



Anatomy of an Eye Diagram

Application Note

Abstract

This paper describes what an eye diagram is, how it is constructed, and common methods of triggering used to generate one. It then describes different ways that information from an eye diagram can be sliced to gain more insight. It also discusses some basic ways that transmitters, channels, and receivers are tested. It is designed to give an engineer new to this field a basic grasp of the concepts commonly used.




Figure 1. Overlaying of bit sequences to form an eye diagram.

Eye Measurements Basics

Eye diagrams are a very successful way of quickly and intuitively assessing the quality of a digital signal. A properly constructed eye should contain every possible bit sequence from simple 101's and 010's, through to isolated ones after long runs of consecutive zeros and other problem sequences that often show up weaknesses present in system design.

What does it show?

Eye diagrams show parametric information about the signal – effects deriving from physics such as system bandwidth health, etc. It will not show protocol or logical problems – if a logic 1 is healthy on the eye, this does not reveal the fact that the system meant to send a zero. However, if the physics of the system mean that a logic one becomes so distorted while passing through the system that the receiver at the far end mistakes it for a zero, this should be shown in a good eye diagram.

Common ways of characterizing an eye are to measure the rise times, fall times, jitter at the middle of the crossing point of the eye, the overshoot present and many other numerical descriptions of eye behavior in order to compare devices being measured. Instruments usually offer automated measurements that simplify and speed up the taking of such measurements.

The Effects of Triggering on Eye Diagrams

Many eye diagrams are constructed on test equipment that uses test patterns that repeat, such as PRBS patterns generated by a BERT (Bit Error Ratio Tester). Such equipment can commonly generate a variety of trigger signals:

1. A clock trigger at the same rate as, and synchronous with the data signal
2. A divided clock trigger at some divide ratio of the data rate often related to a power of 2 such as $\div 4$, $\div 16$, etc.
3. A pattern trigger – a signal that provides a trigger once per pattern repetition.
4. The data itself can be used as a trigger.
5. The last option is to derive the trigger signal by using clock recovery on the data signal. See Figure 2.

Each method provides different results when used to construct a waveform.

Clock Trigger provides a classical eye diagram containing all possible bit transitions in one display.

Divided Clock Trigger also produces an eye diagram, and this can be useful when the instrument being used to construct the eye has a trigger input bandwidth narrower than the data rate of the signal being viewed. This method can produce a good eye unless the pattern length divided by the divide ratio yields an integer – for example a 128-bit pattern viewed with a $\div 4$ clock. In this case, the trigger signal will coincide with the same bits in the pattern each time while consistently missing other parts of the pattern – leading to an incomplete eye. See Figure 3.




Figure 2. An eye diagram formed with a full rate clock trigger.




Figure 3. An eye diagram formed with a divided clock trigger.



Figure 4. A bit sequence captured using a pattern trigger.

Pattern Trigger: Used to display individual bits in the pattern. To view the whole pattern, the user must scroll through by altering either the time base of the scope or the pattern trigger position. Scrolling using the extent of a scope time base can lead to increased apparent jitter on the displayed signal due to weaknesses in time base circuitry; incrementing a pattern trigger, either within the BERT or within some more sophisticated scopes is preferable to avoid this problem. See Figure 4.




Figure 5. An incomplete eye diagram formed by triggering on data.




Figure 7. An eye diagram triggered from a clock recovered from the data signal using a wide loop bandwidth clock recovery scheme.

Triggering on Data is the least satisfactory method of constructing an eye and should only be used as a quick look-see. Long runs of identical characters provide no transitions to trigger from, and so a complete eye is almost impossible to achieve. See Figure 5.

Triggering from a Recovered Clock: While increasing complexity, this method has some advantages:

- In some situations a clock signal is not available and so must be derived
- In other cases, particularly long distance fiber optic communications, the relationship between the clock at the transmit end and the data at the receive end may be corrupted by rapidly time-varying effects in the transmission path.
- Lastly in cases where the receiver uses clock recovery and the eye as seen by the receiver needs to be examined – this can be required by some standards, particularly for jitter testing.

Circuits used for recovering clock typically have a loop bandwidth, or filtering function, that removes from the clock signal some of the jitter that was present on the data signal. Depending upon the measurement being made, this can be helpful or hurtful, but needs to be understood.




Figure 6. An eye diagram triggered from a clock recovered from the data signal using a narrow loop bandwidth clock recovery scheme.




Figure 8. An eye diagram triggered such that the delay between jittered clock and jittered data destructively interferes.

Narrow loop bandwidth clock recovery tends to give a rock solid clock trigger signal as the reference, and any jitter, or movement of edges with time, in the data eye diagram that is present will be displayed. This is a useful absolute measure but might not properly represent the jitter seen by a real system if the receiver uses clock recovery to track some of the jitter out. See Figure 6.

Wide bandwidth clock recovery tends to let more of the jitter that was present on the data signal through on to the clock. This can mean that as the data jitters by moving edges in one direction, then the other; the recovered clock tracks it, and the resulting eye appears to have very little jitter present on it. This tracking function is the way many system receivers work to reduce the jitter passed on through the system. See Figure 7.

Conditions can also conspire to create the opposite effect – where the delay between data signal and trigger signal is such that when the data edges are moving to their furthest extent in one direction, the recovered clock signal being triggered from it is moving to its furthest extent in the other, and the resulting eye shows as much as twice the jitter that was present on the data signal. See Figure 8.

While the last situation is seldom desirable, the first two situations, where all jitter is shown and where most jitter is tracked out, have their uses depending upon what information is being sought with the measurement. Luckily, most standards specify what triggering scheme is required to make measurements.

Eye Diagrams and BER

While eye diagrams provide an accessible and intuitive view of parametric performance, systems ultimately are judged on their ability to pass bits faithfully, and without error. The BER, Bit Error Ratio, or Bit Error Rate as it is sometimes called, is a ratio of the number of bits received incorrectly (errors) divided by the total number of bits received. This provides an overall score for how well a system is performing, but provides little help on why performance might be below expectations. It should be noted that BER tests logical problems as well as parametric ones – whether the correct bit was sent in the first place.

So why don't eye diagrams and BER easily link together? A perfect eye diagram would show all parametric aspects of all possible bit sequences, irrespective of how infrequently some effects show up. In other words, it would have a high information depth. Typically, eye diagrams are composed of voltage/time samples of the original data, acquired at some sample rate that is orders of magnitude below the data rate. For sampling oscilloscopes, this can be 105 samples per second at a 10 Gb/s (10¹⁰ bits/second) rate. This means that most eye diagrams are composed of shallow amounts of data.

This becomes a problem when issues arise that are infrequently occurring. These can be pattern related, noise related or deriving from other effects such as crosstalk and other forms of interference. These may not be visible in an oscilloscope eye diagram, but prevent link performance to desired levels. For example, links are often required to work to better than one error in a million, million bits (1×10^{-12} BER), while eye diagrams struggle to show events with probabilities below an occurrence of 1×10^{-5} .

This leaves an information gap that can be solved in a number of ways. The first is to construct the eye diagram from voltage/time measurements taken with an architecture that has a higher sampling efficiency. Systems such as the BERTScope are capable of perhaps three orders of magnitude more depth than sampling scopes.

The second solution is to use the fact that a BERT samples data directly at the line rate, and so is better equipped to see rare events that will limit system performance.




Figure 9. A receiver decision point in the center of the eye.

Slicing the Eye

Typical receivers are designed to make a decision at an instant in time as to whether the signal is above or below a particular threshold voltage. See Figure 9.

From this it decides whether the incoming signal is a data 1 or data 0. Sensible system designers place this decision point as far as possible from rising edges, falling edges, high level and low level – in other words in an unobstructed part of the eye, usually the center. Most BERT instruments have the ability to move this decision point away from the optimum position in time and/or voltage. By moving the decision point, it is possible to probe other parts of the eye, and by measuring the errors encountered, profile the anatomy of the eye diagram.




Figure 10 (a), (b), (c). Moving the BERT decision point in time to trace out the BER profile of the crossing point.




Figure 11 (a), (b). Moving the BERT decision point in voltage to trace out the BER profile of the one and zero levels.

Probing Jitter is a common test carried out with a BERT, done by probing through the crossing point of the eye. This measurement has been given various names including BERTScan, Bathtub Jitter, and Jitter Peak. It is more fully described in MJSQ¹. This measurement taken by a BERT has the advantage that the BERT sees every bit, and so is most likely to capture rare jitter events. See Figure 10.

Probing Signal-to-Noise: A second common eye test that arose out of the fiber optics industryⁱⁱ is Q-Factor. This is particularly useful in systems whose performance is limited by noise. Here the decision point is used to probe a vertical slice through the eye halfway through the bit period. Measuring the way the BER rolls off towards the middle of the eye gives an indication of how much noise is present and how it will impact the link. See Figure 11.




Figure 12 (a), (b). Slicing the eye along multiple axes to form an eye contour.




Figure 13. Block diagram of a reference receiver.

BER Contour is a superset of the last two measurements, and also the perfect linkage between BER and eye diagrams. Here the decision point is stepped around the inside of the eye and the BER profile mapped out on a series of angled slices around the eye. This provides a much more complete and clearer picture of where system parametric problems may be lurking. The BERTScope has been optimized to make this a fast and revealing measurement. See Figure 12.

The Need for Speed

The practicalities of these measurements are that a single point BER measurement of 1×10^{-12} at 10 Gb/s takes several minutes. Given that many systems are expected to function to better than 1×10^{-15} BER, where a single point could take months to measure, the ability to take quicker measurements and then extrapolate is valuable. The key is that BER contour measurements can be made quickly that are already down to near the levels where results are required and so the inaccuracies in extrapolating can be greatly minimized.

Testing Transmitters

Eye diagrams are commonly used for testing transmitters. As test equipment input characteristics vary, a standardized method of test, called a reference receiver, has been devised by international standardsⁱⁱⁱ such as the ITU for equipment used to test optical transmitters. Similar use of the input filter is also appearing now in some electrical standards. See Figure 13.




Figure 14. An eye mask test.

The aim is that the overall frequency response of the measuring system be well controlled, following a 4th order Bessel-Thompson curve within a defined tolerance window, with a systems -3 dB point at 0.75 of the bit rate. In theory, this means that measurements taken with different instruments should be comparable.

Mask testing is an abbreviated eye diagram test for the quick testing of transmitters in manufacturing. Rather than measuring all parametric aspects of the eye, mask testing defines key areas in the eye that are deemed to be “keep-out” areas – if any of the signal is detected to be in this region, then the device fails. Mask testing is often carried out in a few seconds. It should be noted that mask testing only detects gross problems in transmitters – it simply does not have the data depth in a short time to trap rare events, although instruments such as the BERTScope acquire significantly more data in a few seconds than most instruments. See Figure 14.

Testing Channels

Channels can be tested as an independent element, or as a combination of transmitter and channel. Testing of the channel as an independent entity is often carried out by measuring parametric characteristics such as loss, attenuation, reflection, and dispersion. This applies in optical links cases and in electrical links, although in short distance electrical cases such as back planes, s-parameter measurements are often used as an accurate description of all of the parameters mentioned above. One challenge with characterizing a channel as an independent entity is how to translate the measurements taken into what the eye and BER will be like at the end of a link. Modeling programs such as StatEye[®] attempt to convert parametric characterization into a predicted BER contour.

Another approach is to measure the channel with a representative transmitter. This has the disadvantage of presenting a challenge to de-convolve the contributions from transmitter and channel, but has the advantage of allowing direct measurement of eye characteristics and BER performance. This can be useful as a way of generating a measured BER contour that can then be compared with the modeled results of a program like StatEye.




Figure 15 (a), (b), (c). (a) Eye Diagram, (b) BER Contour, (c) Eye Bowl.

BER Contours in Three Dimensions – the Eye Bowl

A good visualization of the health of the eye, including the effects of rare events, can be obtained by plotting BER contour in three dimensions, as in Figure 15.

Testing Receivers

Traditionally receiver testing has been a BER test rather than an eye-related test. The input signal to the receiver is reduced in size and often impaired in other ways so that it represents the limit of the conditions that the device under test will see. If the receiver is able to function error-free then it passes.

An interesting twist on receiver testing is in measuring jitter tolerance, or the ability of the receiver to make the correct decision on each incoming bit even when the signal data edges are moving erratically in time. As discussed earlier, clock recovery will often remove much of this jitter. In the SONET/SDH world, jitter tolerance testing was achieved by deterministically moving the data edges in time using an injected sinusoid that was varied in amplitude and frequency according to a template. More recently, the concept of a Stressed Eye has been introduced, where the edges of the




Figure 16. A stressed Eye.

data are corrupted by various different means. This has developed beyond just a jitter test, as impairments are also added in the amplitude dimension, as shown elsewhere on this poster. The general idea is that the stressed eye is set to be the worst case operating conditions that the receiver would ever see, and if it is able to operate error-free in the presence of such a signal, the device passes. See Figure 16.

Conclusion

There can be an immense amount of information stored in an eye diagram if it is taken correctly, and to enough data depth. This can tell a designer a lot about the parametric performance of his design, and a manufacturing engineer whether parts will cause problems in the field later.

References

- i MJSQ - Methodologies for Jitter and Signal Quality Specification is a document written as part of the INCITS project T11.2. <http://www.t11.org/index.htm>
- ii N.S. Bergano, F.W. Kerfoot, and C.R. Davidson, "Margin Measurements in Optical Amplifier Systems," in IEEE Photonics Technology Letters, vol. 5, no. 3, pp. 304-306, Mar. 1993.
- iii TIA/EIA-526-4-A Optical Eye Pattern Measurement Procedure (<http://www.tiaonline.org/standards/>) Also measuring optical SDH transmitters: ITU G.957, Similarly for SONET transmitters: Bellcore GR-253-CORE
- iv Stateye modeling of channel responses to predict BER Contour: <http://www.stateye.org>

Contact Tektronix:

- ASEAN / Australasia (65) 6356 3900
- Austria* 00800 2255 4835
- Balkans, Israel, South Africa and other ISE Countries +41 52 675 3777
- Belgium* 00800 2255 4835
- Brazil +55 (11) 3759 7600
- Canada 1 (800) 833-9200
- Central East Europe, Ukraine and the Baltics +41 52 675 3777
- Central Europe & Greece +41 52 675 3777
- Denmark +45 80 88 1401
- Finland +41 52 675 3777
- France* 00800 2255 4835
- Germany* 00800 2255 4835
- Hong Kong 400-820-5835
- India 000-800-650-1835
- Italy* 00800 2255 4835
- Japan 81 (3) 6714-3010
- Luxembourg +41 52 675 3777
- Mexico, Central/South America & Caribbean 52 (55) 56 04 50 90
- Middle East, Asia and North Africa +41 52 675 3777
- The Netherlands* 00800 2255 4835
- Norway 800 16098
- People's Republic of China 400-820-5835
- Poland +41 52 675 3777
- Portugal 80 08 12370
- Republic of Korea 001-800-8255-2835
- Russia & CIS +7 (495) 7484900
- South Africa +27 11 206 8360
- Spain* 00800 2255 4835
- Sweden* 00800 2255 4835
- Switzerland* 00800 2255 4835
- Taiwan 886 (2) 2722-9622
- United Kingdom & Ireland* 00800 2255 4835
- USA 1 (800) 833-9200

*** If the European phone number above is not accessible,
please call +41 52 675 3777**

Contact List Updated 25 May 2010

For Further Information

Tektronix maintains a comprehensive, constantly expanding collection of application notes, technical briefs and other resources to help engineers working on the cutting edge of technology. Please visit www.tektronix.com



Copyright © 2010, Tektronix. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.

09/10 EA/WWW

65W-26042-0

