Channel Model Ad Hoc:
Status Summary and the Path Forward

Channel Model Ad Hoc Teleconference

October 21, 2004

Adam Healey

Agere Systems

ahealey@agere.com

Schedule of Events

- Teleconference: Thursday, October 21 (10am PDT)
  - Status summary
  - Problem statement and proposed resolution process.
- Teleconference: Thursday, November 5 (10am PST)
  - “Proposal Preview”
- Wednesday, November 10 (midnight EST)
  - Deadline for requests for presentation time.
- Tuesday, November 16 – Thursday, November 18
  - IEEE P802.3ap Task Force Meeting
  - Wyndham St. Anthony, San Antonio TX

Meeting Agenda

- Project Timeline
- The story so far…
  - Project Justification (PAR, 5 Criteria)
  - Project Objectives
  - Link Model
  - Definition of “Improved FR-4”
  - Informative Mask Set Observations
- The path ahead…
  - “Augmented Practices” vs. “Current Practices”
  - Normative Specification Methodology

Process Overview

- Adopt proposals for the formation of a baseline
- Baseline is used to create Draft 1.0
- Draft 1.0 is reviewed by the Task Force
  - Comments on the draft are generated and resolved
  - Multiple iterations
  - When Task Force review is completed and the draft is technically complete…
- …ask the 802.3 working group to authorize a working group ballot (Draft 2.0)…

Project Timeline

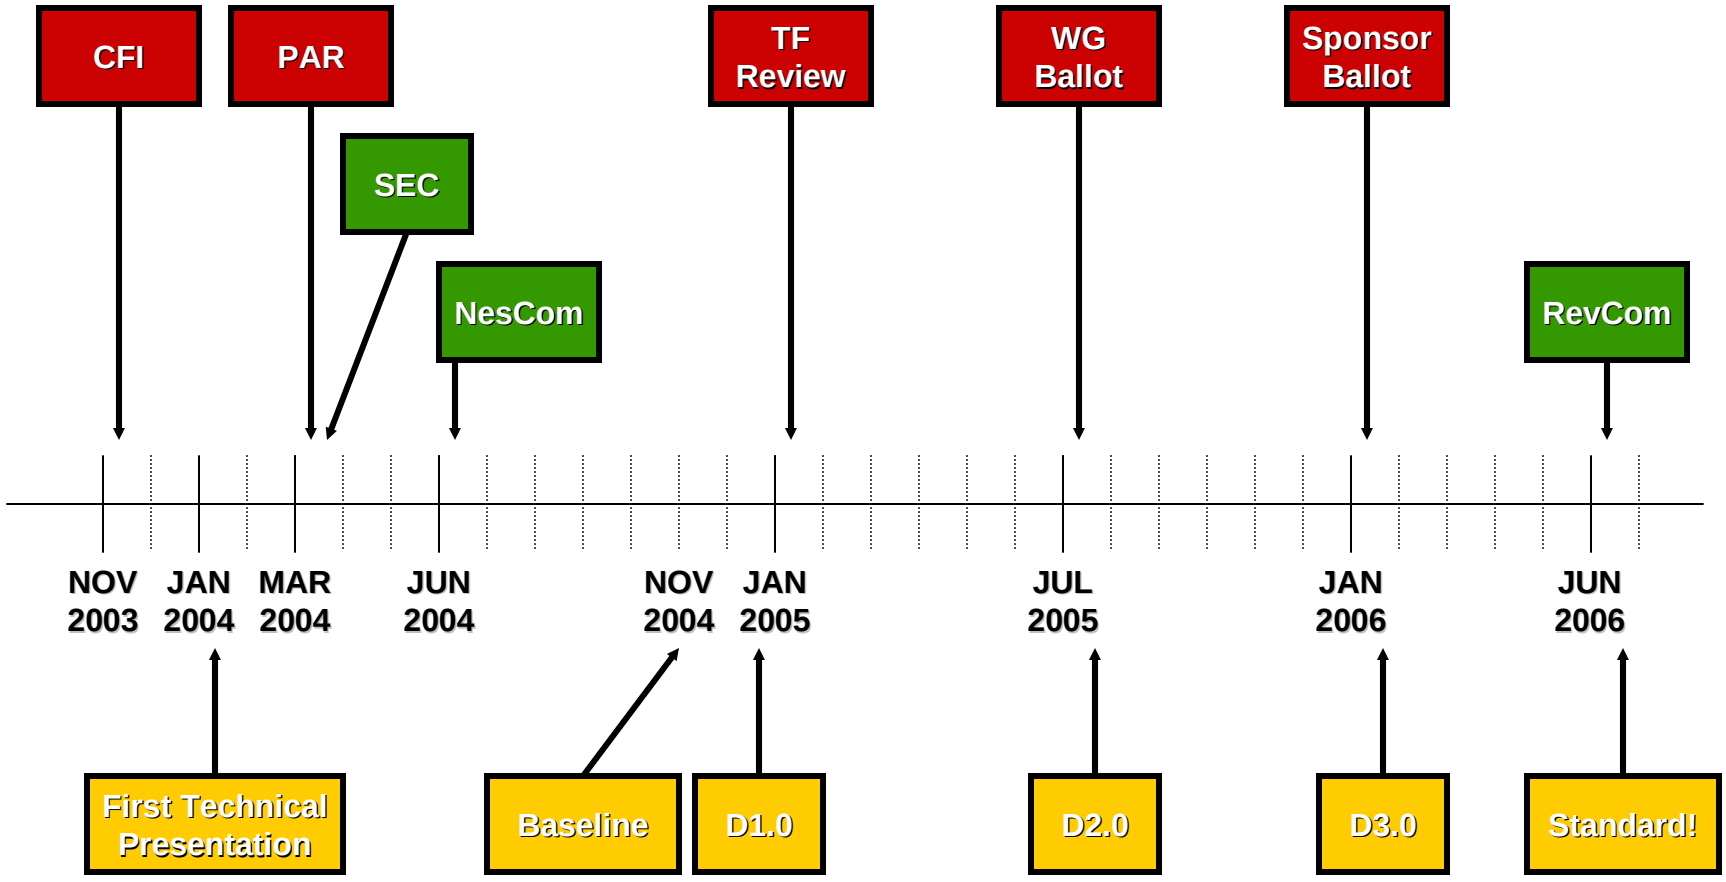

Timeline Detail

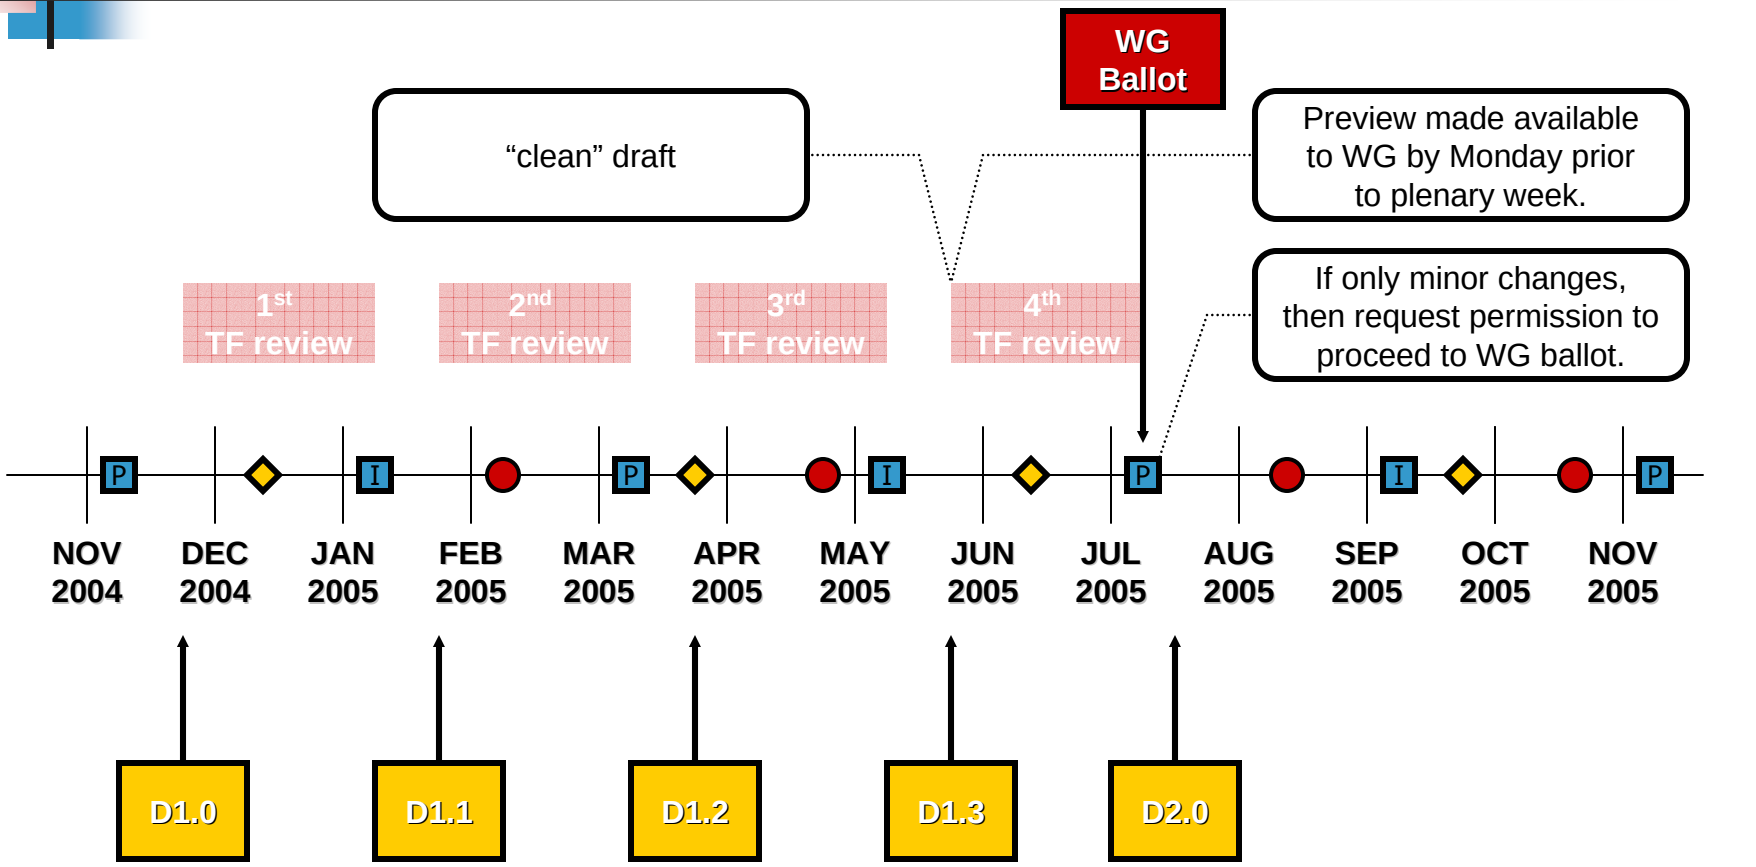

If the schedule slips, it slips 4 months

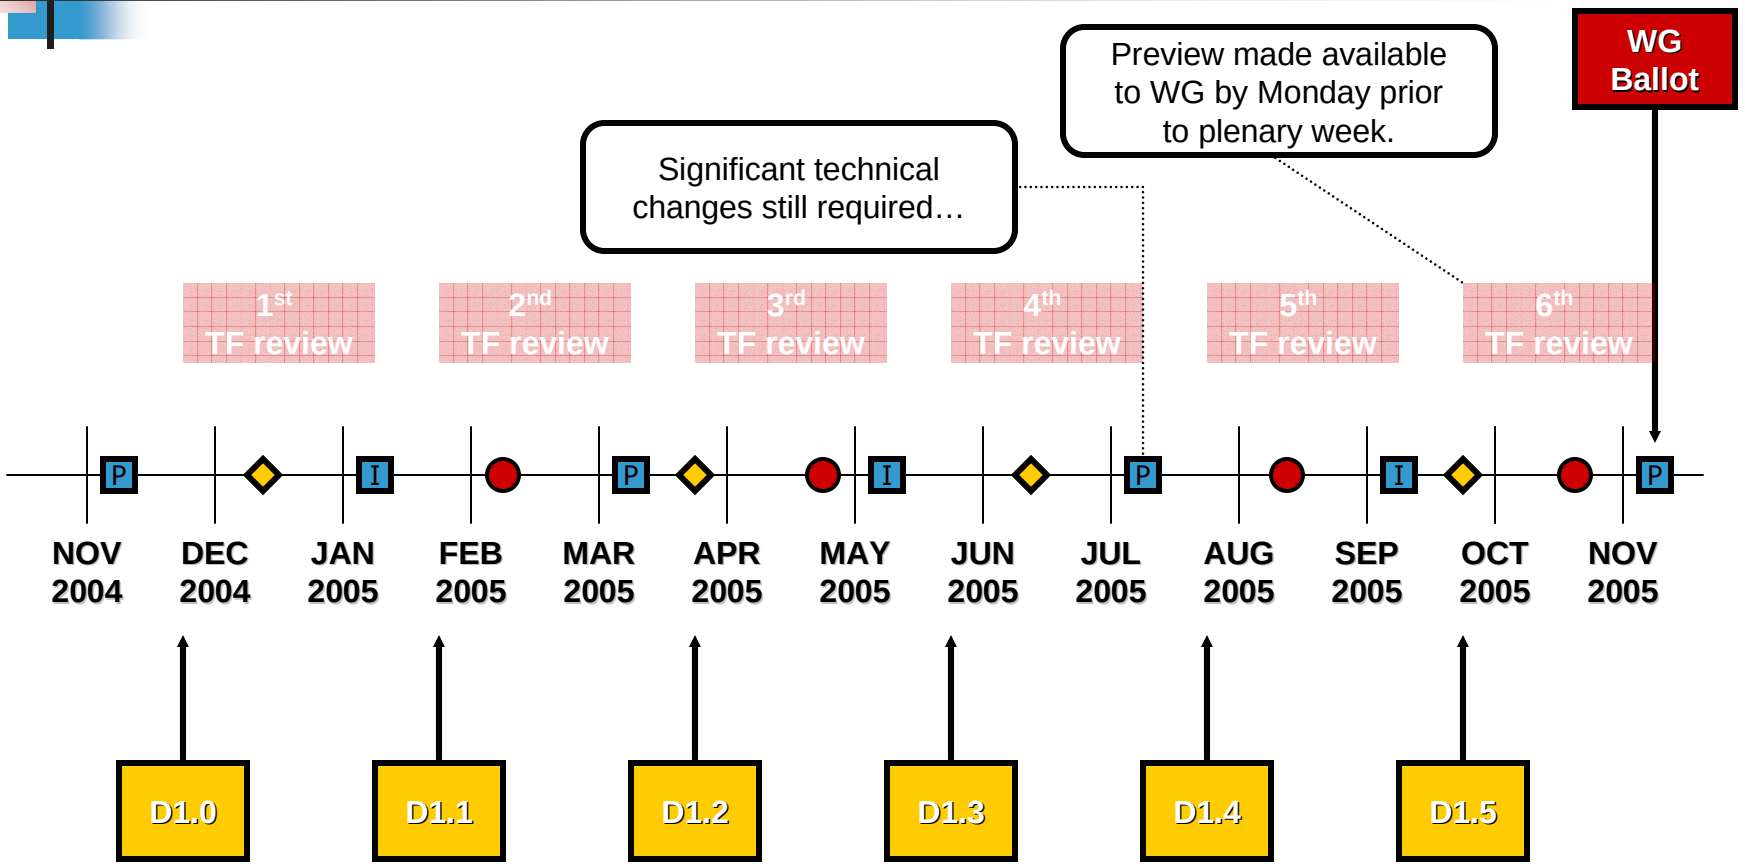

The story so far…

Project Justification (1/2)

- Project Authorization Request
  - **13. Scope of Proposed Project:**

    The scope of this project is to specify additions to and appropriate modifications of IEEE Std 802.3 *to specify operation at 1000 Mb/s and 10 Gb/s across an electrical backplane leveraging the existing MAC.*
  - **14. Purpose of Proposed Project:**

    The purpose of this project is *to provide standards based Ethernet interconnection of server and telecommunication blades over a modular platform backplane*. Industry trends for LAN, SAN and other applications are migrating to backplane interconnects, and this project will optimize Ethernet operation for backplanes.

Project Justification (2/2)

Broad Market Potential

Broad set(s) of applications

Multiple vendors, multiple users

Balanced cost (LAN vs. attached stations)

- **Ethernet has become widely deployed as a preferred backplane solution. Examples include Modular Servers and Enterprise and Telecom Network Equipment. Quantitative presentations have been made to the 802.3 Backplane Ethernet Study Group indicating significant market opportunities for these applications.**
- **Rapid growth of network and internet traffic is driving the need for higher performance over backplanes. Currently, IEEE 802.3 does not address this application with a formal standard.**
- **156 participants attended the Ethernet Over Backplane call-for-interest, representing at least 33 companies, and indicated that they plan to participate in the standardization of Ethernet Over Backplane. This level of commitment indicates that a standard will be developed by a large group of vendors and users.**
- **A standardized Ethernet interface on blades will maintain the balanced cost for backplane applications.**

SG Discussions (References)

- Call for Interest
  - http://ieee802.org/3/bladesg/public/nov03/BackplaneRMG.pdf
    - Modular server and ATCA market projections
- Market Drivers
  - http://ieee802.org/3/bladesg/public/jan04/hegde_01_0104.pdf
    - Modular server market projections
- Project Scope
  - http://ieee802.org/3/bladesg/public/jan04/lerer_02_0104.pdf
    - Recommended scope includes “Existing ATCA Backplanes”

System Requirements Presentations

- Modular Server Requirements
  - http://ieee802.org/3/ap/public/may04/koenen_01_0504.pdf
- Enterprise/Telecom Requirements
  - http://ieee802.org/3/bladesg/public/mar04/goergen_01_0304.pdf
  - http://ieee802.org/3/ap/public/jul04/mandich_01_0704.pdf
- 10G Serial PHY Requirements
  - http://ieee802.org/3/bladesg/public/mar04/altmann_01_0304.pdf
    - Compatible with current connectors and routing (example ATCA)
  - http://ieee802.org/3/bladesg/public/mar04/palkert_01_0304.pdf
    - Support ATCA connector and trace density
- ATCA Requirements
  - http://ieee802.org/3/ap/public/may04/kundu_01_0504.pdf

Requirements Summary

Description	N1 (mm)	N2 (mm)	B (mm)	H (mm)	Total (mm)	No. Connectors	AC / DC Coupling	Source
Blade Server								
Proposed Worst-Case	76	102	533	127	838	3	AC	*koenen_01_0504.pdf*
ATCA								
Full Mesh (max)	0	127	533	127	787	2	AC	*(note 1)*
Switch / Router								
2 to 3 chassis/rack (min)	0	152	51	305	508	2	AC	*goergen_01_0304.pdf (note 2)*
2 to 3 chassis/rack (max)	0	152	559	305	1016			
5 to 8 chassis/rack (min)	0	127	51	229	406			
5 to 8 chassis/rack (max)	0	127	432	229	787			
					700	2	AC or DC	*mandich_01_0704.pdf*
					1000		AC	
ATCA Example (Star)								
min(B)	0	102	28	102	231	2	AC	*peters_01_0504.pdf*
max(B)	0	102	244	102	447			

Note 1: From PICMG 3.0 R1.0 AdvancedTCA Specification, December 30, 2002 (8.4.2.1 and 8.2.4.3).

Note 2: Based on LC-2/SF-2. For minimum values, fabric position is assumed to be in the middle of the line cards. For maximum values, fabric position is assumed to be at the top of the line cards.

No *specific* guidelines in terms of materials or the use of stub-reduction techniques. System requirements stated in terms of the *solution* cost/power relative to XAUI.

Project Objectives

IEEE P802.3ap Objectives

- Preserve the 802.3/Ethernet frame format at the MAC Client service interface.
- Preserve min. and max. frame size of current 802.3 Std.
- Support existing media independent interfaces.
- Support operation over a single lane across 2 connectors over copper traces on improved FR-4 for links consistent with lengths up to at least 1m.
  - Define a 1 Gb/s PHY
  - Define a 10 Gb/s PHY
- Define a 4-lane 10Gb/s PHY for operation over the 802.3ap channel model.
- Consider auto-negotiation.
- Support BER of 10^-12 or better.
- Meet CISPR/FCC Class A.

Test Points: Historical View

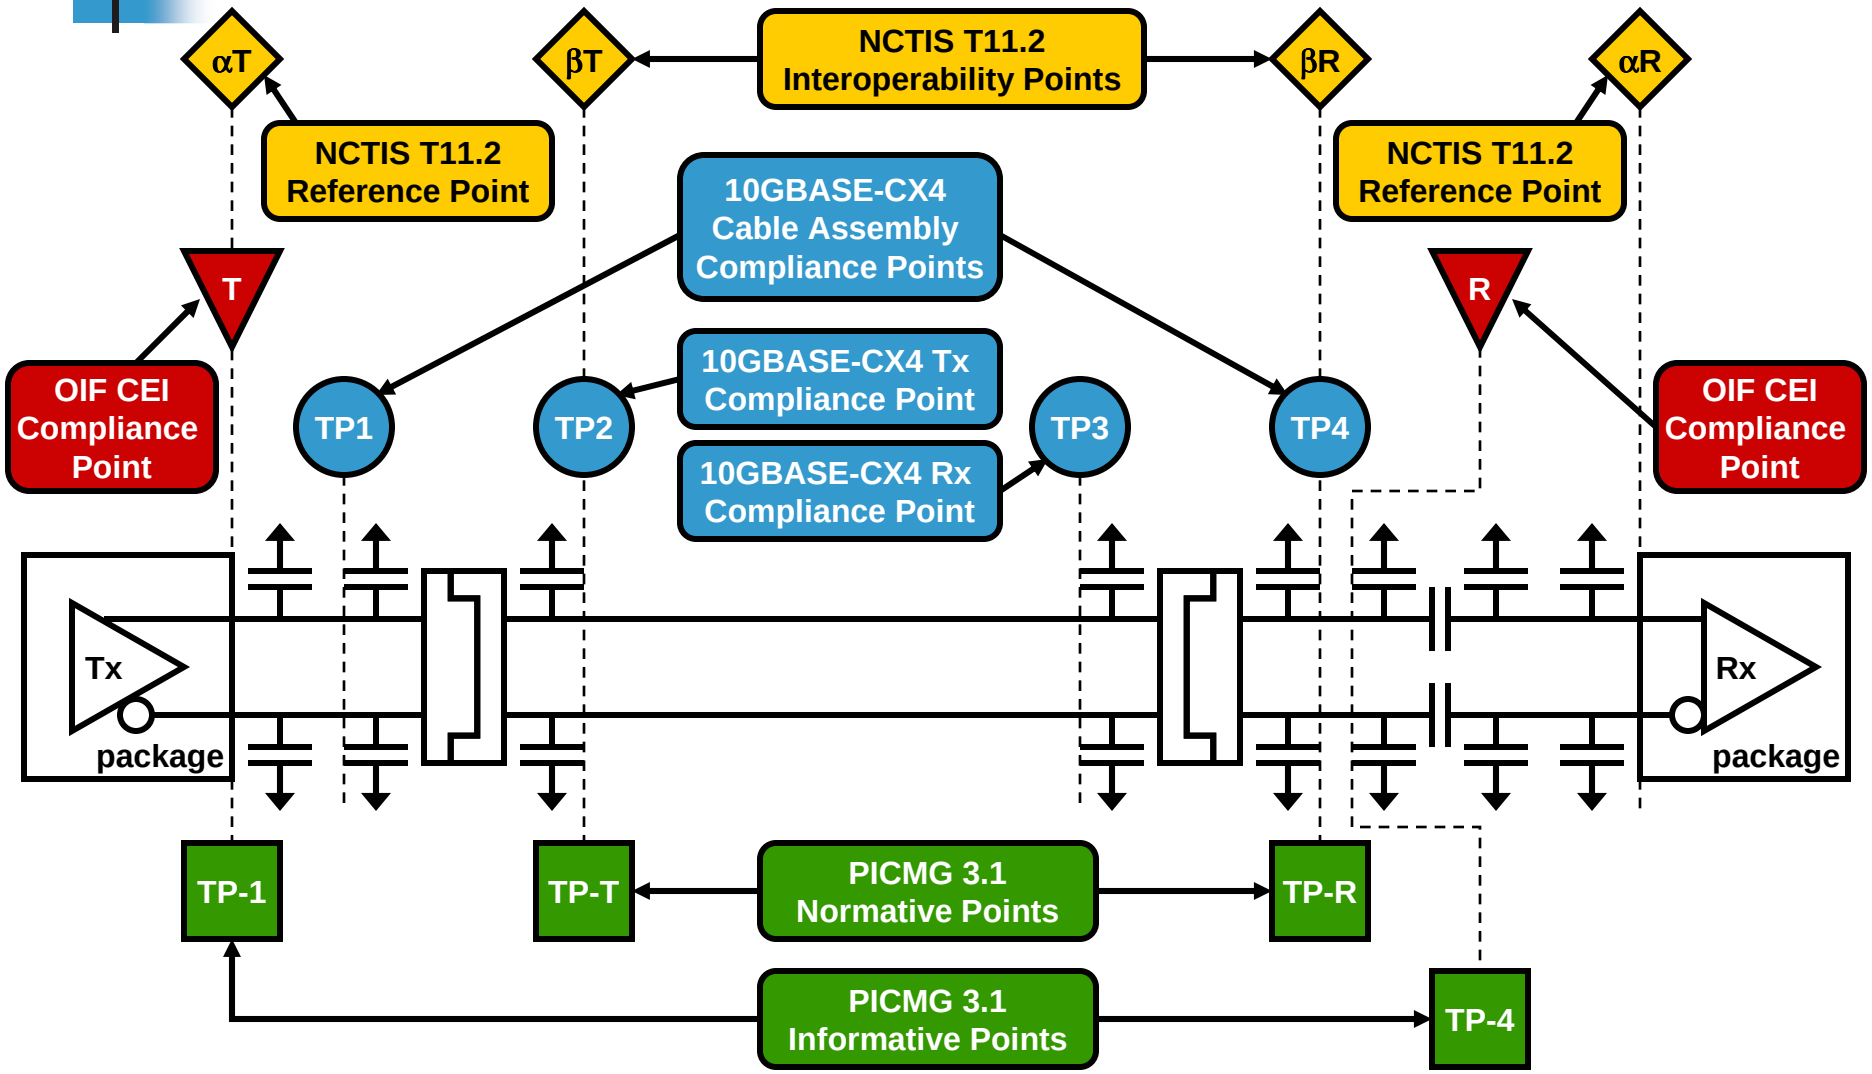

IEEE P802.3ap Link Model

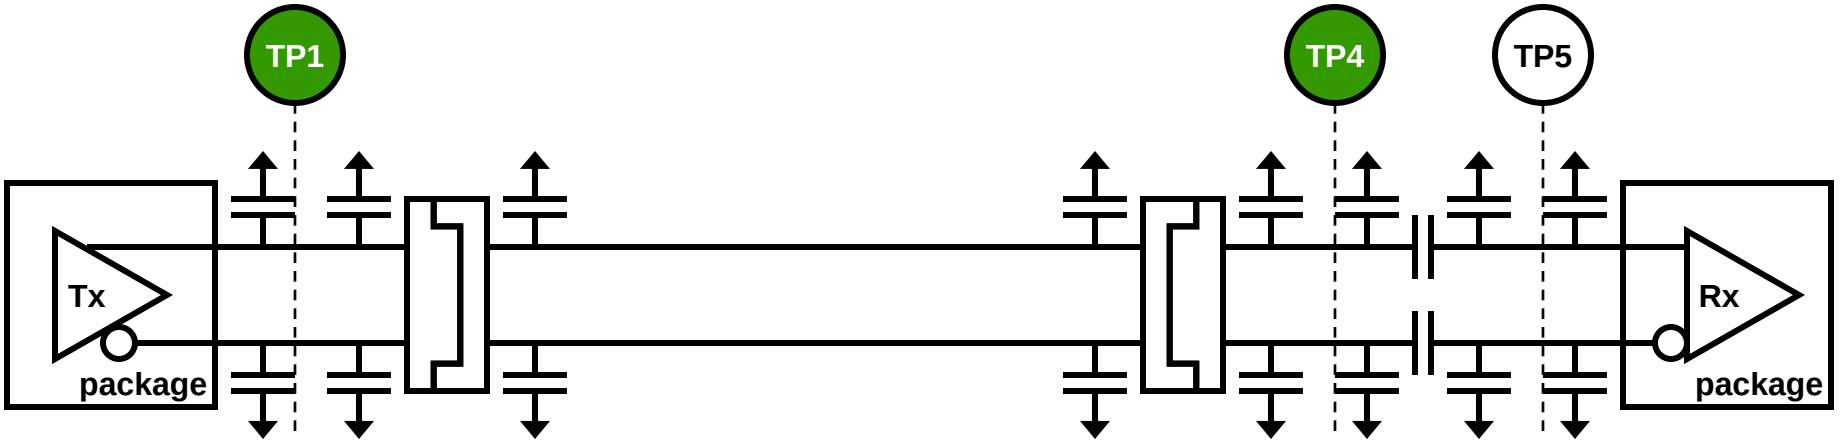

X = Normative

Y = Informative

Note 1: This definition is consistent with conventions adopted in XAUI, OIF TFI-5 and CEI, and PICMG 3.1

Note 2: While only two connectors are shown, a three connector topology may also reside between TP1 and TP4, so long as the channel requirements are met.

Definition adopted via TF Motion July 2004 (Y:32, N:2, A:21)

Definition of "Improved FR-4"

Material Map

- **Results from Nelco N4000-13SI**
  - **improved FR-4**

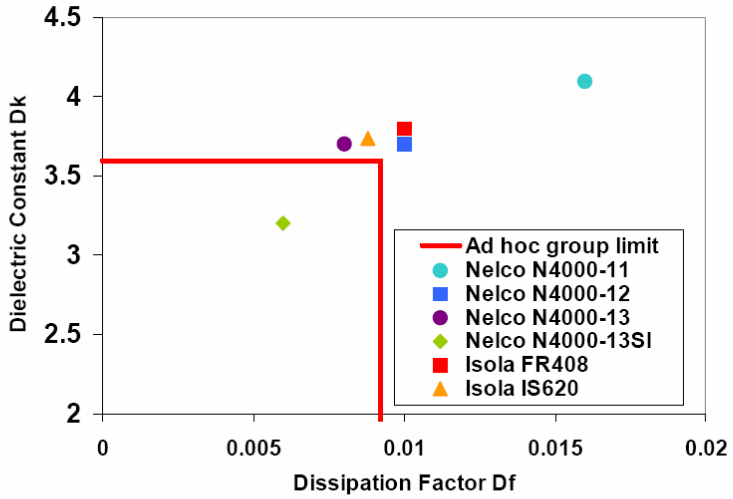

@1GHz

My Thoughts on 'Improved FR-4' in reference to IEEE802.3ap

- Improved FR-4 (Mid Resolution Signal Integrity):
  - 100Mhz: Dk ≤ 3.60; Df ≤ .0092
  - 1Ghz: Dk ≤ 3.60; Df ≤ .0092
  - 2Ghz: Dk ≤ 3.50; Df ≤ .0115
  - 5Ghz: Dk ≤ 3.50; Df ≤ .0115
  - 10Ghz: Dk ≤ 3.40; Df ≤ .0125
  - 20Ghz: Dk ≤ 3.20; Df ≤ .0140
- Temperature and Humidity Tolerance (0-~~55degC~~ **70degC**, 10-90% non-condensing):
  - Dk:+/- .04
  - Df: +/- .001
- Resin Tolerance (standard +/-2%):
  - Dk:+/- .02
  - Df: +/- .0005

12 July 2004 9:30 PM

Definition adopted via TF Motion July 2004 (Y:41, N:0, A:6)

Cost Implication of "Improved FR-4"

Relative Cost Data

- **N4000-13SI**
  - Unprocessed cost ~3.4 x unprocessed FR-4
  - Processed cost ~1.8 x processed FR-4

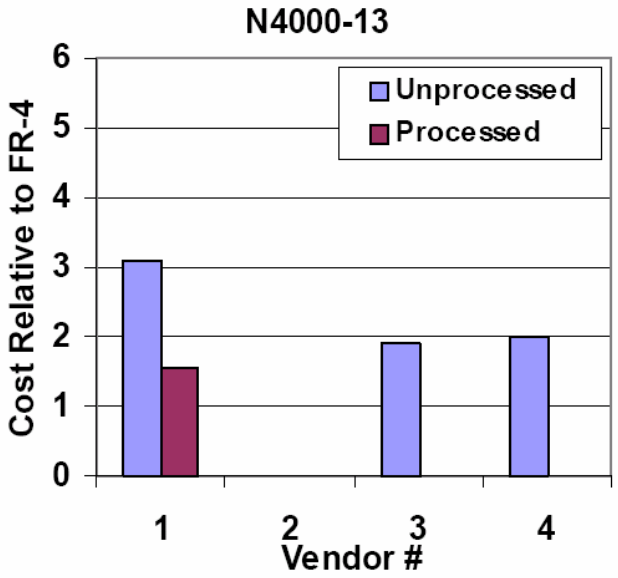

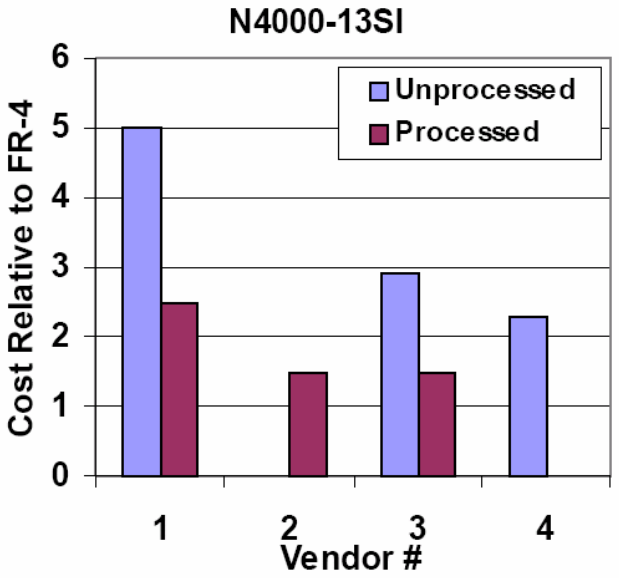

molex

IEEE Backplane Ethernet

SDD21 Magnitude Mask (Proposed)

- TP1 to TP4
- Based on a set of assumed design practices:
  - 1m total channel length
  - 20” Backplane
  - 10” Node/Hub Cards
  - “Improved FR-4”
  - W = 6 mils
  - Stubs not exceeding 30 mils

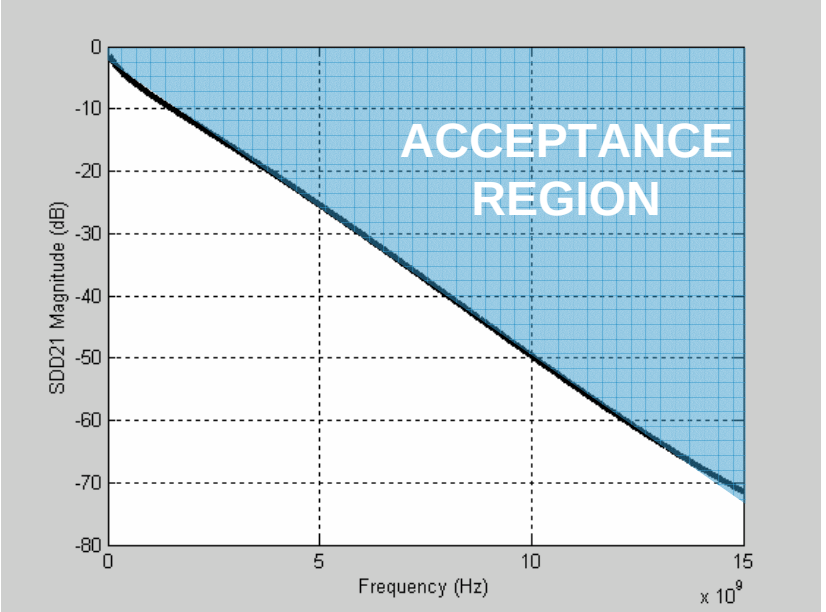

Relationship to Other Specifications

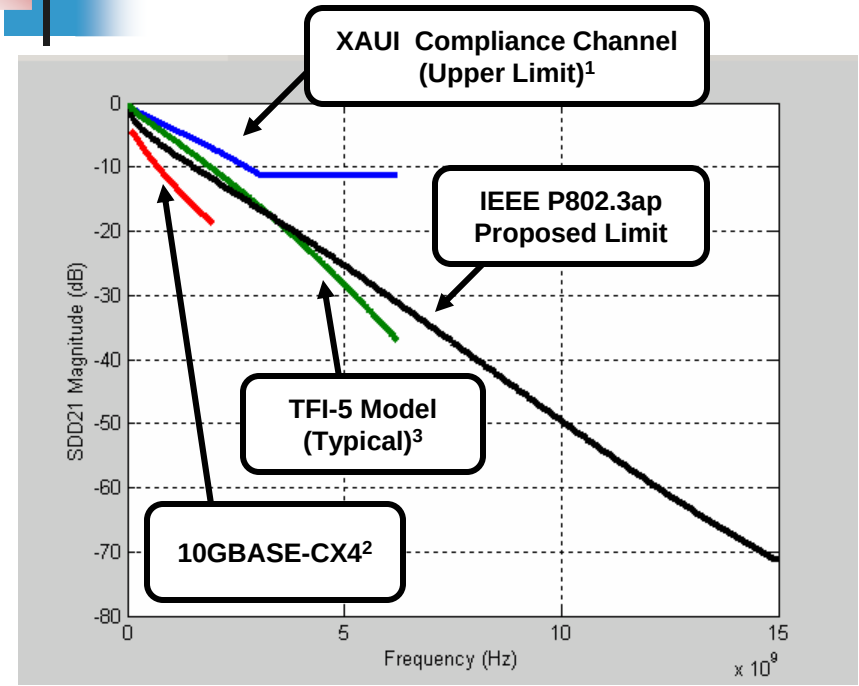

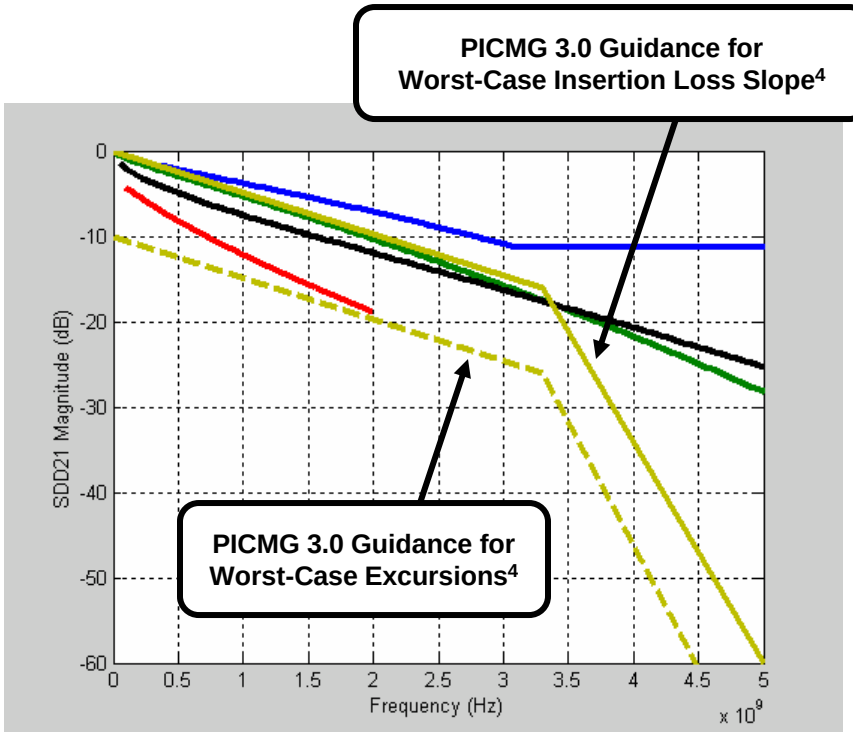

[1] IEEE 802.3ae-2002, 47.4.1
[2] IEEE 802.3ak-2004, 54.7.2
[3] OIF-TFI5-0.1.0, http://oiforum.com/public/documents/OIF-TFI5-01.0.pdf
[4] PICMG 3.0 R1.0 (December, 30, 2003), Figure 8-7

Implementation Examples (minimal stub)

"XAUI-like" (50cm FR-4)

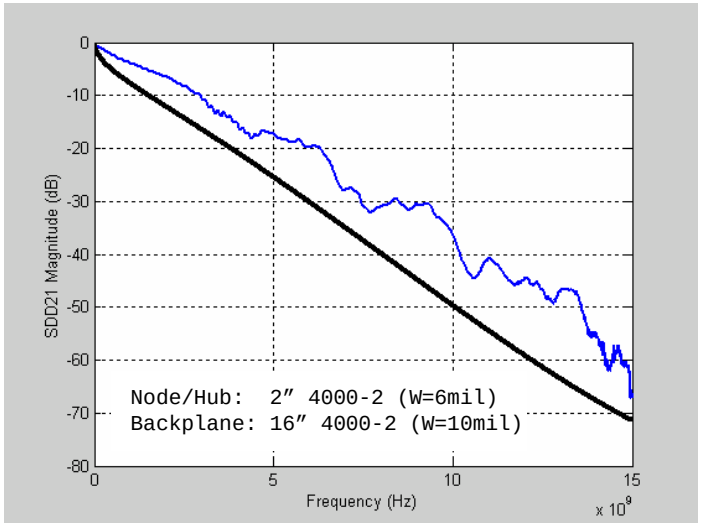

"TFI-5-like" (75cm FR-4)

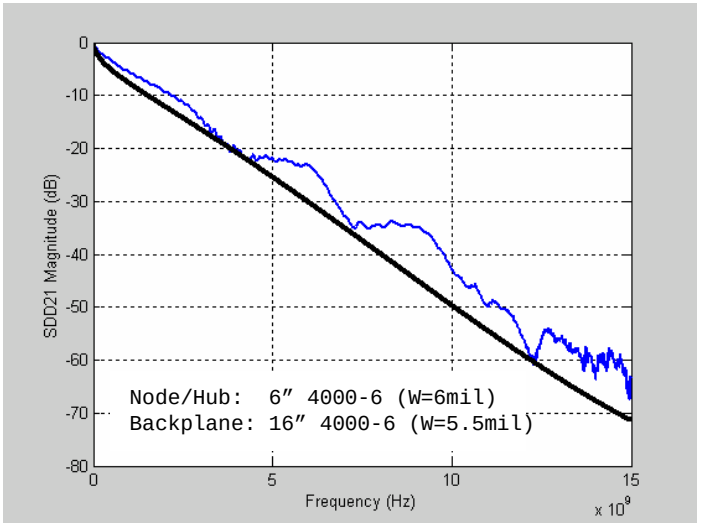

"ATCA-like"

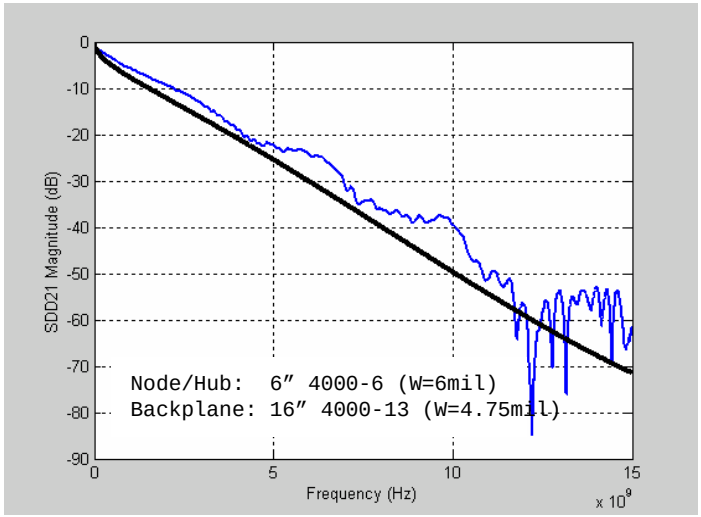

1m "Improved FR-4"

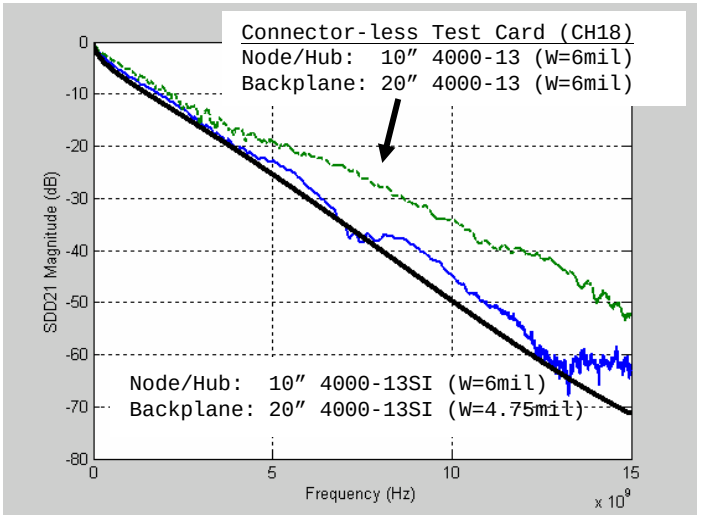

Test Cases 5, 6, and 7

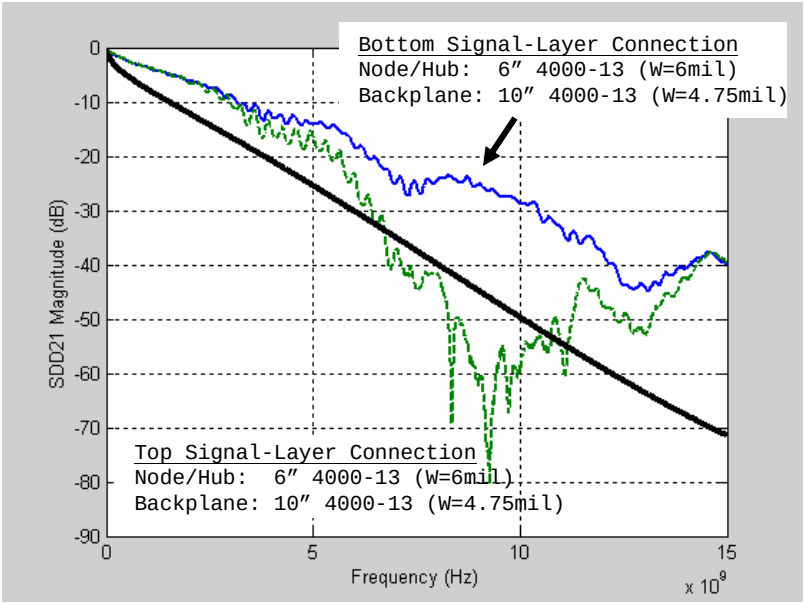

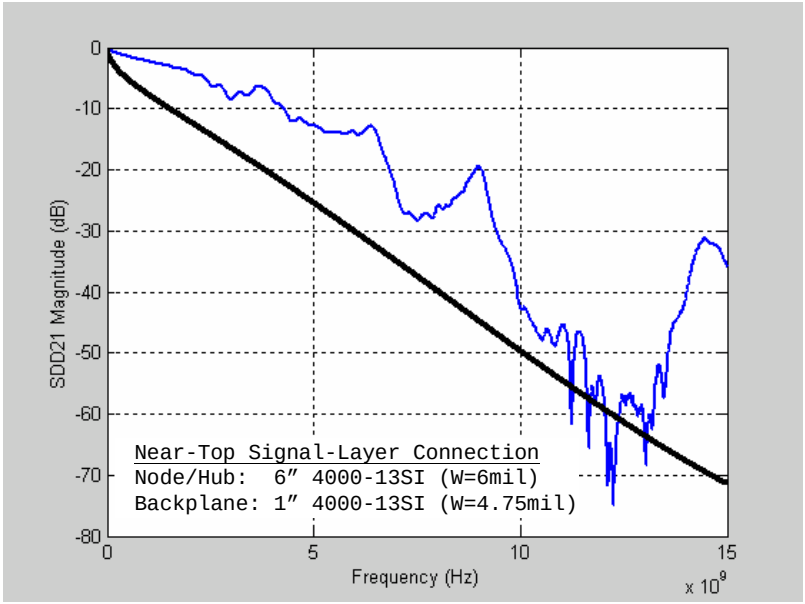

Current mask fails channels with significant stub-related resonances.

ATCA on our minds…

- http://ieee802.org/3/bladesg/public/jan04/seemann_01_0104.pdf
  - Simulation results showing open eyes on an ATCA backplane channel
- http://ieee802.org/3/bladesg/public/mar04/dambrosia_01_0304.pdf
- http://ieee802.org/3/bladesg/public/mar04/oltmanns_01_0304.pdf
  - Cost impact of enhanced materials, counter-boring, with ATCA-based examples
- http://ieee802.org/3/ap/public/sep04/sinsky_01_0904.pdf
- http://ieee802.org/3/ap/public/sep04/abler_01_0904.pdf
- http://ieee802.org/3/ap/public/sep04/liu_01_0904.pdf
  - Performance simulations using ATCA-based models

Observations

- To meet the proposed SDD21 mask at 40”, use of “Improved FR-4” is required.
  - Lower cost materials can be shown to satisfy the mask at shorter, but relevant, distances.
  - Proposed mask is in-line with standards representing “legacy” applications (at 3GHz and below).
- However, proposed SDD21 magnitude mask is not tolerant of stub effects.
  - Implication is that some stub-reduction technique must be applied.
- Interest in supporting ATCA-based applications.
  - Support ATCA = support stub effects?
- **The core issue:** To what degree do we support stubs and what impact does this decision have on Broad Market Potential.
- Recent focus has been completion of the informative mask set.
  - This is appropriate, but eventually, a normative specification must be prepared for inclusion into the draft.

The path ahead…

“Augmented” vs. “Current” Practices

- “Augmented” Practices Model
  - Basis of the channel model ad hoc’s work to date (refer to slide 19).
  - Flexible model, except significant stub effects not tolerated.
- “Current” Practices Model
  - Consistent with what is being done in ATCA systems today.
  - Shorter (than 1m) channels:
    - 20” dual-star, hub cards centered in the chassis
    - up to 31” full-mesh, or dual star with hub cards located at either end of the chassis
  - Lower-cost materials may be employed.
  - More significant stub effects than those currently allowed in “augmented” practices model.
  - Does not necessarily represent the entirety of the installed base.

Relationship to the “Installed Base”

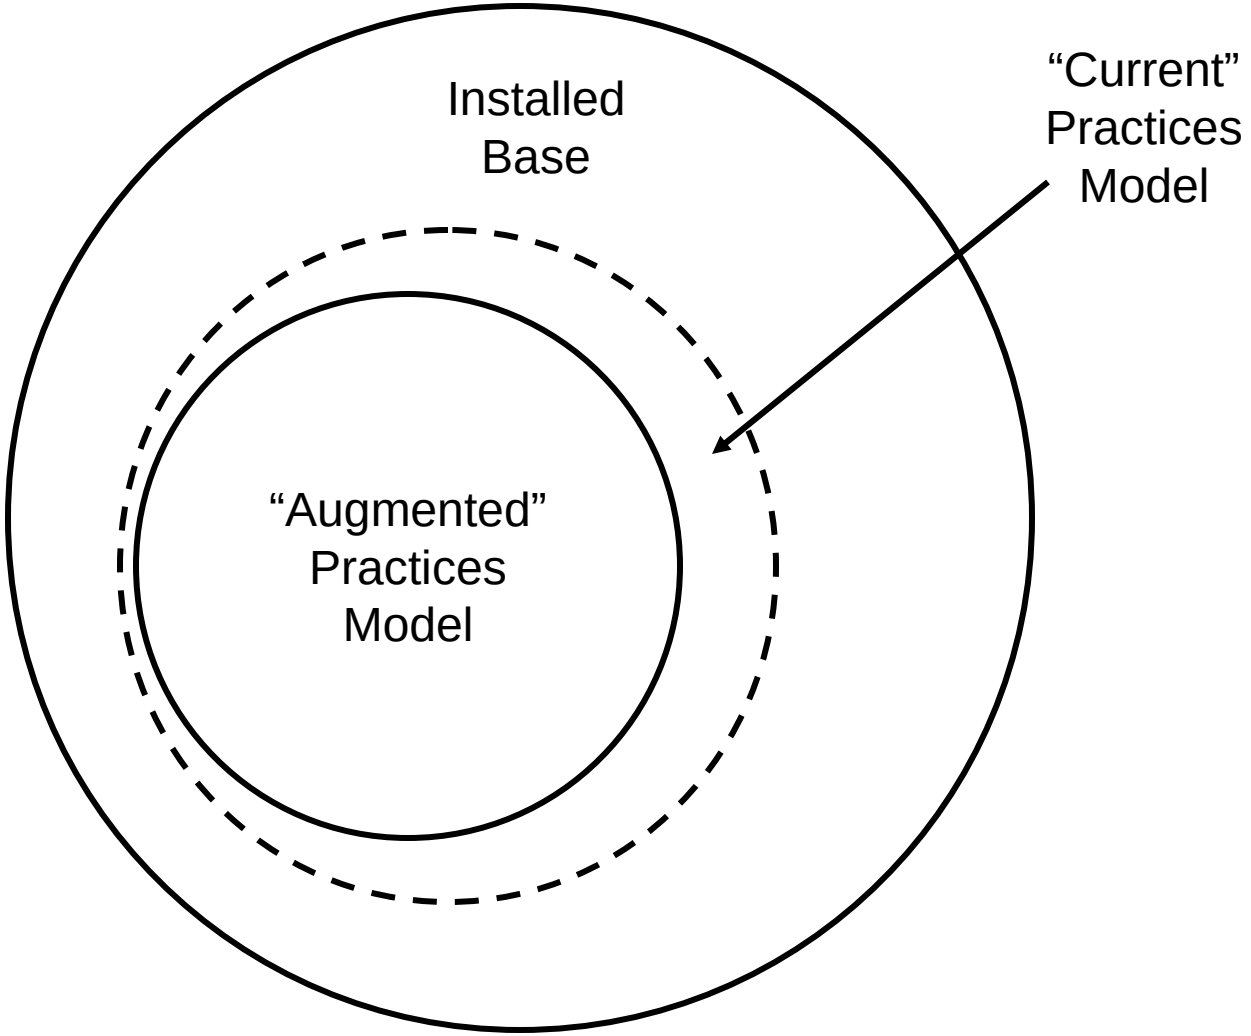

Frequency Range Considerations

- Proposed limits specify channel performance from 0.1MHz to 15GHz.
- For 1000BASE-KX and 10GBASE-KX4, it can be argued that much this frequency information is not relevant.
  - Example, 10GBASE-CX4 channels only specified to 2GHz
- Investigation into the required specification frequency range is warranted.
  - Including 10GBASE-KR…

Straw Poll Questions

- Should 1000BASE-KX support “current practices”?
- Should 10GBASE-KX4 support “current practices”?
- Should 10GBASE-KR support “current practices”?

Normative Methodologies in Practice

- Mask Set
  - Used to describe copper cabling (twisted pair, 10GBASE-CX4)
- Loss Budget
  - Used for optical links where medium is well-defined and penalties are predictable as a function of length.
- “Implicit” (XAUI)
  - Only transmitter and receiver are defined.
    - Compliance channel defined as transmitter test tool, but does not constitute a channel specification.
  - Permissible channels are those that interoperate with the transmitter and receiver.

Proposed Methodologies (1/2)

- Mask Set
  - Minimum data processing
  - High comfort level
  - Limited ability to explore design trade-offs
    - However can be augmented to allow trade-offs (example, ACR)
  - May leave margin on the table, or admit channels that are difficult to handle
    - Again, it is possible that set may be modified or augmented to close “loopholes”
- Pulse Response
  - Medium data processing
  - Relatively “untested” technique
  - Considers channel magnitude and phase information and can provide a more complete account of stub effects

Proposed Methodologies (2/2)

- Statistical Eye
  - New technique still under evaluation
  - Maximum data processing
  - Complicated methodology; difficult to rigorously document
  - Considers channel magnitude and phase information and can provide a more complete account of stub effects
    - Also included channel crosstalk
  - Takes into account the ability of the transmitter and receiver to compensate for channel distortion
  - Many design trade-offs available.
    - In principle, margin due to stacking of corner cases minimized
  - Results far removed from input data
    - No “intuition” regarding whether the correct result was obtained or not

Process Overview

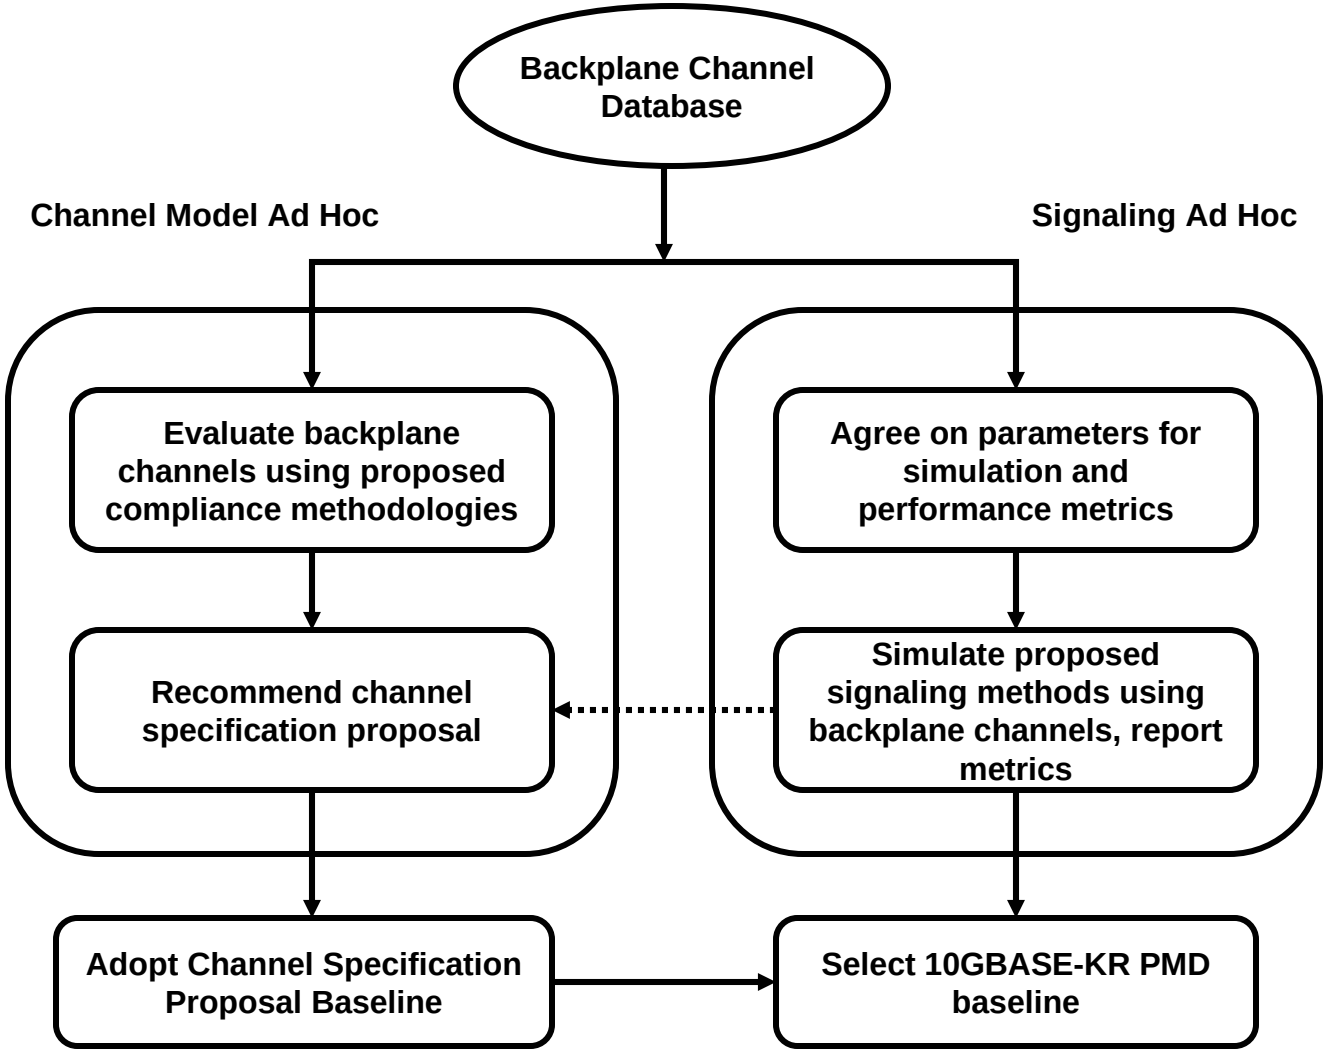

Soliciting Proposals

- Completion of the “augmented practices” model.
  - Close on return loss methodology and specifications.
  - Close on crosstalk specifications.
- Definition of the “current practices” model.
  - Justification for the support of stub effects and proposed methodology for specification.
  - Return loss, crosstalk modified if necessary.
- Required frequency range for specification.
- Normative specification method.
  - Make mask set normative?
  - …or develop one of the alternative approaches?
- Requests that concepts be presented at November 5 call for preliminary review and feedback.

Thank you!