Bit Error Ratio (BER) Test Time
A Statistical Confidence Level Approach

Justin D. Redd, Ph.D.

Note: Versions of this article have been previously published in Lightwave Magazine
https://www.lightwaveonline.com/home/article/16647704/explaining-those-ber-testing-mysteries and as Maxim Integrated Products technical article "HFTA-05.0."

The ultimate function of the physical layer in any digital communication system is to transport bits of data through a medium (such as copper cable, optical fiber, or free space) as quickly and accurately as possible. Hence, two basic measures of physical layer performance relate to the speed at which the data can be transported (the data rate) and the integrity of the data when they arrive at the destination. The primary measure of data integrity is called the bit error ratio, or BER.

This article reviews the BER requirements common to telecommunication and data communication protocols, provides an overview of the equipment used to test BER performance, and examines the trade-off of test time versus BER confidence level.

1. BER Specifications

The BER of a digital communication system can be defined as the estimated probability that any bit transmitted through the system will be received in error, e.g., a transmitted "one" will be received as a zero and vice versa. In practical tests, the BER is measured by transmitting a finite number of bits through the system and counting the number of bit errors received. The ratio of the number of bits received erroneously to the total number of bits transmitted is the BER. The quality of the BER estimation increases as the total number of transmitted bits increases. In the limit, as the number of transmitted bits approaches infinity, the BER becomes a perfect estimate of the true error probability.

In some texts, BER is referred to as the bit error *rate* instead of the bit error *ratio*. Most bit errors in real systems are the result of random noise, and therefore occur at random times as apposed to an evenly distributed rate. Also, BER is an estimate formed by taking a ratio of errors to bits transmitted. For these reasons, it is more accurate to use the word ratio in place of rate.

Depending on the particular sequence of bits (i.e., the data pattern) transmitted through a system, different numbers of bit errors may occur. Patterns that contain long strings of consecutive identical digits (CIDs), for example, may contain significant low-frequency spectral components that might be outside the pass-band of the system, causing deterministic jitter and other distortions to the signal. These pattern-dependent effects can increase or decrease the probability that a bit error will occur. This means that when the BER is tested using dissimilar data patterns, it is possible to get different results. A detailed analysis of pattern-dependent effects is beyond the scope of this article, but it is sufficient to note the importance of associating a specific data pattern with BER specifications and test results.

Most digital communication protocols require BER performance at one of two levels. Telecommunications protocols, such as SONET, generally require a BER of one error in 10^{10} bits (i.e., BER = $1/10^{10} = 10^{-10}$) using long pseudo-random bit patterns. In contrast, data communications protocols like Fiber Channel and Ethernet commonly specify a BER of better than 10^{-12} using shorter bit patterns. In some cases, system specifications require a BER of 10^{-16} or lower.

It is important to note that BER is essentially a statistical average and is therefore only valid for a sufficiently large number of bits. It is possible, for example, to have more than one error within a group of, say 10^{10} bits, and still meet a 10^{-10} BER specification when the total number of transmitted bits is much greater that 10^{10}. This could happen if there is less than 1 error in 10^{10} bits for subsequent portions of the bit stream. Alternately, it is possible to have zero errors within a group of 10^{10} bits and still violate

a 10^{-10} standard if there are more errors in subsequent portions of the bit stream. In light of these examples, it is clear that a system that specifies a BER better than 10^{-10} must be tested by transmitting significantly more than 10^{10} bits in order to get an accurate and repeatable measurement. A natural and common question is "How many bits do I need to transmit through the system in order to prove BER compliance?" The answer to this question is the subject of section 3.

2. Equipment and Procedures

The conventional method for BER testing utilizes a pattern generator and an error detector (figure 1). The pattern generator transmits the test pattern into the system under test. The error detector either independently generates the same test pattern or receives it from the pattern generator. The pattern generator also supplies a synchronizing clock signal to the error detector. The error detector performs a bit-for-bit comparison between the data received from the system under test and the data received from the pattern generator. Any differences between the two sets of data are counted as bit errors.
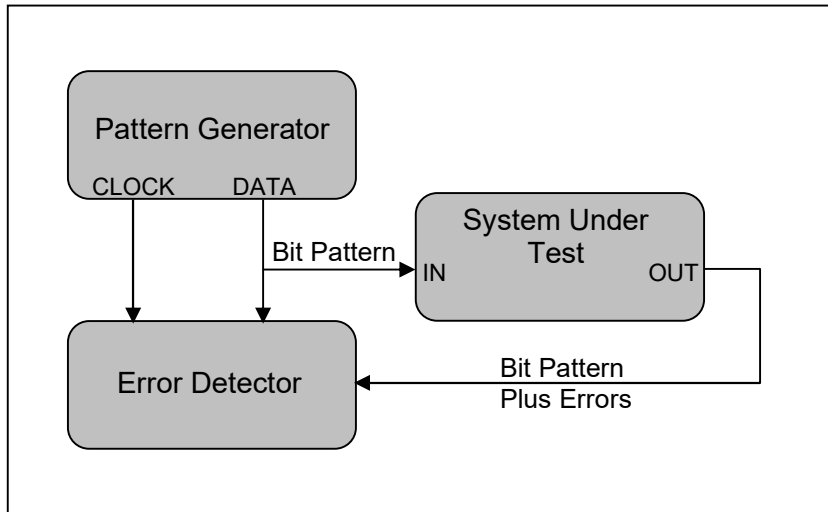

Figure 1. Test equipment set-up for BER test

As noted in the previous section, digital communication standards typically specify the data pattern to be used for BER testing. A test pattern is usually chosen that emulates the type of data that is expected to occur during normal operation, or, in some cases, a pattern may be chosen that is particularly stressful to the system for "worst-case" testing. Patterns intended to emulate random data are called pseudo-random bit sequences (PRBSs) and are based on standardized generation algorithms. PRBS patterns are classified by the length of the pattern and are generally referred to by such names as "2^7-1" (pattern length = 127 bits) or "2^{23}-1" (pattern length = 8,388,607). Other patterns emulate coded/scrambled data or stressful data sequences and are given names like "K28.5" (used by Fibre Channel and ethernet), etc. Commercially available pattern generators include standard built-in patterns as well as the capability to create custom patterns.

In order to accurately compare the bits received from the pattern generator to the bits received from the system under test, the error detector must be synchronized to both bit streams and it must compensate for the time delay through the system under test. A clock signal from the pattern generator provides synchronization for the bits received from the pattern generator. The error detector adds a variable time delay to the pattern generator clock to allow synchronization with bits from the system under test. As part of the pre-test system calibration, the variable time delay is adjusted to minimize bit errors.

3. How Many Bits?

In a well designed system, BER performance is limited by random noise and/or random jitter. The result is that bit errors occur at random (unpredictable) times that can be bunched together or spread apart. Consequently, the number of errors that will occur over the lifetime of the system is a random variable that cannot be predicted exactly. The true answer to the question of how many bits must be transmitted through the system for a perfect BER test is therefore unbounded (essentially infinite).

Since practical BER testing requires finite test times, we must accept less than perfect estimation. As previously noted, the quality of the BER estimation increases as the total number of transmitted bits increases. The problem is how to quantify the increased quality of the estimate so that we can determine how many transmitted bits are sufficient for the desired estimate quality. This can be done using the concept of statistical confidence levels. In statistical terms, the BER confidence level can be defined as the probability, based on E detected errors out of N transmitted bits, that the "true" BER would be less than a specified ratio, R. (For purposes of this definition, true BER means the BER that would be measured if the number of transmitted bits was infinite.) Mathematically, this can be expressed as

$$CL = \text{PROB } [BER_T < R] \text{ given } E \text{ and } N \tag{1}$$

where CL represents the BER confidence level, PROB[ ] indicates "probability that," and BER_T is the true BER. Since confidence level is, by definition, a probability, the range of possible values is 0 to 100%. Once the BER confidence level has been computed, we may say that we have CL percent confidence that the true BER is less than R. Another interpretation is that, if we were to repeatedly transmit the same number of bits, N, through the system and count the number of detected errors, E, each time we repeated the test, we would expect the resulting BER estimate, E/N, to be less than R for CL percent of the repeated tests.

As interesting as equation 1 is, what we really want to know is how to turn it around so that we can calculate how many bits need to be transmitted in order to calculate the BER confidence level. To do this we make use of statistical methods involving the binomial distribution function and Poisson theorem. Details of the calculations and derivations are beyond the scope of this article, but they can be found in reference [1]. The resulting equation is

$$N = \frac{1}{BER}\left[-\ln(1-CL) + \ln\left(\sum_{k=0}^{E}\frac{(N \times BER)^k}{k!}\right)\right] \tag{2}$$

where E represents the total number of errors detected and ln[ ] is the natural logarithm. When there are no errors detected (i.e., $E = 0$) the second term in equation 2 is equal to zero and the solution to the equation is greatly simplified. When E is not zero, equation 2 can still be solved empirically (using a computer spreadsheet, for example).

As an example of how to use equation 2, let us assume that we want to determine how many bits must be transmitted error free through a system for a 95% confidence level that the true BER is less than 10^{-10}. In this example $E = 0$, so the second term (the summation) is zero, and we only have to be concerned with CL and BER. The result is N = 1/BER x [-ln(1-0.95)] $\approx$ 3/BER = 3 x 10^{10}. This result illustrates a simple "rule of thumb," which is that the transmission of three times the reciprocal of the specified BER without an error gives a 95% confidence level that the system meets the BER specification. Similar calculations show that $N = 2.3$/BER for 90% confidence or 4.6/BER for 99% confidence if no errors are detected.

Figure 2 illustrates the relationship between the number of bits that must be transmitted (normalized to the BER) versus confidence level for zero, one, and two bit errors. Results for commonly used confidence levels of 90%, 95%, and 99% are tabulated in Table 1. To use the graph of figure 2, select the desired confidence level and draw a vertical line up from that point on the horizontal axis until it intersects the curve for the number of errors detected during the test. From that intersection point, draw a horizontal line to the left until it intersects the vertical axis to determine the normalized number of bits that must be transmitted, N x BER. Divide this number by the specified BER to get the number of bits that must be transmitted for desired confidence level.

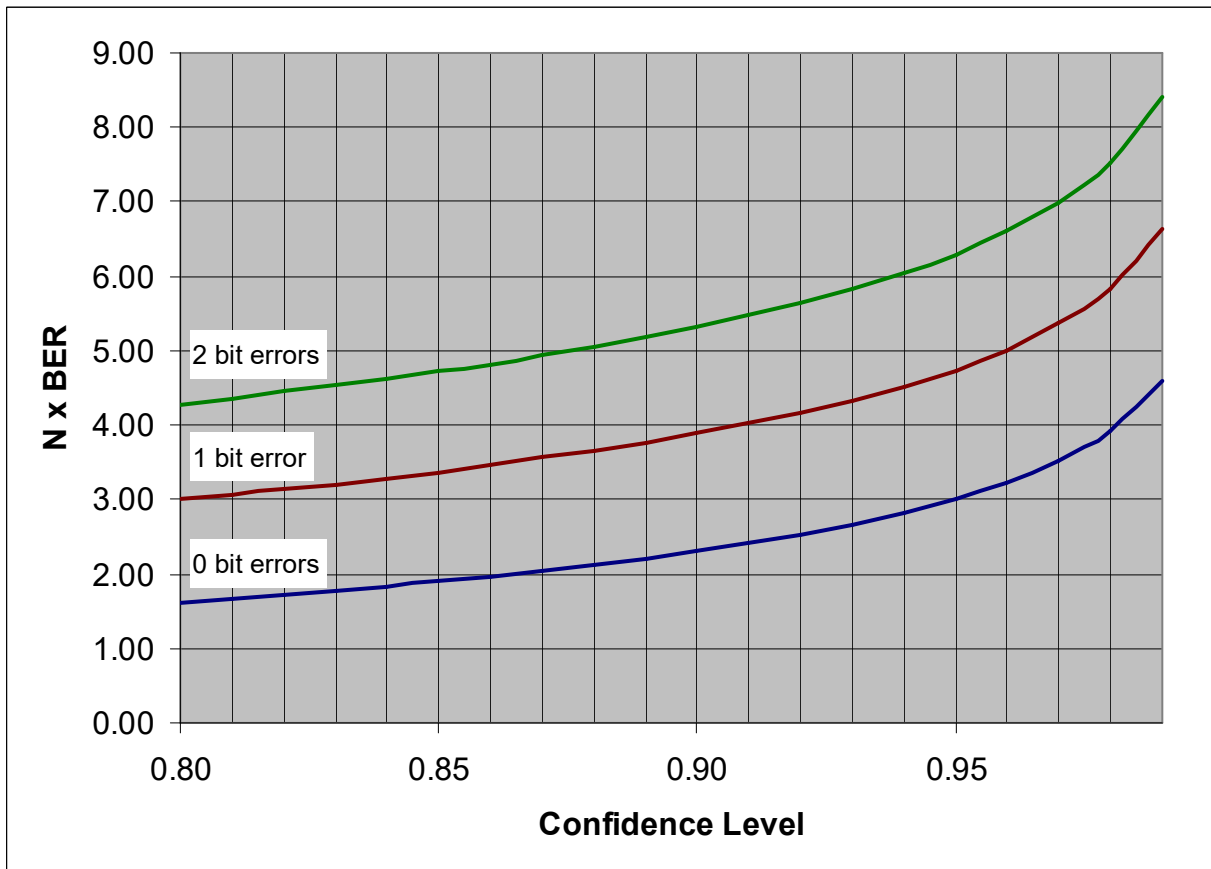Figure 2. Transmitted bits (normalized to the BER) versus confidence level for 0, 1, and 2 bit errors

TABLE 1: N x BER for confidence levels of 90%, 95%, and 99%

Errors	N x BER		
	CL = 90%	CL = 95%	CL = 99%
0	2.30	3.00	4.61
1	3.89	4.74	6.64
2	5.32	6.30	8.40

4. Reducing Test Time

Tests that require a high confidence level and/or low BER may take a long time, especially for low data rate systems. Consider a 99% confidence level test for a BER of 10^{-12} on a 622Mbps system. From Table 1, the required number of bits is 4.61 x 10^{12} for zero errors. At 622Mbps, the test time would be 4.6 x 10^{12} bits/622 x 10^6 bits/sec = 7,411 sec, which is slightly more than two hours. Two hours is generally too long for a practical test, but what can be done to reduce the test time?

One common method of shortening test time involves intentional reduction of the signal-to-noise (SNR) of the system by a known quantity during testing. This results in more bit errors and a quicker measurement of the resulting degraded BER (see reference [2]). If we know the relationship between SNR and BER, then the degraded BER results can be extrapolated to estimate the BER of interest. Implementation of this method is based on the assumption that thermal (Gaussian) noise at the input to the receiver is the dominant cause of bit errors in the system.

The relationship between the SNR and BER can be derived using Gaussian statistics and is documented in many communications text books [3]. While there is no known closed-form solution to the SNR-BER relationship, results can be obtained through numerical integration. One convenient method to compute this relationship is to use the Microsoft Excel™ standard normal distribution function, NORMSDIST[ ]. Using this function, the relationship between SNR and BER can be computed as:

BER = 1- NORMSDIST(SNR/2)   (3)

Figure 3 is a graphical representation of this relationship.

Figure 3. The relationship between BER and SNR

To illustrate this method of accelerated testing, we refer to the example at the beginning of the section. In that example a 99% confidence level test for a BER of 10^{-12} on a 622Mbps system would take over two hours. From figure 3 we see that a BER of 10^{-12} corresponds to an SNR of approximately 14. In the communication system under test, we can interrupt the signal channel between the transmitter and the receiver and insert an attenuator. Since the signal is attenuated prior to its input to the receiver, then, based on the assumption that the dominant noise source is at the receiver input, we will attenuate the signal but not the noise. Therefore the SNR will be reduced by the same amount as the signal. (Note, however, that it is important to ensure that the signal is not attenuated below the noise level of the channel.) For this example, we reduce the SNR from 14 to 12 by using a 14.3% (0.67dB) attenuation. From figure 3 we note that reducing the SNR to 12 corresponds to changing the BER to 10^{-9}. For a 99% BER confidence level at a BER of 10^{-9}, we need to transmit 4.61 x 10^9 bits (a factor of a 1,000 less that the original test) for a test time of 7.41 seconds. So, if we test for 7.41 seconds with no errors using the attenuator, then, by extrapolation, we determine that when the attenuation is removed, the BER should be 10^{-12}. Sounds great, right?

As in all things, reducing the test time through SNR reduction and extrapolation does not come without a price. The price is reduced confidence level after the extrapolation, and the reduction in confidence level becomes more significant as the extrapolation distance becomes larger. To visualize this effect, consider a test where SNR attenuation results in a reduction in BER by a factor of 100. If the SNR attenuated test is done to a 99% confidence level with zero errors, then we would expect that repeating the test 100 times would result in 99 tests with zero errors and one test with a single error. Now, if we concatenate all of the received bits from the 100 repeated tests, we would have 100 times as many bits with one error. Extrapolation of results from the 100 repeated tests to the original, non-reduced BER level gives one error in 1/BER bits or N x BER = 1.0. Using equation 2, we see that the corresponding confidence level is only 63%, low enough that it is off the chart of figure 2 and nowhere near the 99% confidence level we began with.

In light of the example above, the SNR should be attenuated as little as possible to achieve a practical test time. It must be realized that extrapolation will reduce the confidence level. Also, measurements and calculations must be performed with extra precision since errors introduced due to rounding, measurement tolerances, etc. will be multiplied when the results are extrapolated.

5. Conclusion

The bit error ratio of a digital communication system is an important figure of merit used to quantify the integrity of data transmitted through the system. Testing for a finite length of time yields an estimate of the probability that a bit will be received in error. The quality of the estimate improves as the test time increases and this quality can be quantified using statistical confidence level methods. Ideas for reducing the test time have been published, but they should be used with great care since they can significantly decrease the confidence level.

_____

REFERENCES:

[1] J. Redd, "Calculating Statistical Confidence Levels for Error-Probability Estimates," *Lightwave Magazine*, pp. 110-114, April 2000. (Available on the web at http://lw.pennnet.com).

[2] D.H. Wolaver, "Measure Error Rates Quickly and Accurately," *Electronic Design*, pp.89-98, May 30, 1995.

[3] B. Sklar, *Digital Communications: Fundamentals and Applications*, Englewood Cliffs, New Jersey: Prentice Hall, pp. 773-743, 1988.