Analog Receiver Performance Comparison of PAM-4 and PAM-2

Adee Ran
Intel

IEEE P802.3bj 100 Gb/s Backplane and Copper Cable Task Force
January 2012

Contributors and supporters

- Vasudevan Parthasarathy, Broadcom
- Kent Lusted, Intel
- Hamid Rategh, Inphi
- Richard Mellitz, Intel
- Howard Frazier, Broadcom
- Matt Brown, Applied Micro
- Sudeep Bhoja, Broadcom
- David Chalupsky, Intel
- William Bliss, Broadcom
- Dariush Dabiri, Applied Micro
- Beth Kochuparambil, Cisco

Preface

- There is an effort by the NRZ supporters to persuade us that "NRZ is better for backplanes than PAM-N"
- To support their claim, they are citing previous work and providing supposed theoretical arguments
- Choosing what to stress and what to neglect is important
- Simulation results for PAM-2 (AKA NRZ) and PAM-4 with realistic assumptions show that for several important scenarios, PAM-4 is comparable or better than PAM2

Claims heard against PAM-4

And responses

- "PAM-4 is more sensitive to timing errors"
  - Actually, PAM-4 has twice the timing budget of PAM-2 to start with, so it can tolerate jitter better – as explained in bliss_01_1111 (and demonstrated here).
  - Jitter does not scale down with clock frequency (so % UI is a wrong measure). In practice, jitter might be harder to control at higher clock frequencies.
- "PAM-4 is more complex, power hungry, won't scale"
  - Feasibility demonstrated by other presentations.
  - We are looking for the best solution for the current project. Comparison should be done (and would be done by the market) on real products, if the standard defines both methods.
  - Following common wisdom, x2 faster frequency requirements usually more than doubles power consumption ($P \propto fV^2$, faster switching requires higher V).

Claims heard for “PAM-2 only”

And responses

- “DFE overcomes the channel attenuation without amplifying noise, so there is enough bandwidth for PAM-2”
  - DFE with multiple large taps causes error propagation, increases the dynamic range and sensitivity required from analog front end, and reduces jitter tolerance.
  - PAM-2 typically requires x2 DFE taps and higher coefficients than PAM-4.
  - Good enough SNR is still required over the required bandwidth; ICR at top of spectrum might become a practical bottleneck.
- “Modern backplanes with low loss materials can work with PAM-2 with large margins”
  - Volume produced boards have non-negligible deviations from simulated ones.[1]
  - Edge effects (BGA vias, package, termination) not taken into account.
  - Crosstalk and reflections are harder to control at high frequencies.
  - Requirements for PAM-2 may rule out cost-sensitive products.[2]

1) See kochuparambil_01_0112, beers_01_0112
2) See goergen_01_0112

Additional Claims

And responses

- "PAM-4 has poor results" or "doesn't work"
  - Flaws in previous comparisons were pointed out:
    - Receiver architectures should be optimized differently. PAM-4 needs "high enough SNR" whereas PAM-2 needs "high enough timing margin".
    - Parameters tuned for PAM-2 obviously need to be changed when using PAM-4 (e.g. bandwidth filter, CTLE).
    - Doubling the jitter for PAM-4 (by using same % UI values as in PAM-2) is not technically justified.
  - In our simulations it does work. Can we close the gap?
- "Other communication standards already use NRZ at these speeds, doing something else would impact time to market"
  - 802.3bj objectives do not include "compatibility with other standards". Other standards have different requirements, and sometimes niche markets.
  - Assumed market timing for 100 Gb/s over backplane enables development of new products. A two-PHY approach enables using NRZ solutions if/when they are available.

Simulation environment

- Transmitter:
  - 3-tap adaptive FFE, online adaptation by RX
  - Pulse shape: 2nd-order Butterworth @ 0.8 Baud rate
  - Pattern: PRBS31 (aggressors use similar with different seed)
    - For PAM4 – packed into PAM4 symbols and 1/(1+D) precoded
  - Jitter: suggested TX spec values – RJ RMS=0.37 ps, DJ PTP=3.7 ps (sinusoidal @ 200 MHz), DCD PTP=1%
    - Simulating TX jitter makes "jitter amplification" effects appear, and shows fixed-DFE effects
- Receiver:
  - Analog DFE, LMS adaptation
    - 16 taps for PAM4, 32 taps for PAM2
  - Jitter: Algorithmic (real dynamic CDR)
  - Input filter: 4th-order Butterworth @ 0.6 Baud rate
  - Input noise: -154 dBm/Hz AWGN at RX filter input
- Thru and crosstalk channels:
  - Imperfect termination – 350 fF S/E on TX and RX (creating re-reflection effects)
  - Package model added on RX side only
  - Same parameters (amplitude, equalization, and jitter) used for all TX

Simulation environment (cont.)

- ~2 Mbits simulated after initial adaptation
  - Enables BER measurement of 1e-6 with CL≈86%
- Metrics:
  - SNR (normalized vs. “inner” symbol level so PAM-2 and PAM-4 figures are comparable)
  - NEVO (see ran_01a_1111 + backup slide) measured online @ 1e5
    - “Eye height in %” – shows margin and dynamic range requirements
  - BER counter (error propagation can occur)
- Test cases:
  - IBM 1 meter backplane – Thru channel and 8 aggressors (Channel data, reference)
  - TEC short backplane (14” Megtron6, TinMan connectors) – Thru, 3 NEXT, 2 FEXT
  - TEC long backplane (42.8” Nelco 4000-6, Whisper connectors) – Thru, 3 NEXT, 4 FEXT (Channel data, Reference)

Simulated case #1:
IBM 1 m backplane (Thru channel and 8 aggressors)

- This channel was claimed to work well with PAM-2 without FEC, using a 15-tap DFE and a CTLE (patel_01_0911)
- ILD and RL are better than Clause 72 limits
- ICR better than Clause 72 but drops to ~10 dB and lower above 10 GHz (about 40% of the RX bandwidth)

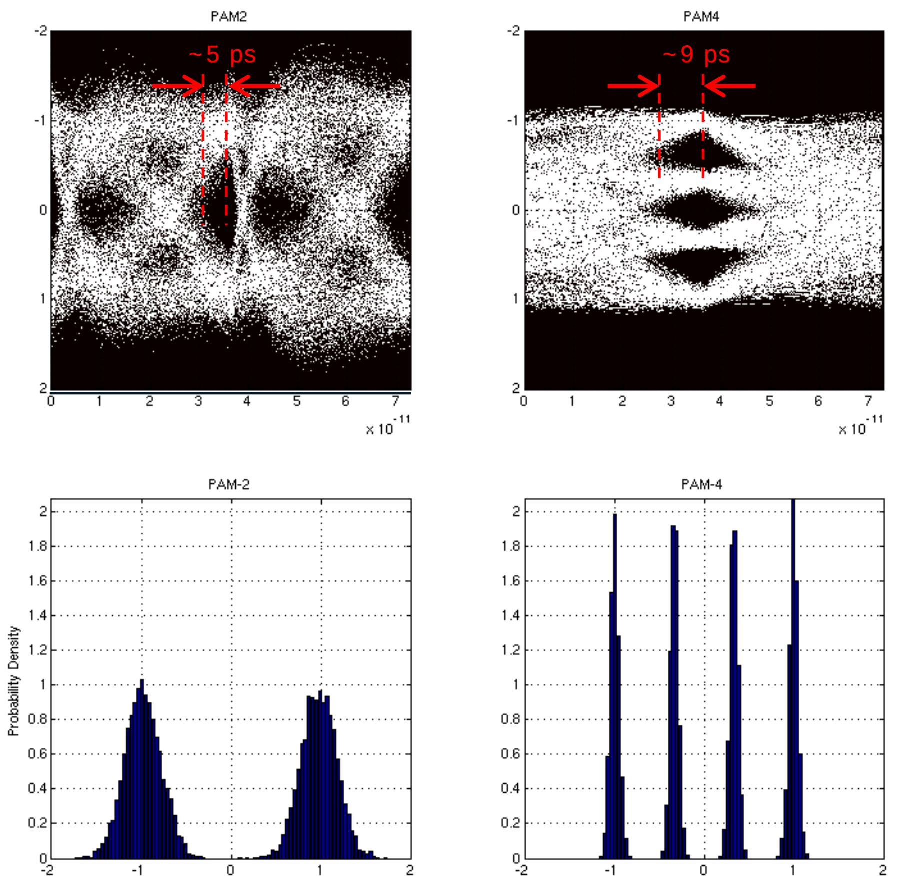

Sampled voltage distributions

Result	PAM-2	PAM-4
SNR [dB][1]	13.8	16.7
BER	3e-6 (6 errors observed)	<1e-6 (0 errors observed)
NEVO @ 10^5 UI	28%	53%

[1] PAM-4 SNR is vs. inner level or “per eye”, thus directly comparable to PAM-2 result.

Simulated case #2:

TEC short backplane – 14” Megtron6, TinMan connectors (Thru, 3 NEXT, 2 FEXT)

- Low IL; RL & ILD better than Clause 72 limits; ICR touches limit line
- Represents a short reach reflection- and crosstalk-dominated channel
- “future legacy” case (most current backplanes have some short channels)

Result	PAM-2	PAM-4
SNR [dB][1]	13	14.1
BER	2e-5 (44 errors observed)	<1e-6 (0 errors observed)
NEVO @ 10^5 UI	20%	33%

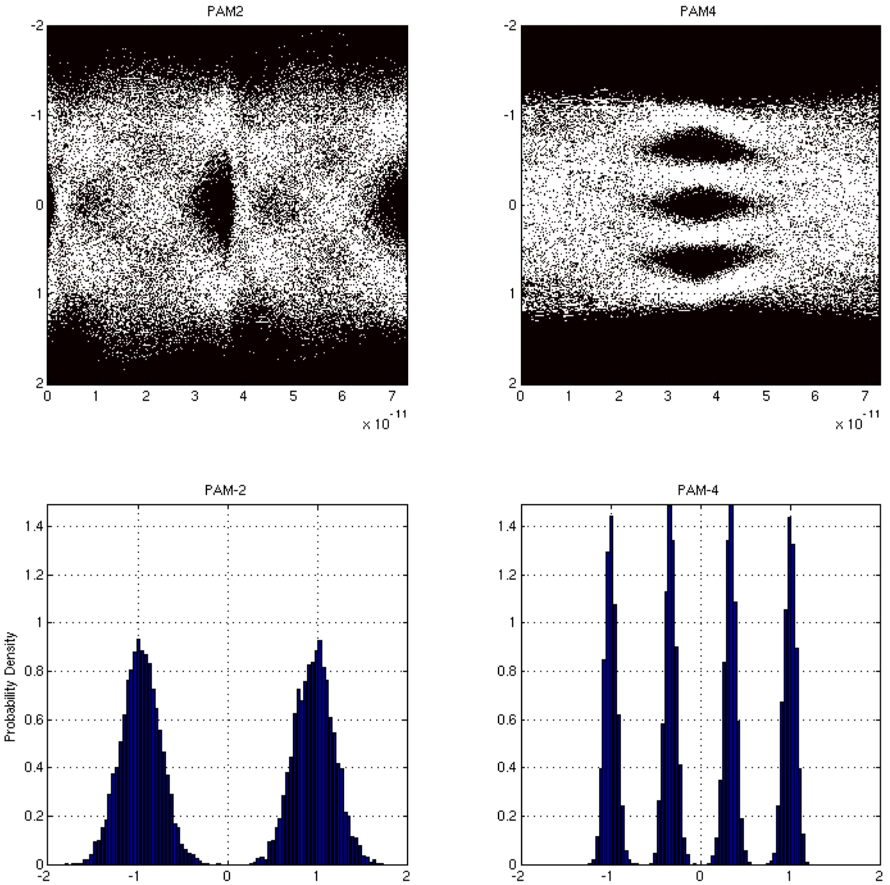

Sampled voltage distributions

[1] PAM-4 SNR is vs. inner level or “per eye”, thus directly comparable to PAM-2 result.

Simulated case #3:

TEC long backplane – 42.8" Nelco 4000-6, Whisper connectors (Thru, 3 NEXT, 4 FEXT)

- This channel was claimed to be marginal at PAM-2 with a 5-tap DFE and a CTLE (li_01_0511)
- Represents a "long reach" case with Tier-1 material

Result	PAM-2	PAM-4
SNR [dB][1]	9.4	13.8
BER	4e-3 (>5000 errors observed)	<1e-6 (0 errors observed)
NEVO @ 10^5 UI	1%	33%

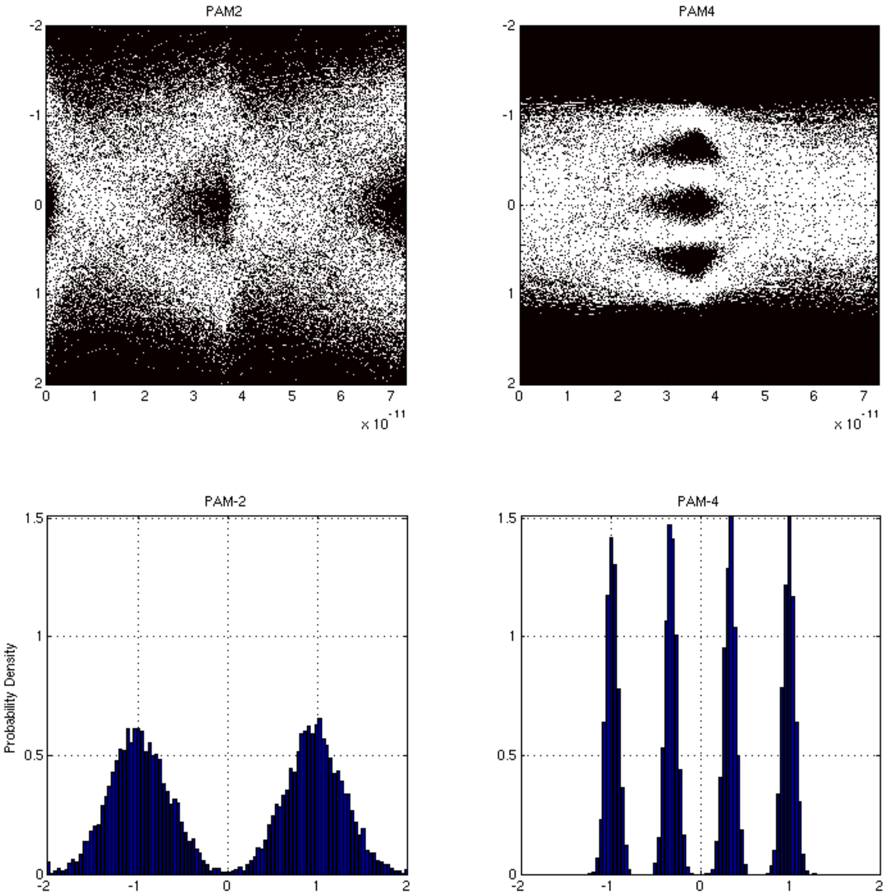

Sampled voltage distributions

[1] PAM-4 SNR is vs. inner level or "per eye", thus directly comparable to PAM-2 result.

Notes on simulation

- Time domain simulation
  - Is not as elegant as mathematical analysis
  - Enables modeling real life effects – so real troubles not shown by some analysis methods may be revealed (vs. ignorance ➔ optimism)
  - Simulating a few million bits is feasible.
- Assuming strong FEC (as currently proposed) requires raw BER of ~1e-6 at decoder input – which can be measured in time domain simulation.
- For this BER target, Gaussian assumptions are justifiable:
  - Total noise crest factor is high enough
  - Extreme events are rare and likely to be caught by FEC
  - SNR using noise power is meaningful

Observations on PAM-2

- The “long and strong“ DFE required by PAM-2 can increase noise margin (eye height) with optimal sampling, but has little effect on timing margin (eye width), since the coefficients are fixed. The theoretical SNR advantage of PAM-2 vanishes with modest TX jitter.
  - For the 3 channels examined, simulated jitter causes SNR degradation of 1-2 dB.
- The DFE cancellation leaves very small signal level that is susceptible to crosstalk and environmental noises (modeled as AWGN).
  - For the 3 channels simulated, the simulated AWGN and crosstalk each cause SNR degradation of 1-2 dB. (For the short channel, AWGN has little impact)
- Mismatched termination and package causes reflections and low-pass filtering which isn’t negligible at PAM-2 bandwidth.
  - Package and termination cause SNR degradation of 4-5 dB for the long channels and ~2 dB for the short channel.
  - Should we expect pad capacitance and package losses to scale down to 40% of today’s technology?

Conclusions

- Channel and endpoint effects (ILD, reflections, crosstalk, jitter, and environmental noise) which were considered negligible at 10G become intolerable at 25G with PAM-2.
- PAM-4 is less sensitive to timing and high-frequency noises and allows less "optimal" channels to work, or more budget to play with.

Thank You

Backup

New performance metric: Normalized Eye Vertical Opening

- Motivation:
  - With non-Gaussian noise distribution, SNR is not a useful metric (classical analysis yields over-pessimistic results)
  - Eye width is an ill-defined concept for a DFE receiver; but eye height is still useful
  - Eye height (EH) in voltage units is misleading, since the receiver typically amplifies the signal (and the noise); the receiver gain (or “target” signal level) is a free parameter, making it hard to compare EH results
- Dividing the EH (at the desired probability) by the ideal height (signal separation) yields the Normalized Eye Vertical Opening (NEVO), which ranges from 0 to 1
  - An intuitive figure of merit, enabling easy comparison of results
  - Useful for determining allowed implementation penalties: e.g. if a receiver has a signal separation of 400 mV and the NEVO is 25%, then the total noise margin is 400*25%=100 mV

Ref: ran_01_1111

Illustration of NEVO

This is a simulated eye pattern for PAM-4. Sampling instant is at x=0.5 (EH is maximized there by the DFE).

The signal level separation (ideal eye height) is 0.4.

The EH (for the measured period – prob ≈10^{-5}) is ~0.27.

NEVO (@10^{-5}) is 0.27/0.4=68%.

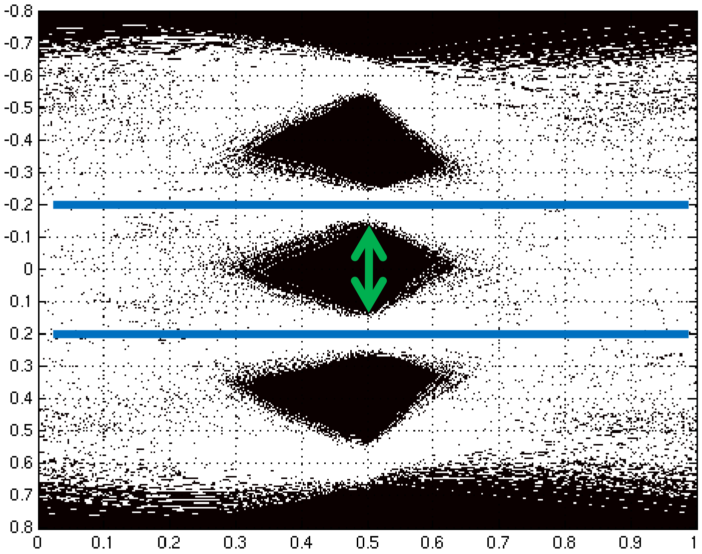

Ref: ran_01_1111

Short Backplane Channel

Data provided by Megha Shanbhag and Nathan Tracy, TE Connectivity

SIMULATION SET-UP

DAUGHTER CARD x 2
- Board material = Megtron 6* or Improved-FR4
- Trace length = 5"
- Trace geometry = stripline
- Trace width = 6 mils
- Trace spacing = 9 mils

CONNECTOR x 2
- STRADA Whisper or Z-PACK TinMan
- Vertical Header
- Right Angled Receptacle

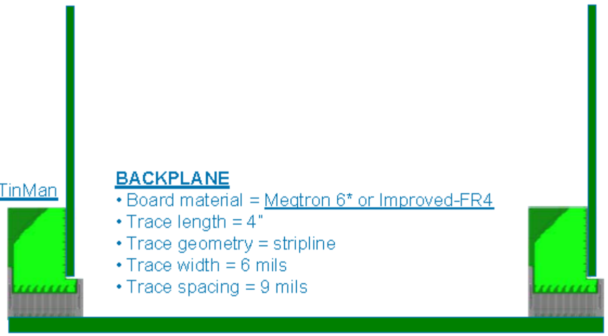

BACKPLANE
- Board material = Megtron 6* or Improved-FR4
- Trace length = 4"
- Trace geometry = stripline
- Trace width = 6 mils
- Trace spacing = 9 mils

DAUGHTER CARD FOOTPRINT x 2
- PCB Thickness = 150 mils
- 16 layers
- Bottom Signal Layer route
- ~140mil via with **10mil stub** for victim
- All noise aggressor vias go THRU board

BACKPLANE FOOTPRINT x 2
- PCB Thickness = 250 mils
- 26 layers
- Bottom Signal Layer route
- ~240mil via with **10mil stub** for victim
- All noise aggressor vias go THRU board

- Board Material Assumptions
  - Megtron 6* HVLP: Er=3.48, TanD=0.0062 @ 15GHz [Svensson/Djordjevic Loss Model]
  - Improved FR4, std: Er=3.40, TanD=0.0140 @ 10GHz [Svensson/Djordjevic Loss Model]
  - Surface roughness is assumed to be built in to effective TanD, No additional surface roughness added

- All connector and footprint models were generated using Ansoft HFSS software
- All traces were modeled as coupled striplines using Agilent ADS software
- The above models were concatenated within ADS to get the full channel model

- ~0.43meters = 16.8" = 2*(5" daughtercard traces) + 2*(0.15" daughtercard footprint) + 2*(~1" connector) + 2*(0.25" backplane footprint) + 4" backplane traces

Short Backplane Channel – Crosstalk

PAIR CONFIGURATION

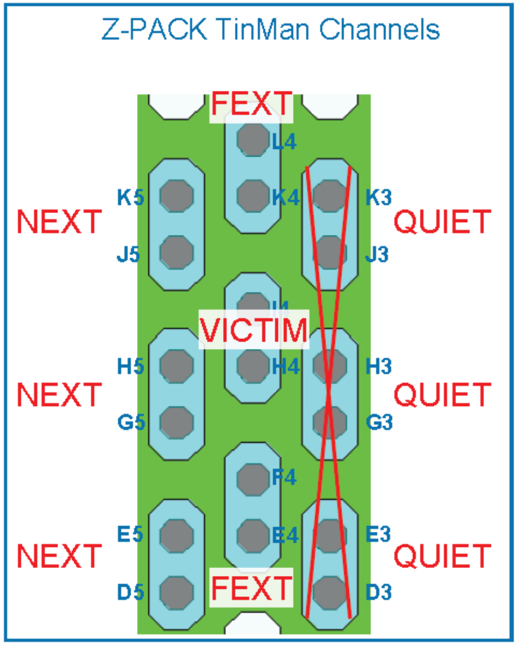

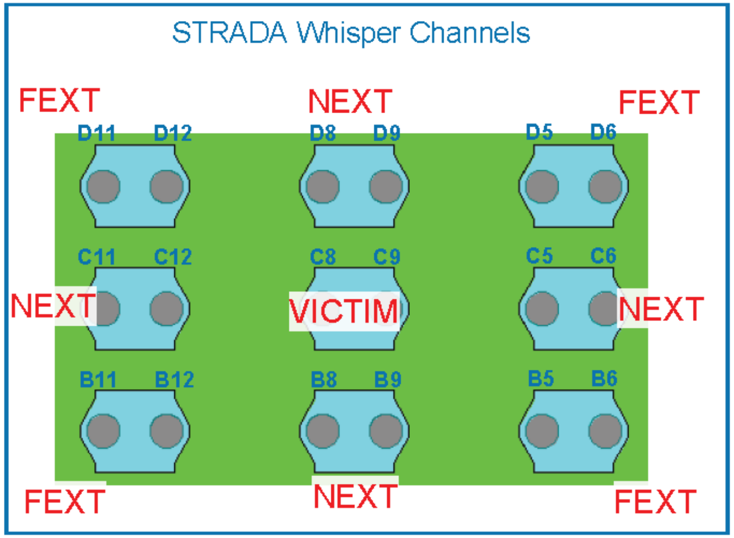

Short Backplane Channel – Spectral data