OmniBER 718 communications performance analyzer

Product Note

Protection-switching time measurements:

Understanding the OmniBER analyzer's measurement technique and the alternative test methods

Agilent Technologies

Introduction

The OmniBER analyzer's service disruption measurement is designed to accurately measure the time taken by a transmission system to perform an automatic protection switch when a transmission defect is detected by the system. The measurement supports testing of all protection switching architectures deployed in today's transmission networks. Result accuracy and reliability is based on a simple test principle, coupled with a specially designed error detector that measures the duration of error bursts associated with a protection switch event.

Contents

Test configuration and setup

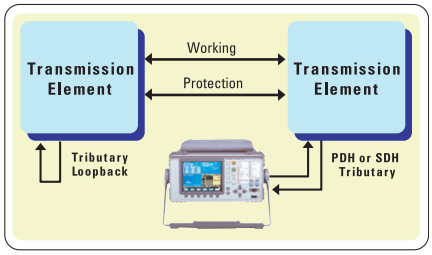

Figure 1: Test configuration for measuring protection-switching time

As shown in figure 1, the OmniBER analyzer measures protection-switching time from the tributary-side of the transmission system under test. This point is significant for two reasons:

1. The measurement is totally independent of the protection switching architecture and all optional settings – it therefore supports all configurations
2. The performance of the system under test cannot be affected by the test set since results are obtained through passively monitoring a PRBS for errors

The measurement is performed by inserting a PRBS test pattern into a tributary feeding the transmission system under test, looping this signal at the associated drop-side tributary, then monitoring the PRBS for errors. This PRBS can be inserted either as part of a PDH signal or as a mapped service within a SDH signal.

Measurement principle

The protection-switching time measurement is an out-of-service test that is typically performed during verification, installation and commissioning of new transmission systems. As such, the measurement principle assumes that the system under test will be operating error-free before the test is performed. This assumption is important since the presence of background errors can affect test results (see Understanding the Test Results for details).

The measurement is performed by, first, connecting the OmniBER analyzer as previously described and verifying error-free reception of the PRBS test pattern. Then, invoking a protection switch on a working section of the transmission system that is transporting the PRBS. Typical methods used to trigger this protection switch are disconnecting a fibre[1, 2] to simulate a fibre-break, or removing power from a transmission element to simulate a node-failure. As will be shown in the following STM-64 Ring case study, there is value in verifying the system's performance when tested separately using a simulated fibre-break and a simulated node-failure.

1 Exercise extreme caution when disconnecting an optical fibre – follow your organization's standard safety procedures.

2. Refer to Appendix 1 for a discussion on the issues associated with generating a loss-of-signal by disconnecting a fibre.

Irrespective of how the protection switch is triggered, it results in the PRBS test pattern being corrupted for a short period.

As shown in figure 2, the duration of this corruption is controlled by three factors:

1. The system's fault detection time
2. The protection-switching time
3. The time taken by the OmniBER analyzer to re-align to the pointers (SDH tributary only) and test pattern

Since the objective is to measure a transmission system's protection-switching time (specified as less than 50 ms in ITU-T standards), it is essential that the contributions made by these additional factors are minimized. For fault detection time, this is achieved by triggering the protection switch using a failure that results in a LOS defect. Although ITU-T G.783 (2000) defines LOS detection time as being "in the province of regional standards", it provides an example based on a value of less than 100 µs (less than 0.2% of the maximum acceptable protection-switching time). In the case of pointer and pattern acquisition, the required times are as shown in figure 2. These values represent the following percentage errors for PDH and SDH tributaries (with 140M & 2M mappings) when related to the maximum acceptable switching time:

PDH: +0.1%
SDH (140M): +0.35% to +0.85%
SDH (2M): +1.35% to 3.85%

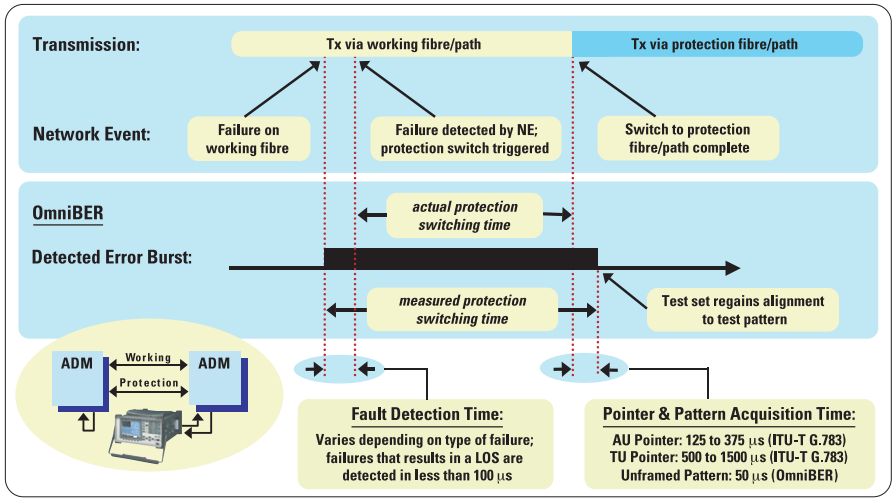

Figure 2: Contributors to the protection switching time as measured by OmniBER

When measuring a system's protection-switching time, the above discussion shows that the total systematic error associated with the OmniBER analyzer's service disruption measurement can be restricted to between +0.3% to +4.05% of the maximum acceptable switching time. Consequently, it can be relied on to accurately evaluate this important system specification.

Service disruption test results

Three separate results are provided by the OmniBER analyzer's existing service disruption measurement:

Longest: The duration of the longest error burst detected during the test
Shortest: The duration of the shortest error burst detected during the test
Last: The duration of the most recent error burst detected during the test

These result fields are reset to 0 ms at the beginning of a protection-switching time test by pressing the instrument's START button. When the protection switch is triggered, the duration of the resulting error burst is measured and displayed. For the system under test to pass, a single[1] error burst of duration less than 50 ms should be detected. Detection of a single error burst is indicated by an identical value being displayed in the three result fields.

Why are three separate results provided? During a detailed investigation into the measurement requirements, it was found that some transmission systems exhibited a characteristic similar to switch-bounce during a protection-switching event. This results in multiple distinct error bursts being present on the received test pattern. By providing three separate results, the OmniBER analyzer's service disruption measurement clearly identifies the presence of this unwanted operating characteristic.

New developments in the OmniBER 718 analyzer have been designed to make analysis of a protection-switch faster and easier, including:

Timestamping: Displaying a relative timestamp of the beginning of each service disruption event.

History: A record of the first ten service disruption measurements.

While an Alarm Indication Signal (AIS), is not a pure protection-switching measurement, it is closely related. AIS is activated as a result of any physical layer failure, such as a fibre break. The OmniBER 718 can now provide details of AIS duration measurements, timestamping and history. When used in conjunction with the OmniBER analyzer's service disruption measurements, it becomes possible to show a relationship between alarms within a device and automatic protection switches in a network. This means you can quickly and accurately debug network elements. For example, using the timestamping and AIS duration measurements, makes it easy to see if a device fails due to AIS not being raised quickly enough, or taking too long to be removed after the switch has taken place.

Understanding the test results

As is true for any measurement, having a working understanding of how the measured values are derived is helpful when attempting to identify the root-cause of unexpected results. In the case of the OmniBER analyzer's service disruption measurement this means understanding the rules associated with the analysis of error-burst duration.

The OmniBER analyzer's service disruption test measures the elapsed time between the first and last error in an error-burst that consists of two or more errors. The error-burst is taken as having ended when no errors are detected during a period of greater than 200 to 300 ms following the last error. Single errors that are separated by more than 200 to 300 ms are not considered as being part of an error-burst (no result is returned).

Figure 3 illustrates the affect these simple rules have on measurement results when different error distributions are present in the received test pattern. In case 1 and case 2, there are single errors due to a low background error rate in the transmission system, plus an error-burst associated with a protection-switching event.

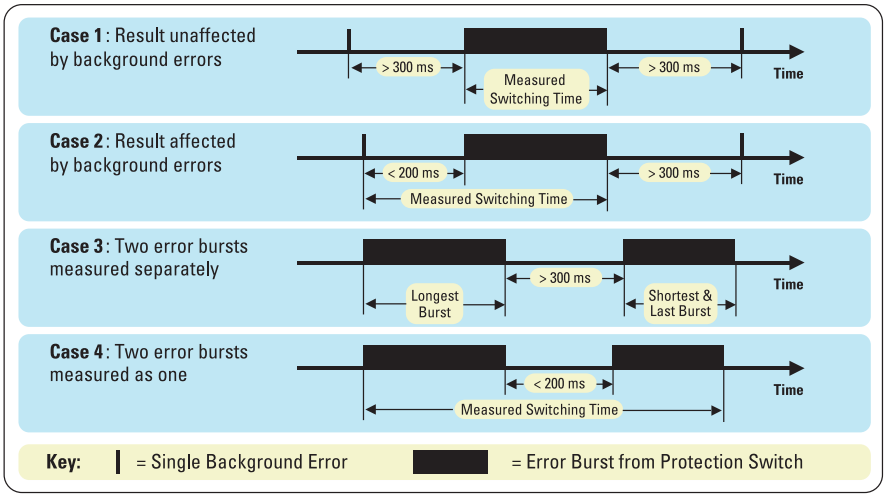

Figure 3: The OmniBER analyzer's error burst analysis

However, in case 1 the measured protection-switching time is not affected by the background errors since these occur outside the 200 to 300 ms period used to define the end of the error-burst. In contrast, the result obtained in case 2 is affected due to a single background error being present less than 200 ms before the error-burst actually starts. This leads to an artificially high protection-switching time being reported, and consequently emphasizes the importance of ensuring that the system under test is error-free before performing the measurement.

Cases 3 and 4 are based on a scenario where the system under test generates two error-bursts when performing a protection switch. The point being illustrated here is that the results obtained will be affected by the separation of these two error-bursts. In case 3 a result for each error-burst will be reported (since they are more than 300 ms apart), while in case 4 only a single high value will be reported (since they are less than 200 ms apart). However, in both cases the reported results will indicate that a problem exists in the system under test.

Case study of protection-switching time measurements on a STM-64 dual-ring

The following case study illustrates how the OmniBER analyzer is enabling the protection-switching time performance of STM-64 systems to be evaluated before deployment in today's operational networks. This particular example documents results obtained during an evaluation of the dual-ring transmission system shown in figure 4.

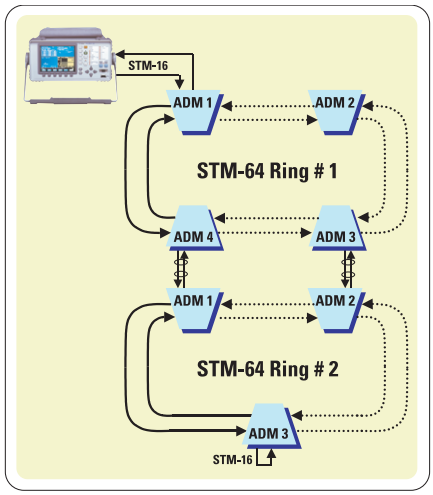

Figure 4: STM-64 System Under Test

System configuration

- MSPRING protection switching
- Wait-to-Restore period set at 5 minutes
- Under normal working conditions, the OmniBER analyzer's test pattern is transmitted via the links shown as solid lines on the diagram
- The two rings are connected via STM-16 tributaries ports

OmniBER setup

- Connected at an STM-16 tributary port
- Unframed PRBS test pattern inserted as a mapped 140 Mb/s signal within a selected VC-4 channel
- The OmniBER analyzer measurement restarted after each protection-switching event

Summary of test procedure and results

Protection Switch Trigger Event	Transmission Direction Of Test Pattern [1]	Measured Results
Normal operation	R1A1→ R1A4 → R2A1 → R2A3	-
Node failure [2]: Ring #1, ADM 4	R1A1→ R1A2 → R1A3 → R2A2 → R2A1 → R2A3	100 ms [4]
Node failure [2]: Ring#2, ADM 1	R1A1→ R1A2→ R1A3 → R2A2 → R2A3	30 ms
Node Recovery [3]: Ring #1, ADM 4	R1A1→ R1A4→ R1A3→ R2A2 → R2A3	174 & 18 ms [5] then 34 ms [6]
Node Recovery [3]: Ring #2, ADM 1	R1A1→ R1A4→ R2A1→ R2A3	14 & 11 ms [5] then 34 ms [6]
Fibre-disconnect: Between ADM 1 & 4 on Ring#1	R1A1→ R1A2→ R1A3→ R1A4 → R2A1→ R2A3	30 ms
Fibre-reconnect: Between ADM 1 & 4 on Ring#1	R1A1→ R1A4→ R2A1→ R2A3	30 ms

Notes:

1. Indicates the direction of transmission of the test pattern after completion of the protection switch [Rn = Ring number (1 or 2); An = ADM number (1 to 4 on R1; 1 to 3 on R2)]. Go and return paths for the test pattern are the same (i.e. bi-directional protection switching is used)
2. Caused by removing power from the selected ADM
3. Power restored to failed node
4. Customer accepted this as being acceptable due to need for switching events to occur on both rings in order to restore error-free transmission
5. These results were measured approximately 90 seconds after power was restored to the ADM – before the wait-to-restore period was complete. On further examination using OmniBER's graphical measurement results, it was found these unexpected results were caused by an AU-AIS alarm being output for a short time period at the tributary connected to OmniBER's receiver. The only possible explanation for this is that the ring's operation is being corrupted during the process of the failed node re-establish itself as an active node on the ring.
6. This result was measured 5 minutes after power was restored to the ADM – it is therefore the actual protecting-switching result for this test.

Initially, the wide variation in the results obtained from these tests caused questions to be raised concerning the reliability of the OmniBER's analyzer's service disruption measurement. When examined more closely, however, with an understanding of the measurement technique (as explained in this Product Note), it became clear that the analyzer was accurately measuring the error burst activity associated with each test. And in doing so, the OmniBER analyzer's service disruption measurement provided the evaluation team with a detailed knowledge of the actual protection-switching performance supported by the system under test.

Appendix 1:

Protection switching time test methods and associated error sources

Protection switching time measurement – the problem

Many of today's high-reliability transmission networks are founded on SDH/SONET technology, which provides built-in fault restoration known as Automatic Protection Switching (APS). This general term covers a range of different protection schemes designed for use in Linear and Ring network topologies, and includes linear Multiplex Section Protection (MSP), Multiplex Section Protected Rings (MSPRING) and Path Protection. However, regardless of the network topology and specific protection scheme, the basic principles behind the OmniBER analyzer's service disruption measurement, and its application in verifying a transmission system's protection switching time, remain valid.

In Figure A1, two network nodes (e.g. ADMs) are shown with a single working circuit and a single protection circuit between them. In linear systems, working and protection circuits may be paired (1+1 protection as shown), or one protection circuit may be shared among several working circuits (1:n protection).

This example shows the state of the nodes after a switch has taken place. The typical sequence of events is:

- The tail-end node detects the failure and signals the head-end to request a protection switch.
- The head-end performs a bridge or bridge and switch operation, and sends back an acknowledgement.
- The tail-end receives the acknowledgement and performs a bridge and switch operation, then finishes by sending a status message to the head-end.
- The head-end finishes by performing a switch operation if necessary.

Following a failure, full service is not restored until all the bridge and switch operations are completed. A key design goal for Network Equipment Manufacturers (NEMs) is to keep this service disruption a short as possible, as their customers (Network Operators) will demand that all system deployed in the network meet or exceed the specification published by the governing standards body (Telcordia or ITU-T). This appendix deals with the challenge of making meaningful and repeatable measurements of protection switch time.

Figure A1

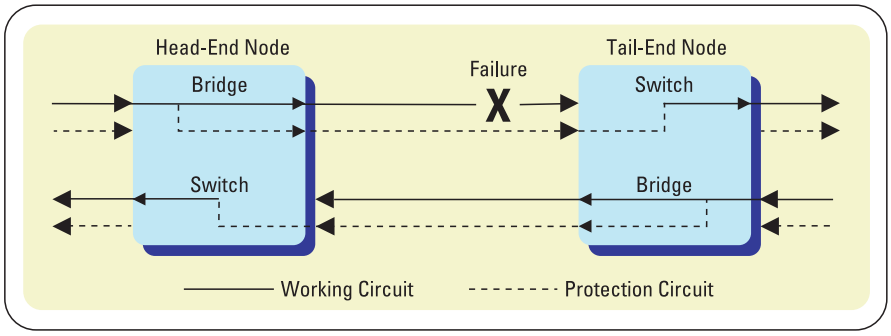

Figure A2 shows the two main components of a service disruption following any kind of failure:

The first part is the time taken to detect a failure. Protection switching can be initiated by either of two events:

1. Signal Fail (SF): usually loss of signal, loss of framing, or a very high error ratio such as 1 x 10^{-3} or greater.
2. Signal Degrade (SD): a persistent background error rate that exceeds a provisioned threshold in the range 1 x 10^{-5} to 1 x 10^{-9}. Note that, at the multiplex section level, ITU-T G.806 (October 2000 draft) specifies the 'detection time' for these error rates as 1 second for 1 x 10^{-5} to 10,000 seconds for 1 x 10^{-9}.

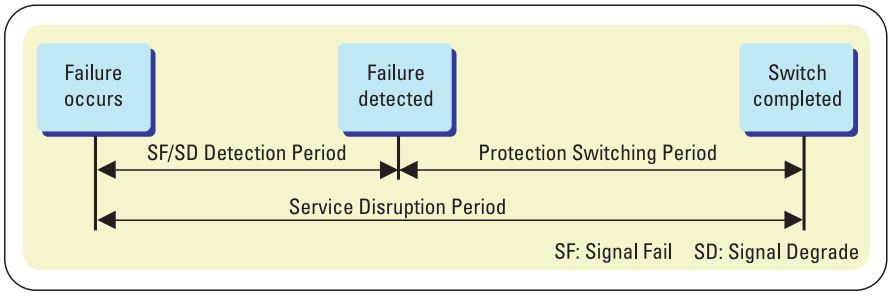

The second part is the time taken for the actual switching process to complete as described above. This is normally dominated by the protocol processing time at each node on the Protection Circuit.

Figure A2

Although users of protected services are more interested in the *total* service disruption time, the ITU-T standards provide separate specifications for protection switch time and the 'detection times' associated with the various SF and SD conditions. While this division of service disruption time into its component part may seem unhelpful, it is necessary due to the wide variation in detection time associated with the different SF/SD conditions. These range from approximately 100 µs for a LOS failure to 10,000 seconds for a signal degrade that has a provisioned threshold of 1 x 10^{-9} error rate. Another good reason for treating detection time separately is that the nature of some fault conditions can be very unpredictable. For example, when a fibre is damaged during construction work it may not break cleanly. Instead, the optical signal may fade over several tens of milliseconds or vary erratically before finally disappearing. So the ITU-T standards require that, once SF/SD is detected, a protection switch event must be completed *in 50 milliseconds or less*. This is a tough requirement, but if it is met, end-users will not normally notice a protection switch event even allowing for a realistic SF/SD detection time.

Now, having looked at the general processes associated with a Protection Switch event, it is time to address the central question: *How can the* ***potection switch time*** *of a transmission system be measured to ensure that it meets the ITU-T specification (equal to or less than 50 millisecond)?*

Ideally this would be achieved by measuring the time from "SF/SD detected" to "switch completed". However, the "SF/SD detected" event cannot be seen, and is difficult to infer, from signals outside the Network Elements (NE's) in the system under test. So the question remains, 'how can you obtain a reliable measure of a system's protection switching time'.

Protection switching time measurements - the alternative solutions

Two different approaches can be used to obtain measurement results that are (or should be) closely related to a system's protection switching time, namely, measure the service disruption time associated with a SF/SD condition that either:

1. *Minimizes* the 'detection time' (create a LOS failure – typically detected in less than 100 µs), or
2. *Eliminates* the 'detection time' (generate control parity errors* on the entity being protected)

* B2 errors for multiplex section protected systems; HP-B3 errors for High-order Path protected systems;
LP-B3/BIP-2 errors for Low-order Path protected systems

This document now provides a discussion on the merits of each of these approaches.

(1) Service disruption time with simulated LOS failure (minimize detection time)

As discussed earlier in this product note, the principles behind the service disruption measurement are simple:

- Insert a PRBS test pattern at a tributary port feeding the protected transmission system
- Perform a loopback on this test pattern at the associated drop-side tributary port
- Induce a failure on the 'working' line system that affects the test pattern
- Measure the duration of the resulting error-burst in the PRBS test pattern at the tributary port.

Although ITU-T G.783 (2000) defines LOS detection time as being "in the province of regional standards", it provides an example based on a value of 100 µs or less. This is less than the detection times specified for all other SF conditions. Consequently, if it is assumed that the system under test detects LOS within this 'example' time, then inducing an ***'instantaneous'*** LOS failure and measuring the resulting service disruption time will provide an accurate estimate of the system's protection switching time. In theory, this method will yield a result that over-estimates the protection switching time by a maximum of 100 µs, since the measured service disruption time will include both the LOS detection time and the protection switching time.

Figure A3

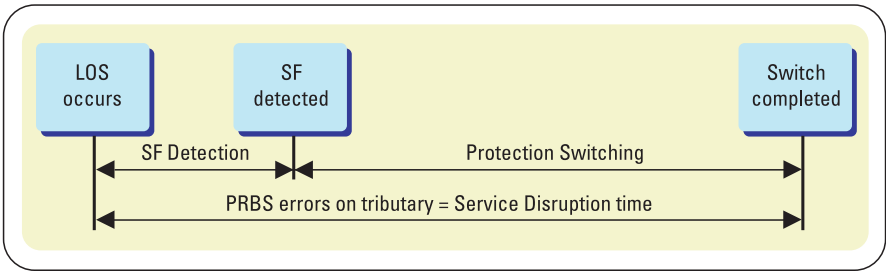

While the theory behind this test method is simple, there is a significant practical issue that must be addressed, namely, how do you induce an 'instantaneous' LOS. This is more difficult than it appears at first glance.

As illustrated in Figure A4, three different methods can be used to generate a LOS failure. The simplest of these is to manually disconnect a fibre on the working circuit (**Warning**: Exercise extreme caution when disconnecting an optical fibre – follow your organization's standard safety procedures). While this approach is attractive due to its simplicity, it does not result in the generation of an 'instantaneous' LOS condition. This is due to the finite time associated with the process of manually disconnecting the fibre.

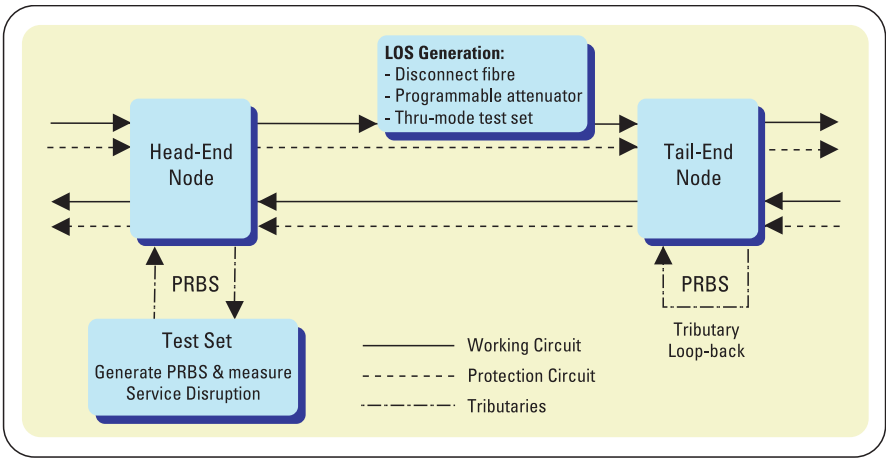

Figure A4

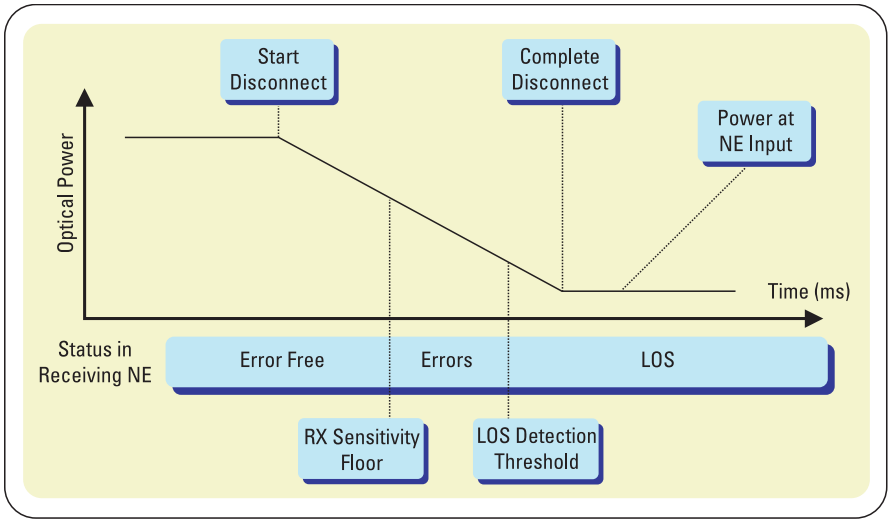

Figure A5

As shown in Figure 5, the optical power at the receiving end of the link will not disappear instantaneously when a fibre is disconnected. Instead the power level will roll-off over the time taken to perform the disconnection. A consequence of this is that the receiving NE will experience a short period of time where errors are generated before it enters the LOS condition. And since these errors will be measured as part of the service disruption time, the period of time they persist is a source of 'over-estimation' of a system's protection switching time. It should be noted here that the duration of this error period will, in many cases, be related to the time taken to disconnect the fibre – the faster the disconnection, the shorter the error period. Consequently, variation in the 'speed' of manual disconnection can lead to poor result repeatability.

Inserting a programmable optical attenuator in to the working circuit provides a more predictable method of inducing a LOS condition. In this case the LOS condition is induced by switching-in high attenuation in the transmission path. While this approach will almost certainly provide repeatable results, it may not fully address the issue of 'measurement error' due to the optical power level rolling-off over a finite period of time. This particular point relates to the fact that most programmable optical attenuators control the level of attenuation by mechanically controlling the position of a prism. And as with any mechanical process there will be a finite response time associated with this control. Note that 'response time' is a typical technical specification for most programmable optical attenuators.

The final method discussed here of inducing a LOS condition is via a test set connected in thru-mode on the working circuit. In this case the LOS condition is induced by either switching off the test set's laser transmitter or using its alarm generation controls to transmit LOS. Both of these controls will yield a predictable and instantaneous LOS condition, and consequently enable repeatable and accurate protection switching time measurements to be performed. In principle, the only source of measurement error associated with this method should be due to the LOS detection time being included in the service disruption time result.

From the above discussion, it is clear that the method used to induce a LOS condition will affect both the repeatability and accuracy of protection switching time results as measured by the service disruption time technique. However, in all cases the measurement error leads to results that are an 'over-estimate' of a system's actual protection switching time performance. Consequently, if the measured service disruption time is less than 50 milliseconds, then you can have full confidence that the system under test truly complies with the published standards for protection switching time.

(2) Service disruption time with SF trigger by excessive errors (eliminate detection time)

In this method, a test set connected in thru-mode is used to inject a high-rate of errors in the parity-check byte(s) associated with the protection system under test. In the case of a multiplex section protected system B2 parity errors are used, while HP-B3 and LP-B3/BIP-2 parity errors are used for high-order path and low-order path protected system respectively. For simplicity, the following discussion will assume that the system under test is protected at the multiplex section level.

The technique discussed here for measuring protection switching time is based on creating a Signal Fail in the system under test that is caused by an excessive error condition. In our multiplex section protected system this means injecting a B2 error rate that exceeds the receiving NE's provisioned threshold for the excessive error condition. To ensure that you always exceed the provisioned error threshold, set the thru-mode test set to inject the maximum error rate supported by the parity-check bytes (in our case – continuously error all bits of all B2 bytes). Since these errors are only injected in to the B2 parity bytes they will not affect the traffic being carried in any of the payload channels. Consequently, no errors will be added to the PRBS test pattern that is transmitted and monitored by the second test set connected at a tributary port of the system under test, and used to measure service disruption time.

Figure A6

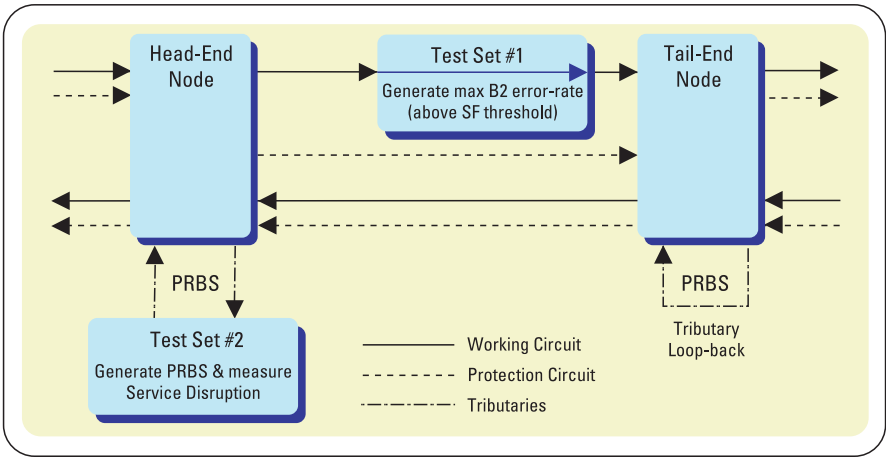

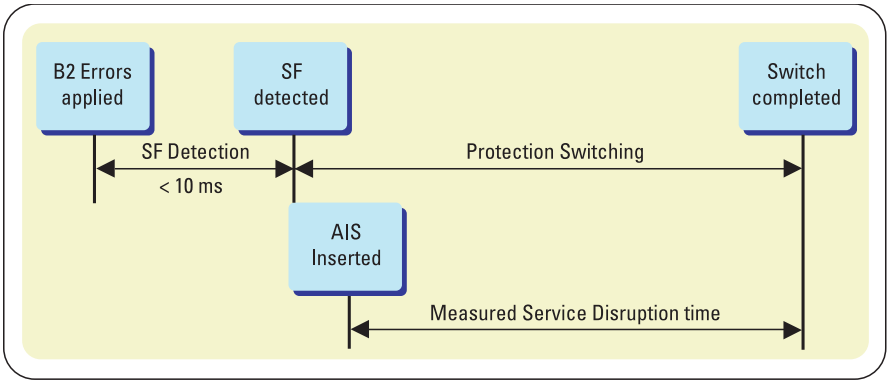

Figure A7

Within 10 ms of injecting the B2 errors, the tail-end node (the NE receiving the B2 errors) will detect the excessive error condition. This causes the NE to declare a SF and to initiate a protection switch sequence. In addition, the tail-end node is required to insert an AIS alarm in all down-stream traffic channels within 250 µs of declaring SF. And since this AIS will overwrite the PRBS test pattern that is transmitted and monitored by test set#2, it causes the service disruption measurement to be triggered (started).

For standards compliant network elements, this method will yield accurate and repeatable results when measuring protection switching time. Its main advantage over the 'LOS methods' discussed earlier is that it eliminates the 'SF detection time' error from the measured result. The only technical drawback is that its results slightly under-estimate a system's protection switching time – but only by up to 250 µs (assuming that the tail-end node inserts the downstream AIS within the 250 µs period specified in ITU-T G.783). Possibly the most serious 'drawback' associated with this measurement method is a commercial one – it requires two transmission test sets (one covering the required tributary rates, the other covering required line rates).

Summary and conclusions

While service disruption measurements can be measured as described in this note, it is worthwhile looking at the overall picture from a system perspective, as illustrated in figure A8. This represents an idealized model of what occurs in the physical layer during a switch. Following error free operation, a network element detects an event such as a fibre break that may give reason to perform a protection switch. Following the event detection, an AIS may be initiated by processes within the network element. The physical switch takes place, then a few moments later the AIS is removed. After a period of synchronization on the protection signal path, error free operation is resumed.

The service disruption measurement within the OmniBER analyzer is defined as being the time between the end of error free operation, before the switch takes place, to the beginning of error free operation after the switch has occurred.

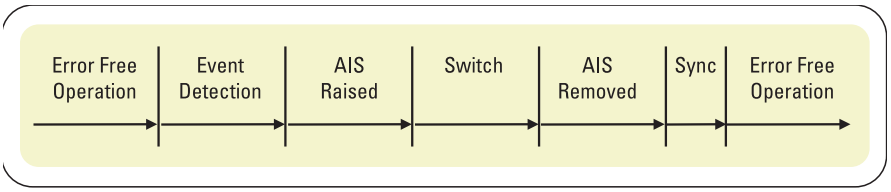

Figure A8: Protection switch processes

The OmniBER analyzer's service disruption measurements are carried out by inserting a PRBS on the tributary side of the device under test, looping it back on itself on the corresponding drop side tributary, and monitoring this PRBS for errors as a switch occurs. ITU-T recommends that a protection switching process (i.e. the "switch" section of figure A8) should be 50 milliseconds or less. While this is a difficult standard to meet, a large part of the problem is in initiating the protection switch. There are two methods to do this effectively:

- Create a LOS failure, which will typically be detected in less than 100 µs.
- Generate control parity errors on the protection system.

While each method has its own advantages, it is worth remembering that whatever device you are trying to varify, the OmniBER analyzer provides the best protection-switch analysis capabilities currently available.

Agilent Technologies' Test and Measurement Support, Services, and Assistance

Agilent Technologies aims to maximize the value you receive, while minimizing your risk and problems. We strive to ensure that you get the test and measurement capabilities you paid for and obtain the support you need. Our extensive support resources and services can help you choose the right Agilent products for your applications and apply them successfully. Every instrument and system we sell has a global warranty. Support is available for at least five years beyond the production life of the product. Two concepts underlie Agilent's overall support policy: "Our Promise" and "Your Advantage."

Our Promise

Our Promise means your Agilent test and measurement equipment will meet its advertised performance and functionality. When you are choosing new equipment, we will help you with product information, including realistic performance specifications and practical recommendations from experienced test engineers. When you use Agilent equipment, we can verify that it works properly, help with product operation, and provide basic measurement assistance for the use of specified capabilities, at no extra cost upon request. Many self-help tools are available.

Your Advantage

Your Advantage means that Agilent offers a wide range of additional expert test and measurement services, which you can purchase according to your unique technical and business needs. Solve problems efficiently and gain a competitive edge by contracting with us for calibration, extra-cost upgrades, out-of-warranty repairs, and on-site education and training, as well as design, system integration, project management, and other professional engineering services. Experienced Agilent engineers and technicians worldwide can help you maximize your productivity, optimize the return on investment of your Agilent instruments and systems, and obtain dependable measurement accuracy for the life of those products.

By internet, phone, or fax, get assistance with all your test & measurement needs

Online assistance:
www.agilent.com/find/assist

Phone or Fax
United States:
(tel) 1 800 452 4844

Canada:
(tel) 1 877 894 4414
(fax) (905) 282 6495

China:
(tel) 800 810 0189
(fax) 1 0800 650 0121

Europe:
(tel) (31 20) 547 2323
(fax) (31 20) 547 2390

Japan:
(tel) (81) 426 56 7832
(fax) (81) 426 56 7840

Korea:
(tel) (82 2) 2004 5004
(fax) (82 2) 2004 5115

Latin America:
(tel) (305) 269 7500
(fax) (305) 269 7599

Taiwan:
(tel) 080 004 7866
(fax) (886 2) 2545 6723

Other Asia Pacific Countries:
(tel) (65) 375 8100
(fax) (65) 836 0252
Email: tm_asia@agilent.com

Printed in USA, 8 March, 2002
5988-5122EN

CE

Agilent Technologies