Fibre Channel Consortium

FC-PI-5 Clause 5
Bit Error Rate Test Suite
Version 1.0

Technical Document

Last Updated: September 30, 2014

Fibre Channel Consortium
InterOperability Laboratory
University of New Hampshire

121 Technology Drive, Suite 2
Durham, NH 03824
Phone: +1-603-862-0090
Fax: +1-603-862-4181
http://www.iol.unh.edu/consortiums/fc

Table of Contents

Modification Record

- August 14, 2014 – Version 0.1

  Daniel Gray: Created test suite for testing the Bit Error Rate defined in FC-PI-5. This document is based on the FC-PI-4 test suite, modified to align to current specifications defined by FC-PI-5.

Acknowledgments

The University of New Hampshire would like to acknowledge the efforts of the following individuals in the development of this test suite.

Daniel Gray	University of New Hampshire
Daniel Reynolds	University of New Hampshire

Introduction

Overview

The University of New Hampshire's InterOperability Laboratory (UNH-IOL) is an institution designed to improve the interoperability of standards based products by providing an environment where a product can be tested against other implementations of a standard. This particular suite of tests has been developed to help implementers evaluate the Physical Layer functionality of their optical Fibre Channel products.

These tests are designed to determine if a Fibre Channel product conforms to specifications of the **FC-PI-5 Rev 6.10** Fibre Channel Standard (hereafter referred to as "FC-PI-5"). Successful completion of all tests contained in this suite does not guarantee that the tested device will operate with other devices. However, combined with satisfactory operation in the IOL's interoperability test bed, these tests provide a reasonable level of confidence that the device under test (DUT) will function properly in many Fibre Channel environments.

Organization of Tests

The tests contained in this document are organized to simplify the identification of information related to a test and to facilitate in the actual testing process. Each test contains an identification section that describes the test and provides cross-reference information. The discussion section covers background information and specifies why the test is to be performed. Tests are grouped in order to reduce setup time in the lab environment. Each test contains the following information:

Test Number

The Test Number associated with each test follows a simple grouping structure. Listed first is the Clause followed by the Test Group Number followed by the test's number within the group. This allows for the addition of future tests to the appropriate groups of the test suite without requiring the renumbering of the subsequent tests.

Purpose

The purpose is a brief statement outlining what the test attempts to achieve. The test is written at the functional level.

References

This section specifies all reference material *external* to the test suite, including the specific subclauses references for the test in question, and any other references that might be helpful in understanding the test methodology and/or test results. External sources are always referenced by a bracketed number (e.g., [1]) when mentioned in the test description. Any other references in the test description that are not indicated in this manner refer to elements within the test suite document itself (e.g., "Appendix 6.A", or "Table 6.1.1-1")

Resource Requirements

The requirements section specifies the test hardware and/or software needed to perform the test. This is generally expressed in terms of minimum requirements, however in some cases specific equipment manufacturer/model information may be provided.

Last Modification

This specifies the date of the last modification to this test.

Discussion

The discussion covers the assumptions made in the design or implementation of the test, as well as known limitations. Other items specific to the test are covered here.

Test Setup
The setup section describes the initial configuration of the test environment. Small changes in the configuration should be included in the test procedure.

Procedure
The procedure section of the test description contains the systematic instructions for carrying out the test. It provides a cookbook approach to testing, and may be interspersed with observable results.

Observable Results
This section lists the specific observable results that can be examined by the tester in order to verify that the DUT is operating properly. When multiple values for an observable are possible, this section provides a short discussion on how to interpret them. The determination of a pass or fail outcome for a particular test is often based on the successful (or unsuccessful) detection of a certain observable.

Possible Problems
This section contains a description of known issues with the test procedure, which may affect test results in certain situations. It may also refer the reader to test suite appendices and/or whitepapers that may provide more detail regarding these issues.

Group 1: Bit Error Rate

Overview:
This group of tests verifies the Electrical and Optical Bit Error Rate (BER), as defined in Clause 5 of FC-PI-5. These tests cannot provide a guarantee that the device conforms to the Bit Error Rate, in all conditions, as defined in Clause 5 of FC-PI-5, but can provide confidence that the device and the Bit Error Rate Tester (BERT) conform to the applicable Bit Error Rate.

Test #5.1.1: Electrical Bit Error Rate (BER)

Purpose:

- To verify that the Bit Error Rate of the DUT is within the conformance limit.

References:

[1] FC-PI-5 - Clause 5.1
[2] Ibid – Clause 9.1
[3] FC-MJSQ – Annex A
[4] FC-MSQS – Clause 9

Resource Requirements:

- Electrical J-BERT Capable of handling up to 16GFC signaling
- (4) SMA/SMA cables

Last Updated: May 04, 2009

Discussion:

In order to ensure the overall quality of the Fibre Channel link, it is important to establish a maximum allowable Bit Error Rate (BER) that devices must achieve. The FC-FS-3 protocol is defined to operate across connections having a bit error ratio (BER) detected at the receiving port of less than 10^{-12}. The BER objective is 10^{-12}. A TxRx Connection bit error rate (BER) of $\leq 10^{-12}$ as measured at its receiver is supported. The basis for the BER is the encoded serial data stream on the transmission medium during system operation.

All Fibre Channel TxRx Connections described in reference [2] shall operate within the BER objective (10^{-12}). Patterns for this test are defined by FC-MJSQ [3] for 4GFC and FC-MSQS [4] for 8GFC and 16GFC. The implementation of specific patterns is mentioned in the Procedure, and is determined by the speed of the device under test.

Test Setup:

The DUT should be setup as defined in Appendix A. Configure the DUT for the appropriate speed. The DUT should be transitioned into the monitoring/active state. If the DUT supports scrambling: the scrambler / de-scrambler must be disabled.

Procedure:

a) L_Port/Nx_Port/Fx_Port Device
   1) Instruct the J-BERT to transmit a LPB to the DUT.
   2) Instruct the J-BERT to begin sourcing Idle continuously.
   3) Measure the Bit Error Rate over the transmission of 3×10^{12} transmitted bits.
   4) Repeat steps 1-3 for each port.
   5) For 4G and 8G (not supporting scrambling) Devices
      1. Repeat steps 1-4 for CJTPAT, CRPAT and CSPAT.
   6) For 8G Devices supporting scrambling
      1. Repeat steps 1-4 for JSPAT and JTSPAT
   7) For 16G Devices
      1. Repeat steps 1-4 using PRBS9, PRBS31, and Scrambled Idle

b) N_Port/F_Port Device
   1) Setup the DUT to retransmit unaltered traffic through the TxRx Connection.
   2) Instruct the J-BERT to begin sourcing Idle continuously.
   3) Measure the Bit Error Rate over the transmission of 3×10^{12} transmitted bits.
   4) Repeat steps 1-3 for each port.
   5) For 4G and 8G (not supporting scrambling) Devices
      1. Repeat steps 1-4 for CJTPAT, CRPAT and CSPAT.
   6) For 8G Devices supporting scrambling
      1. Repeat steps 1-4 for JSPAT and JTSPAT
   7) For 16G Devices
      1. Repeat steps 1-4 using PRBS9, PRBS31, and Scrambled Idle

Observable Results:

The Bit Error Rate (BER) over all traffic shall not exceed 10^{-12} (i.e. not errors shall be detected over the 3×10^{12} transmitted bits).

Possible Problems: None

Test #5.1.2: Optical Bit Error Rate (BER)

Purpose:

- To verify that the Bit Error Rate of the DUT is within the conformance limit.

References:

[1] FC-PI-5 - Clause 5.1
[2] Ibid – Clause 6.1
[3] FC-MJSQ – Annex A
[4] FC-MSQS – Clause 9

Resource Requirements:

- Optical J-BERT Capable of handling up to 16GFC signaling
- (1) Fiber Cable

Last Updated: May 04, 2009

Discussion:

In order to ensure the overall quality of the Fibre Channel link, it is important to establish a maximum allowable Bit Error Rate (BER) that devices must achieve. The FC-FS-3 protocol is defined to operate across connections having a bit error ratio (BER) detected at the receiving port of less than 10^{-12}. The BER objective is 10^{-12}. A TxRx Connection bit error rate (BER) of $\leq 10^{-12}$ as measured at its receiver is supported. The basis for the BER is the encoded serial data stream on the transmission medium during system operation.

Fibre Channel links shall not exceed the BER objective (10^{-12}) under any compliant conditions. Patterns for this test are defined by FC-MJSQ [3] for 4GFC and FC-MSQS [4] for 8GFC and 16GFC. The implementation of specific patterns is mentioned in the Procedure, and is determined by the speed of the device under test.

Test Setup:

The DUT should be setup as defined in Appendix A. Configure the DUT for the appropriate speed. The DUT should be transitioned into the monitoring/active state.

Procedure:

a) L_Port/Nx_Port/Fx_Port Device
   1) Instruct the J-BERT to transmit a LPB to the DUT.
   2) Instruct the J-BERT to begin sourcing Idle continuously.
   3) Measure the Bit Error Rate over the transmission of 3×10^{12} transmitted bits.
   4) Repeat steps 1-3 for each port.
   5) For 4G and 8G (not supporting scrambling) Devices
      1. Repeat steps 1-4 for CJTPAT, CRPAT and CSPAT.
   6) For 8G Devices supporting scrambling
      1. Repeat steps 1-4 for JSPAT and JTSPAT
   7) For 16G Devices
      1. Repeat steps 1-4 using PRBS9, PRBS31, and Scrambled Idle

b) N_Port/F_Port Device
   1) Setup the DUT to retransmit unaltered traffic through the TxRx Connection.
   2) Instruct the J-BERT to begin sourcing Idle continuously.
   3) Measure the Bit Error Rate over the transmission of $3x10^{12}$ transmitted bits.
   4) Repeat steps 1-3 for each port.
   5) For 4G and 8G (not supporting scrambling) Devices
      1. Repeat steps 1-4 for CJTPAT, CRPAT and CSPAT.
   6) For 8G Devices supporting scrambling
      1. Repeat steps 1-4 for JSPAT and JTSPAT
   7) For 16G Devices
      1. Repeat steps 1-4 using PRBS9, PRBS31, and Scrambled Idle

Observable Results:
The Bit Error Rate (BER) over all traffic shall not exceed 10^{-12} (i.e. not errors shall be detected over the $3x10^{12}$ transmitted bits).

Possible Problems: None.

Appendix A: Test Setup

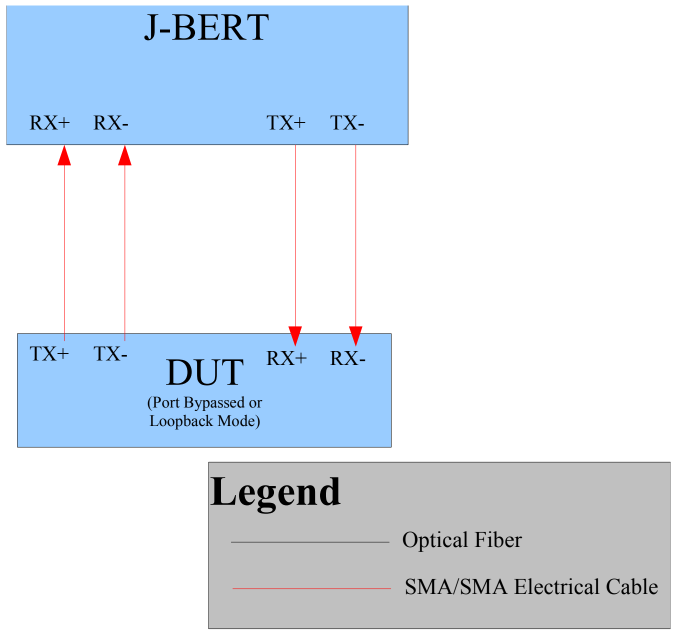

Figure 1: Electrical Test Setup

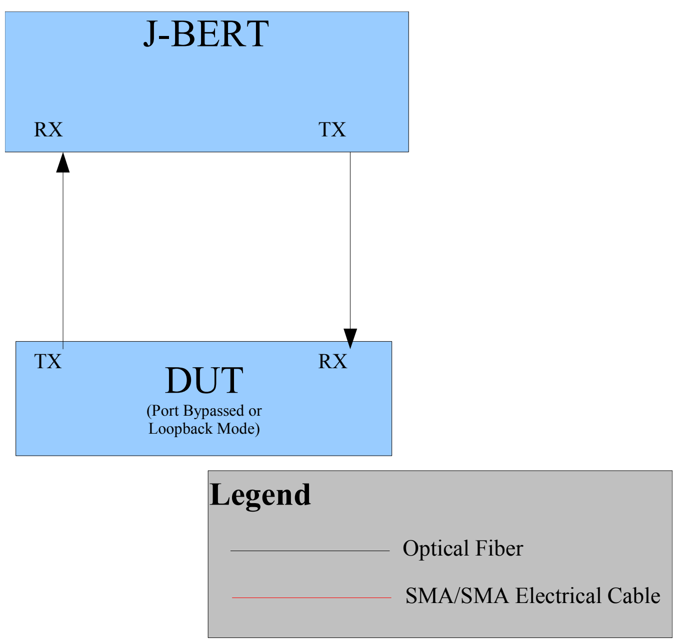

Figure 2: Optical Test Setup

Appendix B: Test Patterns

References:
[1] MJSQ – Table A.9, A.11, A.13
[2] MSQS – Clause 9.1.2, 9.1.3

For 16GFC devices, patterns used are defined by Clause 9 of FC-MSQS. 4GFC and 8GFC patterns used are as follows:

	Primitive				Count
(Idle)	BC	95	B5	B5	6
(SOFn3)	BC	B5	36	36	1
	7E	7E	7E	7E	41
	7E	7E	7E	74	1
	7E	AB	B5	B5	1
	B5	B5	B5	B5	12
	B5	5E	4A	7E	1
	7E	7E	7E	FE	1
(CRC)	F5	2E	F6	DD	1
(EOFn)	BC	B5	D5	D5	1

Table 1 - CJTPAT (JTPAT in a FC compliant frame format)

	Primitive				Count
(Idle)	BC	95	B5	B5	6
(SOFn3)	BC	B5	36	36	1
	BE	D7	23	47	
	6B	8F	B3	14	16
	5E	FB	35	59	
(CRC)	EE	23	55	16	1
(EOFn)	BC	B5	D5	D5	1

Table 2 - CRPAT (RPAT in a FC compliant frame format)

	Primitive				Count
(Idle)	BC	95	B5	B5	6
(SOFn3)	BC	B5	36	36	1
	7F	7F	7F	7F	512
(CRC)	F1	96	DB	97	1
(EOFn)	BC	95	D5	D5	1

Table 3 - CSPAT (SPAT in a FC compliant frame format)

Table F.1 – Scrambled jitter pattern (JSPAT)

D1.4	D16.2	D24.7	D30.4	D9.6	D10.5
0111010010	0110110101	0011001110	1000011101	1001010110	0101011010
D16.2	D7.7	D24.0	D13.3	D23.4	D13.2
1001000101	1110001110	0011001011	1011000011	0001011101	1011000101
D13.7	D1.4	D7.6	D0.2	D21.5	D22.1
1011001000	0111010010	1110000110	1001110101	1010101010	0110101001
D23.4	D20.0	D27.1	D30.7	D17.7	D4.3
0001011101	0010110100	1101101001	1000011110	1000110001	1101010011
D6.6	D23.5	D7.3	D19.3	D27.5	D19.3
0110010110	0001011010	1110001100	1100101100	1101101010	1100100011
D5.3	D22.1	D5.0	D15.5	D24.7	D16.3
1010010011	0110101001	1010010100	0101111010	0011001110	1001001100
D1.2	D23.5	D20.7	D11.7	D20.7	D18.7
0111010101	0001011010	0010110111	1101001000	0010110111	0100110001
D29.0	D16.6	D25.3	D1.0	D18.1	D30.5
1011100100	0110110110	1001100011	1000101011	0100111001	1000011010
D5.2	D21.6				
1010010101	1010100110				

Figure 3: JSPAT (scrambled jitter pattern) [2]

Table F.2 – Jitter tolerance scrambled pattern (JTSPAT)

D1.4	D16.2	D24.7	D30.4	D9.6	D10.5
0111010010	0110110101	0011001110	1000011101	1001010110	0101011010
D16.2	D7.7	D24.0	D13.3	D23.4	D13.2
1001000101	1110001110	0011001011	1011000011	0001011101	1011000101
D13.7	D1.4	D7.6	D0.2	D21.5	D22.1
1011001000	0111010010	1110000110	1001110101	1010101010	0110101001
D23.4	D20.0	D27.1	D30.7	D17.7	D4.3
0001011101	0010110100	1101101001	1000011110	1000110001	1101010011
D6.6	D23.5	D7.3	D19.3	D27.5	D19.3
0110010110	0001011010	1110001100	1100101100	1101101010	1100100011
D5.3	D22.1	D5.0	D15.5	D24.7	D16.3
1010010011	0110101001	1010010100	0101111010	0011001110	1001001100
D1.2	D23.5	D29.2	D31.1	D10.4	D4.2
0111010101	0001011010	1011100101	0101001001	0101011101	0010100101
D5.5	D10.2	D21.5	D10.2	D21.5	D20.7
1010011010	0101010101	1010101010	0101010101	1010101010	0010110111
D11.7	D20.7	D18.7	D29.0	D16.6	D25.3
1101001000	0010110111	0100110001	1011100100	0110110110	1001100011
D1.0	D18.1	D30.5	D5.2	D21.6	D1.4
1000101011	0100111001	1000011010	1010010101	1010100110	0111010010
D16.2	D24.7	D30.4	D9.6	D10.5	D16.2
0110110101	0011001110	1000011101	1001010110	0101011010	1001000101
D7.7	D24.0	D13.3	D23.4	D13.2	D13.7
1110001110	0011001011	1011000011	0001011101	1011000101	1011001000
D1.4	D7.6	D0.2	D21.5	D22.1	D23.4
0111010010	1110000110	1001110101	1010101010	0110101001	0001011101
D20.0	D27.1	D30.7	D17.7	D4.3	D6.6
0010110100	1101101001	1000011110	1000110001	1101010011	0110010110
D23.5	D7.3	D19.3	D27.5	D19.3	D5.3
0001011010	1110001100	1100101100	1101101010	1100100011	1010010011
D22.1	D5.0	D15.5	D24.7	D16.3	D1.2
0110101001	1010010100	0101111010	0011001110	1001001100	0111010101
D23.5	D27.3	D3.0	D3.7	D14.7	D28.3
0001011010	1101100011	1100010100	1100011110	0111001000	0011101100
D30.3	D30.3	D7.7	D7.7	D20.7	D11.7
0111100011	1000011100	1110001110	0001110001	0010110111	1101001000
D20.7	D18.7	D29.0	D16.6	D25.3	D1.0
0010110111	0100110001	1011100100	0110110110	1001100011	1000101011
D18.1	D30.5	D5.2	D21.6		
0100111001	1000011010	1010010101	1010100110		

Figure 4: JTSPAT (Jitter tolerance scrambled pattern) [2]

Appendix C: Bit Error Rate Measurement

Purpose: To develop a procedure for bit error rate measurement through the application of statistical methods.

References:

[1] Miller, Irwin and John E. Freund, Probability and Statistics for Engineers (Second Edition), Prentice-Hall, 1977, pp. 194-210, 240-245.

Last Modification: November 4, 2004 (Version 1.0)

Discussion:

B.1 – Introduction

One key performance parameter for all digital communication systems is the bit error rate (BER). The bit error rate is the probability that a given bit will be received in error. The BER may also be interpreted as the average number of errors that would occur in a sequence of n bits.

While the bit error rate concept is quite simple, the measurement of this parameter poses some significant challenges. The first challenge is deciding the number of bits, n, that must be sent in order to make a reliable measurement. For example, if 10 bits were sent and no errors were observed, it would be foolish to conclude that the bit error rate is zero. However, common sense tells us that the more bits that are sent without error, the more reasonable this conclusion becomes. In the interest of keeping the test duration as short as possible, we want to send the smallest number of bits that provides us with an acceptable margin of error.

This brings us to the second challenge of BER measurement. Given that we send n bits, what reasonable statements can be made about the bit error rate based on the number of errors observed? Returning to the previous example, if 10 bits are sent and no errors are observed, it is unreasonable to say that the BER is zero. However, it may be more reasonable to say that the BER is 10^{-1} or better. Furthermore, you are absolutely certain that the bit error rate is not 1.

In this appendix, two statistical methods, hypothesis testing and confidence intervals, are applied to help us answer the questions of how many bits we should be sent and what conclusions can be made from the test results.

B.2 – Statistical Model

A statistical model for the number of errors that will be observed in a sequence of n bits must be developed before we apply the aforementioned statistical methods. For this model, we will assume that every bit received is an independent Bernoulli trial. A Bernoulli trial is a test for which there are only two possible outcomes (i.e. a coin toss). Let us say that p is the probability that a bit error will occur. This implies that the probability that a bit error will not occur is (1-p).

The property of independence implies that the outcome of one Bernoulli trial has no effect on the outcomes of the other Bernoulli trials. While this assumption is not necessarily true for all digital communications systems, it is still used to simplify the analysis.

The number of successful outcomes, k, in n independent Bernoulli trials is taken from a binomial distribution. The binomial distribution is defined in equation B-1.

$$b(k;n,p)=C_{n,k}\,p^{k}(1-p)^{n-k}$$ (Equation B-1)

Note that in this case, a successful outcome is a bit error. The coefficient $C_{n,k}$ is referred to as the binomial coefficient or “n-choose-k”. It is the number of combinations of k successes in n trials. Returning to coin toss analogy, there are 3 ways to get 2 heads from 3 coin tosses: (tails, heads, heads), (heads, tails, heads), and (heads, heads, tails). Therefore, $C_{3,2}$ would be 3. A more precise mathematical definition is given in equation B-2.

$$c_{n,k}=\frac{n!}{k!(n-k)!}$$ (Equation B-2)

This model reflects the fact that for a given probability, p, a test in which n bits are sent could yield many possible outcomes. However, some outcomes are more likely than others and this likelihood principle allows us to make conclusions about the BER for a given test result.

B.3 – Hypothesis Test

The statistical method of hypothesis testing will allow us to establish a value of n, the number of bits to be sent, for the BER measurement. Naturally, the test begins with a hypothesis. In this case, we will hypothesize that the probability of a bit error, p, for the system is less than some target BER, P_0. This hypothesis is stated formally in equation B-3.

$$H_0\colon p \leqslant P_0$$ (Equation B-3)

We now construct a test for this hypothesis. In this case, we will take the obvious approach of sending n bits and counting the number errors, k. We will interpret the test results as shown in table B-1.

Table B-1: Acceptance and rejections regions for H_0

Test Result	**Conclusion**
k = 0	H_0 is true
k > 0	H_0 is false

We now acknowledge the possibility that our conclusion is in error. Statisticians define two different categories of error. A type I error is made when the hypothesis is rejected even though it is true. A type II error is made when the hypothesis is accepted even though it is false. The probability of a type I and a type II error are denoted as α and β respectively. Table B-2 defines type I and type II errors in the context of this test.

Table B-2: Definitions of type I and type II errors

Type I Error	$k > 0$ even though $p \leq BER$
Type II Error	$k = 0$ even though $p > BER$

A type II error is arguably more serious and we will define n so that the probability of a type II error, β, is acceptable. The probability of a type II error is given in equation B-4.

$$\beta = (1 - p)^n < (1 - P_0)^n \quad \text{(Equation B-4)}$$

Equation B-4 illustrates that the upper bound on the probability of a type II error is a function of the target bit error rate and n. By solving this equation for n, we can determine the minimum number of bits that need to sent in order to verify that p is less than a given P_0 for a given probability of type II error.

$$n > \frac{\ln(\beta)}{\ln(1 - P_0)} \quad \text{(Equation B-5)}$$

Let us now examine the probability of a type I error. The definition of α is given in equation B-6.

$$\alpha = 1 - (1 - p)^n \leq 1 - (1 - P_0)^n \quad \text{(Equation B-6)}$$

Equation B-6 shows that while we increase n to make β small, we simultaneously raise the upper bound on α. This makes sense since the likelihood of observing a bit error increases with the number of bits that you send, no matter how small bit error rate is. Therefore, while the hypothesis test is very useful in determining a reasonable value for n, we must be very careful in interpreting the results. Specifically, if we send n bits and observe no errors, we are confident that p is less than our target bit error rate (our level of confidence depends on how small we made β). However, if we do observe bit errors, we cannot be quick to assume that the system did not meet the BER target since the probability of a type I error is so large. In the case of $k > 0$, a confidence interval can be used to help us interpret k.

B.4 – Confidence Interval

The statistical method of confidence intervals will be used to establish a lower bound on the bit error rate given that $k > 0$. A confidence interval is a range of values that is likely to contain the actual value of some parameter of interest. The interval is derived from the measured value of the parameter, referred to as the point estimate, and the confidence level, $(1-\alpha)$, the probability that the parameter's actual value lies within the interval.

A confidence interval requires a statistical model of the parameter to be bounded. In this case, we use the statistical model for k given in equation B-1. If we were to compute the area under the binomial curve for some interval, we would be computing the probability that k lies within that interval. This concept is shown in figure B-1.

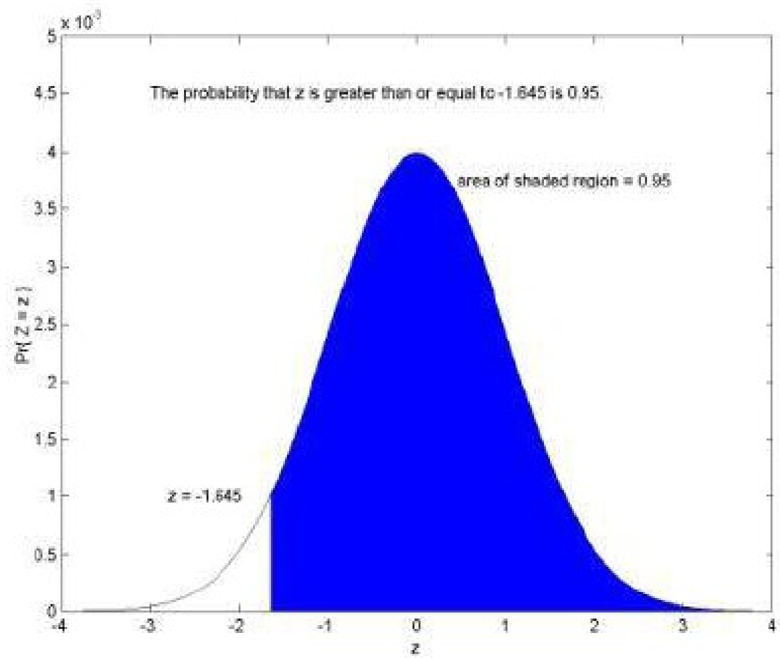

Figure B-1: Computing the probability that $z \geq -1.645$ (standard normal distribution).

To compute the area under the binomial curve, we need a value for the parameter p. To compute a confidence interval for k, you assume that k/n, the point estimate for p, is the actual value of p.

Note that figure B-1 illustrates the computation of the lower tolerance bound for k, a special case where the confidence interval is $[k_l, +\infty]$. A lower tolerance bound implies that in a percentage of future tests, the value of k will be greater than k_l. In other words, actual value of k is greater than k_l with probability equal to the confidence level. Therefore, if k_l/n is greater than P_0, we can say that the system does not meet the target bit error rate with probability $(1-\alpha)$. By reducing α, we reduce the probability of making a type I error.

To determine the value of k_l, it is useful to assume that the binomial distribution can be approximated by a normal (Gaussian) distribution when n is large. The mean and variance of this equivalent distribution are the mean and variance of the corresponding binomial distribution (given in equations B-7 and B-8).

$$\mu_k = np \qquad \text{(Equation B-7)}$$

$$\sigma_k^2 = np(1-p) \qquad \text{(Equation B-8)}$$

Now, let α be the probability that $Z \leq z_\alpha$ where Z is a standard normal random variable. A standard random variable is one whose mean is zero and whose variance is one. The random variable K can be standardized as shown in equation B-9.

$$Z = \frac{K - \mu_k}{\sigma_k} \qquad \text{(Equation B-9)}$$

Note that Z is greater than z_α with probability (1-α), the confidence level. We apply this inequality to equation B-9 and solve for K to get equation B-10.

$$\begin{aligned} K &> \mu_k + z_\alpha \sigma_k \\ K &> np + z_\alpha \sqrt{np(1-p)} \end{aligned} \qquad \text{(Equation B-10)}$$

As mentioned before, we assume that p is k/n. We can now generate an expression for k_l, the value that K will exceed with probability (1-α). This expression is given in equation B-11.

$$k_l = k + z_\alpha n \sqrt{\frac{(k/n)(1-k/n)}{n}} \qquad \text{(Equation B-11)}$$

Finally, we argue that if K exceeds k_l, then the actual value of p must exceed k_l/n. Therefore, we can generate an expression for p_l, the value that p will exceed with probability (1-α), and compare it to the target bit error rate. By applying this comparison (given in equation B-12) the probability of a type I error can be greatly reduced. For example, by setting z_α to –1.645, the probability of a type I error is reduced to 5%.

$$P_0 \geq p_l = \frac{k_l}{n} = \frac{k}{n} + z_\alpha \sqrt{\frac{(k/n)(1-k/n)}{n}}$$ (Equation B-12)

B.5 – Sample Test Construction

We now compress the theory presented in sections B-2 through B-4 into two inequalities that may be used to construct a bit error rate test. First, we take equation B-5 and assume that $\ln(1-P_0)$ is $-P_0$ (valid for P_0 much less than one). The result is equation B-13.

$$n > \frac{-\ln(\beta)}{P_0}$$ (Equation B-13)

Second, we examine equation B-12. Assuming that (1-k/n) is very close to 1 and substituting $-\ln(\beta)/P_0$ for n, we get equation B-14.

$$-\ln(\beta) \geq k + z_\alpha \sqrt{k}$$ (Equation B-14)

The largest value of k that satisfies equation B-14 is k_l. The benefit of these two equations is that a bit error rate test is uniquely defined by β and α and that the test scales with P_0. Table B-3 defines n and k_l in terms of β and α.

Table B-3: n and k_l as a function of β and α.

β	**-ln(β)**	**n**	**α**	**z_α**	**k_l**
0.10	2.30	$2.30/P_0$	0.10	-1.29	5
0.10	2.30	$2.30/P_0$	0.05	-1.65	6
0.05	3.00	$3.00/P_0$	0.05	-1.65	7
0.05	3.00	$3.00/P_0$	0.01	-2.33	10
0.01	4.60	$4.60/P_0$	0.05	-1.65	9
0.01	4.60	$4.60/P_0$	0.01	-2.33	13

As an example, let us construct a test to determine if a given system is operating at a bit error rate of 10^{-12} or better. Given that a 5% chance of a type I error is acceptable, the test would take the form of sending 3×10^{12} bits and counting the number of errors. If no errors are counted, we are confident that the BER was 10^{-12} or better.

Given that a 5% chance of a type II error is acceptable, we find that k_l is 7. If more than 7 errors are counted, we are confident that the bit error rate is greater than 10^{-12}. However, what if between 1 and 7 errors are counted? These cases may be handled several different ways. One option is to make a statement about the bit error rate (whether it is less than or greater than 10^{-12}) at a lower level of confidence. Another option would be to state that the test result is success since we cannot establish with an acceptable probability of error that the BER is greater than 10^{-12}. Such a statement implies that we failed to meet the burden of proof for the conjecture that the BER exceed 10^{-12}. Of course, the burden of proof could be shifted to the device under test which would imply that any outcome other than k = 0 would correspond to failure (the device under test failed to prove to us that the BER was no more than 10^{-12}). If neither of these solutions are acceptable, it is always an option to perform a more vigorous bit error rate test in order to clarify the result.