Backplane and copper cabling objectives – wording and technical decisions

IEEE 802.3 100Gbps Single-lane Electrical Study Group

George Zimmerman

CME Consulting, Inc./Aquantia

Supporters

- Jon Lewis, Dell+EMC
- Ramin Farjadrad, Aquantia

Objectives

- Meant to describe the problem being solved, not the solution
- Presentations show multiple paths to solutions
- Objective is to define the problem
- Our problems are connectivity in a system:

Connectivity in systems

Various constructions:

(heck_100GEL_01_0118.pdf)

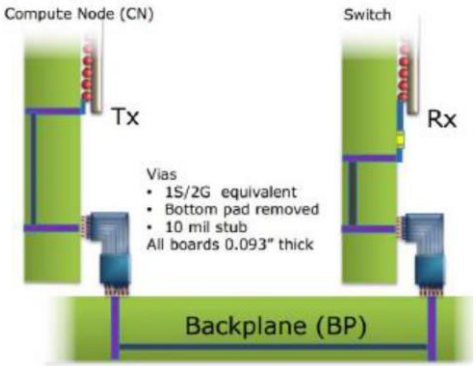

(stone_100GEL_01_0118.pdf)

Cabled backplane

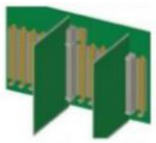

Traditional orthogonal backplane

Facebook / OCP

Orthogonal mid-plane

BROADCOM

802.3 100GEL Study Group, Geneva

BUT: 3 Primary Applications:

(ghiasi_100GEL_01a_0118.pdf)

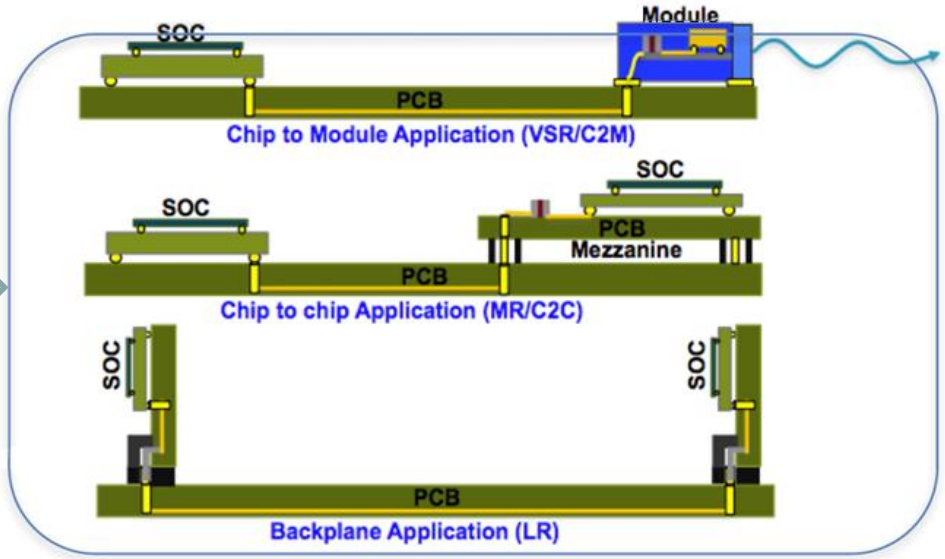

How to state the problem

- 802.3ba defined the problem physically:

- Provide Physical Layer specifications which support 40 Gb/s operation over:
  - at least 10km on SMF
  - at least 100m on OM3 MMF
  - at least 7m over a copper cable assembly
  - at least 1m over a backplane

- 802.3bj: Used loss and frequency
  - Implies the form of the solution
  - Multiple cases allowed the TF to get away with it

802.3bj – the breaking point

- Study Group defined the problem generally:

  (July 2011)

  - Define a 4-lane 100 Gb/s backplane PHY for operation over links consistent with copper traces on “improved FR-4” (as defined by IEEE P802.3ap or better materials to be defined by the Task Force) with lengths up to at least 1m.
  - Define a 4-lane 100 Gb/s PHY for operation over links consistent with copper twin-axial cables with lengths up to at least 5m.

- Task Force chose solutions and refined with frequencies and losses

  (March 2012)

  - Define a 4 lane PHY for operation over a printed circuit board backplane with a total channel insertion loss of <= 35 dB at 12.9 GHz**
  - Define a 4 lane PHY for operation over a printed circuit board backplane with a total channel insertion loss of <= 33 dB at 7.0 GHz**
  - Define a 4-lane 100 Gb/s PHY for operation over links consistent with copper twin-axial cables with lengths up to at least 5m.

802.3by and 802.3cd used existing SERDES and assumed the channel losses.

Now, many assume we are settled on PAM4 losses & frequencies

BUT – we have a harder problem here.

Can we really assume losses & frequencies?

- Ad hoc shows challenges exist for 100GEL
- Something has to change - variations have been shown to effect loss budgets and frequencies:
  - Packaging (holden, ghiasi)
  - Joint encoding (holden)
  - Loss budgets (lim, zhang, haser, ofelt)
  - Board materials (lim, slide 4)
  - Backplane constructions (zhang, heck, sakai)
  - Precoding/TX FFE (sun)
  - FEC and Decision error rate (sun)
  - Single vs. Dual-duplex (farjarad)
  - Receiver complexity vs. Material Complexity

Example – reported losses

- COM calculates Zero-Forcing DFE
  - Not implementation-independent, not optimal
  - May be influenced by crosstalk (MMSE-DFE & others are better in crosstalk)

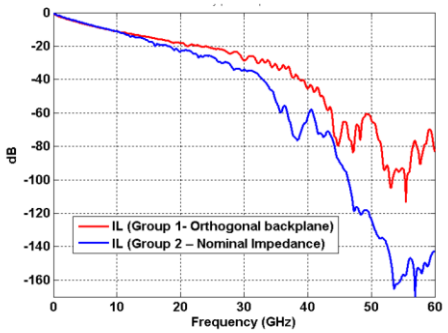

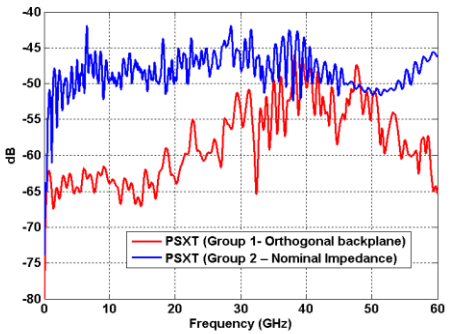

"Group 1" and "Group 2" from zhang_100GEL_adhoc_01_022618.pdf

- DFE theory gives us bounds
  - Salz SNR allows us to compare ultimate performance due to the channel (even if we ARE complexity limited)
- Shows nearly 10dB margin difference due to channel, BUT, choice of implementation can account for about as much!

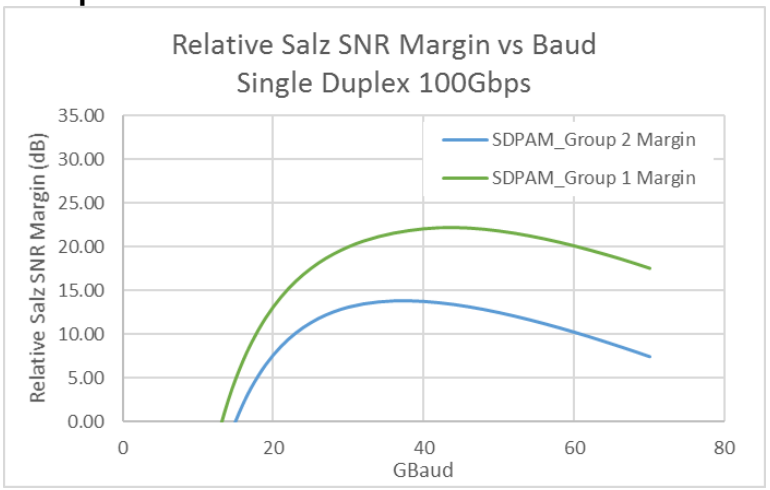

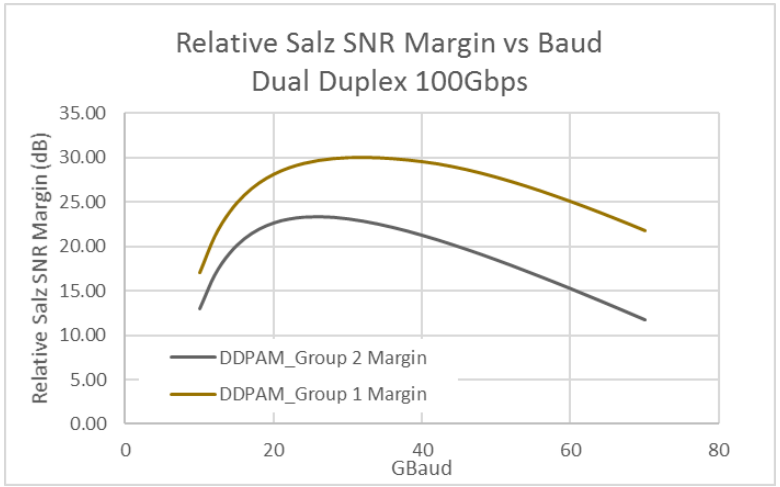

By focusing on frequency and loss, objective assumes PHY choices

50 GBd PAM-4 Single-duplex

- 2X signaling rate in PHY
  - 2X 50Gbps speed DFE ~ 4X 50G length
  - 2X 50Gbps speed FFE ~ 4X 50G length
- High performance channel
  - Connectors, Cabling, Package
  - New SI considerations
- Balance of complexity shifts toward board materials – impacts ALL circuits/cost balance
- Lower expectations from system
  - Shorter DAC (if at all)
  - Shorter trace lengths
  - Overall less margin
- Loss budgeted at 26 GHz

25 GBd PAM-4 Dual-duplex

- Well established equalizer from 50Gbps Ethernet
- Adds Echo canceller
  - PAM-4 -> multiply-free structure
  - Similar to extra DFE taps except simpler - no need to close timing loop
- Robust, known system expectations
  - SI considerations similar to 50Gbps PAM4 (25 GBd)
- Balance of complexity shifts to PHY
  - Don’t need to upgrade board materials everywhere
  - Ability to use with high performance channel de-risks the approach
- Loss budgeted at 13.28125 GHz

The proposal in a nutshell

- Don’t assume the solution in the initial objectives
- Let the Task Force do its work
- Then... Refine the objectives if needed

The detail:
Chip to Chip and Copper Cabling Objectives

- For chip-to-chip objective & twinaxial cabling, form agreed in January is OK, makes no assumptions on solution
- For twin-ax copper cables objective, fill in TBD with 3m:
  - Define a single-lane 100 Gb/s PHY for operation over twin-axial copper cables with lengths up to at least 3 m.
    - In line with prior projects (802.3cd objectives) PHYSICAL needs – racks aren't getting smaller!
  - Presentations have shown feasibility, either through signal processing (e.g., farjadrad_100GEL_01c_0118.pdf and follow-ups); OR by advanced materials and budgeting less loss on board/chip
  - Several presentations support at least 2m as feasible with various choices of techniques – none of which pulled out all the stops. (e.g., haser_100GEL_adhoc_01_022618.pdf showed several solutions > 2m)
  - Combining media, budgeting and signal processing will get us beyond 2m
- The form of the objective is unambiguous and general with multiple ways to do it, and leaves technical decisions for the Task Force

Proposal: Backplane Objective

- For backplane objective, replace 802.3bj form from January:
  - "Define a single-lane 100 Gb/s PHY for operation over electrical backplanes supporting an insertion loss ≤ TBD dB at TBD GHz."
  - With 802.3ba form:
    - "Define a single-lane 100 Gb/s PHY for operation over electrical backplanes supporting up to at least 1m? over a backplane." (802.3ba form)

- It's unambiguous and general – it describes the problems we are trying to solve...
- Multiple ways to do it (media or signal processing) and leaves technical decisions on the table for the Task Force
  - Combining solutions will only improve the performance and robustness

- And, it leaves technical decisions on the table

THANK YOU!