Towards a 200G/lane Backplane Objective – An Update

Kent Lusted, Intel
Howard Heck, Intel
Rich Mellitz, Samtec
Tom Palkert, Samtec
Jim Weaver, Arista
Jason Chan, Arista
Nathan Tracy, TE
Megha Shanbag, TE
Sam Kocsis, Amphenol
Priyank Shukla, Synopsys
Matt Brown, Huawei

March 2023

v01a

Introduction

- This contribution proposes a form for a set of backplane PHY objectives
- Intend to have a straw poll to measure support for the direction
  - No motion to adopt at this meeting

Adopted Physical Layer Objectives

This Presentation's Focus

Ethernet Rate	Assumed Signaling Rate	AUI	BP	Cu Cable	MMF 50m	MMF 100m	SMF 500m	SMF 2km	SMF 10km	SMF 40km
200 Gb/s	200 Gb/s	Over 1 lane		Over 1 pair			Over 1 Pair	Over 1 Pair		
400 Gb/s	200 Gb/s	Over 2 lanes		Over 2 pairs			Over 2 Pair			
800 Gb/s	100 Gb/s	Over 8 lanes	Over 8 lanes	Over 8 pairs	Over 8 pairs	Over 8 pairs	Over 8 pairs	Over 8 pairs		
	200 Gb/s	Over 4 lanes		Over 4 pairs			Over 4 pairs	1) Over 4 pairs 2) Over 4 λ's		
	TBD								Over single SMF in each direction	Over single SMF in each direction
1.6 Tb/s	100 Gb/s	Over 16 lanes								
	200 Gb/s	Over 8 lanes		Over 8 pairs			Over 8 pairs	Over 8 pairs		

Leverage existing or work-in-progress 100 Gb/s per lane (e.g. 3cu, 3ck, 3db) to higher lane counts

Develop 200 Gb/s per lane electrical signaling for 1/2/4/8 lane variants of AUIs and electrical PMDs

Develop 200 Gb/s per optical fiber for 1/2/4/8 fiber based optical PMDs and 4 lambda WDM optical PMD

Potential for either direct detect and / or coherent signaling technology

Making it all work together

https://www.ieee802.org/3/B400G/public/21_1028/B400G_overview_c_211028.pdf

Background

- Initial thoughts on 200G/lane Backplane Objectives were provided in a January 2023 contribution
  - https://www.ieee802.org/3/dj/public/23_01/23_0116/lusted_3dj_01a_230116.pdf
- Strong support per Straw Poll #1
- Feedback received included:
  - Confusion between objective and specification details
    - “Die-die” reference
    - Package aspects
  - Clarify the difference from high loss AUIs
  - Request for channels to study

Straw Poll #1

I am interested in backplane PHY objectives for 200Gbps/lane rates

Y: 58 , N: 19 , A: 34

https://www.ieee802.org/3/dj/public/23_01/motions_3dfdj_a_2301.pdf

Objective form

- We propose to reshape the backplane <u>objective</u> format to be inclusive of the package structures. In other words, “die-to-die” insertion loss
  - Test points and compliance methods would be a subject for baseline proposals

Add a one-lane 200 GbE, a two-lane 400 GbE, a four-lane 800 GbE, and an eight-lane 1.6 TbE backplane objective of the form:

- “Define a physical layer specification that supports [n*200] Gb/s operation over [n] lanes over electrical backplanes supporting a die-to-die insertion loss ≤ X dB at 53.125 GHz”

Agree upon “X” later

Insertion Loss Allocation

- Having a die-to-die insertion loss objective does not contradict the ability, need or intention to budget
- Die-to-die budget consists of the following elements:
  - Transmit Package
  - Transmit board/line card
  - Backplane
  - Receiver side board/line card
  - Receiver side package
- The budget of each element is a tradeoff amongst the other elements
  - This can be addressed in a baseline proposal

Slide courtesy of Liav Ben-Artsi

Specification form

- The test points (e.g. TP0 to TP5) could be defined like the transmitter and receiver characteristics and compliance methodologies in IEEE Std. 802.3ck-2022 Clause 163
  - This would be a subject for the baseline proposal candidate specification
  - The baseline might consider test points at the die

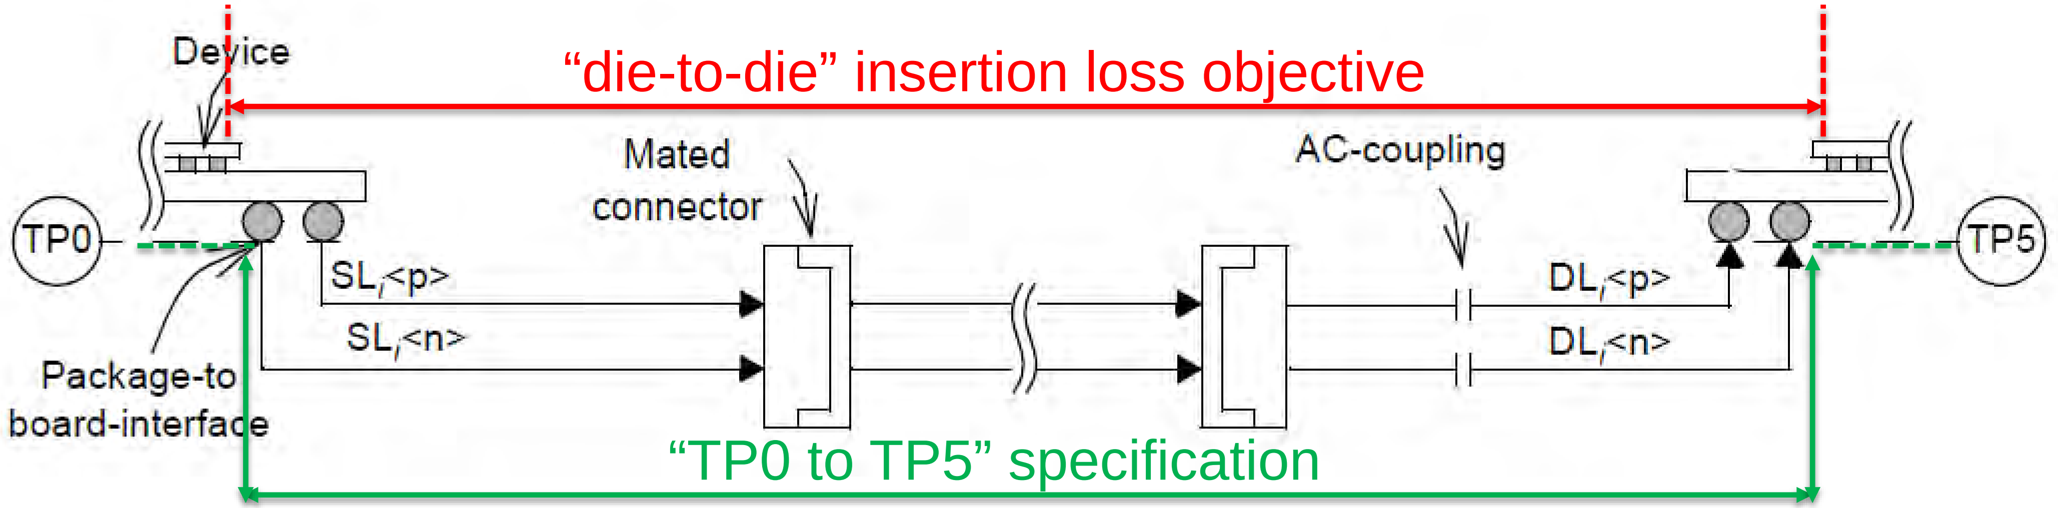

Some differences from a AUI C2C

- PMDs are architecturally different from an AUI, even if all electrical specifications are the same
- Needs Cl 73 Auto-negotiation
- Needs in-band PMD control function (“link training”)
  - Link performance optimization with a partner “outside the box” (i.e. not under the same management domain)
  - AUI C2C has a choice of in-band or out-of-band “link training”
- Potentially higher IL target
  - May make use of higher complexity reference receiver
  - Exact IL number subject to discussion later, potentially coupled to passive copper cable
- Likely different BER target
  - AUI BER targets are typically lower than PMD BER targets
  - AUI BER targets are not yet adopted

Much is dependent on what happens with AUIs

Stack Comparison

120A.5 Partitioning examples supporting 200GBASE-CR2/KR2 and 400GBASE-CR4/KR4

Figure 120A–8 depicts an example of 200GBASE-CR2/KR2 and 400GBASE-CR4/KR4 PMA layering with a single 200GAUI-*n* or 400GAUI-*n* interface.

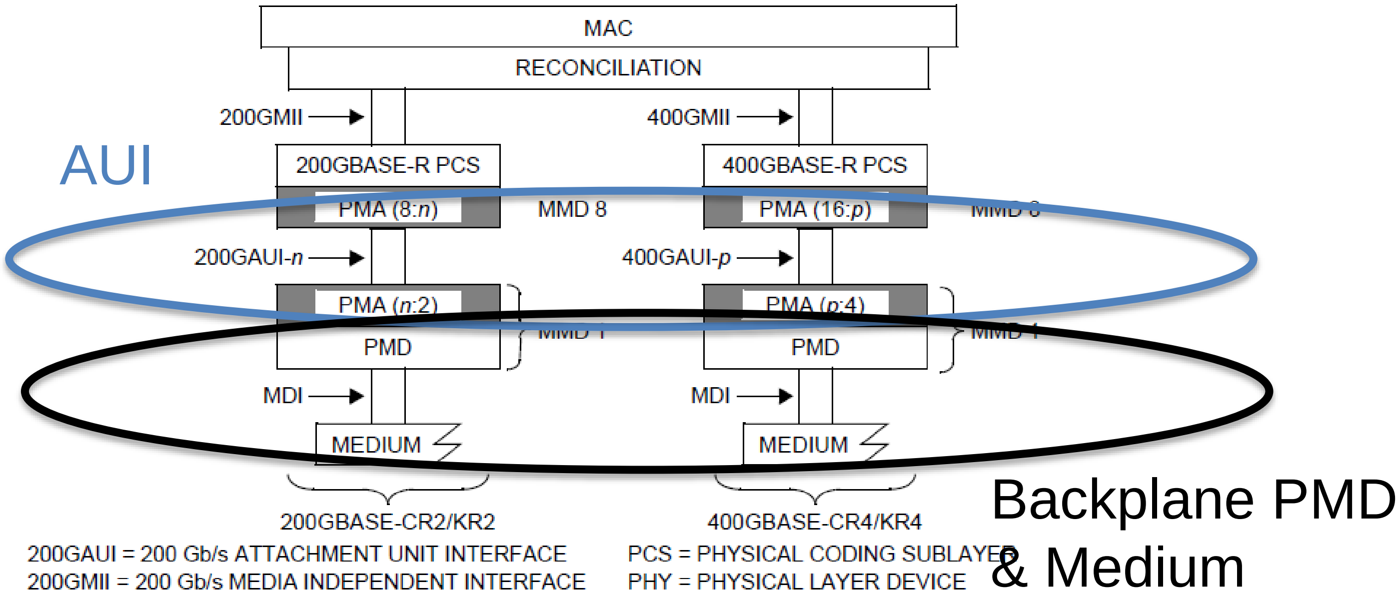

200GAUI = 200 Gb/s ATTACHMENT UNIT INTERFACE
200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE
400GAUI = 400 Gb/s ATTACHMENT UNIT INTERFACE
400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
MMD = MDIO MANAGEABLE DEVICE
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
n = 2, 4, or 8
p = 4, 8, or 16

Figure 120A–8—Example 200GBASE-CR2/KR2 and 400GBASE-CR4/KR4 PMA layering with single 200GAUI-2/4/8 or 400GAUI-4/8/16 C2C interface

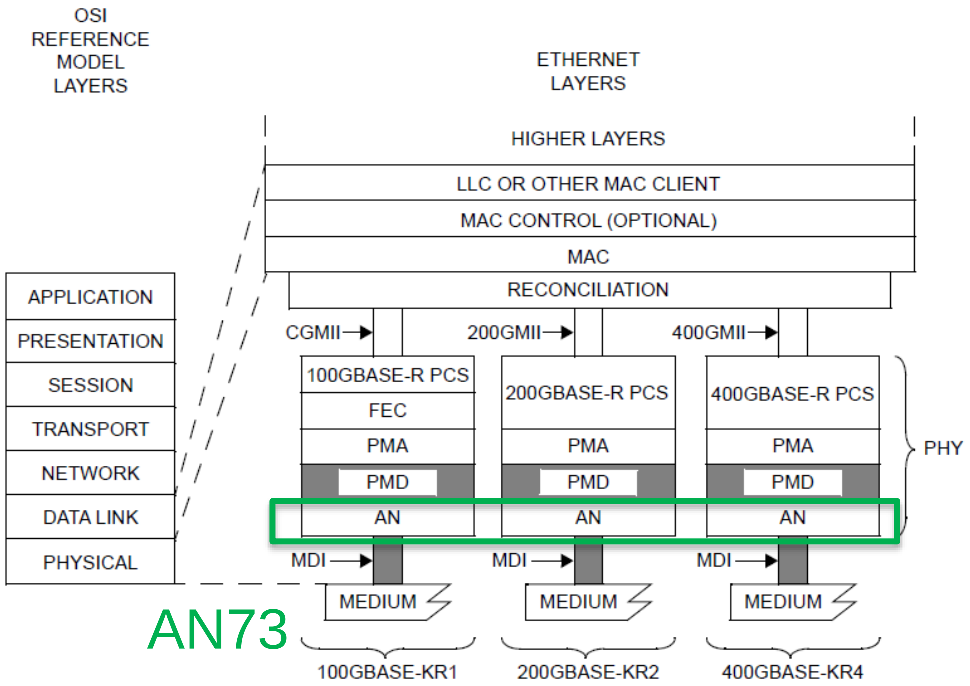

200GMII = 200 Gb/s MEDIA INDEPENDENT INTERFACE
400GMII = 400 Gb/s MEDIA INDEPENDENT INTERFACE
AN = AUTO-NEGOTIATION
CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
FEC = FORWARD ERROR CORRECTION
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
MAC = MEDIA ACCESS CONTROL
PCS = PHYSICAL CODING SUBLAYER
PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT

Figure 163–1—100GBASE-KR1, 200GBASE-KR2, and 400GBASE-KR4 relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

More Channels Available for Study

- Since January 2023, individuals created more backplane channels for study
- Channels and supporting material are available on the Task Force “Tools and Channels” website
  - (https://www.ieee802.org/3/df/public/tools/index.html)
- Summaries were provided on the next slides for reference
  - Questions about the channels should be directed to the email reflector or the channel contributor(s)
  - Detailed review planned for future electrical ad hoc meetings
- These channels can help determine the value of the insertion loss target “X” in the objectives and help to form the baseline proposals

Backplane Channel Summaries (1/4)

- TP0 to TP5 insertion losses range from 23.5dB to 27.7dB in five different model variants

KR Backplane Cable Assembly + Host

- PCB Composition
  - BGA & NPC Breakout Footprints
    - ~ 3mm PTH breakout depth
    - 8 mil vias with 5 mil stubs
    - Conforms to current PCB fab design rules
    - Nothing exotic: no skip layers, no microvias
  - Host Breakout Trace
    - Fanout length to NPC's: ~ 3 inches
    - Loss: ~ 1.25 dB/in @ 53.125 GHz
    - 90 ohm @ 6 mil line width
    - Room Temperature
- Cable Assembly Composition
  - Near Packaged Copper (NPC)
    - 95 ohm 29 AWG Twinax lengths
      - 200mm, 250mm, 300mm, 350mm, 400mm
    - Room Temperature
    - Assumes symmetric lengths on both sides of channel
  - BP Cable Connector + Twinax
    - 95 ohm 27 AWG
    - Twinax length: 800mm
    - Room Temperature

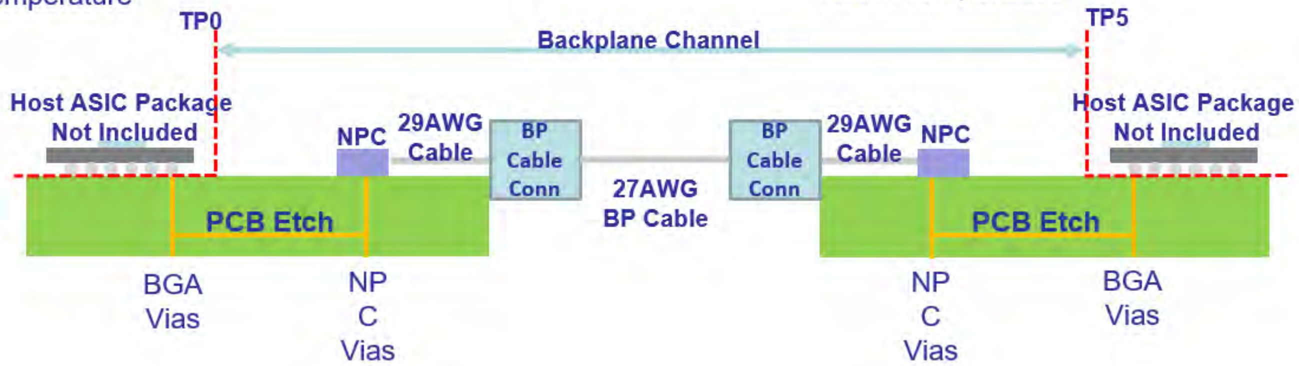

5

Jim Weaver, Jason Chan

Contribution: https://www.ieee802.org/3/dj/public/23_03/weaver_3dj_01_2303.pdf
Channels: https://www.ieee802.org/3/df/public/tools/KR/weaver_3dj_02_2303.zip

Backplane Channel Summaries(2/4)

- 27 channels
  - 16.2 dB to 33.8 dB @ 53.125 GHz
  - Two flavors of crosstalk

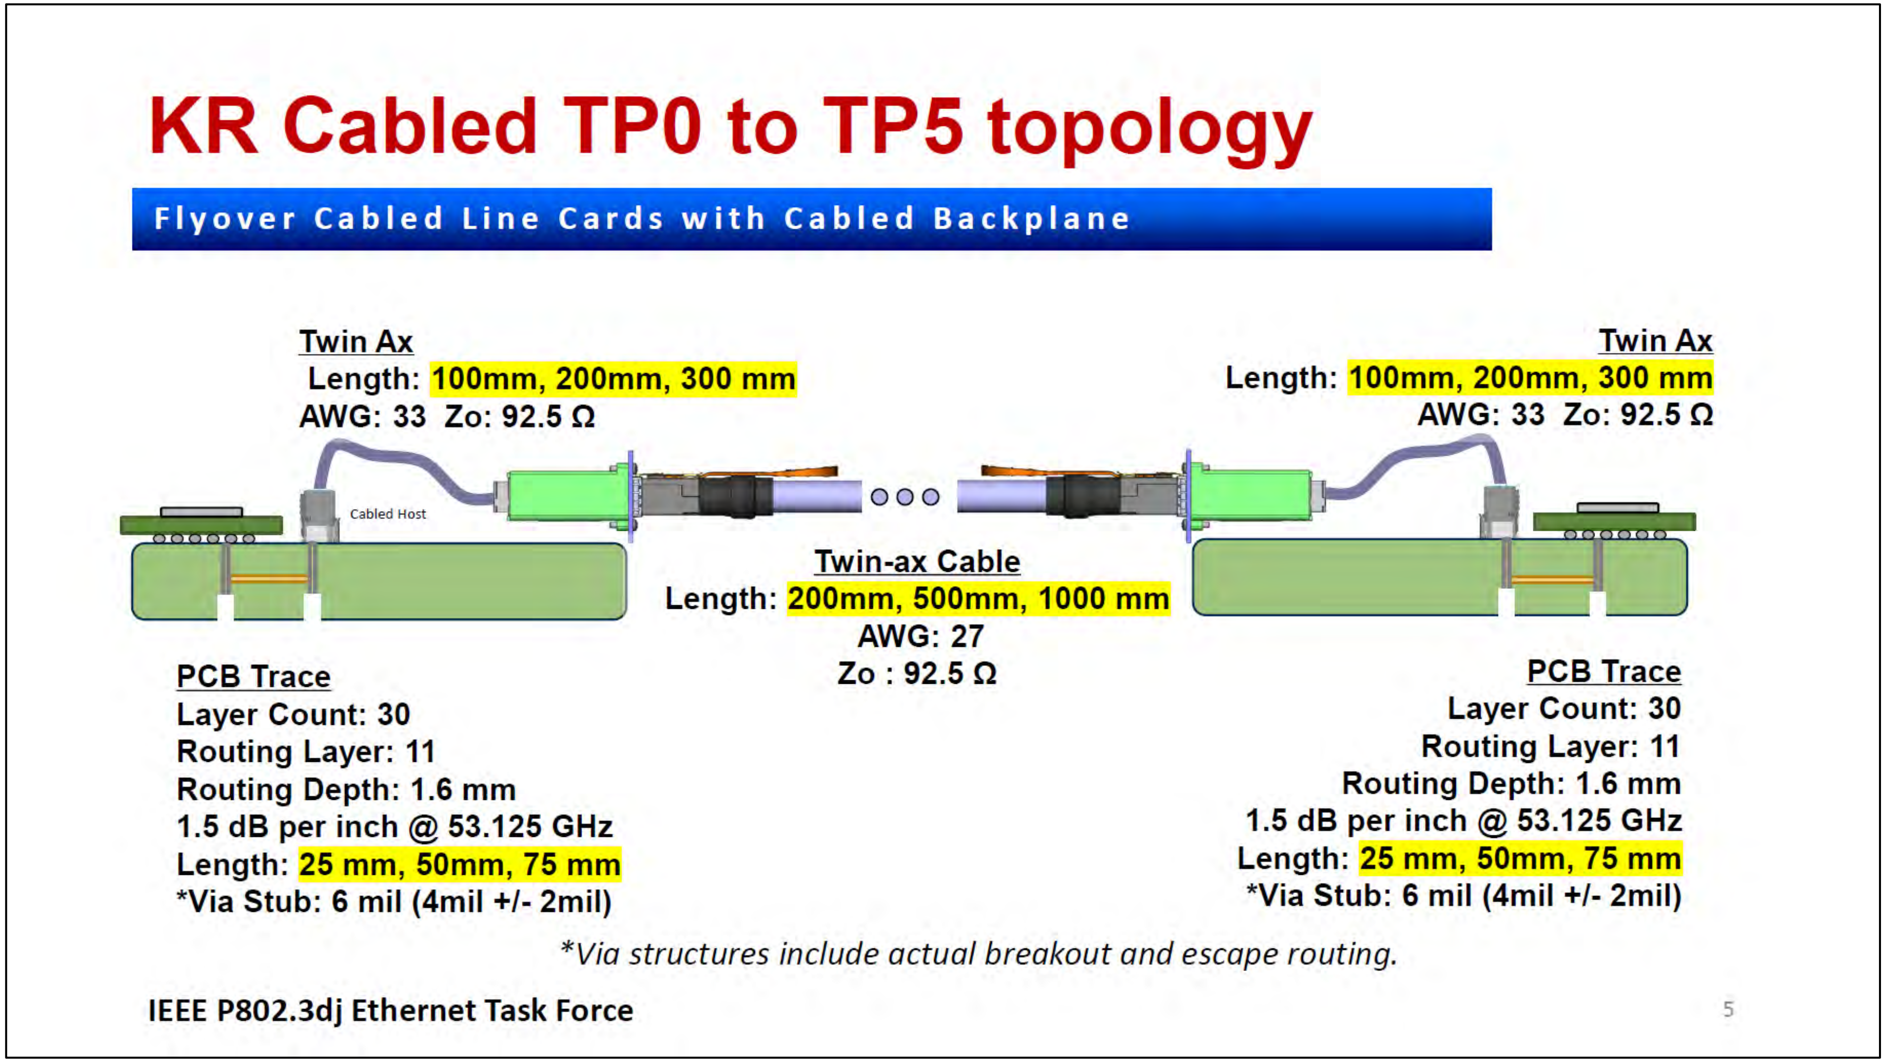

Contribution: https://www.ieee802.org/3/dj/public/23_03/mellitz_3dj_01_2303.pdf
Channels: https://www.ieee802.org/3/df/public/tools/KR/mellitz_3dj_02_2303.zip
https://www.ieee802.org/3/df/public/tools/KR/mellitz_3dj_03_2303.zip

Backplane Channel Summaries (3/4)

- 25 channels
  - 7.7 dB and 29 dB @ 53.125 GHz
  - Crosstalk not available yet

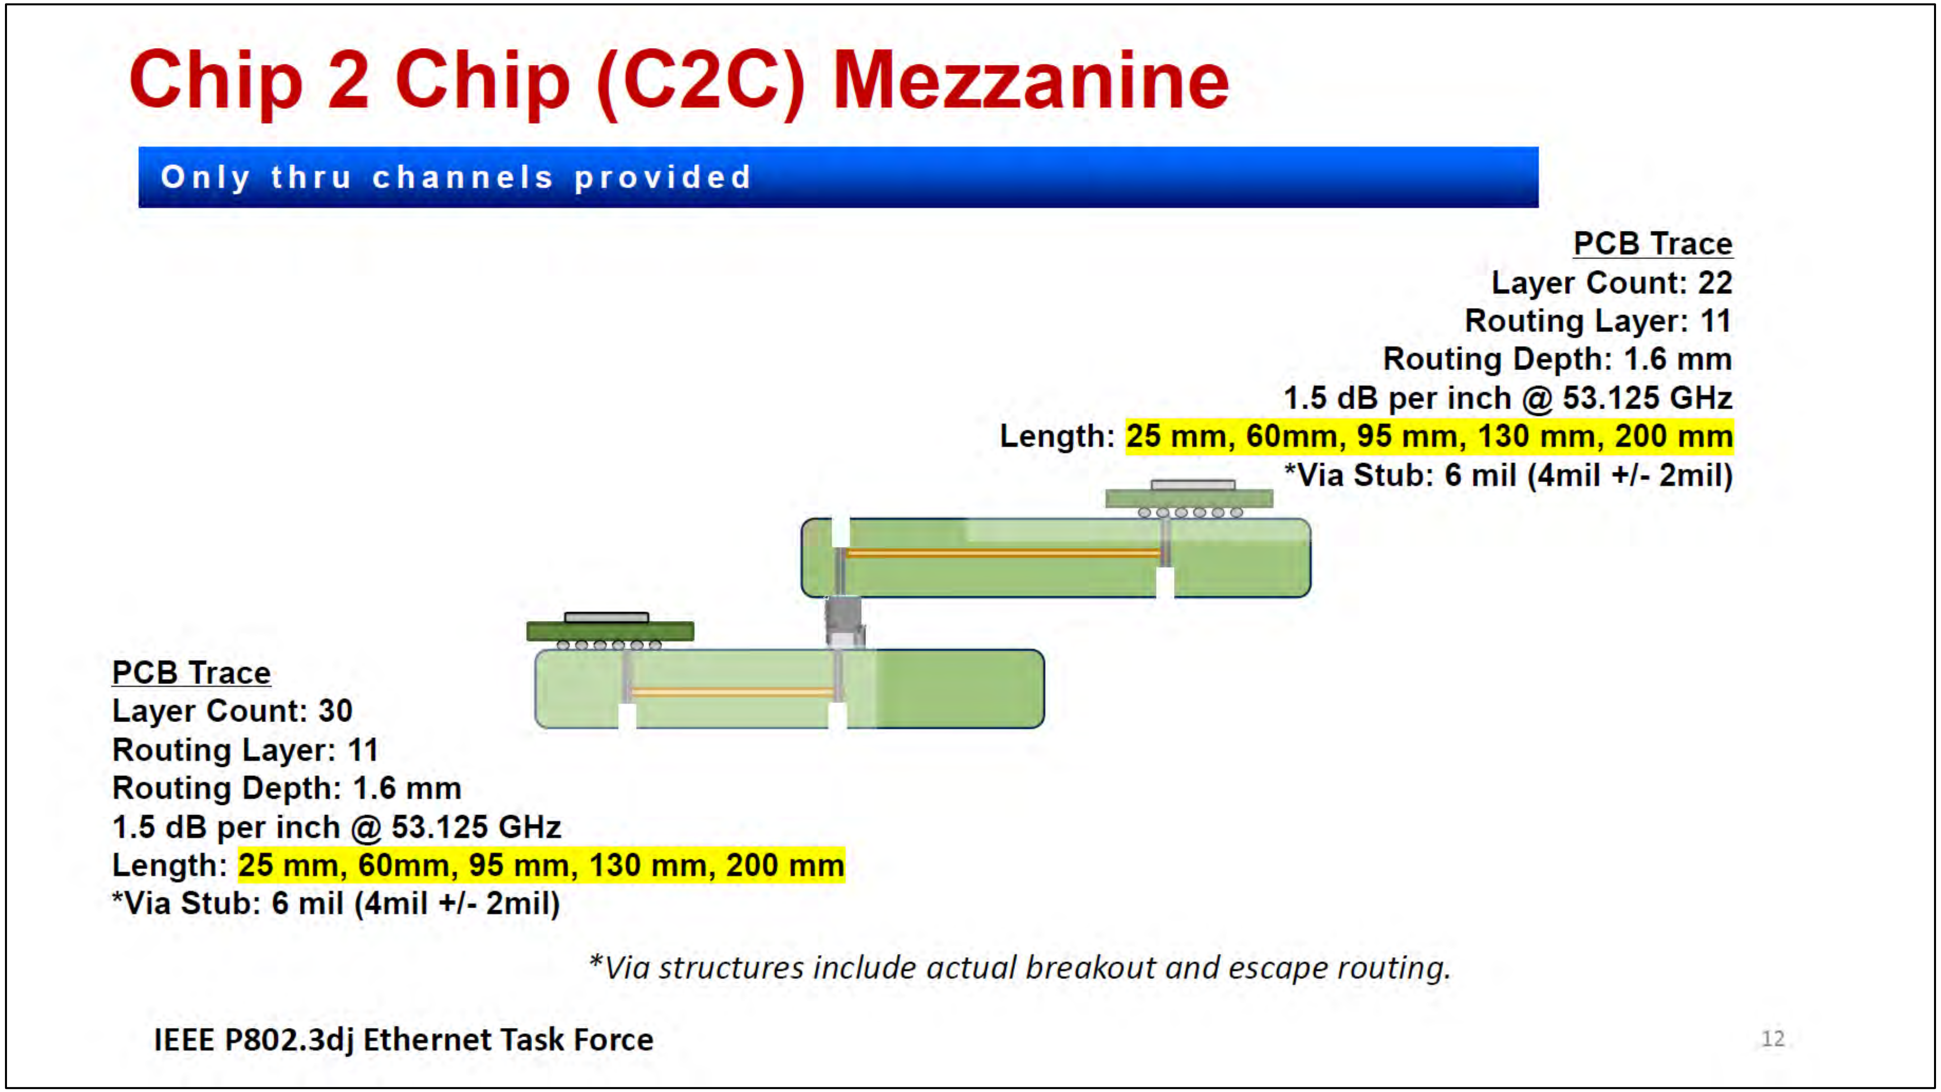

Backplane Channel Summaries (4/4)

- Both PCB right angle and cabled connector versions
- 2.7 dB and 7 dB loss host traces

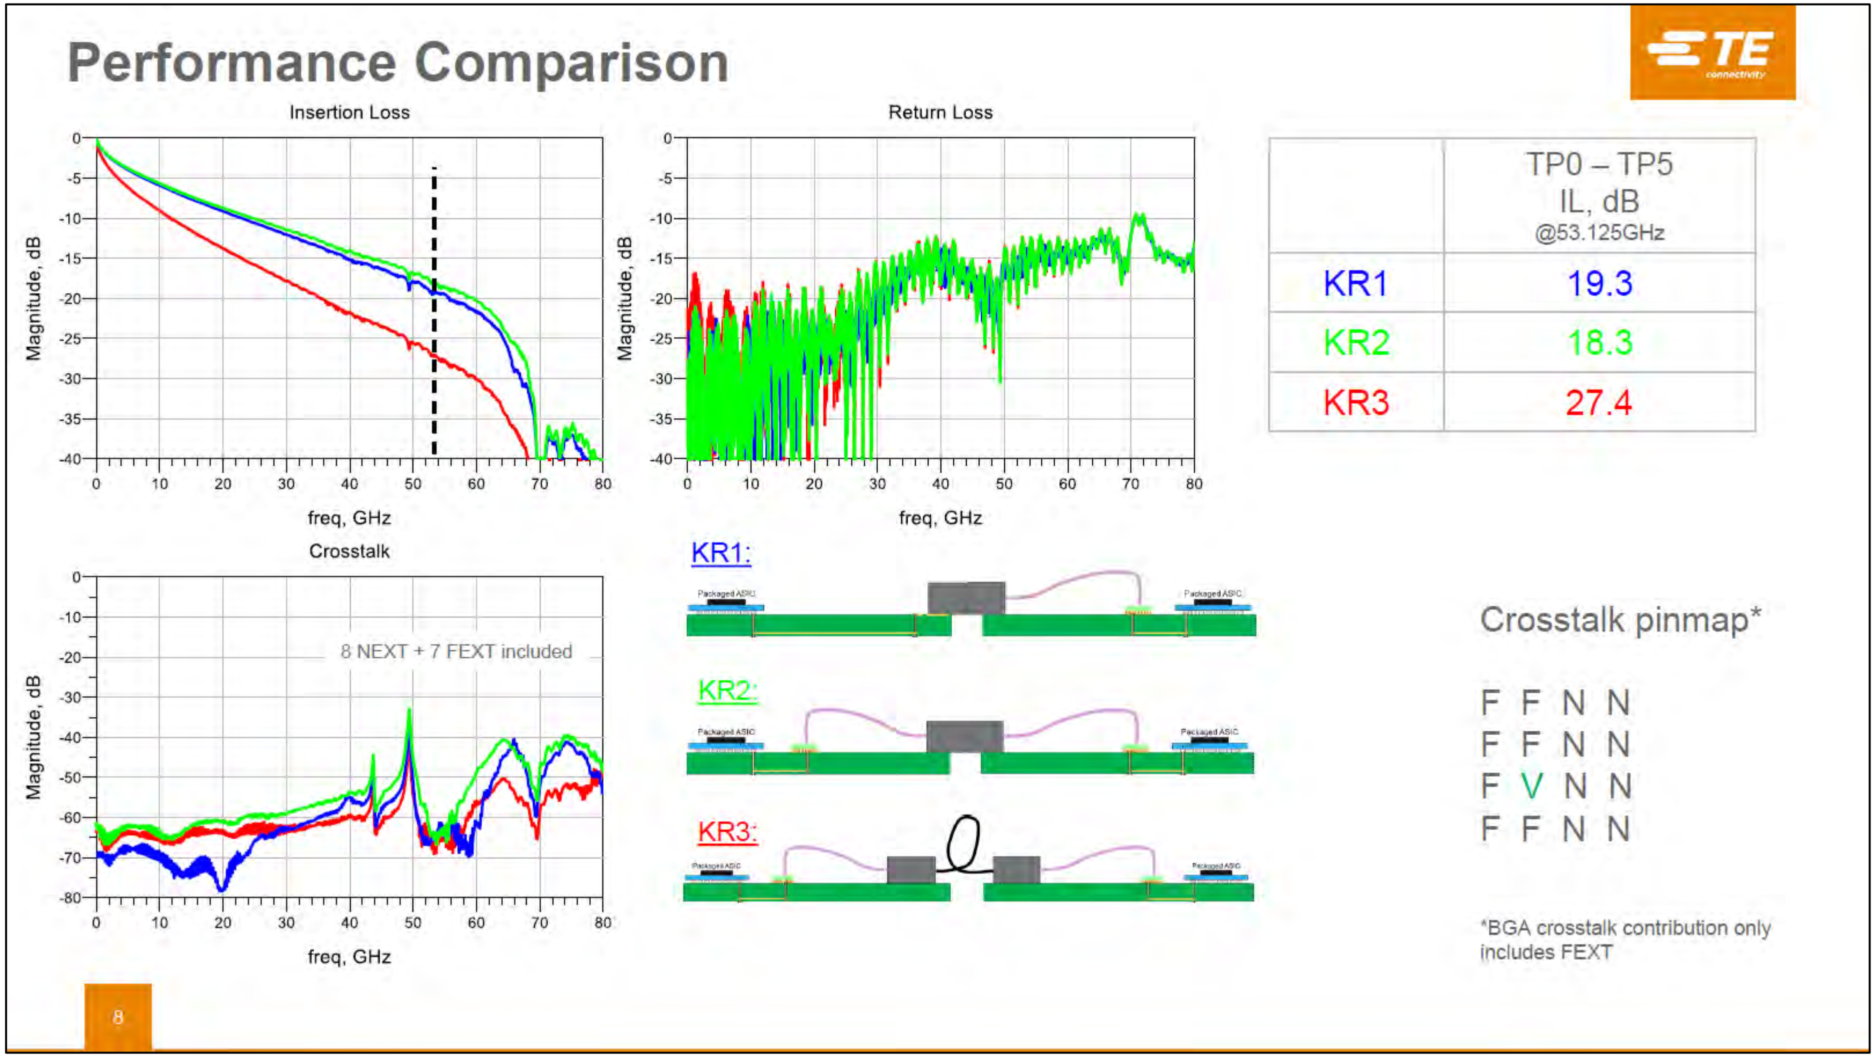

Nomenclature

- Once the objectives are adopted, nomenclature is needed for these PHY types
- We propose to use: 200GBASE-KR1, 400GBASE-KR2, 800GBASE-KR4, 1.6TBASE-KR8

Summary

- This contribution is trying to get an agreement on the objective form for backplane PHYs
  - The insertion loss target "X" is the subject of a future contribution
  - Other items, such as test point definitions and allocation of the budget, is a subject for baseline proposals
- The PMDs are distinct from AUI C2C in various ways
- More channels are available on the TF website
- Various physical instantiations of backplanes have been illustrated

Next Steps

- Provide analysis of the channels
- Agree on a value for insertion loss target "X" in the objectives
- Adopt the objectives
- Develop and adopt baseline proposals

THANKS!

Proposed Straw Poll

- I would support adopting a one-lane 200 GbE, a two-lane 400 GbE, a four-lane 800 GbE, and an eight-lane 1.6 TbE backplane objective of the form:
  - “Define a physical layer specification that supports [n*200] Gb/s operation over [n] lanes over electrical backplanes supporting a die-to-die insertion loss ≤ X dB at 53.125 GHz”
- Y: , N: , NMI:, A:

APPENDIX

Objectives as documented in P802.3ck

- Define a [n-lane, 100G/lane] PHY for operation over twin-axial copper cables with lengths up to at least 2 meters

- Define a [n-lane, 100G/lane] PHY for operation over electrical backplanes supporting an insertion loss ≤ 28dB at 26.56GHz