

Using wavelengths outside of the Telecom spectrum

Stefan Plug
Remy de Boer

February 9, 2013

Preface

This report is the product of a four week research project done at BeetleFiberOptics, and as part of the Master programme 'System and Network Engineering' at the University of Amsterdam[6].

BeetleFiberOptics is part of the Tallgrass company in Amsterdam and is a manufacturer of Coarse Wavelength Division Multiplexing, CWDM, equipment. A CDWM is a device which aggregates multiple wavelengths of light onto a single fibre. BeetleFiberOptics would like us to find out if, and how, it could use an extra wavelength outside of the traditional CDWM wavelengths to both monitor an optical fibre and as an extra out-of-band Ethernet network which might be used for network management. Primarily BeetleFiberOptics would like to know if we can use relatively cheap hardware, which has the potential of being integrated into their CDWM devices, to achieve this.

We would like to thank the people of BeetleFiberOptics and Tallgrass for their support and knowledge during this project.

Abstract

For a month we have focussed on the following research question:

What applications can the unused wavelengths outside of the Telecom spectrum be used for?

In the first stages of the research we decided to focus on the 1625nm band, as it is closest to the regular C- & L-bands and has the least loss of all the wavelengths in the U-band. We started researching the feasibility of monitoring and out-of-band management links using this frequency.

We tested the 1625nm wavelength to see how it compares to the 1550nm and 1610nm wavelengths. We performed the following tests:

- OTDR test
- Attenuation test per wavelength
- Link stability test using iperf

Our tests showed the 1625nm wavelength suffers from considerably more attenuation than the 1550nm and 1610nm bands do. Even though the attenuation was considerable, the link quality was still good enough to establish a stable Ethernet connection. We used this Ethernet connection to let Raspberry Pis at either end of the link communicate and monitor the attenuation of the link.

The result of our research is that the 1625nm wavelength has shown to be capable of sustaining a stable connection, well suited for out-of-band management and monitoring.

Contents

1	Research question	5
2	Telecom spectrum	6
2.1	Choosing a wavelength	6
3	Proof of Concept	8
3.1	Inteno XG6746	8
3.2	Primary test network	10
3.3	Attenuation test network	11
4	Tests	12
4.1	Ethernet link reliability test	13
4.2	Attenuation test	13
4.3	OTDR tests	16
5	Conclusion	17

1 Research question

It is common practice today to use Wavelength Division Multiplexing, WDM, devices on network links to increase the total amount of bandwidth that a single optical network link can carry. The WDM accomplishes this by assigning each input data stream its own unique light wavelength channel, λ . Not every wavelength is suitable for heavy traffic usage because of the physical characteristics of these channels. Our hypothesis is that these channels could be used for lower speed applications such as monitoring and out of band management. To test our hypothesis we will look at the possibilities of the unused wavelengths outside of the Telecom spectrum. The main research question will be as follows:

What applications can the unused wavelengths outside of the Telecom spectrum be used for?

It is important to be able to have an active monitoring system in place which can send out warnings when it detects either degradation of the link over time or a sudden change in the links characteristics, for example when someone places a tap on the link.

To help us understand how we can effectively monitor an optical link we will ask the following sub-question:

What optical link characteristics should we monitor?

It could happen that the main traffic interface on a switch shuts down which could potentially shut down that section of the network. It would be good to have a separate out of band network link up on another interface which could be used to still manage that device.

During this project we will focus on link monitoring and out of band management, but other usages may arise during our research. If this should happen we shall try to document them.

2 Telecom spectrum

When we look at the wavelengths used for telecommunication we can see that the ITU-T has defined the following bands as shown in table 1

Band	Descriptor	Range [nm]
O-band	Original	1260 to 1360
E-band	Extended	1360 to 1460
S-band	Short wavelength	1460 to 1530
C-band	Conventional	1530 to 1565
L-band	Long wavelength	1565 to 1625
U-band	Ultra-long wavelength	1625 to 1675

Table 1: ITU-T Bands[2, p. 134]

The ITU-T further describes the U-band as:

In some cases it is desirable to perform a number of maintenance functions (preventive, after installation, before service and post-fault) on fibre cables in the outside plant. These involve surveillance, testing, and control activities utilizing optical time domain reflectometer (OTDR) testing, fibre identification, loss testing, and power monitoring. A wavelength region, that is intended to be never occupied by transmission channels, may be attractive for maintenance, even if enhanced loss occurs. The U-band has been defined exclusively for possible maintenance purposes.

The ITU-T's intended use of the U-band corresponds with our own objectives, it is therefore obvious to us that we will further concentrate on this band.

2.1 Choosing a wavelength

The U-Band reaches from 1625nm to 1675nm. When choosing a wavelength we need to make sure that we have the least amount of energy loss or attenuation when transmitting light. The amount of attenuation depends both on the fibre type and wavelength which is used.

The main reason of our research is to look at potential uses of the unused channels at the edge of the L-band in a WDM system. Because of this we can safely assume that the fibre optical cable used is a single mode fibre which supports C+L band WDM systems. If the link was not used in a WDM set up, then there would be no reason to use the unconventional U-band.

The ITU-T has several standards ranging from G.652 to G.657 which specify single mode optical fibres[2, chap. 6-7], of which none have official support for the U-band. The U-band channel we choose should for this reason be as close to the L-band as possible to keep the wavelength closer to the cable's specifications.

Figure 1 on page 7 shows attenuation as a function of the wavelength of an unbent fibre. We can see that in the case of the U-band there is a gradual upward slope as the wavelength increases.

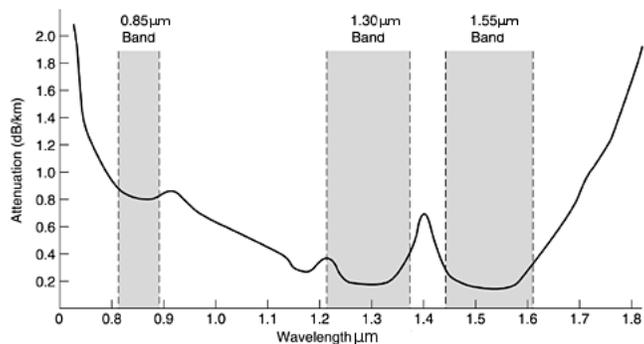


Figure 1: Attenuation of light through fibre in the infra-red region.[9, fig 2-6]

Figure 2 on page 7 shows the effect of fibre bends on attenuation. Here we can see again that it is better to choose a wavelength as close to the L-band as possible.

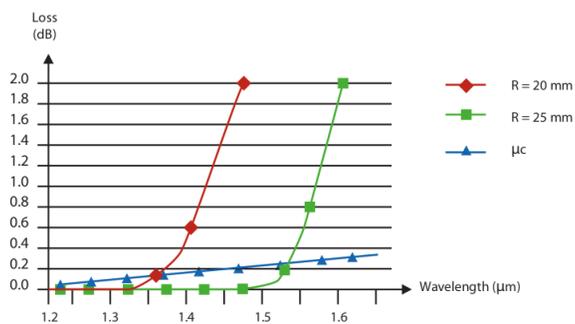


Figure 2: Effect of fibre bends on attenuation.[4, p. 27]

Because the U-band begins with 1625nm we think that this would be the best wavelength to use.

3 Proof of Concept

To test our hypotheses we were initially provided with two Inteno XG6746[3] switches and four 1625nm SFP optics by BeetleFiberOptics. It is in the interest of BeetleFiberOptics that we use this relatively cheap equipment instead of specialised, and usually more expensive fibre monitoring equipment. The Inteno XG6746 switches have one modular SFP fibre optic port in which we can place the optics to be tested.

We would like the fibre optic switches to communicate with each other so they can exchange statistics about their fibre optics and use these statistics to create MRTG[5] graphs that could show interesting trends which might be an indicator of fibre performance over time.

3.1 Inteno XG6746

The Inteno XG6746 runs on the software image `XG6746_4.02ITT24.01_20121022`. Because of the closed nature of the software we were not able to directly install programs or scripts which we needed to generate traffic on the fibre link directly. Before the start of our research project we had been informed that customizable software would be made available for these switches so we could write our own testing modules. This, however, never happened and as such we had to come up with another solution.

We tried to install several open-source operating systems on the switches, unfortunately none of them would work. To still be able to run diagnostic tests we decided to attach a Raspberry Pi to the switch via Ethernet which can both generate traffic and do periodic polling via SNMP. Ideally all the functions performed by the Raspberry Pi would be integrated into the Inteno XG6746, however to simplify matters we chose to use a Raspberry Pi as it boasts a full-fledged Linux environment, Raspbian[8].

The Inteno XG6746 allows us to read out the following information about the optic:

- Vendor Name : EOPTOLINK INC
- Vendor OUI : 00 00 00
- Vendor PN : EOLS-1603-29SD
- Vendor rev : 1.0
- Temperature : 62.559 C
- Voltage : 3.19 Volts
- Bias : 8 mA
- Tx Power : -5.65 dBm *probably wrong*
- Rx Power : -0.03 dBm *probably wrong*

As seen above we can read out the optics's Rx and Tx, however, we do not know in what format these values are presented. An example of the readout is shown in table 2 on page 9 in which we show the represented values of two connected Inteno XG6746 switches using 1550nm optics with a 25KM fibre and where we compare the Rx value with that of a separate Optical Power Meter, a Fujikura FID-25R.

Switch	Tx	Rx	OPM measured Rx
Sw1	-4.61 dBm	-0.44 dBm	-9.8 dBm
Sw2	-4.86 dBm	-0.83 dBm	-12.4 dBm

Table 2: Tx and Rx values of two connected Inteno XG6746 switches

The web-based GUI shows us that the values should be in dBm, but this makes no sense as the Tx is always lower than the Rx, this would indicate an increase of power in transit. This effect also increases when we introduce an attenuator on the link. Even if we would assume that this is the result of a simple mix-up of Rx and Tx then these values still make no sense because we get totally different values when measuring the actual output with the optical power meter. Also when we assume that these values are in fact measured in milliwatts then these values still make no sense as the values are negative. When the optics are placed in a different type of switch then we do get expected values which would indicate that this problem does not lie in the optics, but rather in the switch itself. We expect this to be a software bug which can be resolved in the future. For this reason we have decided to still use the Inteno XG6746 for testing out-of-band management. For correct fibre attenuation monitoring we had to use another type of switch, a Zyxel XGS4728F.

Although the Zyxel XGS4728F does show us the correct Tx and Rx values for the 1625nm optics, it does not allow an Ethernet link to actually be brought up. The reason for this is that the optics are only 100-mbit which the Zyxel XGS4728F does not support.

We can calculate how much attenuation is present on the fibre between two switches by calculating the difference between the optical power sent from switch1 and the optical power received at switch2: $Att = Tx(sw1) - Rx(sw2)$. We think it would also be interesting to keep track of the temperature of the Tx optic because this might influence the precise Tx wavelength, and thus attenuation.

Our particular hardware situation requires us to build two different networks. Our primary test network will prove that an Ethernet network using 1625nm optics is possible, and our Attenuation network will prove that we can monitor the attenuation of a fibre link over time.

3.2 Primary test network

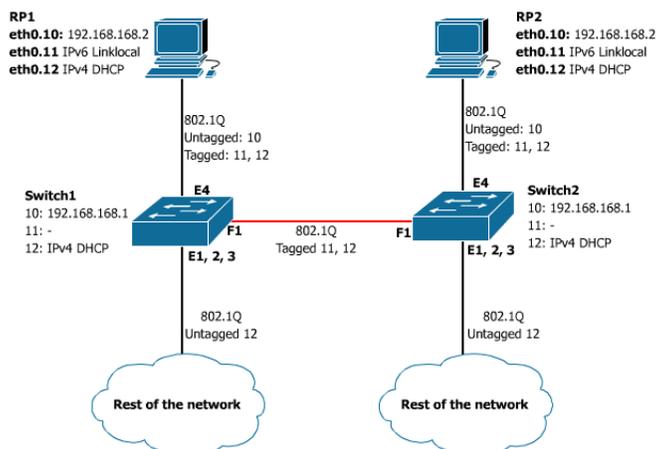


Figure 3: Network diagram of the out-of-band management network

Figure 3 on page 10 shows our primary test network which uses the Inteno XG6746 switches. A secondary goal of this network was to make it as 'plug-and-play' as possible, meaning that an engineer shouldn't be required to do much configuring. This network should prove that we can use the 1625nm wavelength to create a Ethernet link which can be used alongside the regular wavelengths. The RP devices will automatically use SNMP to poll Rx and Tx data from their nearest switch, and send the Rx data to the other RP device for comparison. It is important to note that we use 3 separate VLANs for this.

The purpose of VLAN-10 is to create a single unit out of 1 Raspberry Pi and 1 monitored switch. This way RP1 will only monitor Switch1 although Switch2 uses the same default IPv4 address as Switch1. In the ideal situation where the RP# and Switch# would be combined into 1 device this VLAN would be obsolete.

VLAN-11 is used to exchange data between 2 RP# devices. This VLAN creates a group of 2 VLAN-10 units using the fiber optics in port F1 on Switch#. The RP# devices use an IPv6 Link local address, this is done so there is no need to manually configure an IP address when 2 VLAN-10 units are connected to each other. The 2 RPs can discover each other by sending a ping to IPv6 multicast address FF02::1 which all IPv6 linklocal hosts should listen and reply to[7]. After the RP knows the IPv6 address of a potential neighbour it will do a simple check to see if that neighbour is indeed our polling partner and not some rouge computer. It will do this by sending a SHA256 hash of a common pre-shared password plus its own linklocal IPv6 address which should protect against a simple re-play attack.

VLAN-12 is used to connect our devices to the rest of the network and thus might be used for out-of-band management.

3.3 Attenuation test network

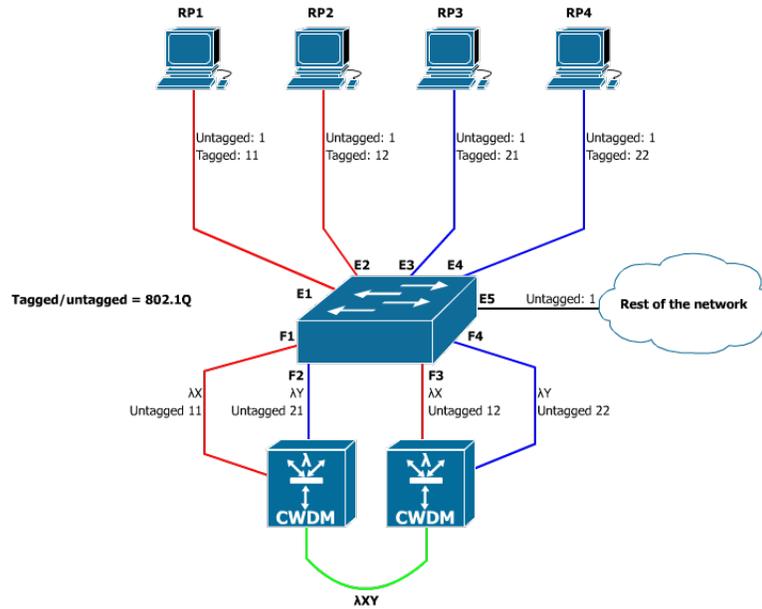


Figure 4: Network diagram of the attenuation test network

Because of the dBm representation problem of the Inteno XG6746 switches we had to do create a alternate network using another switch type which does display the Tx and Rx values correctly. The Zyxel XGS4728F does display the correct values, but there was only one available to us, therefore a new network set-up was needed. The Attenuation test network diagram is shown in figure 4 on page 11.

The Zyxel XGS4728F switch does have one major flaw, it does not support 100-mbit speed optics like our 1625nm ones. Luckily it does still give the correct Tx and Rx values for these optics when they are connected to each other with a fibre.

We can both test our dBm exchange script and poll from the 1625nm optics because we can let the RPs poll not only from the optics in their own VLAN, but also from the 1625nm optics.

4 Tests

For our tests we shall use the network set-ups as described in section 3.2 and section 3.3.

We have performed the following tests:

- Sending a steady stream of UDP packets across the fibre at 10Mb/s.
 - Using the primary network described in section 3.2.
 - This test will show if we can set up a reliable Ethernet link over 1625nm.
- Measuring the difference between Tx and Rx values at either end of the fibre link.
 - Using the attenuation network described in section 3.3.
 - To show that the Zyxel switch's Tx and Rx readout can be trusted we will also use the same optics in a separate Bit Error Rate Tester, BERT, which also shows the optic's Tx and Rx values.
 - We will compare the following wavelengths:
 - * 1550 nm
 - * 1610 nm
 - * 1625 nm
 - We will compare over the following fibre lengths:
 - * 1 m
 - * 25 Km
 - This test will show us how much more attenuation the 1625nm wavelength has in practise in comparison with the theory.

For the attenuation test it is important to bear in mind that we are getting our Tx and Rx values from the digital diagnostic information provided by the optic. This information has an error margin which can be vendor specific, but may be at most 3dBm, as stated in the SFF-8472 Specification for Diagnostic Monitoring Interface for Optical Transceivers[1]

In theory these optics should return the same values regardless of which equipment they are plugged into, although The Inteno switches didn't display the correct values as shown in table 2 on page 9. The optics return their values through the standardised digital diagnostics monitoring interface using i²c, this standard defines which values are measured and how the information is stored.[1]

4.1 Ethernet link reliability test

After establishing that the 1625nm optics seemed to communicate properly, we set up an Ethernet link using two Inteno switches and 1 KM G652.D fibre cables. We then used this link to send a steady flow of UDP data at a speed of 10Mbit/s for 15 hours using the programme iperf.

- RP1: `iperf -suV`
- RP2: `iperf -c fe80::ba27:ebff:fe58:69e2%eth0.11 -V -b 10000000 -t 54000`

We had the following result on RP2:

```
-----  
Client connecting to fe80::ba27:ebff:fe58:69e2%eth0.11, UDP port 5001  
Sending 1470 byte datagrams  
UDP buffer size: 160 KByte (default)  
-----  
[ 3] local fe80::ba27:ebff:fe95:ba76 port 38146 connected with  
fe80::ba27:ebff:fe58:69e2 port 5001  
[ ID] Interval Transfer Bandwidth  
[ 3] 0.0-54000.0 sec 62.9 GBytes 10.0 Mbits/sec  
[ 3] Sent 45918349 datagrams  
[ 3] Server Report:  
[ 3] 0.0-54000.0 sec 62.9 GBytes 10.0 Mbits/sec 0.047 ms 0/45918348  
(0%)  
[ 3] 0.0-54000.0 sec 1 datagrams received out-of-order
```

The reason for using a relatively slow speed of 10Mbit/s is that the optics are only rated at 100Mbit/s. If we were to run tests with higher speeds, there is a greater chance of the packet loss being caused by an overly stressed optic, rather than the behaviour of the 1625nm wavelength along the fibre. We let the stream run all through the night, for 15 hours, and checked the results the next morning. Over the course of the 15 hour test period, no packets were lost. This seems to indicate that the 1625nm wavelength is capable of sustaining stable Ethernet links.

To test under which conditions we would start losing packets, we started bending the fibre cable. No loss occurred until we started bending the fibre cable excessively. Only when very sharp bends, which could permanently damage the fibre cable, were made, did we see packet loss along the link. We also performed the same bending test using the BERT devices. The BERT also showed the BER increasing once sharp bends were made.

4.2 Attenuation test

We started off by using Bit Error Rate Testers (BERT) to read the attenuation of the light signals across different wavelengths and cable lengths. The following wavelengths were tested:

- 1550 nm
- 1610 nm
- 1625 nm

For each of these wavelengths we tested four different optics, to make sure we had somewhat reliable results to accurately portray the attenuation of each wavelength. It would have been preferable to use more than just four optics per wavelength, but during this research we did not have access to additional optics. All of the used optics were also designed for use over a range of 40 kilometres. Each of these optics was then run across a one metre and a 25 kilometre G657.A fibre cable. These tests were then also performed using a different Bit Error Rate Tester, This results in the following tables, showing the average attenuation for each wavelength and fibre cable length:

Wavelength	1 M	25KM
1550nm	-0.05 dB	8.36 dB
1610nm	0.09 dB	7.98 dB
1625nm	2.28 dB	14.33 dB

Table 3: Wavelength attenuation as measured with BERT

Wavelength	1 M	25KM
1550nm	0.44 dB	9.23 dB
1610nm	0.62 dB	9.28 dB
1625nm	2.25 dB	14.96 dB

Table 4: Wavelength attenuation as measured with BERT & OPM

As you can see, the 3 dB error margin for the digital diagnostics monitoring distorts the results of the BERT measurements using the one metre cable quite significantly. For these tests the attenuation measured with the OPM seems more reliable, however only the Rx is measured by the OPM. The Tx is still provided by the DDM interface and may have a larger deviation. We must also mention that one of the four tested 1625nm optics suffered from an additional three to four dB attenuation on the transmitting side across any cable length, this will have raised the average attenuation for the 1625nm in table 4.2 and 4.2.

Regardless of this deviation, it is clear that the 1625nm wavelength suffers from considerably more attenuation than the other tested wavelengths. The 1625nm optics are designed to handle at most 29dB difference between Tx and Rx power, so this test across 25 KM still falls well within the defined range of the optics. A downside to the 1625nm wavelength suffering from such increased attenuation is that more power will be needed to cover larger distances, compared to regular C- and L-band wavelengths.

Besides measuring these attenuation values using the BERT, we also set up two Raspberry Pis to monitor the attenuation along the 1625nm and 1610nm links. The following graphs show changes in attenuation over the course of a day:

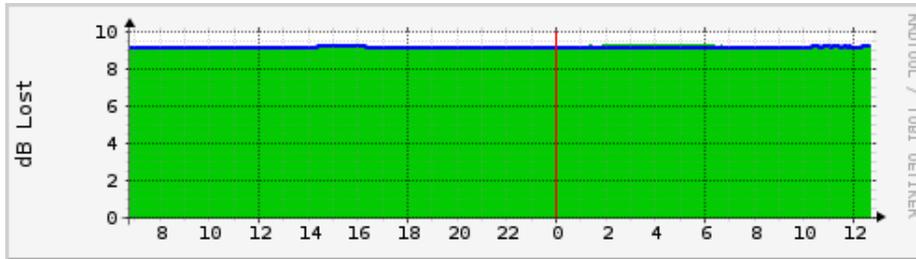


Figure 5: Attenuation for both 1610nm fibre links

As you can see in figure 5 the regular 1610nm attenuation is reasonably stable, only deviating slightly over the course of a day.

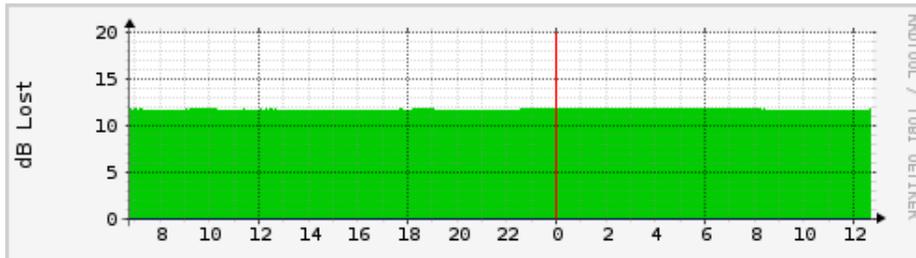


Figure 6: Attenuation for one side of the 1625nm link

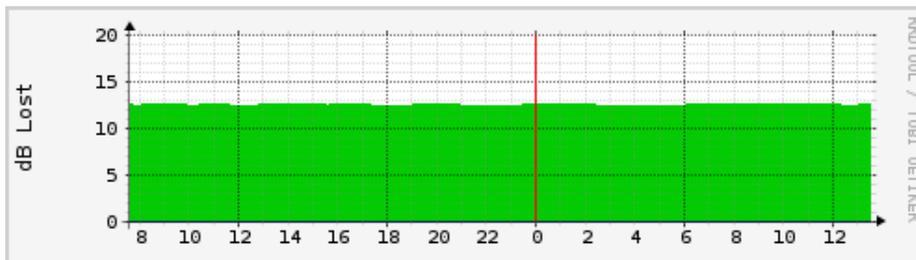


Figure 7: Attenuation for the other side of the 1625nm link

The same goes for the 1625nm links, they fluctuate slightly over the course of the day but there are no significant changes. These results are as expected,

as the fibre cables have not been moved or bent during this monitoring period. We also measured the temperatures of the optics, to see if there was an obvious correlation between the temperature and the amount of attenuation. Every

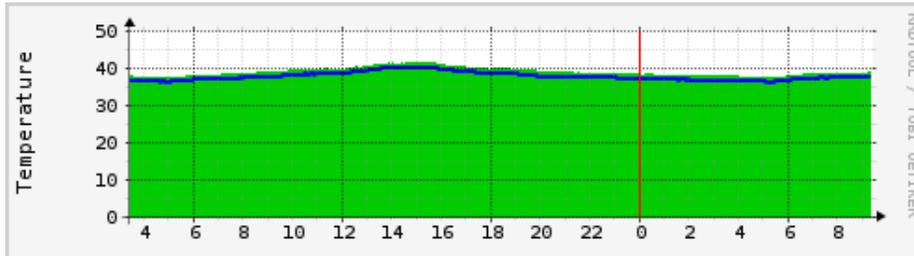


Figure 8: Temperatures for both 1625nm optics

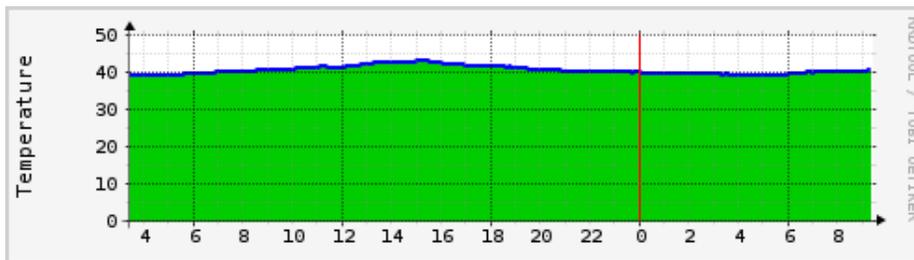


Figure 9: Temperatures for both 1610nm optics

night when the lab is empty the temperatures of the optics drop by roughly two or three degrees Celsius, as shown in figures 9 and 8. This could cause the emitted light wavelength to drift slightly, which could in turn affect the attenuation over a fibre cable.

As it turned out, the statistics did not prove a correlation between the temperature and the attenuation of the signal in our test set-up. Perhaps if the temperature of the optics changed more significantly it could affect the attenuation, but with small temperature changes the attenuation does not seem affected.

4.3 OTDR tests

To test the reliability of our previous measurements using the optics Digital Diagnostics statistics, we measured our 25 KM fibre cable including the CWDM on either end using the 1550nm band on an EXFO FTB-200 OTDR. The result of the measurement showed 7.009 dB of attenuation across the entire fibre path. When this result is compared to our measurements from table 4.2 it shows our measurements are reasonably close to the measurements using the OTDR. We can therefore assume our measurements can be used as fairly reliable results.

5 Conclusion

Up until now the 1625nm band has been used for monitoring, however so far it has been mainly used for out-of-band OTDR testing. Especially in DWDM systems the 1625nm band is used for these OTDR tests to tell if the monitored fibre is still in good condition. What has surprised us somewhat is that 1625nm optics are not readily available, meaning the wavelength is not regularly used for data communications using Ethernet. The tests we have performed during our research have shown the 1625nm band to be a lot more sensitive to attenuation through bending, temperature differences and distance. Regardless of these drawbacks, the 1625nm wavelength has proven to allow for the creation of stable Ethernet links over 25 kilometres. From a monitoring perspective, it may even be an added feature that the 1625nm wavelength is so sensitive to bending, as it will show signs of unwanted manipulation of the fibre sooner than the regular wavelength connections would.

References

- [1] SFF Committee. Sff-8472 specification for diagnostic monitoring interface for optical transceivers. <ftp://ftp.seagate.com/sff/SFF-8472.PDF>, 2012.
- [2] Yoichi Maeda et al. *ITU-T manual 2009, Optical cables, fibers and systems*. ITU-T, 2009.
- [3] Inteno. Xg6746, 2013.
- [4] J. laferriere, G. Lietaert, R. Taws, and S. Wolszczak. *Reference Guide to Fiber Optic Testing, Second Edition, Volume 1*. JDS Uniphase Corporation, 2011.
- [5] Tobi Oetiker. Mrtg, 2013.
- [6] University of Amsterdam. Master programme: System and network engineering. <http://www.os3.nl>.
- [7] S. Deering R. Hinden. Rfc4291: Ip version 6 addressing architecture, 2006.
- [8] Raspbian. Raspbian, 2013.
- [9] Andrew S. Tannenbaum. Computer networks, fourth edition. 2002.