DesignCon 2019

TDECQ for PAM4 Optical Transmitters: Does it really work?

Greg D. Le Cheminant, Keysight Technologies

Abstract

The use of PAM4 and forward error correction led to significant changes in the test methods historically used to characterize optical transmitters used in digital communications systems. TDECQ (transmitter dispersion and eye closure quaternary) is the primary example of this change. Does the TDECQ measurement really provide the results it was intended to yield, specifically the power penalty metric needed to predict how well a transmitter will operate in a real system? This paper will try to document when it does and when it does not.

Author Biography

Greg D. Le Cheminant is a Measurement Applications Specialist for digital communications analysis products in the Internet Infrastructure Solutions Group. He is responsible for development of new measurement applications for the group’s optical communications test products. He represents Keysight on several industry standards committees. Greg's experience at Keysight/Agilent/Hewlett-Packard began in 1985 with five years in manufacturing engineering, and the remainder in various technical marketing positions. He is a contributing author to four textbooks on high-speed digital communications and has written numerous technical articles on test related topics. He holds two patents. Greg earned BSEET and MSEE degrees from Brigham Young University.

Standards-based networks require interoperability to achieve lower costs

The fundamental requirement for the digital communications system is to transmit data with very few bits received in error. The overall system quality is then gauged by the bit-error-ratio (BER). While the design and installation of a telecommunications system often is done by a single company, datacommunications systems usually assume that the transmitter, channel and receiver may come from three different vendors, with perhaps a fourth installing and verifying the overall performance. Each of the three components is then considered in terms of its impact on the system BER. Each of the elements must interoperate with the worst case allowed version of the other two and still achieve the desired BER.

The benefits of the interoperability concept are that there are more players, resulting in competition, fostering higher levels of innovation and lower costs. The goal of interoperability dominates the process of setting component specifications. The overall cost of the link must be managed. Allowing poorly performing receivers by requiring an expensive high-performance transmitter may result in a system that does not meet cost objectives. Thus, the burden of achieving the BER must be carefully allocated to the three system components in a way that they can be economically produced. The common approach to specifying a receiver is to verify that it can achieve the required BER when receiving a signal that represents "just good enough". This is referred to as stressed receiver sensitivity or SRS. The channel, which is typically an optical fiber must meet specific attenuation and dispersion metrics. There are many parameters that can be used to describe the quality of a transmitter. Metrics like edgespeed, jitter, amplitude and eye opening are among the many that could be considered. The difficulty lies in how can measurable transmitter parameters be related to the key system level metric, BER. One approach to solving this problem is through a metric known as the transmitter dispersion penalty (TDP). The basic concept of TDP is that the transmitter being tested is compared to an ideal transmitter in a test bed where BER is measured. Rather than trying to assess the absolute quality of the transmitter, we determine the relative quality of that transmitter.

Looking back at legacy specifications

For non-return-to-zero (NRZ) systems, a minimum optical modulation amplitude (OMA, the difference between the logic 1 and logic 0 power levels) is required. But the quality of the modulation must also be considered, as a signal might have sufficient modulation power, but if there is also significant intersymbol interference (ISI) and noise, the signal might not be accurately received. The TDP measurement accounted for this, as the BER performance is determined relative to an ideal reference transmitter. The ideal reference transmitter is noise and ISI free, so a comparison to it provided the system specification designers a metric that accounts for eye closure.

As has been documented [2,5,6] the TDP concept worked well on paper but was difficult to use in practice. One problem is that it is not practical to physically realize an ideal reference transmitter. What represents 'ideal' and how do you know when it has been achieved? If you could build it, it would likely be expensive. Also, BER-based tests take a long time and the test equipment to perform a BER analysis is considered expensive. It has become common to forego TDP testing and instead do an eye-mask test. In the eye-mask test, polygons are placed above, within, and below the eye diagram representing regions where the transmitter waveform may not exist. The eye-mask test existed long before the TDP test. It coexists in many datacommunication standards along with TDP. Unlike TDP, there is no direct correlation between the eye-mask and system level performance. The shape and dimensions of the eye mask have become somewhat arbitrary. To trace the origins of the eye mask would perhaps require interviewing an engineer that retired from Bell Laboratories several decades ago. The original eye-mask dimensions have been trimmed and squeezed by standards developers, not so much with the system performance in mind but rather to ensure that most transmitters could pass the test.

But it works! Millions and millions of transmitters have been manufactured and deployed without the benefit of the key TDP test, and the systems they are a part of work as expected. This is likely due to the fact that the eye-mask is a decent indicator of NRZ eye quality, and in addition to OMA specs, prevent bad transmitters from being shipped. It is also likely that systems have some margin, with channels and receivers never being as bad as the specifications allow.

When PAM4 technology was accepted as a viable modulation scheme, transmitter specifications needed to be defined and the obvious method was to take what worked for NRZ systems and adapt them to PAM4. OMA and extinction ratio were easily modified. But a major obstacle to creating a PAM4 mask is not that there are three eyes, but rather the PAM4 eyes have very little eye opening. Waveform trajectories, as seen by sample points, frequently transect the eye opening. An example is shown in figure 1.

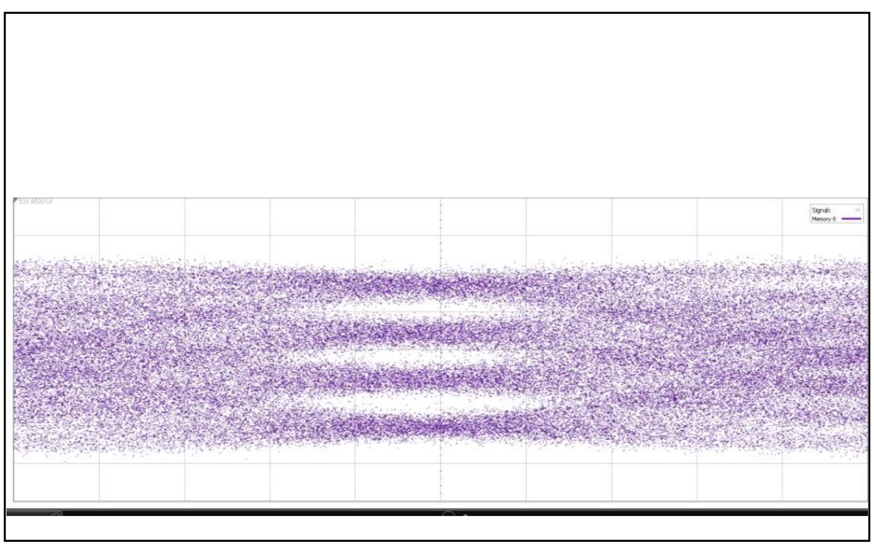

Figure 1: The PAM4 eye diagram shows little or no 'open' regions

It is possible that very small polygons could be placed within the eyes, but likely would not be useful. Another important issue is that the receivers used along with PAM4 transmitters will employ equalization. Thus, the waveform at the transmitter output is not what the receiver circuitry will actually see. Rather than try to figure out a meaningful eye-mask method, the choice was made to improve the TDP method. The result was TDECQ. The TDECQ method has been previously documented [3,4,5,6], but essentially operates are follows:

- The transmitter waveform is captured with an oscilloscope
- A virtual ideal transmitter is created having the same OMA
- Noise is mathematically added to the ideal transmitter while the symbol error rate (SER) is estimated mathematically
- Noise is also mathematically added to the transmitter waveform. The waveform is also passed through a virtual equalizer with the tap settings adjusted to reduce the SER and allow additional added noise

The two added noise values are compared as a ratio, expressed in dB. The result is a power penalty metric representing the additional power required from the transmitter compared to the ideal transmitter, to compensate for any impairment in the transmitter waveform.

TDECQ seems to be here to stay: What have we learned so far?

As TDECQ was designed in the IEEE 802.3bs 400G project, the test and measurement providers built early prototypes of the measurement. This allowed transmitter manufacturers to verify their transmitter performance. It also provided critical feedback on the TDECQ measurement itself. The TDECQ specification was 'tight' at roughly 2 dB. Many transceiver vendors reported that they could create a working link with their transmitters but that they failed the TDECQ transmitter specification. It was common to hear complaints that the specification was too difficult and should be relaxed.

The most important and effective question was "Does the virtual receiver used in the TDECQ method represent the receivers that transmitters will be paired with?". Recall that the goal of interoperability-based transmitter tests is to ensure that a transmitter will work with the worst case allowed receiver. The original TDECQ receiver model made no allowance for the receiver to adapt decision thresholds (no tolerance for modulation non-linearity) and no allowance for the receiver to optimize where in the symbol interval decisions were made. The majority opinion was that this receiver design placed too heavy of a burden on the transmitter and that most if not all receivers were more capable than the virtual receiver. The TDECQ virtual receiver was changed.

Where are we now?

The TDECQ virtual receiver now allows the decision thresholds to be adjusted up to 1 % of the signal OMA. This may seem small, but the impact was significant. More transmitters ccould achieve the TDECQ specification. The decision sampling time is now also allowed to be optimized anywhere in the PAM4 eye. This also increased transmitter yield against the TDECQ spec. The burden on the transmitters was eased. But the burden had to go somewhere. With the changes in the TDECQ reference receiver, complementary changes were required in how a receiver is tested. Essentially, the stressed-receiver sensitivity test needed to include a signal that emulated a lower performing transmitter.

As the TDECQ test evolved, the design and performance of the PAM4 transmitters steadily improved. More vendors reported better TDECQ results and links that met the standard specifications. While this was good news, it does not necessarily mean that the TDECQ method was working. Similar to NRZ testing, it is possible to have loose specifications and working links as long as no element of the system approaches a worst case allowed value. How can the TDECQ method be validated?

Recall the intent of TDECQ: Predict the extra power required from a transmitter to overcome any impairments in the transmitter waveform. For a 'good/compliant' receiver, if it were paired with two different transmitters A and B, A with a TDECQ of 2 dB, and B with a TDECQ of 3.2 dB, we would expect that transmitter B would need to produce 1.2 dB more power at the receiver than transmitter A for both to achieve the expected SER.

In the IEEE 802.3cd project, as the TDECQ receiver definition stabilized, confidence in the TDECQ method improved. Reports documenting TDECQ performance coupled with actual link performance showed important progress. One of the most thorough experiments was reported in the July 2018 IEEE 802.3 meeting (7). By permission, the results are show here.

Three transmitters with TDECQ values of 2.5, 3.4, and 4.0 dB are each paired with a receiver and three receiver sensitivity curves are generated, by attenuating the transmitter signal at the receiver and measuring the SER for several input power levels. The system is required to achieve an SER of 4.8 e-4. The power level to achieve this SER is recorded for each transmitter. If TDECQ is operating as intended, the change in receiver input power should have been predicted by an equal change in TDECQ. When receiver input power is plotted verses TDECQ, we see a 1:1 correspondence. TDECQ is achieving the intended result. (Note in this case the SECQ metric is used to indicate that a dispersive channel was not used in the link).

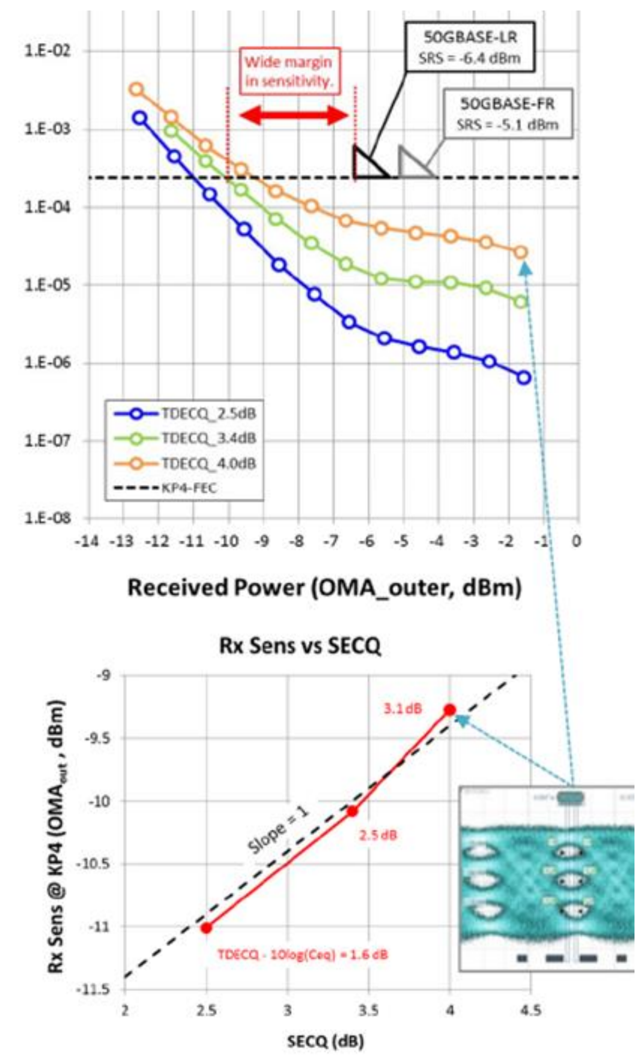

Figure 2: TDECQ relationship to receiver sensitivity

In the same report, another similar experiment was performed on a set of four transmitters with TDECQ values of 2.5, 3.1, 3.9, and 4.2 dB. Receiver sensitivity curves were generated and TDECQ was plotted against receiver input power. Again, the 1:1

correspondence was seen indicating that TDECQ predicted the relative transmitter penalties.

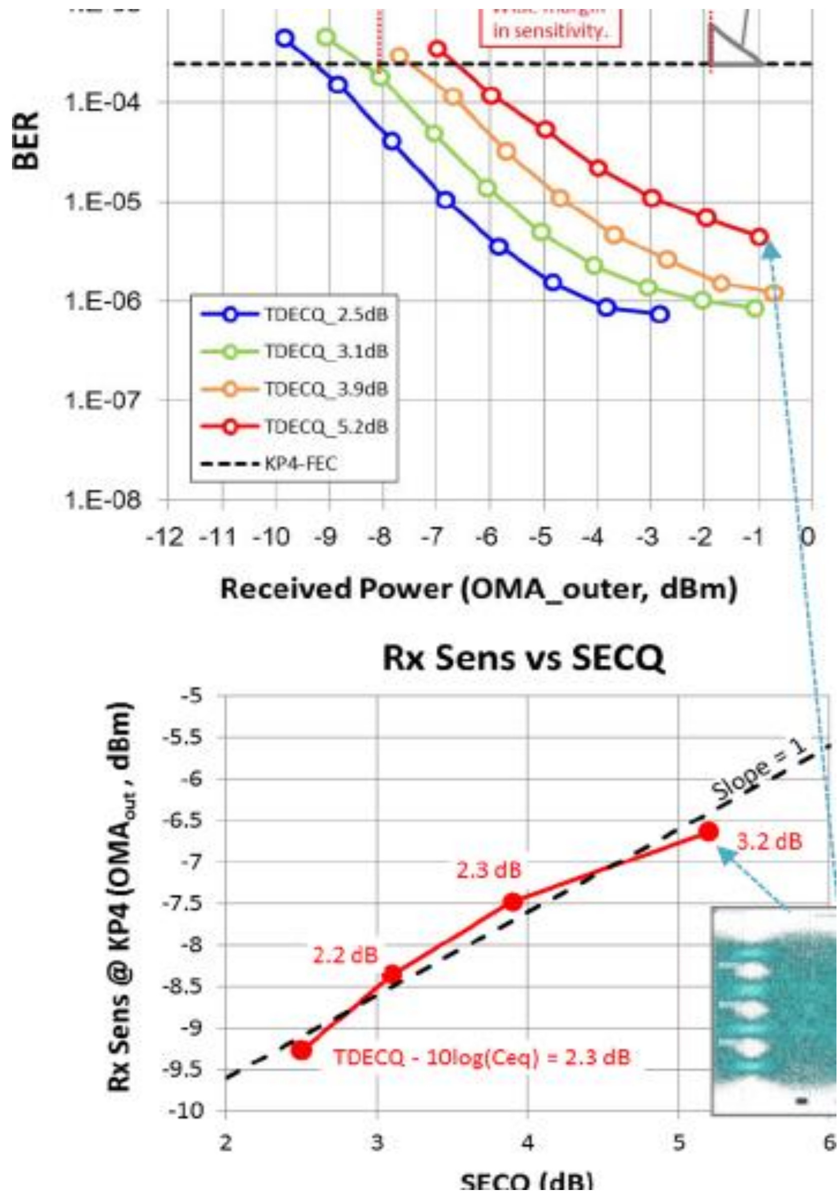

Figure 3: TDECQ relationship to receiver sensitivity #2

An important consideration for the receiver power versus TDECQ curve is that while the 1:1 slope of the curve is an excellent indicator of the *relative* performance of different transmitters, we do not know if the 2.5 dB TDECQ accurately represents the *absolute* power penalty that should be assigned to that transmitter. One way to confirm this is if an 'ideal' 0 dB TEDCQ transmitter could be constructed to peg the 0-dB location on the curve. This effectively takes us full circle back to the old TDP method, where a physical ideal transmitter was required. On the other hand, the plots of figure 2 and 3 could be used to predict the very best receiver sensitivity that could be achieved. In theory, the SER is dictated by the internal receiver noise impairments and occur when the input powers are less than ~ -13.5 dBm (figure 2) and ~ -10.5 dBm (figure 3) and a 0 dB TDECQ transmitter were available.

Another important consideration is the likelihood that a transmitter will be paired with a receiver that performs better or worse than the TDECQ virtual receiver. If the actual system receiver is better than the TDECQ receiver, the TDECQ penalty will likely be pessimistic. Similarly, if the system receiver is 'weaker' than the TDECQ receiver the TDECQ penalty will be an optimistic predictor of actual system performance. How well do we expect real receivers to perform? Very little information is available on receiver performance relative to the IEEE specifications. The test system is available but not used as often as the transmitter test systems. If interoperability becomes a problem as more systems are deployed, receiver test could increase in importance.

What should we expect going forward?

PAM4 transmitter are not yet being produced in volume. Very few interoperability tests have been performed. Are we in a better place than we are with the NRZ specification methodology, where the eye mask dominates but has loose correlation to actual system level performance? TDECQ has been shown to be a good predictor of system power penalty, but the number of experiments that validate this is small. As more devices are built and deployed, the use of TDECQ will likely evolve. Like mask testing, margins may be set to ensure that not only must transmitters meet the basic TDECQ specification, but it should exceed (be lower than) the specification by some value. As we learn more about receiver performance we may discover there is margin that allows some relaxation of TDECQ performance. The opposite could also be true.

References

1. IEEE Std 802.3bm(TM)-2015: Physical Layer Specifications and Management Parameters for 40 Gb/s and 100 Gb/s Operation over Fiber Optic Cables

2. J. Petrilla, P. Dawe, G. Le Cheminant , ”New metric offers more accurate estimate of optical transmitters impact on multimode fiber-optic links” Design Con 2015.

3. IEEE 802.3bs (TM): Standard for Ethernet Amendment: Media Access Control Parameters, Physical Layers and Management Parameters for 200 Gb/s and 400 Gb/s Operation

4. IEEE P802.3cd Draft Standard for Ethernet Amendment: Media Access Control Parameters, Physical Layers and Management Parameters for 50 Gb/s and 100 Gb/s Operation

5. J. King, D. Leyba, G. Le Cheminant, “TDECQ (Transmitter Dispersion Eye Closure Quaternary) Replaces Historic Eye-mask and TDP Test for 400 Gb/s PAM4 Optical Transmitters”, Optical Fiber Communications conference, March 2017

6. D. Leyba, G. Le Cheminant “Understanding Compliance Testing for PAM4 Optical Transmitters” Design Con 2018

7. K. Tamura, contribution to IEEE 802.3cd project, http://grouper.ieee.org/groups/802/3/cd/public/July18/tamura_3cd_01c_0718.pdf