100Gb/s Backplane/PCB Ethernet
Two Channel Model and Two PHY Proposal

IEEE P802.3bj **100Gb/s Backplane and**

Copper Cable Task Force

Newport Beach

January 2012

Rich Mellitz, Intel Corporation

Kent Lusted, Intel Corporation

Supporters and Contributors

- Matt Brown, AppliedMicro
- Dariush Dabiri, AppliedMicro
- Beth Kochuparambil – Cisco
- Joel Goergen, Cisco
- Brad Booth, Dell
- Howard Frazier, Broadcom
- Vasudevan Parthasarathy, Broadcom
- William Bliss , Broadcom
- Sudeep Bhoja, Broadcom
- Hamid Rategh, Inphi
- Sanjay Kasturia, Inphi
- Adee Ran, Intel
- David Chalupsky, Intel
- Ilango Ganga, Intel
- Bob Grow, Intel

100G Copper Interconnects Span a Complex Roadmap

- Hi-End system interconnect demand performance today
  - Examples: Edge/Core routers, switches, HPC compute nodes
- Mid level and low level systems demand cost effective solutions
  - Examples: ATCA, X86 blade servers, switches, single board computers, highly integrated silicon products
  - Chip and board level ingredients are trending towards desktop/laptop collateral.
  - ATCA & X86 blade servers have seen year over year growth and presently represents a significant market presence
    - Compare to when KR was originally proposed in 2003
- Two PHYs provide
  - better coverage over all types of medium while still allowing optimization for a particular market
  - an optimized evolution of cost reduction in board material vs. silicon cost

Backplane landscape at a glance illustrates a variety of loss boundaries

System Class	Line Card	Backplane	Loss	Freq.
Hi End	Meg6_LowSR–Wide	Meg6_LowSR–Wide	< 35dB	12.9GHz
Mid level	ImpFR4_HighSR–Narrow & FR4-MC	ImpFR4_HighSR–Narrow	< 33dB	7GHz
Low Level	FR4-MC	ImpFR4_HighSR–Narrow	< 33dB	7GHz

In the near future, low end system performance will exceed today's high end server performance

dB/in from Kochuparambil_01_0112.pdf

Attenuation* (dB/in) at:	1 GHz	6.5 GHz	7 GHz	12.89 GHz	14 GHz
Meg6_LowSR – Wide	0.0951	0.4159	0.4433	0.7562	0.8127
Meg6_LowSR – Narrow	0.1466	0.5849	0.6205	1.0152	1.0847
Meg6_HighSR – Wide	0.1175	0.5960	0.6367	1.0891	1.1688
Meg6_HighSR – Narrow	0.1856	0.8971	0.9557	1.5924	1.7020
ImpFR4_LowSR – Wide	0.1202	0.6096	0.6541	1.1772	1.2734
ImpFR4_LowSR – Narrow	0.1717	0.7794	0.8323	1.4410	1.5512
ImpFR4_HighSR – Wide	0.1427	0.7904	0.8484	1.5158	1.6367
ImpFR4_HighSR – Narrow	0.2106	1.0930	1.1692	2.0283	2.1813

PROPOSED PARAMETERS; GRAPHS ON PREVIOUS SLIDE

Traditional FR4 added as moderate controlled or managed loss controlled

FR4-MC	0.2673	1.3844	1.5369	2.6779	2.8798

Low loss channels required for high end systems can use NRZ

- Backplane channels can be constructed with < 35dB loss at ~12.9 GHz.
  - Channels using newer connectors, low-roughness copper foil, and “low-loss” PCB materials
    - i.e. (Meg6, LowSR-HC)

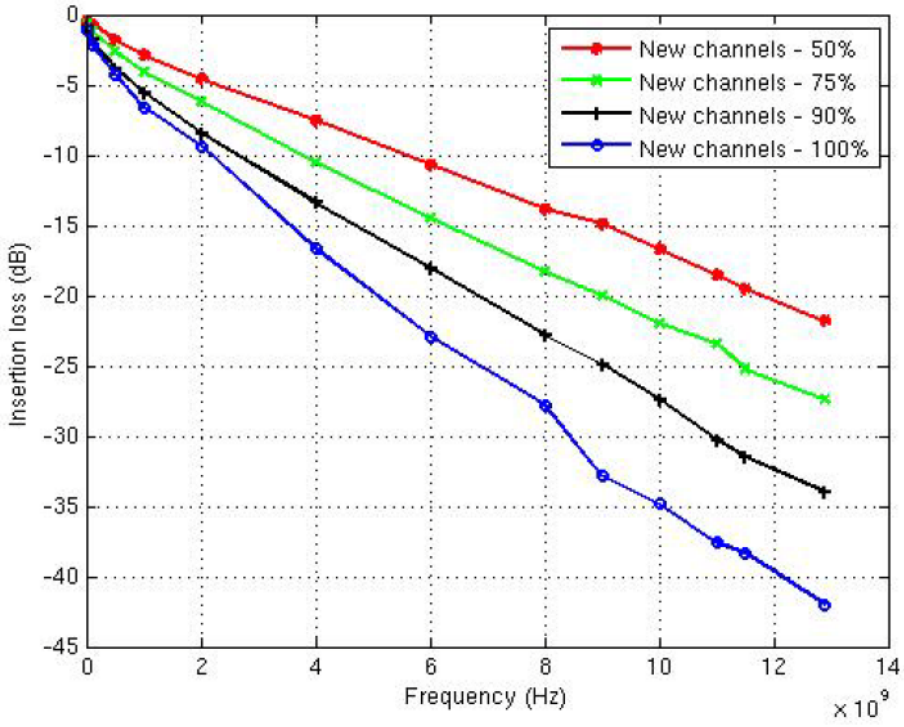

[1] Vasu Parthasarathy and Howard Frazier, *Channel Analysis at ~12.5GHz*, IEEE 802.3 100 Gb/s Backplane and Copper Cable Task Force, January 2012

Low/Mid Level Systems:
Can't Reach The Hi End Channel Bar, but Can Deliver 25Gbps Performance With PAM4

- Backplane channels can be constructed with <33dB loss at ~7 GHz.
  - Designed by out-sourced ODM resources
  - Higher channel variation and less sophisticated manufacturing
  - Can be managed with IPC-TM-650, Method 2.5.5.12 2011
    - IPC PCB loss test method
- The question **is not**
  - Is PAM4 better than NRZ?
- The question ***is***
  - What designs facilitate the low/mid systems as well the hi end system?
- PAM4 also facilitates backplane reuse with other interfaces designed for 10Gb/s signaling per lane:
  - 10GBASE-KR
  - 40GBASE-KR4
  - Future low/mid level platforms

The suggested insertion loss limit for low/mid level channels enables today's volume and future low end products

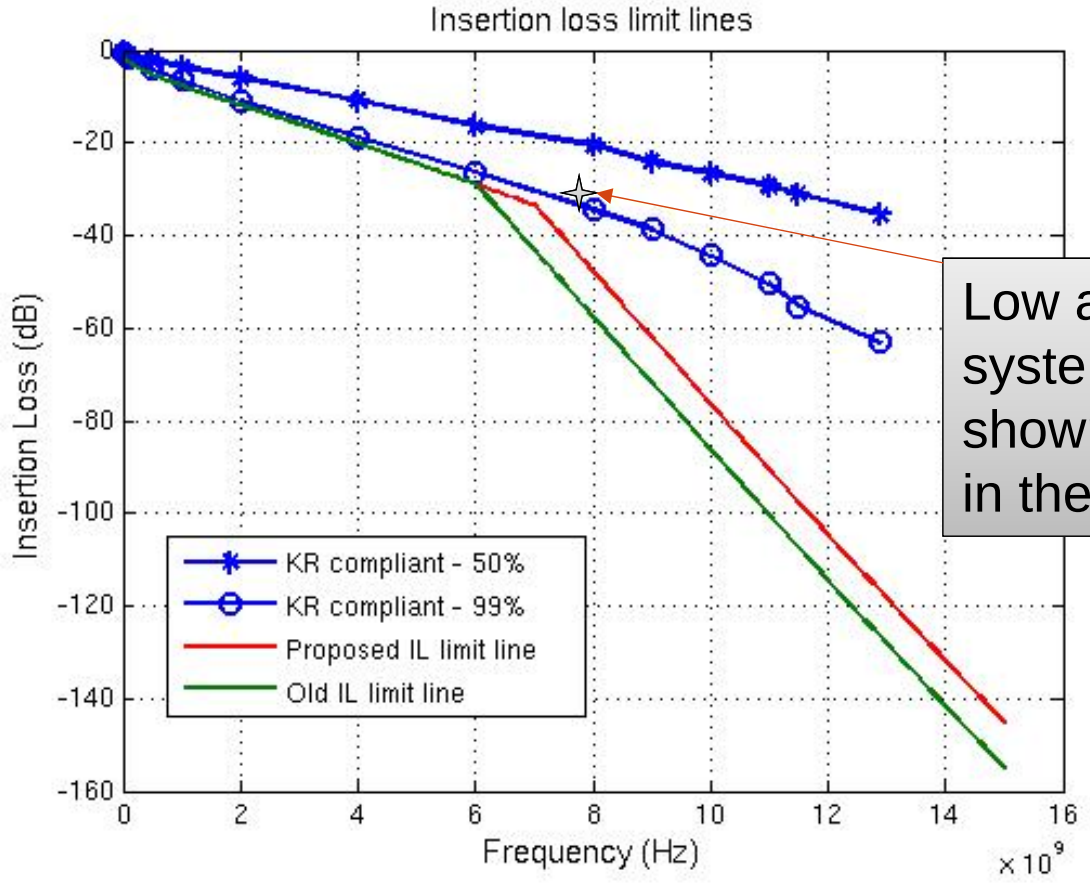

Low and mid level system IL budget shown later are with in these limits

[1] Rich Mellitz and Vasu Parthasarathy, *Rough channel targets for 4 x 25 Gb/s operation on existing backplanes*, IEEE 802.3 100 Gb/s Backplane and Copper Cable Study Group, May 2011

Server Platform is Optimized for High Volume

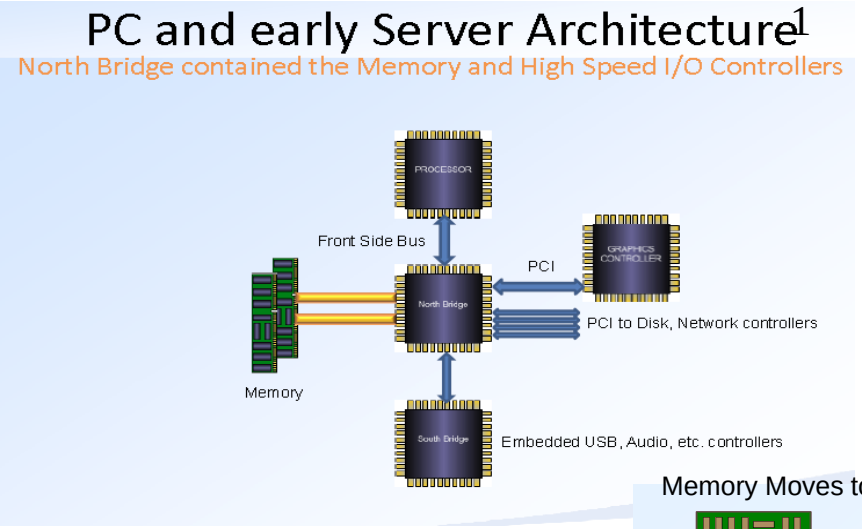

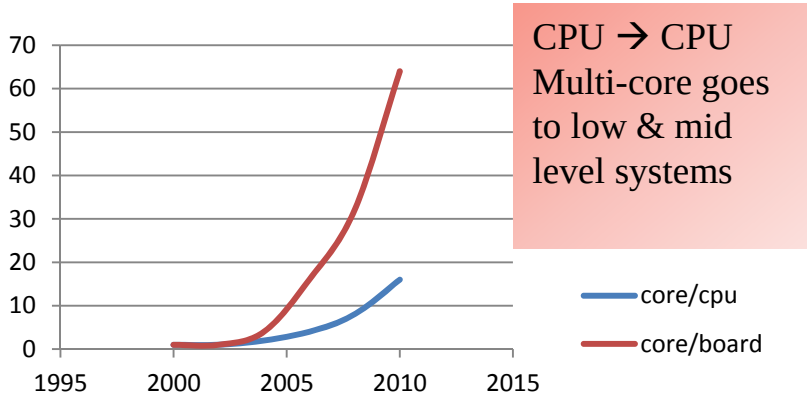

CPU → CPU Multi-core goes to low & mid level systems

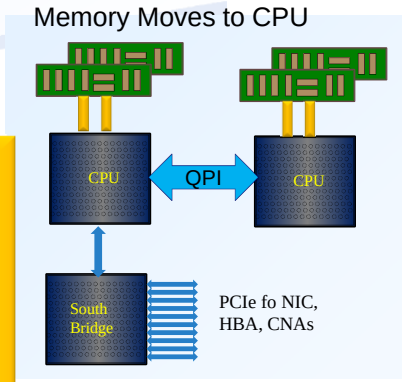

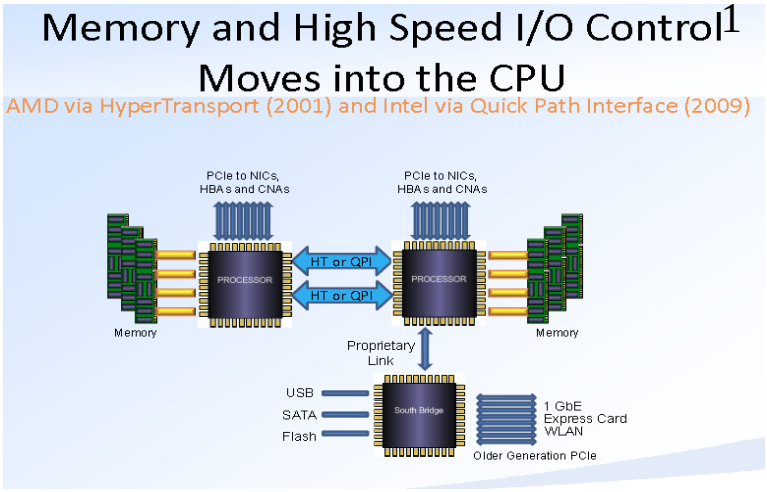

There is strong desire for the friendly Ethernet integration on low and mid level platforms

[1]Server Bandwidth Scenarios
Signposts for 40G/100G Server Connections
Presented by Kimball Brown
kimball@lightcounting

Low and Mid Level PCBs are very cost sensitive

- PCB technology is still standard FR4-class materials
  - Loss is starting to be managed
- Typical server motherboards are large
  - 130-150 sq. inches, 12-16 layer, 62 to 125 mils thick
  - Will not use 802.3ap spec'd "improved" FR4 materials
- Most volume server designs are outsourced
- Other interfaces drive PCB requirements:
  - BTW: None are 100 ohm differential and most have +/- 15% impedance tolerance or worse
- Hyper market segmentation is creating a large number of system form factors

BGA vias are important to consider for low/mid level systems

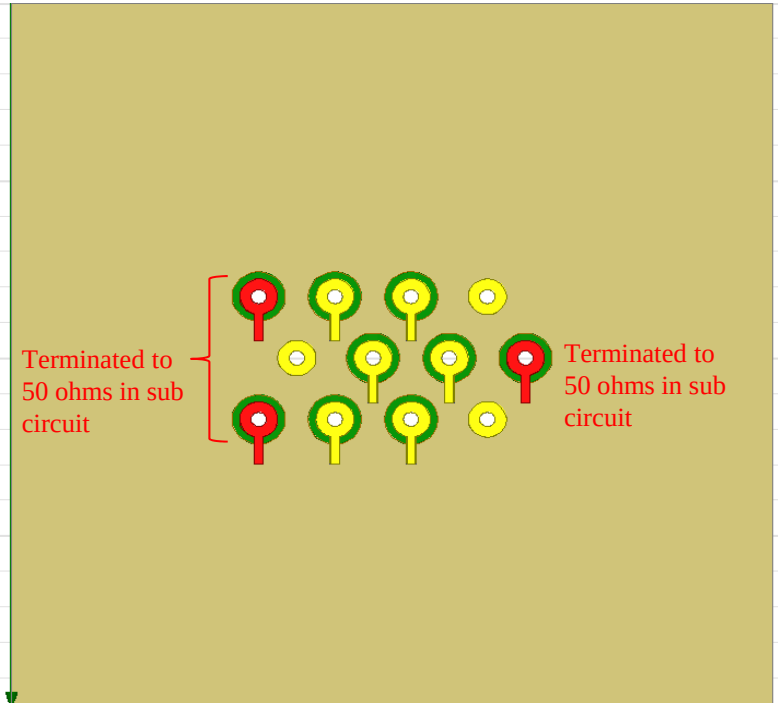

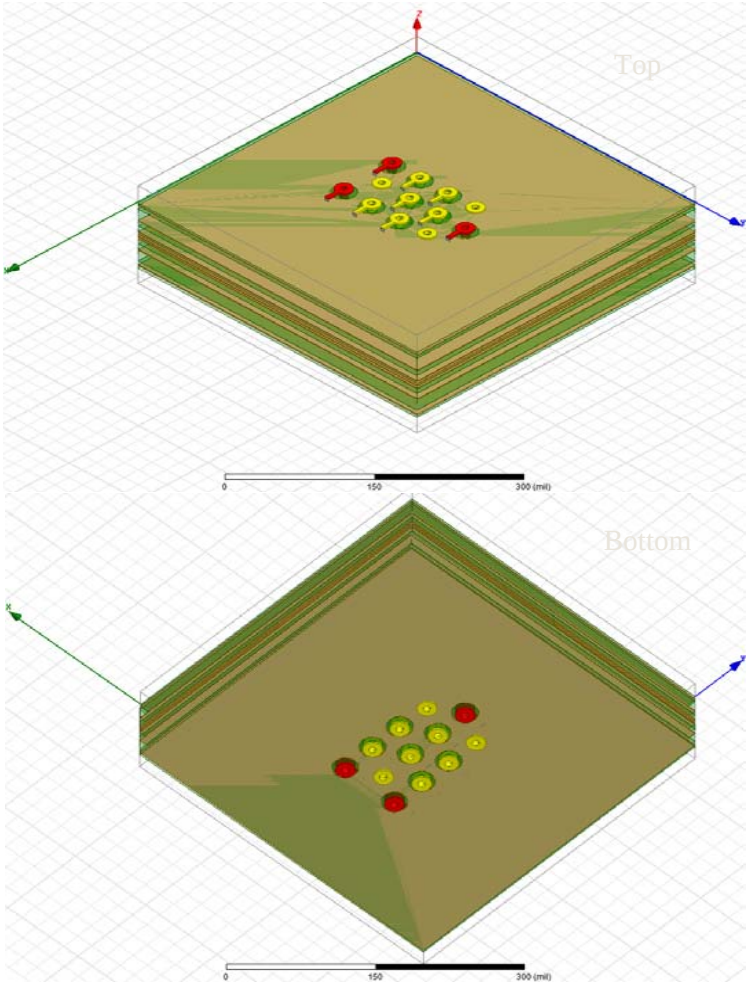

Drill diameter = 10 mils
Pad diameter = 20 mils
Antipad diameter = 28 mils
Pitch = 20 x 34.7 = 40mils
Coupling: 40 mils

Loss of typical thru via at BGA is high using only the bottom 2 stripline layers

BGA crosstalk is also a trade off with the via IL

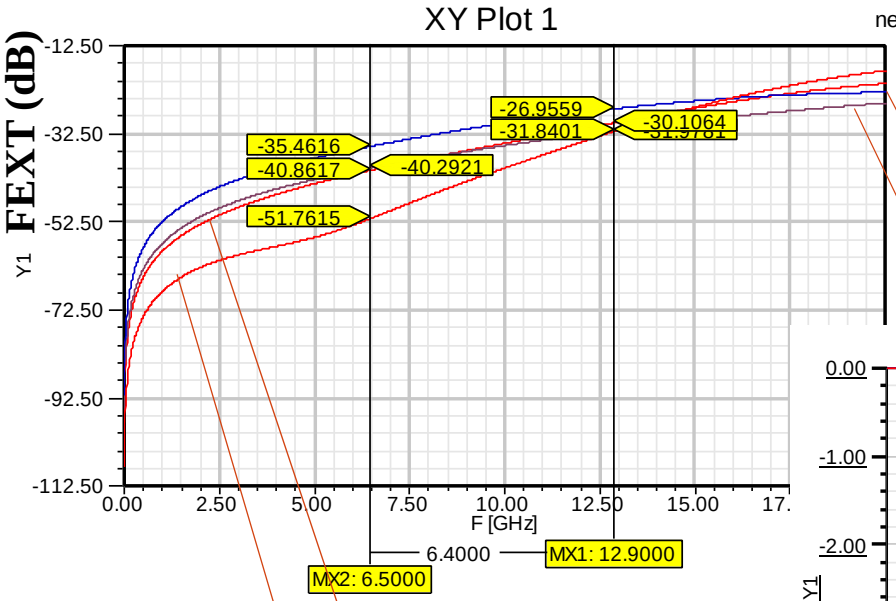

062 mil thick 8 layer board
Layer 6 exit

093 mil thick 14 layer board
Layer 12 exit

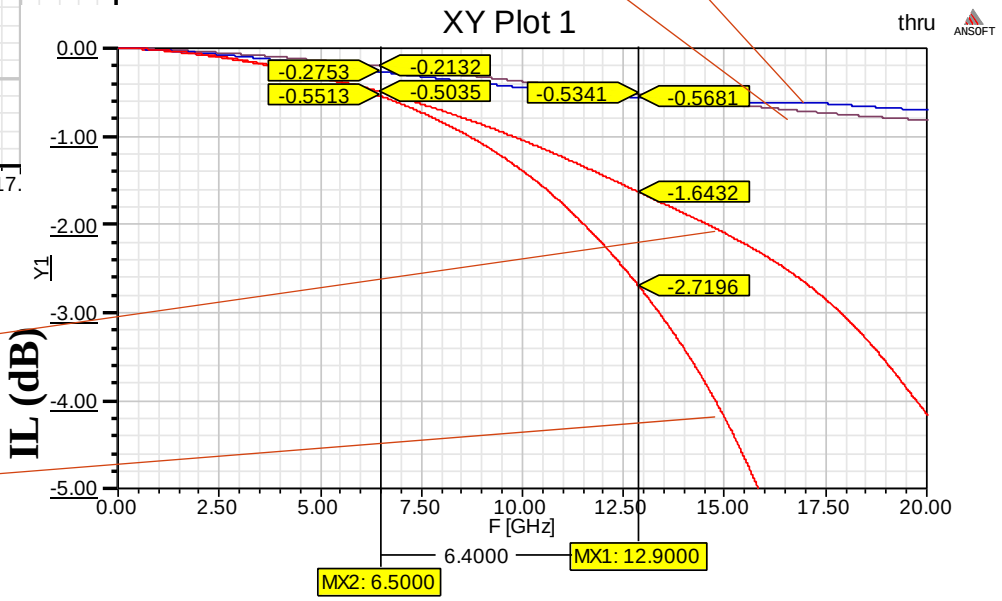

093 mil thick 14 layer board
Layer 10 exit

062 mil thick 8 layer board
Layer 4 exit

Back drill in BGA region is problematic for mid and low-end systems

Simplified typical loss budgets for low/mid level systems after 7 GHz exhibit a non log-linear loss

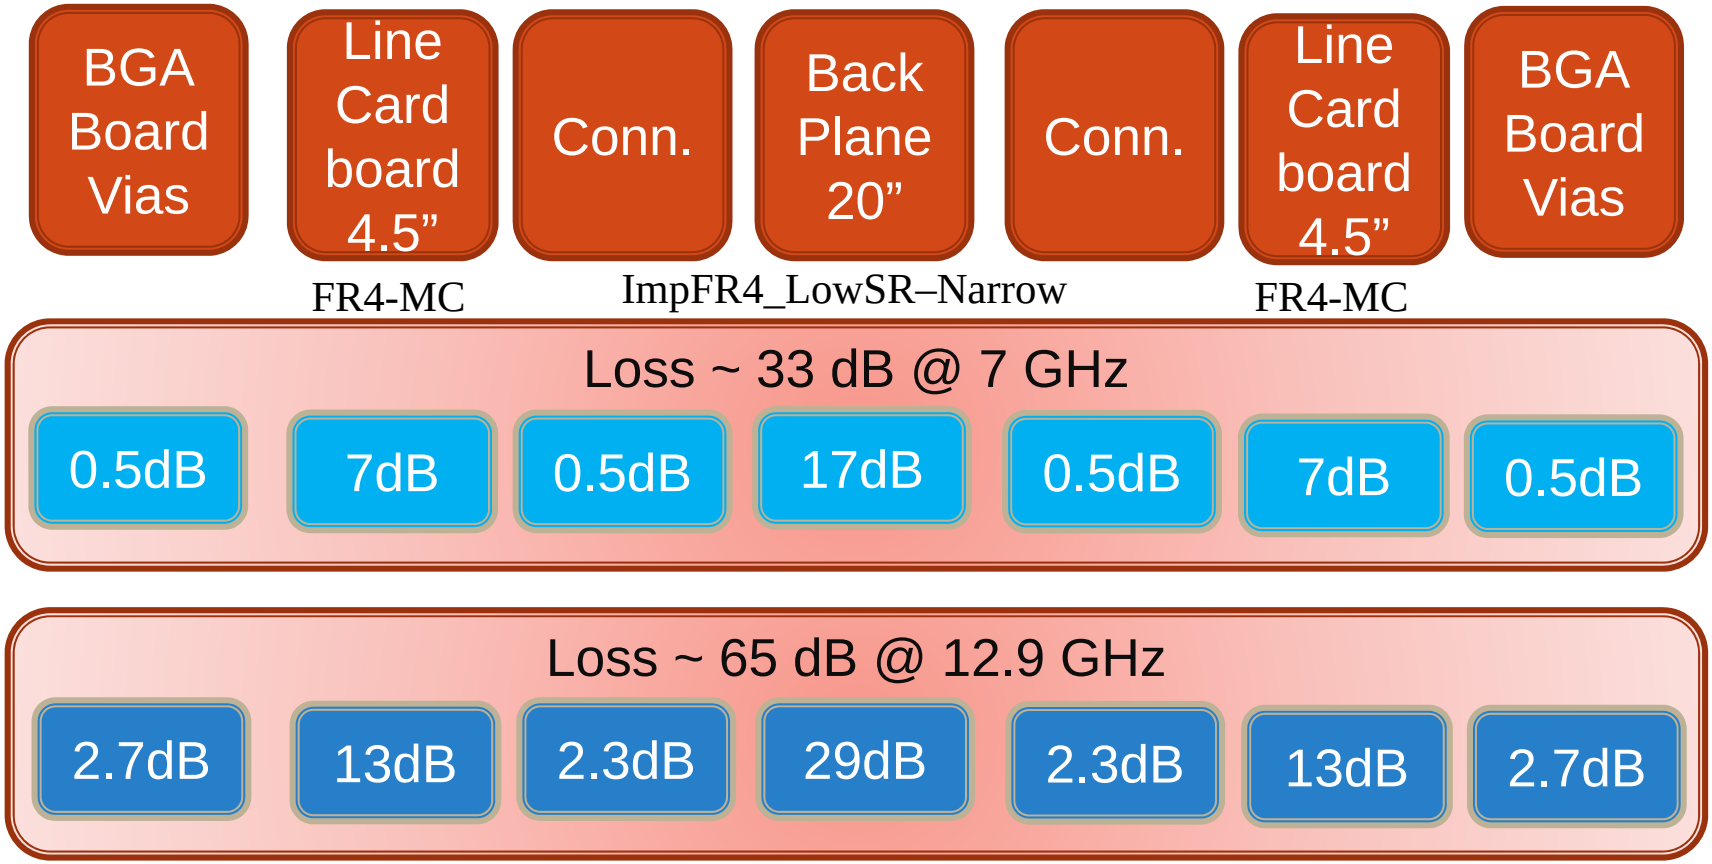

Losses computed using DkDf Algebraic Model Tool, http://www.ieee802.org/3/bj/public/tools.html

Summary

- 100G copper interconnects span a variety of loss boundaries
- Channel loss varies widely from high end to low end due to design/mfg. choices
- Mid and low end server platforms are high volume, enabled by large scale integration
- The low/mid level channels have a non log linear loss behavior after 7 GHz
- Two PHYs provide an optimized solution for:
  - Low loss channels in high end systems
  - Manufacturing and volume channel requirements in mid/low level systems

Thank You!