APPLICATION NOTE

TDECQ Measurement Update

Improvements and insights for PAM4 optical transmitter testing

KEYSIGHT

Table of Contents

TDECQ Overview

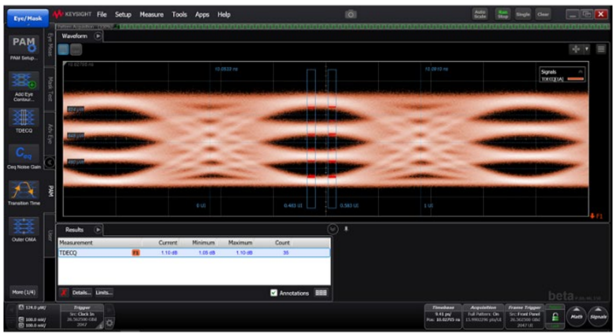

Transmitter Dispersion and Eye Closure Quaternary (TDECQ), first developed in 2016 in the IEEE 802.3bs 200/400 Gb/s project, has become the primary metric for gauging the quality of an optical transmitter using PAM4 modulation. Virtually every standard using PAM4 optical transmitters relies on the TDECQ metric, and this is expected to continue indefinitely. This paper will discuss how the measurement of TDECQ has evolved and improved, and what users of Keysight systems should expect for accuracy and consistency for a wide variety of test scenarios.

TDECQ measurement basics

TDECQ is a power penalty metric that determines how much extra power is required at the receiver from the transmitter under test, relative to an ideal transmitter, to achieve a specific symbol-error-ratio. This extra power is necessary to compensate for both non-ideal transmitter waveforms and the impact of chromatic and/or modal dispersion. It is like the transmitter dispersion penalty (TDP) metric used in NRZ optical communications standards. It was developed in the IEEE 802.3bs standard and was improved and modified in the IEEE 802.3cd, .cu and .db standards. It is a convenient way to specify the performance of a PAM4 transmitter, as it aggregates many transmitter attributes into a single metric, and the metric directly indicates the impact a transmitter's imperfections will have on overall system-level performance.

The step-by-step process to perform a TDECQ measurement can be found in "TDECQ Part 1: Making Accurate and Repeatable Compliance Measurements" found at Keysight.com https://www.keysight.com/us/en/assets/7018-06485/white-papers/5992-3635.pdf. Those that are not familiar with making TDECQ measurements should consider reading this document.

In summary:

- Acquire the pattern-locked waveform using the SSPRQ test pattern with a DCA oscilloscope channel configured in a TDECQ reference receiver bandwidth appropriate for the transmitter data rate
- Pass the waveform through the virtual TDECQ equalizer/receiver
- Perform the TDECQ measurement, with equalizer tap-weights automatically adjusted to yield the minimum observed TDECQ penalty

Note that there is a similar measurement, TDEC, available for NRZ measurements. For IEEE 802.3 bm, there is no requirement for pattern locking or virtual equalization. For ITU-PON G.9804.3, a pattern-locked equalized measurement is performed and discussed later in this paper.

TDECQ evolution

The intent of TDECQ is to predict how well the transmitter will operate in a real system. It is assumed that the receiver used in a real system will be able to operate with good sensitivity even if the transmitter is not ideal. One of the ways that the receiver can compensate for a non-ideal transmitter is through equalization. The decision timing of the receiver can also adapt to non-ideal signals. These attributes are included in the ‘virtual receiver’ used to assess TDECQ.

There has been a steady improvement in the electrical circuits used in optical receivers. Has this been accounted for in the virtual receivers used for TDECQ? From the original TDECQ method developed in 802.3bs, the 802.3cd project added the ability to adapt the receiver decision thresholds to compensate for the non-ideal spacing of PAM4 signal levels. Through the equalizer built into the virtual equalizer, the TDECQ method can significantly reduce the observed undershoot and overshoot that can be present on a highly peaked transmitter. A real receiver may not be able to manage this type of distortion as the receiver front end could become overloaded before any equalization is applied. In addition to the TDECQ metric, the 802.3cu standard added overshoot and undershoot limits. The VCSEL transmitter technology used for short-span multimode links is expected to require more extensive receiver equalization to achieve working links at 100 Gb/s per lane. In response, the 802.3db project extended the length of the virtual TDECQ equalizer from five taps to nine. (This is currently only used for the multimode transmitters specified by 802.3db).

If TDECQ is working as expected, the observed TDECQ penalty should indicate the impact the transmitter has on system level receiver sensitivity. If one transmitter has a TDECQ of 2 dB and another a TDECQ of 3.2 dB, receiver sensitivity should degrade by 1.2 dB if the former transmitter is replaced with the latter. In general, the virtual equalizing receiver used in 802.3 standards is assumed to represent the ‘worst case allowed’ receiver that a transmitter might be paired with. Real receivers likely perform significantly better than the virtual receiver used for TDECQ, with continual improvements to compensate for transmitter imperfections. The result is that TDECQ can be a conservative estimate of transmitter performance when used in a real system.

Improvements in TDECQ measurement speed

The time required to perform the TDECQ measurement is based on three primary components: Acquiring the transmitter waveform, optimizing the tap weights of the TDECQ virtual equalizer, and performing the TDECQ analysis. The time required to acquire a waveform is based on the DCA sample rate, the number of symbols in the test pattern (pattern length), and how many samples are acquired for each symbol. The sample rate of the DCA is fixed at 250 Ksamples/second. TDECQ is specified to be tested using the SSPRQ test pattern, with a length of 65535 symbols. The only variable for acquisition time is then the number of samples per symbol. The best approach to setting this value is to let FlexDCA set this parameter automatically. (Setup/Acquisition Setup with Samples/UI set to automatic):

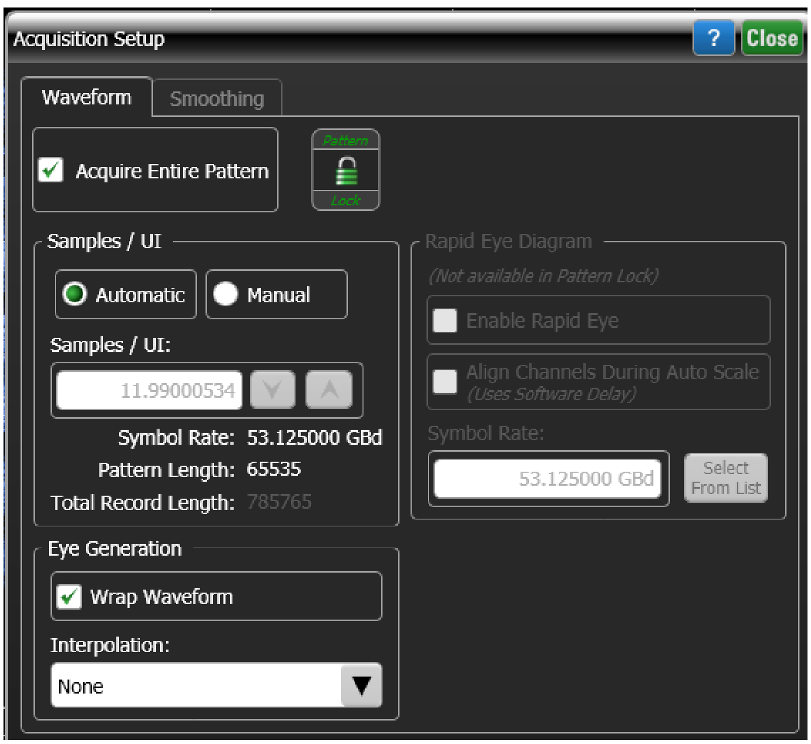

Figure 1. Use the Automatic setting to configure Samples/UI

In the automatic setting, the value may seem peculiar, just slightly less than an integer value. In the example above the samples/UI value is set to 11.99000534. Why not round up to 12? Note that the product of the samples/UI and the pattern length (66535) is equal to the total record length (total number of samples recorded for one complete acquisition of the SSPRQ pattern) and in this case the record length is 785765 samples. If the samples/UI value were an integer value, such as 12, the eye diagram created by overlaying all the samples into a single unit interval would have all samples positioned at the same 12 locations across the unit interval (see figure 2).

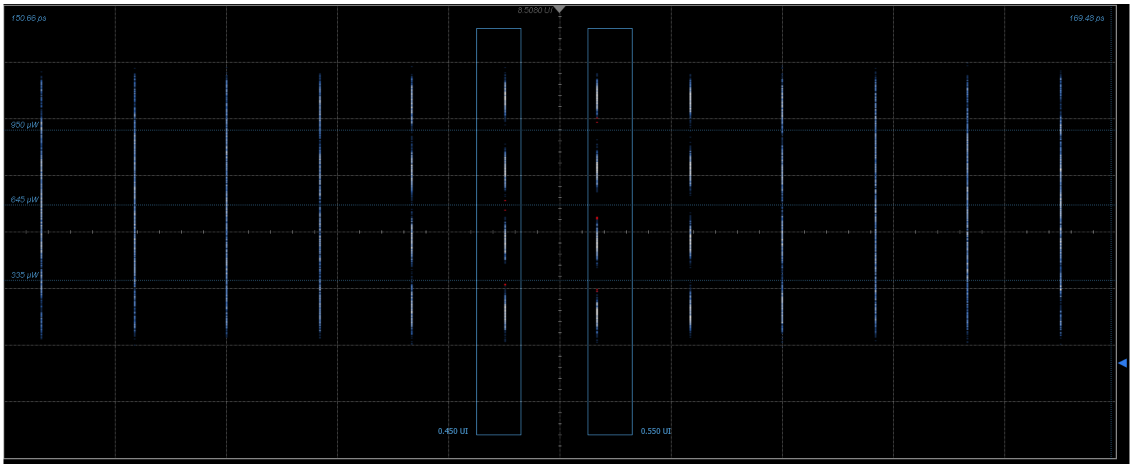

Figure 2. An integer Samples/UI setting results in an ineffective grouping of samples

Consider that TDECQ is expected to be performed with data spread across histogram widths of 0.04 UI. This is not achieved with the 12 samples/UI configuration. Obviously this could be mitigated by setting the samples/UI value to a very large number. A more efficient approach is to set the samples/UI to a non-integer value. The resulting eye diagram aggregated from the long waveform pattern has a very high density with effectively no gaps. In fact, the automatic setting of the samples/UI value is determined with the pattern length known and the objective of an accurate TDECQ analysis through the acquisition of only a single waveform. The time to acquire a waveform for TDECQ analysis is the record length divided by the sample rate, which in this example is 3.14 seconds.

Gauging the time required for the optimization of the virtual equalizer tap weights is complex. IEEE 802.3 defined the TDECQ method with the key details found in Clause 121.8.5.3. “When the ….. value … cannot be (improved) by further optimization of the equalizer tap coefficients, then TDECQ is calculated”. This effectively implies that every possible combination of tap weight values should be considered before making the final TDECQ calculation. The intent of this statement is to acknowledge that any combination of tap weights can yield a valid TDECQ result. From the perspective of designing an automatic test, this is problematic. Tap weights can be adjusted with very high resolution, leading to an extremely large number of tap weight combinations to consider, potentially leading to long test times. Another problem is that this is not something that a real system receiver would do, and the TDECQ reference receiver should be attempting to emulate real receiver capability.

In the IEEE 802.3db project (100G lanes over multimode fiber) a 9-tap virtual equalizer is used for TDECQ and it was acknowledged that test times could significantly increase due to longer tap-weight optimization. While IEEE 802.3 clause 121 remains intact, clause 167.8.6 indicated that optimization methods such as minimum mean squared error (MMSE), may be used to determine equalizer tap weights to reduce test time. The non-iterative MMSE method should not yield lower TDECQ values compared to an ‘infinite’ search, but usually is within 0.1 dB. This is true for both single-mode and multimode transmitters.

FlexDCA allows a user to choose optimization methods. The TDECQ reference equalizer setup panel can be accessed in two ways. Select Math (in lower right screen) /signal processing or select Measure/Waveform Signal processing/Signal processing. In either scenario, pull down the TDECQ operator and click on it to see the following:

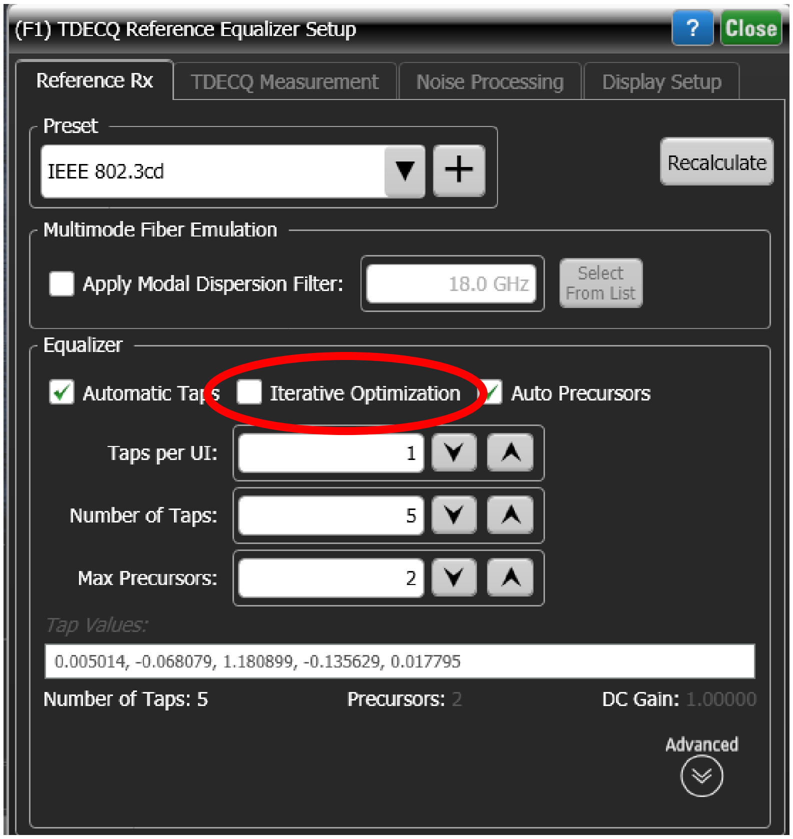

Figure 3. Setup panel for the TDECQ equalizer and measurement

In the equalizer section, there is a selection titled "Iterative Optimization" (firmware revision A.05.70 and above). In the default state this selection is 'off'. When this is checked, the tap-weight optimization will approach a 'full search' of all possible tap weights. When this is not checked, an MMSE approach is used to optimize tap weights. For any TDECQ analysis, the recommended method is to have the iterative optimization disabled. See Figure 3. In early versions of the TDECQ measurement built into FlexDCA, the iterative method could take over a minute to complete. The MMSE method was much faster at about 5 seconds. The optimization methods have been steadily improved over time. The iterative approach can now require 15 to 30 seconds. Improvement in the MMSE optimization speed (iterative optimization disabled) has been even more significant, with tap-weight values determined in less than 1 second.

Once the waveform has been acquired and the tap weights optimized, the time required to assess the TDECQ penalty is almost instantaneous. The total test time to perform a TDECQ measurement, including waveform acquisition, equalizer optimization, and analysis is now under five seconds. (Setting up and locking the clock-data recovery (CDR) system can add 1 to 2 seconds to the overall test process).

TDECQ Measurement Accuracy and Uncertainty

Assessing the accuracy of a TDECQ measurement is complicated. There is no national laboratory-based TDECQ standard reference signal for determining the accuracy of a TDECQ measurement system. Therefore, there is no TDECQ measurement accuracy specification in the data sheets for DCA's. However, it is possible to determine the repeatability of TDECQ analysis and the relative accuracy of test systems.

As stated above, the intent of the TDECQ measurement is to assess the impact transmitter impairments have on a link budget. The TDECQ metric effectively compares a transmitter's performance to an ideal transmitter. It would be easy to interpret that a 0 dB TDECQ value indicates a transmitter has no impairments. This could be true but is likely not the case. Consider that the waveform from the transmitter is passed through a virtual equalizer, which can mitigate impairments and significantly improve the signal. The bandwidth of the DCA channel is set to half of the baud rate. For example, a 53 Gbd "100 Gb/s" transmitter is observed in a 26.56 GHz bandwidth. An *ideal* signal passed through the half-baud bandwidth (without equalization) will incur some inter-symbol interference resulting in a TDECQ of approximately 0.8 dB. (802.3 TDECQ limits have been compensated to account for this). TDECQ values, with virtual equalization, observed directly from transmitters are typically in the 1 to 2 dB range.

One way to assess the accuracy of a TDECQ measurement would be to observe the measured transmitter in a real system. If one transmitter has a TDECQ of 1.5 dB and a second transmitter has a TDECQ of 2.2 dB, we would expect the system receiver sensitivity, measured at a symbol error ratio (SER) of 4.8e-4 (the SER used to determine TDECQ) to be 0.7 dB (2.2-1.5) better (operates at a lower input power) when the first transmitter is used, compared to when the second transmitter is used. This type of experiment was conducted and contributed to the IEEE 802.3 cd project: https://grouper.ieee.org/groups/802/3/cd/public/July18/tamura_3cd_01c_0718.pdf. Pages 7, 9 and 10 show a close 1:1 relationship between TDECQ and receiver sensitivity. Two important points should be made. First, the receivers used in the experiments were close in design to the virtual receiver used in the TDECQ analysis, with equalizers slightly longer than the 5-tap TDECQ equalizer. Second, the comparisons are for relative and not absolute values of TDECQ. If the TDECQ values shown on page 7 had all been 0.5 dB lower at 1.1, 2.0, and 2.6 dB rather than 1.6, 2.5, and 3.1 dB, we would have still observed the desired 1:1 correlation with receiver sensitivity, even though the set of data reduced by 0.5 dB would imply better-performing transmitters. Similar experiments documenting correlation between TDECQ and receiver sensitivity were documented in the IEEE 802.3 dj project: https://grouper.ieee.org/groups/802/3/dj/public/23_09/lecheminant_3dj_01b_2309.pdf .

Without a laboratory-grade reference signal with a known TDECQ value, the best we can do is to assess the relative accuracy of TDECQ. TDECQ measurements have been available and in a stable form since 2017. The Keysight solutions have been used extensively over that time and have effectively become a standard reference measurement. It is then useful to document and avoid the mechanisms and situations that can degrade the relative measurement accuracy of the Keysight systems to ensure the best possible results are achieved.

Low signal power

The TDECQ measurement process is designed to maintain good accuracy in the presence of noise, specifically noise from the DCA. DCA channel noise is measured directly and recorded whenever a user module calibration is performed. The channel noise is then accounted for in the TDECQ calculation (see 802.3 clause 121.8.5.3 equation 121-11, where σs represents the noise of the oscilloscope).

TDECQ should not change as the signal level is reduced. If we attenuate a signal, we should see a constant TDECQ signal versus signal power. However, the measurement process has limits. If the oscilloscope noise is large relative to the transmitter signal power, the process of adding virtual noise (802.3 121.8.5.3 paragraph 5) begins to fail and the DCA will report "SER?". The signal power level at which this occurs depends on both the noise of the DCA channel and the quality of the transmitter. If TDECQ is low, measurements can be made at lower powers than a transmitter with a high TDECQ.

For any DCA configuration, whether it is a low-noise N1092 or the ultra-wide bandwidth N1032, as the signal power is steadily attenuated, TDECQ accuracy eventually degrades. The typical result is that the reported TDECQ begins to increase and eventually SER? is reported. "SER?" is intended to indicate that the eye closure is so severe that the TDECQ SER target is exceeded prior to adding any virtual noise. This implies that the power penalty for that transmitter is infinite. When this signal is connected to a receiver, the target SER will never be observed, even at very high input powers. However, for very small signals the SER? result is likely due to eye closure from the DCA channel noise. (Note that SER? is very different from a result that has a '?' annotation, such as TDECQ: ? 1.33 dB. The value annotated by '?' indicates that the measurement has been performed, but something about the setup was questionable. A common example would be if TDECQ was measured on a pattern other than SSPRQ, and the expected run of consecutive symbols to compute OMA was not observed. The measurement is correct but not performed exactly as specified by IEEE 802.3. Full details are available by selecting "Details" in the measurement results section of the display).

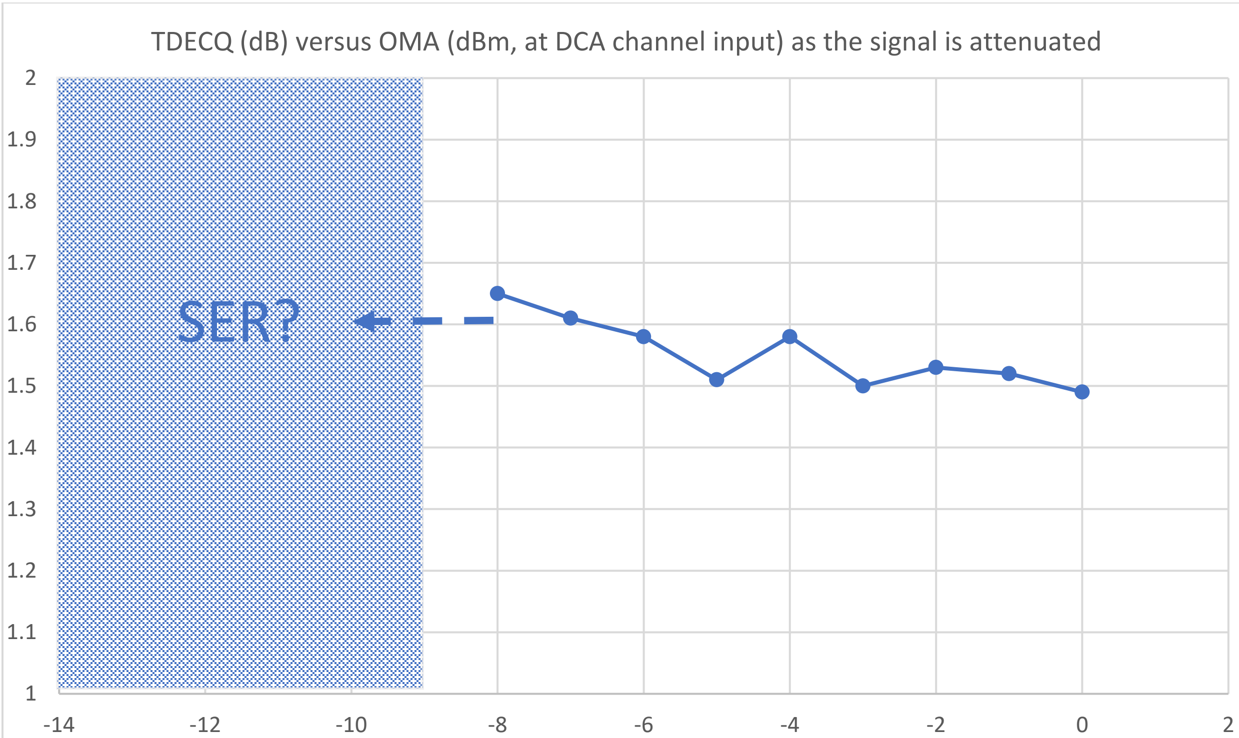

Figure 4. As a signal is attenuated eventually the DCA will report SER? rather than a true value

Figure 4 shows the TDECQ results for a signal as it was attenuated. Ideally TDECQ should have been constant. Measurement repeatability on the order of 0.1 dB is common. For OMA levels below -8 dBm the DCA reports SER?

FlexDCA firmware revision A.07.50 provides new capability to improve the accuracy of TDECQ analysis on low power signals. In the TDECQ reference receiver, under the Advanced tab, the Low-SNR (low signal-to-noise) feature has been added:

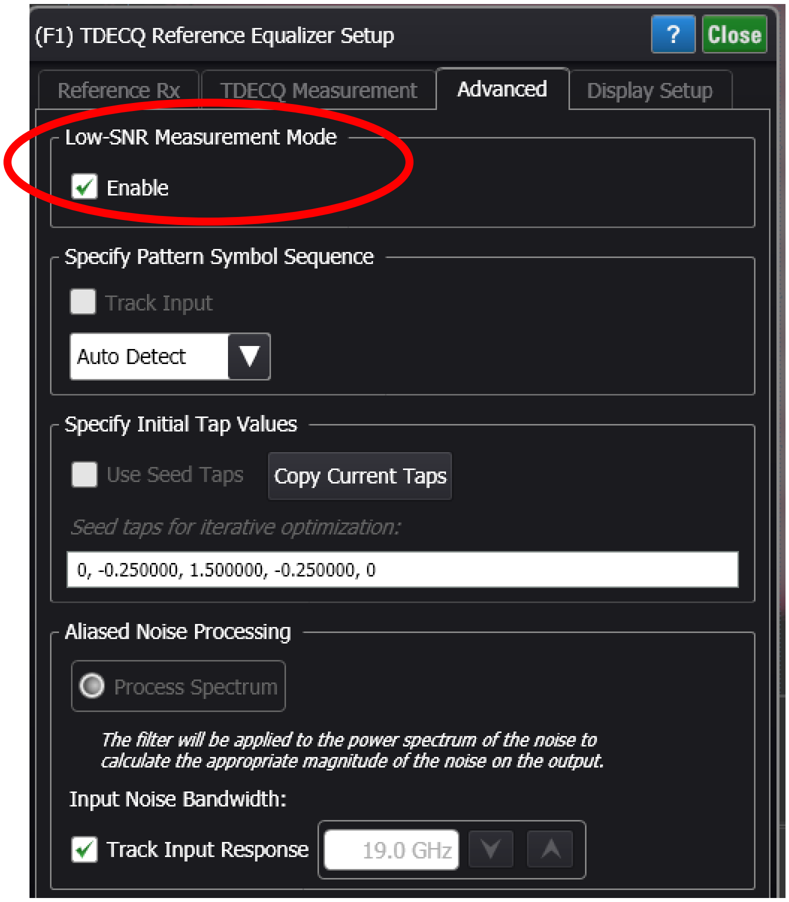

Figure 5. Low-SNR measurement setup

Low-SNR, enabled by default, improves the DCA's ability to configure and execute the TDECQ measurement when eye diagrams are closed by DCA channel noise, due to either low signal power or high DCA noise. The system is then better able to mathematically remove the DCA noise, as designed in the IEEE 802.3 TDECQ method.

The impact of Low-SNR can be enhanced by increasing the Samples/UI setting from the automatic value to something in the range of 100 Samples/UI. (Setup/Acquisition setup/Samples/UI -> Manual. Do not enter a number. Use the 'up arrow' key to increase the value. For example, the automatic value may be set to 7.98998993, and can be manually incremented in four steps to 127.98998993).

The following three graphs show the impact of the Low-SNR feature using the N1092 DCA-M. Measurements were made with Low-SNR disabled (green), with Low-SNR enabled (Blue), and with Low-SNR enabled and the Samples/UI increased (yellow). Measurements were made on a typical single-mode transmitter, a high TDECQ single-mode transmitter, and a multimode transmitter. In all cases signals were attenuated and values reported until the SER? limit was observed. The charts show the reported TDECQ (dB) versus the OMA (dBm) value seen at the DCA channel input.

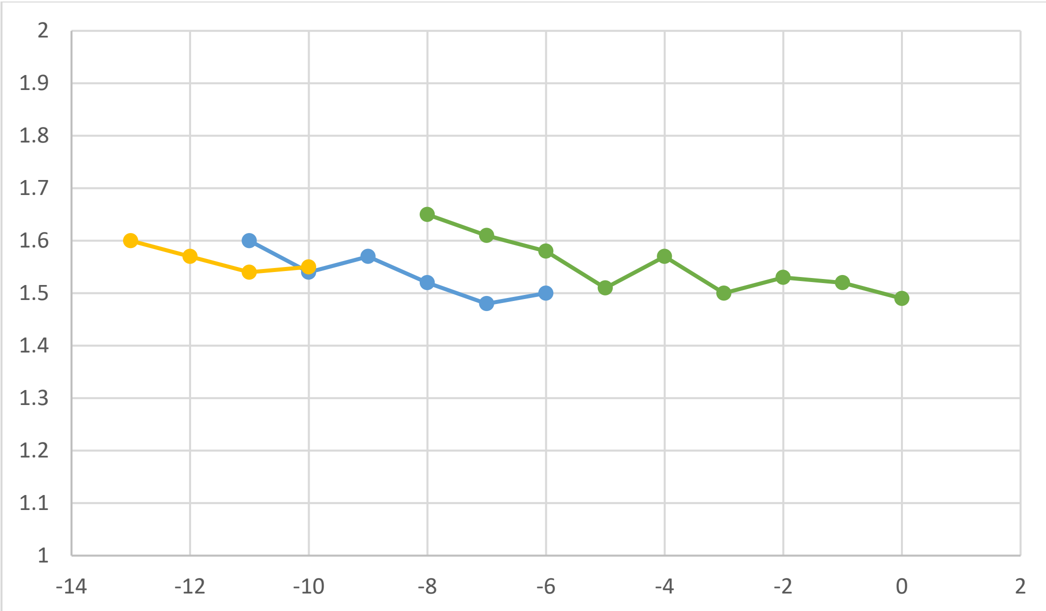

Figure 6. TDECQ measurement range is extended with Low-SNR enabled: Low TDECQ single-mode example

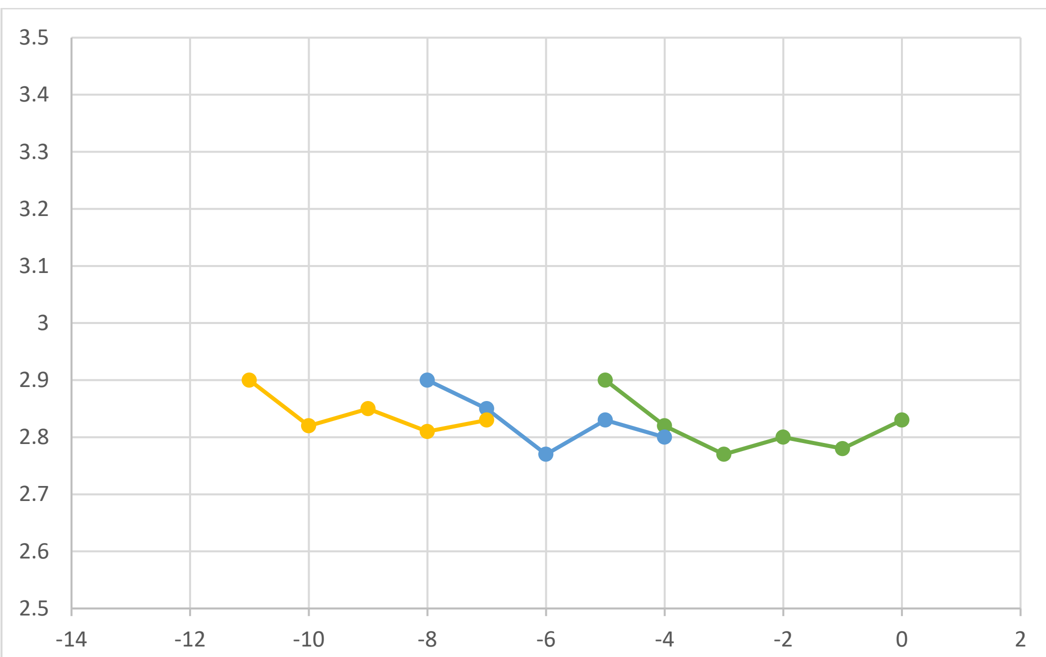

Figure 7. TDECQ measurement range is extended with Low-SNR enabled: High TDECQ single-mode example

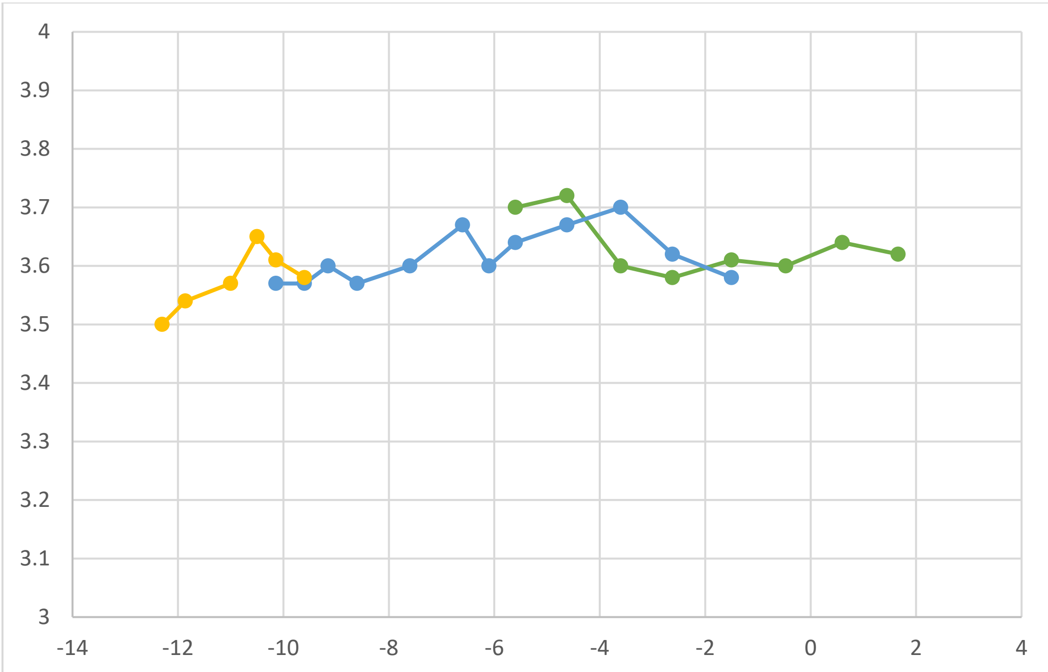

Figure 8. TDECQ measurement range is extended with Low-SNR enabled: Multimode example

For the cases above, the range over which valid TDECQ measurements can be made is extended 3 to 4 dB with the Low-SNR feature enabled. The measurement range is extended further, 2 to 3 dB by having both Low-SNR enabled and the Samples/UI setting increased above 100. Ideally TDECQ should be the same for all settings and all powers. Most of the change in TDECQ can be attributed to typical measurement repeatability (~0.1 dB). The sensitivity limit of any DCA system is dependent both on the noise of the DCA channel and the quality of the signal being measured. This type of experiment, enabling and disabling Low-SNR, can be performed on any DCA system to gauge the impact of Low-SNR. Note that the tradeoff for increasing the Samples/UI is a significant increase in waveform acquisition time. Overall measurement time will be on the order of 30 seconds rather than the typical 3 to 5 seconds. Increasing the Samples/UI should only be used for testing very small signals. Note that the Low-SNR feature does not and should not be expected to yield lower TDECQ results versus when the feature is disabled. Low-SNR is intended to provide correct results at much lower input powers into the DCA.

Measurement inaccuracy due to measurement algorithms

It was shown above that as a signal was attenuated eventually the reported TDECQ changed when it should have remained constant. A common misconception is that when two systems are measuring the same signal, the system reporting a lower TDECQ value is *always* more accurate than the system reporting a higher value. The opposite has been observed. Scenarios exist where TDECQ is incorrectly reported lower than the true value through poorly designed waveform analysis and TDECQ calculation. While the IEEE 802.3 standard provides extensive guidance on the TDECQ measurement method, some elements of the measurement process are left to the developer and are subject to errors that can yield optimistic TDECQ values. One example occurred when TDECQ “improved” as signal levels were attenuated. The value should not change with attenuation. It is believed that as the waveform power was reduced, samples that ‘crossed the decision threshold’ and were represented as symbol errors, were now incorrectly being interpreted as non-errored symbols resulting in a false reduction in TDECQ. Another example was seen where signal noise was incorrectly processed and reduced through the TDECQ equalizer. The error can be exposed by adding noise to the signal, which should result in a degraded TDECQ. In the faulty system, TDECQ did not change with added noise, indicative of a system that incorrectly reported low TDECQ values.

Measurement inaccuracy due to transmitter stability

If the transmitter signal is not stable (inconsistent coupling, power variation due to control loops etc.) the eye diagram will smear and TDECQ will be high or unmeasurable. The signal being tested needs to be stable over the entire waveform acquisition time, typically three to five seconds. Prior to equalization it can be hard to differentiate a good signal observed with an incorrect test setup from a bad signal that has been setup correctly. A good measurement practice before making any TDECQ measurement is to verify that the signal is correctly pattern locked and is stable. Simply switch the DCA from Eye Mask mode to Oscilloscope mode. A stable sequence of symbols should be observed.

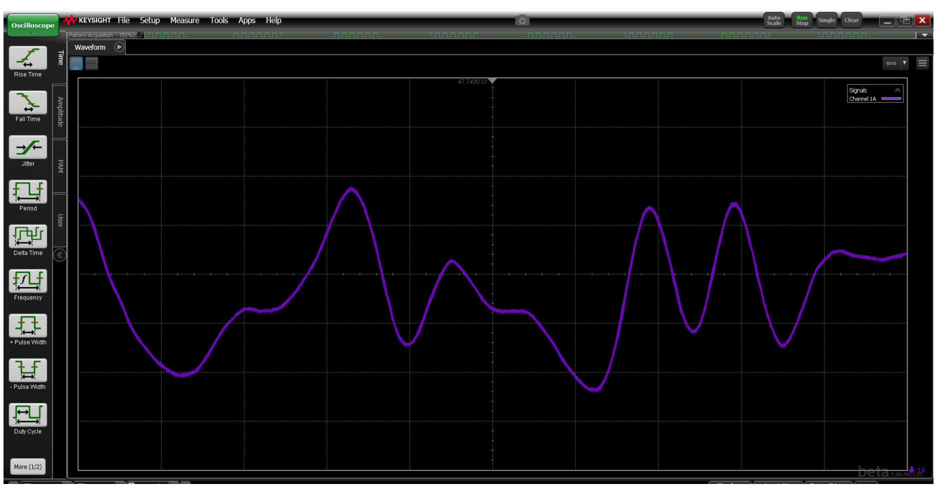

Figure 9. Observing the waveform in Oscilloscope Mode can verify the DCA setup and the stability of the waveform

The impact of CDR on measurement accuracy

A CDR system has been an integral component of IEEE 802.3 optical transmitter test schemes for many years. The intent of the CDR is to track and remove low frequency jitter from the observed transmitter waveform, similar to the behavior of a real receiver. CDR is achieved through some form of a phase-locked loop (PLL). The PLL will have a loop bandwidth, and jitter with frequencies within the loop bandwidth will be tracked by the PLL and passed to the recovered clock. Jitter frequencies above the loop bandwidth will not be passed to the recovered clock. When the timing reference used to trigger the DCA is derived from the signal being observed, jitter that is common to both the trigger and the waveform will not be observed on the displayed waveform. Since the PLL passes low frequency jitter to the recovered clock triggering the DCA, from a measurement perspective this achieves a jitter high-pass function. In a real system, it is expected that the receiver will be able to track and tolerate low frequency jitter, thus the transmitter test system, using CDR, will not penalize a transmitter for having low frequency jitter. It is common for the 3 dB jitter high-pass bandwidth to be 4 MHz with a 20 dB per decade slope.

Recent generations of optical transmitters often multiplex lower rate electrical lanes and use retiming to generate the higher rate optical signal. This retiming can result in decorrelation between the optical signal and the electrical modulation. When optical signals travel through a dispersive channel, the optical signal can lose correlation from the original timing at the transmitter. A clock from the electrical pattern generator may not provide an adequate signal to trigger the DCA. Both scenarios can be aided by extracting the DCA trigger from the optical signal with a CDR system.

Generally, a wider loop bandwidth results in lower jitter on the displayed waveform and a reduced TDECQ. It is important to have an accurate setting to avoid overly optimistic or pessimistic TDECQ results. In practice, instrumentation that produces the desired 4 MHz jitter high-pass function is complicated to build. Precisely achieving the 4 MHz loop bandwidth requires a sophisticated tuning scheme. This has been steadily improved in both CDR hardware and the controlling software. Significant improvements in CDR loop bandwidth accuracy were enabled with FlexDCA firmware A.06.40. The CDR loop bandwidth is dependent on the input signal. In the FlexDCA CDR controls, a “Tune on Relock” feature was added to facilitate better control of CDR loop bandwidth for a wide range of input signals and data patterns. To achieve the most accurate loop bandwidth and jitter high-pass frequency, this feature should be enabled.

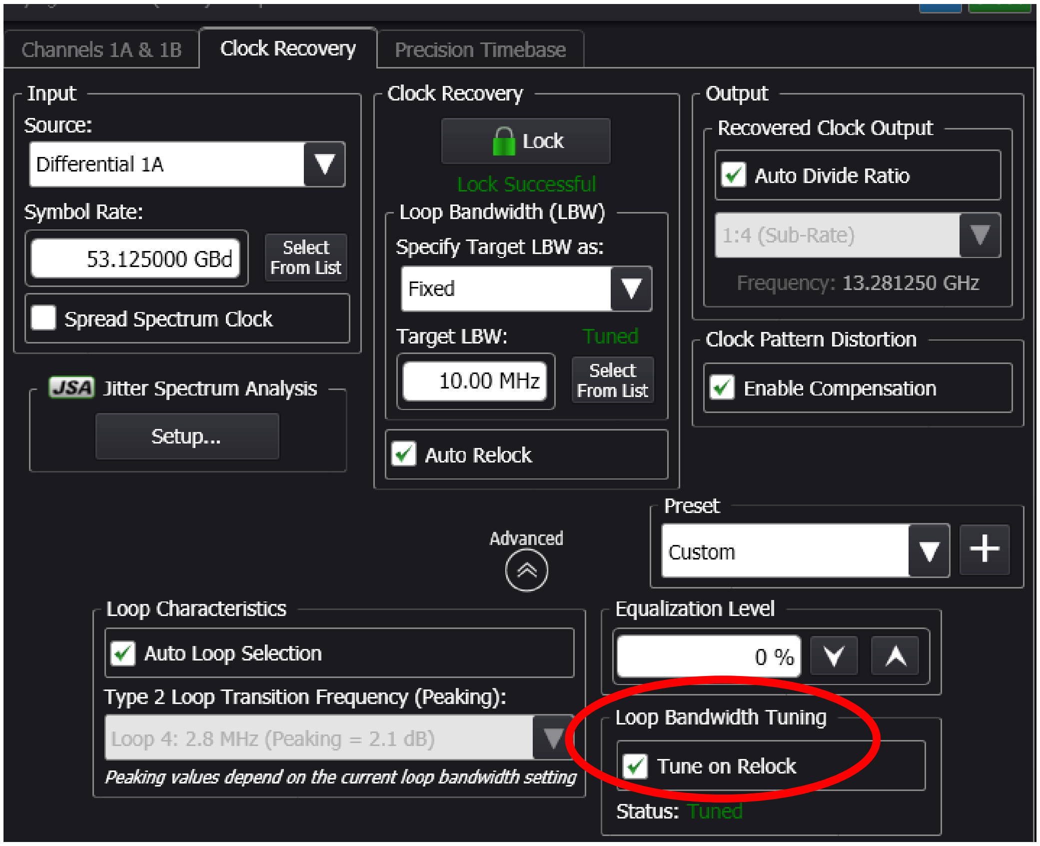

Figure 10. Clock recovery loop bandwidth tuning setup

Industry standards treat CDR as a simple jitter high-pass function and there is no other guidance on how the CDR system should operate. In practice, the overall performance of the CDR hardware, and its impact on TDECQ goes beyond simply setting the loop bandwidth accurately. TDECQ measurements can be influenced by the CDR beyond the loop bandwidth setting. To maintain TDECQ measurement integrity, Keysight CDR systems are designed to operate over a very wide range of data rates, patterns, signal powers, and signal quality. But there are scenarios where even a high-performance CDR system can be 'stressed' by a degraded signal. The result can be a distortion within the CDR system that is not representative of the timing of the signal being measured. This translates to spurious energy that, when within the loop bandwidth of the CDR, will be transferred to the recovered clock and degrade the observed signal and resulting TDECQ.

One method to determine if CDR based distortion exists is to reduce the loop bandwidth of the CDR system. As discussed above, the observed jitter should be reduced when the loop bandwidth is increased. In most scenarios loop bandwidth reduction should then yield a higher TDECQ. If CDR distortion occurs, its impact on TDECQ can be suppressed when the distortion is 'pushed out of the system' by reducing the loop bandwidth. In this scenario a reduction in loop bandwidth results in a lower TDECQ.

Reducing the loop bandwidth below the standard 4 MHz value to reduce reported TDECQ should still be considered a standards compliant measurement, as the test system is removing less jitter from the observed waveform than is allowed and there is no risk of a 'false positive' result. However, this type of distortion is often systematic, repeatable, and identifiable by the DCA system. This allows the system to compensate and minimize the impact of the distortion while maintaining the common 4 MHz value.

In firmware revision A.07.00 a feature labeled ‘Clock Pattern Distortion’ was added to the CDR setup menu. The clock pattern distortion and its correction (CPDC) were improved and updated in firmware revision A.07.40:

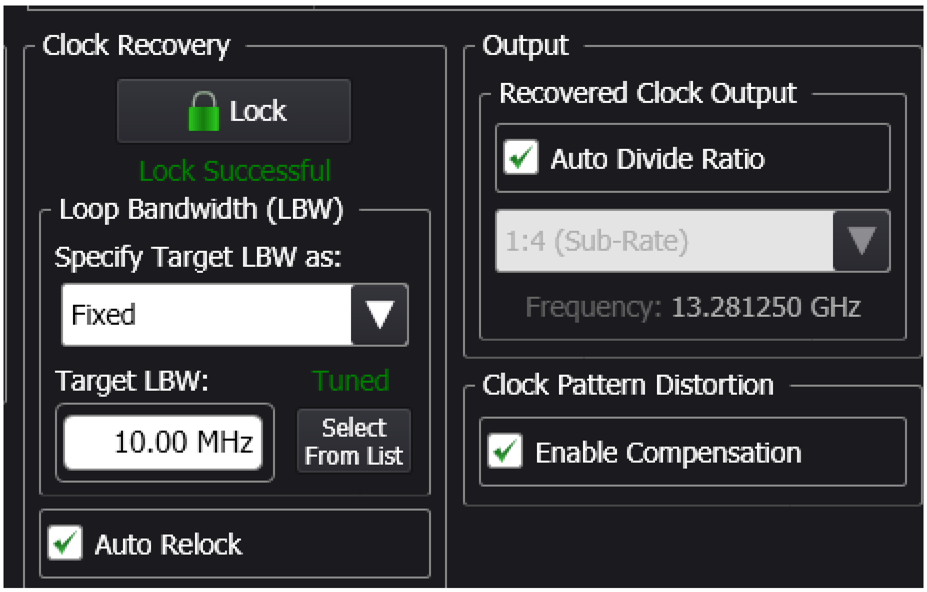

Figure 11. Clock pattern distortion correction setup

This CPDC feature is defaulted in the enabled state. It will not impact TDECQ results except in the case where there is systematic distortion of the recovered clock signal. (This can be verified by comparing TDECQ values obtained with and without CPDC enabled). The CPDC is achieved without any user intervention or calibration process. The correction is determined in real time and is unique for the signal being observed. CPDC will compensate for the non-ideal deterministic behavior of the CDR due to internal ISI and non-ideal behavior of the CDR in response to variations in edge and symbol density throughout the pattern.

CPDC corrects the displayed waveform and is not a mathematical correction of the TDECQ calculation. CPDC is expected to have the largest impact on test scenarios where the test system is ‘stressed’ such as with low signal power or when testing with high dispersion channels. Note that CPDC cannot differentiate between distortion caused by the test system CDR and similar distortions generated by the device under test. For example, if the signal driving the transmitter employs some form of PLL that is sensitive to the SSPRQ stress pattern, the stress impairment may be compensated for in the TDECQ analysis.

Measurement error due to DCA channel frequency response

The frequency response of the DCA channel when performing a TDECQ measurement should follow a fourth-order Bessel response. The -3dBe (electrical) bandwidth should be at a frequency equivalent to half the transmitter baud rate. For example, a 106 Gb/s PAM4 transmitter operating at 53.125 Gbaud should be tested with the DCA channel bandwidth set to 26.56 GHz. DCA channels can be configured as reference receivers for a variety of data rates. For NRZ eye-mask testing, the DCA analog hardware is physically adjusted to achieve the required bandwidth and frequency response. For example, the N1092 has bandwidth settings for 20.6, 25.8, 26.6, 27.9, and 28.05 Gb/s. The analog frequency response may not be perfect, but always falls within the tolerance limits required by ITU specifications. For TDECQ measurements, a different approach is taken. TDECQ measurements are always pattern locked. That is, the sampling process is synchronized to the pattern repetition rate. This allows the DCA system to mathematically correct its frequency response and create an almost perfect Bessel filter. It is known that variation in frequency response can lead to system-to-system variation in eye-mask results. It is difficult to determine the impact that a non-ideal frequency response will have on TDECQ accuracy. The DCA measurement approach, with its ideal pattern-locked frequency response, assures that this error will be negligible.

Achieving agreement between test systems

Keysight offers several systems that are capable of TDECQ analysis on transmitters operating at 26, 53, and 112 Gbaud ranges. The most common systems for 26 and 53 Gbaud are the N1092 combined with the N1078 CDR or the N1092 with an internal CDR system (currently available for single-mode signals only). The N1092 and N1077B are available for multimode signals. The systems use very similar CDR hardware, but the systems are not identical. With the N1092/N1078 system, the signal splitting used to provide a reference for the CDR is typically 50/50 while the N1092/N1077B is 70/30. For the N1092 with internal CDR, the signal splitting occurs after optical to electrical conversion. The split is electrical, and the split ratio is designed to balance the available power to the DCA channel and the power available for CDR, knowing that if either is insufficient, measurements can be degraded.

The N1092/N1078 system was the first complete system available for 100G/53 Gbaud TDECQ analysis and became an industry standard. As the N1092 integrated CDR system was developed, a key objective was for it to yield similar results as the N1092/N1078 system. Compared to the N1092/N1078 system (using a 50/50 optical splitter), the integrated CDR system provides higher power to the DCA channel and lower power to the CDR. Both systems maintain the DCA channel frequency response required for TDECQ analysis. The CDR signal path experiences more equalization in the N1092/N1078 system.

The signal characteristics that can result in TDECQ measurement variation are varied and complex. Although significant work was done to achieve agreement between the integrated CDR and the N1092/N1078 systems, there are scenarios where a systematic difference can be observed. In these scenarios the general trend is that the N1092/N1078 system will yield a slightly lower TDECQ value, on the order of 0.3 to 0.5 dB. Analysis indicated that the CDR system can in some cases be stressed by the SSPRQ pattern used for TDECQ analysis. Since this effect is systematic, this presents the opportunity for compensation. The use of the CPDC feature discussed above can reduce the systematic differences between two systems, below the ~0.3 dB difference observed without CPDC. Use of CPDC can also improve measurement to measurement repeatability on a single system.

Configuring a test system for the highest accuracy

As discussed above, a variety of factors can impact the accuracy of a TDECQ measurement. Some of these factors can be influenced by how a test system is configured. For example, the N1092 can be configured with an internal CDR. Alternatively, an external CDR can be used. It is important to provide sufficient power to the oscilloscope channel to minimize the impact of channel noise. The CDR also needs sufficient power to be able to reliably lock to the test signal and provide a quality timing reference. With the internal CDR, signal splitting to the CDR is done to provide the best balance of signal power to the oscilloscope channel and signal power to the CDR. The split ratio is set based on an expected range of power from the transmitter. It is important to realize that the required input levels are dependent on signal quality. That is, as the TDECQ value increases, more signal power can be required for the measurement channel and/or for CDR. What might be an optimum split ratio for one signal might not be optimum for another signal. The N107X external CDR solutions can be configured with an internal splitter, or with no splitter, (allowing the user flexibility in choosing the split ratio using their own splitter). In unusually low power or high TDECQ scenarios, experimentation on split ratios may yield a more robust measurement system.

TDECQ for 200G Lanes and Analysis in the Presence of High Dispersion

TDECQ analysis quantifies the impact that transmitter eye closure will have on the system link budget. It also is intended to quantify the impact of dispersion. Due to chromatic dispersion, CDR can be essential for transmitter waveform characterization at the end of a long fiber. For standards designed to operate at spans of 10 kilometers and longer, transmitter performance when observed at the end of the fiber span can be difficult to measure. Channel attenuation reduces signal power and chromatic dispersion can significantly distort the waveform. Low signal power and degraded waveform quality can impact the CDR process and degrade TDECQ measurement accuracy. Analysis of low power signals is improved using the Low-SNR feature. The impact of non-ideal CDR, discussed above, can be reduced using the clock pattern distortion correction (CPDC) feature.

In LR4 systems (4 wavelengths on a single fiber) TDECQ values are sometimes reduced by extracting the clock trigger from a low dispersion lane. This assumes that all lanes are synchronous and timed with a common clock. This has the benefit of providing more power to the DCA channel and providing a lower stress signal to the CDR system. This technique varies from the general TDECQ method described in IEEE 802.3 but is useful when extracting a clock from the lane being tested does not yield a useful TDECQ value. Again, long-span measurements that might have resulted in SER? may now yield valid results though the use of Low-SNR and CPDC without resorting to extracting clock from an adjacent lane.

It is not unusual to observe high TDECQ values in high dispersion cases. 200G lanes may also exhibit high TDECQ values when measured using the legacy TDECQ reference receiver. Yet these transmitters are often capable of interoperating with a receiver in a real system. The virtual receiver used in TDECQ analysis is expected to emulate the worst-case receiver expected in real systems. It is likely that latest generation system receivers are more robust and have stronger equalization than the ‘worst-case’ model, allowing them to interoperate with transmitters that may fail the general TDECQ test limit.

When waveform dispersion is high there are scenarios where a valid TDECQ value cannot be observed, and the “SER?” result is given, indicating an infinite power penalty for the transmitter when observed with the common TDECQ reference receiver. In addition to invoking Low-SNR and CPDC, there are modifications in the TDECQ reference receiver that may yield TDECQ values rather than “SER?”. In some configurations may not be compliant to the standards, but can be useful to have practical TDECQ values and see how changes in transmitter performance result in measurable changes in waveform quality and a subsequent TDECQ penalty.

Several attributes of the TDECQ reference receiver are defined in IEEE 802.3 standards. FlexDCA will set these parameters according to the specific standard selected. However, these parameters can be changed for a customized reference receiver. These parameters are:

TDECQ Reference Receiver:

- Number of FFE taps
- Maximum number of FFE precursors

Figure 12. TDECQ virtual equalizer setup

TDECQ measurements:

- Target SER
- Histogram properties
  - Histogram width
  - Histogram spacing
  - Threshold optimization adjustment limit

Figure 13. TDECQ measurement setup

Number of FFE taps and precursors: This sets the length of the linear equalizer. Most standards use a 5-tap equalizer, while the 802.3 db standard extended the length to 9 taps. The 800G MSA has specified a 21-tap equalizer. As of January 2024, the IEEE 802.3 dj standard for 200G lanes will use a 15-tap equalizer. The number of precursors that will be allowed is currently unknown, but also easily overridden.

Target SER: The TDECQ measurement process determines how much noise can be added to a signal until the target SER is reached. If the SER limit is increased, it follows that more noise can be added to the transmitter signal. But this does not necessarily mean that a lower TDECQ penalty will be observed. Recall that the added noise is compared to the noise that could be added to an ideal signal, which also will be analyzed at a higher SER limit. However, if the transmitter being measured has significant eye closure even after equalization, the DCA can report a SER? result rather than a numeric TDECQ value. This indicates that the SER limit, typically 4.8e-4, has been exceeded before any virtual noise has been added to the signal. By increasing the SER limit, the transmitter signal may have enough margin for a TDECQ measurement to be constructed. For 200 Gb/s lanes, discussions in IEEE 802.3dj indicate that a more powerful FEC will be required to allow a higher pre-FEC SER. This would imply that the target SER used for TDECQ would also be increased. Increasing the target SER from 4.8 e-4 to a value in the 5e-3 to 9e-3 range are being considered in 802.3dj.

Histogram properties: The waveform samples used to construct the TDECQ measurement are typically from two vertical slices of the PAM4 eye diagram, 0.04 UI in width, and separated by 0.1 UI. The histogram width is kept narrow, but not so narrow that there will be a small population of samples acquired. Two histograms are acquired to emulate the likelihood that a real receiver will not be making decisions at an ideal time instant. The TDECQ algorithm is allowed to locate the two histogram slices anywhere in the UI, but the two slices are always ganged together, separated by the histogram spacing parameter. The 0.1 UI histogram spacing value was defined in 2016. It is difficult to know if this emulates current receiver behavior. However, TDECQ values do improve, sometimes significantly, with small reductions in the histogram spacing, perhaps 0.07 UI rather than 0.1 UI.

In the following example, a waveform has a SER? TDECQ result:

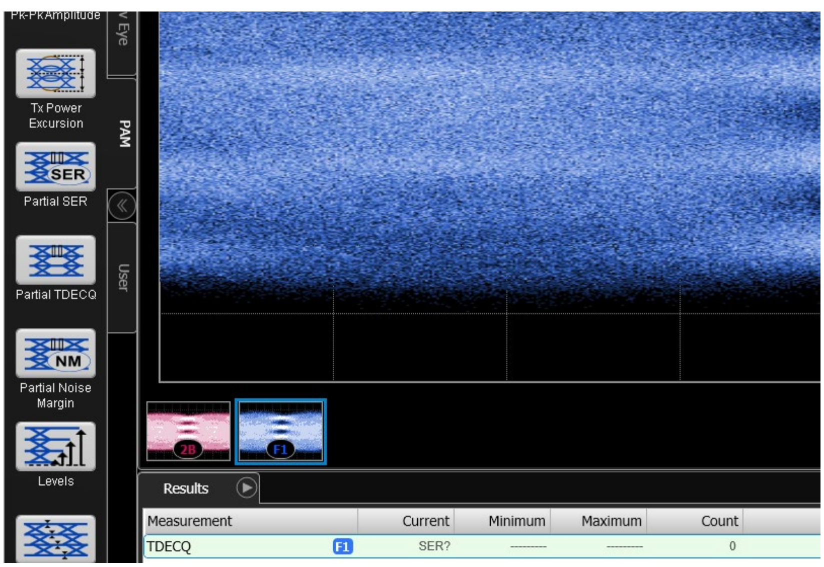

Figure 14. TDECQ result reported as SER?

As the histogram spacing is gradually reduced a TDECQ value is eventually reported, at 5 dB with an 0.07 UI spacing and less than 4 dB with an 0.05 UI spacing. This is not a compliant metric but can be useful compared to an SER? value.

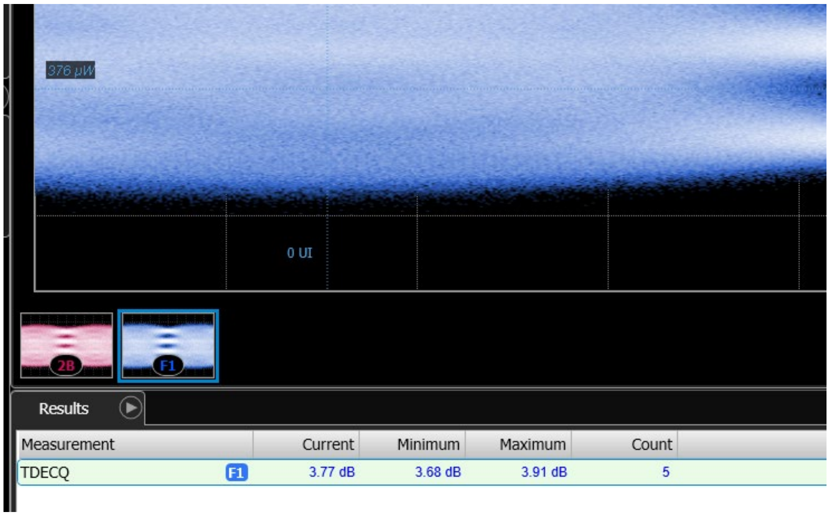

Figure 15. TDECQ result changes from SER? to a useful metric through reducing the histogram spacing

Threshold optimization adjustment limit: This parameter is intended to emulate a receiver's ability to adjust its three decision thresholds to manage the non-ideal linearity (spacing) of the PAM4 signal levels. Increasing this parameter should result in a lower TDECQ value if the waveform has significant nonlinearity.

Unique issues for TDECQ analysis of 200 Gb/s (100 Gbaud) signals

200 Gb/s lanes using 100 Gbaud PAM4 transmitters are under development now. IEEE 802.3 Standards using these transmitters will likely not be finalized until sometime in 2026. Yet it will be important to develop the technology required for these systems in advance of completion of the standard. It has been decided that TDECQ will continue to be the primary transmitter metric. Eventually the standards will set specific TDECQ reference receiver parameters. When stable, these will be deployed in the FlexDCA interface, initially as a beta setting in the TDECQ setup menu, and a final setting when the standard is completed. Until then, manual adjustments of the TDECQ reference receivers will be necessary.

TDECQ analysis requires a half-baud channel bandwidth for the DCA, which would be in the 56 GHz range for a 224 Gb/s signal. This can be achieved using either the N1030 or N1032 DCA channels. The N1032 channel has the added benefit that it can be configured as a very wide bandwidth channel (over 100 GHz (-3 dBo)) allowing accurate analysis of the raw waveform. However, this comes at the cost of reduced responsivity compared to the N1030. This is seen in the noise of the channels. Configured as a TDECQ reference receiver, the N1030 has more than 2 dB lower noise than the N1032. If the priority is 224 Gb/s TDECQ analysis, the N1030 has better sensitivity.

Setting the correct bandwidth for a 200G TDECQ measurement is straightforward using the following steps:

- Connect the signal to the DCA and confirm that some signal is present on the display
- Pattern lock the signal (this assumes a repeating pattern such as SSPRQ is being transmitted. (A valid pattern lock can be confirmed by observing a stable sequence of symbols with the DCA in Oscilloscope Mode. See figure 9)
- With a valid pattern lock, the DCA channel menu will activate the SIRC menu. Check the SIRC box:

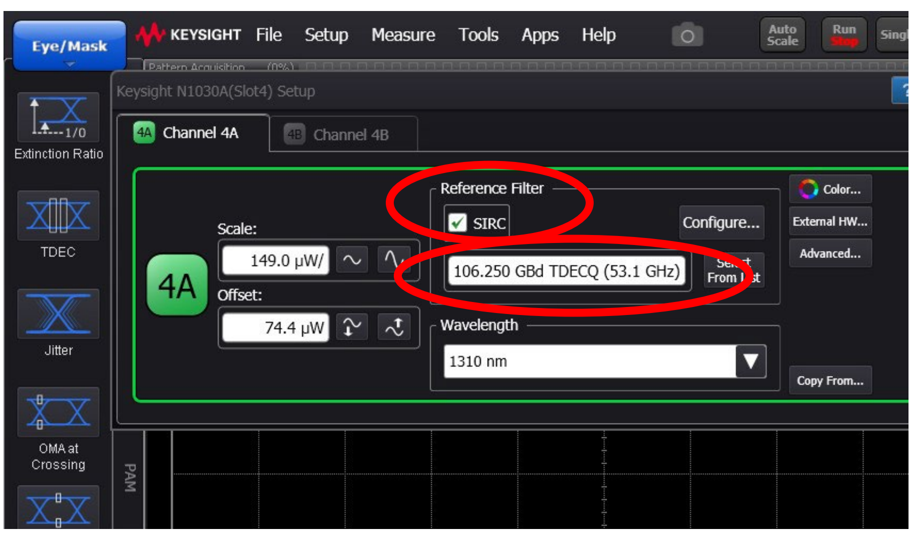

Figure 16. Using the pattern locked SIRC menu to configure TDECQ channel filter

- With the SIRC function active, a large selection of possible filter settings are available, including TDECQ filters that have a bandwidth equal to half of the baud rate, such as 106.25 Gbaud/53.125 GHz
- When operating at a baud rate that is not included in the available filters, a custom filter can be created. Select “Configure” to display the menu. Manually enter a value in the Channel Bandwidth. For example, if operating at 113.44 Gbaud, the half-baud bandwidth is 56.72 GHz. Note that the reference filter rate will be 75.62667 Gbaud and not 113.44 Gbaud. This is due to the legacy NRZ filter definitions, based on 75% of the baud rate. The filter 56.72 GHz configuration is still appropriate for 113.44 PAM4 TDECQ analysis.

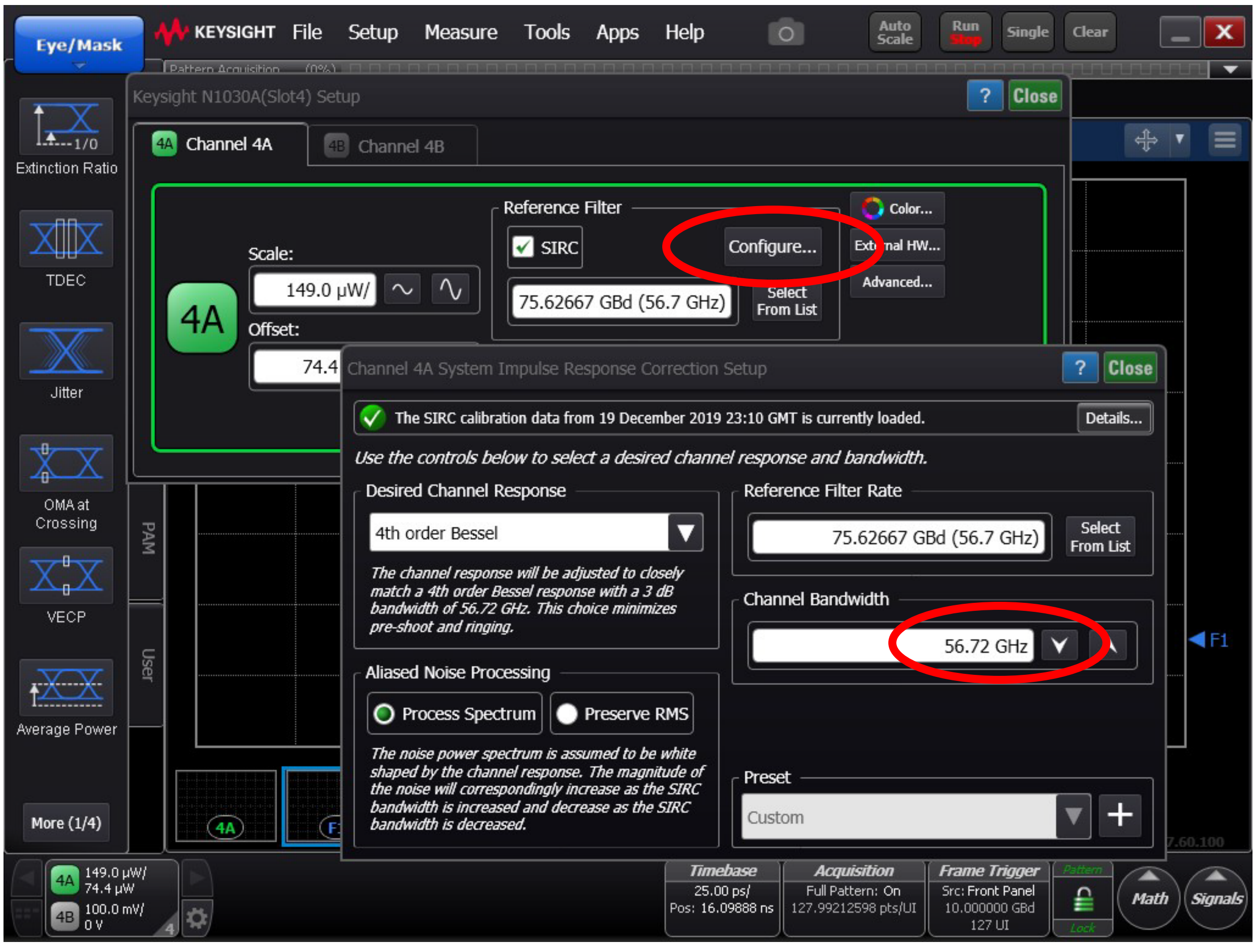

Figure 17. Creating custom TDECQ channel bandwidths

Once the DCA channel bandwidth is set correctly, the TDECQ measurement is made similar to legacy TDECQ tests, with the exception that parameters such as the equalizer length (number of taps) and target SER may need to be adapted to the latest standard definitions for 200G lanes.

Noise removal techniques

The extra bandwidth required to observe a 200 Gb/s waveform results in extra noise in the DCA channel compared to 100 Gb/s DCA channels. This can obscure the view of waveforms. Waveform averaging is a common technique used to remove random noise and jitter from a signal. (Setup/Acquisition setup/Smoothing). Note that this can remove the DCA random noise added to a signal, but also the random noise and random jitter of the transmitter. TDECQ analysis with averaging may yield an optimistic result. Also, measurement throughput is significantly reduced. If 8 waveforms are used for averaging, it will take 8 times longer to complete the acquisition and report a TDECQ value.

An alternative method to remove DCA channel noise relies on the fact that a sampling oscilloscope does not synchronously sample random noise. Noise signals are aliased and are perceived as being very high frequency. This does not present any problems in displaying eye diagrams and general waveforms. Noise is displayed accurately. However, if the waveform is subject to digital signal processing, noise, if not managed properly, is removed from the processed waveform. For this reason, math blocks in FlexDCA allow control of how noise is managed:

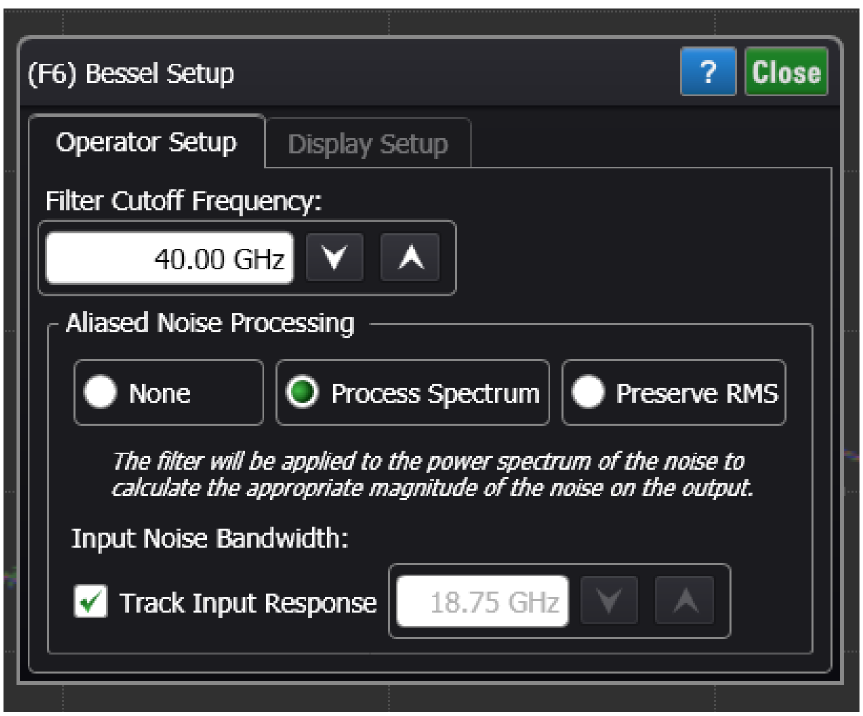

Figure 18. Processing noise through Math filters

Managing noise through Math filters

There are several mathematical filters available in the FlexDCA signal processing operators. Almost all allow some flexibility in how noise is managed with Preserve RMS, Process Spectrum, and None as available settings:

Aliased Noise Processing: None

The option **None** does not make any special considerations for aliased noise. This is an appropriate choice when the waveform is known to have no aliased components, such as from a real-time or simulated source. This also is the preferred choice when dealing with waveforms where the desired outcome is to *remove noise and interference* components that are not correlated to the trigger signal.

Aliased Noise Processing: Process Spectrum Selection

The option **Process Spectrum** works by applying the transfer function of the filter to the power spectrum of the noise to determine the appropriate magnitude of the noise on the output signal.

By tracking the accumulated effects of the filtering operations, accurate noise processing can be done. In addition to the sampled waveform, information about the acquisition channel and noise power spectrum are maintained in each signal and appropriately processed by the filter function.

The default behavior of the **Process Spectrum** noise processing option is to use the noise power spectrum of the input signal. If the input signal is a sampling scope channel with SIRC active, this spectrum will be established by the measured hardware response of the channel. For other channels, the response will be assumed Gaussian with a 3 dB frequency corresponding to the nominal channel bandwidth. This behavior can be overridden by clearing the **Track Input Response** checkbox and manually entering a bandwidth. If this option is utilized the response will be presumed Gaussian with the selected 3 dB bandwidth.

Aliased Noise Processing: Preserve RMS Selection

The option **Preserve RMS** is appropriate when the noise bandwidth is very low relative to the channel and filter bandwidth. It is also appropriate when the aliased components are known to have most of the power within a range of in-band frequencies. This could be from laser RIN, for example, or from intentionally modulated interference/crosstalk. When *Preserve RMS* is selected, the RMS magnitude of the aliased components will be scaled by the DC gain (the sum of the taps) of the filter. For the low-pass filters (Bessel, Butterworth, Gaussian, Sinc), the DC gain is unity. For the equalizers and embedding/de-embedding operators, the DC gain depends upon the settings of the operator.

A subtle but important aspect of noise management when ‘None’ is selected, is the noise is not managed and is then *effectively removed from the signal*. This can be a very efficient method to ‘clean up’ a signal and display only the deterministic elements of waveforms.

For a 224 Gb/s measurement, a simple method to remove noise from the signal is to place a Sin X/X or “Brickwall” filter in front of the TDECQ equalizer:

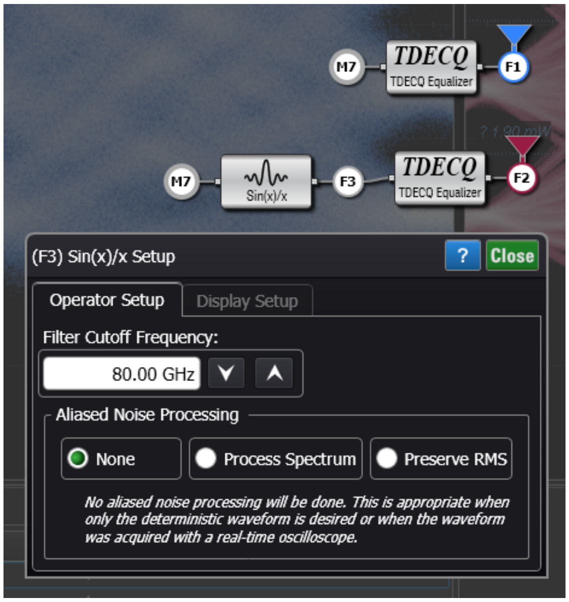

Figure 19. Configuring a filter to eliminate noise

It is important to set the filter bandwidth and noise processing correctly. If the primary intent of the filter is only to remove noise, the bandwidth should be set to a value such that the filter has no impact on the waveform shape. In this example, the waveform is a 224 Gb/s (112 Gbaud), and the DCA channel bandwidth is 56 GHz with a fourth order Bessel response. The signal itself does not have significant energy beyond 56 GHz. The Math filter bandwidth is set well above that rate, at 80 GHz. Normally a brickwall filter is not used for DCA waveforms, but when used specifically for noise removal, this is the most effective filter. To remove noise, the noise processing is set to None.

Compare the two waveforms with (blue) and without (red) the 80 GHz filter. The results are dramatic:

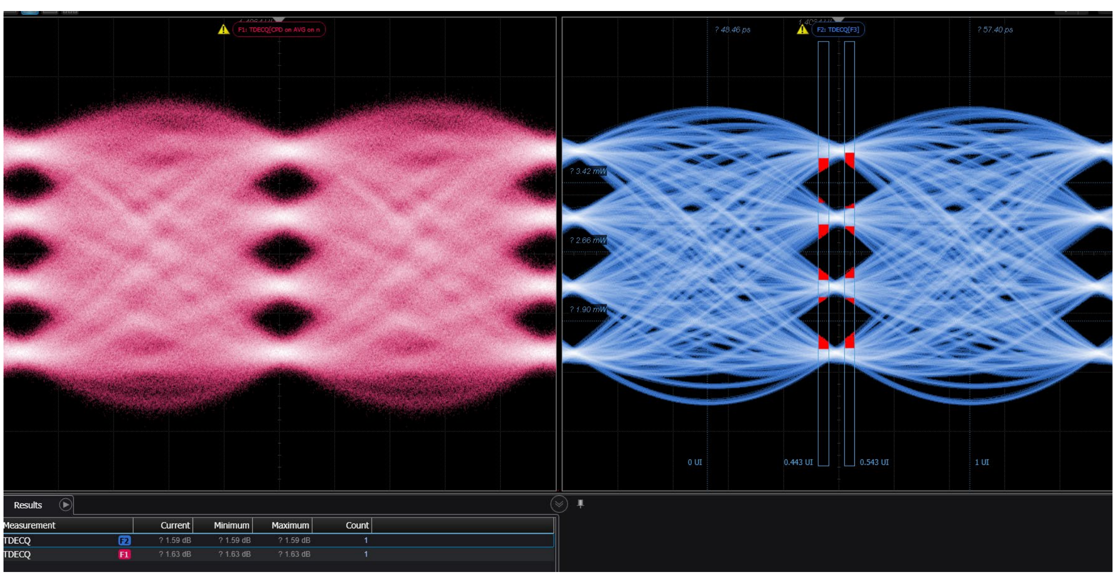

Figure 20. With noise removed deterministic signal features are easily observed

The efficiency of the noise filtering can be enhanced by increasing the Samples/UI in the acquisition setup. Instead of the automatic value such as 7.99…. Samples/UI, switch to manual and use the arrow keys to increase the value to perhaps 127.99….. Samples/UI. More noise will be removed at the expense of the time required to acquire a waveform. Experimentation is useful to determine the ideal setting.

It is important to gauge what impact the noise filtering has on TDECQ, knowing that both the DCA channel noise and the transmitter noise will be removed. When TDECQ is measured on the waveform above, on the left (no noise filtering) the value is 1.6 dB. With the removal of the noise the TDECQ (for the waveform on the right) is 1.4 dB. While the noise filtering has a large impact on the displayed waveform, it has only a small impact on the TDECQ metric. One of the advantages of the TDECQ method is that the DCA channel noise is measured and mathematically removed from the calculation even if the DCA channel noise is not removed from the displayed waveform (on the left). In this example, the dominant noise contribution appears to be from the DCA channel noise. There is only a small reduction in TDECQ when the transmitter noise is also removed through filtering. Recall that the LSNR feature, without use of noise filtering, provides the most accurate method to measure TDECQ on low power or high noise scenarios.

TDEC for ITU-PON 50G Systems

The ITU standard for 50G systems, G.9804.3, uses NRZ modulation, and the TDECQ measurement discussed above does not directly apply. However, the general method for TDECQ was first developed in the IEEE 802.3bm standard as an NRZ measurement called TDEC (transmitter dispersion and eye closure). In addition to being a PAM4 metric, another key addition made by TDECQ was the use of an equalizing reference receiver. To reduce cost the ITU G.9804.3 standard will use 25G class components operating at 50 Gb/s. This will require significant equalization by the system receiver. Thus, the ITU-PON TDEC measurement is a hybrid of the TDEC from 802.3bm and PAM4 TDECQ. PON receivers also commonly use APD-based receivers. The noise performance of the APD is signal level dependent. This must be accounted for in the TDEC reference receiver.

The process for making the ITU-PON TDEC measurement is similar to a TDECQ measurement. First set up the DCA channel with the correct bandwidth. G.9804.3 specifies a “fourth order Bessel-Thomson response with a 3 dB bandwidth of 18.75 GHz”. This filter is available in the Channel menu once the system has been pattern locked and the Reference Filter SIRC feature is enabled. There are two filters for 49.7664 Gbd, one with 37.3 GHz bandwidth (shown below) which is used for the eye-mask test and one with an 18.75 GHz bandwidth (for TDEC). Be sure to select the 18.75 GHz setting for TDEC and switch to 37.3 GHz for the eye-mask test:

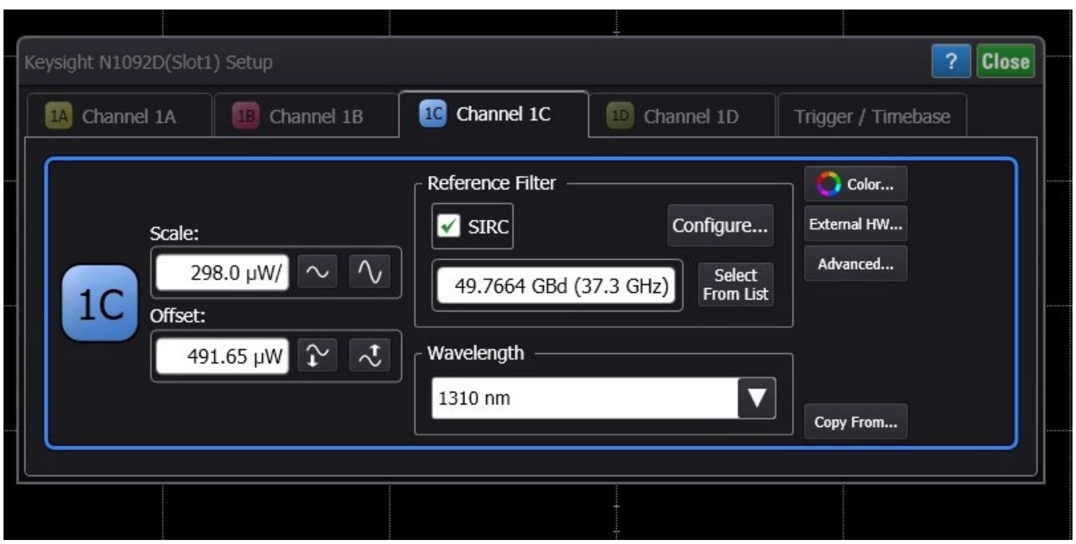

Figure 21. Configuring the DCA channel for 50G PON test (eye-mask setting)

The TDECQ reference receiver block from the Math/signal processing menu is set up. FlexDCA should recognize the signal as being NRZ, and this results in a modified TDECQ setup. The reference receiver tab allows the selection of the ITU PON 13 tap equalizer:

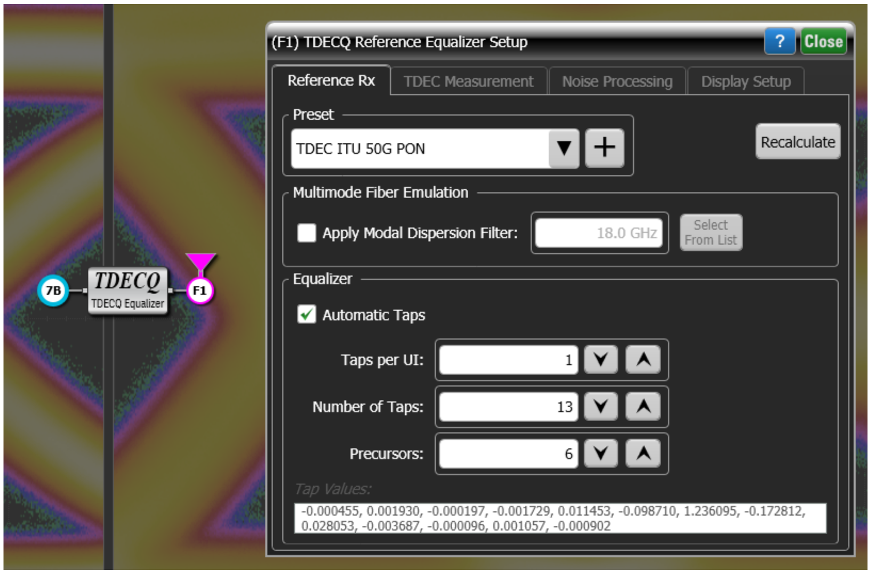

Figure 22. Configuring the TDECQ/TDEC reference receiver for 50G PON TDEC

The TDECQ measurement setup tab will now be active for a TDEC measurement. While the TDEC measurement is for an NRZ signal, the user interface leverages the PAM4 measurement interface. If the menu still shows a TDECQ measurement tab rather than TDEC, it is possible that the signal detection is not in autodetect mode and is set to PAM4. Either select Eye/Mask/PAM setup/Signal Types to autodetect or set the signal to NRZ:

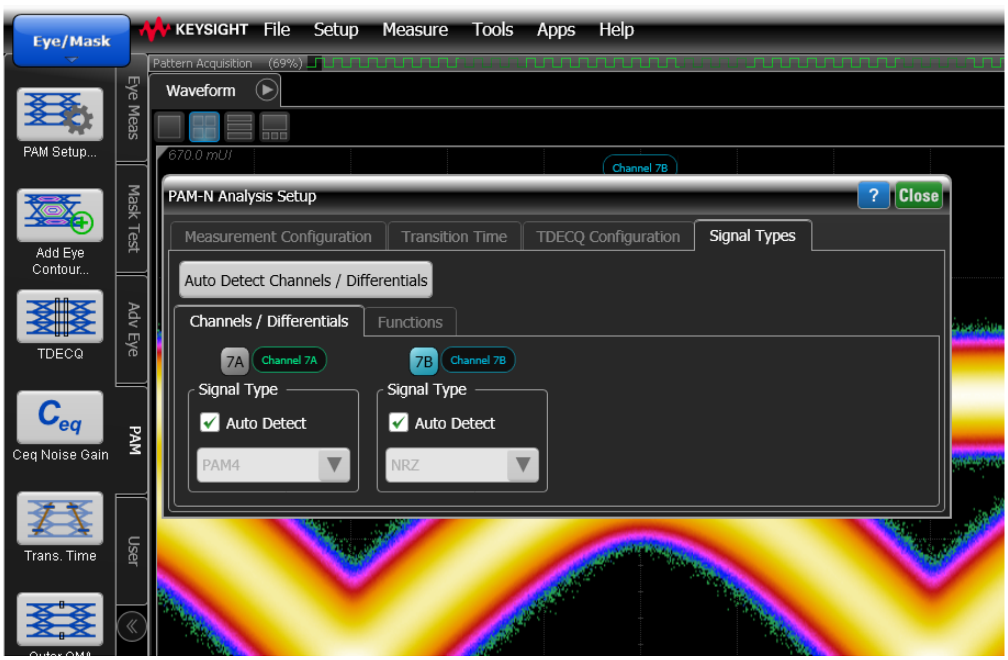

Figure 23. Ensuring the DCA is configured for and NRZ signal for TDEC analysis

With the test signal recognized as NRZ, the TDEC measurement tab is seen and the ITU-PON preset can be selected. Note the additional noise parameters in the TDEC calculation:

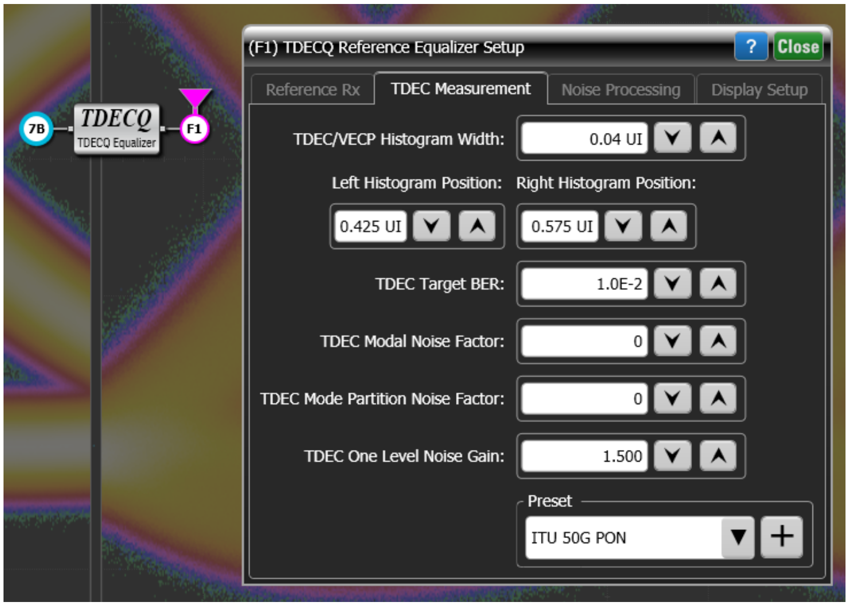

Figure 24. Configuring the TDEC measurement for 50G PON

To execute the TDEC measurement, FlexDCA needs to be switched from the Eye/Mask 'PAM ' configuration tab to the 'Eye Meas' tab which will display the TDEC measurement:

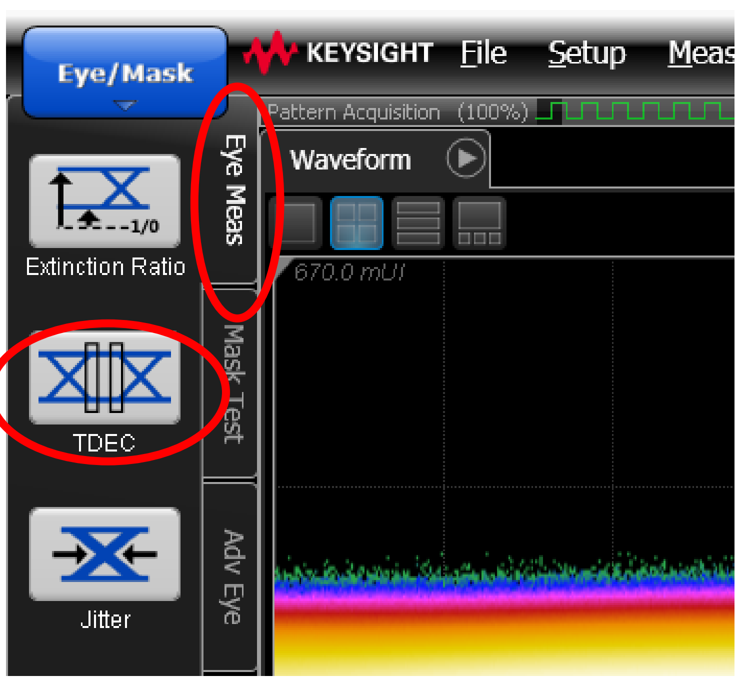

Figure 25. The 50G PON TDEC measurement is in the Eye Meas menu

The TDEC measurement is applied to the output of the TDECQ equalizer block (configured for TDEC) with the reference OMA value acquired from the waveform:

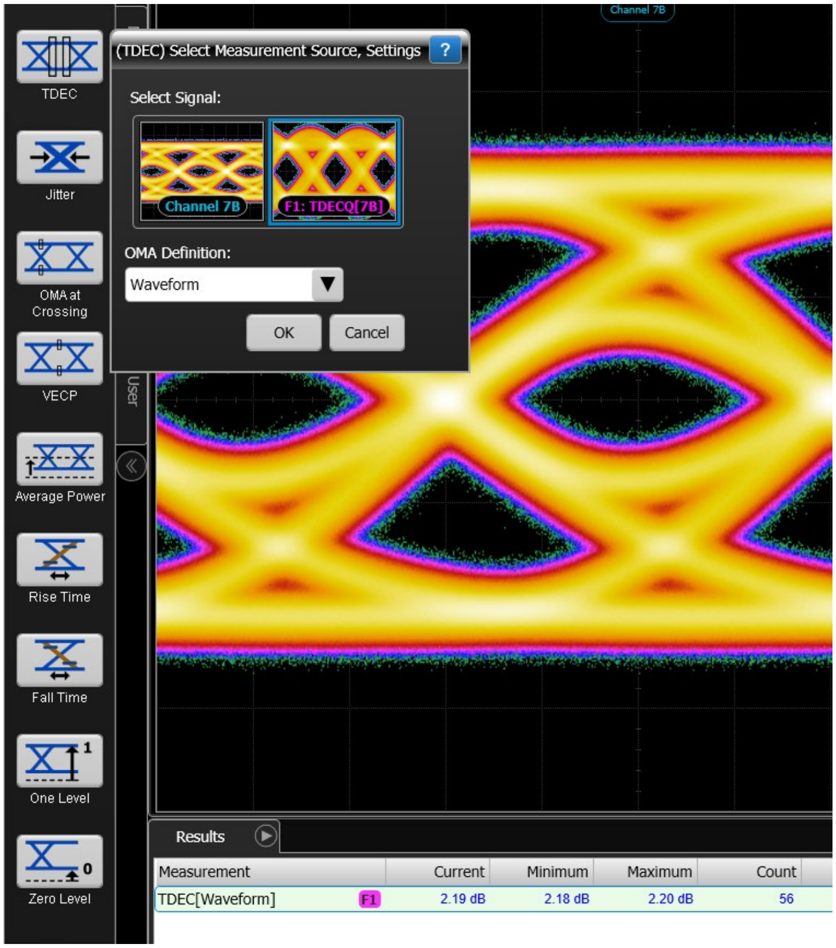

Figure 26. The 50G PON TDEC measurement is applied to the output of the TDEC reference equalizer (Math function)