Go Long!

Sending weird signals long distances over existing optical infrastructure

Ben Cartwright-Cox

Hi

- I am Ben Cartwright-Cox
- You may know me from such things like:
  - bgp.tools
  - BGP Battleships
  - Using the weird SFP's with computers inside the optic
  - Last years BGP DoS "exploits"
  - 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

Optical capable stuff is everywhere

Optics are a weird world

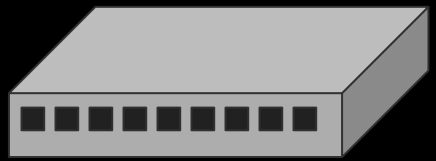

Device

1000baseKX
10000baseKR
25000baseCR
Etc etc

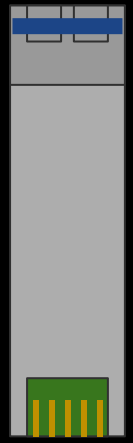

Optic Type

SFP
SFP+
XFP
QSFP
QSFP+
OSFP

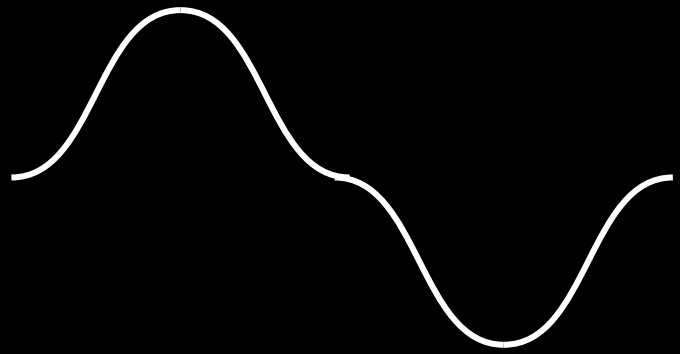

Transmission / Wavelength

LR
LX
SR
SX
ER
BX

850nm
1310nm
1510nm
CWDM
DWDM

Average "boring" 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*

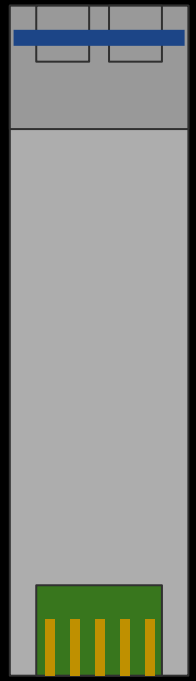

We will ignore weird things that have full CPUs in them

VDSL Modem in SFP

ARM Machine inside SFP

GPON ONT inside SFP

Some optics are surprisingly stupid

This thing is NOT thinking*

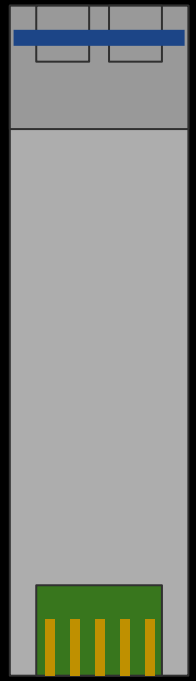

Some optics are surprisingly stupid

This thing is NOT thinking*

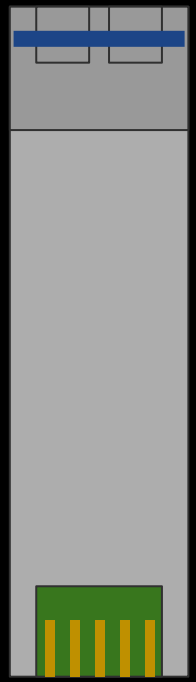

* There is a EEPROM to tell the switch what it is (and have the switch complain to you to buy vendor optics instead)

Some optics are surprisingly stupid

This thing is NOT thinking*

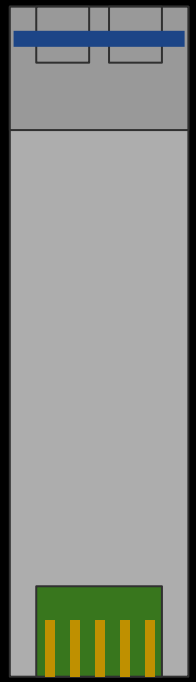

* **There is also** sometimes a tiny microcontroller to read out laser light level metrics and temperature for the switch

Average faster "boring" 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 like "LR1"
- 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 nice ones have to turn 4 "lanes" of 25GBaud data into one "normal" beam 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 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/PSM) the complexity (and price) of the chips are better/worse

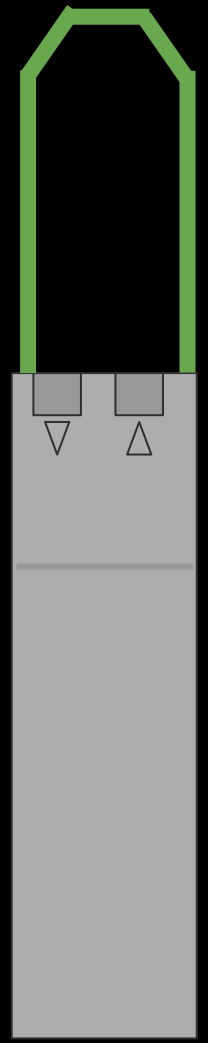

Other optic form factors are available

- 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**

So I was thinking

So I was thinking

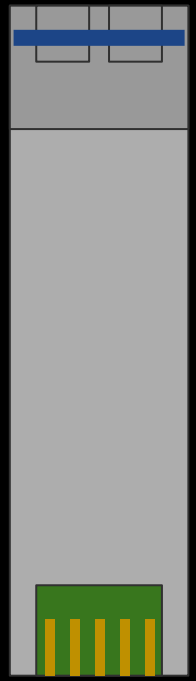

So I was thinking

Too stupid to care

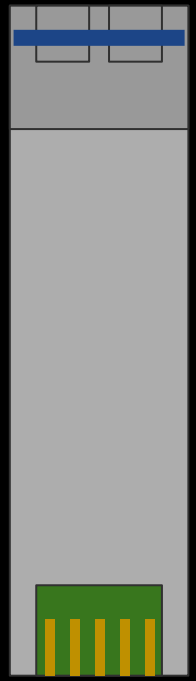

So I was thinking

+

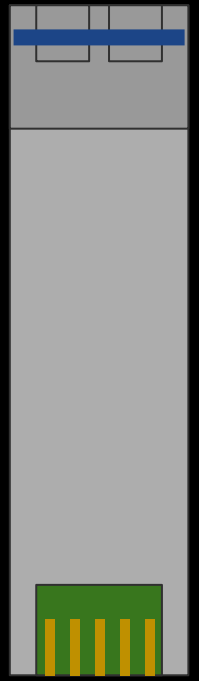

So I was thinking

+

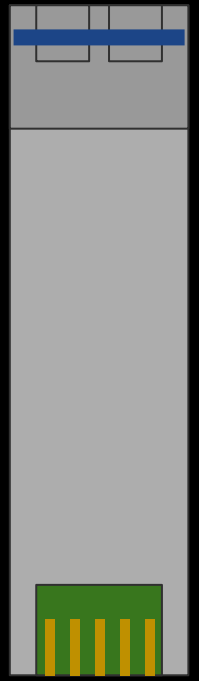

10km TOSLINK?

We are going to need to do some things

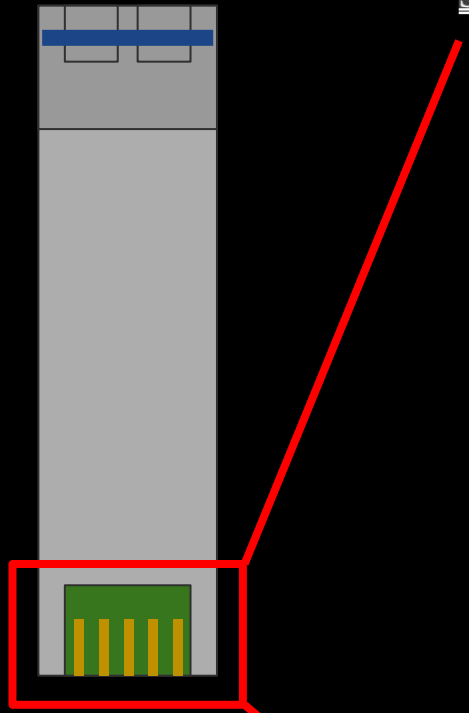

Small Form-factor Pluggable (SFP) Transceiver MultiSource Agreement (MSA)

	Top of Board		Bottom of Board (as viewed thru top of board)
20	VeeT	1	VeeT
19	TD-		
18	TD+	3	Tx Disable
17	VeeT		
16	VccT		
15	VccR		
14	VeeR		
13	RD+		
12	RD-	9	VeeR
11	VeeR	10	VeeR

We are going to need to do some things

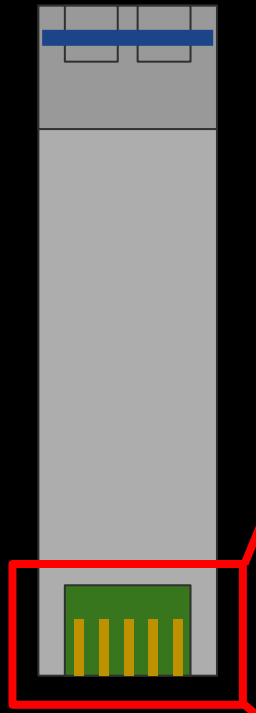

Small Form-factor Pluggable (SFP) Transceiver MultiSource Agreement (MSA)

	Top of Board		Bottom of Board (as viewed thru top of board)
20	VeeT	1	VeeT
19	TD-	2	TxFault
18	TD+	3	Tx Disable
17	VeeT	4	MOD-DEF(2)
16	VccT	5	MOD-DEF(1)
15	VccR	6	MOD-DEF(0)
14	VeeR	7	Rate Select
13	RD+	8	LOS
12	RD-	9	VeeR
11	VeeR	10	VeeR

We are going to need to do some things

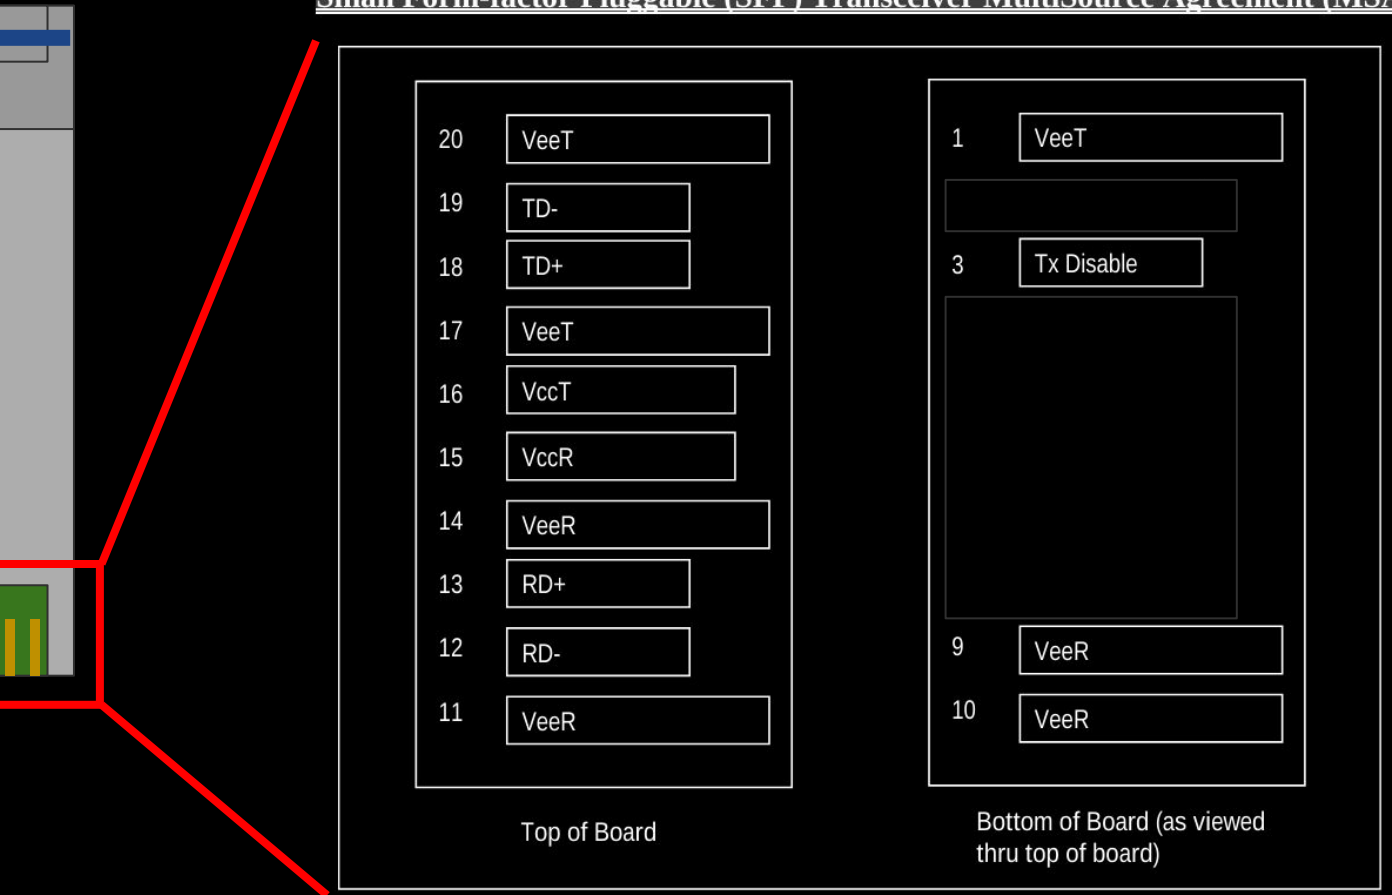

VeeT: TX Ground
VeeR: RX Ground
VccT: Tx Power
VccR: Rx Power

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

Cheating a bit

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

- A 5v to 3.3v power supply
- SFP Connecter + 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 made some Osmocom SFP Experimenter boards, They come with:

- A 5v to 3.3v power supply
- SFP Connecter + 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

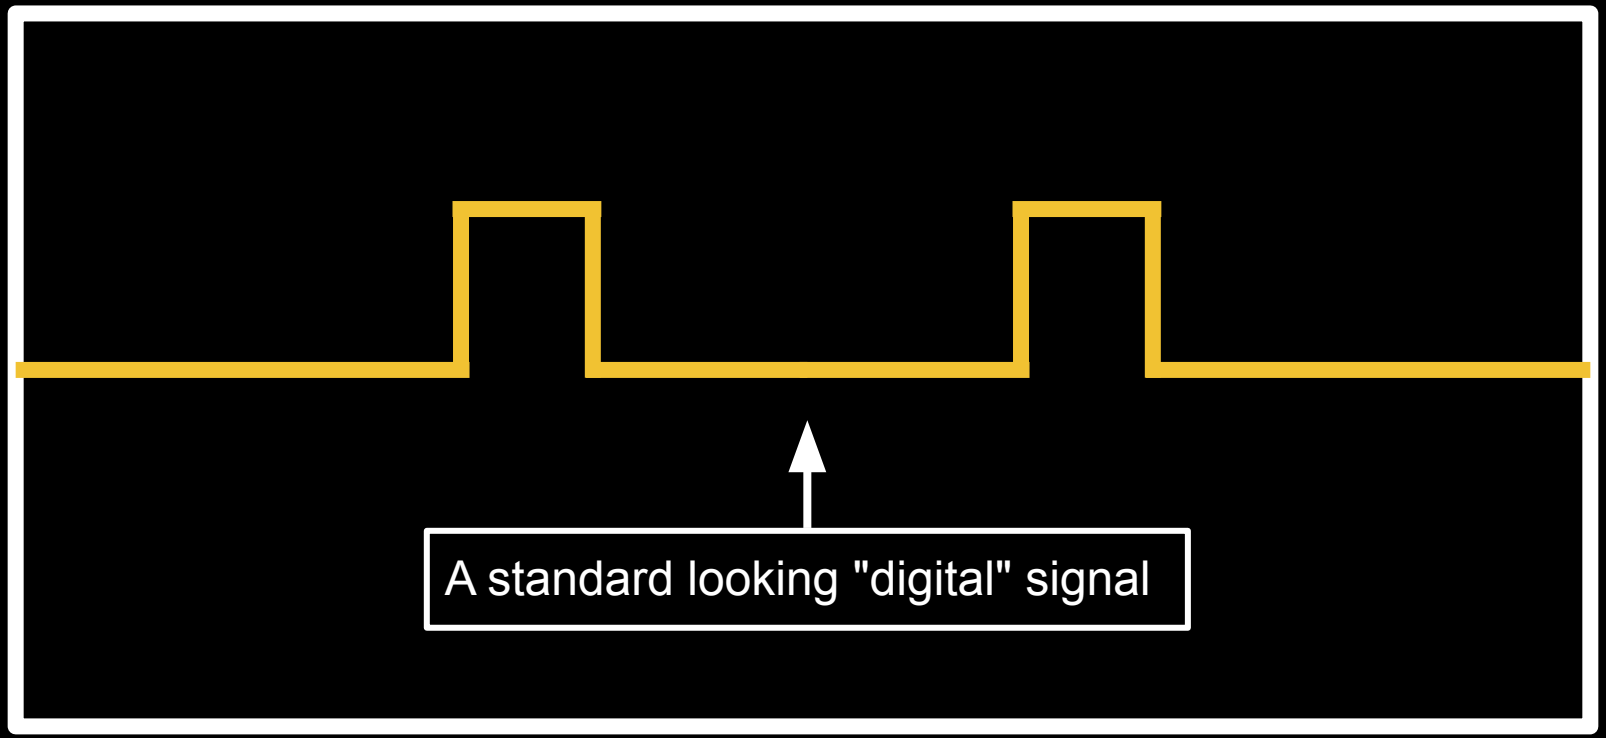

Line drivers make things wobble just right

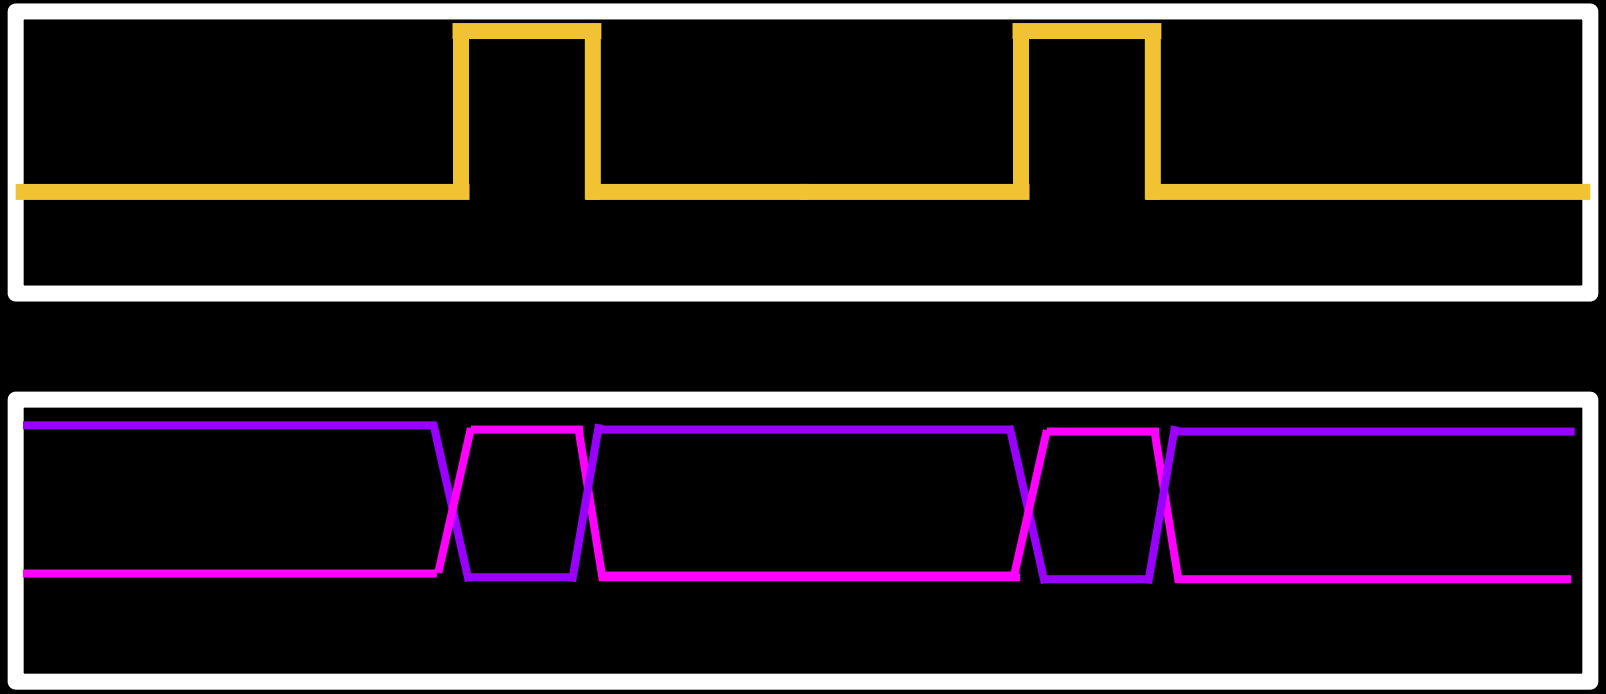

Audio In

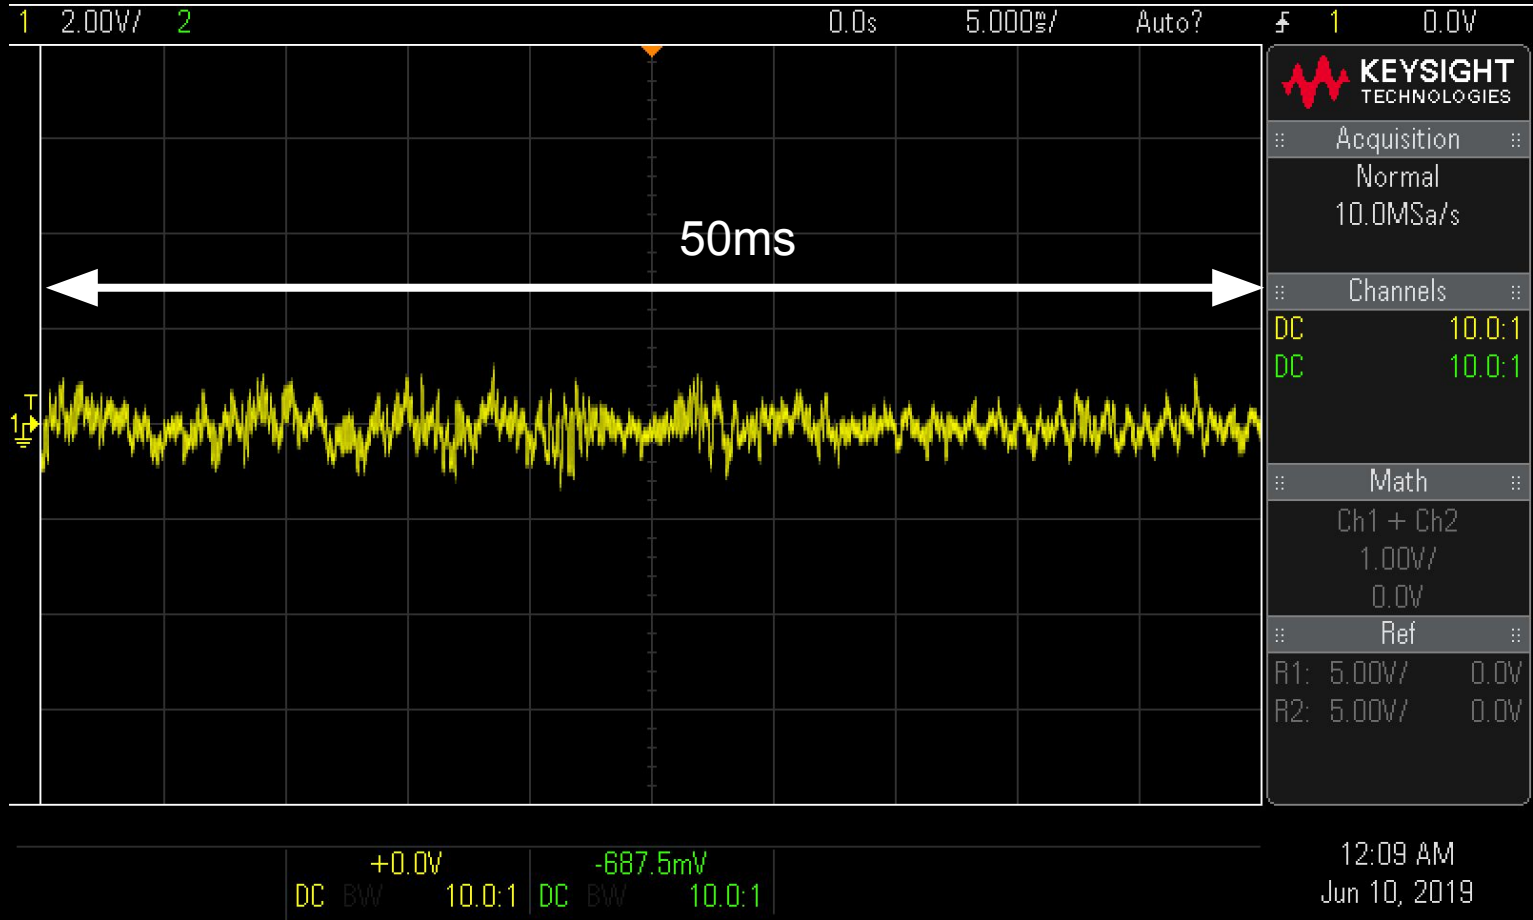
1 2.00V/ 2
0.0s
5.000ms/
Auto?
1 0.0V
50ms
KEYSIGHT TECHNOLOGIES
Acquisition
Normal
10.0MSa/s
Channels
DC 10.0:1
DC 10.0:1
Math
Ch1 + Ch2
1.00V/
0.0V
Ref
R1: 5.00V/ 0.0V
R2: 5.00V/ 0.0V
+0.0V
DC BW 10.0:1
-687.5mV
DC BW 10.0:1
12:09 AM
Jun 10, 2019

Line drivers make things wobble just right

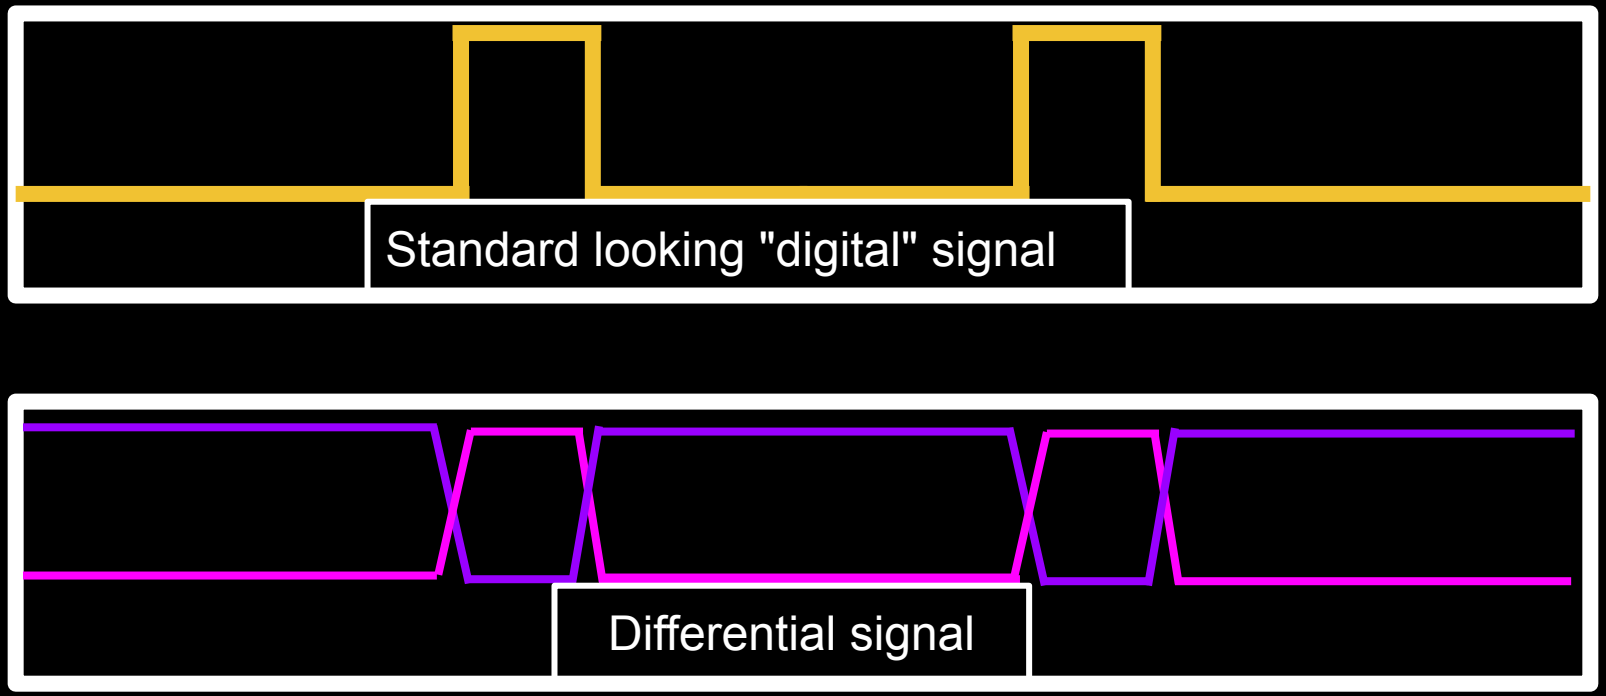

Audio In

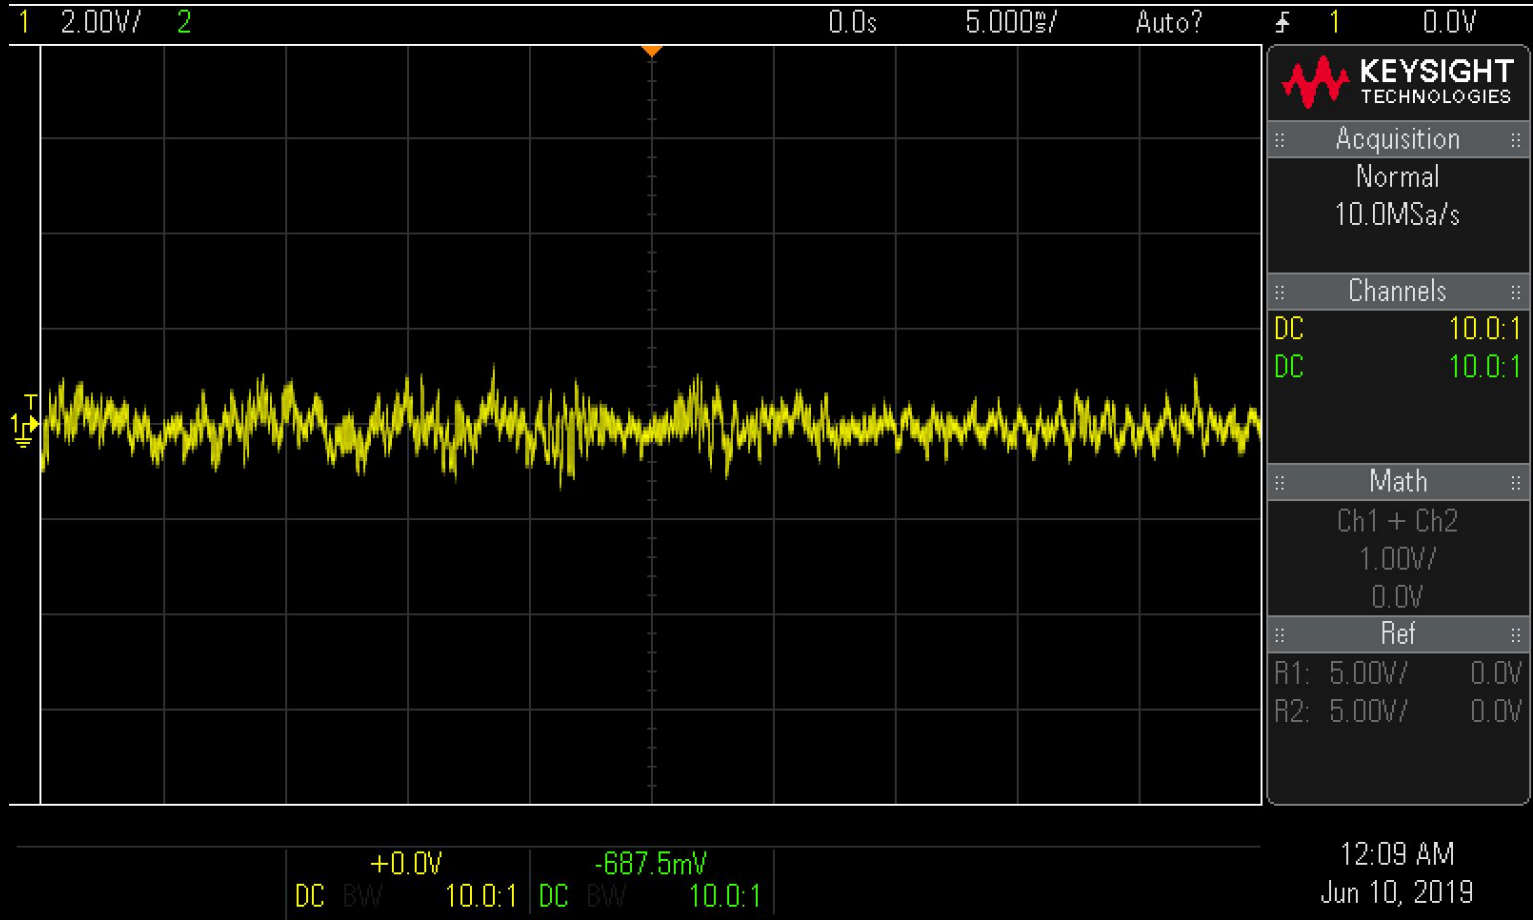
1 2.00V/ 2
0.0s
5.000ms/
Auto?
1 0.0V
KEYSIGHT TECHNOLOGIES
Acquisition
Normal
10.0MSa/s
Channels
DC 10.0:1
DC 10.0:1
Math
Ch1 + Ch2
1.00V/
0.0V
Ref
R1: 5.00V/ 0.0V
R2: 5.00V/ 0.0V
+0.0V
DC BW 10.0:1
-687.5mV
DC BW 10.0:1
12:09 AM
Jun 10, 2019

SPDIF Out

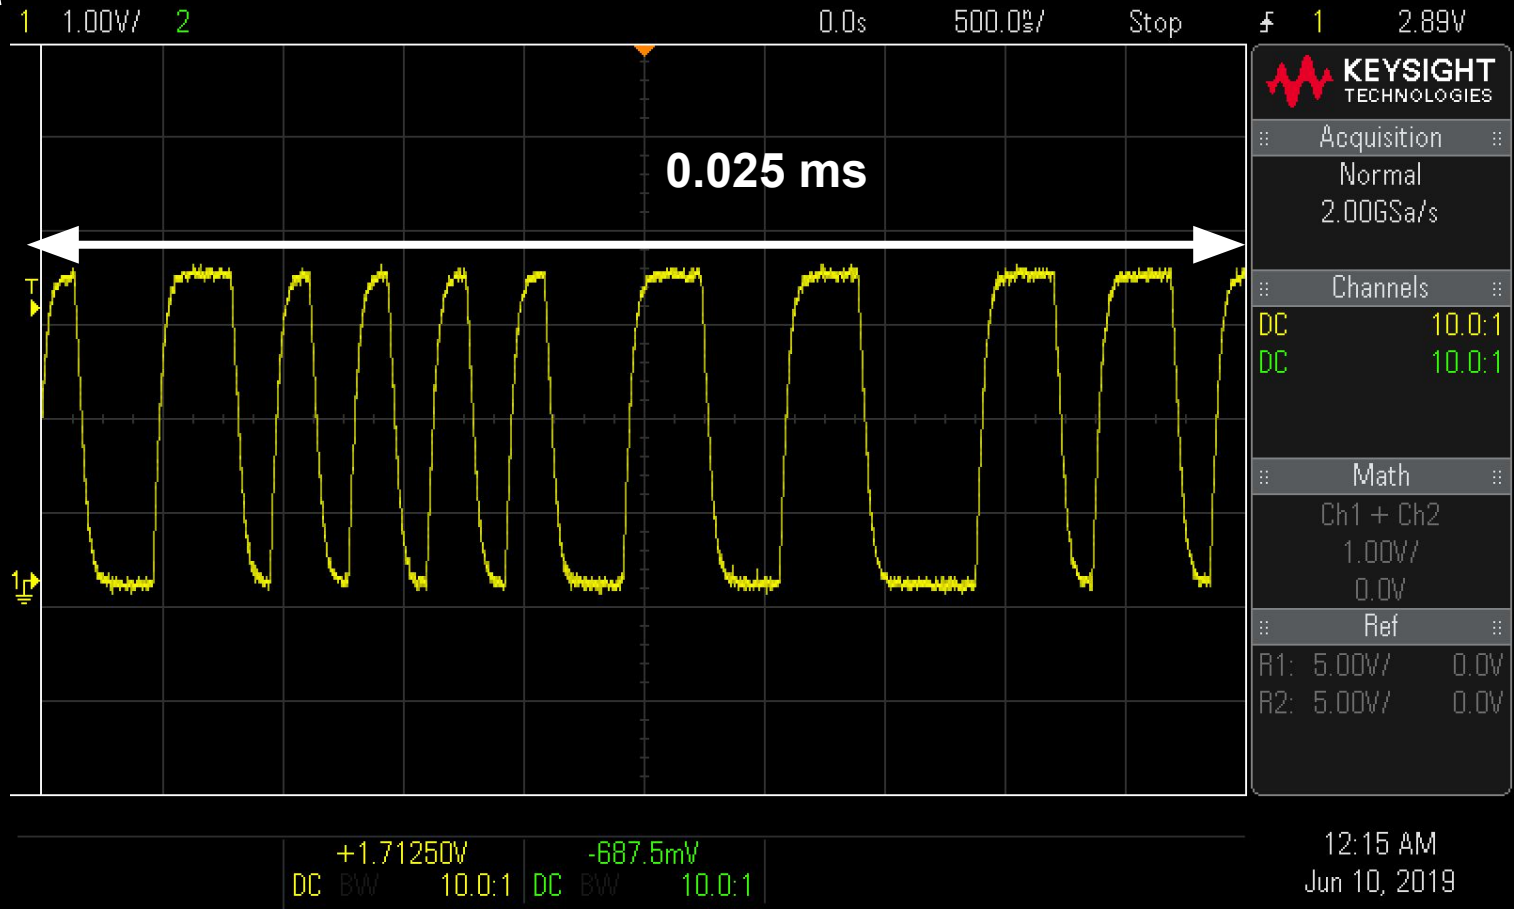
1 1.00V/ 2
0.0s
500.0%/
Stop
1 2.89V
0.025 ms
KEYSIGHT TECHNOLOGIES
Acquisition
Normal
2.00GSa/s
Channels
DC 10.0:1
DC 10.0:1
Math
Ch1 + Ch2
1.00V/
0.0V
Ref
R1: 5.00V/ 0.0V
R2: 5.00V/ 0.0V
+1.71250V
DC BW 10.0:1
-687.5mV
DC BW 10.0:1
12:15 AM
Jun 10, 2019

About 2Mhz~ of data

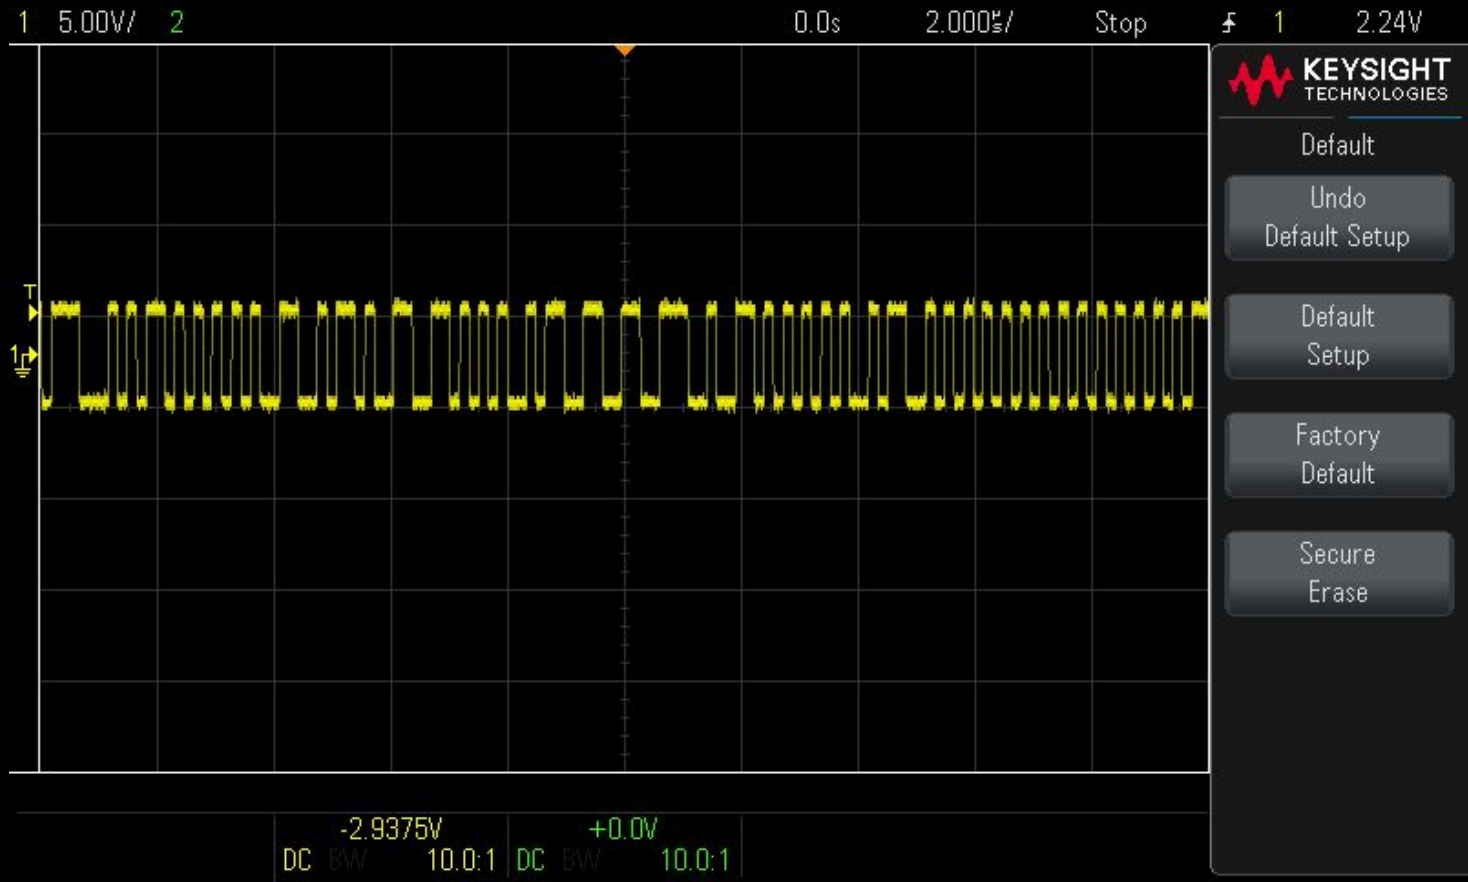

POC Setup

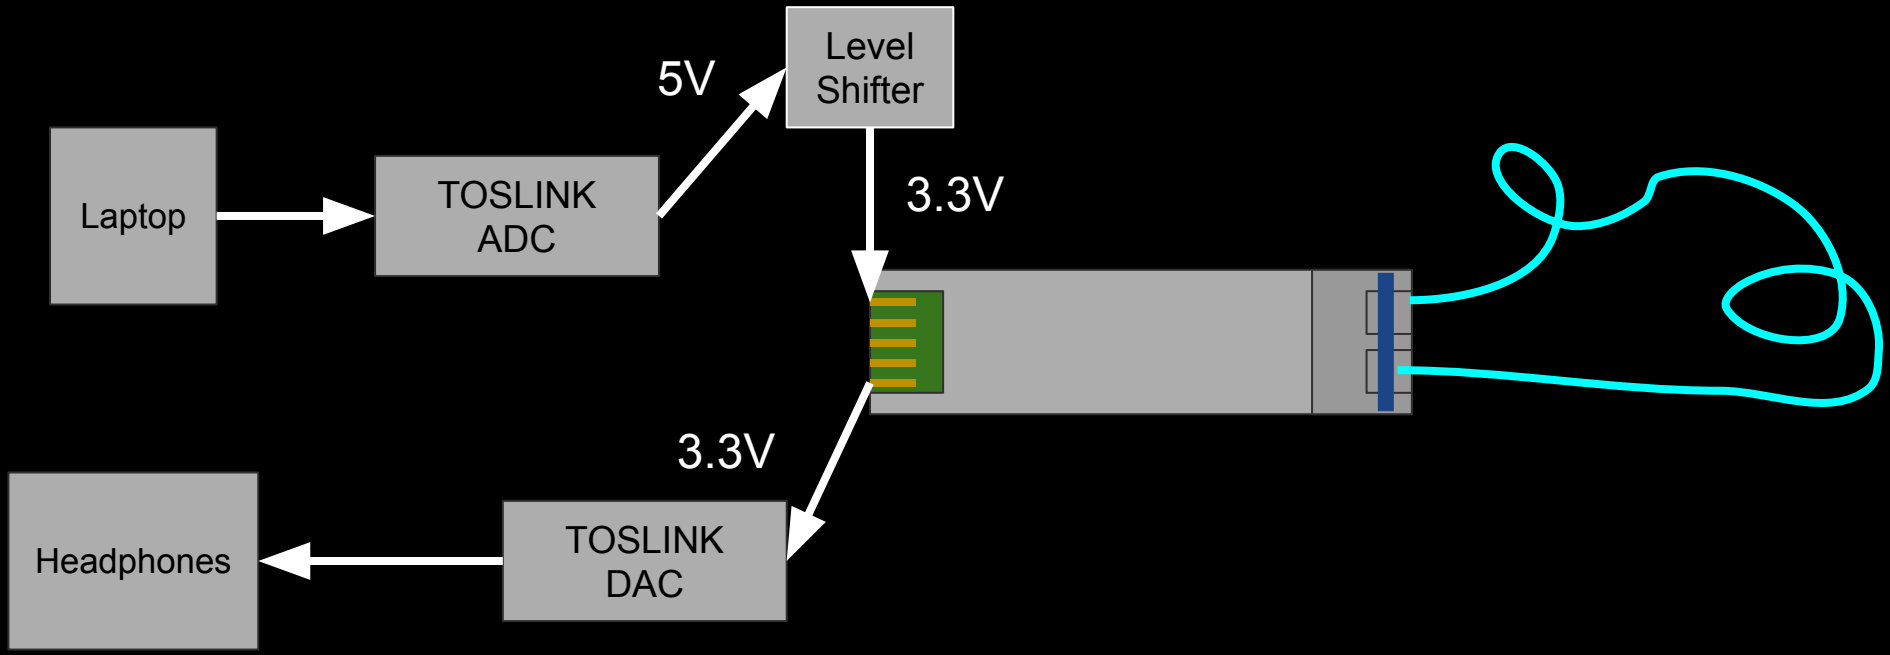
Laptop
TOSLINK ADC
5V
Level Shifter
3.3V
3.3V
TOSLINK DAC
Headphones

Working POC

- Not TSA friendly

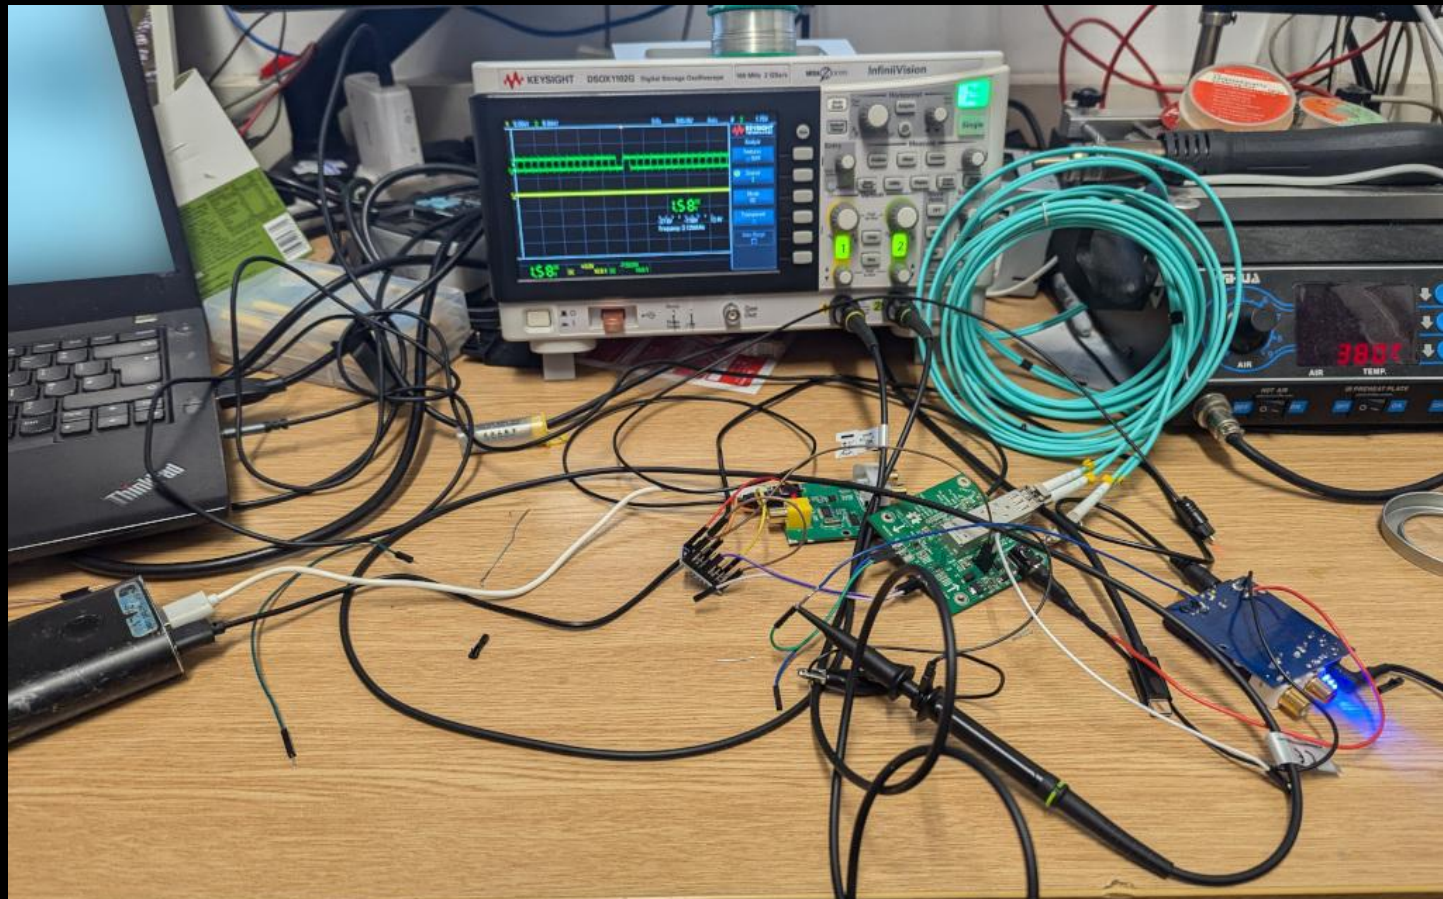

Latency

Adds about 34ns~

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

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

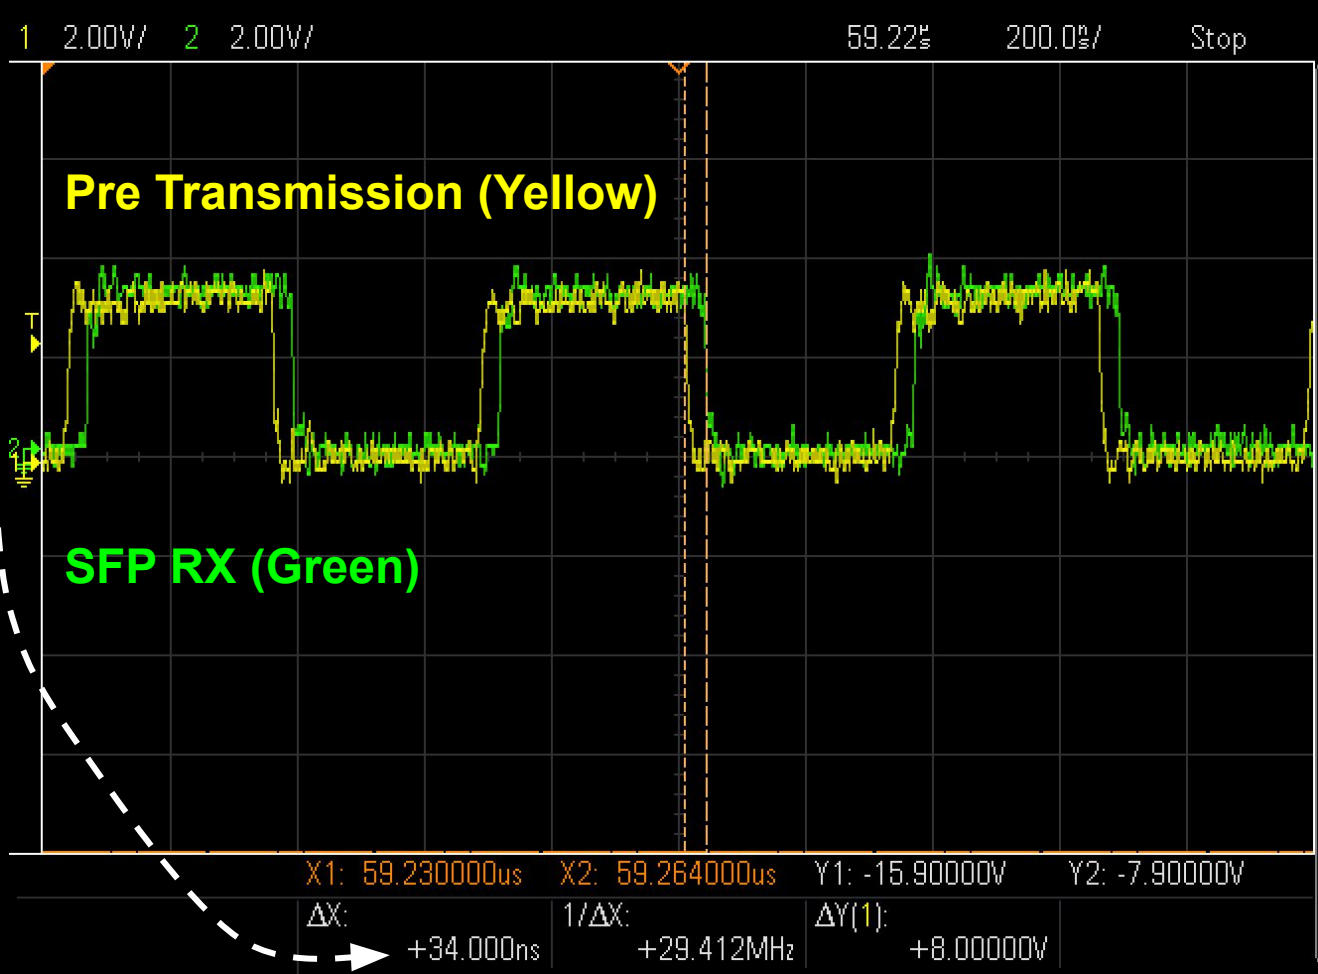

But what if we could go long haul?

About 650
meters~
IP
H
Telehouse
Poplar Riverside
Abbott Rd
Dee St
B125
Blair St
Village Ave
Lanrick Rd
E India Dock Rd
Bow Creek
Bow Creek Ecology Park
Orchard Wharf
Limmo
Peninsula
Ecological
Park
River Lea
Hope St

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

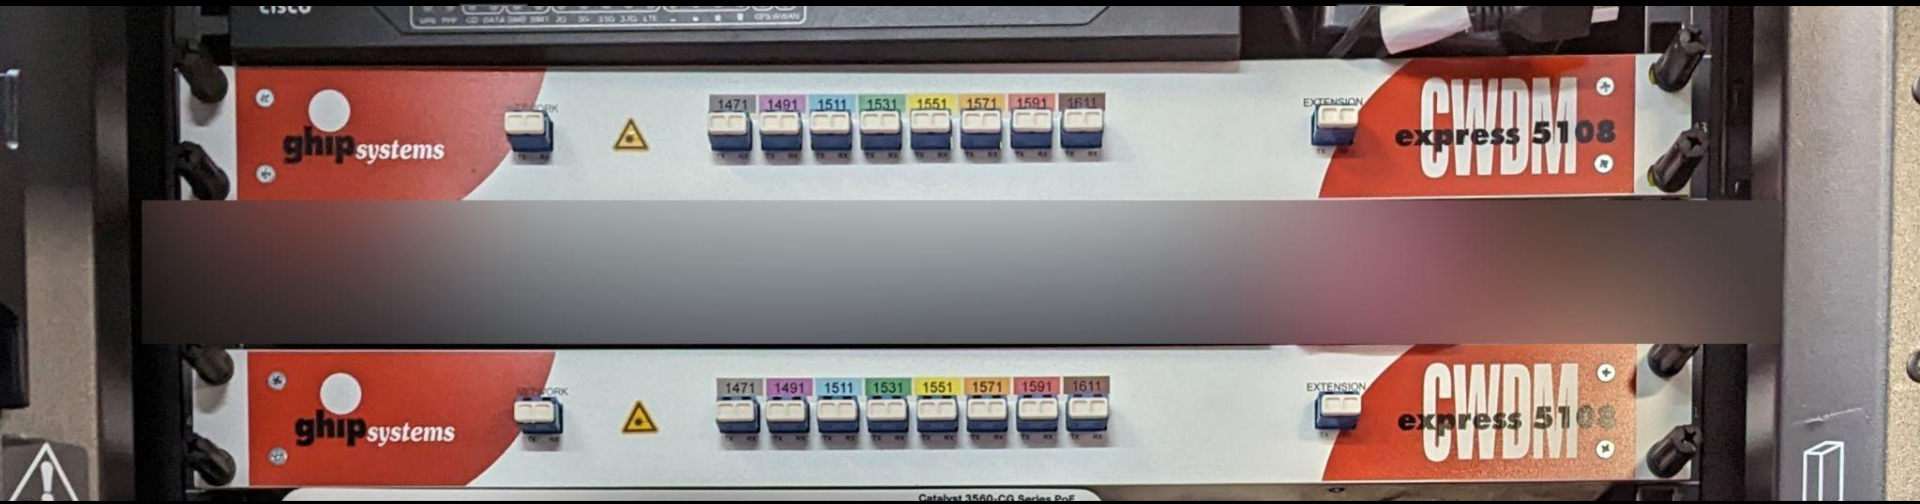

CWDM Channel from Jump Networks ~1km ish

- 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 blow up the CWDM optic!

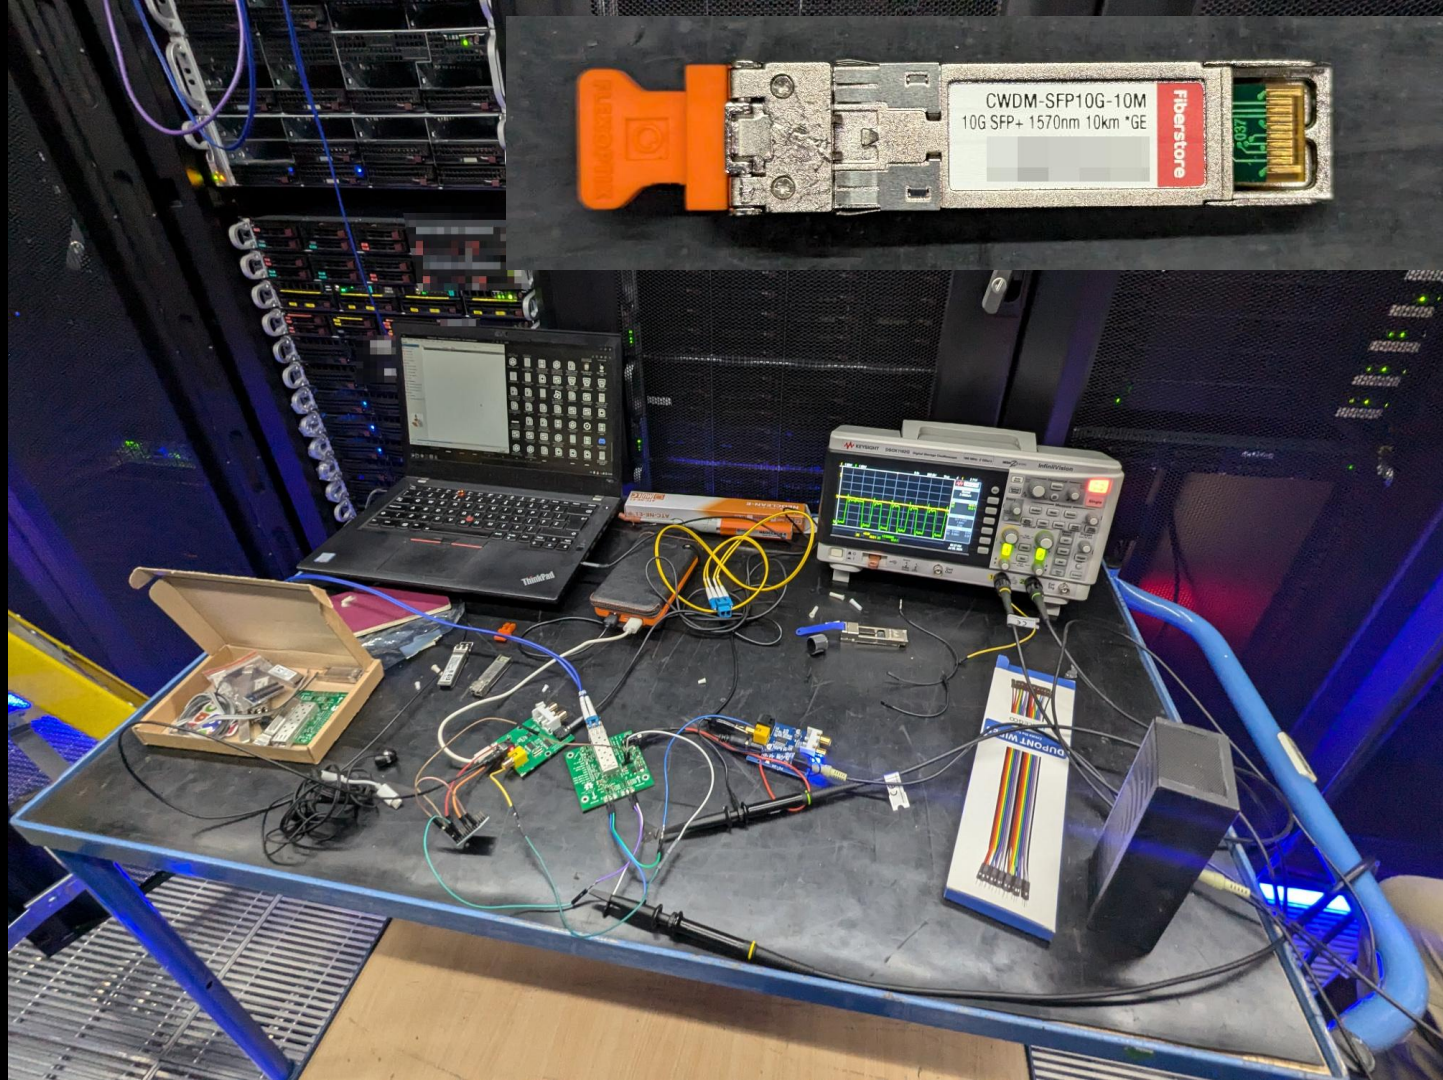

Latency

Harder this time on the scope because *11µs* requires you to zoom out a lot to find the offset

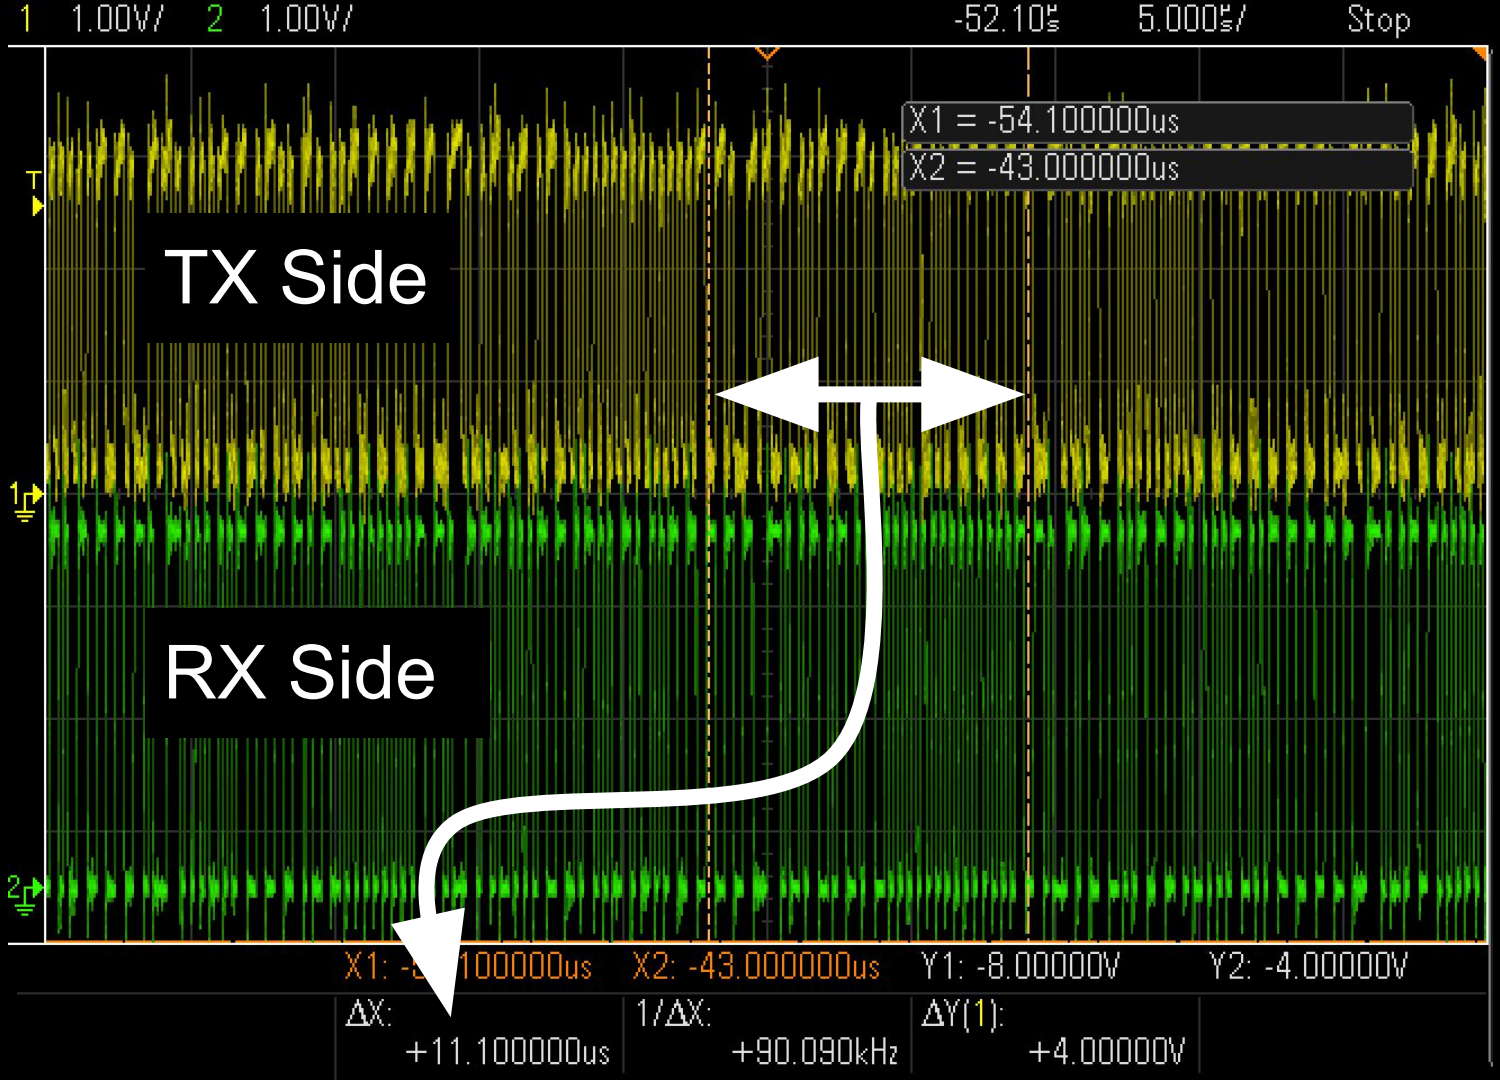

11μs = 2.2km

- Double what a very rough google maps estimate gave

- DC cabling on both sides likely add up a lot
- This is more like the distance if added a circling of both DC buildings (in 2D)

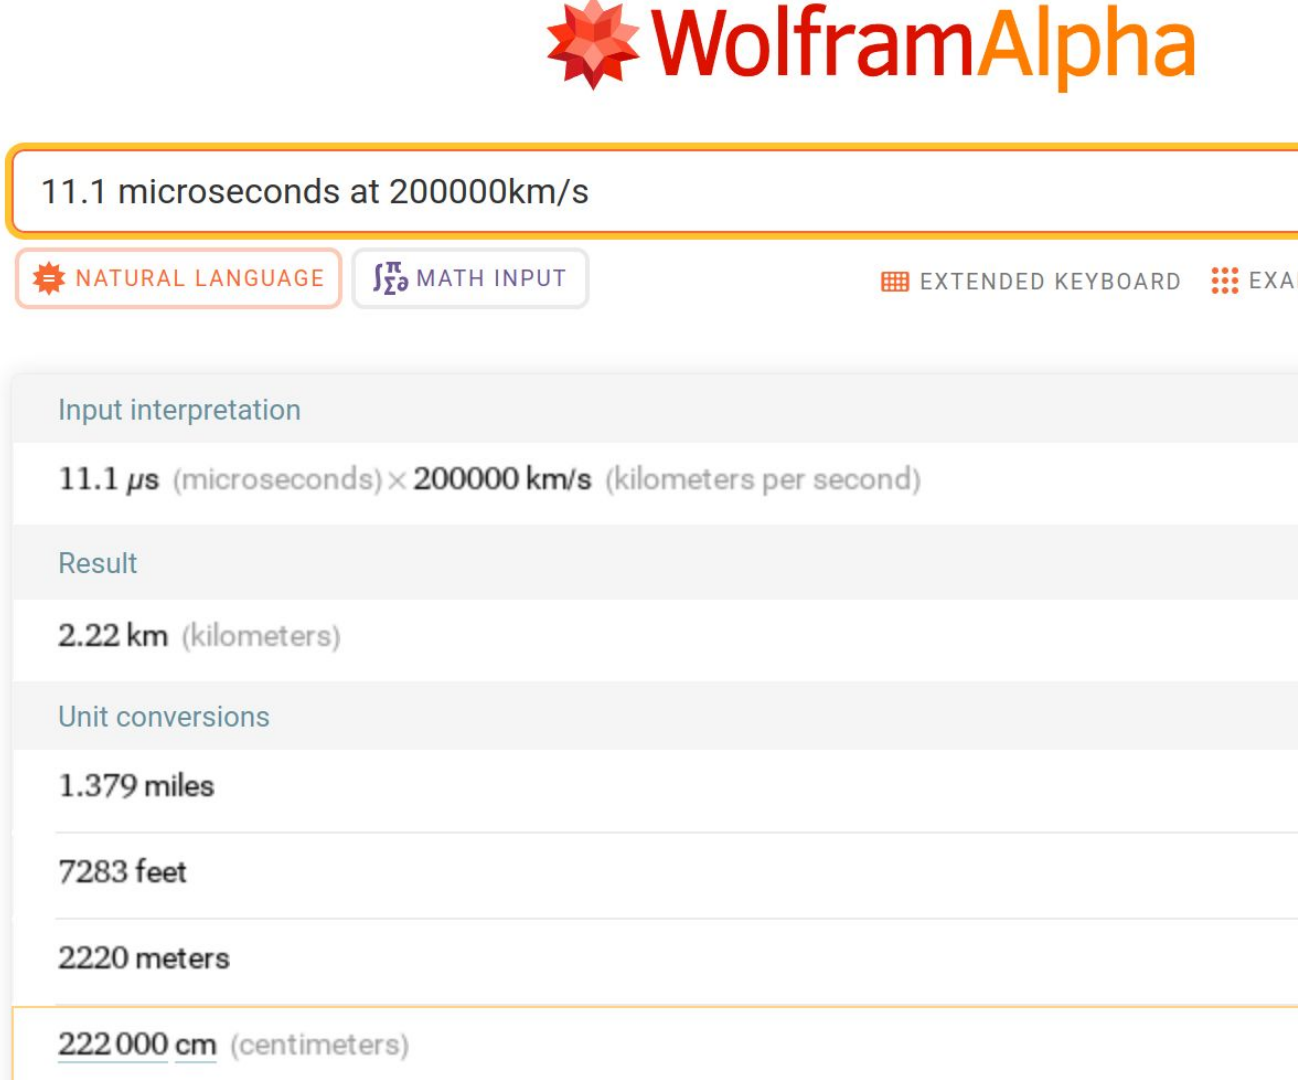

But what if we could go longer haul?

Longer haul options

- Typically, if you were doing 1G/10G you would be doing DWDM at this point
- DWDM is like CWDM, but uh, Denser!
- 8 channels to 32/64~ channels.
- Optics are more expensive because they have cooler/hotter physics inside of them

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 of them
- You can also amplify these to get more distance
    - Must be careful, as all of the input channels have to be close to the same brightness else weird stuff happens to your EDFA amplifier
    - 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 in the other 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 increasingly common**

Inter-city TOSLINK Audio

- LONAP (❤️) had a spare DWDM between:
  - Equinix LD6 (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

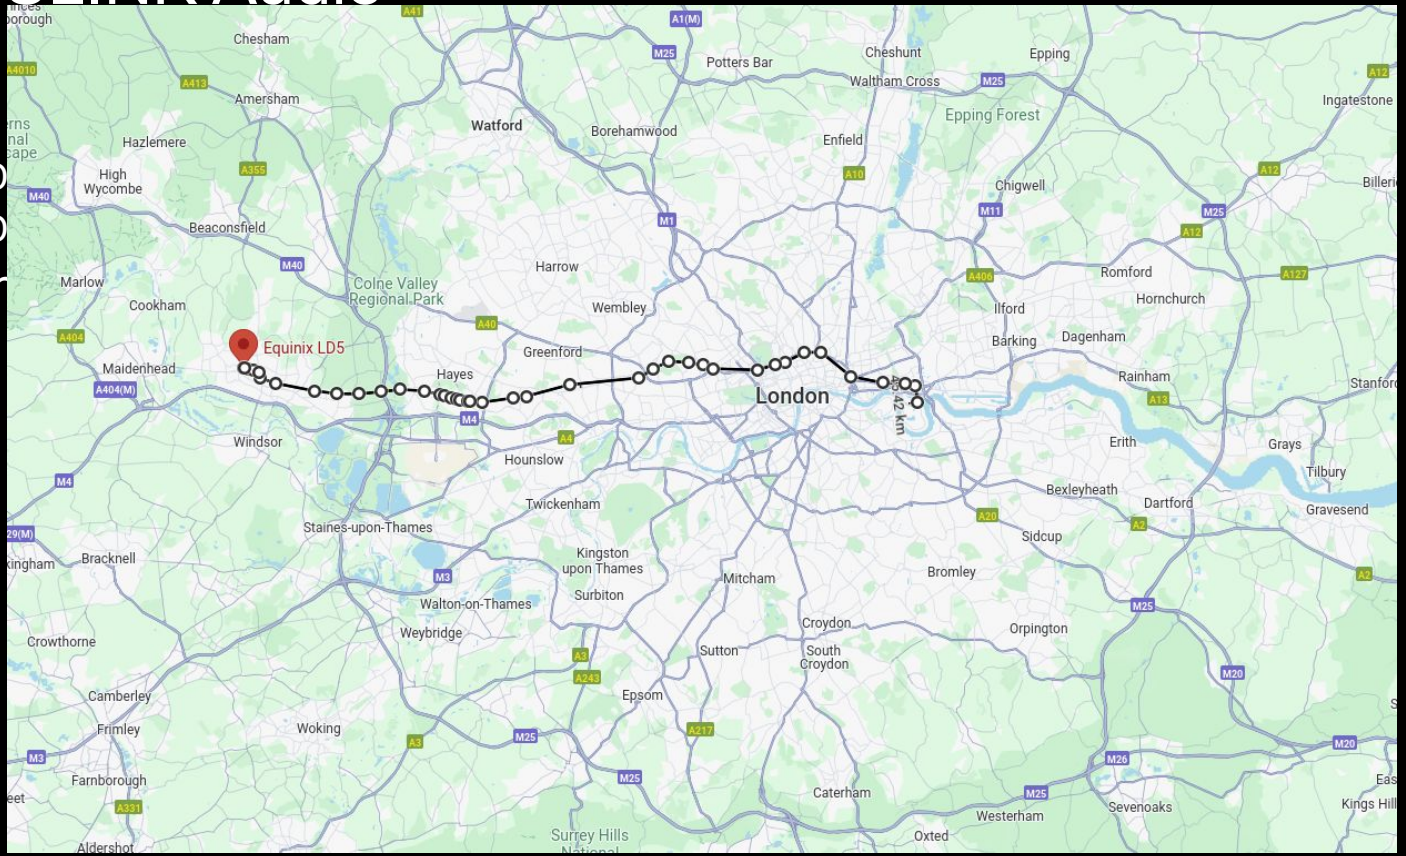

Inter-city TOSLINK Audio

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

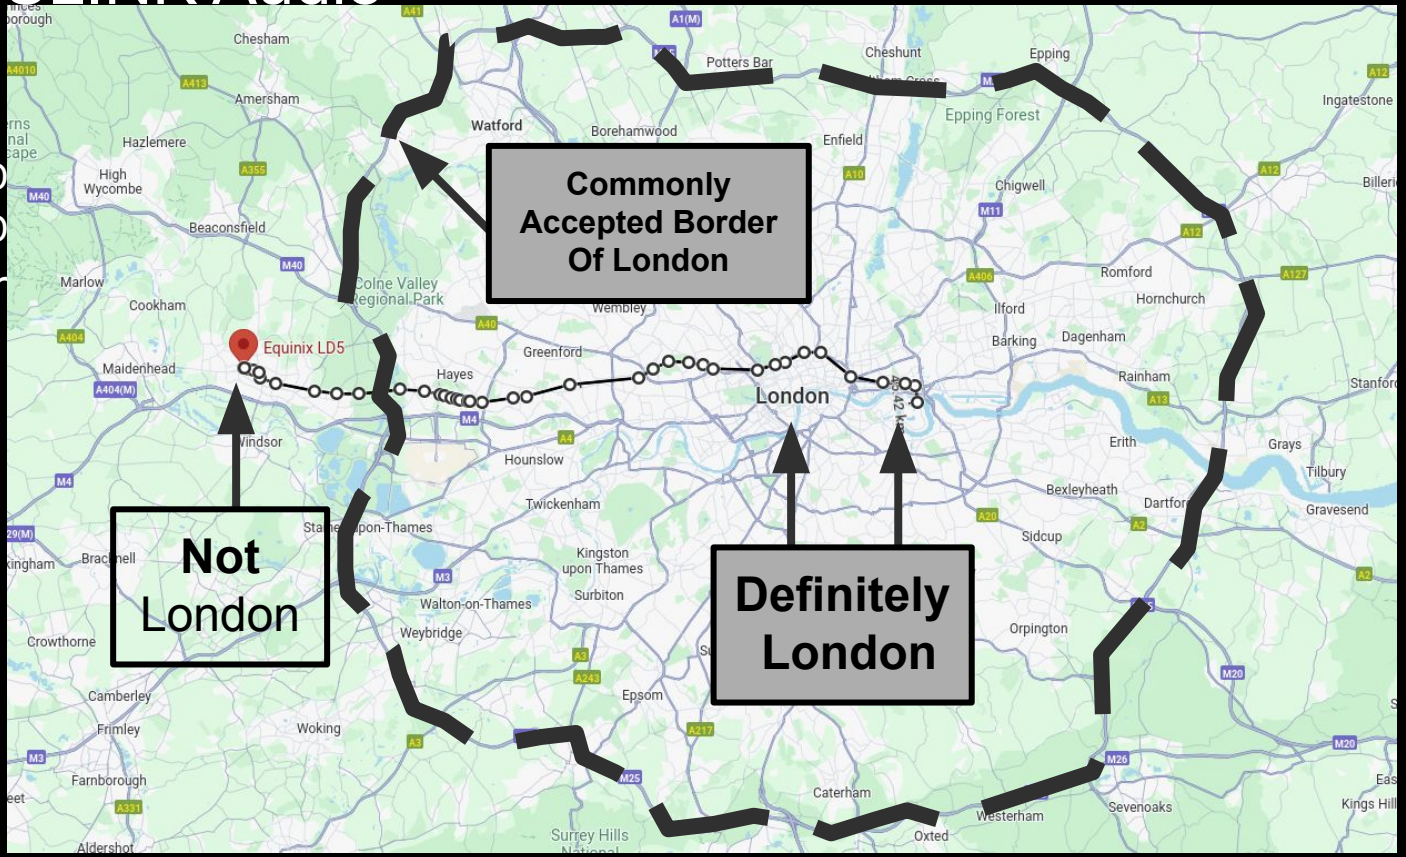

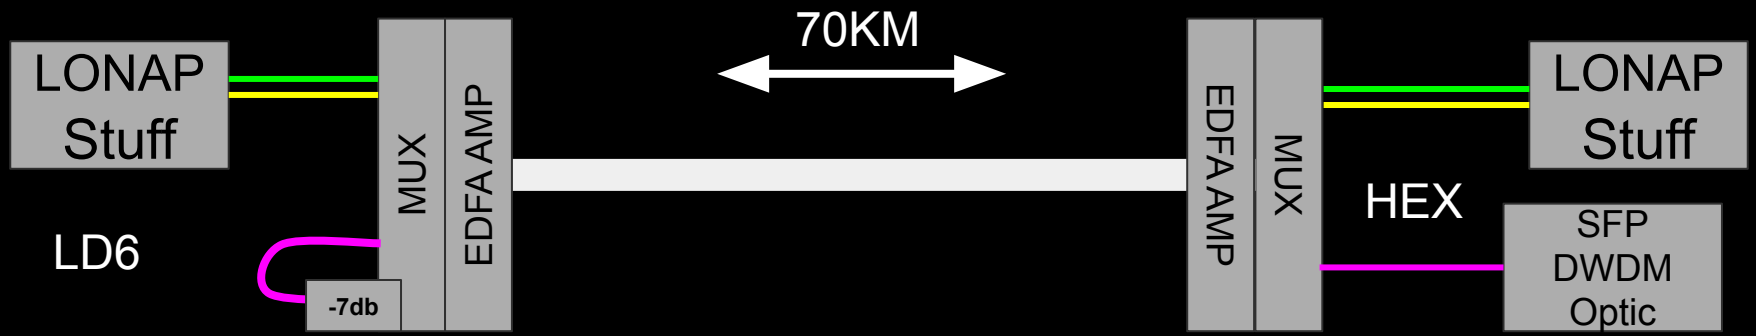
LONAP
Stuff
LD6
MUX
EDFA AMP
-7db
70KM
EDFA AMP
MUX
HEX
LONAP
Stuff
SFP
DWDM
Optic

Inter-city TOSLINK Audio

- Used a boring 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

SmartOptics gives some cool stats too

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

chassis-1

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

admin@hex-eqs-somuxc>show interface

              Status                Rx power  Tx power            Expected
  Interface   [Rx/Tx]     Alarm     [dBm]     [dBm]     Format    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

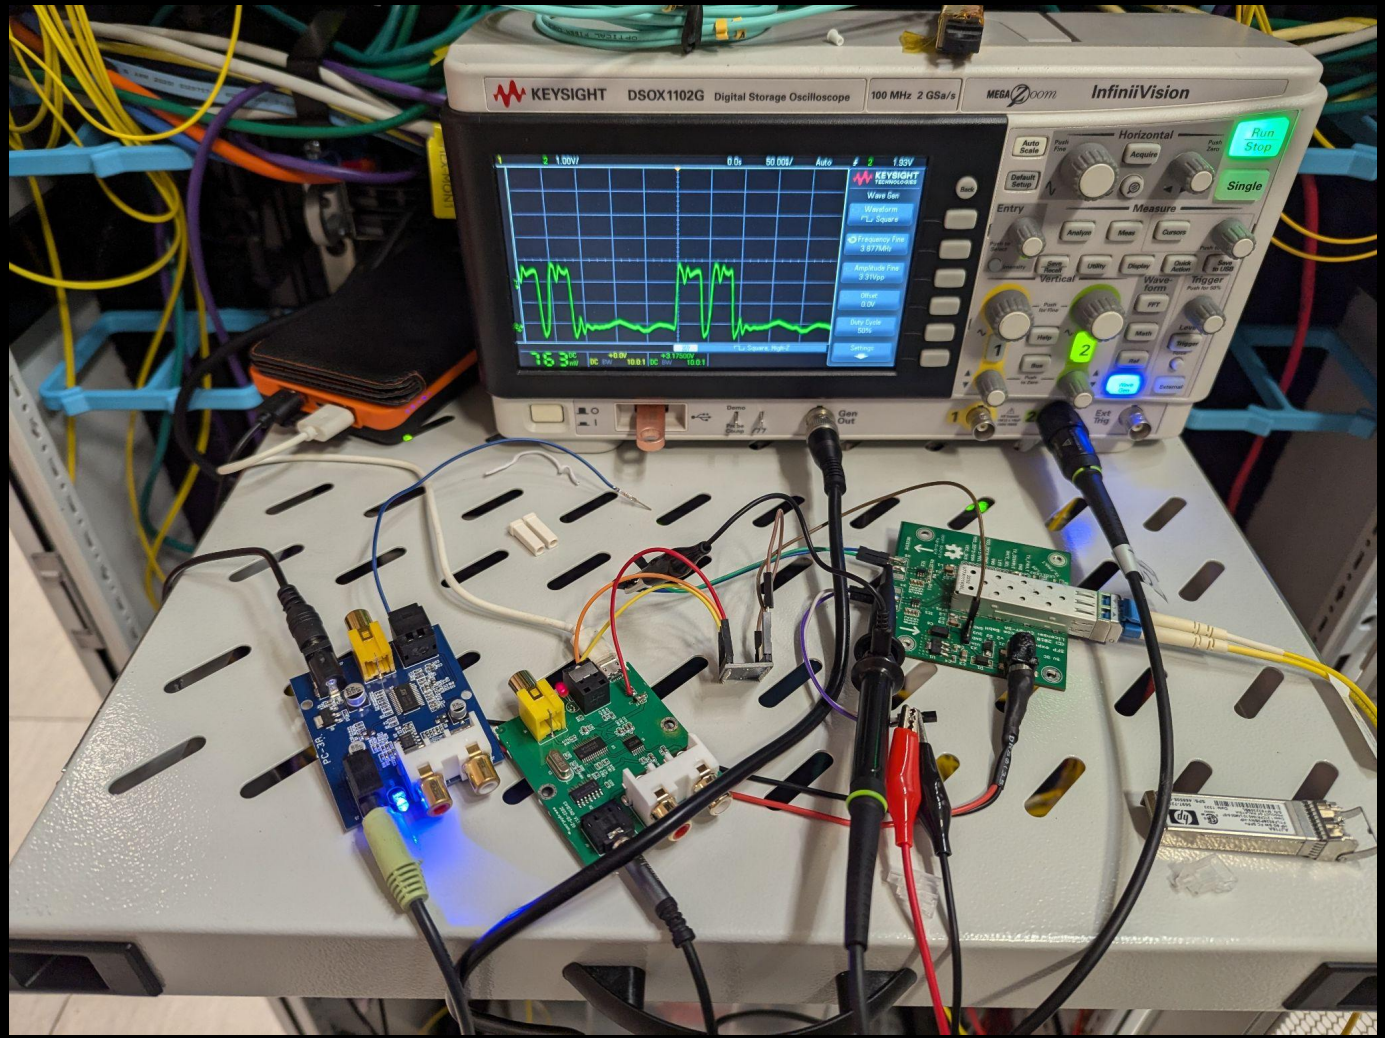
KEYSIGHT
DSOX1102G
Digital Storage Oscilloscope
100 MHz 2 GSa/s
InfiniiVision
Horizontal
Run
Stop
Single
Entry
Measure
Vertical
Trigger

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 the DWDM optic has a much higher bar for input/output speed!

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

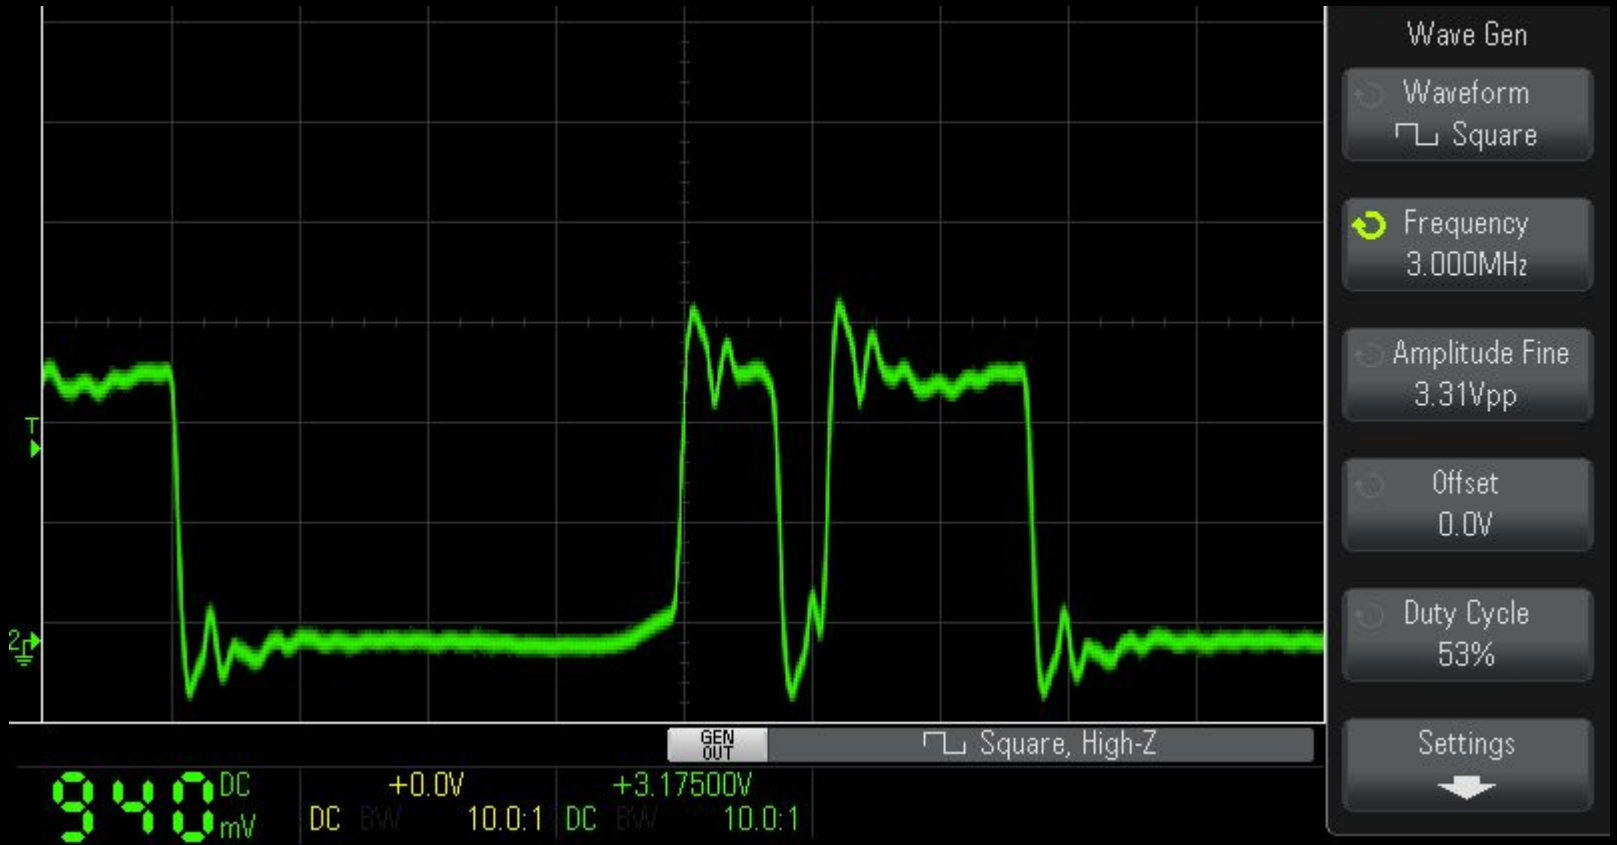

3mhz min?!

- 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

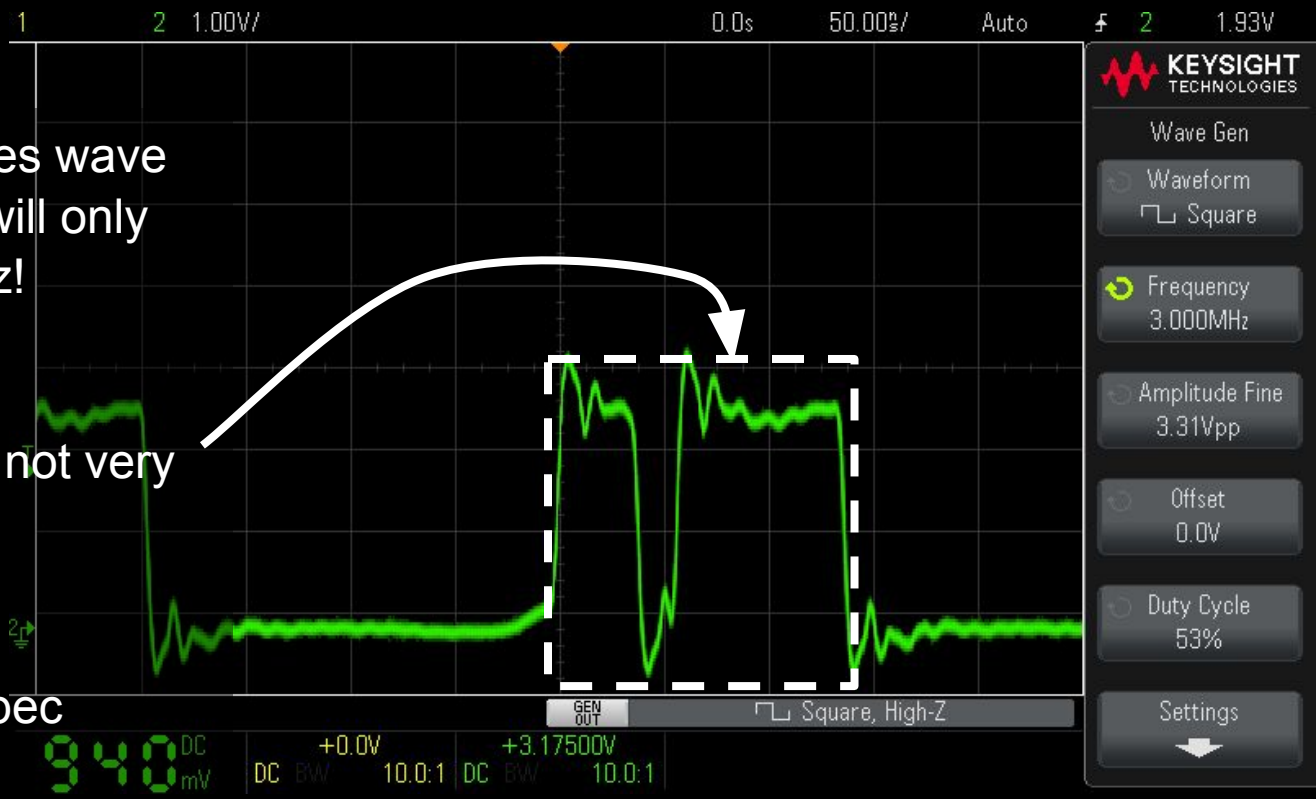

Stable at 6.5MHz (ish)

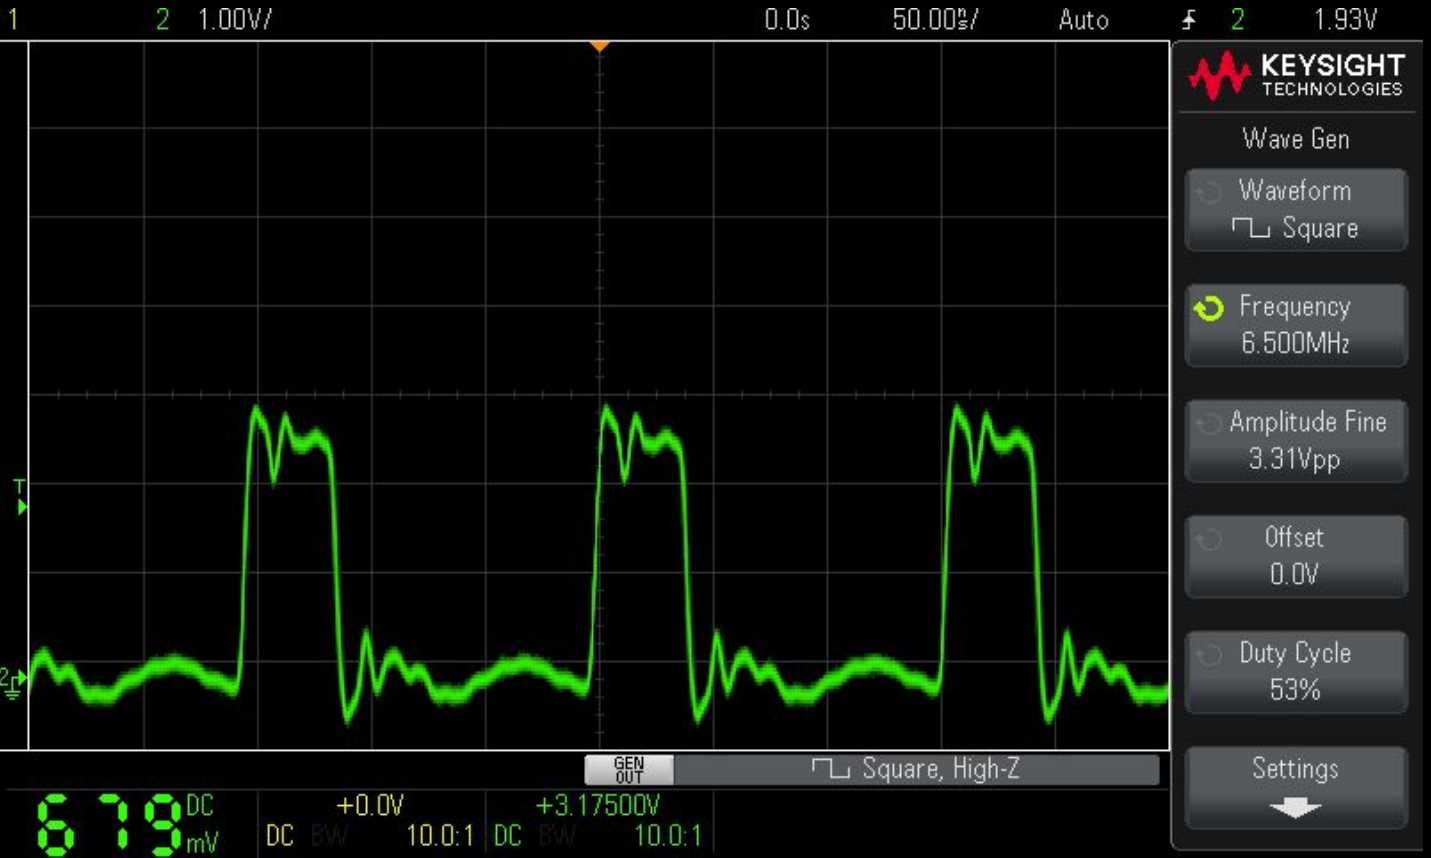

Our TOSLINK signal is 2.2mhz, but the square wave only really becomes square at 6.5mhz

This optic is just picky, sadly it's all I had for DWDM C59

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 ***suspect*** that most of this is some kind of noise protection on the signal, however that is pure speculation
  - LX/SR/CWDM optics start working at 150KHz, so unclear why so much for DWDM
- This optic was also a 10G DWDM, maybe a 1G optic would work?
  - I ran out of time to test this for now

- Maybe one of the wonderful optic module vendors here can speculate? :) ?

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 likely burn the laser out)
  - I hear this is sorta how "RF over Fiber" works

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 likely burn the laser out)
  - I hear this is sorta how "RF over Fiber" works

- It is tempting to attach a "dialup" modem to both sides
  - This would create the greatest waste of a 100 GHz optical channel in terms of 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 fancy 10G stuff seems to start at 6 MHz
- You can even send them very far
- There is likely not a good reason to do this however, 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 work arounds to problems
- I know far more about optical transmission, and general inner optics workings now
  - Maybe you do too!

Questions?

Shy? nanog92@benjojo.co.uk

or find me in the hallway!