The Effects of Multiple Optical-Electrical-Optical Signal Conversions in Long-Distance Optical Networks on the Round-Trip Time and TCP/UDP Performance

Frederik Norberg

April 3, 2011

Preface

This capstone project concludes my studies towards a Master of Science in Internetworking degree at the University of Alberta. The work presented in this report was carried out in the period between September 2010 and March 2011.

I would first and foremost like to thank Nader Rezazadeh for providing me with the project topic, as well as for his valuable input over the course of this project. This is very much appreciated.

I would also like to thank Dr. M.H. MacGregor for providing me access to the equipment necessary to carry out this project.

Finally, I would like to thank Shahnawaz Mir for all his help in the lab.

Contents

List of Figures

List of Tables

Acronyms

ACM	Association for Computing Machinery
AMI	Alternate Mark Inversion
B8ZS	Binary 8 Zeroes Substitution
BLSR	Bidirectional Line Switched Ring
dB	Decibel
DS	Digital Signal Level
ESF	Extended Super Frame
ICMP	Internet Control Message Protocol
IEEE	Institute of Electrical and Electronics Engineers
IR	Intermediate Reach
ISO	International Organization for Standardization
LR	Long Reach
MSPP	Multiservice Provisioning Platform
nm	Nanometer
O-E-O	Optical-Electrical-Optical
OSI	Open Systems Interconnection
PCM	Pulse-Code Modulation
QRSS	Quasi-Random Signal Sequence
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Networking
SPE	Synchronous Payload Envelope
SR	Short Reach
STS	Synchronous Transport Signal
TCP	Transmission Control Protocol
TDM	Time-Division Multiplexing
UDP	User Datagram Protocol
UPSR	Unidirectional Path Switched Ring
UR	Ultra-long Reach
VR	Very-long Reach
VT	Virtual Tributary

Chapter 1

Introduction

Optical fibres have for quite some time been preferred over copper cables for the transport of voice and data traffic in carrier networks, especially over longer distances. One major advantage of optical fibres is that, because of their huge amount of available bandwidth, they can support significantly higher data rates than copper cables. Optical fibres are also less susceptible to undesirable effects such as electromagnetic interference, which make them better suited for use in long-distance networks.

As optical networks are rarely all-optical, a data signal is often converted multiple times between optical and electrical form between two endpoints. In most cases, this is done in order to either process or switch the signal, or in order to regenerate it as it traverses a long-distance link. In this project, the effects of these physical-layer conversions on the performance of a data signal, as measured mainly at the layers further up in the protocol hierarchy, will be investigated.

1.1 Project Objectives

This project has two main objectives. The first one is to investigate how an increasing number of successive *optical-electrical-optical* (o-e-o) data signal conversions in long-distance optical networks affects the round-trip time and TCP/UDP performance. In this performance evaluation, the following metrics will be used:

- **Round-Trip Time** - The two-way delay between two endpoints.
- **Datagram Loss** - Loss of data at the transport layer.
- **Throughput** - The rate at which data can be transferred across the network.
- **Jitter** - The variability of the delay in the network.
- **Bit-Error Rate** - Number of individual bits received in error.

The second objective is to determine the maximum number of successive o-e-o conversions possible before the quality of the data signal becomes unacceptable due to excessive delay, loss or jitter.

1.2 Project Scope

The data signal that will be used in this project is the *Digital Signal Level 1* (DS1), which will be carried across an optical *Synchronous Optical Networking* (SONET) network. All signal processing, including switching, in SONET networks is carried out in the electrical domain. This means that an optical signal will always get converted to electrical form when it is received by a network element. After processing, the signal is either converted back to optical form for transport across an optical link, or it is dropped to some customer equipment. Further information about SONET networks and DS1 signals is given in Appendix C.

The main focus of this project will be on performance testing in a SONET network in the *Unidirectional Path Switched Ring* (UPSR) configuration. However, testing will also be carried out in a SONET network in a terminal (point-to-point) configuration for comparison. These two configurations will for the remainder of this report be referred to as the *main configurations*.

The number of o-e-o conversions can be controlled by looping a DS1 signal back and forth the desired number of times between two endpoints. This looping is achieved by first dropping, and then adding the same DS1 signal at an endpoint. As this procedure also increases the total distance that the DS1 signal must travel, the discussion of the limiting number of o-e-o conversions will be expanded to include this. For the remainder of this report, the focus will therefore be on the number of times the DS1 signal is sent through the network instead of just o-e-o conversions.

1.3 Preliminary Literature Survey

Before starting this project, a literature survey of several scientific and technical databases was carried out in order to ensure that results from similar performance evaluations were not already available. The main databases included in this literature survey were the following:

- **IEEE Xplore** - a database maintained by the Institute of Electrical and Electronics Engineers (IEEE) that provides access to IEEE transactions, journals, magazines and conference proceedings[6].
- **ACM Digital Library** - a database maintained by the Association for Computing Machinery (ACM) that provides access to ACM transactions, journals, magazines and conference proceedings[1].
- **Computer and Information Systems Abstracts** - a database that contains information about both theoretical research and practical applications within the area of computer and information systems from a variety of different sources[3].

The literature survey resulted in the finding of performance evaluations of several different network types. However, these did not target the same systems and configurations considered in this project. This does not guarantee that something similar has not been done previously, but if it has, the results cannot be found amongst the available sources. The work on this project was therefore continued.

Chapter 2

Methodology

As mentioned in Chapter 1, one of the two main objectives of this project is to investigate how an increasing number of successive *optical-electrical-optical* (o-e-o) conversions of a DS1 signal in a long-distance SONET network affects the round-trip time and TCP/UDP performance. The second objective is to determine the maximum number of times the DS1 signal can be sent through the network before its quality becomes unacceptable due to excessive delay, jitter and loss.

The upper limits for delay, jitter and loss used in this project are given in the appropriate sections below, and are based on those recommended by Cisco for networks carrying videoconferencing traffic in [10]. These requirements are even stricter than those commonly used for voice over IP networks.

An introduction to the various hardware components used in this project, including optical network elements, routers, servers and fibre spools, is given in Appendix A.

This chapter is structured in the following way. First, Section 2.1 describes how the number of o-e-o conversions of a DS1 signal, as well as the total distance this signal must travel, can be controlled by sending it back and forth through the same physical network segment one or more times. This is followed by Section 2.2, which introduces the different system configurations used. Finally, Section 2.3 covers the metrics used for performance evaluation in this project. This section also includes the reasoning behind the choice of these specific metrics, as well as information about the actual testing procedures.

2.1 DS1 Signal Looping

As all signal processing in SONET networks takes place in the electrical domain, an optical signal gets converted to electrical form every time it enters a network element. Hence, the total number of times a data signal gets o-e-o converted can be varied by changing the number of network elements it passes through between two endpoints in the network.

Four SONET network elements are available for use in this project. In order to increase the number of o-e-o conversions beyond this, the DS1 signal can be looped back and forth across the same physical network in different timeslots of an OC-12/OC-48 signal. This process is described below and illustrated in Figure 2.1. In this example, a DS1 signal is sent across the same OC-12 link three times before being dropped. Note that for the sake of simplicity, this example only shows the signal going in one direction between the two routers.

1. Router *R1* is connected to port 1 on the DSX-1 panel attached to network element *NE1*. NE1 takes the DS1 signal from R1, places it in a VT1.5 container, and multiplexes it into a timeslot on the OC-12 circuit between NE1 and network element *NE2*.
2. At NE2, the incoming OC-12 signal is demultiplexed, and the DS1 signal is picked out and dropped to port 1 on attached the DSX-1 panel.
3. The DS1 signal is looped back from port 1 to port 2 on the DSX-1 panel using a T1 crossover cable.
4. NE2 then places the DS1 signal in a VT1.5 container and multiplexes it into an available timeslot on the OC-12 circuit between NE2 and NE1.
5. At NE1, the DS1 signal is picked out from the demultiplexed OC-12 signal and dropped to port 2 of the attached DSX-1 panel.
6. The DS1 signal is looped back from port 2 to port 3 on the DSX-1 panel using a T1 crossover cable.
7. NE1 then places the DS1 signal in a VT1.5 container and multiplexes it into an available timeslot on the OC-12 circuit between NE1 and NE2.
8. At NE2, the DS1 signal is once again picked out from the demultiplexed OC-12 signal and dropped to port 3 of the connected DSX-1 panel, which is connected to router *R2*.

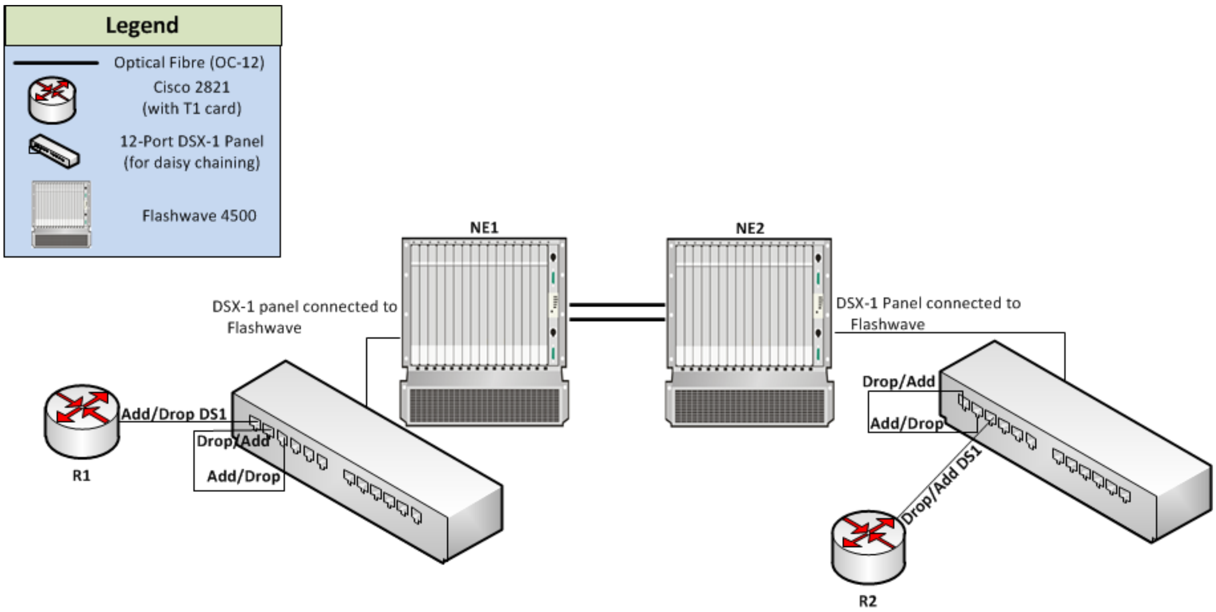

Figure 2.1: Looping of a DS1 signal in order to increase the number of o-e-o conversions as well as the total network length (DSX-1 panel icons are enlarged to show the looping).

Using the equipment available in the lab, a DS1 signal can be sent across the SONET network anywhere from one to twelve times. For each additional time the signal is sent across the network, the distance it has to travel also increases linearly. By using fibre spools, a long-distance network can be set up in the lab.

2.2 Network Configurations

As previously mentioned, performance tests will be carried out for both UPSR ring and point-to-point (terminal) network configurations. For each of these two main configurations, twelve batches of tests will be executed. The first of these will be run when the DS1 signal is sent across the network once, the second when the DS1 signal is sent across the network twice and so on.

In addition to the two main configurations, a batch of tests will also be carried out for a configuration where the SONET part of the network is left out. Here, the two routers that normally add and drop DS1 signals to and from the optical network elements will instead be directly connected using a crossover T1 cable. This configuration is illustrated below in Figure 2.2. These tests will provide baseline results that can be compared to the results from the tests of the two main configurations.

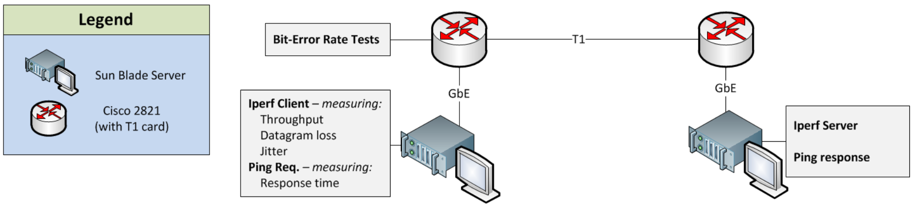

Figure 2.2: Direct electrical T1 connection between two routers.

2.2.1 UPSR Configuration

For the UPSR ring testing, the network configuration shown below in Figure 2.3 will be used. The DS1 signal traveling from the router on the left hand side of the figure to the right hand side, takes the following path:

1. The router is connected to port 1 on the DSX-1 panel attached to network element *NE1*. NE1 takes the incoming DS1 signal, places it in a VT1.5 container, and multiplexes it into a timeslot on the OC-12 circuit between NE1 and network element *NE2*.
2. After traversing the two fibre spools between NE1 and NE2, the OC-12 signal is received by NE2. Here, it is cross-connected onto the OC-48 UPSR ring.
3. After having passed through the ring, the incoming OC-48 signal is demultiplexed in network element *NE5*, and the OC-12 signal containing the DS1 signal is dropped onto the OC-12 circuit between NE5 and network element *NE6*.
4. After traversing the fibre spool between NE5 and NE6, the signal is received by NE6. Here, the incoming OC-12 signal is demultiplexed and the DS1 signal is picked out and dropped, through the attached DSX-1 panel, to the router.

For the DS1 signal going from the right hand side to the left hand side in the figure, corresponding steps are carried out.

However, setting up the optical part of the network exactly as shown in Figure 2.3 requires six Flashwave 4500 network elements. A network that is logically equivalent to the one shown in this figure can be set up using the four available network elements if they are physically interconnected as shown above in Figure 2.4. This figure also shows how the signals must be cross-connected internally in the network elements to make this configuration work.

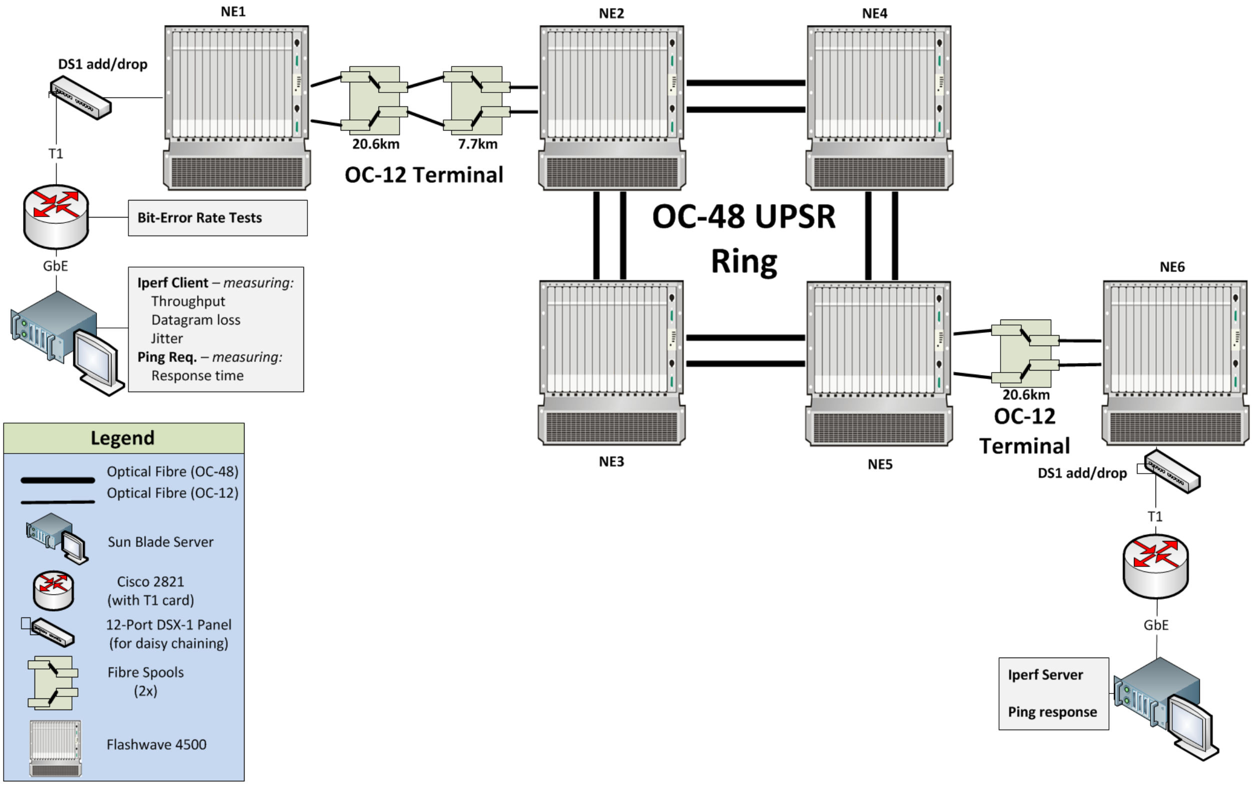

Figure 2.3: UPSR network configuration.

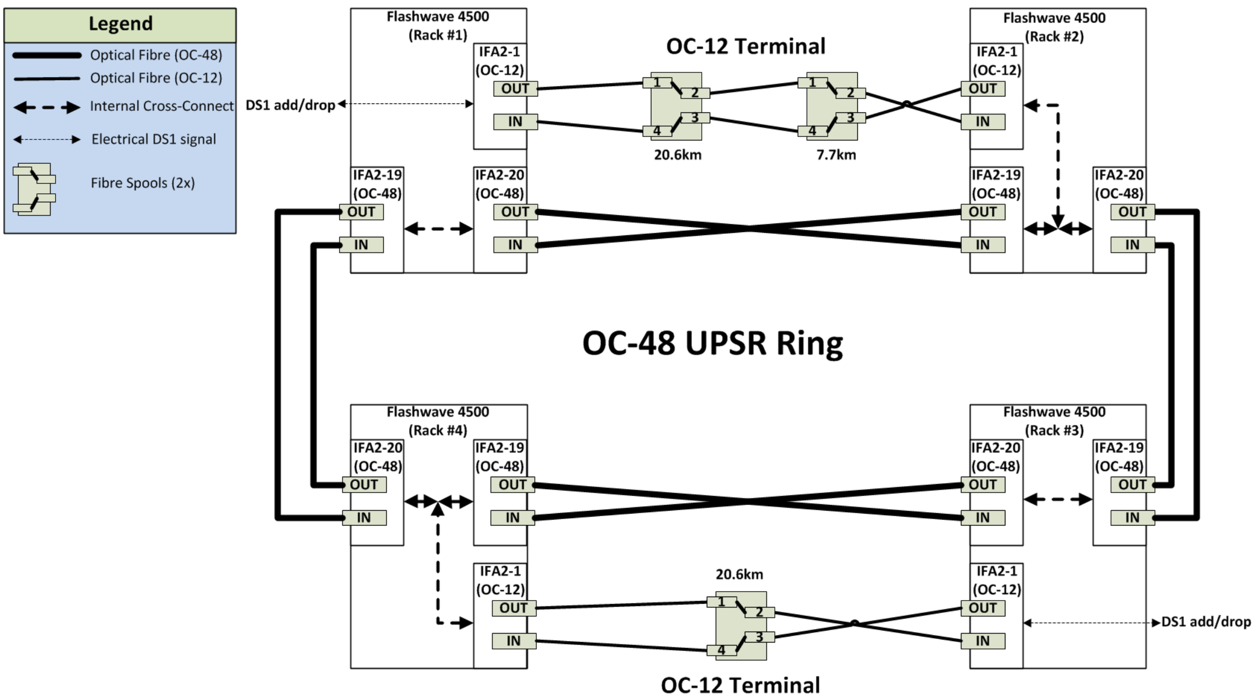

Figure 2.4: Physical layout of the optical part of the network in the UPSR configuration.

In Figure 2.4, the network element in Rack #1 corresponds to NE1 and NE3 in Figure 2.3, Rack #2 corresponds to NE2, Rack #3 corresponds to NE4 and NE6, and Rack #4 corresponds to NE5.

The looping is handled by the two endpoints in the network, NE1 and NE6, as described in Section 2.1.

2.2.2 Terminal Configuration

For performance testing of the network in the terminal configuration, the equipment will be set up as shown below in Figure 2.5. The DS1 signal traveling from the router on the left hand side of the figure to the right hand side, takes the following path:

1. The router is connected to port 1 on the DSX-1 panel attached to network element *NE1*. NE1 takes the incoming DS1 signal, places it in a VT1.5 container, and multiplexes it into a timeslot on the OC-12 circuit between NE1 and network element *NE2*.
2. After traversing the three fibre spools between NE1 and NE2, the OC-12 signal is received by NE2. Here, the incoming OC-12 signal is demultiplexed and the DS1 signal is picked out and dropped, through the attached DSX-1 panel, to the router.

For the DS1 signal going from the right hand side to the left hand side in the figure, corresponding steps are carried out.

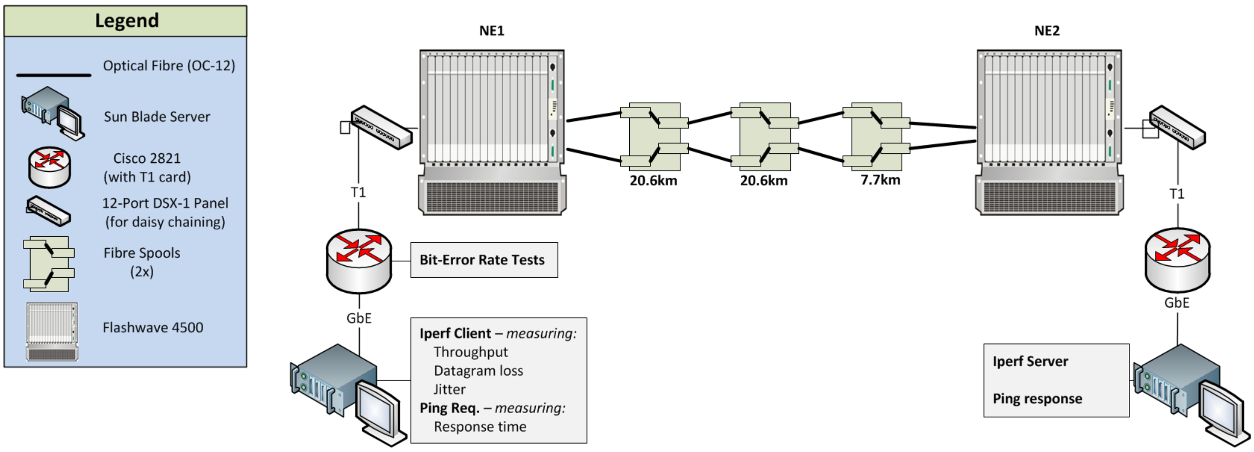

Figure 2.5: Terminal network configuration.

2.3 Metrics and Testing Methodology

With the exception of the bit-error rate, all of the metrics used in this project are at the *internet* or *transport* layers in the *Internet Protocol Suite*. The bit-error rate tests are included because knowing the bit-error rates in the different network configurations can be useful in helping to explain the results obtained in the other tests.

Round-Trip Time

In this project, round-trip time will be measured using the *ping* utility, which comes bundled with most operating systems. Ping uses *Internet Control Message Protocol* (ICMP) packets to measure the round-trip time between two hosts. The round-trip time is measured as the time from when the source transmits an *echo request* packet until it receives an *echo reply* packet from the destination host.

The ping utility can be configured to ping a remote host a certain number of times at given intervals, and when this has been done, it automatically calculates the average round-trip time.

For each batch of tests, ping will run for three times five minutes, sending one ping per second. The average of the results from these three runs will be used. Furthermore, an upper limit of 300 ms for the round-trip delay, as suggested by Cisco in [10], will be used in this project.

Bit-Error Rate

Bit-error rate tests measure the number of individual bits that are received in error by a node in the network. A bit-error rate test works by transmitting a predefined bit pattern across the network, and then comparing the received bits to this pattern. Several such bit patterns are available for testing of T1 circuits. In this project, the *Quasi-Random Signal Source* (QRSS) will be used as it simulates real traffic.

In this project, the bit-error rate tester built into the Cisco 2821 routers will be used. While it would have been ideal to use a stand-alone tester for the best possible results, the built-in tester should be good enough since these tests are not the main focus in this project.

For each batch of test for the UPSR configuration, an 8-hour bit-error rate test will be executed. For tests run when the network is in the terminal configuration, 8-hour bit-error tests will only be executed when the DS1 signal is sent across the network 1 and 12 times. For the other cases, 15-minute tests will be used, due to the limited lab-time available.

Throughput

The *Transport Control Protocol* (TCP) is a connection-oriented transport-layer protocol widely used for in-order, reliable transfer of data across the Internet. It is therefore natural to run the throughput tests in this project over a TCP connection. These tests will be performed using version 2.0.5 of the commonly used *Iperf* [7] network testing tool. Iperf is open source and has a client and a server component.

For throughput testing, Iperf allows the user to specify the TCP *window size*. This setting determines how much unacknowledged data TCP can have *in flight* before it stops transmitting and waits for acknowledgments. The ideal value for this parameter is dependent on the capacity of a link and its two-way delay. Setting the TCP window size too low limits the throughput of the link, whereas setting it too high can lead to long delays when retransmissions due to loss are necessary. Tests were carried out with several different window sizes, and it was found that the default values of 16 Kbytes for the client and 85.3 Kbytes for the server were suitable for the network configurations that will be used in this project.

Trials that were run in the lab show that it typically takes Iperf a couple of seconds to reach maximum throughput due to TCP's slow start algorithm. It is therefore important that the duration of each throughput test is significantly longer than this, so that the average throughput for the test becomes as accurate as possible. It was found that running each test for five minutes is adequate.

For each batch of tests, three five-minute throughput tests will be run. The average of the results from these three runs will be used.

Jitter

For applications such as voice over IP and videoconferencing, it is important that packets arrive at regular intervals in order to provide a good quality of service. Ideally, the delay for all packets sent over the same path is constant, however, this can be difficult to achieve in real networks. Jitter is therefore used as a measure of the variability of packet delays. A low value for the jitter is good as it indicates that the delay between packets is fairly consistent.

In this project, we will be measuring the jitter for a stream of *User Datagram Protocol* (UDP) datagrams. Like TCP, UDP is a commonly used protocol at the transport layer in the Internet Protocol Suite. However, unlike TCP, UDP is a connectionless best-effort protocol. UDP is often used in delay-sensitive applications where the retransmissions caused by TCP in the event of packet loss are not desirable.

To measure jitter, the Iperf testing tool will be used. Iperf allows the user to specify the size of the datagrams, as well as the sending rate. It is necessary to specify the sending rate because unlike TCP, UDP does not have flow control. Iperf sends datagrams of the specified size, at the specified rate for the specified amount of time, and then automatically calculates the jitter.

For the datagram size, the default value of 1470 bytes was found to be appropriate. For the sending rate, two different values will be used. These are:

- **1.48 Mbps** - This sending rate was chosen because it lies very close to the maximum capacity of the T1 line used in this project. Tests run with this setting will therefore provide jitter measurements for the channel when it is heavily loaded. Lab trials showed that setting the UDP sending rate much higher than 1.48 Mbps would cause a significant number of packets to be dropped as the channel becomes overloaded.
- **1.00 Mbps** - This value was chosen well below the maximum rate of the channel to provide jitter measurements for a lightly loaded channel.

For each batch of tests, three 10-minute UDP tests will be executed for each of the two sending rates. For each of the two sending rates, the average of the results from the three runs will be used. Furthermore, an upper limit of 10 ms for the one-way jitter, as suggested by Cisco in [10], will be used in this project.

In addition to the tests that will be run over the UPSR and terminal configurations, 30 10-minute tests, for both sending rates, will also be run over the configuration shown in Figure 2.2. The main purpose of these tests is to help determine whether the jitter measured in the UPSR and terminal configurations originate inside or outside the optical part of the network.

Datagram Loss

UDP sends data in units called datagrams. These are sometimes lost during transport for a variety of reasons such as errors in the transmission of bits across a link, buffers overflowing, links failing and so on. It is therefore important to measure these losses as part of this performance evaluation.

In this project, datagram losses will also be measured using the Iperf network testing tool. Iperf records datagram losses as part of the same set of tests that are used for measuring jitter. The settings and procedures for the datagram loss tests will therefore be the same as for those that measure jitter. Furthermore, an upper limit of 0.05% for the datagram loss rate, as suggested by Cisco in [10], will be used in this project.

Chapter 3

Test Results and Discussion

This chapter presents the performance tests results of the different network configurations, as well as a discussion of these.

Two important differences between the UPSR and terminal network configurations that were used in this project are the number of network elements, as well as the lengths of the individual fibre spans, in each of these. For each one-way trip the data signal takes through the network in the UPSR configuration, it passes through five Flashwave 4500 network elements. In the terminal configuration, on the other hand, it only passes through the two of these network elements. Hence, the optical signal gets o-e-o converted fewer times when it passes through the terminal network than it does passing through the UPSR network. In terms of fibre span lengths, the single 48.9 km OC-12 link in the terminal network is significantly longer than the 28.3 and 20.6 km OC-12 links in the UPSR network. However, the total length of the network in each of these two configurations is almost identical, around 49 km. The only difference lies in the cables that are used for creating the ring in the UPSR configuration, which have a total length of around 30 meters.

In this chapter, the results of the round-trip time measurements are presented first in Section 3.1. These are followed by the results of the bit-error rate tests in Section 3.2. Then, the results of the datagram loss tests are covered in Section 3.3, before the throughput test results are covered in Section 3.4. Finally, the results of the jitter tests are given in Section 3.5.

3.1 Round-Trip Time

The round-trip times for the different number of passes the data signal made through the UPSR and terminal networks are illustrated in Figure 3.1.

As one might expect, there was an almost perfectly linear increase in the round-trip time for each additional time the data signal was sent through the network in both the UPSR and terminal configurations. In the UPSR configuration, one additional pass through the network, on average, increased the round-trip time by 0.84 ms. In the terminal configuration, one additional pass through the network, on average, increased the round-trip time by 0.79 ms.

As can be seen from Figure 3.1, the round-trip times in the UPSR network were slightly higher than those in the terminal network. This can be explained by the number of network elements the data signal had to pass through between its source and its destination. Going one way in the UPSR network, the signal passed through five network elements, whereas it only passed through two in the terminal configuration. Each network element adds a very small amount of delay due to signal processing and switching.

The difference in the distance traveled when the data signal passed through the network in the UPSR configuration versus the terminal configuration was very small relative to the total length of these networks. All three fibre spools were used in both configurations, so the only difference in terms of the meters of fibre traversed comes from the fibres in the OC-48 core ring in the UPSR configuration. However, these fibres were only a few meters long and did therefore not contribute much to the overall round-trip time.

In terms of the previously set round-trip time limit of 300 ms, it can be seen that all the tested configurations were well below this. Assuming a continued linear increase in the round-trip times, the data signal could potentially be sent through the network, in both the UPSR and terminal configurations, over 350 times without reaching this limit.

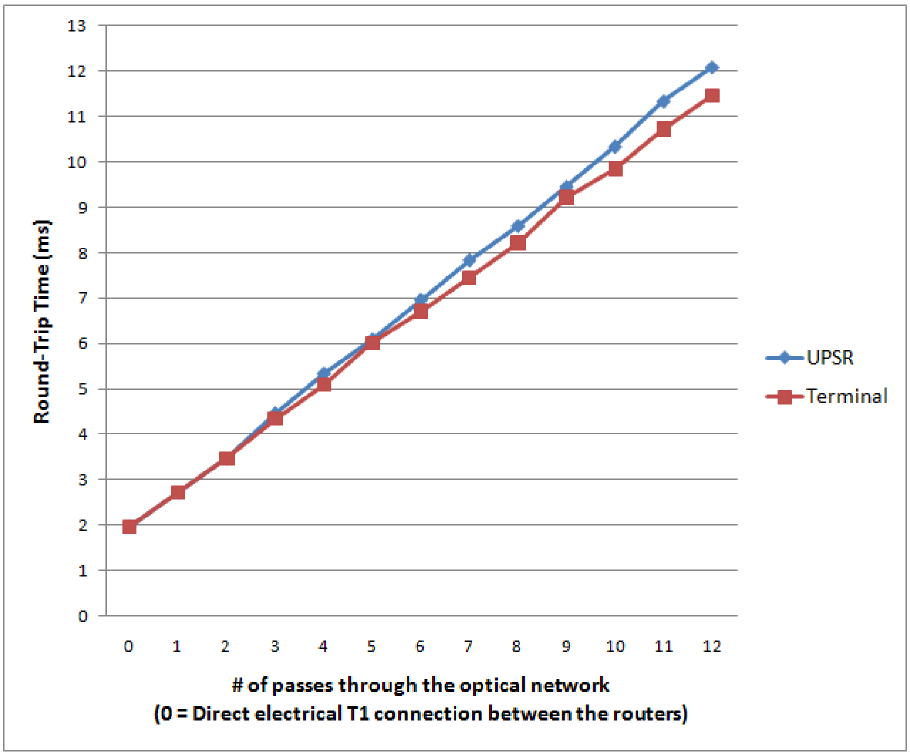

Figure 3.1: Round-trip times in the UPSR and terminal network configurations.

3.2 Bit-Error Rate

As previously mentioned, bit-error rate tests (BERTs) were included because they might help to explain the results obtained in the other tests. Initially, a 24-hour bit-error rate test was run between the two Cisco 2821 routers in the network when they were directly connected, using a crossover T1 cable, as shown in Figure 2.2. This was done in order to check whether any bit errors were introduced by the equipment outside of the optical part of the network. In this test, not a single bit was received in error. The short cable, which was around 10 meters long, used to interconnect the two routers, was most likely one of the main reasons for this. The short cable allowed the two routers to receive strong, high-quality signals from each other.

3.2.1 UPSR Configuration

The results of the bit-error rate tests of the UPSR network configuration are given in Table 3.1. These show that when the data signal was sent through the network once, there were four bit errors (an error rate of $9.26x10^{-11}$), and when it was sent through the network 11 times, there was only a single bit error (an error rate of $2.31x10^{-11}$). For all of the other tests run in the UPSR network, there were no bit errors.

One likely explanation for the very low bit-error rates is that all of the individual links in the network were relatively short compared to what the different equipment is capable of handling. The fibre cables in the OC-48 core ring were only a few meters long, whereas the line cards that were used can handle fibre spans of up to 3 km. The two OC-12 spans were 28.3 km and 20.6 km long, which is well below the 60 km loss limit of the line cards that were used. Additionally, the length of each of the electrical T1 cables was only a few meters.

For the range of values tested in this project, the results given in Table 3.1 do not indicate that increasing the number of o-e-o conversions in the network, as well as its total length, effects the bit-error rate as long as the quality of the optical signal received by each network element is good. This is logical as the optical signal gets regenerated every time it gets converted between optical and electrical form in a network element. If the quality of the received signal is good, it can be correctly interpreted and retransmitted.

# of Passes	Test Runtime	Errored Bits	Total Bits	Bit-Error Rate
1	8 hours	4	43,202,829,000	$9.2587x10^{-11}$
2	8 hours	0	43,202,828,000	0
3	8 hours	0	43,202,825,000	0
4	8 hours	0	43,202,826,000	0
5	8 hours	0	43,202,825,000	0
6	8 hours	0	43,202,825,000	0
7	8 hours	0	43,202,826,000	0
8	8 hours	0	43,202,826,000	0
9	8 hours	0	43,202,831,000	0
10	8 hours	0	43,202,831,000	0
11	8 hours	1	43,202,830,000	$2.3147x10^{-11}$
12	8 hours	0	43,202,830,000	0

Table 3.1: Bit-error rate test results for the UPSR configuration.

3.2.2 Terminal Configuration

As previously mentioned, most of the bit-error rate tests of the network in the terminal configuration were only 15 minutes long because of the limited lab time available. A summary of the results from these tests is given in Table 3.2.

Most of the tests in the terminal configuration resulted in zero or insignificant bit-error rates. However, when the data signal was sent across the network five and seven times, the bit-error rates were $4.45x10^{-9}$ and $1.04x10^{-8}$, respectively. One potential explanation for this is that more errors were introduced because of the longer OC-12 link, compared to those used in the UPSR configuration, but this is most likely not the case. A more likely explanation is that these two results are just random anomalies, which is supported by the fact that the two 8-hour test only resulted in just a single bit error, and that all of the other 15-minute tests produced zero bit errors. Furthermore, re-running the two bit-error rate tests in question also resulted in zero bit errors.

As with the results of the tests that were run in the UPSR network configuration, there is no clear trend showing that increasing the number of o-e-o conversions in the network as well as its total length increases the bit-error rate for the range of values tested in this project. However, in order to verify this, all 15-minute tests should be replaced with longer-running tests that would be more accurate.

# of Passes	Test Runtime	Errored Bits	Total Bits	Bit-Error Rate
1	8 hours	1	43,202,831,000	$2.3147x10^{-11}$
2	15 minutes	0	1,349,216,000	0
3	15 minutes	0	1,349,216,000	0
4	15 minutes	0	1,349,216,000	0
5	15 minutes	6	1,349,066,000	$4.4475x10^{-9}$
6	15 minutes	0	1,349,216,000	0
7	15 minutes	14	1,349,216,000	$1.0376x10^{-8}$
8	15 minutes	0	1,349,216,000	0
9	15 minutes	0	1,349,216,000	0
10	15 minutes	0	1,349,216,000	0
11	15 minutes	0	1,349,216,000	0
12	8 hours	0	43,202,830,000	0

Table 3.2: Bit-error rate test results for the terminal configuration.

3.3 Datagram Loss

The results of the datagram loss tests for the different network configurations are given in Table 3.3. The first row of results is from a series of 30 tests run over a configuration where the two Cisco 2821 routers were directly connected, as shown in Figure 2.2. In these tests, not a single datagram was lost for either of the two sending rates.

# of Passes	UPSR @1.48 Mbps	UPSR @1.00 Mbps	Terminal @1.48 Mbps	Terminal @1.00 Mbps
0	0	0	0	0
1	0	0	0	0
2	$4.41x10^{-6}$	0	0	0
3	0	0	0	0
4	0	0	0	0
5	0	0	0	0
6	$4.41x10^{-6}$	0	0	0
7	$4.41x10^{-6}$	0	0	0
8	0	0	0	0
9	0	0	0	$6.53x10^{-6}$
10	$4.41x10^{-6}$	0	0	0
11	$4.41x10^{-6}$	0	$4.41x10^{-6}$	0
12	$4.41x10^{-6}$	0	0	$1.31x10^{-5}$

Table 3.3: Datagram loss rates for all configurations. The first row gives the results of the tests run when the two Cisco 2821 routers were directly connected.

Looking at the results of the tests run in both the UPSR and terminal network configurations, and for both sending rates, it can be seen that a few datagrams were lost when either of the two optical networks were inserted between the two routers. However, these datagram loss rates are extremely low, much lower than the threshold of 0.05% introduced in Chapter 2. A loss rate of $4.41x10^{-6}$ means that one out of 230,000 datagrams were lost.

One important factor that contributed to these very low datagram loss rates is the low bit-error rates in all of the different network configurations, as discussed in Section 3.2. Furthermore, when the testing was carried out, there was no other traffic in the network that might otherwise have caused IP packets to be dropped because of congestion.

Although some datagrams were lost in both the UPSR and terminal network configurations, the loss rates are too low and randomly distributed to provide concluding evidence as to whether increasing the number of o-e-o conversions in the network, as well as its total length, effects the datagram loss rate. Because of the very low datagram loss rates, each of the tests would most likely have to run for hours, if not days, in order to reach a firm conclusion.

3.4 Throughput

As Table 3.4 shows, the TCP throughput stayed consistently around 1475 kbps for all of the tests that were run in both the UPSR and terminal network configurations. The following two factors were important in maintaining this throughput when the number of o-e-o conversions in the network, as well as its total length, increased:

1. As discussed in Section 3.2, all measured bit-error rates in the different network configurations were very low. Additionally, the test results discussed in Section 3.3 show that very few UDP datagrams were lost in these networks. These two factors indicate that few TCP segments, the data unit used by TCP, were lost as well. This is important because TCP retransmissions in networks with higher packet losses can reduce the throughput substantially.
2. The TCP window sizes, i.e. the send and receive buffers, were large enough that TCP did not have to stop transmitting to wait for acknowledgments for *in flight* segments. This can be a problem in high-capacity links with long two-way delays.

# of passes	**UPSR**	**Terminal**
0	1475 kbps	1475 kbps
1	1475 kbps	1475 kbps
2	1475 kbps	1475 kbps
3	1475 kbps	1475 kbps
4	1475 kbps	1475 kbps
5	1475 kbps	1475 kbps
6	1475 kbps	1475 kbps
7	1475 kbps	1475 kbps
8	1475 kbps	1475 kbps
9	1475 kbps	1475 kbps
10	1475 kbps	1475 kbps
11	1475 kbps	1475 kbps
12	1475 kbps	1475 kbps

Table 3.4: Throughput for all configurations. The first row gives the results of the tests run when the two Cisco 2821 routers were directly connected.

One point that should be noted is that the achieved TCP throughput of 1475 kbps is somewhat lower than the raw user data capacity of the DS1 signal, which is 1536 kbps. The reason for this is that some of the capacity of this signal is taken up by various protocol overheads.

3.5 Jitter

In this section, the results of the performance tests of a lightly loaded T1 circuit, using a sending rate of 1 Mbps, are presented first. This is followed by the test results for a more heavily loaded T1 circuit, using a sending rate of 1.48 Mbps.

During the initial trial runs of the jitter tests, there was significant variation between the results obtained over the same network configuration when the sending rate was set to 1.48 Mbps. It was therefore decided to run 30 tests for each of the two sending rates over the configuration shown in Figure 2.2, where the two Cisco 2821 routers are directly connected. These are meant as reference measurements that can be compared to the results from the remaining tests where the two different optical networks are inserted between the routers.

3.5.1 Sending Rate: 1 Mbps

The results of the reference tests that were run between the two directly connected routers are given in Figure 3.2. In this configuration, there was very little jitter, and it only varied slightly in the range between 0.004 ms (0.23% of the average ping time) and 0.010 ms (0.58% of the average ping time). In these tests, the average jitter was 0.0065 ms and the sample standard deviation was 0.0014 ms (21.2% of the average jitter).

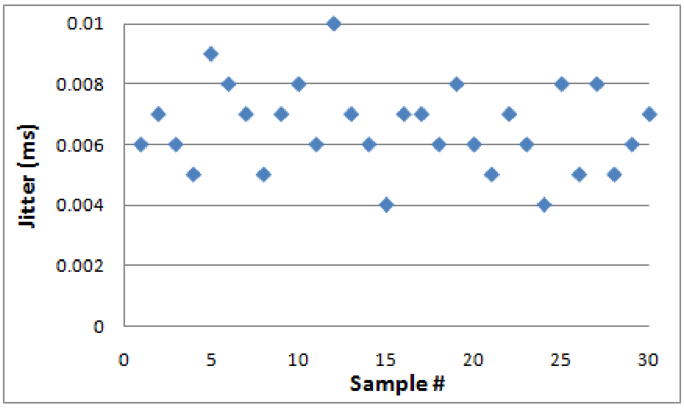

Figure 3.2: Jitter in the configuration shown in Figure 2.2 for a sending rate of 1 Mbps.

The results of the performance tests in the UPSR and terminal network configurations are illustrated in Figure 3.3. With the exception of a single measurement, when the data signal was sent through the terminal network once, the results for both configurations were consistently within the 0.005 ms to 0.01 ms range. This is the same range as the results from the tests that were run between the two directly connected routers. This indicates that what little jitter there was, was introduced outside of the optical part of the network, i.e. by the servers that ran the testing tools or the two routers, and that increasing the number times the data signal was sent through the network did not introduce additional jitter for the range of configurations and settings used in this project.

The previously mentioned result that deviates significantly from the others is most likely a random outlier. This is supported by the fact that the results on both sides in the graph (in Figure 3.3) are significantly lower. Additionally, every single point on the graph is the average of three successive tests. The value in question is the average of 0.07 ms, 0.006 ms and 0.009 ms. Going through the logs from the testing tool shows that the reason why this first result if so much higher than the other two is because of a short spike in the jitter during the last 10-15 seconds of the 10-minute test.

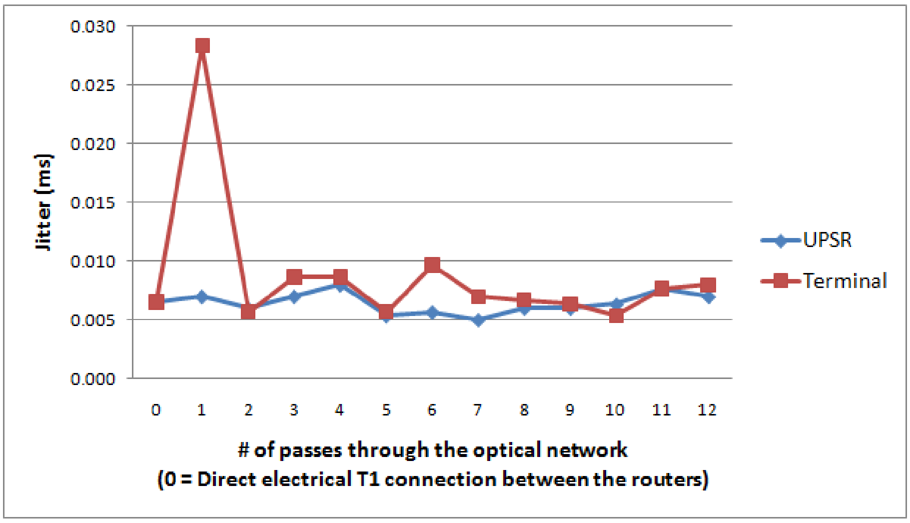

Figure 3.3: Jitter in the UPSR and terminal configurations for a sending rate of 1 Mbps.

3.5.2 Sending Rate: 1.48 Mbps

The results of the reference tests that were run between the two directly connected routers are given in Figure 3.4. These results show that increasing the sending rate from 1 Mbps to 1.48 Mbps significantly increased both the amount, as well as the variability of the jitter. The jitter varied between 0.021 ms (1.22% of the average ping time) and 0.735 ms (42.7% of the average ping time). In these tests, the average jitter was 0.25 ms and the sample standard deviation was 0.21 ms (85.94% of the average jitter), which is significantly more than for the same set of tests run using a sending rate of 1 Mbps. This increase was most likely a result of the fact that we were at that point approaching the maximum capacity of the channel, and that we were seeing the results of buffering in the two routers along the path.

The results of the tests that were run in the UPSR and terminal network configurations are graphed in Figure 3.5. Based on these results, the jitter in both of these network configurations appears to be randomly distributed in the range from 0.05 ms to 0.45 ms. There is also no apparent pattern indicating that increasing the number of times the data signal was sent through the network increased the amount of jitter for the range of configurations and settings tested in this project.

However, all of the data points in the graph in Figure 3.5 fall within the same range as those in Figure 3.4. This indicates that the jitter in both the UPSR and the terminal network was introduced outside of the optical part of these networks.

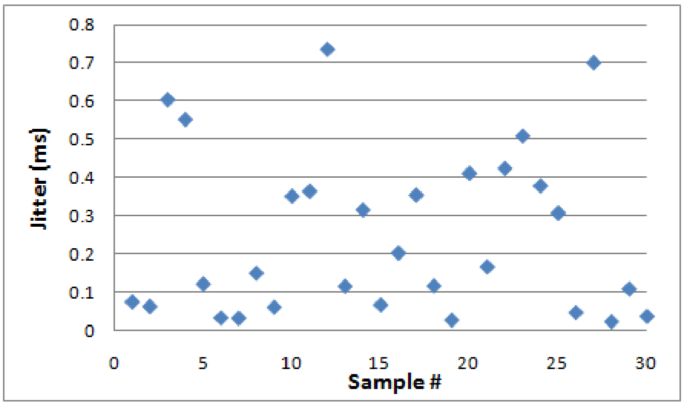

Figure 3.4: Jitter in the configuration shown in Figure 2.2 for a sending rate of 1.48 Mbps.

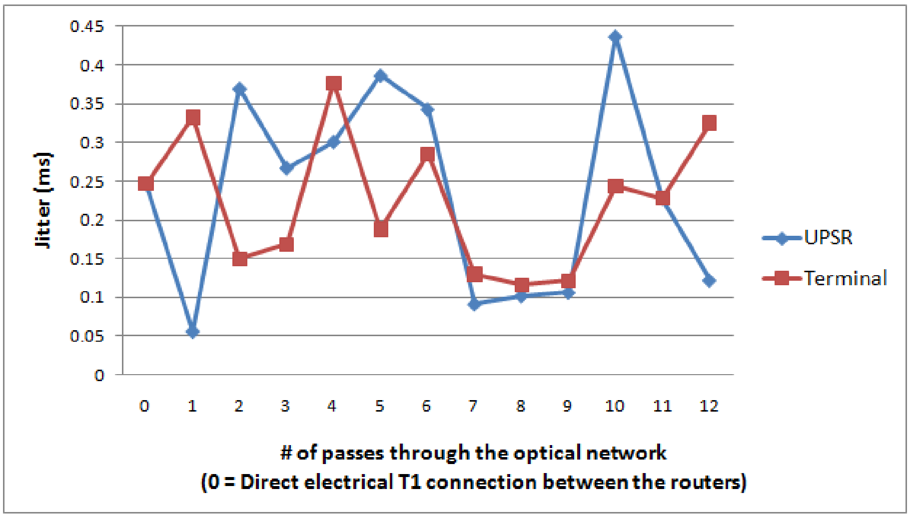

Figure 3.5: Jitter in the UPSR and terminal configurations for a sending rate of 1.48 Mbps.

Based on the test results for both sending rates of 1 and 1.48 Mbps, it appears that the jitter in the two main configurations was highly dependent on the sending rate. When the sending rate was low, there was little jitter, but when the sending rate approached the maximum capacity of the channel, the jitter increased significantly. Increasing the number of times that the data signal was sent through the network, on the other hand, did not appear to introduce much (if any) additional jitter in the range of configurations tested in this project.

The measured jitter in all of the tests that were run in both the UPSR and terminal network configurations, for both sending rates, was well below the threshold value of 10 ms defined in Chapter 2. As the test results give no clear indication that increasing the number of o-e-o conversions, as well as the total length of the network, introduces additional jitter, it is not possible to extrapolate a maximum number of times that the signal can be sent through either of these networks before the jitter reaches the threshold value.

Chapter 4

Conclusion

This project had two main objectives. The first one was to investigate how an increasing number of successive optical-electrical-optical (o-e-o) conversions of a DS1 signal in a long-distance SONET network affected the round-trip time and TCP/UDP performance. The second objective was to determine the maximum number of times this DS1 signal could be sent through the network before the signal quality became unacceptably low. In order to meet these objectives, a series of performance tests were carried out. The most important findings from these tests, grouped by metric, are given below:

Round-Trip Time

- For each additional time that the data signal was sent through the network in either of the two main configurations, there was an almost perfectly linear increase in the round-trip time. On average, this increase was 0.84 ms in the UPSR configuration, and 0.79 ms in the terminal configuration. As each optical network element adds a small amount of delay, due to signal processing and switching, the higher number of network elements in the UPSR configuration compared to the terminal configuration was most likely the cause of this difference.
- All measured round-trip times in this project were well below the upper limit, which was set at 300 ms. Assuming that we can extrapolate the results from the round-trip time tests, it would be possible to send the data signal through the network in the UPSR and terminal configurations approximately 350 and 375 times, respectively, before reaching this limit.

Bit-Error Rate

- The bit-error rates in both the UPSR and terminal configurations were very low, typically 10^{-11} or lower.
- Increasing the number of times that the data signal was sent through the network in either of the two main configurations incrementally from one to twelve did not appear to increase the bit-error rate.

Datagram Loss

- The datagram loss rates in both of the main network configurations were very low, typically 10^{-6} or lower. This was the case for both the 1 and 1.48 Mbps sending rates.
- As the few datagram losses that did occur appear to be randomly distributed, it is difficult to determine whether or not increasing the number of times the data signal is sent through

the network increases the datagram loss rate. In order to reliably answer this question, more and longer tests would be necessary.

- As all of the measured datagram losses were far below the upper limit of 0.05%, it should be possible to send the data signal through the network a significantly higher number of times than what was done in this project before reaching the upper limit. However, because of the previously mentioned random distribution of the results, it is not possible to extrapolate these to determine a likely estimate for this number.

Throughput

- The result of every single throughput test, for both the UPSR and terminal network configurations, was approximately 1475 kbps. Low bit-error rates, as well as minimal amounts of packet loss (indicated by low datagram loss rates) were most likely important in avoiding time-consuming retransmissions. Additionally, appropriately sized TCP window sizes contributed to the high utilization of the T1 circuit.

Jitter

- For both sending rates of 1 and 1.48 Mbps, increasing the number of times that the data signal was sent through the network in either of the two main configurations incrementally from one to twelve did not appear to introduce additional jitter (most likely because capacity for the DS1 data signal was statically allocated in the very precisely timed SONET networks used in this project). On the contrary, the test results indicate that most, if not all of the jitter was introduced outside of the optical part of the network.
- For the configurations tested, the amount of jitter in the network appears to be mostly dependent on the datagram sending rate. Using a low sending rate, there was hardly any jitter, and there was little variation within the results. Using a sending rate close to the maximum capacity of the channel, there was a significant increase in the amount, as well as the variability of the jitter.
- All measured values for jitter were well below the upper limit, which was set at 10 ms. Based on the test results, it is not possible to determine how much, if any, additional jitter is introduced when the number of times the data signal is sent through the network is increased. Therefore, it is not possible to estimate how many times the signal can be sent through the network before the jitter reaches the 10 ms limit.

It should be noted that in order to ensure high-quality optical signals between the network elements, the lengths of all fibre spans were kept well within the limitations of the equipment. Had the tests been carried out in a network where the quality of the signals was significantly lower, the results would most likely have turned out very differently.

Although the network topologies are different for the UPSR and terminal configurations, the total length of the fibres used in each of these were kept as similar as possible, at approximately 49 km. Some additional fibres, around 30 meters in total, were necessary to create the core ring in the UPSR configuration. However, this is negligible compared to the total length of the two networks. The number of network elements the data signal passes through on a one-way trip across the network is also different for each of the two configurations. In the terminal configuration, the data signal passes through two network elements, whereas it passes through five in the UPSR configuration.

Bibliography

[1] *ACM Digital Library*. Association for Computing Machinery. URL: `http://www.acm.org/dl/`.

[2] John C. Bellamy. *Digital Telephony*. Third Edition. John Wiley & Sons, 2000. Chap. 4.

[3] *Computer and Information Systems Abstracts*. ProQuest. URL: `http://csaweb108v.csa.com/factsheets/computer-set-c.php`.

[4] *FLASHWAVE®4500*. Release 10.1, Issue 1. Fujitsu. 2008.

[5] *FLASHWAVE®4500 Multiservice Provisioning Platform - Interactive Datasheet*. Fujitsu. 2008.

[6] *IEEE Xplore*. Institute of Electrical and Electronics Engineers. URL: `http://www.ieee.org/ieeexplore/`.

[7] *Iperf*. National Laboratory for Applied Network Research/Distributed Applications Support Team. URL: `http://iperf.sourceforge.net/`.

[8] R. Ramaswami, K. Sivarajan, and G. Sasaki. *Optical Networks: A Practical Perspective*. Third Edition. Morgan Kaufmann, 2009. Chap. 6 and 9.

[9] *Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria (GR-253-CORE)*. Issue 4. Telcordia Technologies. 2005.

[10] Tim Szigeti et al. *Cisco TelePresence Fundamentals*. Cisco Press, 2009. Chap. 3.

[11] Andrew S. Tanenbaum. *Computer Networks*. Fourth Edition. Prentice Hall, 2003. Chap. 1 and 2.

Appendix A

Equipment Descriptions

This appendix is intended to provide a brief introduction to the main hardware components used in this project. More in-depth information is available in the referenced documentation. First, Section A.1 presents the Fujitsu Flashwave 4500 Multiservice Provisioning Platform (Flashwave 4500). This is followed by Section A.2, which covers the remaining hardware used in this project, such as routers, servers and fibre spools.

A.1 The Fujitsu Flashwave 4500

The Flashwave 4500 is a versatile multiservice optical transport platform commonly deployed in carrier infrastructures today. As a network element in a SONET network, it can be configured to serve as a terminal multiplexer, an add/drop multiplexer, or as a digital cross-connect. Its modular design, which supports a variety of plug-in units, also makes it highly customizable. These plug-in units can be chosen by the user to provide the necessary services and interface rates. The technical information presented in the remainder of this section is based on [4] and [5].

As can be seen in Figure A.1, the Flashwave 4500 is logically divided into the following four sections:

- **Flexible Interface Slots** - The traffic interface plug-in units installed in this section send and receive signals to and from external equipment. Plug-in units with rates ranging from DS1 to OC-192 are supported. Internally in the Flashwave 4500, signals are routed between these plug-in units and the two main switch fabrics.
- **STS Switching** - The two main switch fabrics (working and protect) are installed in this section. These handle the cross-connecting of signals between the different traffic interfaces, as well as protection switching. The main switch fabrics operate on electrical signals at the STS-1 level.
- **Synchronization** - The plug-in units installed in this section provide timing for the system. The timing can be received from an external *building integrated timing source* signal, recovered from an incoming OC-N signal or based on an internal clock (stratum 3).
- **System Management** - The plug-in units installed in this section handle overall command and control functions for the system. One of these critical functions includes storing and executing the system software.

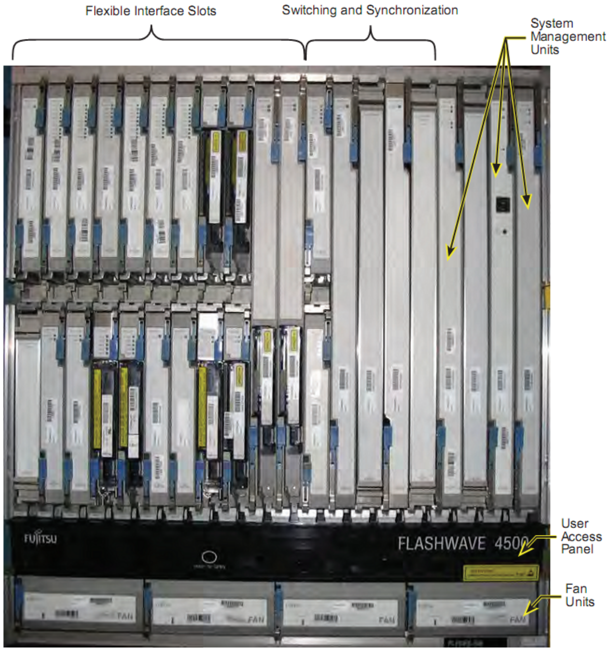

Figure A.1: Front view of the Flashwave 4500 (from [4]).

In this project, four single-shelf Flashwave 4500s were used. These were all of the SHF2 type, running release 10.1 of the system software. Timing for the network elements was provided by their internal clocks. The remainder of this section provides a short introduction to the most important plug-in units in these network elements.

A.1.1 Plug-in Units

Main Switch Fabric (SFA2-SF21)

Each of the Flashwave 4500s used in this project has two *SFA2-SF21* non-blocking switch fabrics installed; one working and one protect. Each switch fabric can cross-connect up to 1344x1344 (in x out) STS-1s, which gives it a total switching capacity of 70 Gbps. These switch fabrics also support protection switching for terminal, linear and ring applications.

VT Switch Fabric (IFA2-VF11)

As the main switch fabrics operate at the STS-1 level, *virtual tributary* (VT) switch fabrics are needed in order to cross-connect the DS1 signals used in this project. *IFA2-VF11* non-blocking VT switch fabrics were therefore installed in two of the Flashwave 4500s that were used. Each of these can cross-connect up to 1344x1344 VT1.5s, for a total switching capacity of 2.5 Gbps.

OC-48 Optical Interface (IFA2-L8B1)

Each of the Flashwave 4500s has two *IFA2-L8B1* 1-port OC-48 short-reach optical interface units installed. These interfaces operate at a wavelength of 1310 nm, and are used with single-mode fibres. Because they are dispersion limited at 3 km, these interfaces can only support fibre spans of 3 km or less.

OC-12 Optical Interface (IFA2-L2C5)

Each of the Flashwave 4500s has one *IFA2-L2C5* 4-port OC-12 long-reach optical interface plug-in unit installed. These interfaces operate at a wavelength of 1310 nm, and are used with single-mode fibres. They are loss limited at 60 km, which means that the longest fibre span these interfaces can support is 60 km.

DS1 Traffic Interface (IFA2-D1V2)

In order to be able to add and drop DS1 signals, *IFA2-D1V2* units were installed in the same two Flashwave 4500s (one in each) that were equipped with VT switch fabrics. Each of these units supports the adding and dropping of up to 14 DS1 signals.

A.2 Other Equipment

A.2.1 The Cisco 2821 Router

The Cisco 2821 is an integrated service router with a modular design that allows for several different types of interface modules to be installed. Each of the two routers used in this project are equipped with two Gigabit Ethernet ports, four Fast Ethernet ports, as well as a *VWIC-1MFT-T1* 1-port T1 module. The T1 module provides physical layer T1 connectivity, and supports both data and voice applications. Both routers used are running Cisco IOS version 12.4(11)

In this project, these Cisco 2821 routers served two purposes. The first was to connect the servers that ran the network performance testing software, which are only equipped with Ethernet interfaces, to the core TDM network. Secondly, the routers' built in bit-error rate testers were used to measure the bit-error rates for the different network configurations.

A.2.2 The Sun Blade Server

The network performance testing software used in this project was run on two Sun blade servers with 1 GHz AMD Opteron processors and 1 GB RAM. These are running the 64-bit Ubuntu Linux version 6.06.1 LTS operating system.

A.2.3 Fibre Spools

The long-distance links in the network were set up using three fibre spools, each consisting of two single-mode fibres. Two of these spools contain 20.7 km of fibres each. The fibre type used in these is *Corning SMF 28e+*, which has an attenuation of 0.330 dB/km at an operating wavelength of 1310 nm. The fibres in the third spool are 7.7 km long. The fibre type used in this is *Corning SMF 28*, which has an attenuation of less than 0.3 dB/km at an operating wavelength of 1310 nm.

Appendix B

Equipment Configurations

This appendix describes how the main system components in this project were configured. First, Section B.1 shows the relevant cross-connects in the Fujitsu Flashwave 4500 network elements. This is followed by the Cisco 2821 router configurations in Section B.2. Finally, the commands used for configuring the network interface cards in the Sun blade servers are given in Section B.3.

B.1 Flashwave 4500 Cross-Connects

This section does not show every single cross-connect from all four network elements. However, those included should be sufficient to describe how the two SONET networks were configured.

B.1.1 Cross-Connects in the UPSR Configuration

The cross-connects shown in this section are those that were used when the network was set up in the configuration shown in Figure 2.4 (in Chapter 2). For the remainder of this section, network elements will be referred to by the rack they are installed in. These rack numbers correspond to those in Figure 2.4.

Figure B.1 shows a VT1.5 cross-connect in Rack #1. Here, the T1 circuit from port 1 on the DSX-1 panel is cross-connected with the OC-12 circuit between Rack #1 and Rack #2.

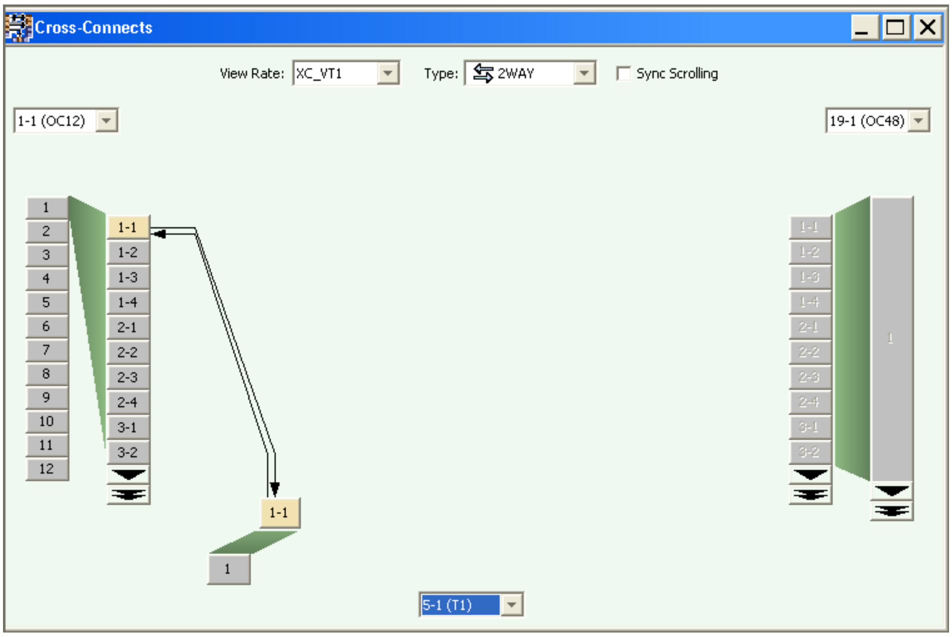

Figure B.1: Rack #1 - UPSR configuration: Cross-connect between a T1 and an OC-12 circuit.

Figure B.2 shows the OC-48 cross-connect between the two OC-48 line cards in Rack #1, which is part of the UPSR ring.

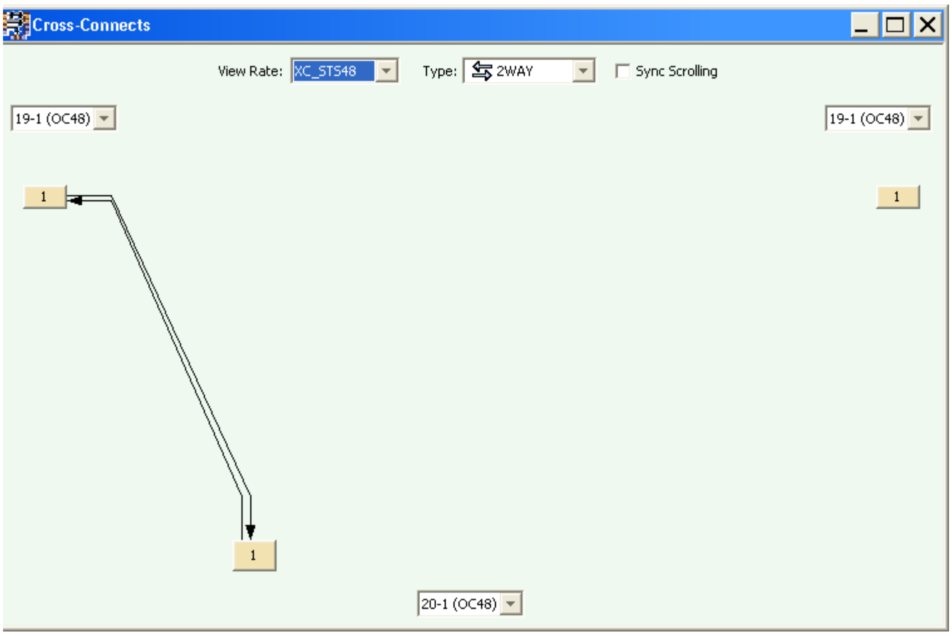

Figure B.2: Rack #1 - UPSR configuration: OC-48 pass-through cross-connect

Figure B.3 shows the *cross-connects view* for Rack #1 when the DS1 signal is sent through the UPSR network four times. The cross-connects marked XC_VT1 represent adding and dropping of T1s, whereas those marked XC_STS48 are part of the OC-48 UPSR ring. Note that this is a list of one-way cross-connects, which is why it shows eight XC_VT1s and two XC_STS48s.

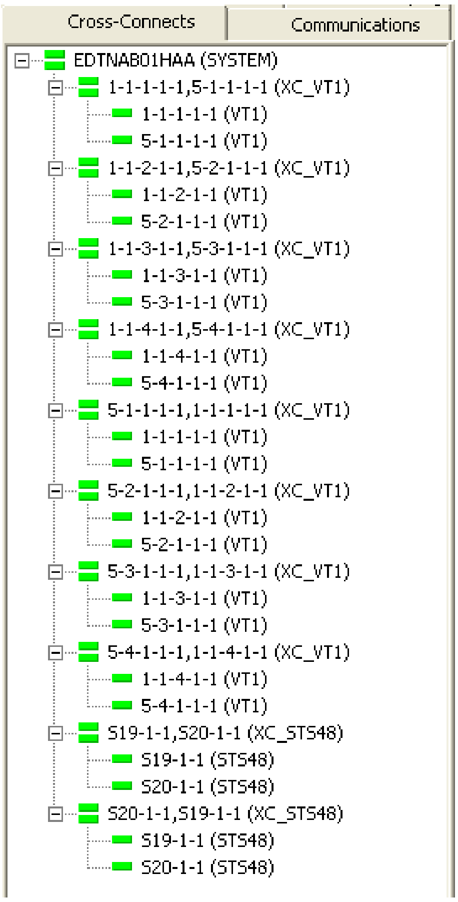

Figure B.3: Rack #1 - UPSR configuration: Cross-connects view when the DS1 signal is sent through the network four times.

Figure B.4 shows how the OC-12 circuit between Rack #1 and Rack #2 is cross-connected onto the ring in Rack #2. Because the UPSR configuration is used, the OC-12 signal is sent in both directions around the ring.

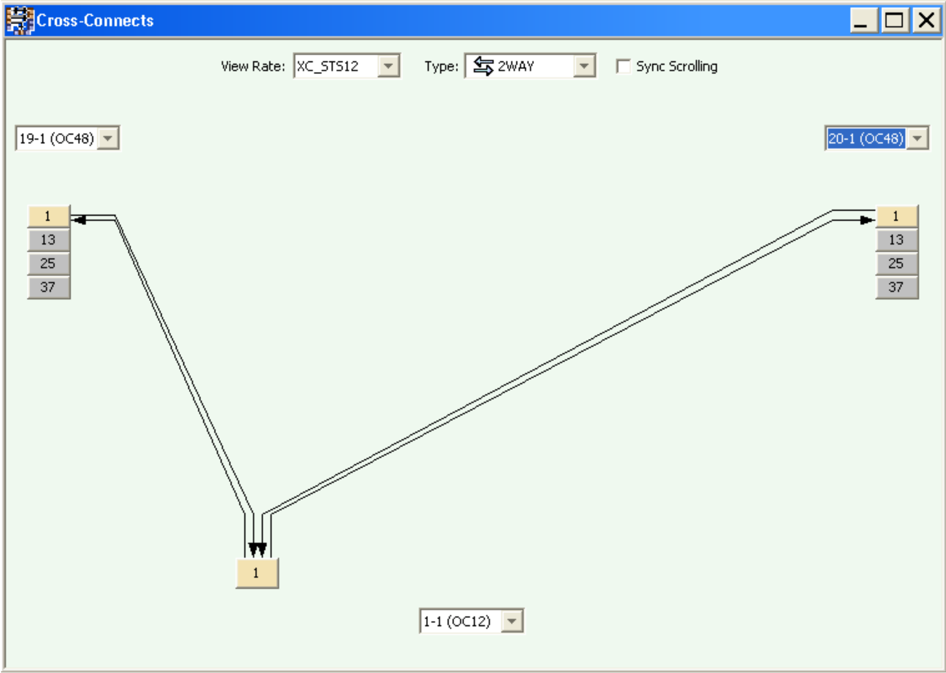

Figure B.4: Rack #2 - UPSR configuration: Cross-connection between the OC-12 and OC-48 circuits.

The cross-connects in Rack #3 are not shown here because these are the same as for Rack #1. Similarly, the cross-connects in Rack #4 are not shown because these are the same as in Rack #2.

B.1.2 Cross-Connects in the Terminal Configuration

The cross-connects discussed in this section are those that were used when the network was set up in the terminal configuration, as shown in Figure 2.5 (in Chapter 2). The network element names used in this figure corresponds to those used in this section.

As the cross-connects are the same in both network elements, NE1 and NE2, Figure B.5 is sufficient to illustrate these. The T1 circuit from port 1 on the DSX-1 panel is cross-connected with the OC-12 circuit between NE1 and NE 2.

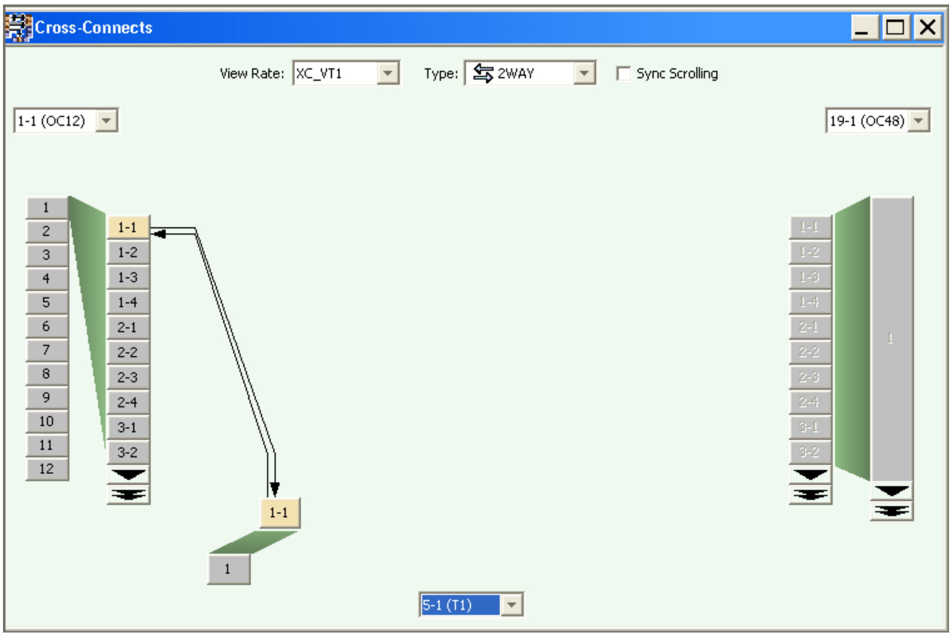

Figure B.5: NE1 and NE2 - Terminal configuration: Cross-connect between a T1 and an OC-12 circuit.

B.2 Cisco 2821 Configurations

The relevant parts of the running configurations for the two Cisco 2821 routers that were used are shown below. Configuration commands for unused interfaces, as well as default values have been removed.

Router in rack #2

```
hostname C2800-R2
!
controller T1 0/3/0
 framing esf
 clock source internal
 linecode b8zs
 cablelength short 133
 channel-group 0 timeslots 1-24
!
interface GigabitEthernet0/1
 ip address 10.0.0.1 255.255.255.0
 duplex auto
 speed auto
!
interface Serial0/3/0:0
 ip address 192.168.1.1 255.255.255.0
```

```
!
ip route 10.0.1.0 255.255.255.0 Serial0/3/0:0
!
end
```

Router in Rack #4

```
hostname C2800-R4
!
controller T1 0/3/0
 framing esf
 clock source line
 linecode b8zs
 cablelength short 133
 channel-group 0 timeslots 1-24
!
interface GigabitEthernet0/1
 ip address 10.0.1.1 255.255.255.0
 duplex auto
 speed auto
!
interface Serial0/3/0:0
 ip address 192.168.1.2 255.255.255.0
!
ip route 10.0.0.0 255.255.255.0 Serial0/3/0:0
!
end
```

B.3 Sun Blade Server

This section provides the commands necessary to set up the network interfaces on the Sun blade servers, which are running Ubuntu Linux. Interface eth3 is the *GbE management* port on the blade servers.

Blade Server #2

```
sudo ifconfig eth3 10.0.0.4 netmask 255.255.255.0 up
sudo route add -net 10.0.1.0 netmask 255.255.255.0 eth3
```

Blade Server #4

```
sudo ifconfig eth3 10.0.1.4 netmask 255.255.255.0 up
sudo route add -net 10.0.0.0 netmask 255.255.255.0 eth3
```

Appendix C

Background Material

This appendix provides an introduction to a few topics that are relevant to this project. First, Section C.1 briefly explains *Time Division Multiplexing* (TDM). This is followed by Section C.2, which covers the *T1* carrier system, and the associated *Digital Signal Level 1* (DS1) signalling scheme. Finally, Section C.3 gives an introduction to *Synchronous Optical Networking* (SONET).

C.1 Time Division Multiplexing

In computer networks, it is often desirable to transport multiple lower-rate data streams together over a high-rate shared medium. In wired networks, the motivation for doing this is often to lower costs by reducing the number of cables needed between nodes in the network.

At the sender side, the lower-rate streams are combined into a composite higher-rate stream for transport. This process is called *multiplexing*, and is handled by a device called a *multiplexer*. This composite stream is then sent across the network to the receiver side. Here, the composite stream is split up, and each one of the individual lower-rate streams is extracted. This process is known as *demultiplexing* and is carried out by a device called a *demultiplexer*. There are several different multiplexing techniques available, with some of the most common being *Time-Division Multiplexing* (TDM), *Frequency-Division Multiplexing* (FDM) and *Wavelength-Division Multiplexing* (WDM), which is a variant of FDM. As the network used in this project is based on TDM systems, this is the only technique that will be described further.

In a TDM system, access to the shared medium is divided into periods of time, called *timeslots*, in which data from only one data stream may be transmitted. These timeslots are then assigned to the individual data streams. Timeslots can either be assigned to streams statically in what is called *Synchronous TDM*, or on demand in what is called *Asynchronous TDM*. Synchronous TDM is used in the systems considered in this project, so unless otherwise noted, any reference to TDM in the remainder of this report will be to Synchronous TDM. An example of one such system is shown in Figure C.1.

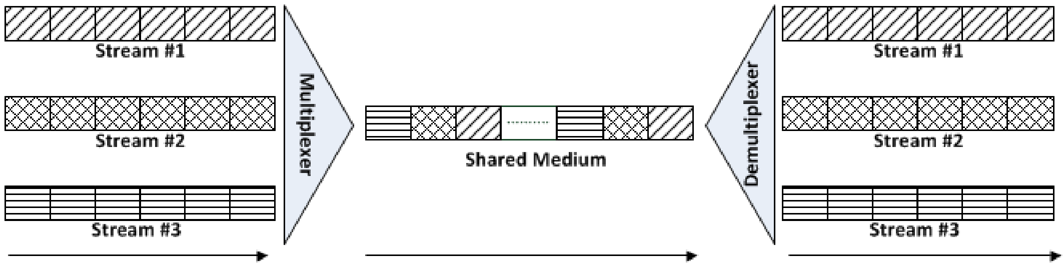

Figure C.1: Example TDM system where three data streams are multiplexed, transported and then demultiplexed.

C.2 The T1 Carrier

The *T1* carrier is part of the T-carrier hierarchy, which was developed several decades ago to allow multiple voice channels to be multiplexed, using TDM, for transport over a single line. A T1 circuit can carry 24 channels, and higher-order T-carriers, such as the 72-channel T2 and the 672-channel T3, are created by multiplexing T1s. Later, as the demand for data services appeared, the T-carrier systems were adapted to also support digital data transmission. Unless otherwise noted, the information in the remainder of this section is based on [11] and [2].

Framing

The signalling scheme used over T1 lines is called the *Digital Signal Level 1* (DS1). In the DS1 signal, data is transmitted in 193-bit units called *frames* 8000 times per second. As can be seen in Figure C.2, each frame consists of 24 8-bit timeslots, one for each channel, plus an additional *framing bit*. The choice of 8 bits per channel and a frame rate of 8000 frames per second was made to accommodate *Pulse Code Modulation* (PCM) coded voice channels. Such channels produce 8-bit samples at a rate of 8000 samples per second, for a total bit rate of 64 kbps. This gives the T1 circuits a gross data rate of 1.544 Mbps, out of which 1.536 Mbps is available for the user payload.

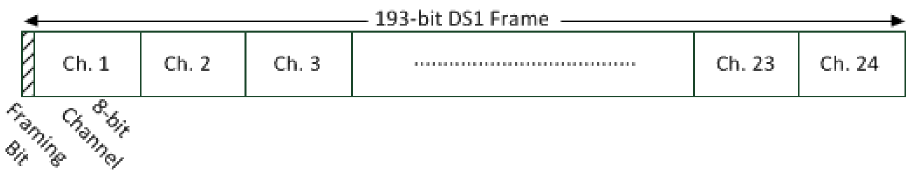

Figure C.2: Illustration of a 193-bit DS1 frame.

The DS1 standard supports several different framing formats. These all use 193-bit frames, but they each have slight variations in how they use some of the bits in the frame. The *Extended Super Frame* (ESF) format is the most recent, and is the framing format used in this project. It is therefore the only framing format that will be covered here. The advantage of using ESF combined with the *B8ZS* line code (described below) is that it provides functionality for performance monitoring, frame synchronization and error checking for circuits carrying digital data without using any of the bits in the 24 8-bit timeslots. These three functions are handled by using only the framing bit in each frame. This means that the all 24 64 kbps channels are available to carry user data.

In the original T1 framing format, the only purpose of the framing bit was to ensure *frame synchronization* between the sender and the receiver. That is, letting the receiver know where a DS1 frame starts and ends by marking the beginning of each frame with specific pattern (alternating ones and zeros). In the ESF format, this bit has been redefined to serve multiple purposes. In ESF, 24 frames are grouped together into what is called a *superframe*, which gives a total of 24 available framing bits per superframe. Out of these, 6 bits are used for frame synchronization, 12 bits are used for interrogating for and reporting of performance statistics and 6 bits are used for error checking.

Line Codes

A *line code* describes how data is to be encoded for transmission over some physical medium. In T1 systems, *Alternate Mark Inversion* (AMI) and *Binary 8 Zeroes Substitution* (B8ZS) are two such options. Both of these are variations of *bipolar* line codes that use three different signal levels to transmit binary data (e.g. +3V, 0V and -3V). A binary 0 is encoded as a zero voltage level, whereas a binary 1 is encoded by alternating positive and negative voltage levels. That is, two successive binary 1s are encoded with opposite voltage levels. If two successive binary 1s are encoded with the same voltage level, it is referred to as a *bipolar violation*.

In order to maintain synchronization between a sender and a receiver, it is important that the received signal contains a sufficient number of transitions between the different voltage levels. These transitions help the receiver to determine the bit boundaries, that is, where each bit starts and ends. Note that this type of synchronization for single bits is not the same as the frame synchronization discussed in the previous section.

Bipolar Line codes are used instead of line codes that always encode binary 1s with the same voltage level because they force high/low and low/high transitions in signals containing long strings of 1s. However, the problem with AMI is that for long strings of 0s, it does not produce any signal transitions, which can potentially cause synchronization problems. In T1 systems using AMI, this is handled by always setting certain bits in each frame to 1. This ensures a minimum number of signal transitions, but because the bits that are used for this purpose could otherwise be used to transport data, it is not an ideal solution.

A better way to solve the problem of long strings of zeroes without affecting the bit rate of the signal, which is used in many T1 systems, is to use the B8ZS line code. Using B8ZS, a sender replaces a string of 8 zeroes with an 8-bit code that purposely contains a bipolar violation. An example of a binary violation is transmitting the binary 11 as two successive positive pulses instead of one positive and one negative (the order depends on the preceding pulse). The receiver recognizes these codes and replaces them with a string of 8 zeroes. B8ZS is the line code used in this project.

C.3 Synchronous Optical Networking (SONET)

The *Synchronous Optical Networking* (SONET) standard for multiplexing and data transmission over optical fibres is extensively used in the carrier infrastructure of North America. A related standard that is very similar to SONET, *Synchronous Digital Hierarchy* (SDH), is usually deployed in Europe, Japan and some other parts of the world. Both SONET and SDH are Synchronous TDM systems, as described in Section C.1. The SONET standard is fairly large and complex, and therefore, only the portions that are relevant to this project are covered in this section. Unless otherwise noted, this section is based on chapters 6 and 9 of [8].

The *Open Systems Interconnection* (OSI) model from the *International Standardization Organization* (ISO) is often used to decompose the functionality of a communication system into seven parts called layers. Each layer in this model provides a set of services to the layer above, and uses the services of the layer below. An illustration of the OSI model is shown in Figure C.3. For an introduction to the different layers in this reference model, see [11]. As SONET deals directly with the transmission of bits over optical fibres, it belongs at the physical layer in this reference model.

SONET can transport signals for a variety of higher-layer protocols. It was originally developed to transport circuit switched applications, such as voice and DS1/DS3 signals. However, data from packet networks (e.g. Ethernet) can also be adapted at the data link layer for transport across a SONET network using the *Generic Framing Procedure* (GFP).

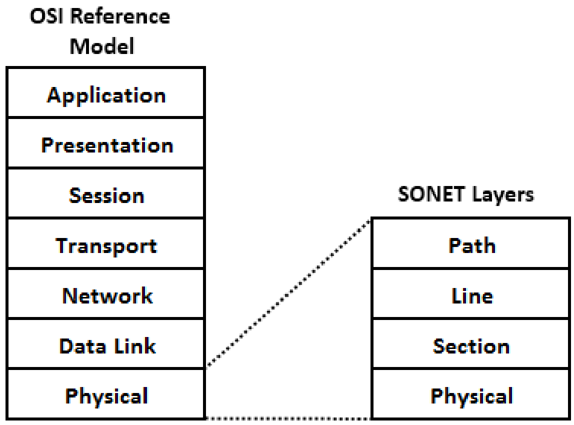

Figure C.3: The seven layers of the ISO OSI reference model and the four SONET layers.

C.3.1 Multiplexing and Frame Structure

The basic SONET rate is the 51.84 Mbps *Synchronous Transport Signal Level 1* (STS-1). SONET also supports a variety of higher-rate signals that are created by multiplexing STS-1s. A table of the currently standardized signal rates for SONET is given in Table C.1. It should be noted that the STS-N signals are electrical, and that these must be converted to optical form before transmission over the fibre. As part of this process, the signal is also scrambled to avoid long strings of 0s or 1s that may otherwise cause synchronization problems. For every STS-N, there is a corresponding *Optical Carrier* (OC-N) with the same bit rate. For instance, the optical signal corresponding to an STS-12 is called OC-12.

SONET transmits data in frames 8000 times per second, regardless of the data rate. This means that as the data rate increases, the frame rate stays the same, but the size of each frame increases. In any TDM system operating at high bitrates, the duration of each individual bit is very short, and precise timing in the network elements is therefore crucial. If the clock rates of two network elements are not synchronized, the receiver of a signal might occasionally sample the channel at the wrong time, and will in that way receive an erroneous bit. In SONET, consistent timing between all network elements can be ensured by synchronizing all clocks to a high-precision master clock.

OC-N Level	STS-N level	Bit rate (Mbps)
OC-1	STS-1	51.84
OC-3	STS-3	155.52
OC-12	STS-12	622.08
OC-24	STS-24	1244.16
OC-48	STS-48	2488.32
OC-192	STS-192	9953.28
OC-768	STS-768	39,814.32

Table C.1: SONET bit rates.

The basic STS-1 frame consists of 810 bytes, logically arranged into 90 columns and 9 rows, as illustrated in Figure C.4. The first three columns contain two types of overhead called *line* and *section* overhead, which is used for signalling and network management. These two types of overhead are together called the *transport* overhead. The remaining 87 columns contain what is called the STS-1 *Synchronous Payload Envelope* (SPE). The STS-1 SPE also contains one column of overhead, called the *path* overhead, as well as two columns with fixed content that cannot be used for data payload. The remaining 84 columns, or 756 bytes, contain the payload. Higher-rate STS-N frames are created by interleaving N STS-1 frames byte by byte. As an example, an STS-3 frame can be viewed as consisting of 270 columns and 9 rows, with the first 9 columns containing the line and section overheads for each of the STS-1 frames. The bytewise interleaving to create the previously mentioned STS-3 frame is done in the following manner, assuming $B_{x,y}$ denotes byte y from STS-1 frame number x: $B_{1,1}$, $B_{2,1}$, $B_{3,1}$, $B_{1,2}$, $B_{2,2}$, $B_{3,2}$ and so on.

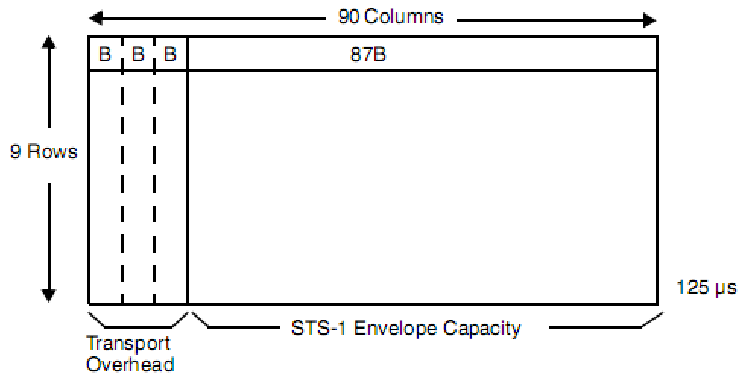

Figure C.4: Illustration of a SONET STS-1 frame (copied from [9]).

Lower-rate non-SONET signals can also be transported within STS-1 frames in so-called *virtual tributary* (VT) structures. The following virtual tributaries have been defined:

- **VT1.5 (1.728 Mbps)** - Carries a 1.5 Mbps signal, like the DS1 signal used in this project
- **VT2 (2.304 Mbps)** - Carries a 2 Mbps signal
- **VT3 (3.456 Mbps)** - Carries a 3 Mbps signal
- **VT6 (6.912 Mbps)** - Carries a 6 Mbps signal

These virtual tributaries are mapped into so-called *VT groups*. A VT group consists of four VT1.5s, three VT2s, two VT3s or one VT6. At the top level, seven VT groups are mapped into the SPE of an STS-1 frame. This mapping is illustrated in Figure C.5.

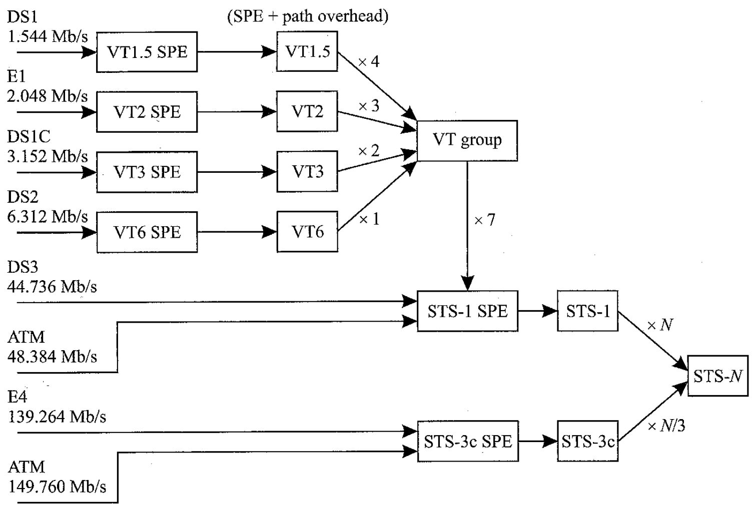

Figure C.5: Mapping of virtual tributaries into an STS-1 SPE (copied from [8]).

C.3.2 Layers and Application Categories

As mentioned previously, SONET belongs at the physical layer in the OSI reference model. Internally, SONET is divided into four layers, shown in Figure C.3. A brief top-down description of these layers, based on [9], is given below:

- **The Path Layer** - Responsible for end-to-end transport between nodes, that is, all the way from source to destination. The path layer maps the various payloads, e.g. DS1s or DS3s, to the format required by the line layer. The path layer also deals with certain types of protection switching.
- **The Line Layer** - Takes care of synchronization and multiplexing of path layer connections onto the physical medium. The line layer is also responsible for certain types of protection switching, which if possible, restores service in the event of link or equipment failure.
- **The Section Layer** - Takes care of transporting STS-N frames across the physical medium using the services provided by the physical layer. The section layer deals with issues such as framing, scrambling and section error monitoring.
- **The Physical Layer** - Transports bits across the physical medium. The physical layer deals with issues such as pulse shaping, power levels and line coding.

SONET can be used for transport over distances ranging from a few meters to a couple of hundred kilometres without amplification. In order to support all of these different distances, a set of broad application categories have been defined. A list of these, based on [9], is given below. As can be seen from this list, the *loss budget* changes depending on the application category. Because an optical signal is attenuated as it traverses the fibre, an increased loss budget usually translates to a greater distance supported between two network elements, assuming that the signal is loss limited and not limited by dispersion. Typically, the type of fibre, transmitter, and receiver used will vary between these application categories.

- **Short Reach (SR)** - optical sections with loss budgets of 0 dB to 4 or 7 dB.
- **Intermediate Reach (IR)** - optical sections with loss budgets of 0, 3 or 6 dB to 11 or 12

dB.

- **Long Reach (LR)** - optical sections with loss budgets 10, 11 or 16 dB to 22, 24 or 28 dB.
- **Very-long Reach (VR)** - optical sections with loss budgets from 22 dB to 33 dB.
- **Ultra-long Reach (UR)** - optical sections with loss budgets of 33 dB to 44 dB.

C.3.3 Network Topologies and Infrastructure Elements

SONET can be used for a variety of applications ranging from providing intraoffice connectivity to long-haul transport in carrier networks. Part of the reason why it has become popular in carrier infrastructures is the built-in features that help provide high *availability*, which is crucial in networks that carry traffic from thousands of users at a time. In order to achieve such high availability, it is important that the network is *survivable*, meaning that it can continue to provide service even in the presence of failures. One important SONET feature that helps ensure network survivability is *protection switching*, which involves having some redundant capacity in the network that can be used to re-route traffic around failures. Using protection switching, service in a SONET network can be restored in less than 60 ms (10 ms to detect the failure and 50ms to do the protection switching) following a failure.

SONET supports several different network topologies that are combined with protection switching to provide different levels of network survivability. Which one of these is chosen depends on the level of protection required, cost and so on. The remainder of this section gives an introduction to each of these topologies.

Before going into the description of the topologies, a couple of terms should be explained. In SONET, there is a distinction between *working paths* and *protections paths*. A working path carries traffic during normal operation of the network, whereas the protection path only carries traffic in the event of failures. Ideally, the working and protection paths are diversely routed to minimize the risk of losing both paths in the event of a failure such as a fibre cut. A distinction should also be made between *unidirectional* and *bidirectional* protection schemes. In this context, unidirectional means each direction of traffic is handled independently. E.g. if the working fibre carrying traffic from A to B is cut, only traffic going in that direction will be switched to the protection fibre. Traffic flowing from B to A will still be carried over the working fibre. In a bidirectional scheme on the other hand, a switch from either working to protection fibres, or the other way around, will affect traffic in both directions.

Terminal and Linear Applications

The simplest SONET topology is the *terminal* application, which is a point-to-point link. Here, all traffic is added at a network element at one end, traverses the link, and is dropped at the network element at the other end. Related to this is the *linear* application, which is similar, but where there may be one or more network elements in-between the endpoints where traffic can be added or dropped.

Both the terminal and linear topologies can be operated in *unprotected* mode where there is only one fibre in each direction between the network elements. In this scenario, failure anywhere along the path will result in loss of service.

In the 1+1 configuration, there are two fibres in each direction between two network elements. One pair serves as the working path and the other as the protection path. Identical signals are transmitted simultaneously on both the working and protection fibres, and the receiver selects the best incoming signal based on certain performance parameters. If a loss of signal or signal degradation is detected for one of the fibres, the receiver can just switch to the other. This method is fast and requires no signalling between network elements.

Another common protection scheme for terminal and linear applications is the 1:1, or more generally, 1:N configuration. Here, one fibre serves as protection for N working fibres. During normal operation, traffic is transmitted on the working fibres and the protection fibres are left unused. In some cases, the protection fibres may also be used for transport of low-priority traffic during normal operation. In the event of a failure, the traffic from the affected fibre is switched to the protection fibre. This method gives a somewhat slower restoration time than the 1+1 scheme as it requires signalling between the two affected network elements. Since this scheme allows several working fibres to share the same protection fibre, it is more cost-efficient than the 1+1 scheme; however, it can only handle single failures.

Both the 1+1 and the 1:N schemes are usually unidirectional, which means that switching traffic between working and protection fibres is made independently for each direction of traffic.

Unidirectional Path Switched Rings

Using SONET rings is a relatively cost-efficient way of providing network survivability and low restoration times, and they are therefore often used in carrier infrastructures today. One of the ring configurations supported by SONET is the *Unidirectional Path Switched Ring* (UPSR), as shown in Figure C.6. Here, every network element in the ring is connected to each of its neighbours using two fibres, one in each direction. This provides two separate paths all the way around the ring with traffic flowing in opposite directions.

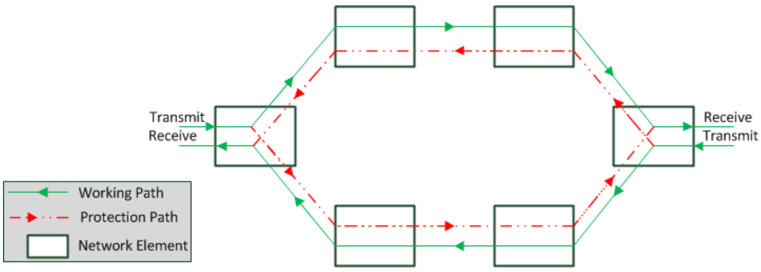

Figure C.6: UPSR during normal operation (adapted from [4]).

Senders transmit identical working and protection traffic simultaneously on both of these paths. Working traffic travels on what is called the working path, while protection traffic travels on the protection path. This ensures that there are two alternate paths between every two network elements in the ring, which protects against the failure of any single node or link. When a node or link fails, the receivers will detect a loss of signal for traffic originating at, or passing through that link/node, and it can just select the signal flowing through the path in the other direction, as can be seen in Figure C.7. As the name suggests, this protection scheme is unidirectional.

The total amount of traffic that can be supported by a UPSR ring equals its line rate. E.g. if

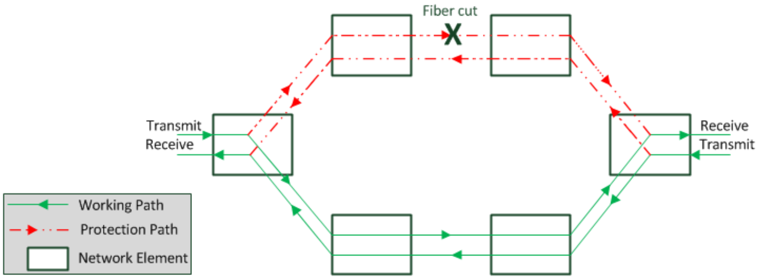

Figure C.7: UPSR with one fibre cut. Notice that only the affected fibre is switched out, the other direction of traffic is unchanged (adapted from [4]).

the line rate of the ring is OC-12, the ring can only support four add/drop nodes if they are each adding and dropping an OC-3 signal. As the protection fibres are always in use, they cannot be used for transporting low-priority data during normal operation. Hence, half the capacity of the ring is reserved for protection purposes. The advantages of UPSR rings are that they are easy to implement and that they do not require any signalling between network elements in the event of failure, and hence have very fast restore times.

Bidirectional Line Switched Rings

In addition to the UPSR, SONET also supports another ring configuration called the *Bidirectional Line Switched Ring* (BLSR). Unlike a UPSR, a BLSR carries working traffic in both directions in the ring, and in most cases traffic between two network elements is routed on the shortest path between them. This means that during normal operation, working traffic does not have to pass through all the other network elements in the ring, and hence a BLSR can in many cases support more aggregate traffic than a UPSR. There are two main types of BLSRs, as described below:

- **Two-Fibre BLSR (2F-BLSR)** - Each of the network elements in the ring is connected to its two neighbours using two fibres, one in each direction. There are no dedicated protection fibres; however, half of the capacity (that is, timeslots) of each fibre is reserved for this purpose. In 2F-BLSR rings, traffic can be completely restored for any single failure. An example of a 2F-BLSR ring under normal working conditions is illustrated in Figure C.8.
- **Four-Fibre BLSR (4F-BLSR)** - Each of the network elements in the ring is connected to its two neighbours using four fibres, two in each direction. One fibre in each direction carries traffic during normal working conditions, while the other two are reserved for protection traffic in the event of a failure.

As with UPSRs, half the capacity in both 2F-BLSRs and 4F-BLSRs is reserved for protection purposes. However, the way this capacity is used is different. During normal operation, the protection capacity in both BLSR types is unused. When a failure occurs, traffic is automatically switched from the working timeslots or fibres to the available protection capacity.

When a failure occurs in a 2F-BLSR, the nodes adjacent to the failed segment will route traffic away from it using the timeslots pre-assigned for protection (called a *ring switch*). This traffic will then flow in the opposite direction of what it normally does around the ring until it reaches the far side of the failed segment. Here, it is switched back to the working timeslots of the fibres, and

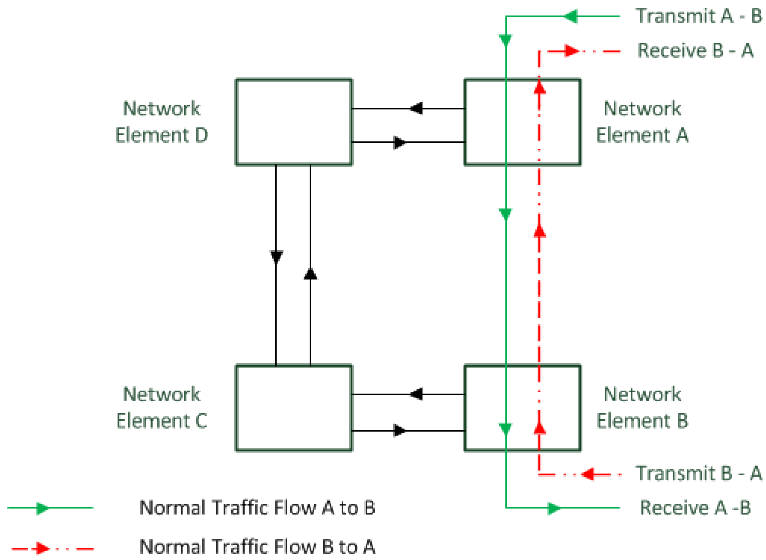

Figure C.8: A 2F-BLSR under normal working conditions (adapted from [4]). The traffic is carried in the working timeslots over both fibres in the span between network elements A and B.

follows its normal path to the destination node from there on. An example of protection switching in a 2F-BLSR is shown in Figure C.9.

When an entire node in a BLSR fails, it is important that the two nodes adjacent to it do not try to re-route connections originating or terminating at the failed node using the protection capacity of the ring. This requires that the two adjacent nodes signal between themselves to determine whether or not they are dealing with a line failure or a node failure before invoking the ring switch for certain connections. The procedure of not restoring connections originating or terminating at a failed node is called *squelching*, and each node in the ring must maintain a *squelching table*, which is a table listing the connections that have to be squelched in the event of a node failure.

As previously mentioned, a 4F-BLSR has two pairs of fibres interconnecting neighbouring network elements, and hence it offers additional protection compared to a 2F-BLSR. If a failure, such as a cut of the working fibres occurs, the adjacent network elements will first try to switch the traffic from the working fibres to the protection fibres on the same span. This is called *span switching*, and does not change the sequence of network elements a specific traffic flow passes through on its way between the sender and the receiver. However, if the span switching fails, e.g. because both fibre pairs on the same span have been cut, or a node itself has failed, ring switching is employed in the 4F-BLSR in a similar manner as in a 2F-BLSR. The difference is that instead of using protection timeslots in fibres on the working path, the dedicated protection fibres are used to reroute traffic around the ring to avoid the failed segment.

Because BLSRs require signalling between network elements to set up alternative paths when failures occur, they are more complex to implement then UPSRs. However, they can support a higher aggregate traffic as the working traffic between two network elements only takes up capacity on the links between them, not necessarily the entire ring, like for UPSRs.

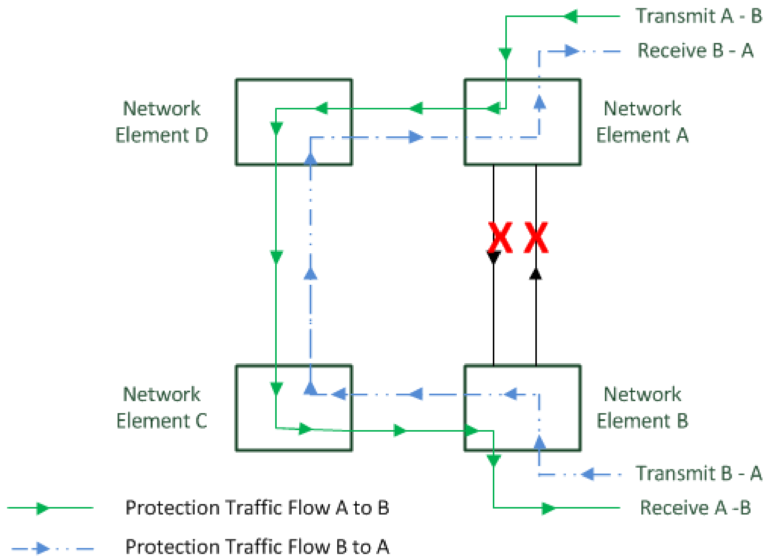

Figure C.9: A 2F-BLSR after fibre cuts in the A to B span (adapted from [4]). The traffic is carried in the protection timeslots over the fibres in the new path between A and B in both directions.

Infrastructure Elements

Two important types of network elements used in SONET networks are *terminal multiplexers* and *add/drop multiplexers*. These are interconnected to form SONET networks, and are also used for connecting client equipment to the network. These two network elements have two types of interfaces that are used for data traffic. The *trunk interfaces* are used to interconnect the multiplexers in the network. The *tributary interfaces*, on the other hand, are used to connect client equipment, such as IP routers, ATM switches and so on to the multiplexers. As the trunk interfaces typically transport aggregate traffic from many tributary interfaces, the trunk interfaces typically support much higher data rates than the tributary interfaces. As an example, a customer may lease a T1 line, which is connected to a tributary interface on a SONET multiplexer. The DS1 signal may then be multiplexed with a number of DS1 signals from other customers into a STS-1 SPE for transport over a OC-48 trunk line.

Terminal multiplexers are the endpoints in terminal (point-to-point) networks, which means that all incoming connections over the trunk interfaces are terminated in these network elements. Outgoing traffic is received from the tributary interfaces, and is added to the signal sent out over one or more trunk interfaces.

Add/drop multiplexers are used in linear and ring networks, and offer more functionality than the terminal multiplexers. Unlike terminal multiplexers, add/drop multiplexers allow high-rate transport connections to pass through them via the trunk interfaces. This is necessary in order to create linear and ring networks. Add/drop multiplexers can pick out individual lower-rate signals from the high-rate transport signals and cross-connect these to tributary interfaces, or to other trunk interfaces. They can also add low-rate signals to the high-speed transport signals. E.g. in an OC-48 ring, an add/drop multiplexer can be set up to add/drop an OC-12 signal to/from the OC-48 signal. One reason for doing this may be to deliver an OC-12 service to a customer connected to one of the OC-12 tributary interfaces.

Another component in SONET networks is the *Digital Cross-Connect*. This type of device is commonly used for interconnecting different SONET network segments. E.g. to connect one ring network to another, or to connect a ring network to a linear network segment.

The final component covered here is the *Regenerator*, which is placed at certain intervals in long-distance network in order to improve the signal quality. This is necessary because an optical signal traversing a fibre will be affected by effects such as attenuation and dispersion, whose effects increase with distance. The regenerators convert the received signals from optical to electrical form, cleans them up and converts them back to optical form for retransmission.

Appendix D

T1 Cable Information

This appendix provides some information about the T1 cables that were used in this project. First, Section D.1 describes how to make the two different types of T1 cables. This is followed by Section D.2, which explains how these cables were physically connected in order to loop a DS1 signal several times through the optical part of the network in order to increase the number of o-e-o conversions, as well as the total distance traveled by this DS1 signal.

D.1 Making the T1 Cables

Two different types of T1 cables were used in this project. Straight T1 cables were used for connecting the T1 card in each of the Cisco 2821 routers to the DSX-1 panels that were connected to the Flashwave 4500 network elements. Crossover T1 cables, on the other hand, were used for looping incoming DS1 signals back across the optical network.

Both of these T1 cable types can be made using two RJ-45 connectors and a category 5 or 6 cable of the desired length. The pin-out for the straight T1 cable is shown in Figure D.1, and the pin-out for the crossover T1 cable is shown in Figure D.2. Note that in both of these figures, the "hook" on each connector is pointing down, and the two connectors are facing each other.

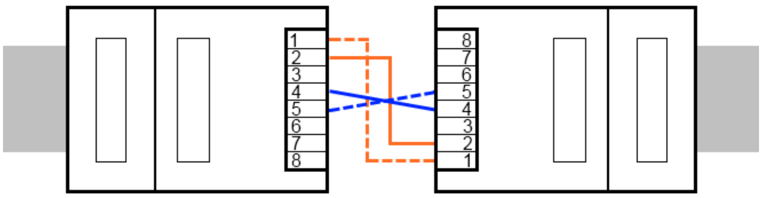

Figure D.1: Pin-out for the straight T1 cable.

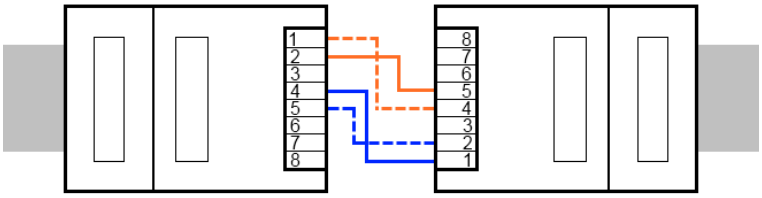

Figure D.2: Pin-out for the crossover T1 cable.

D.2 Connecting the T1 Cables

This section explains how to physically connect the T1 cables, as shown in Figure 2.1 (in Chapter 2). In this example, the DS1 signal is sent through the optical network three times before being dropped.

Left Hand Side (of Figure 2.1)

1. Connect router R1 to port 1 on the DSX-1 panel connected to network element NE1, using a straight T1 cable.
2. Connect one end of a crossover T1 cable to port 2 on the DSX-1 panel, and the other end to port 3.

Right Hand Side (of Figure 2.1)

1. Connect one end of a crossover T1 cable to port 1, and the other end to port 2, on the DSX-1 panel that is connected to network element NE2.
2. Connect router R2 to port 3 on the DSX-1 panel.