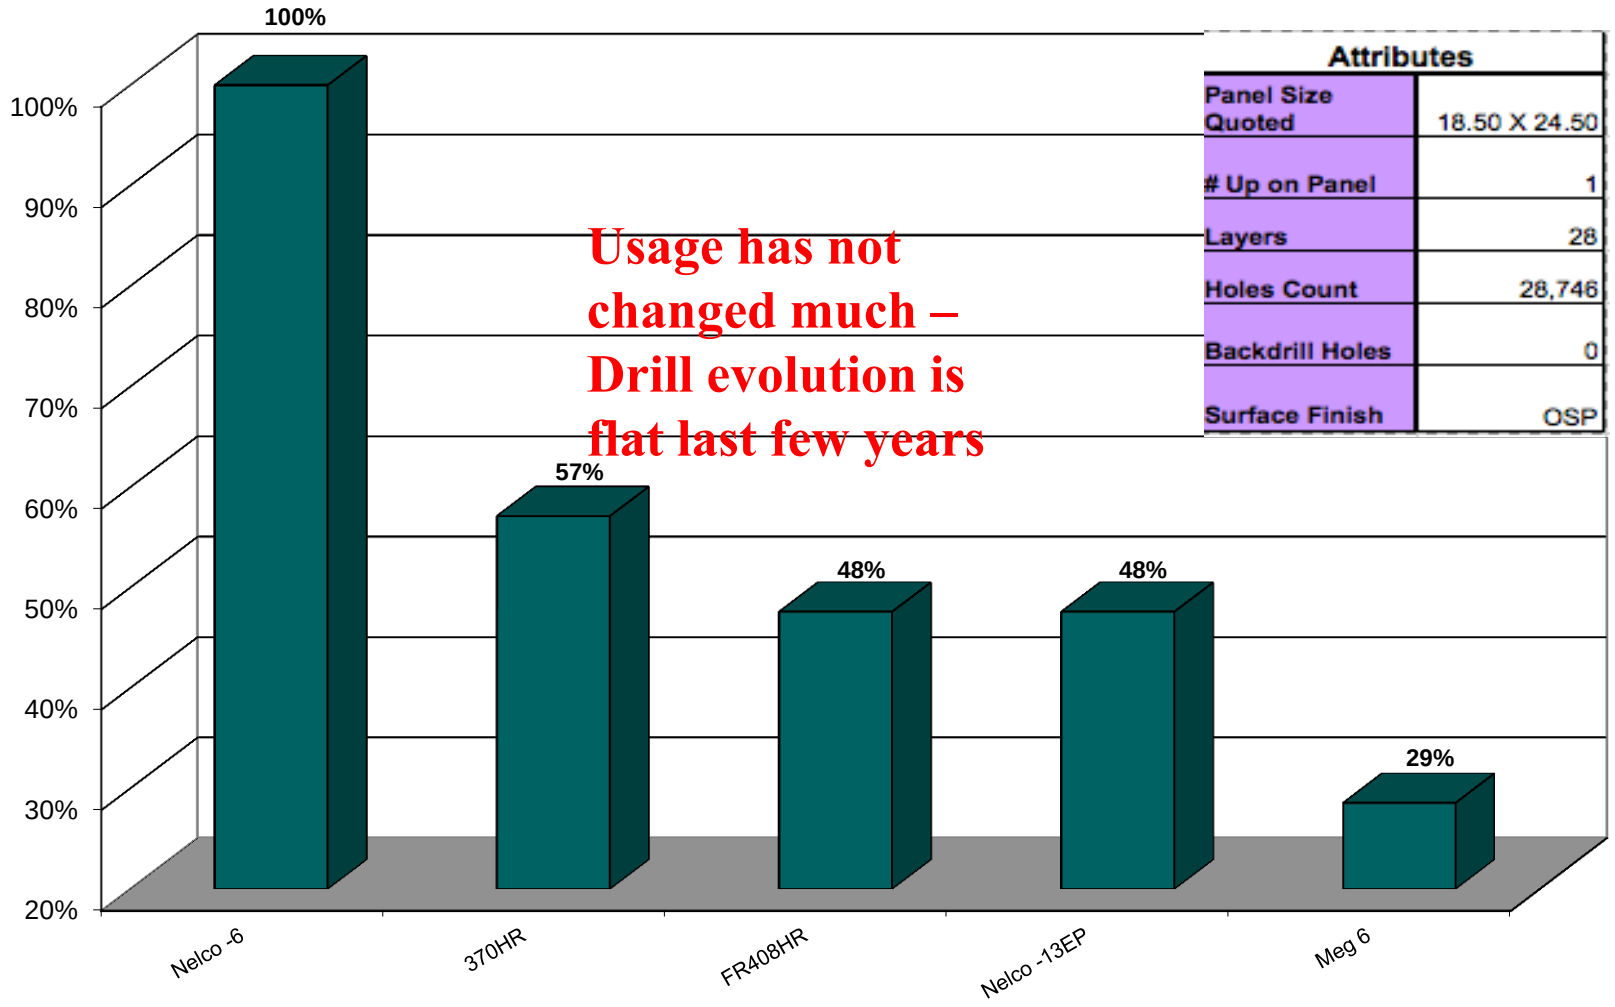

DC Blocking Capacitor Field

- Consumes a large area
- Requires considerable grounding to reduce common mode noise, control cross talk, and limit radiated emissions.
- A lot of solutions in ASICs to replace the blocking cap – but come with system trade-offs and vendor interoperability implementations.

Technology Contribution: VIA in PAD Plated Over

Item	Comment
Signal Integrity	Eliminate dogbones
Routability	Very little – freed up outer layer area, but requires outer layer features and spaces
Reliability (SnPb)	Proven
Reliability (Pb-free)	No issues found yet
Supply Base	Large
Process Complexity	Moderate – additional plating, epoxy fill and planarization
Cost	~20% adder dependant on tech level, layer count, etc
Hidden Cons	Restricted OL feature size Restricted OL spacing

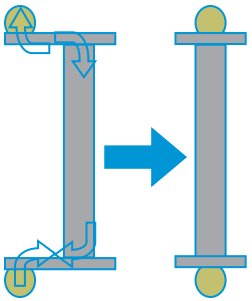

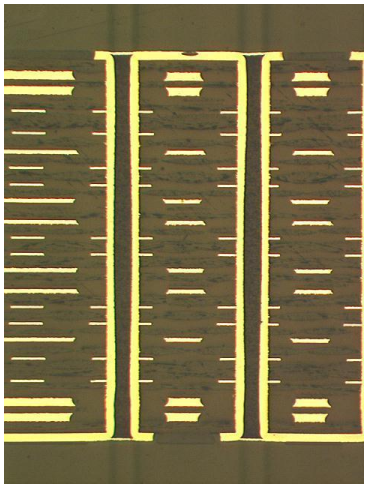

Technology Contribution: Blind VIA / Micro VIA

Item	Comment
Signal Integrity	Small, stubless via
Routability	Freed up space on layer below via
Reliability (SnPb)	Proven
Reliability (Pb-free)	Proven
Supply Base	Large
Process Complexity	Minimal – laser drilling and microvia plating…pretty common technology for most suppliers.
Cost	~5-15% (Conformal plated) ~15-40% (Cu fill plated)
Hidden Cons	Some design tools not 100% optimized

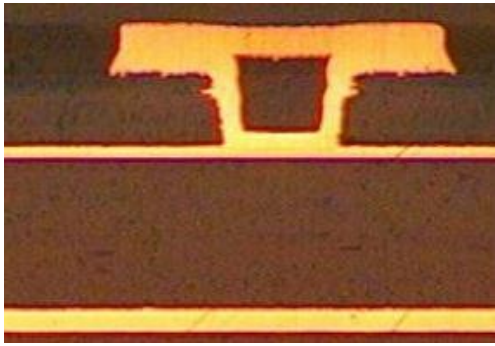

Technology Contribution: Skip VIA

Item	Comment
Signal Integrity	Stubless via connection for high speed signals
Routability	Freed up space on layer below via (1to3 and 1to4)
Reliability (SnPb)	Passed
Reliability (Pb-free)	Passed
Supply Base	Limited (6mil dia std, 4mil is advanced)
Process Complexity	Moderate – complex laser drilling process. Still has limitations.
Cost	~15-20% (Conformal plated) ~30-40% (SKIPPO) Adds processing days
Hidden Cons	Prone to laminate cracking below via

Technology Contribution: Back Drilling

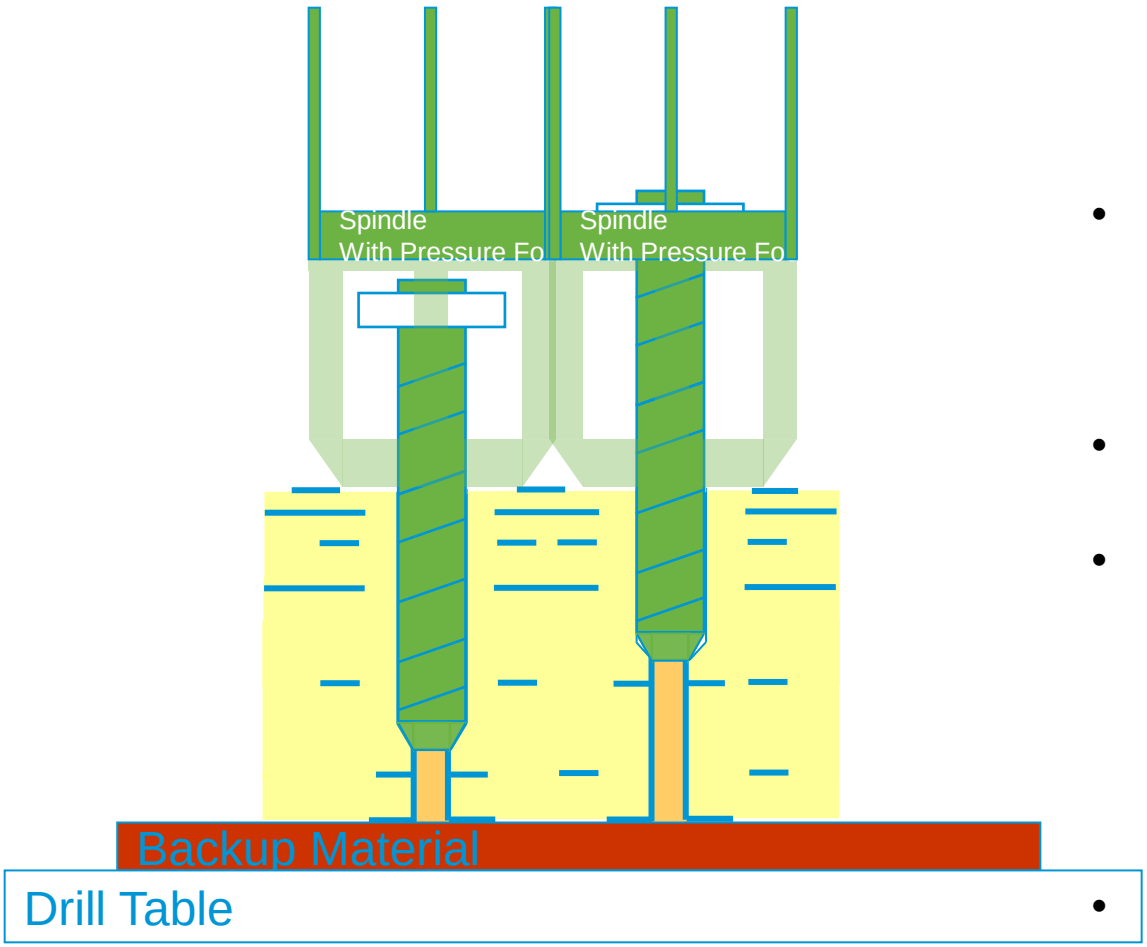

- *Back Drilling is a well defined process with < 4% cost impact on fr4 and < 8% on MEG6.*
- Stop depth tolerance can be as low as +/- 5 mils but often is in the range of +/- 10 mils.
- Removes a significant portion of the stub.
- Don’t be afraid to deploy this fabrication technology. Seldom used in 2000, this technique is used today in almost all high speed designs.
- Depth control is difficult on large, thick panels.

Technology Contribution: Copper Surface Roughness

- Much work has been done here.
- Impact at 10Gbps is not worth the added costs.
- Impact at 25Gbps shows improvement.
- ***Impact at higher speeds is very relevant.***

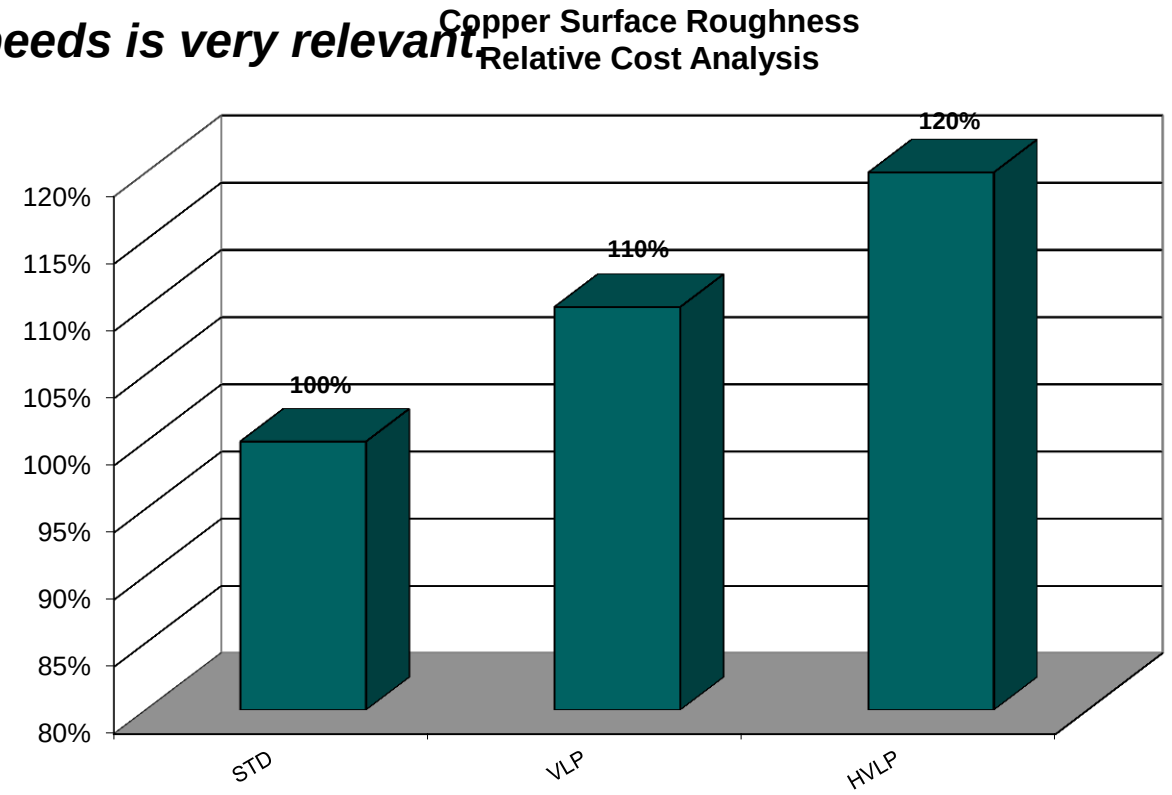

Take Away Points

- Remove the stubs
- Line widths are getting smaller / Pitch is getting smaller
- Cu thickness is getting thinner to accommodate the fine geometries
- Make use of skip vias and micro vias
- Surface roughness, Cu thickness, and via to trace bonding will play key roles in next generation channel models.

Defining the Serial Interfaces
4 Basic Types: C2M / C2C / C2F / C2EO

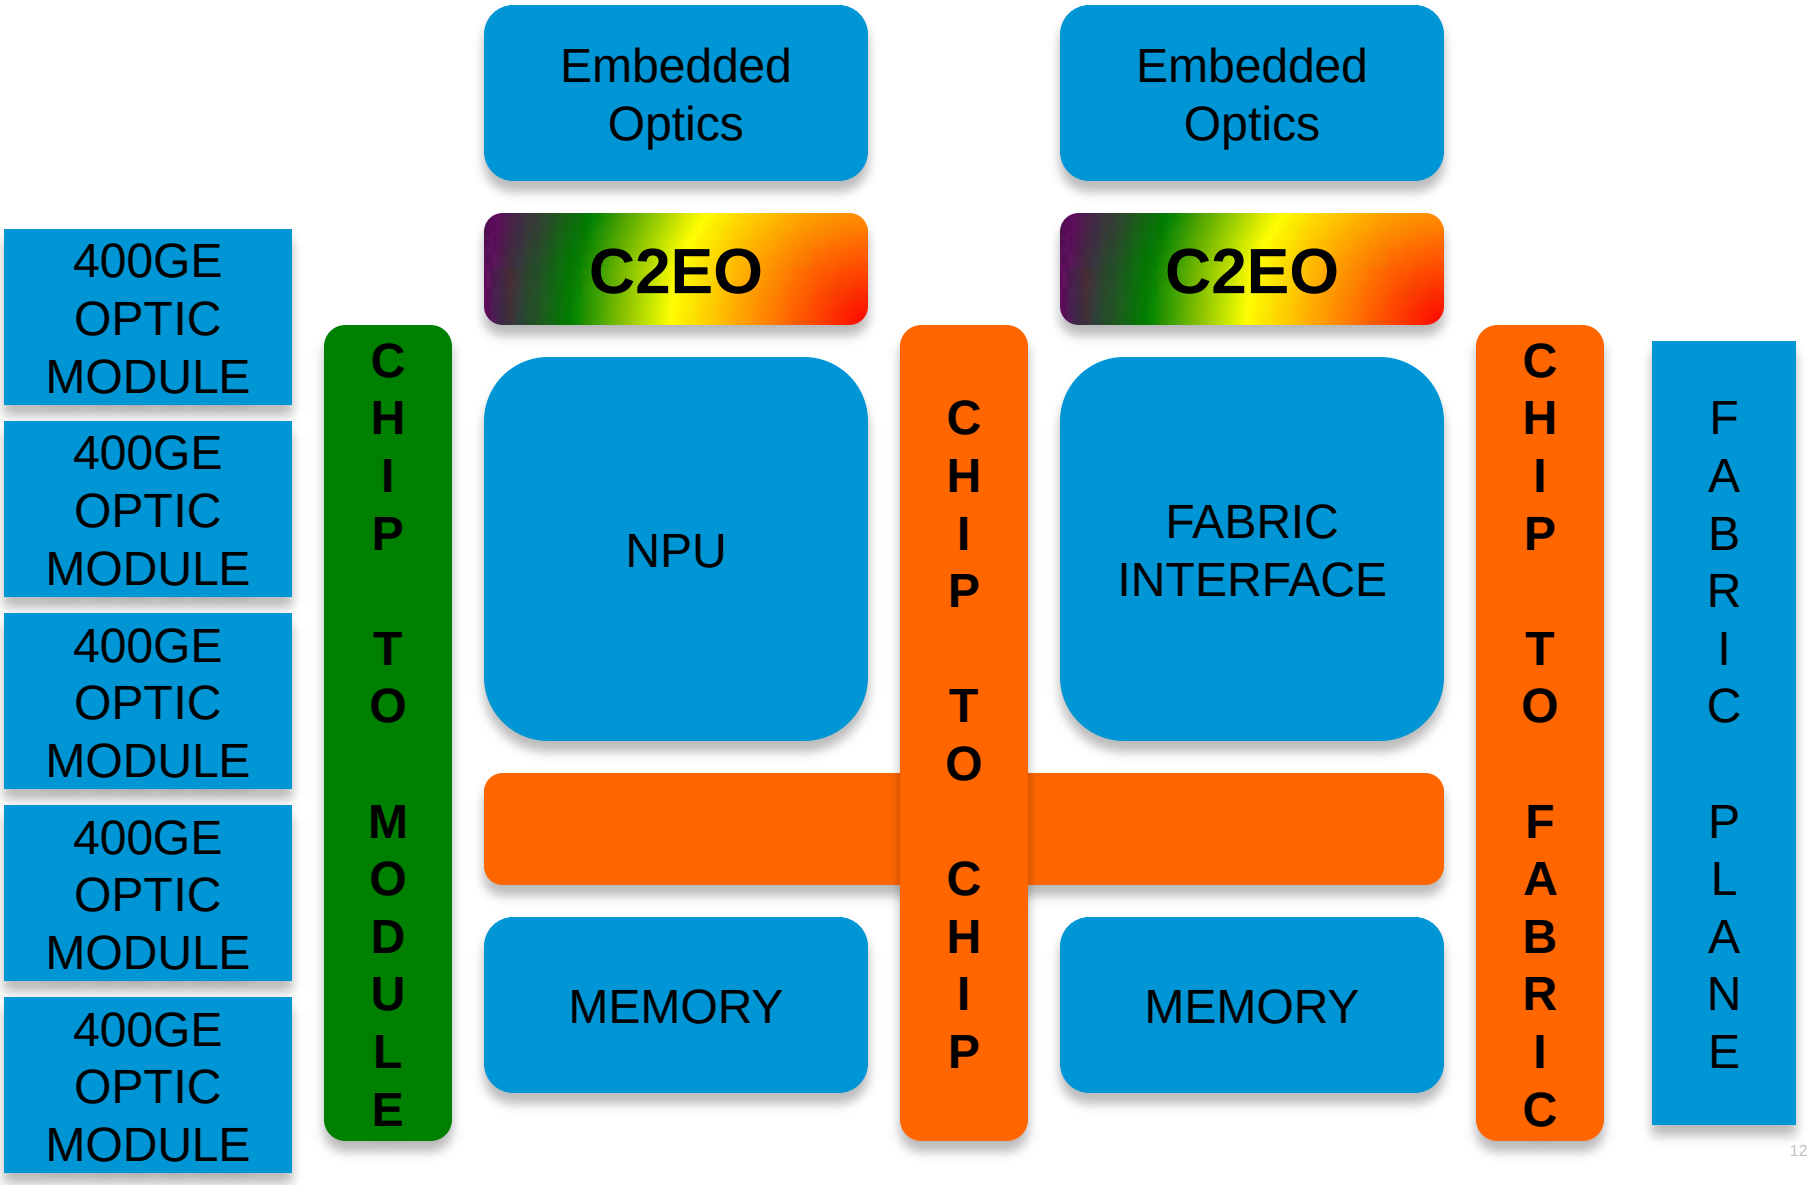

Take Away Points

- C2M: Chip2Module - Well defined interface. Tremendous effort in mechanical / SI analysis for the interconnect and packaging design. Usually about 5cm to 7cm – varies by design.
- C2C: Chip2Chip – Often custom protocols, but use the IEEE definitions. Can run up to 50cm in length. Usually about 5cm to 30cm – varies by design.
- C2F: Chip2Fabric – The fabric interconnect replaced what was once called the back plane or mid plane. This interconnect can be very long and consume a lot of power/bit. Probably outside the objectives here.
- C2EO: Chip2EmbeddedOptics – Emerging interface. Can easily replace all of the other interfaces in the system.

Exploring the Possibilities for Electrical Interfaces: Optimizing Lane Width to Module Rate

#Lanes / Port Rate	40x10G	32x12.5G	25x16G	20x20G	16x25G	10x40G	8x50G	4x100G	2x200G
10GE	40ports	32ports	25ports	40ports 20ports	32ports 16ports	40ports 10ports	X	X	X
25GE **Not IEEE MAC Rate**	X	16ports	X	X	16ports	X	16ports	16ports	X
40GE	10ports	X	X	10ports	8ports	10ports	8ports	8ports	10ports
50GE **Not IEEE MAC Rate**	8ports	8ports	X	X	8ports	X	8ports	8ports	8ports
100GE	4ports	4ports	X	4ports	4ports	X	4ports	4ports	4ports
200GE **Not IEEE MAC Rate**	X	2ports	X	2ports	2ports	2ports	2ports	2ports	2ports
400GE	X	X	X	1port	1port	1port	1port	1port	1port

Keeping the table simple
FEC is NOT shown here

Interface 16x25G / 20x20G
Best: C2M / Not so Best: C2C and C2F

- Lot of lanes to route / Lot of board layers / lot of pins.
- Covers every interface type.
- Available today / lots of solutions and layout guidelines.
- Good use for C2M Interface, but still lots of pins and SI issues.
- For Embedded Optics – C2EO, the short reach interface could be wide. The key is to drive as low pj/b as possible.

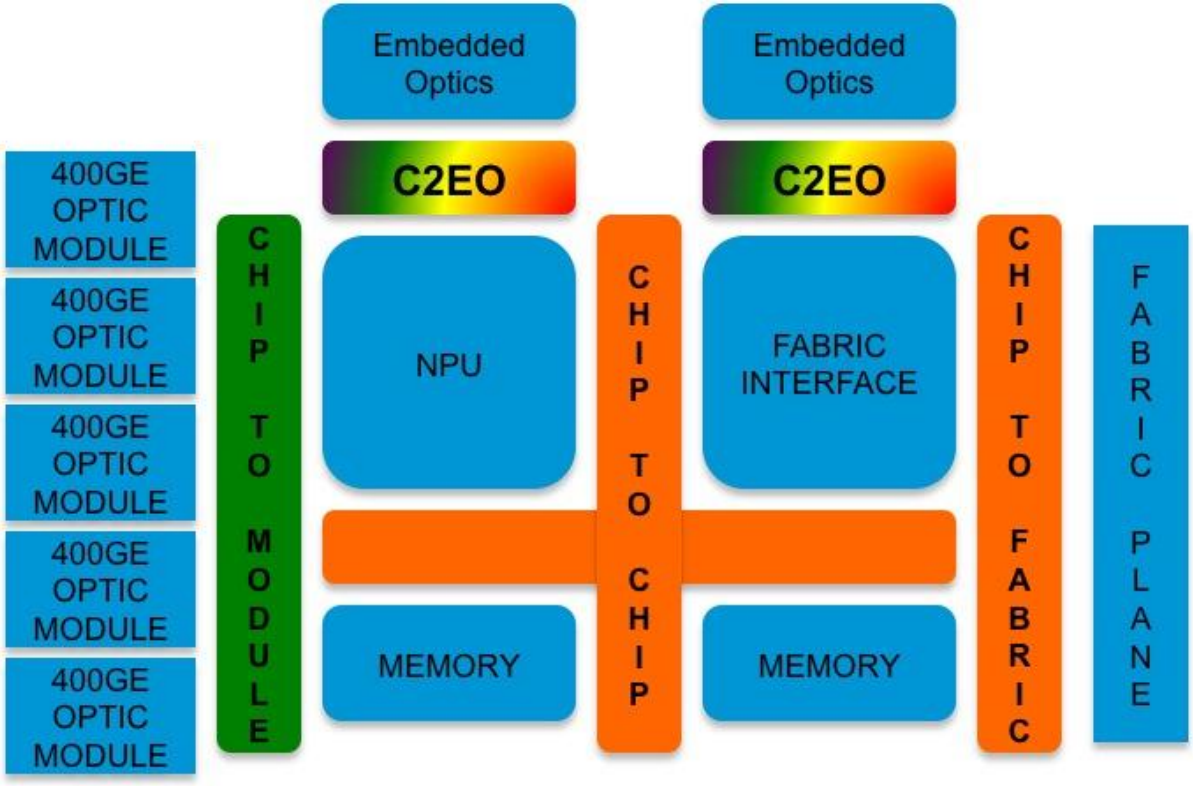

Interface 8x50G
Best: C2M / Good: C2C and C2F

- Good trade-off for routing, layers, and pin count.
- Not optimal for a range of port types.
- Available tomorrow/ lots of solutions in process. Could be a fast follower to 16x25G.
- Works good for C2M / C2C / C2F. There are SI issues.
- For Embedded Optics – C2EO, the short reach interface could be wide. The key is to drive as low pj/b as possible.

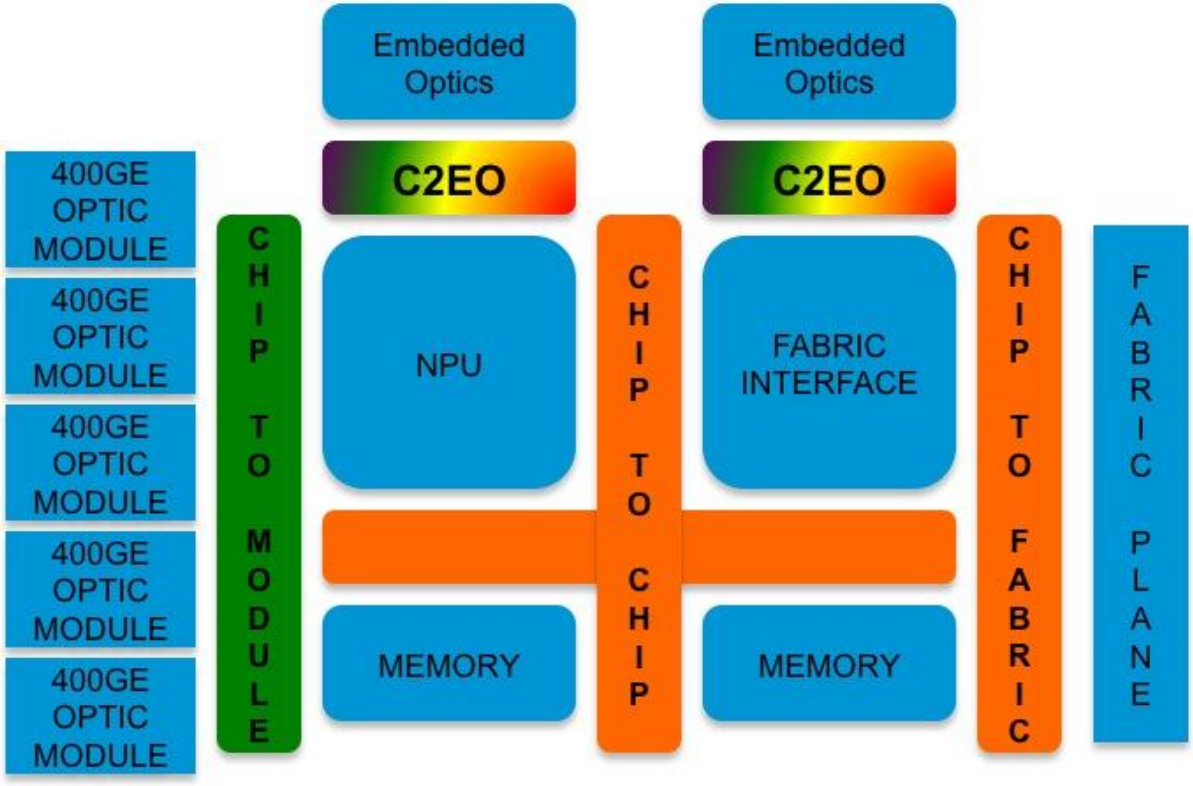

Interface 4x100G
Best: C2M / C2C / C2F / C2EO

- Good use of routing, layers, and pin resources.
- OK for a range of port types.
- Available tomorrow / lots of solutions in process (kidding – but close). This could be a reasonable follower to 16by25G.
- Works good for C2M / C2C / C2F / C2EO.

- For Embedded Optics – C2EO, this narrow short reach interface could work well if pj/b is low.
- Cost optimized by utilizing on chip die area to address channel impairments, as opposed to an embedded circuit that would be used in an optical module.

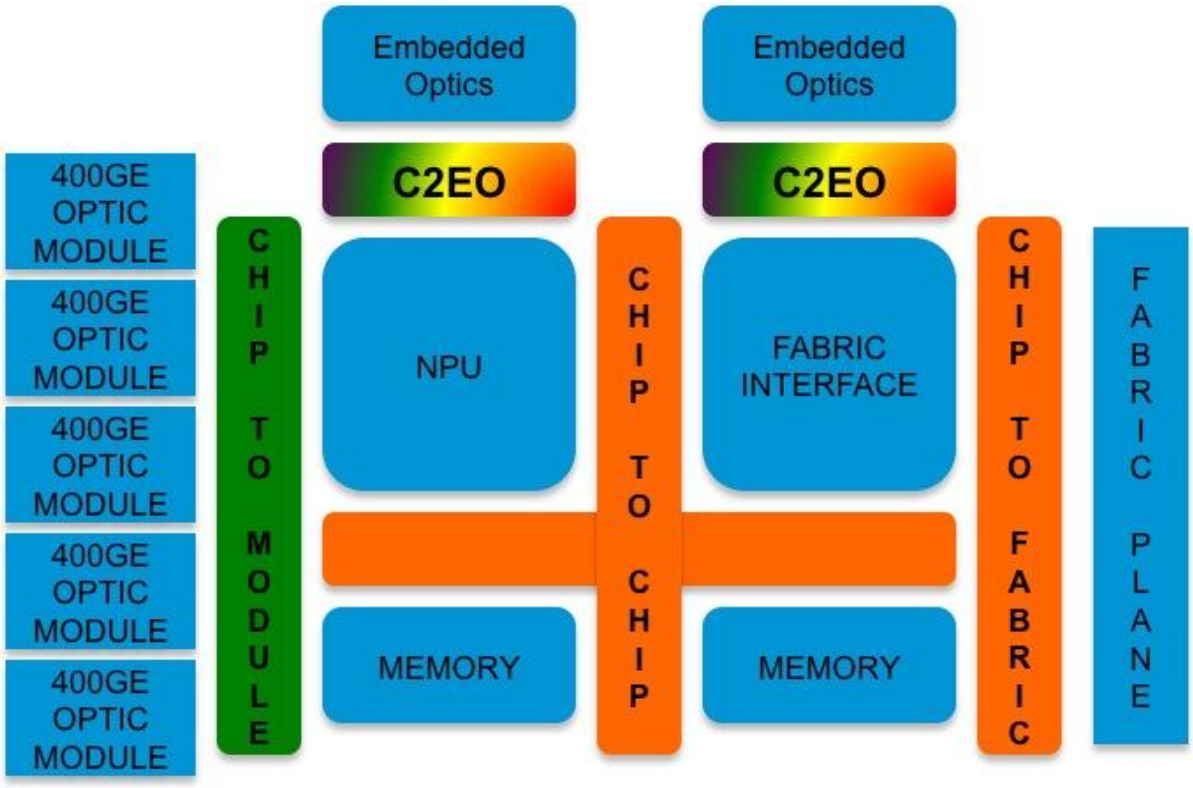

Interface 2x200G
Best: C2M / C2C / C2F

- Great use of routing, layers, and pin resources.
- Good for a range of port types and scales to 1Tbps.
- Available tomorrow / lots of solutions in process (kidding, but has potential).
- Works good for C2M / C2C / C2F / C2EO.
- For Embedded Optics – C2EO, this narrow short reach interface could work well if pj/b is low.
- Cost optimized by utilizing on chip die area to address channel impairments, as opposed to an embedded circuit that would be used in an optical module.

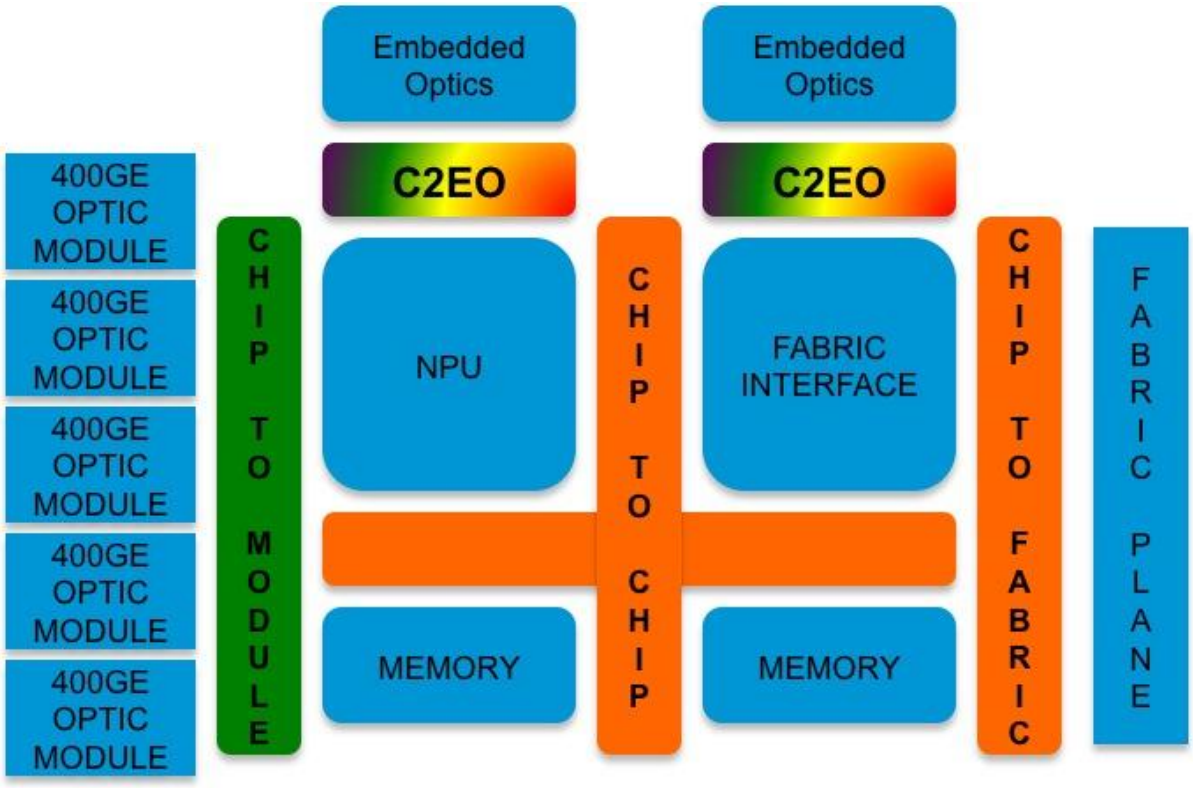

Take Away Points

- All four interface types need to be discussed.
- Width vs BW has to be addressed for each type.
- C2M / C2C / C2EO are key interfaces for next generation systems. C2F is most likely outside scope.
- Given all the uses for an interface, one width can’t address everything.
- 16x25G seems a good fit for early adoption based on past 20G and 25G adoptions.
- 8x50G might be a fast follower, but might be just an intermediate step.
- 4x100G could be a long term follower and may be preferred over 8x50G.

Interconnects and Power are Tightly Coupled

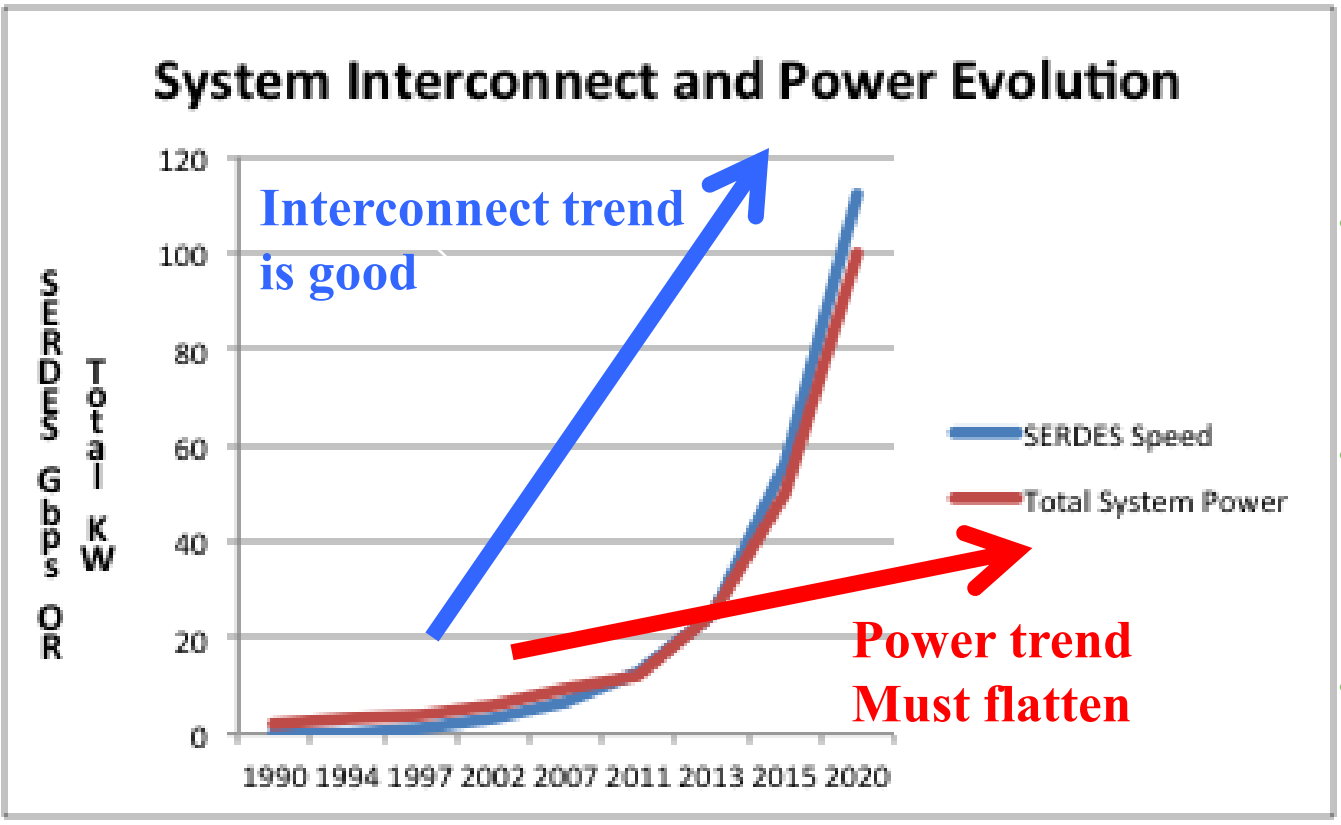

- From 20Mbps interconnects in a 2KW chassis to 25Gbps interconnects in a 25KW chassis today.
- SERDES trend is good, while the total power trend is not so good – needs to be more green friendly.
- Suggests the total interconnect capacity in Gbps/W is in a positive direction
- Implies the density story for interconnects is getting interesting at the front end and the back end.

Wishing for a Power Efficient Interface C2C / C2F / C2EO What There Is Today

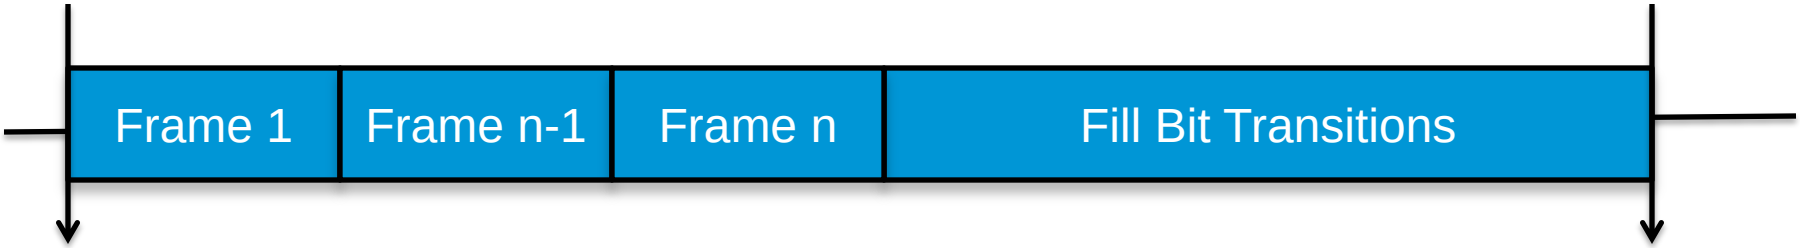

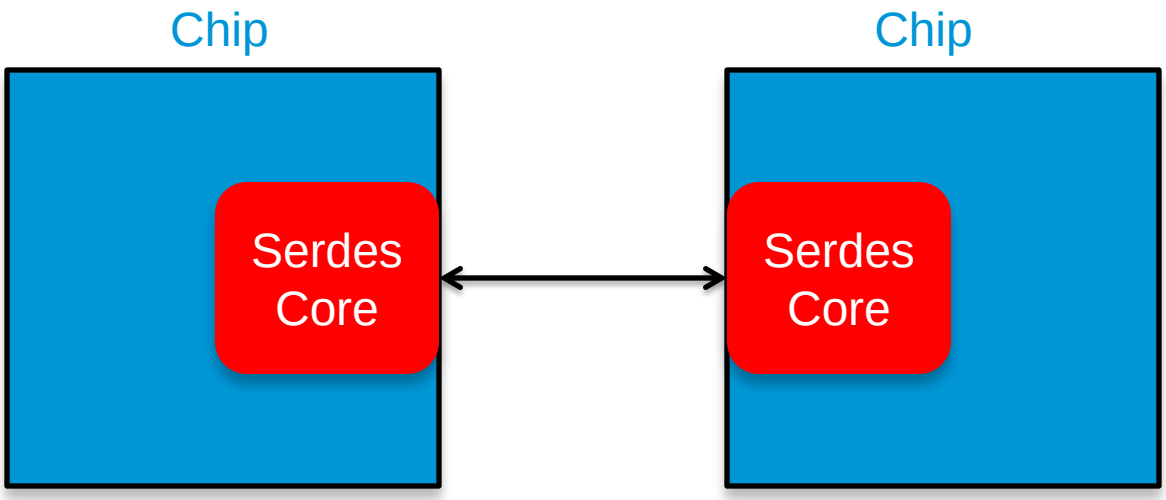

- EEE is great for C2M applications.
- EEE not a great fit for C2C / C2F / C2EO type interfaces.
- Transitions are constant to keep link error and latency minimal.
- Lots of wasted power. With average traffic at 32% or lower, there are a lot of unneeded bits transmitted and received. True even for packed links.

True – these C2C interfaces are often custom protocols

Defining a Power Efficient Interface
C2C / C2F / C2EO
Just Say Green

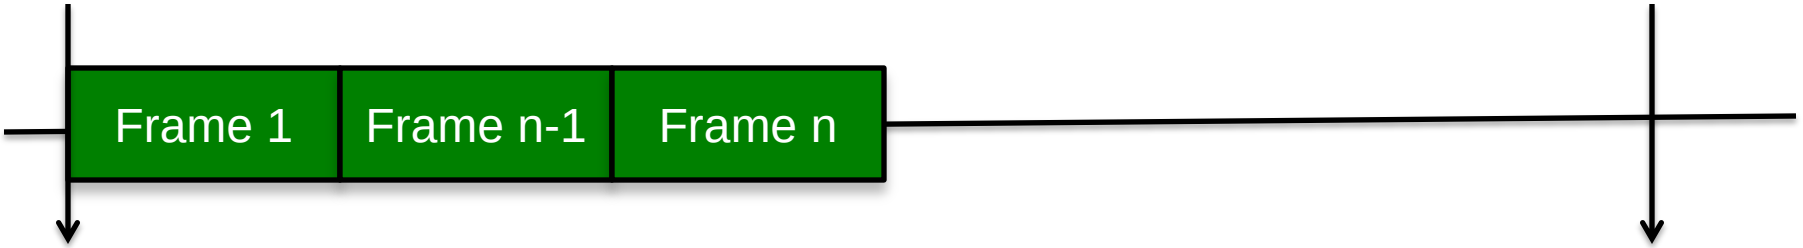

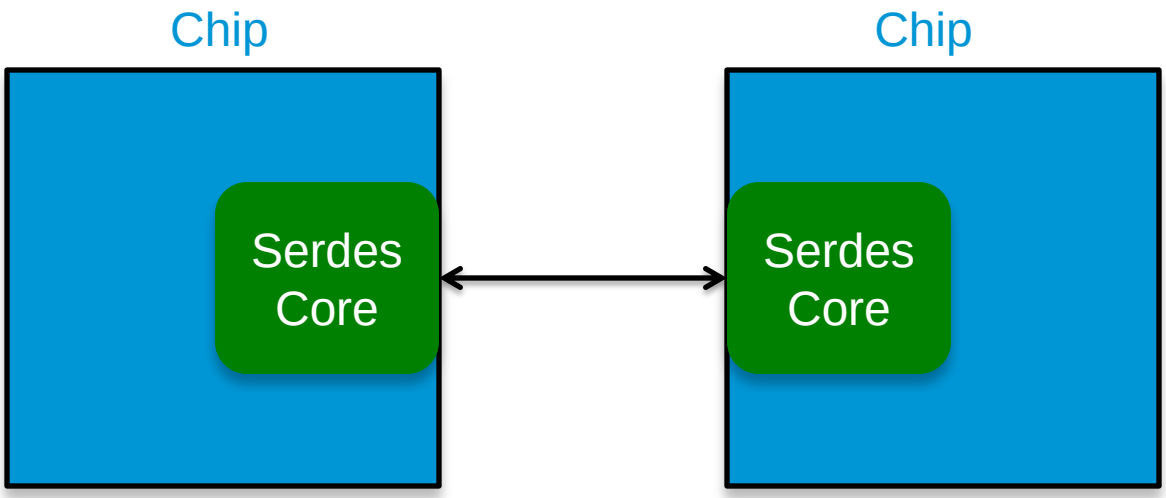

- No bits to send - the link is quiet.
- No bits to send – the serdes core is as asleep as possible.
- “No time needed for recovery” - must cover several jumbo frames.

True – these C2C interfaces are often custom protocols

Interconnects Range in the 1000s Chassis Size 10RU to 44RU Adds Up to A lot fo Power

Single Chassis System

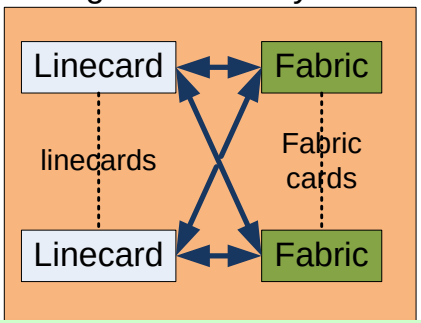

1) Cores from 250mW to 750mW

Multi-Chassis System

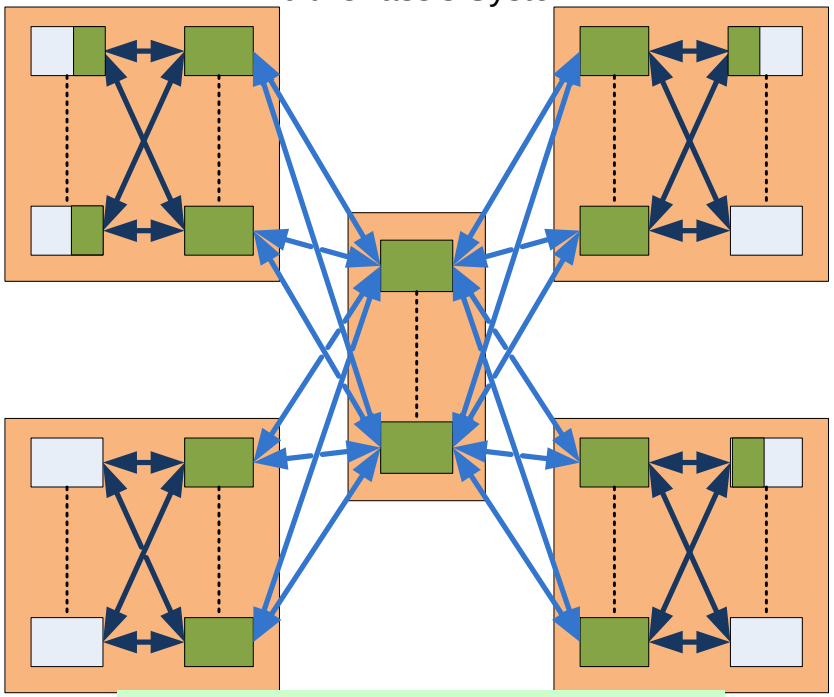

2) That is a Lot of Power

Back-to-back System

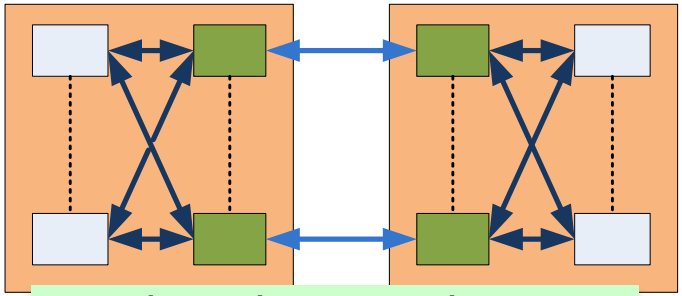

3) Reduce, Reduce, Reduce

Take Away Points

- C2C / C2F / C2EO do not need to be active all the time when average DC traffic, non peak, is 16% to 30%, and average SP traffic, non peak, is 32%. More then 2/3rd the time, the internal system links could be shut off. Up to 2KW savings in a 20KW chassis.
- Define a zero latency and zero recovery time interconnect
- Crucial to reducing power
- Crucial to achieving:

  20% power at zero traffic

  100% power at 100% traffic

Interface Types
Basic Requirements

	BW in GHz	Channel Loss	Coding	FEC	Potential Distance
C2M 28G VSR 56G VSR CAUI-4	??14GHz (today)	10dB	??	??	<15cm Will this be enough given a module 4 wide QSFP??
C2C 28G SR 56G MR CAUI-4	??14GHz (today)	20dB	??	??	<30cm
C2EO	25GHz to 50GHz	5dB	??	??	<1cm
C2F 25G LR .3bj (likely outside scope)	??12.5GHz (today)	35dB	??	??	<75cm

Take Away Points

- Have to use best channel design practices.
- 16x25G seems a good fit for early adoption based on past 20G and 25G adoptions.
- 4x100G could be a long term follower and may be preferred over 8x50G.
- Any interface should address more then C2M. Must include C2C and C2EO.
- Open questions on point to point loss/distance and on channel bandwidth for a given coding/signaling.
- Interface types must address maximum energy efficiency.
- Stay tuned for updates ….