Peak-to-peak voltage definitions

Associated comments: 416, 385, 386, 146, 570, 524, 563, 523

Adee Ran, Cisco

Introduction

- All transmitter/output specifications include differential peak-to-peak voltage (max)
  - In 802.3ck we added specifications for common-mode too; this will be addressed later
- In previous projects the maximum differential values were 1.2 V for KR/CR and 0.9 V (or lower) for C2M
  - Why the different values?
- The definitions of this parameter are outdated; it does not mean what we expect it to mean…
  - The maximum voltage that can occur is higher
  - This can mislead readers, especially component designers, if they interpret it literally
  - The problem is exacerbated by loss to the measurement point
- It is proposed to change the definition and turn it into what readers expect it to be.

10GBASE-KR

Table 72–6—Transmitter characteristics for 10GBASE-KR

Parameter	Subclause reference	Value	Units
Signaling speed	72.7.1.3	10.3125 ± 100 ppm	GBd
Differential peak-to-peak output voltage (max.)	72.7.1.4	1200	mV

72.7.1.4 Output amplitude

The differential output voltage is constrained via the transmitter output waveform requirements specified in 72.7.1.10. For a 1010 pattern, the peak-to-peak differential output voltage shall be less than 1200 mV, regardless of equalization setting. The transmitter output voltage shall be less than 30 mV peak-to-peak when disabled. The differential output voltage test pattern shall consist of no fewer than eight symbols of alternating polarity.

NOTE 2—See Figure 72–8 for an illustration of the definition of differential peak-to-peak output voltage.

Figure 72–8—Transmitter differential peak-to-peak output voltage definition

- The term “peak-to-peak” is only “defined” by the Figure (which is not a very useful definition)
- There is no explicit definition of a test pattern for measurement of this parameter, but the definition says “no fewer than eight symbols of alternating polarity”
  - It isn’t clear that the maximum applies to any pattern; with no equalization, lower frequencies might reach higher voltages.
  - However, with measurement at TP0 after low loss from the source, it is likely that the measured voltage is close to the real peak-to-peak.
- This spec can be interpreted as “the maximum launch voltage is 1.2 Vdpp” (on a 100 Ohm load)
  - Component designers (Tx and Rx) understand it this way (ignoring possible loss in the package or test fixture).
- **The differential peak-to-peak limits the launch voltage; it is not the signal seen by the receiver!**
  - The receiver gets a signal attenuated by the channel and by the Tx equalization; its range is dependent on equalization setting.

40GBASE-CR4

- Table 85–5 points to the definition in 72.7.1.4
  - But the loss to the measurement point (TP0-TP2) is higher – **up to 6.5 dB at Nyquist**
  - A "1010" pattern will be attenuated by the TP0-TP2 channel
  - With "alternating polarity" pattern, 1200 mV at TP0 can become ~570* mV at TP2!
- There is an additional specification of "Transmitter DC amplitude" (equivalent to v_f) with 0.34 min, 0.6 max
  - Measured from linear fit pulse of PRBS9 with a response length Np=8 UI
  - This limits the launch voltage more than the "peak-to-peak" definition in 72.7.1.4; if a device is compliant with this specification, then 1200 mV "peak-to-peak" can't be reached
  - Note that this parameter is based on linear fit, and does not cover non-linear effects or loss at frequencies lower than 20 MHz (period of PRBS9)
- The component requirements are understood to be the same as in 40GBASE-KR4 and 10GBASE-KR, which are specified at TP0

* Rough approximation based on Nyquist loss

100GBASE-CR4 and 100GBASE-KR4

- Table 92–6 (Transmitter characteristics at TP2) and Table 93–4 (Summary of transmitter characteristics at TP0a) both have parameters and values that look similar to those of 40GBASE-CR4
- However, the definitions of “differential peak-to-peak voltage” in 92.8.3.1 and 93.8.1.3 are different from the one in 72.7.1.4 : **they do not refer to the 1010 sequence anymore**
  - Measurement is specified with PRBS9. It is implied that the peak-to-peak of the whole pattern is measured.

92.8.3.1 Signal levels

<...>

The peak-to-peak differential output voltage shall be less than or equal to 1200 mV regardless of the transmit equalizer setting. The peak-to-peak differential output voltage shall be less than or equal to 35 mV while the transmitter is disabled (refer to 92.7.6 and 92.7.7).

<...>

Differential and common-mode signal levels are measured with a PRBS9 test pattern.

93.8.1.3 Signal levels

<...>

The peak-to-peak differential output voltage shall be less than or equal to 1200 mV regardless of the transmit equalizer setting. The peak-to-peak differential output voltage shall be less than or equal to 30 mV while the transmitter is disabled (refer to 93.7.6 and 93.7.7).

<...>

Unless otherwise noted, differential and common-mode signal levels are measured with a PRBS9 test pattern.

100GBASE-CR4 and 100GBASE-KR4 (cont.)

- For these PMDs, the loss to the measurement point is not negligible:
  - For 100BASE-KR4, the nominal TP0-TP0a Nyquist loss is 1.5 dB. With a maximum package loss of ~3.5 dB we can assume ~5 dB at Nyquist.
  - For 100GBASE-CR4, TP0-TP2 Nyquist loss is assumed to be up to ~10 dB, or ~13.5 dB including a maximum package.
- Even with equalization off, the peak-to-peak measured with PRBS9 is lower than the launch voltage, because of the limited low-frequency content.
  - The attenuation depends on the channel (and package), but if a device is compliant with the additional v_f requirement, then 1200 mV will never be reached.
- With Tx equalization enabled, the low frequencies are attenuated by the FFE, while the high frequencies are attenuated by the channel
  - In the "INITIALIZE" setting, the nominal Rpst is 2.57 (low-frequency attenuation of ~8.2 dB)…
  - The measured PtP is much lower than the launch voltage: 1200 mV ➔ 467* mV at TP2

* Rough approximation based on Nyquist loss

50G per lane PMDs (clauses 136 and 137)

- Clause 136 refers to 93.8.1.3 (the KR spec at TP0a) for the definition of Differential pk-pk voltage…
  - but with a footnote: “Measurement uses the method described in 93.8.1.3 with the exception that the **PRBS13Q test pattern is used**”
- Clause 137 refers to Table 120D–1, which also refers to the same definition and has the same footnote.
- The loss to the measurement points is similar to that of the 25G per lane PMDs.
- The period of PRBS13Q is 8191 UI, creating a minimum frequency of ~3 MHz
  - But its longest runs of the outer levels is shorter than to PRBS9 (7 vs. 9)
  - The measured peak-to-peak is still lower than the launch voltage.

100G per lane PMDs (clauses 162 and 163)

- Both Table 162–11 (TP2) and Table 163–5 (TP0v) refer to 93.8.1.3 (the KR spec at TP0a) for the definition of Differential pk-pk voltage.
  - Both specify using PRBS13Q in a footnote.
- The loss to the measurement points is higher than what was assumed for 50G and 25G: TP0-TP2 budget is 11 dB and with the maximum package allowance we get ~15 dB.
  - Clause 163 has TP0v with specified maximum loss of 6 dB – so the effect is less severe.
- The period of PRBS13Q is 8191 UI, creating a minimum frequency of ~3 MHz
  - But its longest runs of the outer levels is shorter than to PRBS9 (7 vs. 9)
  - The measured peak-to-peak is lower than the launch voltage
- Even with equalization off, the measured peak-to-peak with PRBS13Q can be significantly lower from what it can reach with mission data (2^*v_f).
  - **Real life example**: on a CR port with die-to-TP2 loss of just ~9 dB, v_f is measured as 0.48 V, while v_{di} is 0.87 V instead of the expected 0.96 V (10% reduction!)

C2M host output specifications (prior to P802.3dj)

- The C2M host output has a similar specification, Differential peak-to-peak output voltage (max)
  - However, there is no specification of variable Tx equalization; the signal is only specified in whatever equalization setting the host has
  - As a result, the meaning of “peak-to-peak” is quite different
  - It does not match the launch voltage
  - With unknown equalization, it does not indicate the peak-to-peak that can occur with mission data
    - If the Tx equalizer has strong low frequency attenuation, then mission data may have the same peak to peak; but that is not guaranteed
- For C2M, the host loss can be higher than in CR, and with high loss, equalization is practically required (though not directly specified)
- This was addressed in comments against 802.3ck, see ran_3ck_04b_0721
  - To mitigate this concern, v_f was added to the output specifications in 120G

Excerpt from [ran_3ck_04b_0721](ran_3ck_04b_0721)

Simulation experiment

- Simulated channels
  - Channel3 from lim_3ck_01_0319_c2m (-15 dB) – near-maximum C2M channel
  - Same with 31 mm Tx package (-20.7 dB)
  - OSFP Mated Test Fixture from kocsis_3ck_02_0719_MTFosfp (6.5 dB) – representing minimal C2M channel
- Patterns
  - PRBS13Q
  - SSPRQ
- Launch PtP is 1 V, No Tx equalization
  - FFE constraints can only decrease the peak-to-peak voltage
- "True PtP" is the asymptote of the step response, which includes DC loss

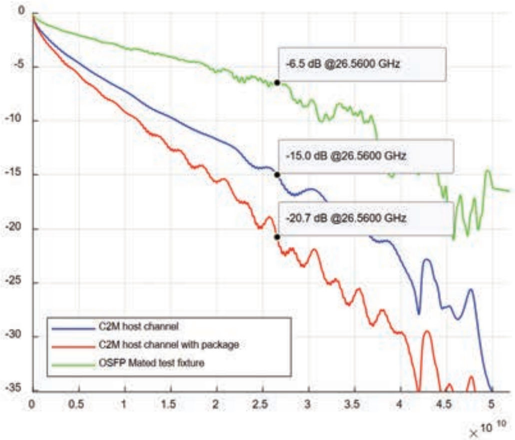

August 2021 P802.3ck 3

PRBS13Q

Channel	min/max with PRBS13Q [mV]	% of true PtP	True PtP if 900 mV is measured [mV]
C2M Host channel	-408, +407	84%	1042
C2M Host channel +COM 31 mm pkg	-386, +386	79%	1100
OSFP Mated Test Fixture	-455, +458	93%	957

Measurement with PRBS13Q is much lower than the launch PtP and is channel dependent.
With scrambled data the signal can reach the launch voltage.
The dynamic range that the receiver will need to handle can't be deduced from the measurement.

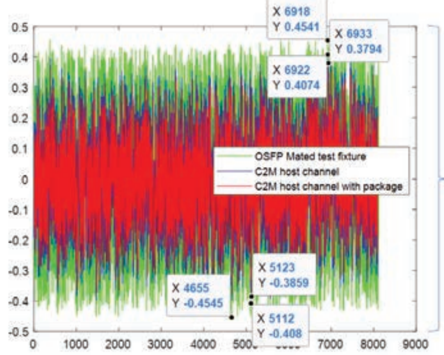

Tx launch PtP

August 2021 P802.3ck 4

SSPRQ

Channel	min/max with SSPRQ [mV]	% of true PtP	True PtP if 900 mV is measured [mV]
C2M Host channel	-436, +436	90%	974
C2M Host channel +COM 31 mm pkg	-424, +424	87%	1001
OSFP Mated Test Fixture	-471, +471	96%	928

Note: SSPRQ includes 14-UI runs of both "0" and "3" symbols, therefore its PtP is expected to be equal to that of PRBS31Q (but is much faster to simulate)

Measurement with PRBS31Q or SSPRQ is still lower than the real PtP that the signal may reach.
But the channel dependence is smaller.

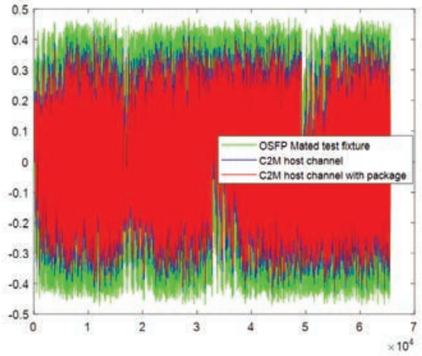

Tx launch PtP

August 2021 P802.3ck 5

Conclusion: measurement with a short test pattern is not representative of the "true" PtP (that can occur with real traffic data), and the "error" is channel-dependent

Why is that a problem?

- When voltages close to the launch voltage appear at the output of a lossy channel, they can exceed the dynamic range of a receiver and cause unexpected effects
- This can occur with patterns that are strongly unbalanced over periods of 100s or 1000s of UIs
  - These are relatively low probability events, but they occur too often to be ignored
  - This is a low frequency effect – when it occurs, it spans many symbols (and can cause correlated errors)
  - The result can be poor post-FEC performance while the average BER is very low
  - Some applications require virtually error-free operation
- Ideally, receivers should tolerate the maximum level that a transmitter can create
  - But it is may not be predictable from observation of typical patterns (even training patterns)
  - It needs to be specified

Returning to P802.3dj…

- All interfaces have specified adjustable Tx equalization, and the corresponding common measurement methodology
  - Specifically, the differential peak-to-peak specification in 176E should be met with any equalization setting – unlike earlier C2M annexes
  - The maximum is not (and should not be) the same as it was in 120G – **see backup slides**
- For C2M, the maximum TP0d-TP1a loss is much higher than we ever had – likely >30 dB
  - Differential peak-to-peak in a measurement with PRBS13Q, even with Tx equalization off, can be much lower than the “true” peak-to-peak
  - If we keep this pattern, the specification will be quite meaningless
- v_f can be thought of as mitigating the problem
  - However, it is based on a linear fit and may not capture all effects of all Tx implementations
  - Also, it may be less intuitive to readers who are used to a peak-to-peak specification
- If we have a differential peak-to-peak specification, it should hold for any valid output, not just PRBS13Q

Related topic – AC common-mode

- AC common-mode is currently specified with a probability of 1e-4 (PMDs) or 1e-5 (AUIs)
- As described in comments 385 and 386, excessive values of common-mode noise can also create errors
  - These errors can be correlated; a probability of 1e-5 with correlated errors is too high
  - We should reduce the probability to something that can be tolerated even if it creates correlated errors.
- The definition of peak-to-peak should be the same for differential and common mode.
  - The probability may be kept as a parameter.

How is a peak-to-peak specification interpreted?

- For a designer of a transmitter, it is quite clear, and can be verified by design
- For a designer of a receiver, it can be used as the dynamic range required at the Rx input (analog circuit and/or ADC)
- For measuring a DUT
  - With rich enough data, the output will have a distribution with possibly long tails; we should define “peak” with respect to some probability (as in jitter).
  - Peak values are measurable from arbitrary data patterns with either real-time scopes or equivalent-time scopes, assuming asynchronous sampling is allowed.
  - Faster measurements and simulations can be made with specific test patterns (such as PRBS31Q or a square wave with long enough period). This may be addressed.

Proposal:

Modify 176E.6.1 based on the text below, with editorial license. Refer to this subclause for differential and common-mode specifications in clauses 178 and 179, and annexes 176D and 176E. Delete the definition common-mode AC in 179.9.4.4.

176E.6.1 Maximum voltages

Differential and common-mode signals are defined in 93.8.1.3.

Peak-to-peak output voltages are defined to a probability P, with respect to the distribution of the output voltage V_{out} sampled at an effective rate of between two and three samples per UI. The sampling may be either synchronous or asynchronous.

A maximum output voltage is defined as the value V_{max} such that the probability that $V_{out} > V_{max}$ is $P/2$. A minimum output voltage is defined as the value V_{min} such that the probability that $V_{out} < V_{min}$ is $P/2$. A peak-to-peak output voltage is defined as $V_{max} – V_{min}$.

Specifications of peak-to-peak output voltage apply with any pattern that appears on the service interface, at any equalizer setting. Since the specification is a statistics-based estimate, measurement should be performed with a typical pattern such as scrambled idle, over a period long enough to enable calculation with the desired level of confidence.

NOTE 1—For short measurement purposes, PRBS31Q or a square wave with a period of at least 128 UI can be used to estimate V_{max} and V_{min}, but the values created by these patterns might be different than those of scrambled data.

Differential peak-to-peak output voltage is defined with P=10^{-9} for the differential output signal.

Full-band AC common-mode peak-to-peak voltage, VCM_{FB}, is defined with P=10^{-7} for the common-mode output signal.

Low-frequency AC common-mode peak-to-peak voltage, VCM_{LF}, is defined with P=10^{-7} for the common-mode output signal filtered by a low-pass filter as defined by Equation (93A–20) with f_r set to 100 MHz.

NOTE 2—The common-mode noise measurement may be sensitive to mismatches between the single-ended paths in the test fixture and the test setup. Careful design and calibration of the test system are recommended.

NOTE 3—The common-mode noise measurement should take into consideration frequencies down to the specified AC-coupling frequency.

Backup

Addressing comments 146, 570, 524, 563, 523

Are the maximum values too high?

- As indicated earlier, **The differential peak-to-peak limits the launch voltage; it is not the signal at the receiver**
- The specifications in previous AUIs do not represent the launch voltage; the value had a different meaning
  - Measurement with PRBS13Q with Tx equalization yields a value that will occur often at the receiver input; but it is not the real maximum voltage
- The peak-to-peak value can only be measured with equalization off (and likely with atypical patterns)
  - It is not the expected operation of the transmitter, but it ensures that higher voltage will not occur
- Even in AUIs (with lower loss), the receiver input is much lower than the maximum PtP
- The receiver can use the ILT protocol to attenuate/equalize the transmitter's signal if necessary
- Higher launch voltages can be useful…
  - To provide performance margin (beyond the Ethernet error ratio specifications)
  - To increase cable or PCB reach (with adequate equalization)
  - Increasing the signal may be easier than reducing the noise
- Reaching the maximum limit is not required!
  - Many existing products do not
  - If a product is sensitive to its own NEXT, then having lower than the maximum allowed vf may be beneficial
- Reducing the maximum vf (and differential peak-to-peak) would unnecessarily limit long-reach applications
  - It may be more justified for AUIs than for PMDs
  - For PMDs, we should consider *increasing the minimum* vf instead

That’s all

Questions?