



Stressed Eye Primer

Primer

Table of Contents

Abstract	3	Adding the Components Together	14
Jitter Tolerance Testing of Receivers — The Basics	3	Conclusion.....	15
The Job of the Receiver	3	References.....	15
The Effect of Jitter	5	Appendix 1: Stressed Eye Standards	
How Clock Recovery Can Help	6	Requirements	15
Testing a Clock Recovery Circuit with		Notes.....	15
Sinusoidal Jitter.....	7	Electrical Stressed Eye Standards.....	16
Sinusoidal Jitter Templates.....	7	Optical Stressed Eye Standards.....	17
The SONET/SDH Way	7	Appendix 2: The Jitter Tree and Stress	
Jitter Tolerance in More Recent Standards	8	Components	18
Impairment Types.....	8		
A Brief Note on BER Contour.....	8		
A Clean Eye	9		
Sinusoidal Jitter.....	9		
Sinusoidal Interference.....	10		
Random Jitter	11		
ISI	13		
BUJ	13		



Figure 1. Conceptual picture of an ideal receiver.




Figure 2. Simplified Receiver Block Diagram.

Abstract

Many high-speed serial interface standards call for a test known as ‘Stressed Eye.’ This paper is an introduction to stressed eye testing, some of the high-speed standards that use it, and how a receiver test using stressed eye is constructed.

Jitter Tolerance Testing of Receivers — The Basics

The Job of the Receiver

The formal name of stressed eye testing is ‘receiver jitter tolerance test’— the ability of a high speed electrical or optical receiver to deal with an impaired incoming signal (with jitter and crosstalk on it) and not make a mistake. In other words, how bad can the incoming signal be and not cause the receiver to misjudge a one to be a zero, or vice-versa, creating errors.

Conceptually most digital communication links are the same, whether they are trans-oceanic fiber optic links or high-speed connections across a circuit board or backplane. A transmitter puts out a good clean signal; the signal passes through a hostile environment that corrupts it in various ways, and the receiver’s task is then to reconstitute the original stream of ones and zeros without making a mistake.

At speeds above 1 Gb/s, it is very common that the data signal being transmitted has a clock signal encoded within it. In order for the receiver to correctly interpret the incoming data, it needs a timing reference that is related to the data. As the hostile environment that the data signal passes through can alter the timing of the signal in unexpected ways, deriving the timing signal from the incoming data makes sense. An oversimplified block diagram of a receiver is shown in Figure 2. While it omits amplifiers and other components, it is useful to get some relevant ideas across.

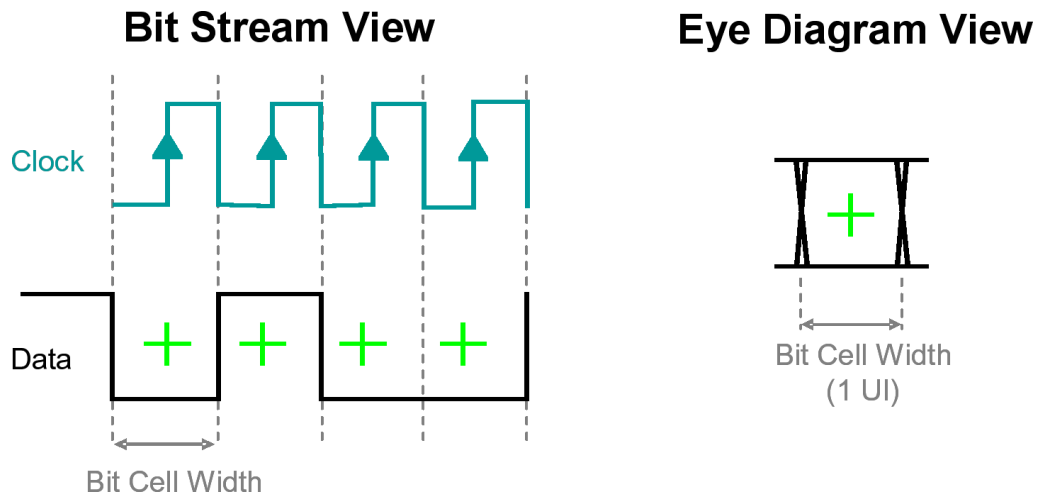


Figure 3. Clock is used to time the execution of a decision on the incoming data between a one and a zero (left). Also shown in eye diagram form (right).

The clock recovery circuit derives a data-related clock signal from the data. The aim is to use this clock to time the execution of decisions on the incoming data — for each bit time slot, whether the signal should be a one or a zero. The decision circuit will look at an instant in time, and at that instant decide whether the data is above or below a threshold point. From this, the circuit will produce a new, clean one or zero, respectively. This is shown in Figure 3. The clock recovery circuit of Figure 2 produces a clock signal, and the delay is then used to position the rising edges of the clock signal in the correct position within the eye to time the decision (where a green cross shows the time instant where a threshold decision is taken).

Also shown in Figure 3 is the picture in eye diagram form — where the bit sequences of the incoming data, such as 0-0-0, 0-0-1, 0-1-0, 0-1-1, etc., are overlaid on a single bit cell to show a combined picture of the system's signal integrity (right side of the figure).

The process described above assumes idealized data coming in. Obviously the real world is not like this, and the signal that reaches a typical receiver is usually of smaller amplitude than needed, contains noise, and has other impairments. We will look at some of these in more detail later. So what is jitter tolerance?




Figure 4. A jittered data stream with some edges moving from their ideal position.

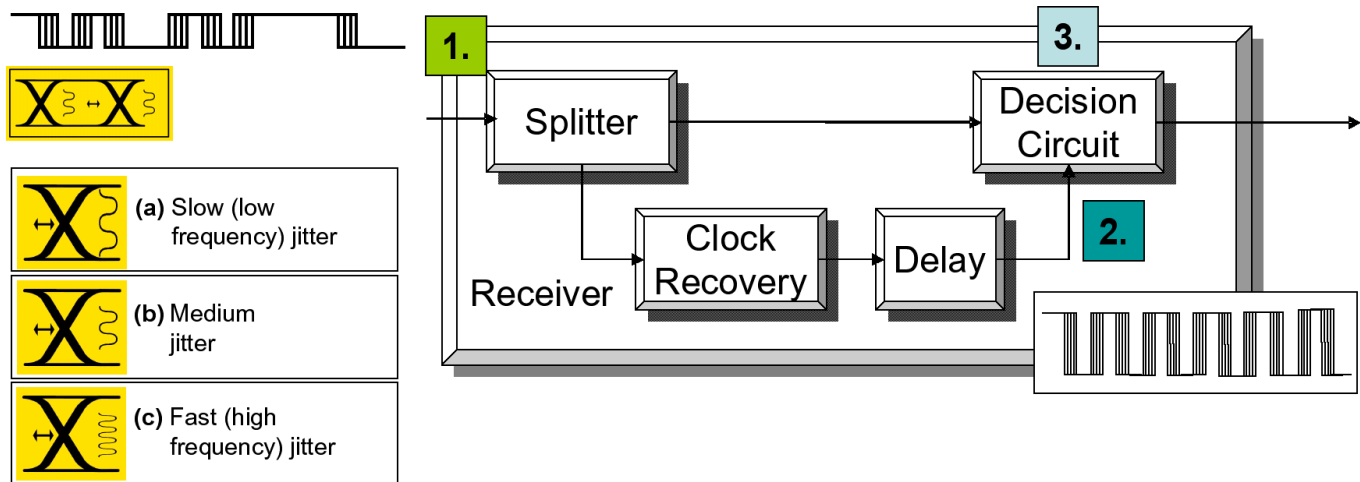


Figure 5. The effect of sinusoidal jitter on a receiver with clock recovery.

The Effect of Jitter

Jitter is often explained as the non-ideal placement of data edges with respect to a clock. Jitter is a complex subject, with many pages written about it elsewhere — we'll keep it simple for this discussion. Figure 4 shows the data signal of Figure 3 undergoing jitter — some edges in the data stream have moved from their ideal position. If this were to happen, it is easy to see that some of those edges might start to interfere with the decision making process, with mistakes increasing the more the edges move across the bit cell.

A major advantage of using encoded clock and then recovering it from the data is that the clock recovery can be used to produce a clock that tracks, or moves with, the data signal even as it moves around in time. It is intuitive that if there is an overall trend for all edges to slowly move around together in time, that the clock recovery delivering a related clock signal could track this slow movement. It is also intuitive that if the movement of edges were very fast or completely erratic bit-to-bit, that a clock recovery circuit would have trouble tracking this. This leads to the idea that a clock recovery circuit can act as a jitter filter, with the ability to track slow-moving jitter up to some limit, resulting in a data stream that has reduced jitter related to the recovered clock. Let's look at how this works.

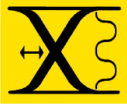
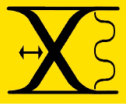
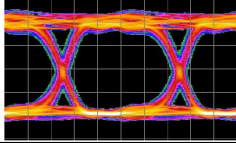
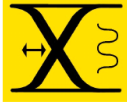
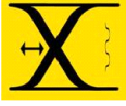
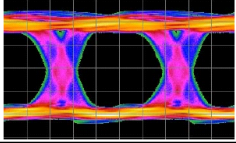
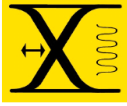

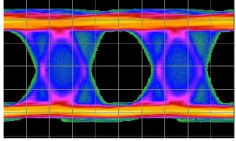
Applied Jitter Frequency 1.	Recovered Clock 2.	Recovered Clock used to make decision on incoming Data: Eye observed at the decision point..... 3.
Low 	 <p>Recovered clock tracks data – as data moves (jitters), clock does too.</p>	<p>Clock and data track – waveform has almost no apparent jitter</p> 
Medium 	 <p>Recovered clock trying to track data, but not completely successfully – jitter attenuated</p>	<p>Clock and data track to some extent – increase in apparent jitter</p> 
High 	 <p>Recovered clock completely unable to track movement (jitter) of data., jitter completely attenuated by PLL BW</p>	<p>Clock is stable, data jitters, so all jitter on data is evident.</p> 

Figure 6. The effect of different frequencies of applied sinusoidal jitter.

How Clock Recovery Can Help

Using the receiver of Figure 2, we will apply a modulation to the edges of the data signal. This will have the effect of moving the edges in one direction, for example making them late an increasing amount, then bringing them back in line again, and then making them early for a time. This will happen over and over again, in a manner that is related to the modulation signal generator we choose for this experiment (rather than in any way related to the pattern of ones and zeroes of the data stream). For this example, we’re going to use a sine wave of different frequencies.

Figure 6 shows the receiver under three different conditions of applied sinusoidal jitter. The first column of the table examines the jittered signal appearing at the input to the receiver (1). The second shows what the clock recovery circuit (often a Phased Locked Loop, or PLL) produces (2). The third column shows

the signal appearing at the receiver decision point. This may not be intuitive. For the slowly varying edges, the clock that is recovered is an exact mirror of the data, with the edges moving at the same rate and by the same amount as the data edges. The net effect is that, each instant a decision is made, the decision circuit sees a bit cell perfectly synchronized with the clock strobe, and the decision is made correctly every time. At high frequencies, beyond the ability of the PLL to cope, it attenuates the jitter present on the recovered clock. The data edges come in heavily jittered, but the clock does not move in sympathy with the data at the decision circuit, and so the net effect seen by the decision circuit is edges that are moving in time with respect to the clocking signal. This obviously increases the chance of errors.

Note that the amount of jitter on the incoming data did not change – only the frequency of the jitter did.

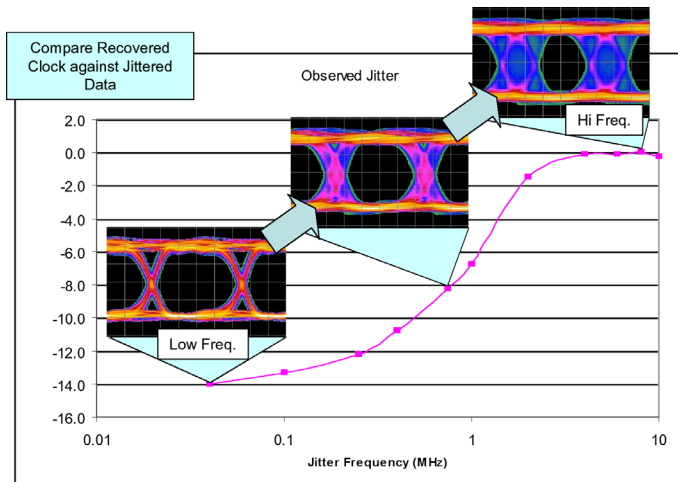


Figure 7. Jitter transfer function graph showing the ratio (dB) of the jitter magnitude out divided by the jitter magnitude in. The left side shows reduced recovered clock to data jitter, due to the recovered clock tracking the jitter in the data.

Testing a Clock Recovery Circuit with Sinusoidal Jitter

At 10 Gb/s, typical clock recovery circuits switch over from tracking the jitter out, to fully passing the jitter on, with a transition in the region of 1 – 4 MHz. The graph in Figure 7 shows the transfer function of one such clock recovery unit.

Sinusoidal Jitter Templates

We've seen that the clock recovery circuit in a receiver can have a big effect on the effectiveness of the receiver in dealing with jitter. Receivers are usually tested for jitter tolerance by applying a jitter modulation of different frequencies and amplitudes to the incoming data signal, and counting when errors occur. Jitter tolerance testing thus requires a BERT error detector and some kind of impaired pattern source. The testing often uses a sinusoidal modulation of the data signal transitions, with instructions on how much modulation should be applied in the form of a jitter template graph. An example is shown in Figure 8. At low jitter frequencies, a receiver is often expected to endure large amounts of jitter — many unit intervals (UI), enough to close the eye completely many times over. At these low frequencies though, the clock recovery is expected to completely remove the effect of the jitter. At higher frequencies, the amount of jitter applied is less than 1 UI.

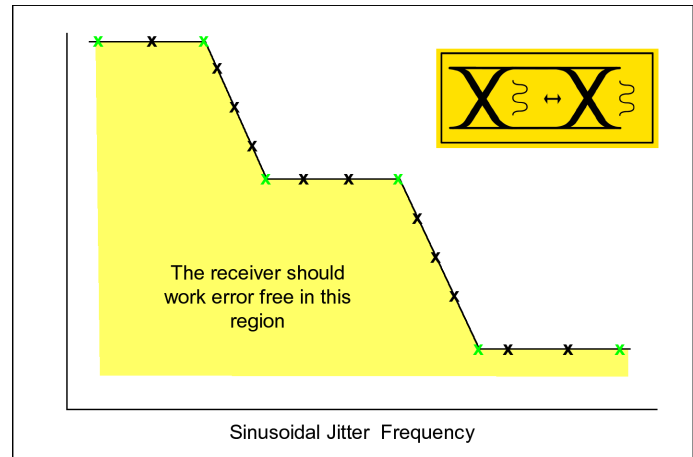


Figure 8. Sinusoidal jitter tolerance template.

When testing, a number of discrete points is chosen on the template, usually a compromise between confidence that the device will comply, and minimizing the test time. Usually more attention is paid to the areas where the template gradient changes (the green crosses in Figure 8).

The SONET/SDH Way

The jitter tolerance method described above, using sinusoidal jitter stepped through a template, was pioneered by the SONET/SDH world. The original implementations of such networks were to have a transmit end of the link and a receive end, and repeaters in between. Each repeater converted the optical signal back to the electrical domain, cleaned it up and amplified it enough to drive the conversion back to the optical domain and the distance to the next repeater (this type of regeneration is less common following the invention of fiber amplifiers that reduced the need for the conversions between optical and electrical domains along the way). In such 'repeated' systems, signal-to-noise was the biggest impairment to system performance, but it was essential that repeaters didn't have the effect of amplifying jitter along the way, which is very possible in such systems.

Jitter Tolerance in More Recent Standards

While it made sense in a repeatered SONET/ SDH environment to measure jitter tolerance with a sinusoidal jitter template, many of the more recent standards have been tackling communications links of a very different sort. One thing that they have in common is that such systems, whether optical, electrical cabling, or across backplanes and circuit boards, is that signal-to-noise is often not the limiting factor in performance. This has led to the creation of more complex recipes for the impairments that are introduced to a receiver being tested. Stressed eye testing usually imposes restriction around the entire eye, not just the time axis, and so is not just a jitter test anymore. As discussed later, a combination of some or all of the following impairments are used to create the stressed signal.

The first standards to use stressed eye were Gigabit Ethernet and 1x, 2x and 4x Fiber Channel between 1998 and 2003. During this period, stressed eye was adopted and modified by the emerging 10 Gigabit Ethernet standard (2002). Since then, different recipes of stressed eye have been adopted in many standards, including high speed buses such as PCI Express and Serial ATA. A table of recipes for common standards is given in Appendix 1. A common theme for many standards is the amount of inferring and guesswork required to interpret what is necessary.

Impairment Types

A Brief Note on BER Contour

The new stressed eye recipes are intended to reflect the likely impairments that a receiver may be faced with. In the following list, the impairments are illustrated with real eye diagrams and BER contours. More detail of BER contours is obtainable from SyntheSys Research, but briefly the contours are like the contours on a geographic map, showing the steepness of mountains. In an eye diagram, the steeper the slopes, the less likely edges are going to intrude on the decision point and cause errors. An advantage of BER contour is that it goes much deeper than conventional oscilloscope eye diagrams in revealing the infrequent events, such as random jitter, that must be controlled for error-free operation to be achieved. In the following examples the outermost (blue) contour shows where the 1×10^{-6} boundary is — that is, if the receiver decision point were to be moved to somewhere along that line, the receiver would achieve a BER of 1×10^{-6} . The innermost (white) contour is, usually, a projected line showing where 1×10^{-16} operation would lie. Placing a receiver decision point inside this boundary is assumed to allow error-free operation. Obviously the more area that lies within the inner boundary, the healthier the eye is considered to be.




Figure 9. Unimpaired Signal displayed as an Eye Diagram (left) and a BER Contour (right) where the outermost (blue) contour represents 1×10^{-6} BER, and the innermost (white) contour represents 1×10^{-16} BER. In this example, the BER contours are tightly spaced because there is little random jitter or noise present in the signal.

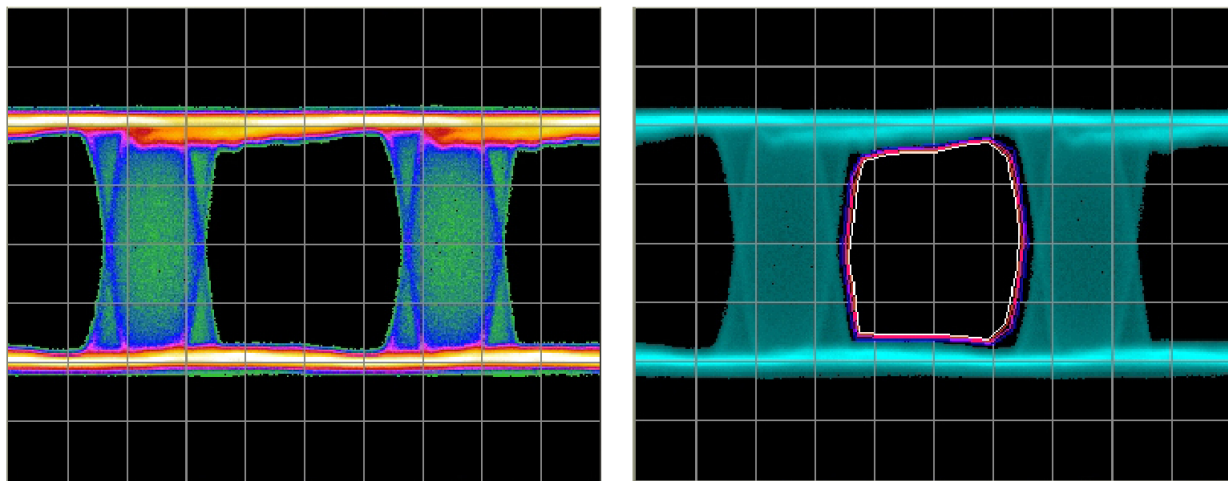




Figure 10. Sinusoidal Jitter shown on an Eye Diagram (left) and BER Contour (right).

A Clean Eye



For reference, Figure 9 shows two views of a clean eye to compare against. This whole sequence is illustrated in more detail on the poster “Anatomy of an Eye Diagram”ⁱⁱ, including other analysis views that give a more complete picture of the effect of each stress type.

Sinusoidal Jitter



As with SONET, sinusoidal jitter is used to verify the jitter tolerance of the clock recovery function in a receiver. Templates similar to those used in SONET are often employed.

It is obvious from the eye diagram that the eye undergoes closure in the time axis, which is the aim in simulating jitter. It is useful to note that the BER contour shows very tightly grouped contours. This means that all of the closure is obvious in the eye, with no infrequent events lurking to catch out the unwary. The jitter modulation is entirely deterministic, completely bounded.

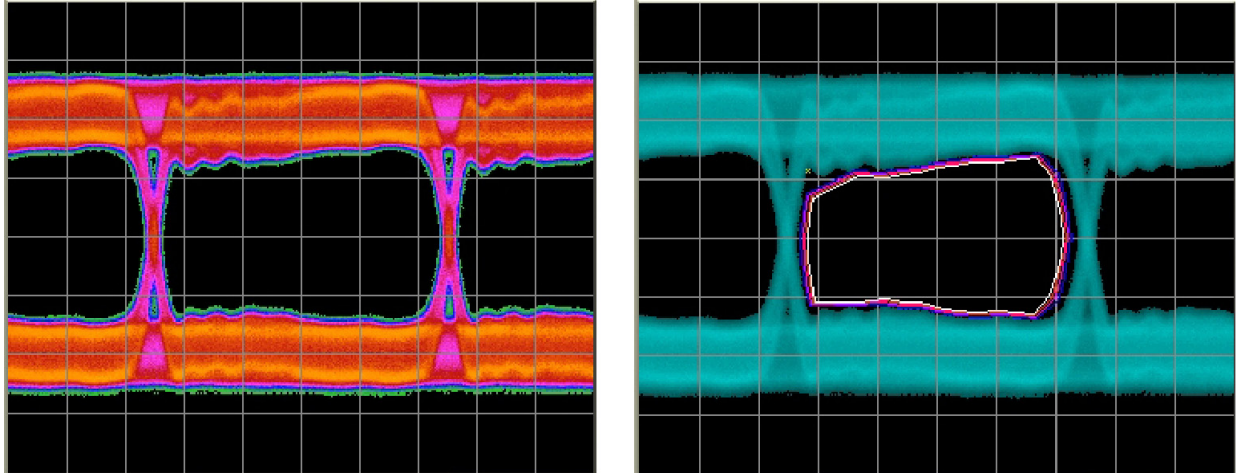


Figure 11. Sinusoidal Interference shown on an Eye Diagram (left) and BER Contour (right).

Sinusoidal Interference



Used to close the eye down on the amplitude axis, this simulates low signal level or other deterministic, bounded impairments, such as poor matching leading to reflections.

The result is similar to using sinusoidal modulation to close down the time domain. The eye looks closed in the amplitude axis, but the BER contour shows the impairment to be entirely well-behaved, leaving plenty of area of operation clear for the receiver decision point.





Figure 12. Random Jitter shown on an Eye Diagram (left) and BER Contour (right). Conventional scope-based eye diagrams give extremely limited visibility to low probability events like random jitter.

Random Jitter



Random jitter is intended to simulate the noise-related effects in the link, such as oscillator phase noiseⁱⁱⁱ. Real-world

causes of random jitter can be any of the following:

- Most often this can be from low cost PLL-synthesizer parts used as clock sources.
- Broadband RJ is not much of an issue in copper links, but strong in optical links through use of a post-amplifier after the receiver.
- Also in optical links, at the optical receiver input there can be a blend of low frequency jitter mixed with high frequency laser and laser/fiber noise. The high frequency transmitter jitter may be filtered in long cable plants.
- 850-nm links based on VCSELs commonly exhibit significant levels of RJ.
- Amplifier noise adds noise to data edges, with the consequence that the slower an edge, the more conversion there is from amplitude to time domain.

Looking at the eye diagram of Figure 12, the eye looks fairly open, with little to worry about. The reality the receiver will experience is shown in the BER contour. This is the first example here of an impairment that has rarely occurring events. This is known as an unbounded process — it is not deterministic, in that imposing the same conditions such as a particular pattern sequence does not necessarily yield the same result in BER terms. In this case, there is a finite probability of an edge appearing in any position across the eye and causing errors. Largely spaced contours such as those shown in Figure 12 are a warning sign. For this reason, if RJ is called for in a recipe, it is important to set it up carefully, preferably using a BER contour or deep jitter measurement to be sure the amount is correct.

Typical test setups will have some inherent RJ anyway (they are composed of active components). Some standards specify adding no more RJ than is already present, others require extra, and this usually means adding amounts until the combination of inherent and added RJ reaches the correct level.




Figure 13. ISI shown on an Eye Diagram (left) and BER Contour (right), caused by a 4th Order Bessel – Thompson Filter with a 3 dB Point at 0.75 of the bit rate, stimulated with a PRBS 2⁷-1.




Figure 14. ISI shown on an Eye Diagram (left) and BER Contour (right), caused by passing the signal through a length of coax with frequency dependent loss (used beyond its designed operating frequency), stimulated with a PRBS 2⁷-1.




Figure 15. ISI shown on an Eye Diagram (left) and BER Contour (right), caused by passing the signal through a high-speed backplane, stimulated with a PRBS 2³¹-1. Conventional eye diagrams cannot see many deterministic effects in a pattern over 2 billion bits in length.





Figure 16. Bounded Uncorrelated Jitter shown on an Eye Diagram (left) and BER Contour (right).

ISI



Inter-Symbol Interference is one contributor to Data-Deterministic Jitter (DDJ). As the name implies, this type of jitter is pattern-related. It simulates bandwidth limitations or other frequency response problems in the system being tested. Sometimes standards are not specific about how the ISI is to be created for a stressed eye, giving only the amount that is to be imposed. Figures 13, 14, and 15 show different ways of achieving ISI, respectively using a filter, a length of coax, and a backplane.


In each case, the effects are entirely deterministic — that is, using the same data patterns in the same way will produce the same eye closure every time. A few things to note:

Figure 13 uses a special kind of filter known as phase linear, producing the symmetrical waveform shape shown. Figure 14, using a poor-quality piece of coax, is asymmetric and shows a capacitive effect. In both cases, the pattern used is a 2^7-1 PRBS — a very short pattern at 127 bits long, and one that doesn't contain particularly aggressive bit sequences. Every single pattern effect is captured in the eye diagrams, with no deterministic rare events lurking, which is why the BER contours of both of these examples are tightly grouped.

Figure 15 is different. The general shape is a variation on a similar theme to Figures 13 and 14, produced here by a backplane. However, a difference becomes immediately clear in the BER contour. The reason for this is that the pattern used was a $2^{31}-1$ PRBS, a much longer pattern than was used in the other two examples. This pattern has much harder pattern sequences, many of which only occur once

per pattern repetition. For a pattern that is over 2×10^9 bits long, this means that sequences that cause edges to move the most in time will not be seen with a conventional sampling oscilloscope. This is important, as many standards require use of long patterns (such as $2^{23}-1$ and $2^{31}-1$ PRBS) to do stressed eye testing. Note that the BER contours in Figure 15 look like they might be caused by some random process such as the RJ case of Figure 12. This isn't quite true — pattern-related effects are bounded. As long as you are able to see all effects in the measurement you are making (such as a BER contour run for long enough) then you will capture the latest edge that will occur. Using a shallower view, infrequent event pattern effects can appear like RJ.

BUJ



Bounded, Uncorrelated Jitter (where the word 'uncorrelated' refers to 'unrelated to the data pattern') is used to simulate the effects of interfering signals. It is usually introduced by modulating the data edges with a short PRBS, somewhere between 2^7-1 and $2^{11}-1$, depending upon the standard. Typical interfering signals can be crosstalk from adjacent channels, which is likely to be at high frequency. It can also be from switching power supply noise, which tends to be in the range 100 kHz to 1 MHz. The standards tend to specify the pattern length, filtered bandwidth, and frequency of modulation required.

The measurement in Figure 16 was made using a 2^7-1 PRBS modulation of the data edges. Because the patterns used are all short and the effect bounded, you would expect to capture the full edge excursions in a typical eye diagram.




Figure 17. A stressed eye mask (left).

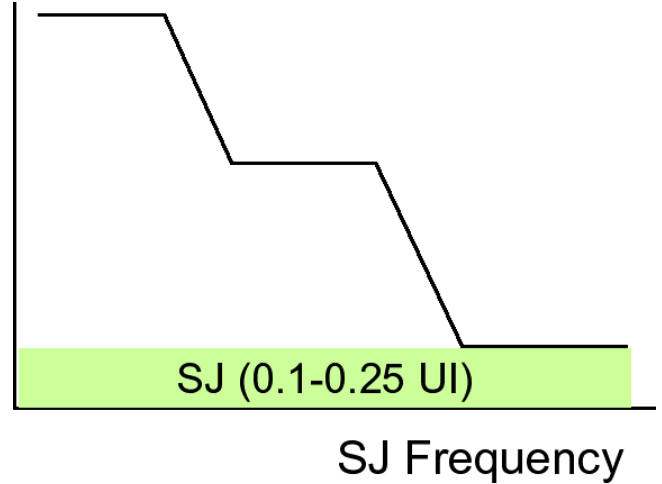


Figure 18. SJ Template (right).

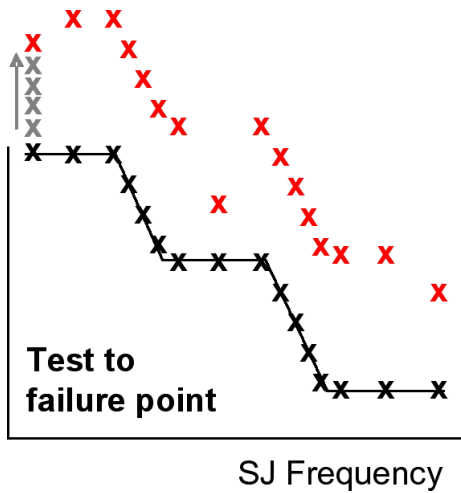


Figure 19. Testing beyond the template to explore device margin.

Adding the Components Together

Typically a stressed eye recipe will have two parts — a sinusoidal jitter template such as the one we saw in Figure 8, used to step through the different sinusoidal modulation frequencies and amplitudes, and an eye mask such as is shown in Figure 17. The idea is that the different components specified are added in to close the eye down (blue area) to the diamond, leaving an assured area of error-free operation in the center (white area within the diamond). The receiver is expected to operate successfully in this small area, despite the impairments.

Setting up components to the template can be achieved in a couple of different ways, depending on the standard. Standards such as XFP require the recipe to construct the eye to the template without SJ being part of it. Once the template is complied with, the SJ template is stepped through. The other approach, as used by standards such as 10 GbE and XAUI, is to construct the stressed eye to the template including SJ, set to the value it has at its highest frequency (see green segment of Figure 18), typically somewhere between 0.1 and 0.25 UI. The other components are then added in to match the mask template, and then the SJ template is stepped through.

In both cases, the aim at the end is to have introduced the stressed eye stimulus to the receiver, step through the SJ template and to count errors. Often engineers will use the template as a basis, but for a given frequency on the template, will step up in amplitude until they hit the limit of performance to assess device margin (or step downwards to explore where a device failed compliance). See Figure 19. Similarly, margin can be added in the eye template by making the diamond smaller.

Conclusion

We've walked through the history and makeup of jitter tolerance testing using a stressed eye. We've seen that stressed eye testing has evolved, appearing in many diverse high speed standards for receiver testing, often exercising the receiver in time for more than just a sinusoidal stimulus, and also the amplitude dimension. Some of the common reasons why standards choose particular stress components have also been explored, along with the construction of a typical stressed eye to a template.

References

- i "Bridging the Gap—A BER Contour Tutorial" Literature Number SR-TN032, www.tektronix.com.
- ii "Anatomy of an Eye Diagram" Poster, www.tektronix.com.
- iii "DJ in stressed eye tolerance test" Tom Lindsay Comment E2O #128 MJSQ, April 2004

Appendix 1: Stressed Eye Standards Requirements

Some common standards are interpreted in the following two tables. We have divided them into optical and electrical requirements, with standards such as Fiber Channel appearing in both.

As stated elsewhere in this paper, interpreting standards in this area is often not straightforward. The tables were correct to the best of our belief when this was written (April 2005).

Notes


Optical: SM = Single Mode, MM = Multimode

Electrical Stressed Eye Standards



	SJ Max	Hi Frequency SJ	DJ	RJ	BUJ	TJ	Notes
Explanations:	A Template is usually used to step through different frequencies and amplitudes - this column is indicating the maximum excursion seen at any frequency on the template.	The level often used to build the eye before stepping through the SJ template.	Sometimes specified to be a specific filter such as a 4th Order Bessel Thompson linear phase filter with a 3dB point set at 0.75 of the bit rate - other times it's not specified, and can be a dispersive cable or any element that provides the correct closure.	Often a contribution that is inferred rather than specified.	Usually one of the standard PRBS polynomials at a defined frequency.	Usually all components including High Frequency SJ, but not other SJ values from a template.	
OIF-CEI	6G: 5 UI 11G SR: 5 UI	6G: 0.05 UI 11G SR: 0.05 UI	"Correlated, bounded, high probability jitter" 6G: 0.45 UI 11G SR: 0.2 UI	6G: "Gaussian Jitter": 0.15 UI 11G SR: "Gaussian Jitter": 0.2 UI	6G: None 11G SR: 0.25 UI	6G: 0.6 UI 11G SR: 0.7 UI	
XFP: XFI interface (Electrical)	Telecom: 15.2 UI Datacom: 5.763UI	0.05 UI	ISI: 0.2 UI min	DCD + RJ + PJ = 0.2 minimum* (*PJ = SJ HF, 0.05 UI)	None at the moment - proposal to add is being discussed.	TJ max B: 0.65 UI D: 0.61 UI TJ - ISI max ("Non-DDJ" Jitter) B: 0.45 UI D: 0.41 UI	Optical ports of XFP module is addressed by relevant Telecom & Enterprise standards. Electrical side (XFI) has different SJ specs depending upon application (telecom or datacom). B = input to XFP electrical receiver D = input to external chip receiver
PCI-Express (2.5 Gb/s)			0.48 UI (193 ps)	0.01 UI (5.4 ps)			Implied receiver tolerance
S-ATA 1i & 1m			5 UI: 0.25 250 UI: 0.35* * Total value made up from variable amounts of PJ, DCD, DDJ	Inferred RJ: 5 UI: 0.18 UI 250 UI: 0.25 UI	None	5 UI measurement: 0.43 UI 250 UI measurement: 0.6 UI	
S-ATA 1x & 2x			0.35 UI	Inferred RJ: 0.3 UI		0.65 UI	$f_{BAUD}/1667$
S-ATA 2i & 2m			0.35 UI	Inferred RJ: 0.11 UI		0.46 UI	$f_{BAUD}/10$
			0.42 UI	Inferred RJ: 0.18 UI		0.60 UI	$f_{BAUD}/500$
Fiber Channel							
2X	All: 1.5 UI	All: 0.1 UI min	Inter (Gamma) Elec: 0.37 UI Inter (Beta) Elec: 0.33 UI Intra (Delta) Elec: 0.39 UI	Inter (Gamma) Elec: 0.3 UI Inter (Beta) Elec: 0.29 UI Intra (Delta) Elec: 0.3 UI	Implied as part of DJ budget	Inter (Gamma) Elec: 0.67 UI Inter (Beta) Elec: 0.62 UI Intra (Delta) Elec: 0.69 UI	Intra (Delta & Beta) (inside a box) - Delta E & O, Beta E only Inter (Gamma) (Between Elements) SERDES (Alpha) (4G FC test for Rx in SERDES) (RJ Values are Inferred)
4X	All: 1.5 UI	All: 0.1 UI min	Inter (Gamma) Elec: 0.37 UI Inter (Beta) Elec: 0.33 UI Intra (Delta) Elec: 0.39 UI SERDES (Alpha) 0.41 UI	Inter (Gamma) Elec: 0.3 UI Inter (Beta) Elec: 0.29 UI Intra (Delta) Elec: 0.3 UI SERDES (Alpha) 0.3 UI	Implied as part of DJ budget	Inter (Gamma) Elec: 0.67 UI Inter (Beta) Elec: 0.62 UI Intra (Delta) Elec: 0.69 UI UI SERDES (Alpha) 0.71 UI	

Optical Stressed Eye Standards



	SJ Max	Hi Frequency SJ	DJ	RJ	BUJ	TJ	Notes
Explanations:	A Template is usually used to step through different frequencies and amplitudes - this column is indicating the maximum excursion seen at any frequency on the template.	The level often used to build the eye before stepping through the SJ template.	Sometimes specified to be a specific filter such as a 4th Order Bessel Thompson linear phase filter with a 3dB point set at 0.75 of the bit rate - other times it's not specified, and can be a dispersive cable or any element that provides the correct closure.	Often a contribution that is inferred rather than specified.	Usually one of the standard PRBS polynomials at a defined frequency.	Usually all components including High Frequency SJ, but not other SJ values from a template.	
10GbE-LX4 (4 Channels of Optical)	5 UI	Somewhere in .05 to .15 UI range that suits	ISI: BT Filter specified	Inherent RJ of optical transmitter only	None	0.3 UI inc HF SJ 0.25 UI w/o HF SJ	Standard specifies use of BT ISI Filter. Optical Tx also has inherent jitter. With all of these components, make eye to fit template.
10GbE serial (Optical)	5 UI	Somewhere in .05 to .15 UI range that suits	ISI: BT Filter specified	Inherent RJ of optical transmitter only	None	0.3 UI inc HF SJ 0.25 UI w/o HF SJ	Standard specifies use of BT ISI Filter. Optical Tx also has inherent jitter. With all of these components, make eye to fit template.
Fiber Channel							
2X	All: 1.5 UI	All: 0.1 UI min	Inter (Gamma) SM: 0.28 UI MM: 0.29 UI Intra (Delta) SM: 0.39 UI MM 0.39 UI	Inter (Gamma) SM: 0.25 UI MM: 0.24 UI Intra (Delta) SM: 0.3 UI MM 0.3 UI	Implied as part of DJ budget	Inter (Gamma) SM: 0.53 UI MM: 0.53 UI Intra (Delta) SM: 0.69 UI MM: 0.69 UI	Intra (Delta & Beta) (inside a box) - Delta E & O, Beta E only Inter (Gamma) (Between Elements) SERDES (Alpha) (4G FC test for Rx in SERDES)
4X	All: 1.5 UI	All: 0.1 UI min	Inter (Gamma) SM: 0.28 UI MM: 0.29 UI Intra (Delta) SM: 0.39 UI MM: N/A SERDES (Alpha) 0.41 UI	Inter (Gamma) SM: 0.25 UI MM: 0.24 UI Intra (Delta) SM: 0.3 UI MM: N/A SERDES (Alpha) 0.3 UI	Implied as part of DJ budget	Inter (Gamma) SM: 0.53 UI MM: 0.53 UI Intra (Delta) SM: 0.69 UI MM: N/A SERDES (Alpha) 0.71 UI	(RJ Values are Inferred)
SONET							
OC-48	106.4 UI	0.15 UI	N/A	N/A	N/A	N/A	N/A
OC-192	446.6 UI	0.15 UI	N/A	N/A	N/A	N/A	N/A

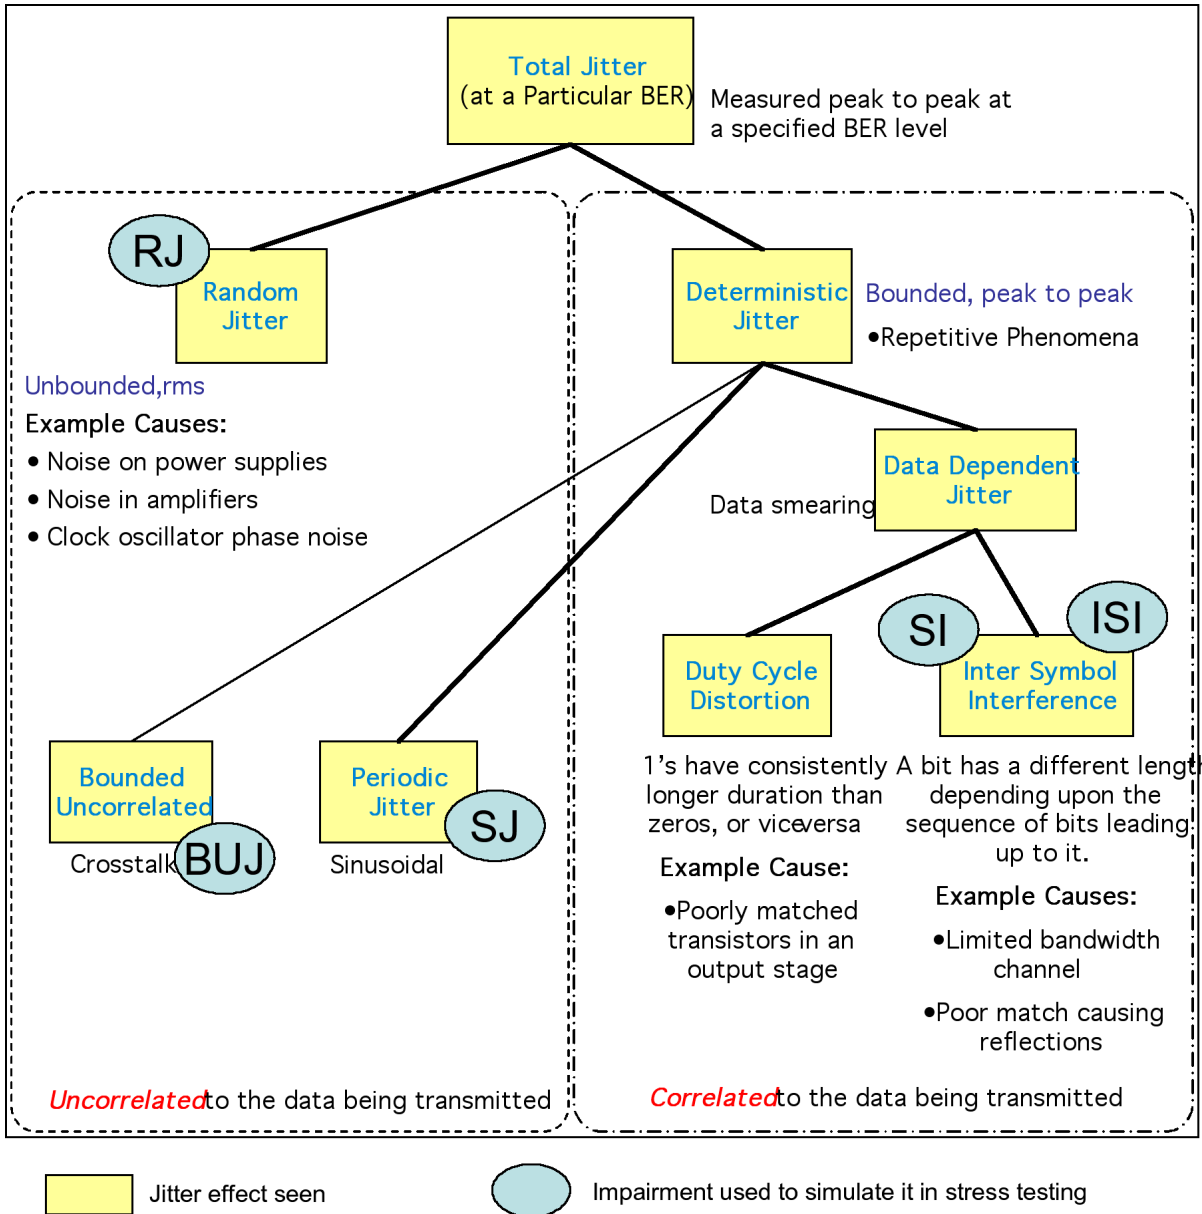


Figure 20. Mapping stress component types onto the enterprise jitter tree model.

Appendix 2: The Jitter Tree and Stress Components

Enterprise jitter is often illustrated in the form of a tree of sub-components.

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-26049-0

