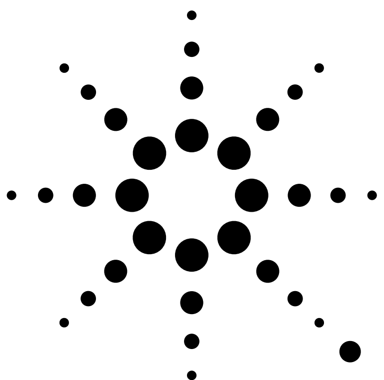

An Introduction to Error Location Analysis

Are all your errors truly random?

Application Note 1550-2

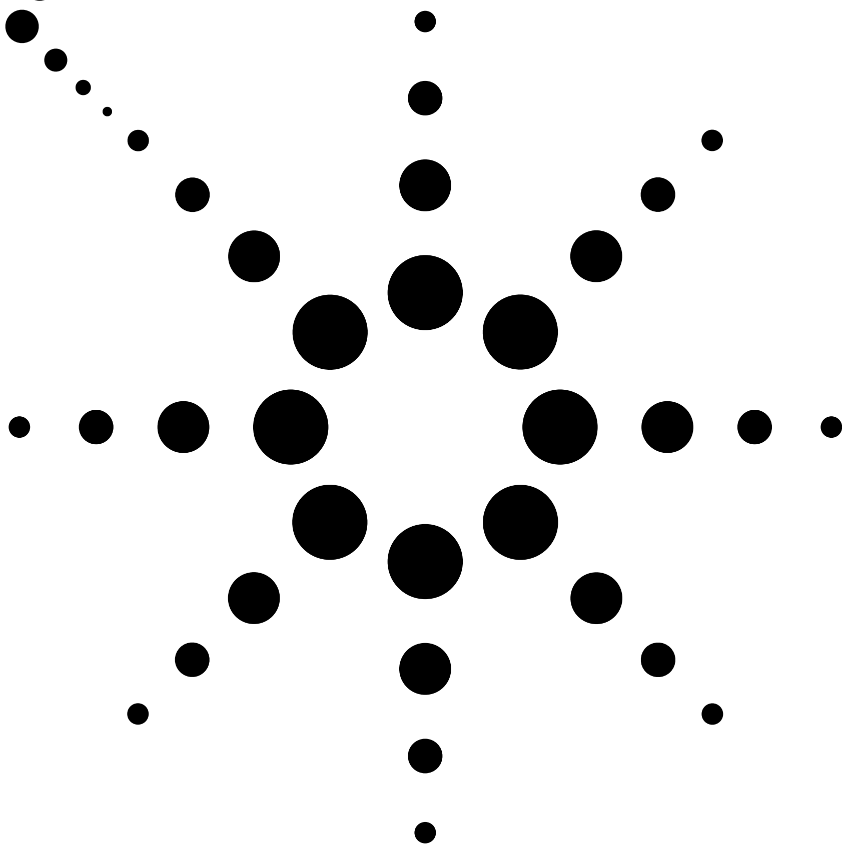

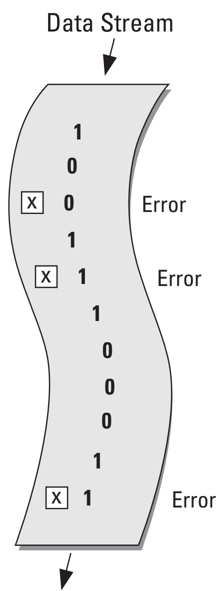

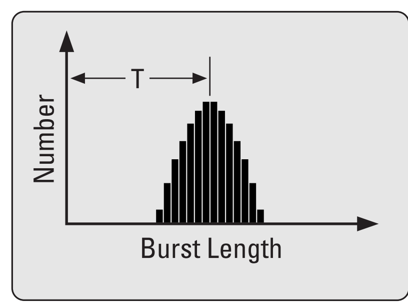

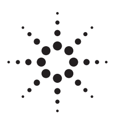

Agilent Technologies

Innovating the HP Way

Introduction

Advancements in test and measurement technology run parallel to the history of communications technology development. Innovations like oscilloscopes, logic analyzers and bit error ratio testers have enabled new, faster and more powerful communications products. Error location analyzers mark the next step in communications test by providing developers detailed error analysis that allows them to isolate causes of bit errors in digital communications systems.

Ideal Applications for Error Location Analysis
- Communications Links with Pattern Sensitivity
- Streams that Suffer Interference
- Packet/Frame-oriented Systems
- Mux/De-Mux Testing
- Burst-prone Communications Links
- Systems with Error Correction Coding
- Electro-mechanical Error Interaction

Developers of digital communications and data storage technology have long used bit error ratio testers in their laboratories and manufacturing floors to measure the quality of their products. Bit error ratio (BER), the result of dividing the number of bit errors by the total number of bits in a stream, can be used to identify the performance of a communications or data storage product. When bit errors exceed the desired limits, engineers must distinguish between the many possible error causes to identify what problems exist and to what degree. Debug and redesign efforts can then be quickly focused on the right activity. Error location analysis can provide this focus.

By using the exact same test connections that are used in traditional bit error ratio testers, error location analysis sorts out the major classifications of errors. For example, design efforts can be directed towards improving equalization and bandwidth issues rather than signal-to-noise related issues or on a systematic hardware logic problem rather than improving error correction capability.

Error Location Analysis

Before starting, it would be helpful to define error location analysis. Error location analysis is a technique of locating the precise position of all errors in a data stream in order to study various aspects of error spacing and relative error position. This error location information improves bit error ratio usefulness in much the same way that memory about data values at various time locations did in digital storage oscilloscopes and logic analyzers. Bit error ratio measurements are a subset of error location analysis allowing traditional BER measurements to be made as well.

Let’s look at an example, in the 40-bit transmit sequence shown in the figure below, we can see that the received sequence had 3 errors in locations 7, 15 and 31.

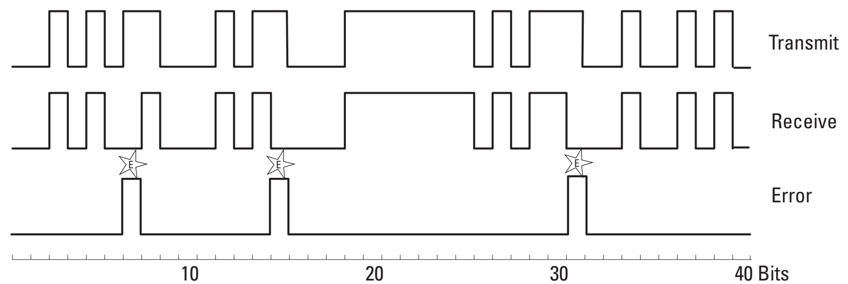

Figure 1. A transmitted bit sequence with errors

These 3 errors out of the total 40 bits transmitted amount to a bit error ratio of 7.5×10^{-2} showing the poor average performance of the communication link. However, the exact error positions hold much more information. For instance by looking further, the errors were spaced in perfect multiples of 8-bits. This can be seen because the error locations 7, 15 and 31 all have the same remainder when divided by the number 8. This casts suspicion on any processing that operates in an 8-bit domain like the low-speed side of an 8-to-1 multiplexer. Magic numbers such as 8-bits or, in general, any other correlation that can be found among detected errors is critical information that will help make better communications products.

In this discussion, we will describe each type of typical error found in digital communications systems and see how error location analysis can be used to identify it. These include:

- Gaussian bit errors
- Pattern sensitive bit errors
- Systematic bit errors
- Systematic burst errors
- Random burst errors

Types of Bit Errors

Bit errors in digital communications can come from many sources. Knowing the source is the first step in solving error problems. Practical classifications of bit error types include: Gaussian bit errors, pattern sensitive bit errors, systematic bit errors, systematic burst errors, and random burst errors. All error types can occur simultaneously and each type of error offers different evidence that it is present. Error location analysis distinguishes between the different types and identifies correlation between systematic error types.

Gaussian Bit Errors

The most popularly discussed source of bit error in digital communications comes from random noise. Bit detectors typically must set a decision voltage threshold above which ones are detected and below which zeros are detected. In the presence of noise, an otherwise zero bit might be distorted just enough to have it be detected as a logical one and visa versa. The probability of this happening is related to distribution of the Gaussian noise being superimposed on the binary data stream.

Bit errors in Binary Channels from Random Noise

An ideal random data stream would have two distinct detected voltage levels, μ_0 and μ_1. All bit decisions would either be μ_0 or μ_1:

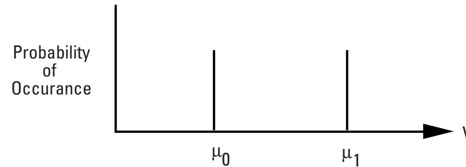

Random noise has constant energy with respect to frequency and adds distortion to traces that make up data bit cells. The probability density function (PDF) of the noise is a Gaussian shape.

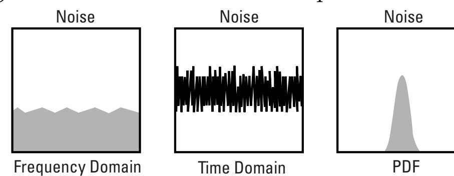

Truly random noise adds uncertainty to the bit levels at the decision point:

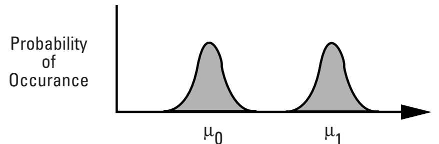

Gaussian tails never have a zero probability of occurring and at some point, the two tails will cross:

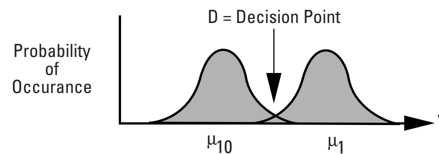

The area underneath the crossing point relates to the random bit errors that occur because of the Gaussian noise:

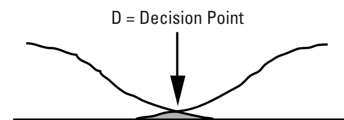

The bit error ratio, p_e, is the probability that a one is detected when a zero is expected (area under the curves to the right of the decision point, D) plus the probability that a zero is detected when a one is expected (area under the curves to the left of the decision point). The area under the tails of a Gaussian is designated with an *erfc( )* function, leaving:

$$p_e = \frac{1}{2} \left\{ erfc \left(\frac{|\mu_0 - D|}{\sigma_0} \right) + erfc \left(\frac{|\mu_1 - D|}{\sigma_1} \right) \right\}$$

Or, more generally put, the bit error ratio of random errors is dependent only on the decision threshold, the average one and zero voltage level and the standard deviation of the Gaussian noise distributions on each of the one and zero voltage rails[1].

[1] For this reason, Q-factor has been used as a quality metric for a digital communications channel that suffers from random errors. Q, is the ratio of distance between the means, i.e. the signal, $(\mu_1-\mu_0)$ divided by the sum of the standard deviations of the Gaussian noise, i.e. the noise $(\sigma_0 + \sigma_1)$. Q can be measured with an oscilloscope or derived from BER measurements. Later we will see that Q is not appropriate for channels that also have other types of error.

By understanding the mechanism by which these types of errors occur, observations can be made. First, it is clear that errors from this source must be able to occur on any bit in a bit sequence[2]. If errors are occurring on only particular bits in a data transmission rather than all bit positions, then the cause of the errors is not only Gaussian noise. Second, the probability of error at any given location is completely independent, so the probability of two or more errors in a row is easily derived as p_e^N where N is the number of adjacent errors. As this is an exponential function, the probability of getting 2, 3 or 4 bit burst errors from Gaussian sources becomes very small. For example, if the probability of getting a random error in a 2.5 Gb/s data link is $1x10^{-7}$, then the probability of getting a 3-bit burst is $1x10^{-21}$ or one every 13,000 years[3]. Large numbers of errors that are observed to be in 3-bit bursts are unlikely to come from Gaussian noise superimposed on the data bits.

Differentiating Bit and Burst Errors

Every bit error ratio application defines burst errors differently because the definition depends greatly on the physics and theory of each underlying communication channel or storage medium. Error location analysis allows for this flexibility by using a burst error-defining model that is configured differently for different applications. The key parameters used relate to the minimum error grouping length required to be of interest as a burst error (the Minimum Burst Length) and the maximum error-free spacing that an application would want to permit inside an error grouping (the Maximum Error Free Interval). Error groupings with an error-free spacing larger than the limit would break apart into two separate groupings. These two parameters are used during error location analysis to assign length to all error groupings and to distinguish burst-related errors from non-burst related errors.

The first step is to assign length to every error grouping. Error location analyzers constantly calculate the distance between every detected error and the previous error. As long as this distance is less than the defined Maximum Error Free Interval, the distance from the first error in the grouping to the current error is accumulated as the error length. As soon as an error is detected that is spaced farther away from the previous error than the Maximum Error Free Interval allows, the previous error grouping is frozen at the length previously accumulated and a new error grouping is opened.

[2] If the decision point D, is closer to the one rail, the detected Gaussian errors would occur as ones turning into zeros; however, all ones would have equal probability of being in error. The same would be true for the zero rail.

[3] Even two-bit errors would only be expected to occur once every 11.5 hours.

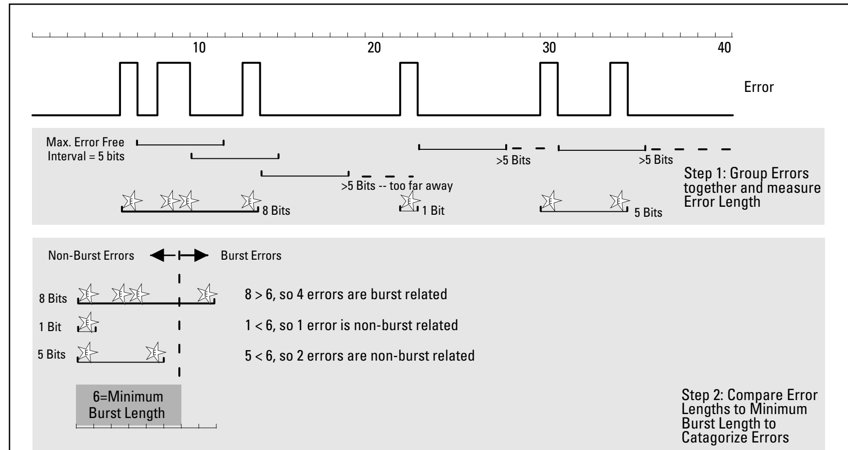

This diagram demonstrates burst and non-burst error separation that is done in error location analyzers with the Maximum Error Free Interval set to 5 bits and the Minimum Burst Length set to 6 bits.

Once the length of all error groupings are known, the analyzer compares these lengths to the desired Minimum Burst Length parameter and add those errors found in groupings which meet the Minimum Burst Length criteria to compute the burst error statistics. All other errors are assigned as non-burst errors.

During this process, all error groupings are assigned unique lengths. Additionally error-free intervals between errors are measured. Individual histograms showing distributions of error grouping lengths or error-free intervals are also produced. It is important to understand the burst error model and the use of Minimum Burst Length and Maximum Error-Free Interval parameters. The same error location data can be analyzed with different parameters to the burst error model yielding different interpretations of the result.

Distinguishing between burst or non-burst errors and determining the length of groups of errors in a communications channel, is done by measuring the distance between the first and last error in an error grouping (see sidebar). It is convenient to histogram all error lengths found in a channel to reveal the distribution of error lengths. This is especially useful as a first step in identifying Gaussian noise related errors. In a channel with only gaussian noise, errors will most probably occur as isolated one-bit errors. It is not true, however, that all channels with one-bit errors suffer from Gaussian noise errors—see Systematic Bit Errors example below.

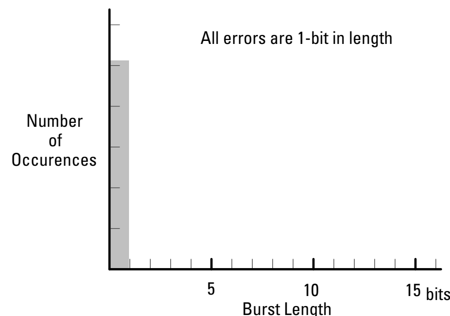

Figure 2. Systematic Bit Errors

Pattern Sensitive Bit Errors

Pattern sensitive bit errors, sometimes called data or pattern dependent errors, occur when certain bit positions of the data being communicated or stored have more errors than other positions within the pattern. This can be caused by a variety of physical phenomenon. For example, in bandwidth limited communications channels, the high-frequency requirements of transmitting an isolated one or zero surrounded by many bits of the opposite sense may cause these types of data bits to be prone to a higher error ratio than other sequences. Similarly, when magnetically recording two bit transitions very close to each other, the magnetism may cause the two nearby fields to attract each other causing the two neighboring bit cells to shrink. This, too, can cause a higher probability of error to occur in these types of data patterns.

Pattern sensitive errors are also very common in systems that exhibit eye diagram distortions. When pattern sensitive errors exist, they are typically responsible for the first bit errors that occur when trying to stress the digital channel. This is because these errors come from the worst-case bit patterns and added stress effects these worst-case positions first. For example, when measuring bit error ratio against decision threshold voltage (often done when making Q-factor measurements), probability theory quickly predicts a certain well-behaved relationship between the decision voltage threshold and the resulting bit error ratio. However, if the bit errors come from pattern sensitive errors then the curve prediction will be inaccurate. This makes it impossible to accurately characterize the overall channel error ratio performance with such methods when pattern sensitive errors are present.

In the case of pattern sensitive errors, error location analysis shows that all bits of the received data stream do not have the same probability of error. This was a key assumption when deriving p_e earlier. That is, that the noise being added to all received bits might cause bit errors. In pattern sensitive data cases, only noise added to a few, or perhaps even only one, bit of the test data sequence will contribute to bit errors. This means that the measured bit error ratio will appear better than that predicted. This is because only a small fraction of all the bits being transmitted fall into the pattern sensitive case while only bits from the pattern sensitive areas of the sequence have any real chance of causing bit errors.

This graph compares the measured bit error ratios near the negative rail of an incoming data signal that suffered from pattern sensitivity. In this example, the decision threshold voltage was varied less than 100 mV. One can see that the bit error ratio projections matched the measured results for decision thresholds that yielded generally poor bit error ratios worse than 1×10^{-2} (decision thresholds of –0.86 mV to –0.81 mV). However, as the decision threshold voltage moved farther away from the negative rail towards the center of the eye, bit error contributions stopped coming from all possible bit positions within the pattern.

Recall that bit error ratio is the ratio of the number of detected errors divided by the number of received bits. In this case, the number of received bits has not changed, but the number of pattern sensitive bit positions that could actually cause a bit error is significantly less. This means that theoretical projections of bit error ratio that rely on using contributions from all received bits would have an inflated number of bits in the denominator compared to the actual number of bit positions that might cause errors by the theoretical model. This would result in smaller bit error ratios. This caused the downward trend seen from decision thresholds of –0.81 mV to –0.79 mV.

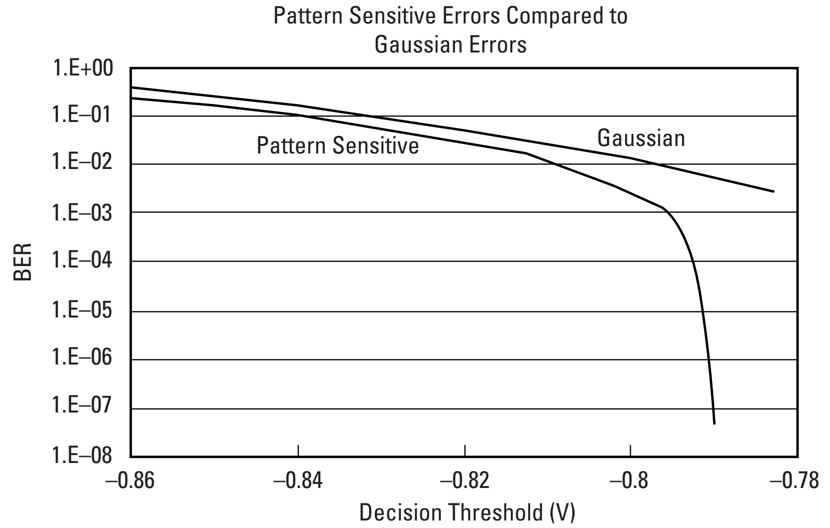

Figure 3. Comparing the effect pattern sensitive and Gaussian errors have on the BER as the decision point is varied.

Identifying pattern sensitive errors may be carried out by attributing every detected bit error to the precise bit location within the test data sequence. Typically, pseudo-random binary sequences are used for such tests. For example, a short pseudo-random sequence, 2^7-1, repeats every 127 bits. Because error location analyzers know the exact bit location of errors, a histogram can be constructed showing the number of times errors were found at each bit position of the 127-bit sequence. A flat distribution would indicate that all bit positions were equally likely to be in error. However, unexpected spikes in this histogram indicate pattern sensitivity.

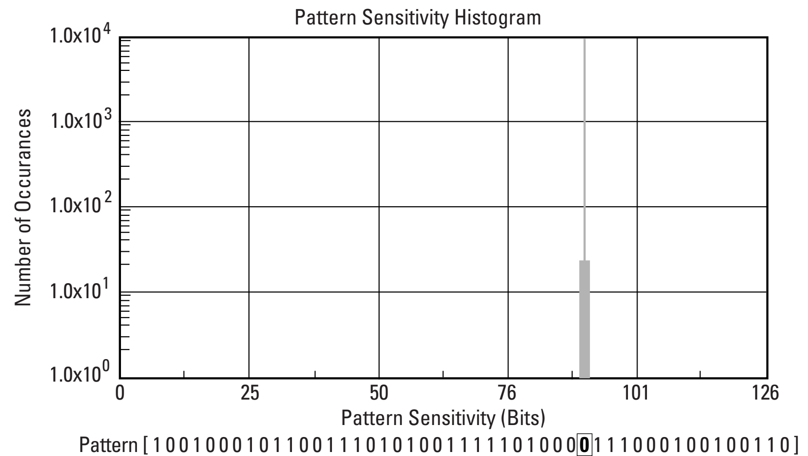

Figure 4. Pattern sensitivity analysis can highlight persistent errors that relate to a particular bit sequence

Solving pattern sensitive issues can be very challenging and depends on the physics of the situation. By comparing incoming data with a previously captured baseline, engineers can quickly see the effect of design changes on performance.

Systematic Bit Errors

Systematic errors are very common in digital communications. The pattern sensitive errors discussed above are a form of systematic error. Systematic errors are not random and, therefore, correlate to some other cause or event. In pattern sensitive errors, the correlation is to the data patterns being used. Non-pattern sensitive systematic errors typically come from interference or digital logic problems. Systematic errors are often very straightforward to detect; solving them can prove to be more difficult.

Error location analysis is a useful tool for finding systematic errors. The goal is to find a repetitive interval or correlation that directly relates to the observed errors. For example, switching power supplies and DC-DC converters trade-off their convenient size and low cost for increased high-frequency switching noise. It is not uncommon to find stray high-frequency voltage spikes on sensitive bit detection circuitry every 4 microseconds stemming from poor filtering on a 250 kHz switching supply. Depending on the size of the bit cell as defined by the data rate and the duration of the interference, one or more data bits can be affected by such interference.

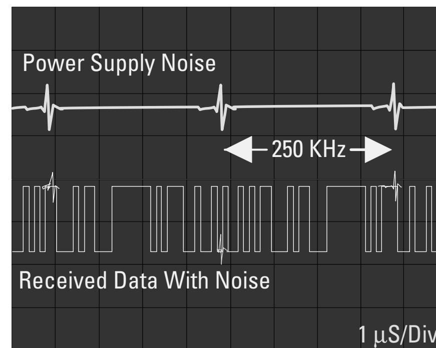

Figure 5. The effect of power supply breakthrough on the incoming data stream.

Because most external interference is asynchronous with the communication channel, systematic one-bit errors often come from electronics inside the communications system. A failing memory transistor cell in a RAM chip used to buffer data packets would cause isolated bit errors each time that cell of the memory chip was referenced. Because the error is short and isolated from other bit errors, it could easily be mistaken to be part of the background Gaussian error except for the fact that the interval between errors would be very predictable.

The multiplexer/de-multiplexer testing application offers an excellent example of probable systematic bit errors that, if not studied correctly, might be incorrectly blamed on the high-speed serial communication interface. In this application errors are almost-automatically attributed to the high-speed serial interface. However, errors can also come from the parallel bit I/O signals. For example, a low-speed signal could be weak, it could suffer from a poor connection, or could have a bad timing relationship with the parallel rate clock causing it to violate setup/hold requirements.

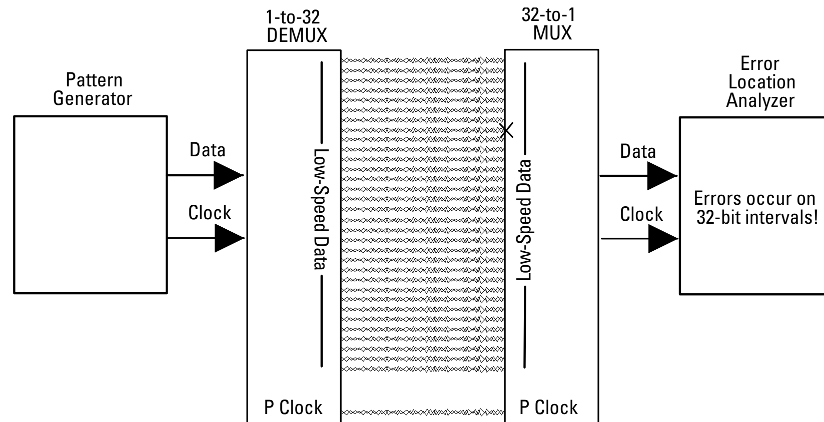

Figure 6. A fault on one line of a parallel interface can be isolated.

Errors caused by a failure on the low-speed interface will have a characteristic signature in the overall serial bit error statistics. These errors will occur in intervals that are multiples of the multiplexer/de-multiplexer width. For example, when testing a 32:1 multiplexer a failing data bit on the 32-bit low-speed data bus would cause errors to come in intervals that are multiples of 32 bits. This evidence quickly shifts suspicion from the high-speed serial logic to the lower-speed interface logic.

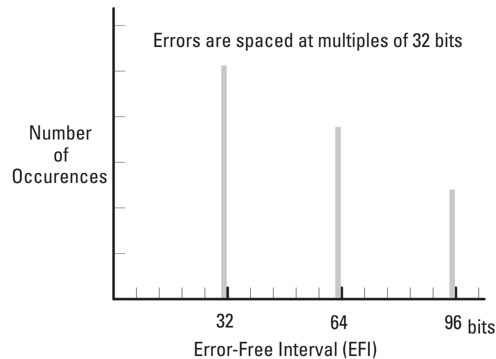

Figure 7. Error free interval display showing structure which indicates a relationship to the number 32.

If error-free intervals of a particular length occur many times, for example at intervals of 32-bits, the errors involved cannot be random. By looking at the histogram of all error-free intervals, error location analysis easily shows when systematic errors are present. A histogram of error lengths can be used in conjunction with a histogram of error-free intervals to determine that errors are both isolated single-bit errors and repetitive. This combination of results would eliminate the conclusion of having a solely Gaussian error cause.

Systematic Burst Errors

Like the systematic bit error case, channels with systematic burst errors also have some error-free intervals that occur more often then others. However, their burst length profiles would show error lengths larger than a single bit.

Systematic interference of any kind can also be categorized as either being synchronous or asynchronous to the communications system. An example of synchronous interference is the multiplexer/de-multiplexer example discussed earlier. In this case, errors occur with precise bit-accurate relationships to each other (for example 32 bits). Other examples include errors caused by electronic

processing of packets, digital counters and pattern sensitivity. Synchronous interference is often rooted inside the communication signal processing itself, as this is the only place packet lengths, multiplexer/de-multiplexer widths or data patterns are known.

Alternatively, asynchronous interference typically comes from outside the basic signal processing of the communications system. Examples here include noise induced from switching power supplies or other electro-mechanical interaction.

Error location analysis can be used to differentiate synchronous interference from asynchronous interference. To accomplish this, error free interval and burst length histogram distributions needs to be studied carefully. If spikes or other anomalies in the histograms occur at very specific intervals and lengths, a synchronous relationship probably occurs. On the other hand, if the histogram results show effects that broadly affect many neighboring histogram entries, then asynchronous interference is suggested.

The difference has great impact when explaining what might have caused the error. A “perfect” spike rarely happens in nature and is probably a result of digital logic.

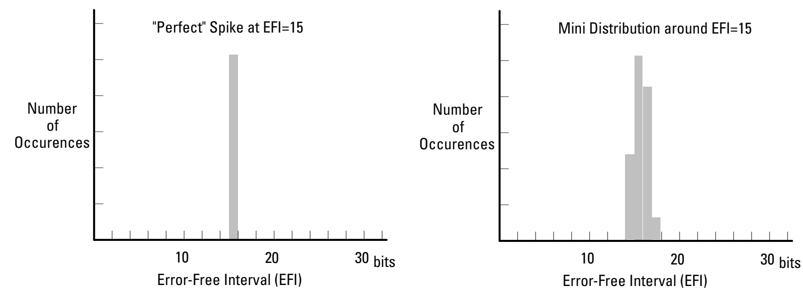

Figure 8. Comparing error free interval plots for a digital logic problem and an external interference related problem.

In the 250 kHz switching power supply example above, the number of bits affected by the injected noise might be quite large if the data rate is high. See Figure 9. In this case the repetitive interference would be seen as small bursts of errors during the worst points of interference in the 4-microsecond cycle. For example, consider that the data rate of the communication was 2.5 Gb/s and that the interference occurred for 5 to 7 nanoseconds. Error bursts of between 12 and 17 bits would be measured corresponding to 5 to 7 nanoseconds of data. Error free intervals would be between 9983 and 9987 bits corresponding to the 4 microsecond repetition rate.

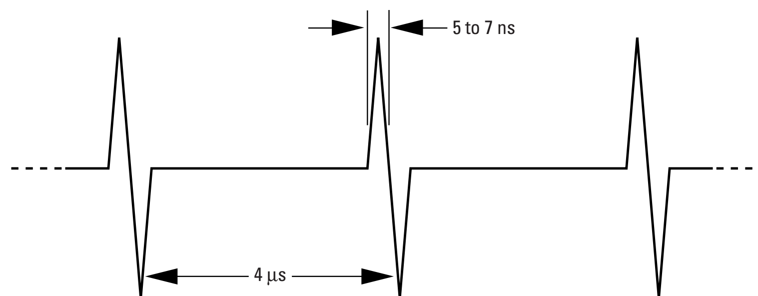

Figure 9. Switching mode power supply breakthrough

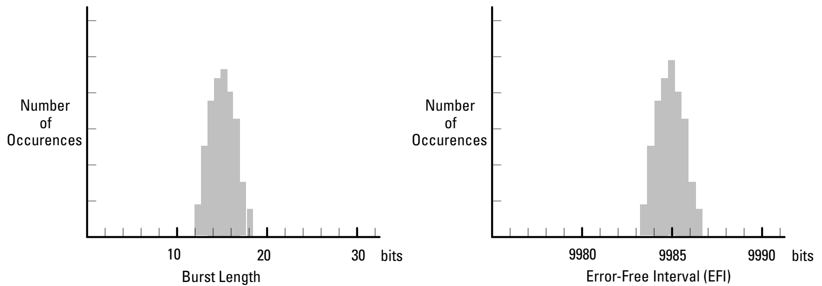

Figure 10. An example burst length plot and error free interval analysis relating to a power supply breakthrough problem.

Such broad peaks in both graphs clearly differentiate asynchronous interference from other causes.

Another form of systematic burst errors occur when one error burst causes another later in the communication. An example of this is certain forms of error correcting channel codes (e.g. Viterbi). One error can propagate into future errors because prior bit decisions in the detector are used to help make future bit decisions. Any bad bit decisions made in the past will inject bad input into future decisions and may cause errors. In these cases, emphasis on correcting the cause of the first error will translate into better system performance.

Errors that are systematically caused by previous errors can be identified by auto-correlating error locations. Error auto-correlation shows how many times errors were found at all bit positions following an error. This view shows the probability of getting an error n-bits away from another error. At first glance, this analysis appears very similar to an error free interval histogram. However, in the presence of high background bit error activity, large error free intervals present from low-frequency interaction might be broken into a series of smaller error free intervals by the high background bit error activity. This would hide a correlation that may exist at larger intervals (i.e. lower frequencies).

Random Burst Errors

Channels with random burst errors have error-free interval histograms that are identical to the random bit error case. That is, a smooth monotonically decreasing occurrence rate of ever-larger error-free intervals. However, by studying the distribution of burst lengths, it is easy to differentiate this type of error.

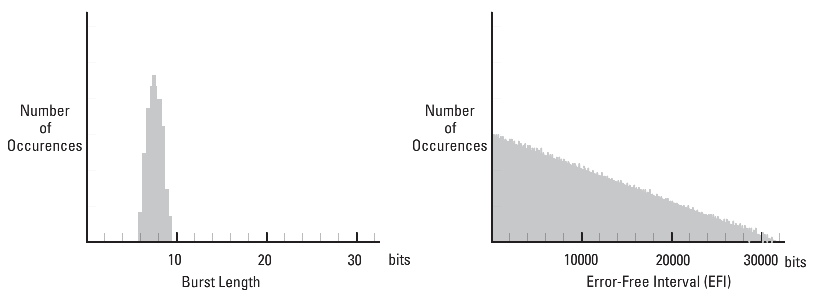

Figure 11. Random burst errors viewed on two different plots.

Random burst errors result from many natural phenomenon including raindrops, snowflakes, debris and media surface sputtering. These errors are characterized by having bursts of characteristic lengths, at intervals that are unpredictable. Burst-prone communication channels include satellite communications, microwave radio links and magnetic/optical recording systems. To improve reliability, these channels are often protected by interleaved block-oriented error correction codes.

Random bursts are differentiated from systematic bursts in order to aid in more quickly identifying the error source. Like the sizes of raindrops or debris particles, random bursts often have a naturally random process as their root error injecting mechanism.

Error Correlation

Identifying and fixing systematic errors deserves special consideration because this is the source of many headaches in both manufacturing and development applications. Once the burst lengths and error free intervals identify that errors are systematic, the next step is to find the exact nature of the systematic interaction. Synchronous bit-accurate internal error interaction, such as memory chip failures, counter chains and packet lengths will correlate to the fundamental divider that is common to the circuits. Interference from asynchronous external mechanisms such as motor noise interference or vibration will correlate to derivable signals such as crank-case top-dead-center or vibration frequency marker.

Error correlation histograms show the number of times errors have occurred at all bit phases of externally applied marker signals or from an internally synthesized fundamental bit-divider. For example, the multiplexer/de-multiplexer example discussed earlier that suffered from errors introduced in the low-speed 32-bit parallel logic side. Correlating the error positions to a synthesized divider of 32 shows the number of errors in each of the 32 bit positions of the low-speed interface.

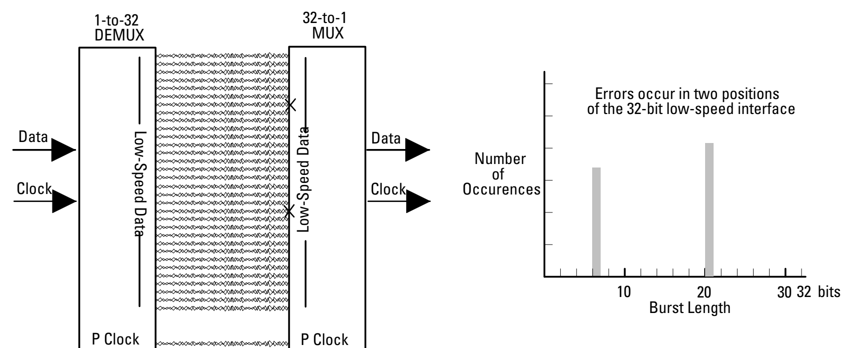

Figure 12. Two error-causing lines on a parallel interface can lead to a burst length plot as shown.

In this example of two weak receivers, we see that two bits on the parallel interface had occasional errors on them. One line was worse than the other. This could not be easily distinguished using the error-free interval analysis as many errors would still be spaced by multiples of 32 bits. However, it shows up clearly when digging deeper into the error correlation. Error correlation analysis operates on relative bit error locations, so the x-axis of this histogram could start at any bit position on the serial interface. This means that the two spikes may occur at different bin locations of the analysis histogram, but the spacing between the bins will remain constant.

This type of correlation analysis is often useful because many high-speed systems achieve their speed by parallel processing and later multiplexing. In these cases, error correlation can quickly identify if errors are present only on certain lower-speed interleaves. For example, in systems that parallel-process two paths for eventual bit-by-bit combination into a serial stream, knowing if the errors are present on the odd or even bits can isolate a failure down to a single path.

Correlating to external markers requires that an external signal act as a trigger for error correlation. Markers such as these are quite familiar to electronics design because they act in much the same way a trigger input acts on an oscilloscope. External markers can be derived from disk rotations, vibration, temperature cycling, tape recording tracks or any other controlled interference. External markers, like oscilloscope triggers, need not occur regularly or at a fixed frequency.

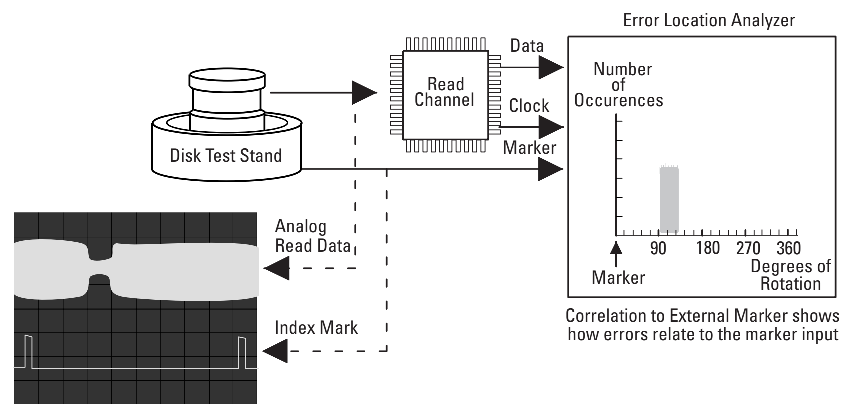

Figure 13. A disk drive example where an external marker is used.

Conclusion

In systems that suffer interference from internal or external sources, quick identification and isolation of the error cause translates directly into reduced debug time. In systems with error correction coding, the nature of bursts and intervals between errors are vital to the success of the error corrector. In this case, being able to quantify these values before correctors are designed is required for efficient development. In optical network and component testing, the likelihood of pattern sensitivity in bit error ratio measurements also requires special attention.

Bit errors come in many varieties. In this application note we have identified five different types of errors and how error location analysis can be used to identify them. Being able to quickly distinguish between the variety means real savings for design and debug efforts.

For more information about Agilent Technologies test and measurement products, applications, services, and for a current sales office listing, visit our web site,

www.agilent.com/comms/lightwave

You can also contact one of the following centers and ask for a test and measurement sales representative.

United States:
Agilent Technologies
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
(tel) 1 800 452 4844

Canada:
Agilent Technologies Canada Inc.
5150 Spectrum Way
Mississauga, Ontario
L4W 5G1
(tel) 1 877 894 4414

Europe:
Agilent Technologies
Test & Measurement
European Marketing Organization
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(tel) (31 20) 547 2000

Japan:
Agilent Technologies Japan Ltd.
Call Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192-8510, Japan
(tel) (81) 426 56 7832
(fax) (81) 426 56 7840

Latin America:
Agilent Technologies
Latin American Region Headquarters
5200 Blue Lagoon Drive, Suite #950
Miami, Florida 33126, U.S.A.
(tel) (305) 267 4245
(fax) (305) 267 4286

Australia/New Zealand:
Agilent Technologies Australia Pty Ltd
347 Burwood Highway
Forest Hill, Victoria 3131, Australia
(tel) 1-800 629 485 (Australia)
(fax) (61 3) 9272 0749
(tel) 0 800 738 378 (New Zealand)
(fax) (64 4) 802 6881

Asia Pacific:
Agilent Technologies
24/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong
(tel) (852) 3197 7777
(fax) (852) 2506 9284

Technical data subject to change

Copyright © 2000
Agilent Technologies

Printed in U.S.A. 4/00
5980-0648E

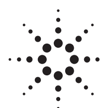

Agilent Technologies
Innovating the HP Way