6 Hints for Better SATA and SAS Measurements

Application Note

SATA and SAS testing can generally be summarized into four test groups:

- Transmitter
- Receiver
- Impedance and return loss
- Protocol

Introduction

Serial ATA (SATA) and Serial Attached SCSI (SAS) are data-transfer technologies for moving data to and from storage devices. The Serial ATA International Organizations (SATA-IO) and Serial Attached SCSI (SAS) T10 standard committees have both released specifications to provide support for 6 Gbps data transfer and will soon be releasing higher data rates up to 12 Gbps. As the speed increases, the validation effort increases exponentially as well. In order for the storage interface to function properly, its signal integrity performance must meet certain minimum requirements. Signal integrity is the key to system interoperability, or the guarantee that products from different vendors will integrate well, when used together. Failures in signal integrity are correlated to other failures, including marginal timing relationships, protocol violations, jitter issues, and errors from other buses. This application note will highlight a few of the methods that can help you characterize, validate and debug your designs faster.

The following six hints take this expert knowledge and provides you with keys to accurate, consistent, and efficient measurements.

Test group	Hint	Description
Transmitter	Hint 1	Remove the loss of your test setup for more accurate measurement
	Hint 2	Quickly find and fix the root cause of jitter failures
Impedance and return loss	Hint 3	Assess the insertion loss and return loss of your cable, connector and chip
Receiver	Hint 4	Test beyond the receiver compliance coverage to avoid any hidden problems
	Hint 5	Characterize the robustness of your receiver to real world spread spectrum clock (SSC) profiles
Host/Device protocol	Hint 6	Inject protocol errors to see if your product can recover

Table 1. Test group and hints

Agilent Technologies

Hint 1 – Remove the loss of your test setup for more accurate measurement

Did you know that your test fixture and cable have an impact on your measurement? In a SATA and SAS test environment, cables and fixtures are lossy and this can mean inaccurate and non-repeatable measurements. The test setup itself can reduce the performance margin of your product. Removing the loss effect of the test setup will allow you to measure the true performance of your design, and return the design margin that could have been lost due to the test setup. Figure 1 shows the characterization of a cable and, its loss compared to the ideal frequency response of a cable for correct measurement. Here, the cable response is corrected using a cable correction filter to generate an ideal cable response.

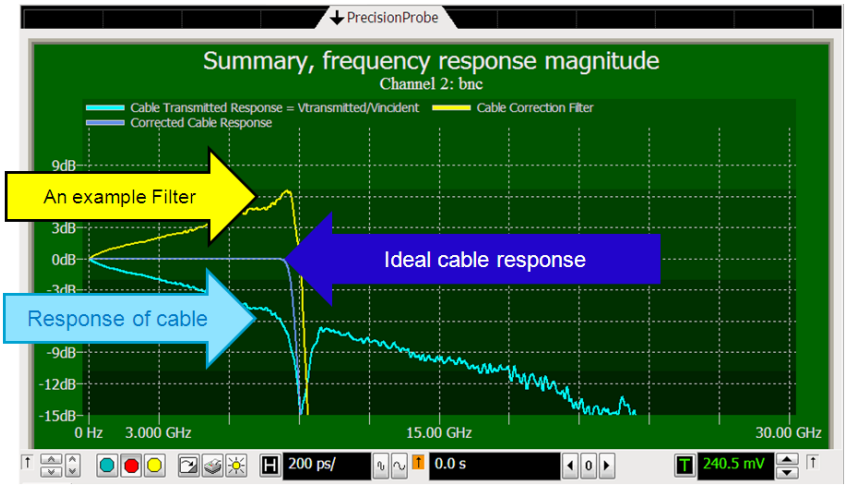

Figure 1. Frequency response and loss profile of a cable

Cable and test fixture loss can affect your measurement in many ways, including: slower rise time, lower amplitude and skew mismatch. All of these can reduce margin and even cause failures. Figure 2 shows three 3 ft 50 Ω cables and their deviation from frequency response (even when 2 and 3 are a "matching" pair). Figure 3 shows a signal as it is measured through a cable and the same signal corrected for cable loss. Here, an amplitude loss of about 30 mV and a rise time loss of 100 ps are corrected. Correction for these losses is important for consistent and accurate measurements.

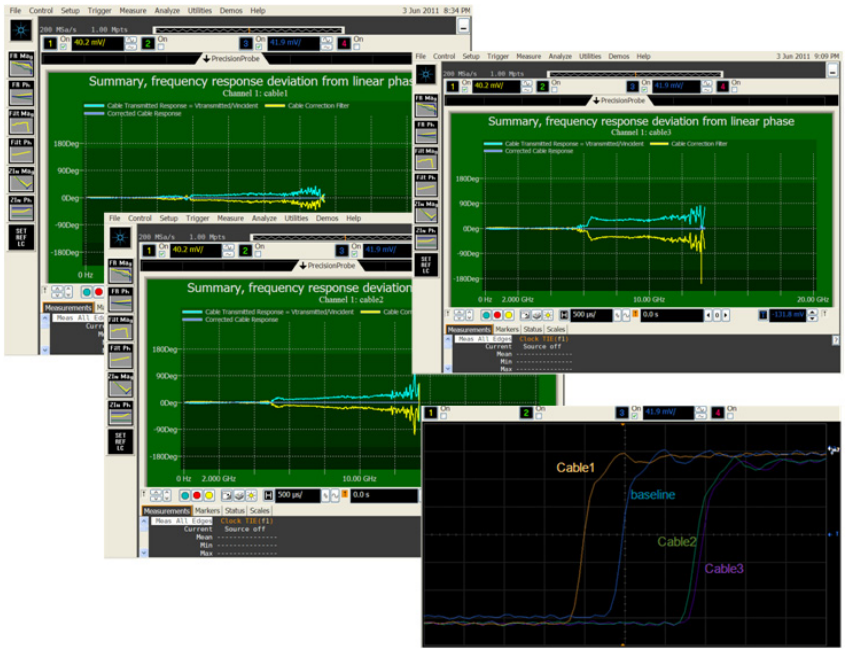

Figure 2. Deviation from frequency response and skew

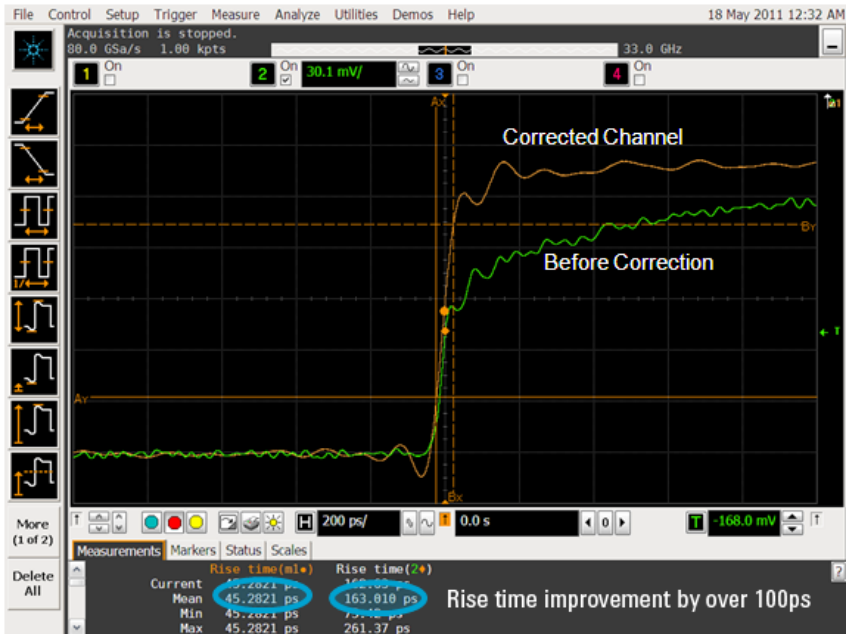

Figure 3. Amplitude and rise time loss due to cable before correction is applied

Hint 2 – Quickly find and fix the root cause of jitter failures

What happens if the signal from your design is failing the SATA or SAS jitter specification? Naturally, you will want to find out the root cause of the jitter issues and fix the problem. This may seem easier said than to done because you may not know where to start.

First, you will really want to find out the type of jitter that is causing the problem. Decomposing the jitter into its different components can help you narrow down the source of the problem. The Total Jitter (TJ) of your system is made up of Random Jitter (RJ), Periodic Jitter (PJ) and Data Dependent Jitter (DDJ), as shown in Figure 4. RJ is caused by thermal or other physical, random processes and the shape of the RJ distribution is Gaussian. PJ is coupled from clock oscillators or power-supply switching into the high speed signal. DDJ is caused by inter-symbol interference (ISI), crosstalk and sub-harmonic distortion. After locating the source of the jitter, the next step is to narrow down which part of your design is causing the problem. An effective way is to use a probe and measure along the signal path to identify which component, channel or source is introducing jitter into the signal path.

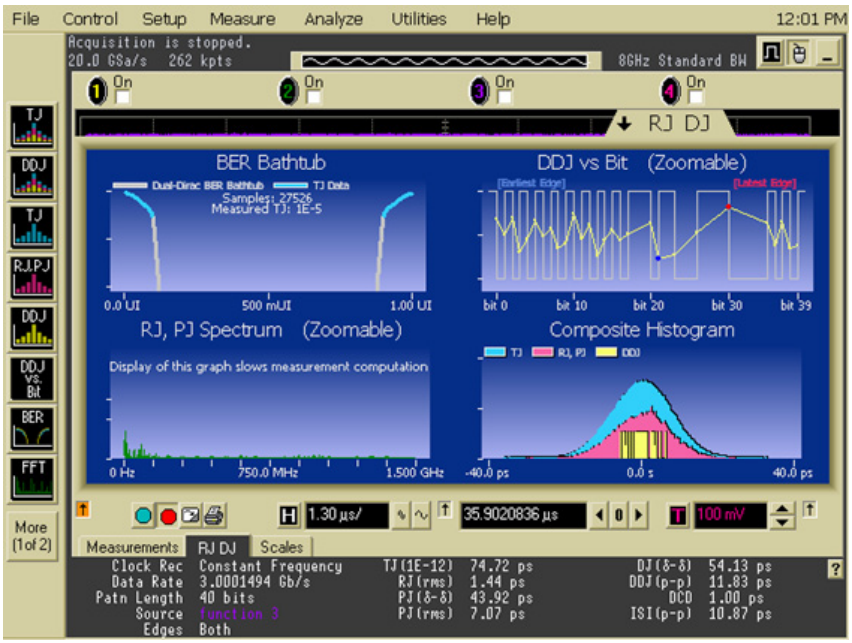

Figure 4. TJ measurement and its components – RJ, PJ, DDJ

Sometimes, the jitter is convolved together so it may be challenging to isolate the source of the jitter further. An example is a low frequency PJ that couples into the signal. In this case, you may need to measure the Time Interval Error (TIE) trend of the signal to identify the PJ, as shown in Figure 5. As can be observed, there are low frequency and high frequency jitters in the signal. Next, you can perform a Fast Fourier Transform (FFT) to the TIE to look at the jitter frequency in the spectrum domain. We can see a spike in the jitter spectrum. Measuring it shows a frequency of 864 kHz, which is the frequency of a clock coupling into the signal. Through this method, we can conclude that insufficient signal isolation is causing the reference clock to couple into the high speed signal. This information can be feedback to board designers to reduce the coupling between the reference clock and high speed signal.

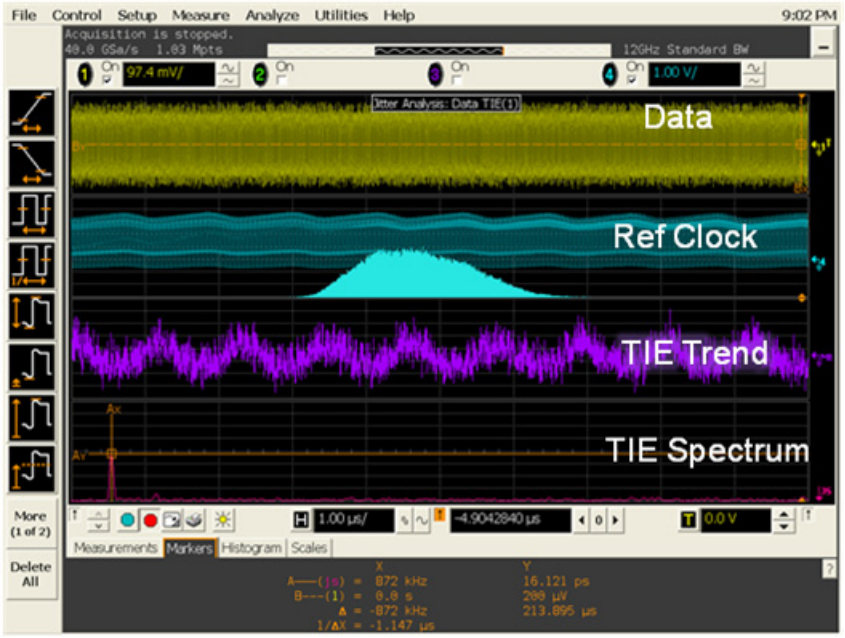

Figure 5. Jitter spectrum – spike at 864 kHz due to reference clock coupling into the high speed signal

Hint 3 – Assess the insertion loss and return loss of your cable, connector and chip

How important is the seemingly simple connector or cable? Using well characterized interconnects is important to guarantee signal quality and achieve high system performance. For SATA and SAS connectors and cable assemblies, eight parameters of time domain, frequency domain and eye diagram measurements are specified to ensure compliance and, more importantly, interoperability. Moreover, impedance mismatch and return loss at the transmitter and receiver of a chip can significantly affect the signal quality of the system and result in eye closure and EMI.

Return loss or reflection occurs when a signal is not absorbed by its termination or an impedance change in its transmission is exceeded. The transmitted signal is reflected back up the transmission line ultimately causing two signals in the transmission line going in opposite directions. These two signals cancel and add along the transmission line at various points. This causes changes in amplitude, phase, jitter, and ultimately, data failure. Changes in physical dimensions, type of insulation and even connectors can cause these reflections.

Insertion loss is the loss of power from an inserted device. The power in does not equal the power out. Transmission line losses are dependent on cable type, operating frequency, and cable length. Insertion loss does vary with frequency and the higher the frequency, the greater the loss. With data rates up to 6 Gbps and soon 12 Gbps, insertion loss plays a key role in the ability to transmit a signal.

As you can see, it is very important to make insertion loss and return loss measurements and find these issues in your transmitters, receivers, cables and connectors. There are two common tools for making these important measurements, a sampling scope or an ENA with TDR option. Figure 6 is an example of making a measurement with an ENA with TDR option.

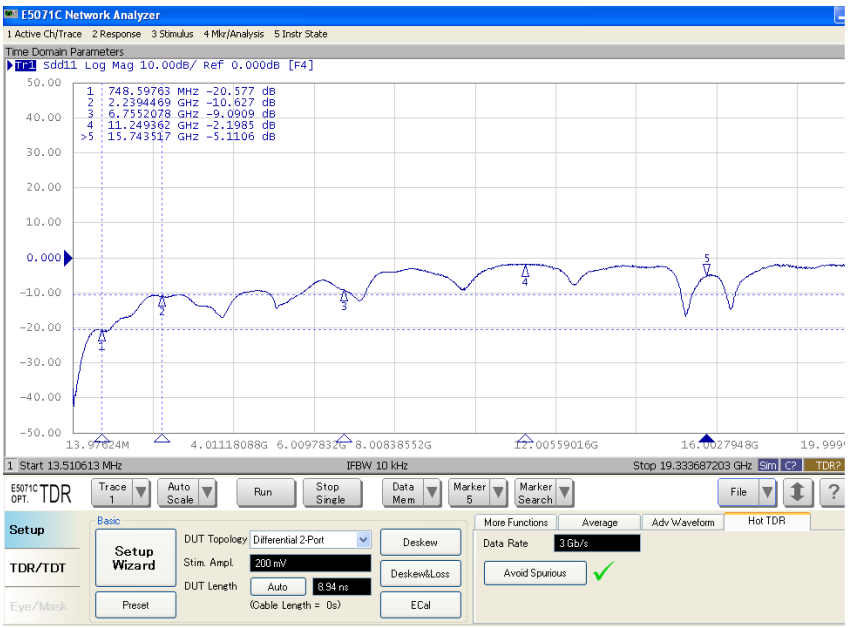

Figure 6. Transmitter differential return loss measurement

Hint 4 – Test beyond the receiver compliance coverage to avoid any hidden problems

Can your receiver handle all the possible signals that it may receive? The receiver jitter tolerance test is designed to tell how robust your receiver is to imperfect signals in the real world. These stresses include noise, jitter, lossy channel and amplitude impairment introduced to the signal. A pattern generator is used to transmit the stressed pattern into the receiver of the product where the product is required to sample and re-transmit the same pattern without error. Error is detected using a CRC Frame Error Rate (FER) counter. Frame error rate is used instead of Bit Error Rate (BER) to check for error because ALIGN patterns could be inserted or removed by the product in the transmitted bit stream.

To comply with the SATA specification, the receiver has to tolerate the combination of stresses up to half of the Unit Interval (UI). SATA-IO only requires the receiver test to be done over four different Sinuisoidal Jitter (SJ) points - 5, 10, 33, and 62 MHz (Figure 7). These frequencies are chosen to stress the Phase Lock Loop (PLL) of the receiver.

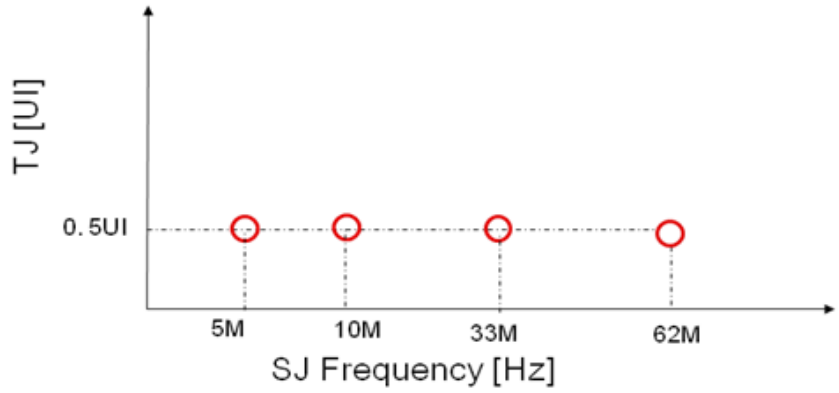

Figure 7. SJ requirements for SATA receiver compliance test

However, this does not tell the whole story of your receiver's performance because the compliance test does not cover all SJ frequency points. There could be failure in between the points that the compliance test does not catch and this could cause problems when your design goes into production.

So, besides the compliance test, you will want to characterize the full spectrum of the SJ frequency points through the jitter tolerance curve test. The test can be created by injecting an increasing amount of jitters to the receiver at a frequency until CRC error is observed at the transmitter. By sweeping the jitter frequency, we can build a curve showing the border where the receiver is passing and failing. Thus, the jitter tolerance curve can verify your receiver performance and the curve is your receiver's susceptibility and tolerance to jitter.

This is actually a more important measurement than the compliance requirement by the SATA compliance standard because only limited points are tested to judge whether your receiver passes or fails. The plot in Figure 8 shows the receiver passing at the compliance point; however, some points have worse performance than the compliance test points.

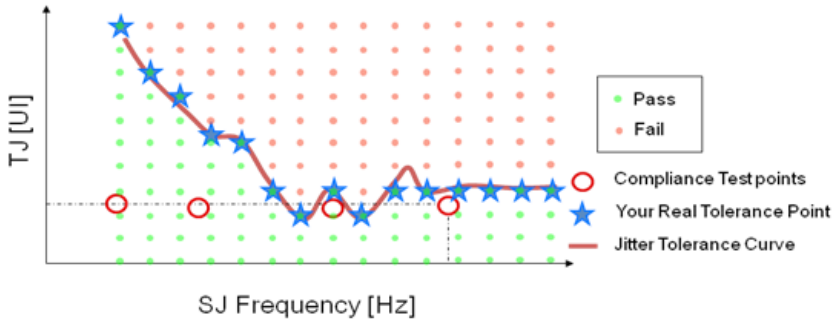

Figure 8. Jitter tolerance curve and results

Hint 5 – Characterize the robustness of your receiver to real world spread spectrum clock (SSC) profiles

One of the issues we've heard a lot recently about storage receiver failures has to do with the strange and distorted SSC profiles from the transmitter of another product. SSC is employed to spread the energy of the transmitted signal to reduce electromagnetic interference (EMI). SATA and SAS specifications recommend a SSC profile of a triangular wave shape, which the receiver's PLL can track easily, as shown in Figure 9. However, many manufacturers are using lower cost parts that generate distorted SSC profiles that the receiver's PLL could find challenging to track. This can lead to an increase in bit errors when the receiver receives the bit stream.

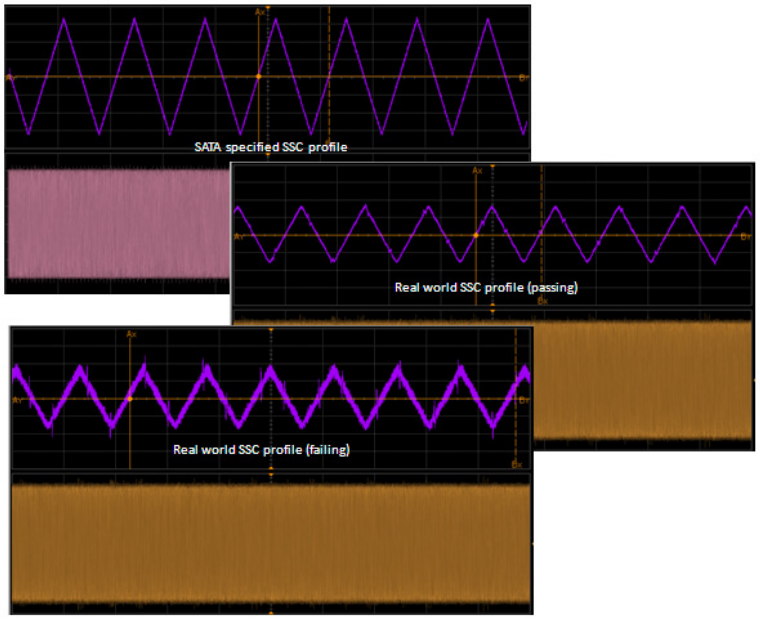

Figure 9. SSC profiles of ideal triangular wave vs. real world

Thus, you will want to be able to characterize your receiver with the various SSC profiles that are observed in shipped products currently on the market. Generating these SSC profiles can be so pattern generator with the functionality to generate arbitrary SSC profiles can speed up the testing. Here is an example of using a J-BERT to insert arbitrary SSC profiles and ability to adjust up to 10000 parts per million (ppm) of clock deviation.

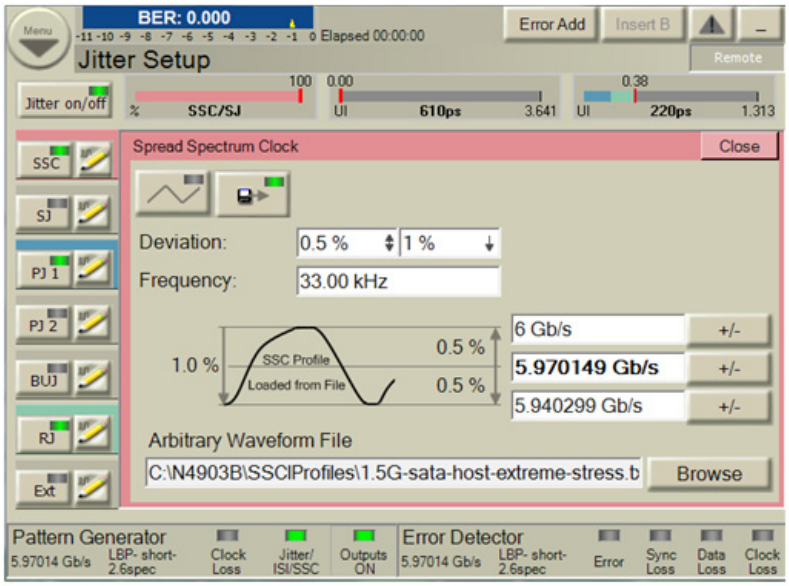

Figure 10. JBERT setup with arbitrary SSC profiles

Hint 6 – Inject protocol errors to check if the product can recover

What if your product cannot recover from an error? Design issues such as error handling and recovery, compatibility, and duplicating issues seen in the field is critical. If your design cannot recover from an error, its use and purpose is a moot point. Injecting errors into the data streams between two devices, thus potentially causing unwanted behavior is helpful during the product development.

Link errors include 8b10b encoding, CRC, link disconnect, running disparity, and ALIGN errors. Protocol/application layers include changing/deleting frames and frame content, inserting and deleting DWORDs into and from a frame, and replacing primitives. Error injection such as starving the buffer, disconnecting a link with outstanding command requests, causing random CRC errors, and changing responses from good to bad is very easy.

In Figure 11, an error injection scenario is created during an FPDMA Queued Write Command (SATA). The jammer changes the fields in the Set Device Bits FIS to reflect a serious error that clears the queue. In State 1, it first looks for a Write FPDMA Queued command and then branches to the next state. In State 2, it identifies the type of frame it is looking for, a Set Device Bits, and specifies an action of modifying the frame fields. The Status Hi, Status Lo, Error, and SActive fields are changed to reflect a new value.

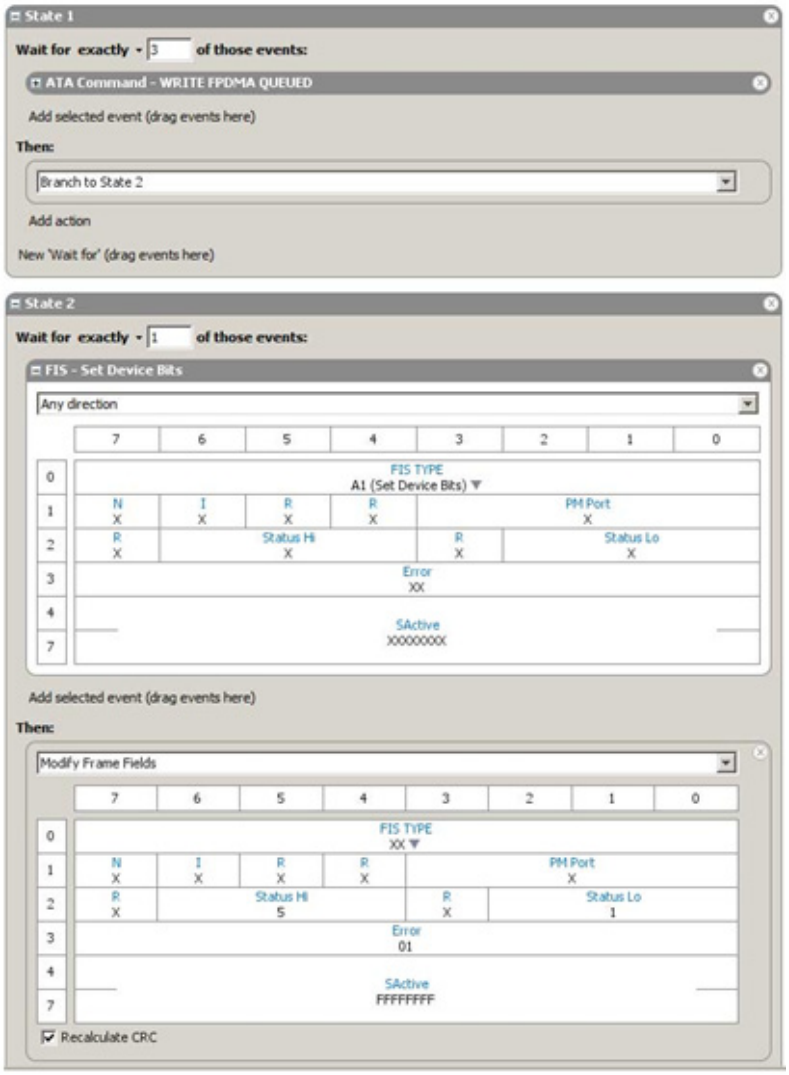

Figure 11. Creating an error injection scenario during an FPDMA queued write command (SATA)

Figure 12 shows a trigger such that the analyzer can zero in on the activity and see the results. Since the error injection was set to change the values of Set Device Bits FIS, the trigger is set to trigger on a Set Device Bits with the SActive field set to FFFFFFF.

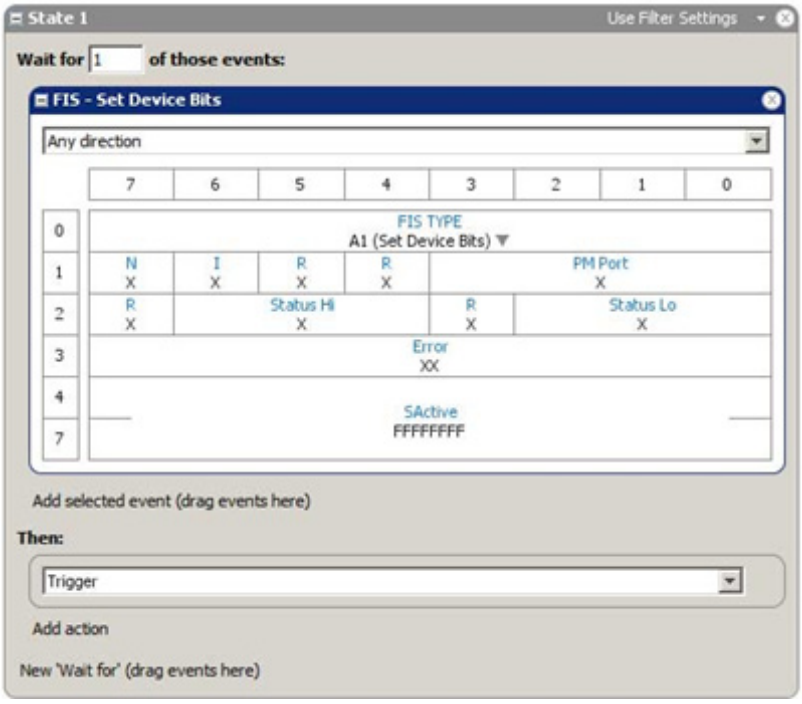

Figure 12. Creating a to find the error condition

Once the protocol analyzer triggers, the protocol view lets the developer see events at the link level. In this example, the Set Device Bits frame is seen before the jammer and after it has been jammed. The trigger point in the trace can also be seen. Frame View confirms the value was changed. In transaction view, the trigger point in relation to the command can be seen. Additionally, in transaction view, there is a read log, which is a command used by SATA to read error logs to help determine the cause of an error. Notice that for other commands there are no corresponding read logs.

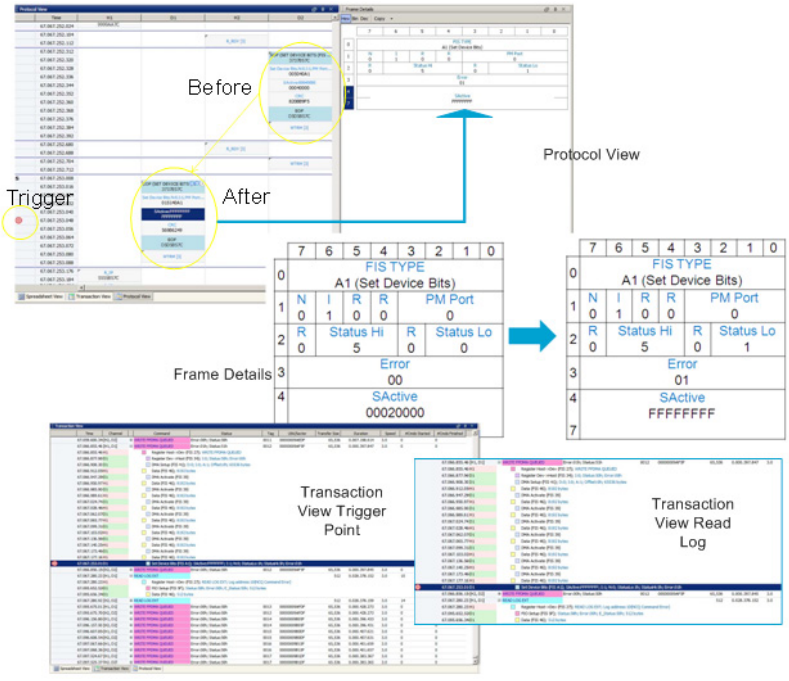

Figure 13. Analyzing the handling and recovery of error injection

Agilent Email Updates

www.agilent.com/find/emailupdates
Get the latest information on the products and applications you select.

LXI

www.lxistandard.org
LAN eXtensions for Instruments puts the power of Ethernet and the Web inside your test systems. Agilent is a founding member of the LXI consortium.

Agilent Channel Partners

www.agilent.com/find/channelpartners
Get the best of both worlds: Agilent's measurement expertise and product breadth, combined with channel partner convenience.

Agilent Advantage Services

Agilent Advantage Services is committed to your success throughout your equipment's lifetime. To keep you competitive, we continually invest in tools and processes that speed up calibration and repair and reduce your cost of ownership. You can also use Infoline Web Services to manage equipment and services more effectively. By sharing our measurement and service expertise, we help you create the products that change our world.

www.agilent.com/find/advantageservices

Agilent Electronic Measurement Group
DEKRA Certified
ISO 9001:2008
Quality Management System

www.agilent.com/quality

For more information on Agilent Technologies' products, applications or services, please contact your local Agilent office. The complete list is available at:

www.agilent.com/find/contactus

Americas	
Canada	(877) 894 4414
Brazil	(11) 4197 3600
Mexico	01800 5064 800
United States	(800) 829 4444

Asia Pacific	
Australia	1 800 629 485
China	800 810 0189
Hong Kong	800 938 693
India	1 800 112 929
Japan	0120 (421) 345
Korea	080 769 0800
Malaysia	1 800 888 848
Singapore	1 800 375 8100
Taiwan	0800 047 866
Other AP Countries	(65) 375 8100

Europe & Middle East	
Belgium	32 (0) 2 404 93 40
Denmark	45 45 80 12 15
Finland	358 (0) 10 855 2100
France	0825 010 700*
	*0.125 €/minute
Germany	49 (0) 7031 464 6333
Ireland	1890 924 204
Israel	972-3-9288-504/544
Italy	39 02 92 60 8484
Netherlands	31 (0) 20 547 2111
Spain	34 (91) 631 3300
Sweden	0200-88 22 55
United Kingdom	44 (0) 118 927 6201

For other unlisted countries:
www.agilent.com/find/contactus

Revised: January 6, 2012

Product specifications and descriptions in this document subject to change without notice.

Published in USA, February 2, 2012
5990-9655EN

Agilent Technologies