

Go Long!

Sending weird signals long distances over existing optical infrastructure

Ben Cartwright-Cox
2024 / 38c3


Hi

- I am Ben Cartwright-Cox (@benjojo@benjojo.co.uk)
- You might remember me from things like:
 - bgp.tools
 - "A deep dive into the world of DOS viruses" - 35C3
 - BGP Battleships
 - Some other blog bits like making {USB drivers, Class E, HDMI Magic boxes}


Hi

- I am Ben Cartwright-Cox (@benjojo@benjojo.co.uk)
- You might remember me from things like:
 - bgp.tools
 - "A deep dive into the world of DOS viruses" - 35C3
 - BGP Battleships
 - Some other blog bits like making {USB drivers, Class E, HDMI Magic boxes}
- I like knowing about how stuff works
 - It makes life more interesting, and I don't like taking any tech for granted
 - My rule of thumb; it's either really smart, or incredibly stupid
 - I'd like to know what one it is before I depend on it

The modern internet backbone is nearly all optical



Networking optical kit is a weird world




Device

1000baseKX
10000baseKR
25000baseCR
Etc etc



Optic Type

SFP
SFP+
XFP
QSFP
QSFP+
OSFP



Transmission / Wavelength

LR
LX
SR
SX
ER
BX

850nm
1310nm
1510nm
CWDM
DWDM

Networking optical kit is a weird world (2)

9 μm




SMF/OS2

Single Mode


50 - 62 μm




OM1



OM2



OM3



OM4

Multi Mode

Average "boring" ISP network

- Uses a SFP+ Port on a switch
 - With a vendor SFP+ optic
 - Or a 3rd party supplier, a lot of the optics are ultimately the same factories anyway!
 - Just juicy vendor markup support costs included
 - The switch chip spits out a XGMII signal towards the optic
 - We will ignore what happens if you need to turn this into "BASE-T"
 - The optic then does some stuff to turn that XGMII signal into laser light
 - The opposite happens on the other end
 - Data bits are exchanged
-
- Cat videos are exchanged or something

Some optics are surprisingly stupid

This thing is NOT
thinking*



We will ignore weird things that have full CPUs in them



VDSL Modem in SFP



ARM Machine inside SFP [1]



GPON ONT inside SFP

[1] <https://blog.benjojo.co.uk/post/smart-sfp-linux-inside>

Some optics are surprisingly stupid

This thing is NOT
thinking*



Some optics are surprisingly stupid

This thing is NOT
thinking*

* There is a EEPROM to tell the switch what it is (and have the switch complain to you to buy vendor optics instead)



Some optics are surprisingly stupid

This thing is NOT thinking*

* **There is also** sometimes a tiny microcontroller to read out laser light level metrics and temperature for the switch




Average faster "boring" ISP network

- Uses a QSFP28 Port on a switch
 - With a vendor QSFP28 optic
 - Depending on the optic (100G is a mess) various fancy chips are used to convert [CK]R4 into something desirable, maybe even coherent for special use cases
 - The opposite happens on the other end
 - Data bits are exchanged
-
- Cat videos are exchanged *at higher speeds* or something

Some optics are surprisingly smart

- The fancy ones have to turn 4 "lanes" of 25GBaud data into one "normal" lane of "data", at 100GBaud speeds
- This (sometimes) requires advanced chips to exist inside the optics
 - Like, your intel laptop likely has a bigger process node than these optics
 - Most optics like this are 7nm/5nm or less
- These optics also consume considerably more power to run both the laser and the chip inside the optic to combine the feeds
- Depending on the type of optic (PLR4/ER/ZR/ZR+/SR/LR4/LR1/PSM/CWDM4) the complexity (and price) of the chips are better/worse



Other optic form factors are available

- XFP / CFP / Various QSFP forks / OSFP / X2 / Xenpack
 - All used in various places and ages of kit, but they are similar enough not to make a difference here

Other optic form factors are available

- XFP / CFP / Various QSFP forks / OSFP / X2 / Xenpack
 - All used in various places and ages of kit, but they are similar enough not to make a difference here
- Home Theater stuff also uses some kind of optical transmission
 - SPDIF / TOSLINK



TOSLINK

- Two kinds of people use TOSLINK/SPDIF
 1. People who have ground loop issues (speaker mains hum)
 2. People who just want to seem cooler by using optical stuff
- Maxes out at ~10mbits/s, But really hangs around 3mbit/s for most setups
- The wavelength is more 700nm (aka red LEDs), and the cables have a super thick (1mm) plastic core

- Quoted max length for a cable is 10m

TOSLINK

- Two kinds of people use TOSLINK/SPDIF
 1. People who have ground loop issues (speaker mains hum)
 2. People who just want to seem cooler by using optical stuff
- Maxes out at ~10mbits/s, But really hangs around 3mbit/s for most setups
- The wavelength is more 700nm (aka red LEDs), and the cables have a super thick (1mm) plastic core

- **Quoted max length for a cable is 10m (due to 1000 μ m core!)**

Networking optical kit is a weird world (2)



SMF/OS2


Single Mode
(9 μ m)




OM1



OM2



OM3




OM4

Multi Mode
(50 μ m~)


Networking optical kit is a weird world (3)





Networking optical kit is a weird world (3)



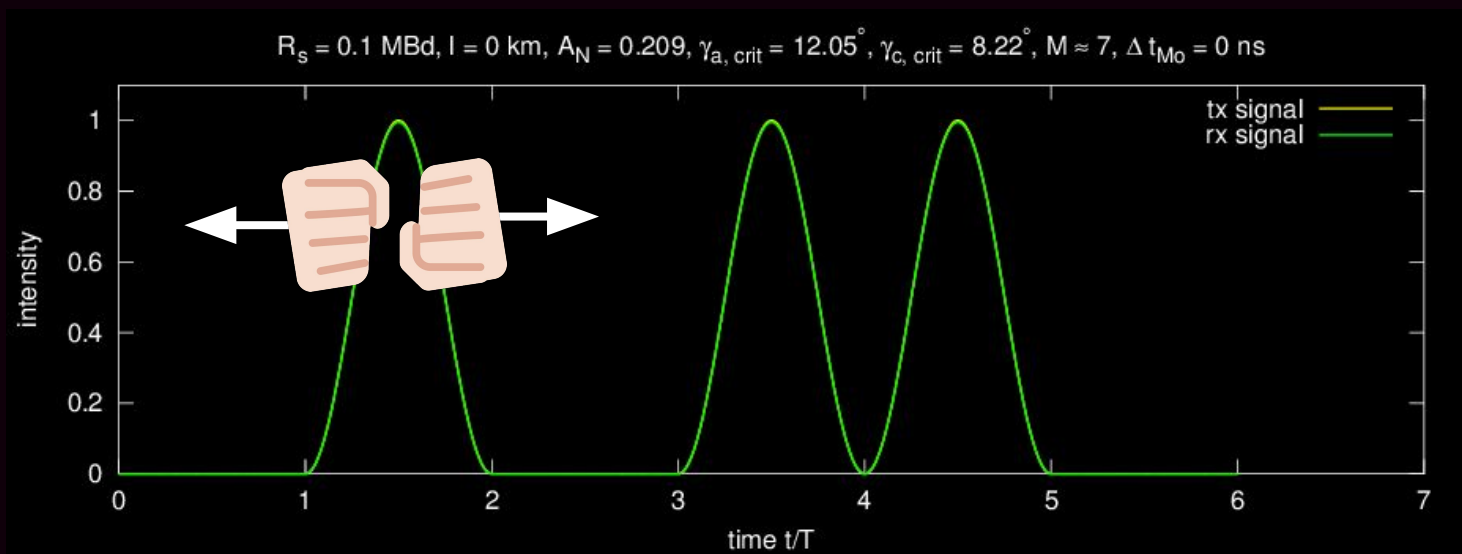

Networking optical kit is a weird world (3)





Optics are a weird world (3)




Optics are a weird world (3)




Optics are a weird world (3)




Networking optical kit is a weird world (3)



Networking optical kit is a weird world (3)



Networking optical kit is a weird world (3)




So I was thinking

So I was thinking



So I was thinking


Too stupid to care



So I was thinking




So I was thinking



10km TOSLINK?

We are going to need to do some things

Small Form-factor Pluggable (SFP) Transceiver MultiSource Agreement (MSA)



20	VeeT
19	TD-
18	TD+
17	VeeT
16	VccT
15	VccR
14	VeeR
13	RD+
12	RD-
11	VeeR

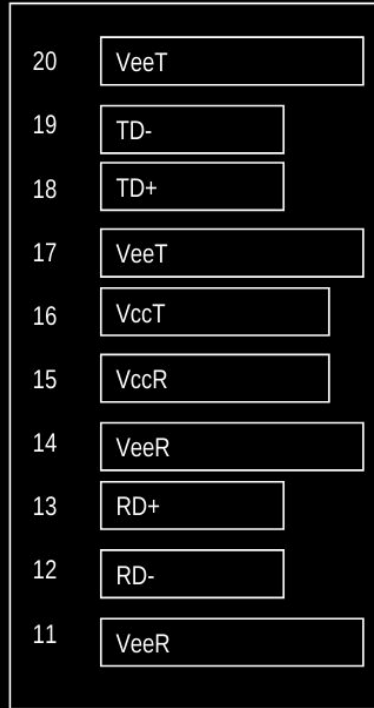

Top of Board

1	VeeT
2	TxFault
3	Tx Disable
4	MOD-DEF(2)
5	MOD-DEF(1)
6	MOD-DEF(0)
7	Rate Select
8	LOS
9	VeeR
10	VeeR

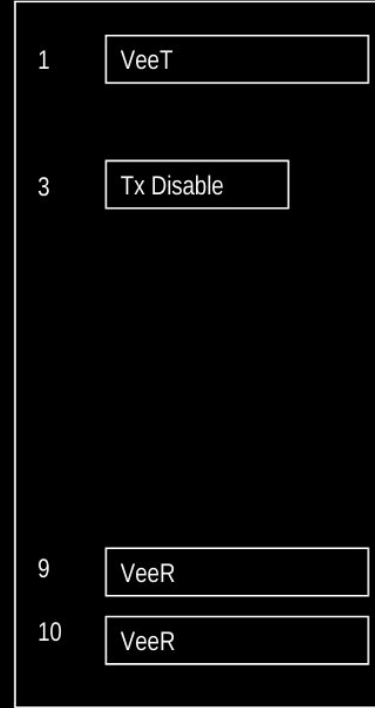
Bottom of Board (as viewed thru top of board)

We are going to need to do some things

Small Form-factor Pluggable (SFP) Transceiver MultiSource Agreement (MSA)



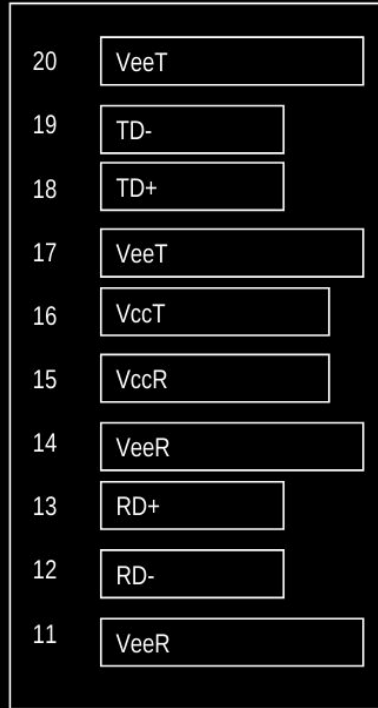
Top of Board



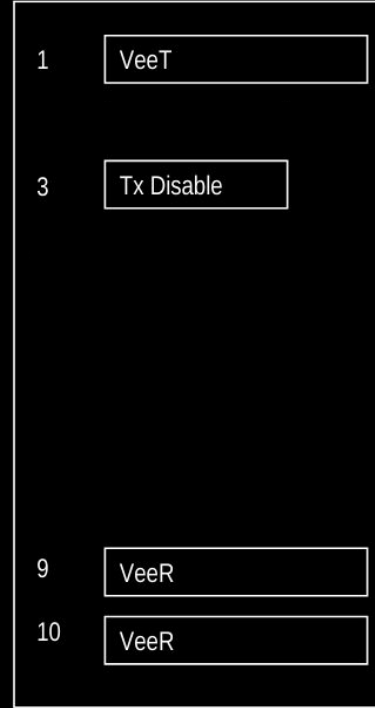
Bottom of Board (as viewed thru top of board)

We are going to need to do some things

Small Form-factor Pluggable (SFP) Transceiver MultiSource Agreement (MSA)



Top of Board



Bottom of Board (as viewed thru top of board)

VeeT: TX Ground
VeeR: RX Ground
VccT: Tx Power
VccR: Rx Power

RD- : RX Signal
RD+ : RX Signal
TD-: TX Signal
TD+: TX Signal

Cheating a bit

A friend[1] made some Osmocom SFP Experimenter boards, They come with:


- A 5v to 3.3v power supply
- SFP Connector + Cage
- Wired up extra pins for various aux signals we want
- LEDs to tell you when you have upset the optic
- Exposed pads to inject your signal
- A pair of line drivers to make the signal friendly to optics



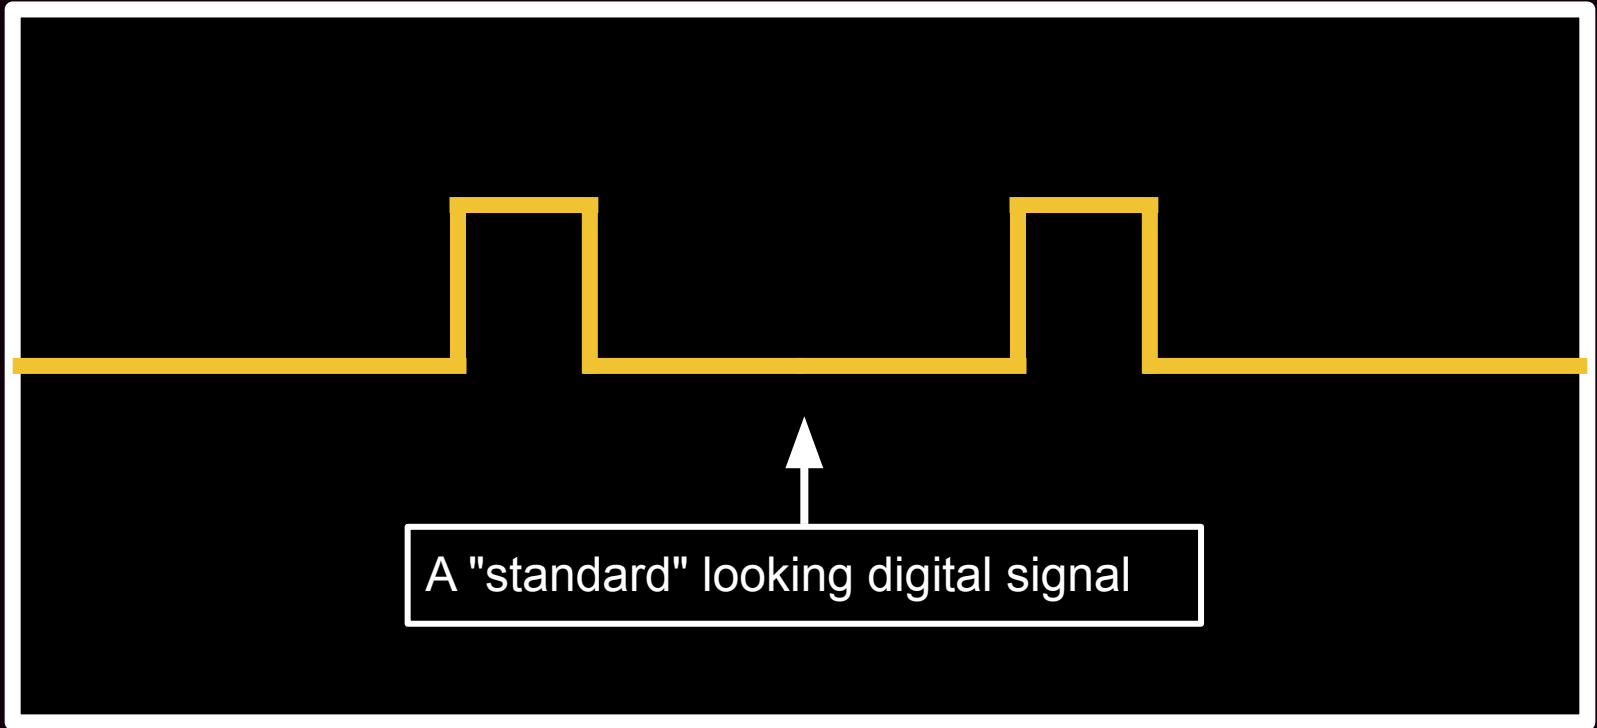
Cheating a bit

A friend[1] made some Osmocom SFP Experimenter boards, They come with:



- A 5v to 3.3v power supply
- SFP Connector + Cage
- Wired up extra pins for various aux signals we want
- LEDs to tell you when you have upset the optic
- Exposed pads to inject your signal
- **A pair of line drivers to make the signal friendly to optics**




Line drivers make things wobble just right




Line drivers make things wobble just right



Line drivers make things wobble just right




Audio In



+0.0V
DC BW 10.0:1 DC BW 10.0:1
-687.5mV

12:09 AM
Jun 10, 2019


Audio In



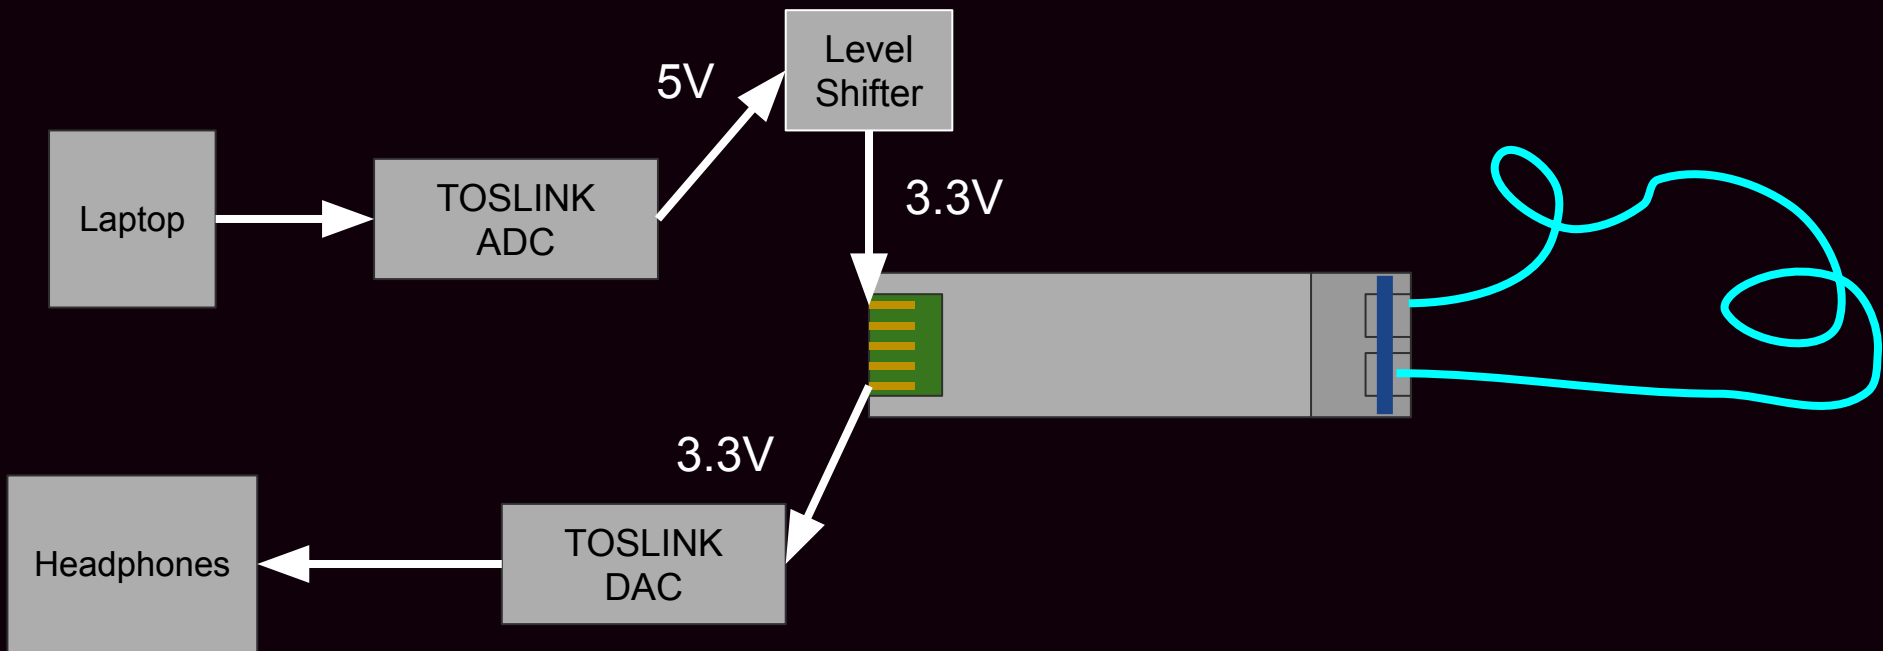
SPDIF Out



About 2Mhz~ of data





POC Setup



Working POC

- Not airport friendly





Latency

Adds about 34ns~

34 nanoseconds /
200000 km per second
is about 680cm.

Not bad! Given the
patch lead is at least
500 cm

But what if we could go long haul?




IP
H

About 650
meters~

Telehouse


CWDM Channel from Jump Networks ~1km ish

- Thanks for James Rice for being trusting enough to be the 1st test case!
- IP-House <-> Telehouse North (London)
- Optical loopback on the telehouse side




CWDM Channel from Jump Networks ~1km ish

- Thanks for James Rice for being trusting enough to be the 1st test case!
- IP-House <-> Telehouse North (London)
- Optical loopback on the telehouse side




It worked!

- Portable speaker was fighting against the data hall noise!
- Did not destroy the CWDM optic!



Latency


Harder this time
on the scope
because $11\mu\text{s}$
requires you to
zoom out a lot to
find the offset



$11\mu\text{s} = 2.2\text{km}$

- Double what a very rough google maps estimate gave
- Building cabling on both sides likely add up a lot
- This is more like the distance if added a lap of both DC buildings (in 2D)

11.1 microseconds at 200000km/s

 NATURAL LANGUAGE  MATH INPUT  EXTENDED KEYBOARD  EXACT RESULTS

Input interpretation

$11.1\ \mu\text{s}$ (microseconds) \times 200000 km/s (kilometers per second)

Result

2.22 km (kilometers)

Unit conversions

1.379 miles

7283 feet

2220 meters

222 000 cm (centimeters)




IP
H

About 650
meters~

Telehouse



Telehouse London
Data Centres



Rosemary Dr

Oregano Dr

Tr
D

F

But what if we could go longer haul?

Longer haul options

- Typically, if you were doing 1G/10G you would be doing DWDM at this point
- DWDM is like CWDM, but uh, **D**enser!
- 8 channels to 32/64~ channels.
- Optics are more expensive because they have cooler/hotter physics inside to adjust their wavelength and are more precise in what wavelength they emit

Longer haul options

- Typically, if you were doing 1G/10G you would be doing DWDM at this point
- DWDM is like CWDM, but uh, Denser!
- 8 channels to 32/64~ channels.
- Optics are more expensive because they have cooler/hotter physics inside to adjust their wavelength and are more precise in what wavelength they emit
- DWDM also amplifies well with fiber amplifiers
 - Side note, The physics involved with EDFA amplifiers are bizarre and feel like a bug in physics

Other Longer Haul options

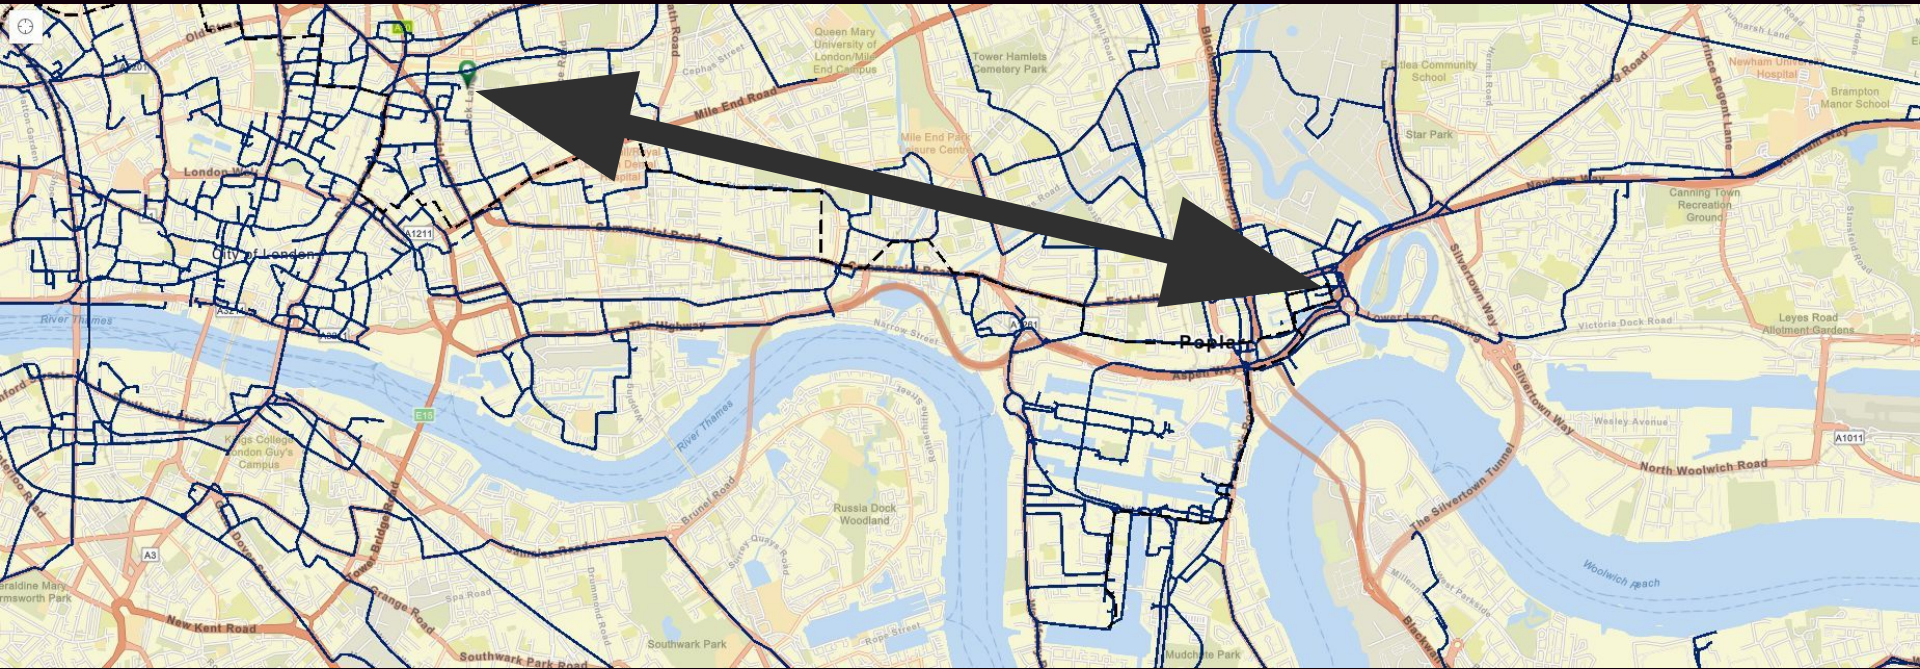
- Most times, you don't want your customers to use your multiplexer directly
 - You probably don't trust them to have the right transmission color/power and risking your other channels if they are too bright
- This is where you use a *transponder*, a transponder is something that basically copies and pastes a signal from one SFP to another
- The idea is that you have your customer connect at 1310nm LR/LX/etc, and you install the magic DWDM/Weird optic on your side!

- My weird signals work fine in transponders, as they are simple devices

Incompatible long haul options

- There are "muxponders" and other fancy "line systems" that do a lot more processing with the input signal
- These things achieve amazing density, by mostly understanding the signals coming in (Ethernet/FiberChannel/Infiniband/etc) and merging into a more denser high speed data streams.
- A Muxsponders understands ethernet etc directly, it "sees" your ethernet frames and packs them into time slots
- For example, a muxponder can have a number of 10G ethernet inputs, and output a single 200Gbit/s (or faster) "beam"
- **You cannot send TOSLINK over these, sadly (for me) these are common with a lot of carriers (and their use is growing)**

DWDM TOSLINK




Physical distance is about ~6km from Interxion LON1 <-> Telehouse North
(Thank you to VeloxServ for letting me borrow a channel)

DWDM works!

- Round Trip $\sim 84\mu\text{s}$
- About 16.8KM of glass tube then
- With 10 db of optical attenuation to avoid blowing the optic up

(Thank you to Brandon Butterworth for lending me the optics)



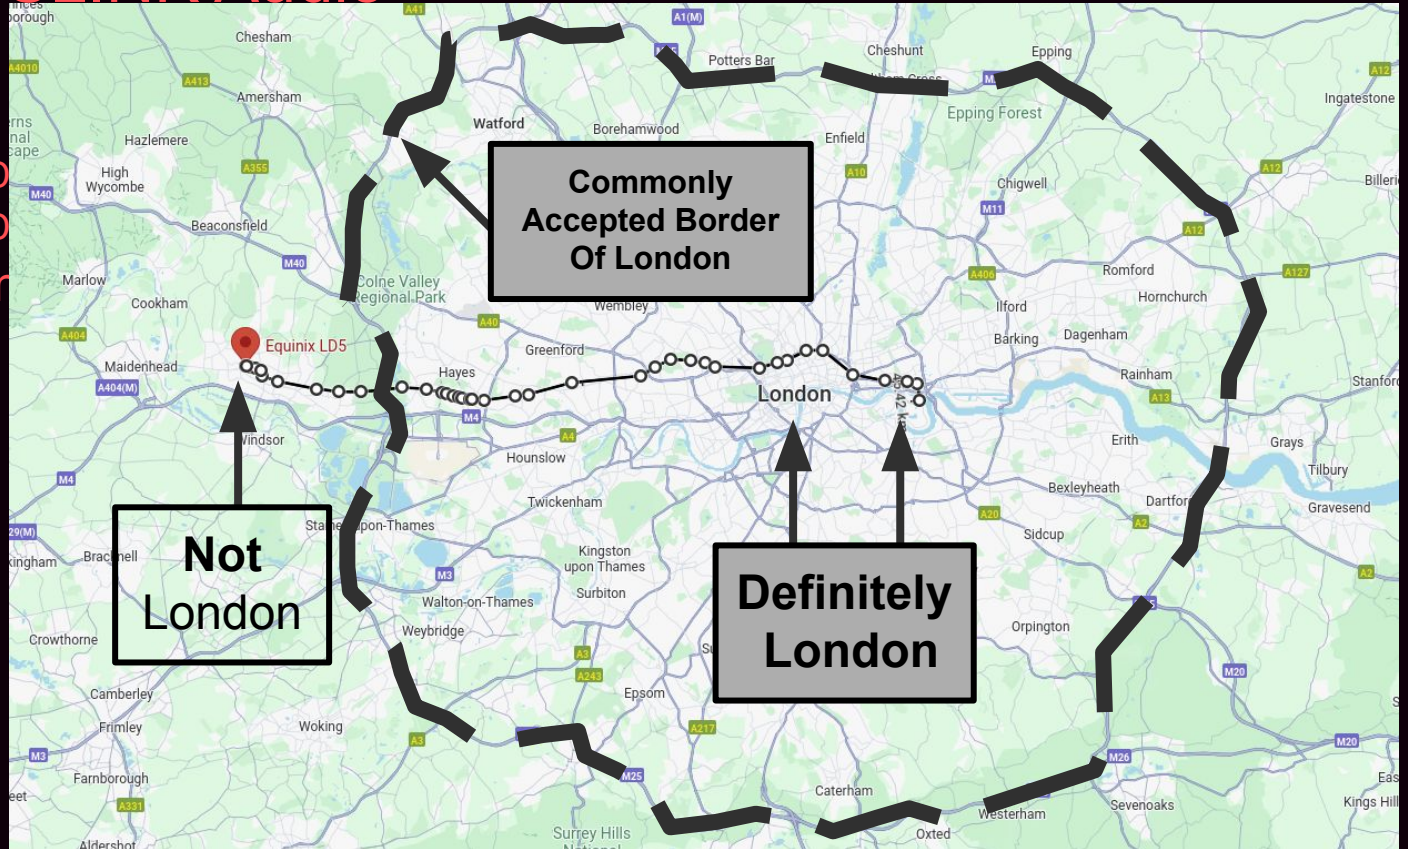
But what if we could go longer haul?


Inter-city TOSLINK Audio

- LONAP (❤️) had a spare DWDM between:
 - Equinix "LD"6 (Actually Slough, not London)
 - Equinix LD8 (London Docklands)
- This is not only intercity, but it's so far that it has to be amplified!
- Estimated to be 73 kilometers one way

Inter-city TOSLINK Audio

- LONAP (❤️)
 - Equinix LD
 - Equinix LD
- This is not or
- Estimated to





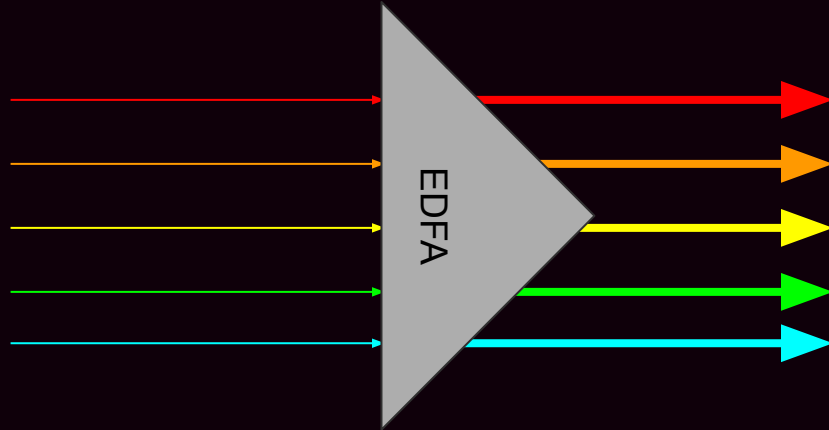
Inter-city TOSLINK Audio

- Used a spare DWDM SFP+, Plugged into a SmartOptics DCP-M40-C-ZR+
- Amplified, and DWDM Mux'd, a loop (with attenuator) was installed on one side so I didn't have to be in Slough for any longer than needed




Why put attenuator on the loop?

- Erbium-Doped Fiber Amplifiers (EDFA) are magic, but they do have their limits



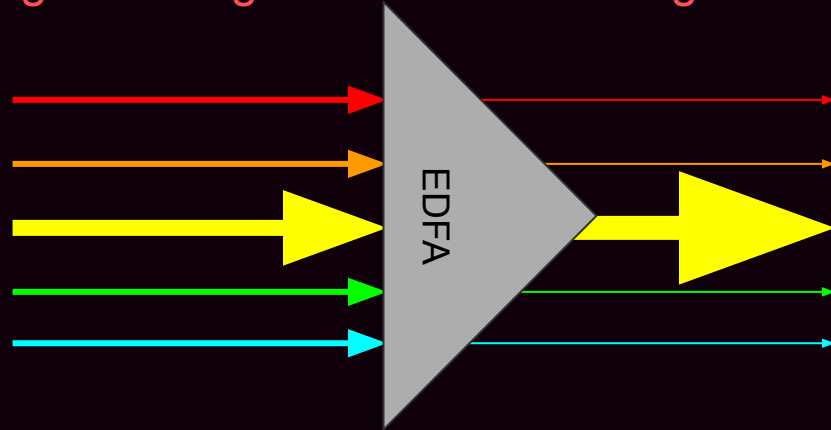
Why put attenuator on the loop?

- Erbium-Doped Fiber Amplifiers (EDFA) are magic, but they do have their limits
- But if the input signal strength on each wavelengths are not mostly the same:



Why put attenuator on the loop?

- Erbium-Doped Fiber Amplifiers (EDFA) are magic, but they do have their limits
- But if the input signal strength on each wavelengths are not mostly the same:



The mux+amp amplifies the signal on the way out as well, so if you have a loopback the input is much brighter than most optics can be

The mux gives some cool stats too

```
admin@hex-egs-somuxc>show linkview
```


```
chassis-1
```

Local system				Fiber			Remote system		
Hostname	Interface	Status	Alarm	Power [dBm]	Loss [dB]	Direction	Power [dBm]	Interface	Hostname
hex-egs-somuxc	if-1/line-tx	up	ok	6.5	21.0	>>>>	-14.5	if-1/line-rx	egs-hex-somuxc
hex-egs-somuxc	if-1/line-rx	up	ok	-13.8	19.7	<<<<	5.9	if-1/line-tx	egs-hex-somuxc

```
admin@hex-egs-somuxc>show interface
```

Interface	Status [Rx/Tx]	Alarm	Rx power [dBm]	Tx power [dBm]	Format	Expected wavelength[nm]	Description
if-1/line	up/up	ok	-13.8	6.5	wdm	n/a	
if-1/9210	up/up	ok	-7.5	-0.7	Coherent	1560.61	
if-1/9220	up/up	ok	-8.3	-0.6	Coherent	1559.79	
if-1/9230	idle/idle	ok	-99.0	-99.0	n/a	1558.98	
...							
if-1/9590	up/up	ok	1.3	-2.3	1-32G	1530.33	


Rigged up



It didn't work!

- Built the setup up using MultiMode optics for testing and confirmed that worked
- Swapped in the DWDM optic, and that didn't work
- DWDM optic + loop confirmed working by using a 1G USB SFP NIC
- Turns out this DWDM optic has a much higher bar for input/output speed!

Square wave via a FS.COM DWDM-SFP10G-80




3MHz min?!

- Tested with the scopes wave generator, the optic will only start working at 3mhz!
- Except the square is not very square at 3 mhz

My signal is very out of spec anyway



Stable at 6.5MHz (ish)



Our TOSLINK signal is 2.2MHz, but the square wave only really becomes square at 6.5MHz

it's all I had for DWDM C59,

Trip failed :(

Why?

- Finding the internal chips inside these optics is very hard
- Even if you know the IC part numbers, getting the datasheet is harder
- I made a guess this is some kind of signal processing for the longer distance?
 - LX/SR/CWDM optics start working at 150KHz, so unclear why so much for this DWDM optic
- This optic was also a 10G DWDM, maybe a 1G optic would work?
- I sent some emails to a number of optics vendors asking about this behaviour

AaAAaaaAAa Retimers?!

- Finding the internal chips inside these optics is very hard
- Even if you know the IC part numbers, getting the datasheet is harder

Another piece of Information that might be helpful for you: **Make sure that the transceiver you want to use does not contain a clock and data recovery IC, sometimes also referred to as an eyeopener or retimer.** Those only can handle specific data rates and even the multirate ones likely won't be capable of dealing with your application. The good news is, **that most of the low data rate transceivers don't have those chips**, they rely only on limiting amplifiers for the RX path and Laser diode drivers for the TX path. These can handle non standard data rates. The long distance Transceivers for multi gigabit applications are more likely to have those ICs and should thus be avoided.

New day, New optics

- Acquired a *brand new* 1G DWDM optic
- Also a pile of 2nd hand DWDM optics that other people no longer needed
- Armed with better testing equipment at home so I can test at short distances before I go to the data center
- Will it work?

Suspect optics



Suspect optics

"HLSPDW-~~XE~~09" is suspect



Tweaked 10G Optics

- Searching for HLSPDW-XE09 brings up a load of 10G DWDM optics from other vendors

MF



Made-in-China.com : In stock
3.2g CWDM SFP Optical Tr...



Made-in-China.com
3.2g CWDM SFP Optical Tr...



Related searches

Tweaked 10G Optics

- Searching for HLSPDW-XE09 brings up a load of 10G DWDM optics from other vendors
- My optic behaves exactly like the 10G DWDM one
- Mine is almost certainly a 10G DWDM optic with it's EEPROM reprogrammed to be "1G"

MF



Related searches



Wait, The optics have a EEPROM?

```
eprom: 0000: 03 04 07 20 00 00 00 00 00 01 00 06 67 00 0a 64 ... ..g..d
0010: 00 00 00 00 51 53 46 50 54 45 4b 20 20 20 20 20 ....QSFP TEK
0020: 20 20 20 20 00 00 0b 40 51 54 2d 53 46 50 2b 2d ...@ QT-SFP+-
0030: 4c 52 20 20 20 20 20 20 42 20 20 20 05 1e 00 c9 LR B ...
0040: 00 1a 00 00 51 54 38 32 33 31 31 31 34 30 36 36 ....QT82 31114066
0050: 20 20 20 20 32 33 31 31 31 35 20 20 68 f0 03 07 2311 15 h...
0060: 00 00 08 8e a9 66 96 16 47 ca 48 cb d8 36 e2 ef .....f.. G.H..6..
0070: 07 13 c6 00 00 00 00 00 00 00 00 00 2e de f1 c5 .....
0080: 5a 00 d3 00 55 00 d8 00 8d cc 74 04 87 5a 7a 76 Z...U... ..t..Zzv
0090: c3 50 13 88 9c 40 27 10 45 77 04 eb 2b d4 07 cb .P...@'. Ew..+...
00a0: 37 2d 00 c8 2b d4 01 3c 00 00 00 00 00 00 00 00 7-...+..< .....
00b0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
00c0: 00 00 00 3f 80 00 00 00 00 00 00 00 01 00 00 00 .....?... .....
00d0: 00 00 00 01 00 00 00 01 00 00 00 00 00 00 64 .....d
00e0: 17 2c 7f 28 40 bf 1a 85 11 a9 00 00 00 00 08 00 ',.(@... .....
00f0: 00 00 00 40 00 00 08 00 00 00 00 00 00 00 00 ...@.....
```



**Tangent
!**

Wait, The optics have a EEPROM?


```
eeprom: 0000: 03 04 07 20 00 00 00 00 00 01 00 06 67 00 0a 64 ... ..g..d
0010: 00 00 00 00 51 53 46 50 54 45 4b 20 20 20 20 20 ....QSFP TEK (Vendor)
0020: 20 20 20 20 00 00 0b 40 51 54 2d 53 46 50 2b 2d ...@ QT-SFP+- (Part)
0030: 4c 52 20 20 20 20 20 20 42 20 20 20 05 1e 00 c9 LR B ....
0040: 00 1a 00 00 51 54 38 32 33 31 31 31 34 30 36 36 ....QT82 31114066 (Serial)
0050: 20 20 20 20 32 33 31 31 31 35 20 20 68 f0 03 07 2311 15 h... (Date)
0060: 00 00 08 8e a9 66 96 16 47 ca 48 cb d8 36 e2 ef .....f.. G.H..6..
0070: 07 13 c6 00 00 00 00 00 00 00 00 00 2e de f1 c5 .....
0080: 5a 00 d3 00 55 00 d8 00 8d cc 74 04 87 5a 7a 76 Z...U... ..t..Zzv
0090: c3 50 13 88 9c 40 27 10 45 77 04 eb 2b d4 07 cb .P...@'. Ew..+...
00a0: 37 2d 00 c8 2b d4 01 3c 00 00 00 00 00 00 00 00 7-...+..< .....
00b0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
00c0: 00 00 00 3f 80 00 00 00 00 00 00 00 01 00 00 00 .....?...
00d0: 00 00 00 01 00 00 00 01 00 00 00 00 00 00 64 .....d
00e0: 7 2c 7f 28 40 bf 1a 85 11 a9 00 00 00 00 08 00 ',.(@...
00f0: 00 00 00 40 00 00 08 00 00 00 00 00 00 00 00 ...@.....
```



Tangent
!


140 km: Round 2

- Armed with a old 1G DWDM optic on a different channel, since it seems that brand new DWDM 1G optics are not being made anymore!

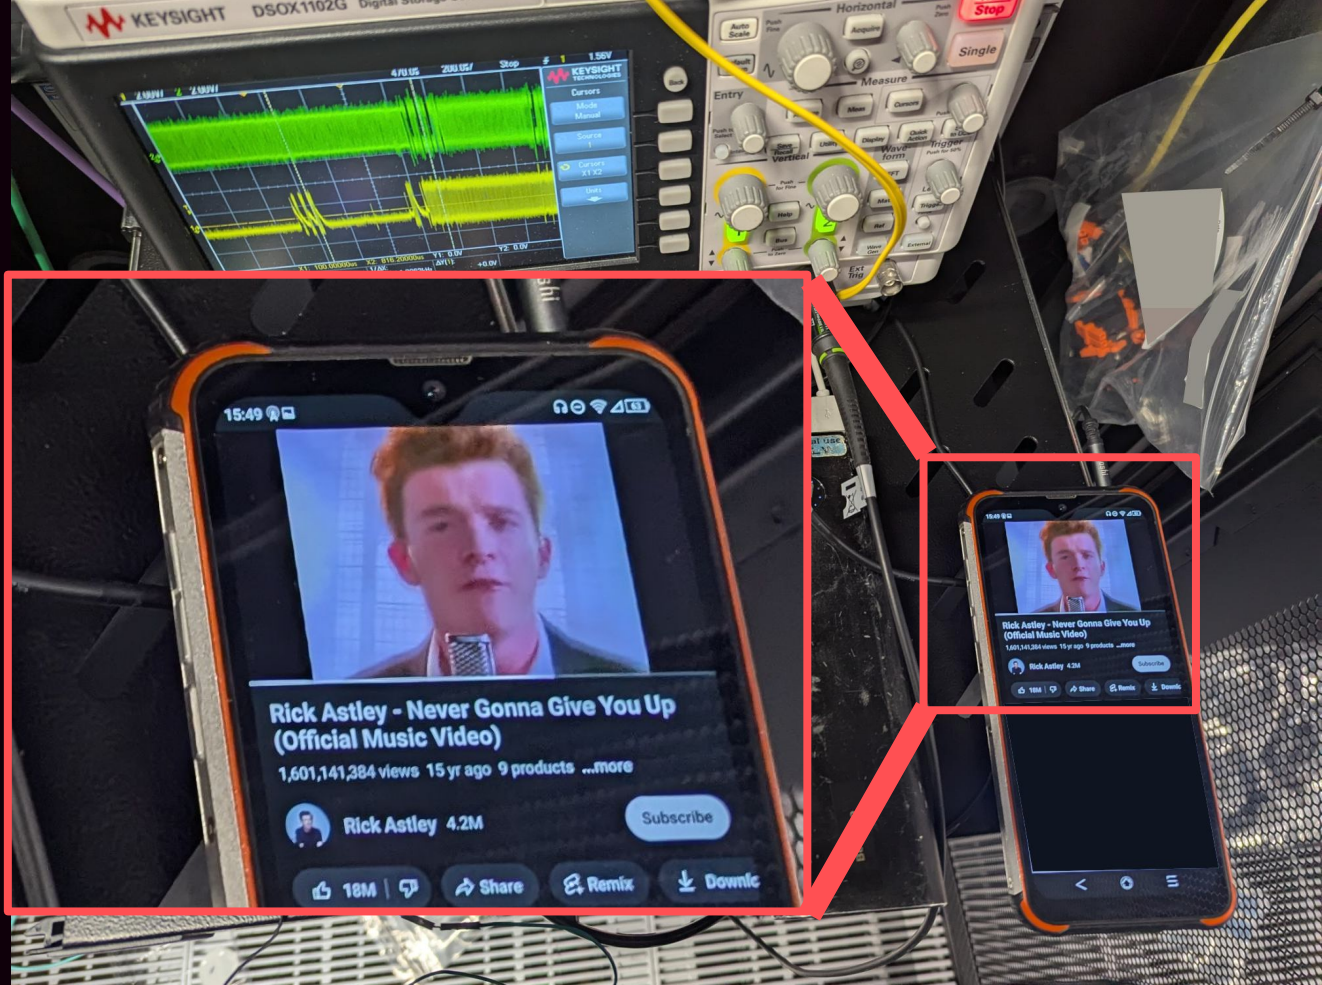


- Harder to measure latency this time by matching bits
- Instead measure "strope time"
- 720 μ s round trip
- ~143.2km

140 km: Round 2



140 km: Round 2



Future (pointless) innovation

- If the SFP pluggable didn't have a LVDS in it, it would be nice to drive the laser directly, send "raw" audio in its full glory (and probably damage the laser)
 - I hear this is sorta how "RF over Fiber" works
 - A friend also suggested YOLO-ing a FM signal might work, I didn't get to test this, but I bet it does

Future (pointless) innovation

- If the SFP pluggable didn't have a LVDS in it, it would be nice to drive the laser directly, send "raw" audio in its full glory (and probably damage the laser)
 - I hear this is sorta how "RF over Fiber" works
 - A friend also suggested YOLO-ing a FM signal might work, I didn't get to test this, but I bet it does
- It is tempting to attach a "dialup" modem to both sides
 - This would create the greatest waste of a 100 GHz optical channel, in that it gives a final output bandwidth of ~40 kbit/s
 - And ***extremely*** confuse some intelligence agency I assume

To conclude

- Yes, you can send "low" speed signals over SFP optics
 - Sometimes, Most optics start working at 150 KHz, the retimers don't like working below 6 MHz for some optics
- You can even send them very far

To conclude

- Yes, you can send "low" speed signals over SFP optics
 - Sometimes, Most optics start working at 150 KHz, the retimers don't like working below 6 MHz for some optics
- You can even send them very far
- There is likely not a good reason to do this. There are better IP solutions now
 - But knowing how things work is important! Knowing something is possible is useful to build bigger, better, **more horrifying** solutions or workarounds to problems
- I know far more about optical transmission, and general inner optics workings now
 - Maybe you do too!

Questions?

Shy? Ask me here instead:

Email 38c3@benjojo.co.uk

Fedi [@benjojo@benjojo.co.uk](https://benjojo@benjojo.co.uk)