The Curious Case Of The
VTRx RECEIVER FAILURES

EP-ESE Electronics Seminars

18 January 2022

The Versatile Link

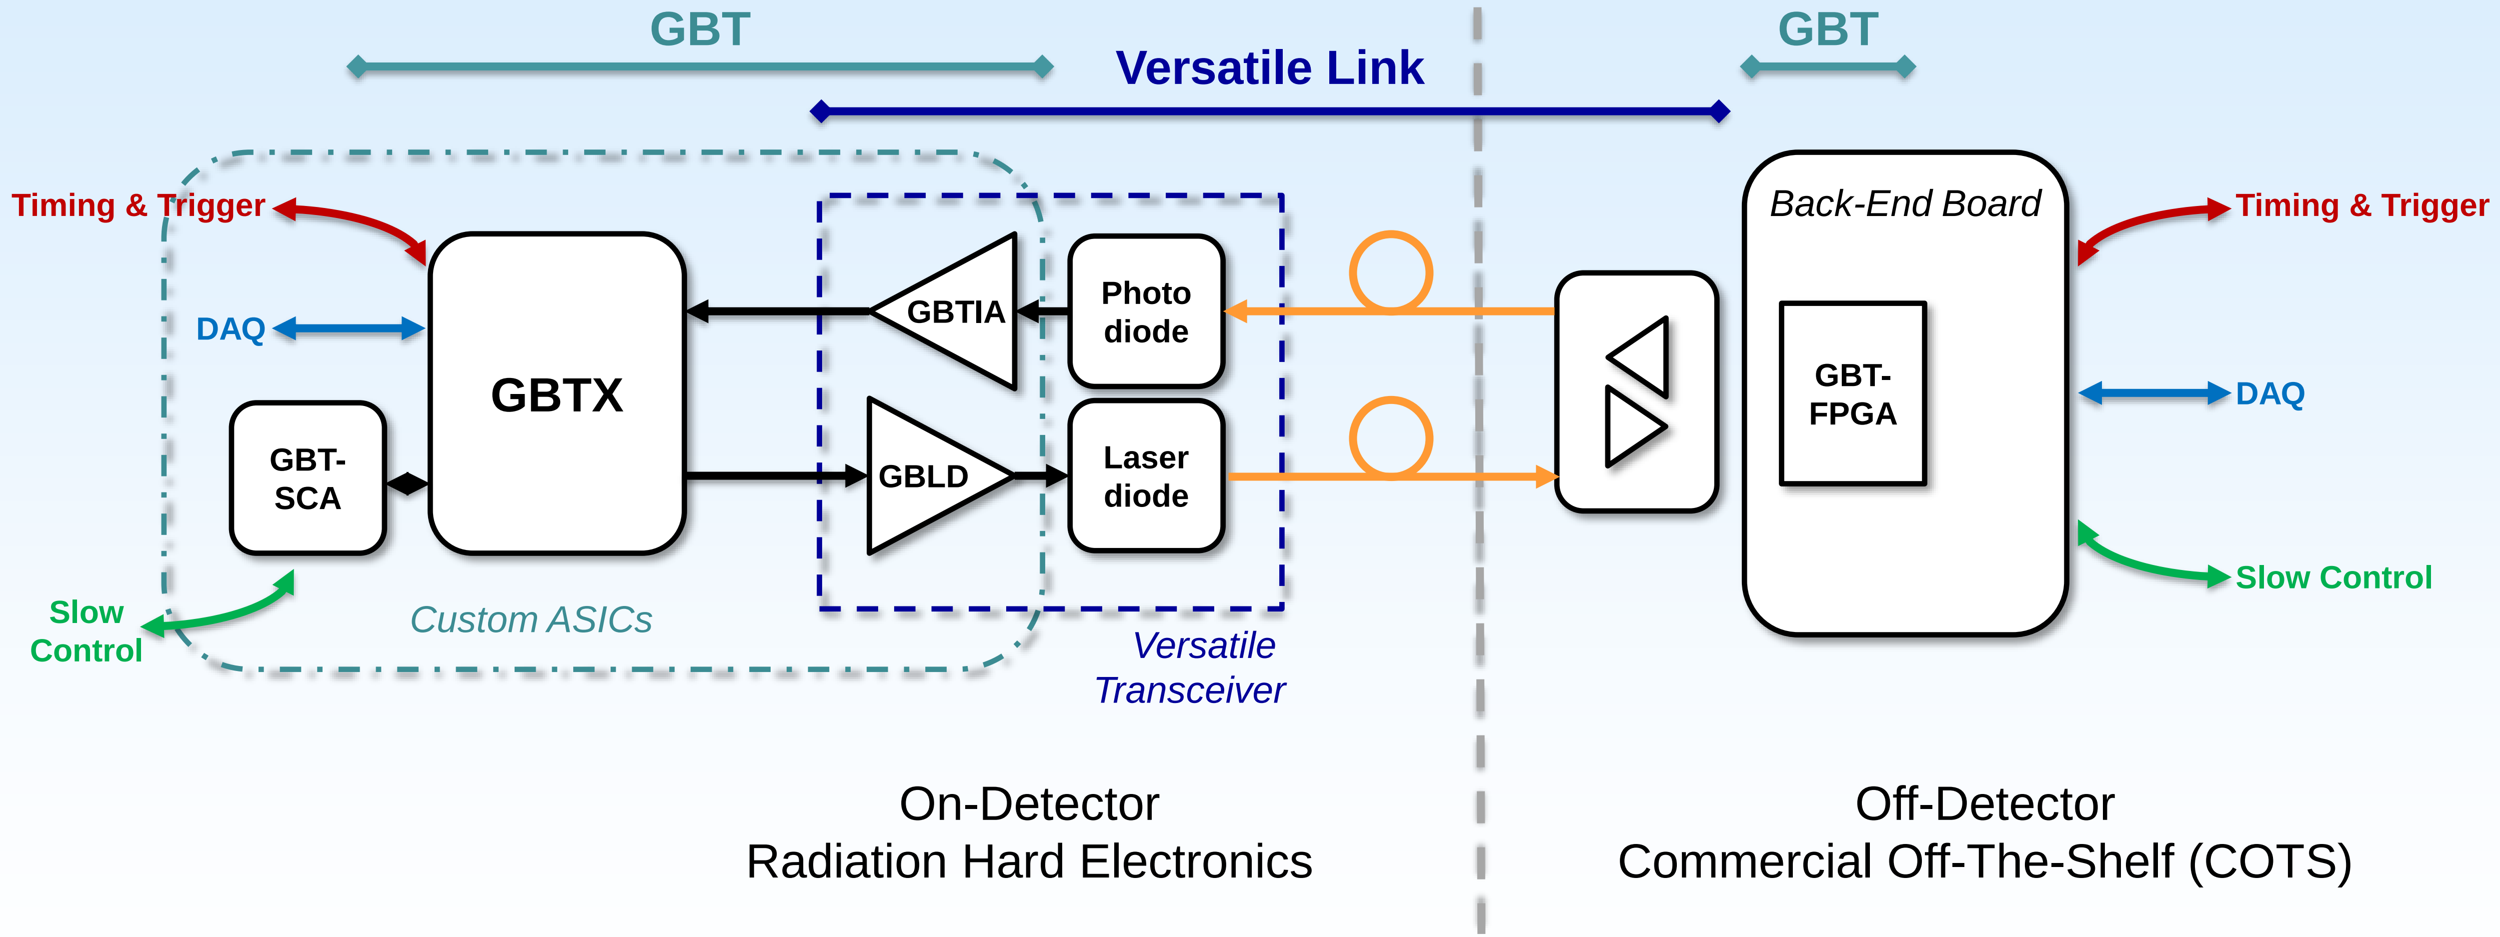

The Versatile Transceiver (VTRx)

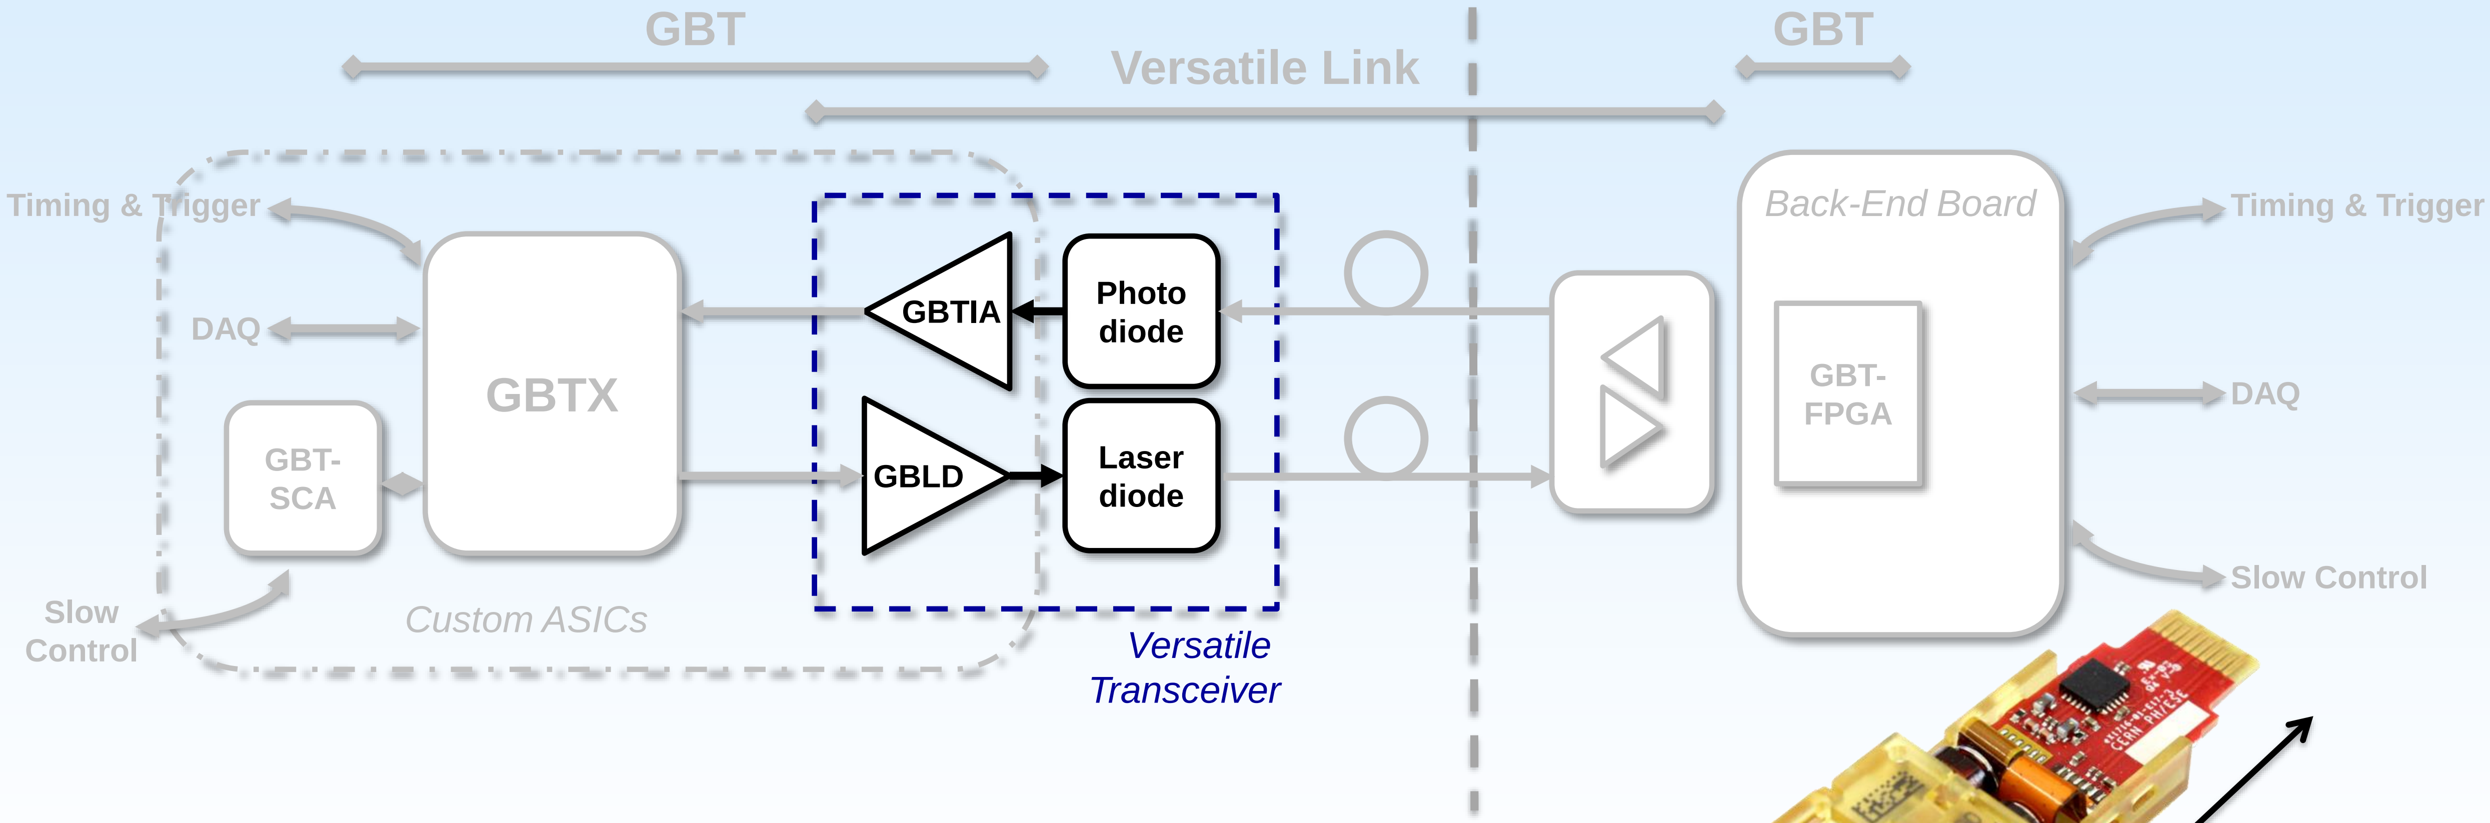

- Optoelectronic transceiver operating at 4.8 Gb/s
- Sources radiation-hard ASICs from GBT project

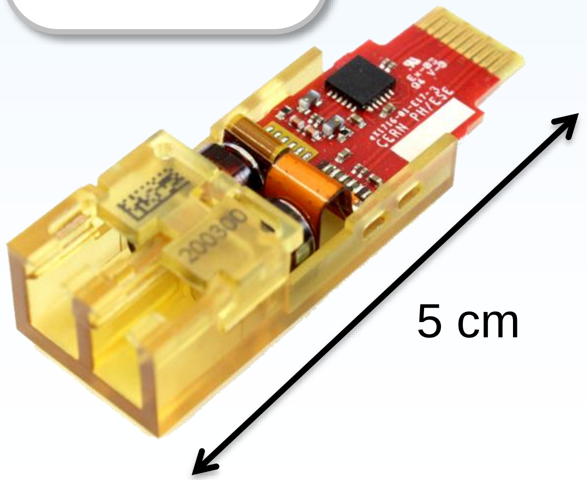

The Versatile Transceiver (VTRx)

- In total over 18,000 VTRx transceivers were produced in 2016-2019
  - Multimode 850nm and single-mode 1310nm versions
  - Twin transmitter VTTx variant
- Users include all four large CERN experiments:
  - ATLAS: 1700 × VTRx
  - ALICE: 6000 × VTRx
  - CMS: 2600 × VTRx
  - LHCb: 2200 × VTRx
- External users, such as
  - Mu2e at Fermilab
  - CBM and Panda at GSI
  - ITER
- Targeting the phase 1 upgrades during the LS2
  - to be completed Autumn 2021

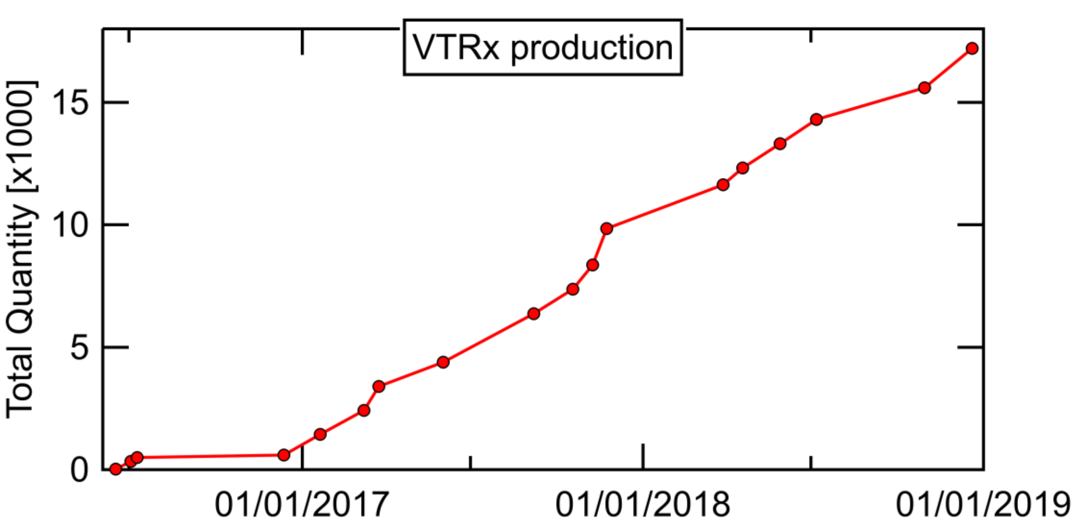

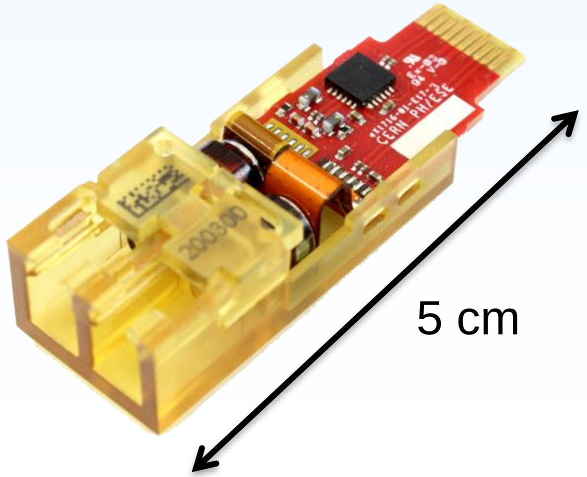

The Versatile Transceiver (VTRx)

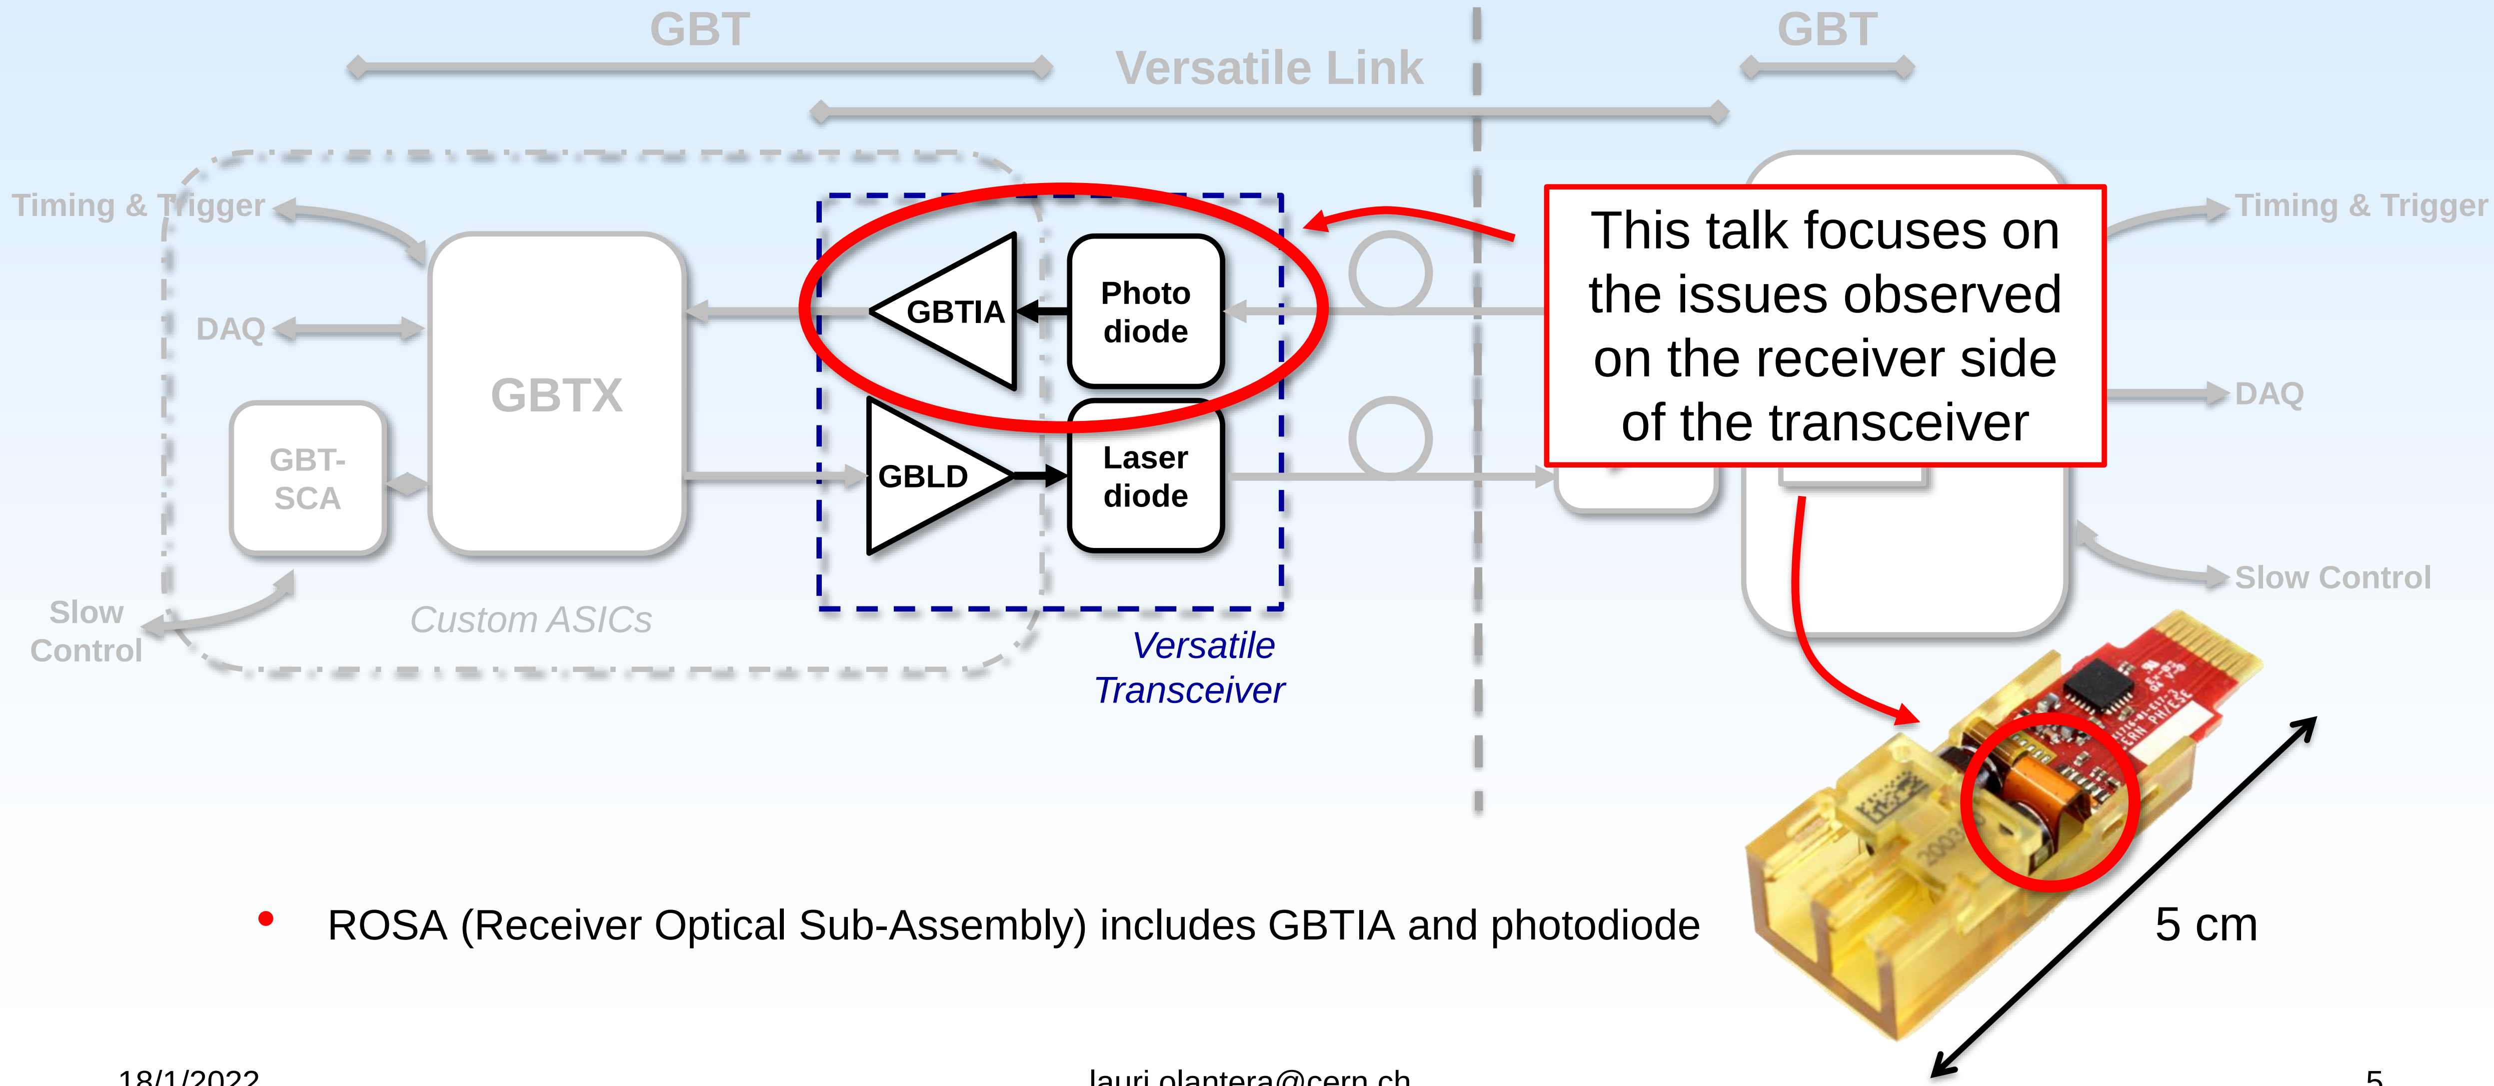

- ROSA (Receiver Optical Sub-Assembly) includes GBTIA and photodiode

RSSI and basic receiver parameters

- The RSSI (*Received Signal Strength Indicator*) output acts as a current sink when connected to the 2.5 V supply with a pull-up resistor
- A current mirror in the GBTIA generates an RSSI current proportional to the current generated by the photodiode
  - If the pull-up resistor has been chosen correctly, the RSSI is a good indicator for the *average* photocurrent

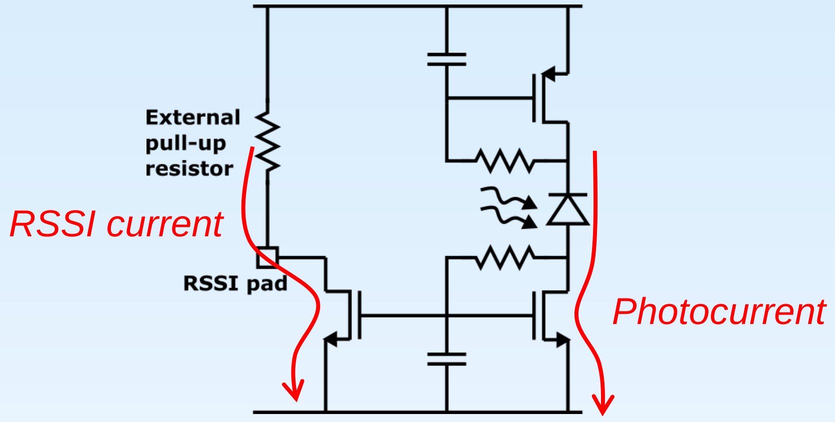

$$\textbf{\textit{Responsivity}} = \frac{\text{photocurrent}}{\text{optical input power}}$$

Sensitivity = minimum optical modulation amplitude that results in error-free* transmission

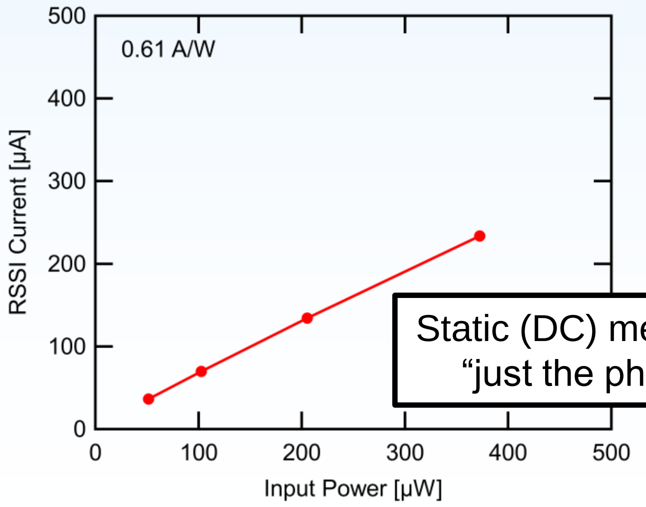

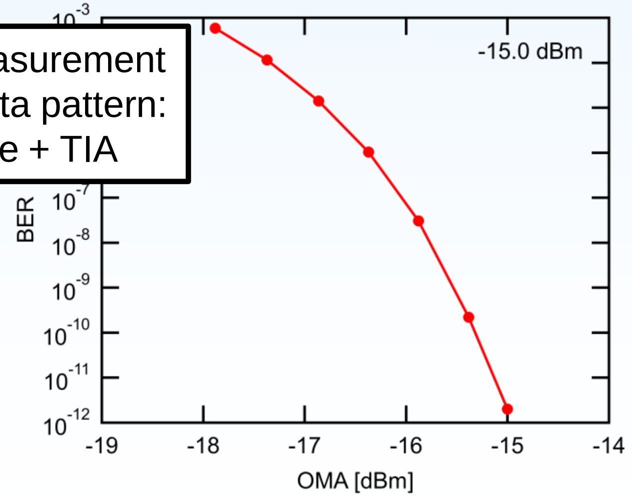

* error-free means a defined bit error ratio, which is small enough for the application in question, e.g., 1e-12

Are you sitting comfortably? Then I'll begin…

The first signs

- CMS HCAL observed communication issues in 2018-2019
  - “Two fallen modules”
  - Unstable RSSI reading and loss of communication
- Issue was difficult to reproduce
  - Testing a failing module from CMS in our lab was not helpful
    - We did not observe any issues
  - CMS, at least initially, had similar problems
    - A module failing in experiment did not fail anymore when moved to another test stand
- HCAL concluded that added cooling element helps
  - This was a strange results for us at the time, because the VTRx had been tested across a wide temperature range
    - Receiver almost completely temperature insensitive
  - Anyways, when it works it works!

See presentation by Grace Cummings:
https://indico.cern.ch/event/1019078/contributions/4444259/attachments/2312222/

The second signs

Versatile Link

- ALICE was the second experiment that brought VTRx receiver issues to our attention

- We got samples returned to us from ALICE ITS in the spring 2020 and tested them after CERN opened again in summer 2020 and this time there were clearly observable issues
  - RSSI connection was bad and resoldering the flex solved most of the issues

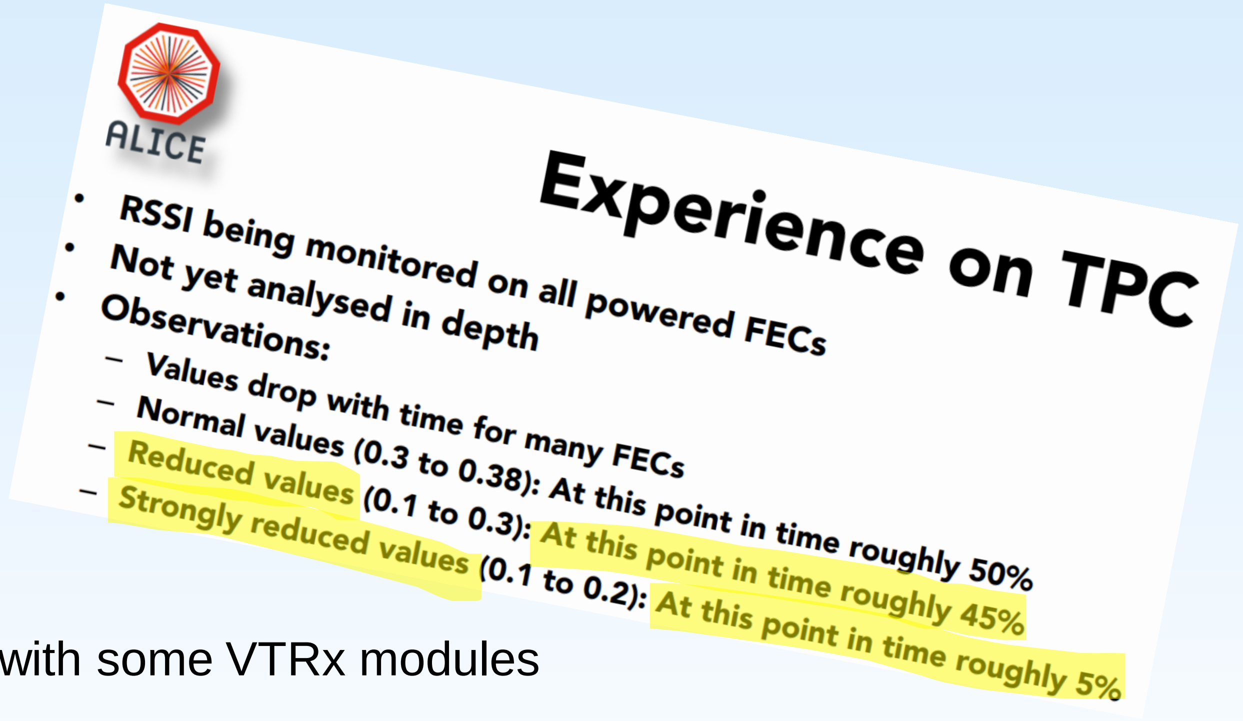

- However, ALICE ITS and TPC kept facing issues with some VTRx modules and continued their own investigations
  - Conclusion from ALICE ITS was that there is contamination on the ROSA ball lens

Now we had easily observable failure modes!

ROSA ball lens contamination

- User observed contamination the surface of the ROSA ball lens
- Concern was that the glue that holds the ball lens in place or some other residue could soften and/or move when the device is operated and therefore cause varying attenuation
- None of the returned samples had enough contaminant to explain observed drifts in the RSSI
- All bad looking samples that had caused problems in the experiment had also been wet cleaned by the user
  - Impossible to say what was residue of the cleaning process and what was real contamination

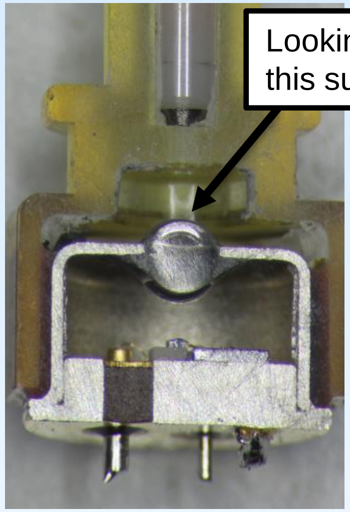

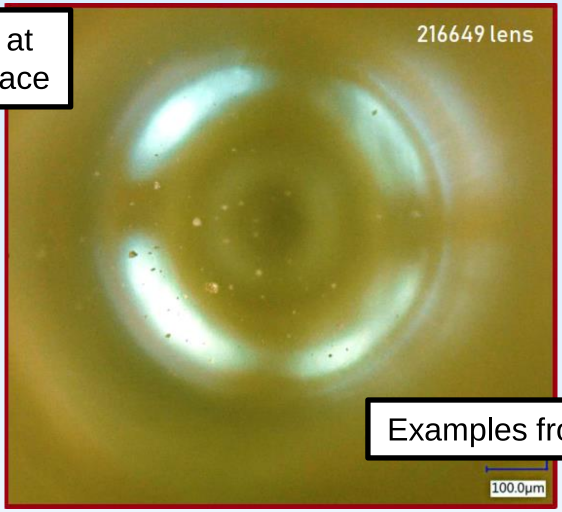

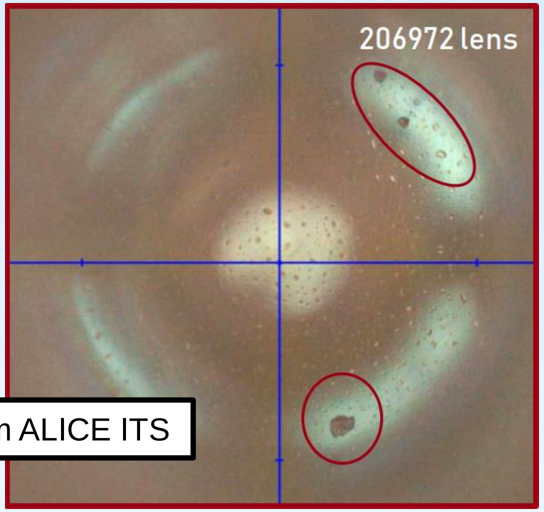

Examples from ALICE ITS

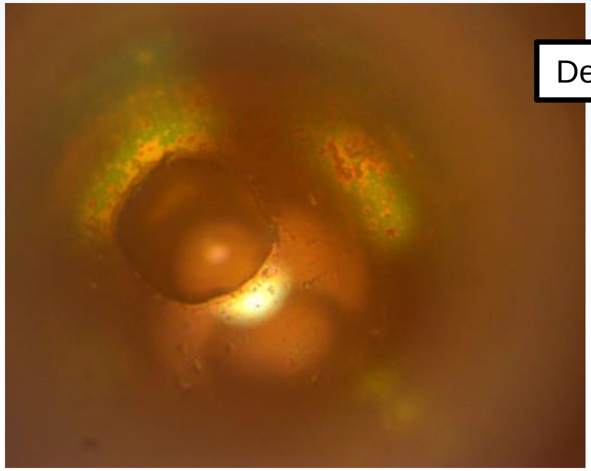

Device returned from ALICE ITS

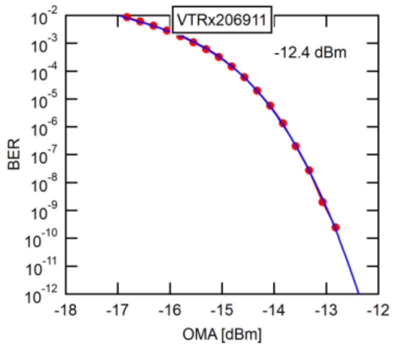

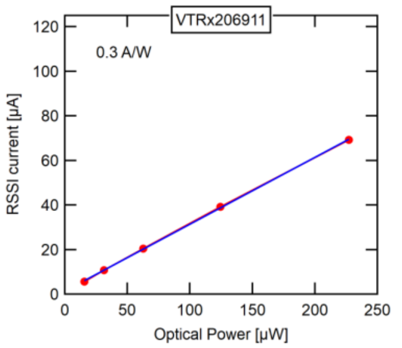

- Dozens of modules imaged and no contamination that could cause significant signal attenuation observed
  - Ball lens receives the expanded beam from the fiber
- 64 ROSAs (4 from each production lot) were imaged, kept at 85C for 120 hours, and imaged again
  - No changes in contamination observed

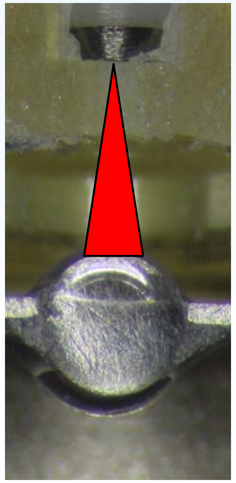

Not enough evidence for ball lens contamination theory

Flex trace issue

- A weakness in the flex design was identified
  - Cracking possible at the edge of coverlay
- 140 VTRxs tested in thermal cycling
  - -40 to +85°C, 500 cycles
  - 12/140 (8.6%) showed issues at some point in the testing
  - 11/12 failures directly due to flex
- New flex designs where the discovered weaknesses fixed
  - 160 VTRXs (3x40 with new flex versions and 40 with old) samples through same temperature cycling
  - No failures
  - New version easier to assemble
    → Preferred for repairs

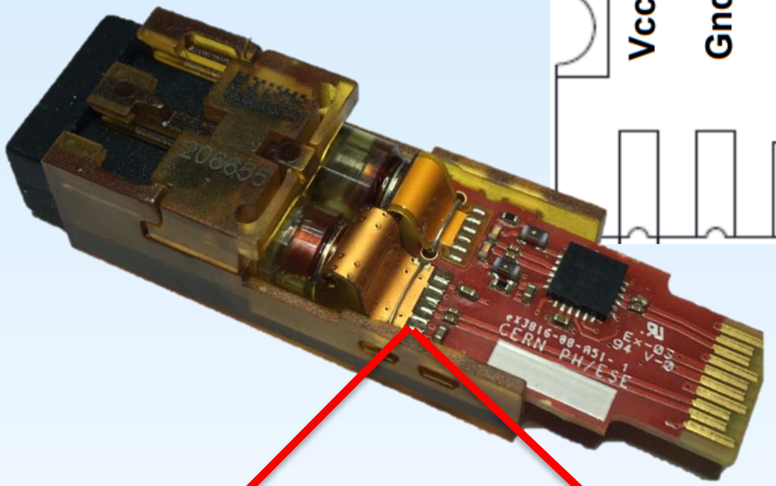

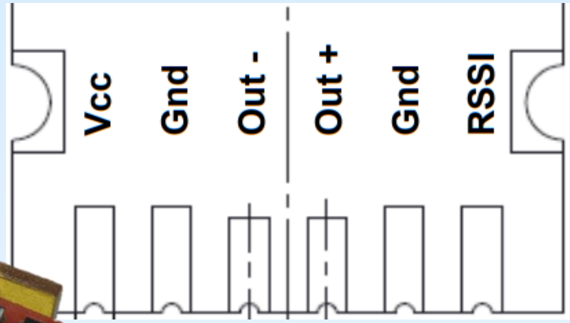

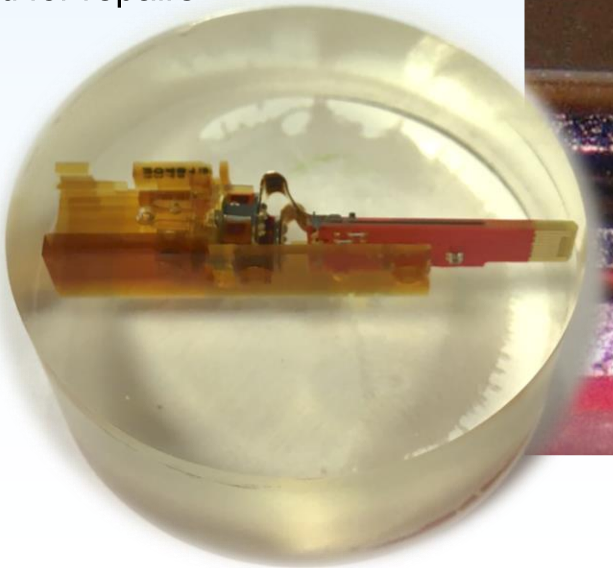

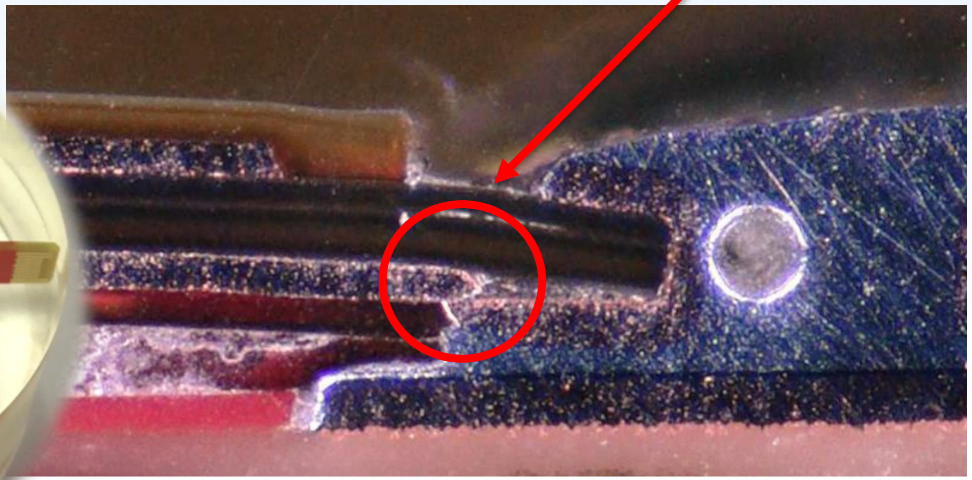

Cross-sectioning

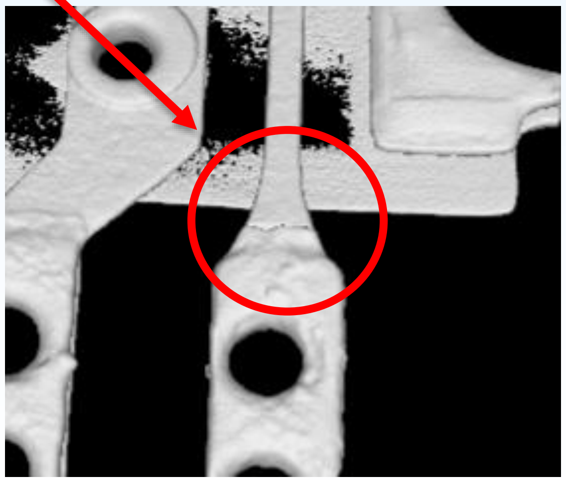

X-ray Tomography (μCT)

lauri.olantera@cern.ch

Electric connection problem?

- Flex cracks and partial cracks could change the supply and RSSI voltage levels

- Very large RSSI voltage drops required for significant changes in RSSI currents
  - This would be only monitoring issue, users have reported loss of communication

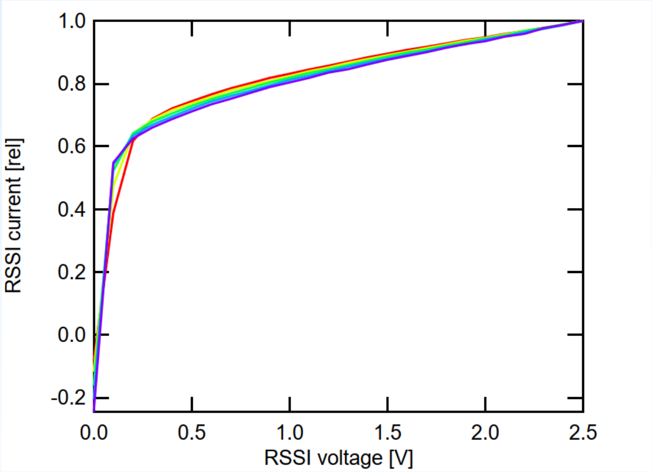

- Problems with supply voltage could cause a loss of link, but reduced supply voltage has very little effect on the RSSI current and when reduced enough RSSI current starts to increase

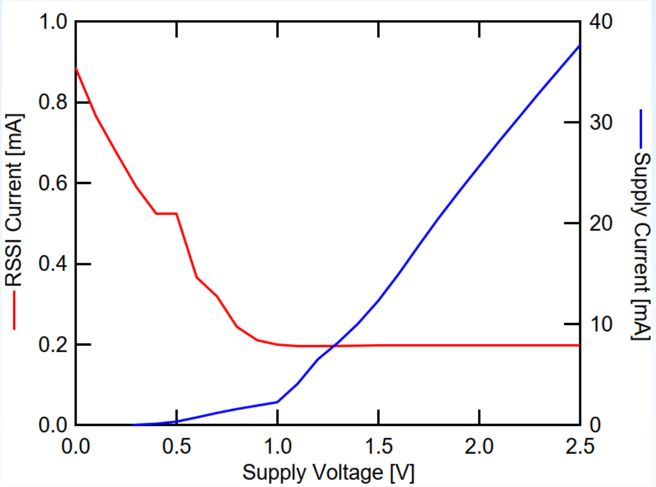

- Users who could monitor RSSI had not observed large jumps in RSSI
- Supply current trace as likely to crack as the RSSI
  → maybe flex issues not that common after all…

No evidence to support flex cracks to be the main cause of the observed problem

GBTIA?

- Issue known to occur at elevated temperatures
  - GBTIA overheating?
- Schematic-level simulation over temperature
  - ASIC works well up to 200°C junction temperature
  - No significant effect on high-speed performance
  - No influence on RSSI output for nominal voltage
  - At lower RSSI pin voltage, the current increases with temperature so cannot explain RSSI drop
- Lab tests pushing GBTIA even out-of-spec conditions did not produce results resembling anything observed by the users

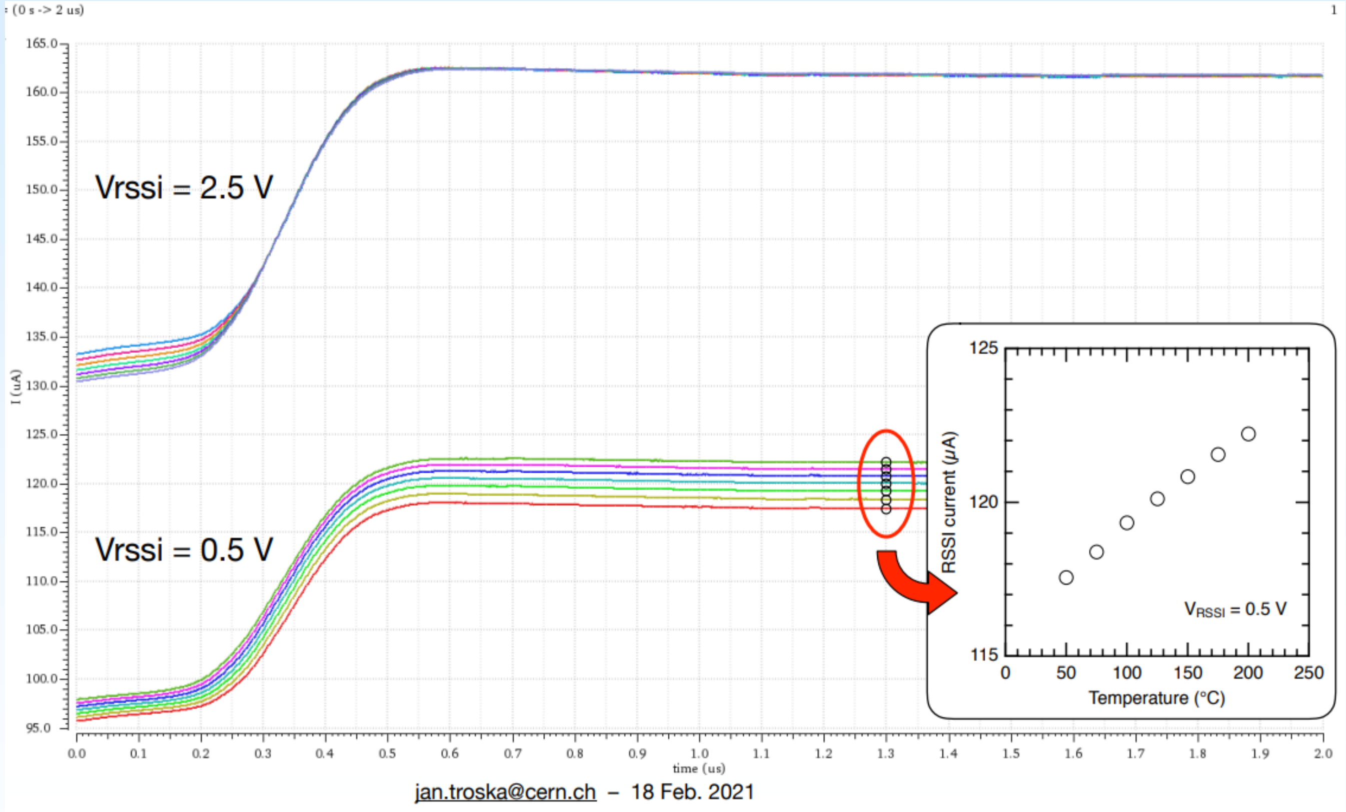

After investigations, there was no evidence to support GBTIA theory

Large-scale setup: Setup 1

- **Reminder:** at this point we were still struggling to reproduce the problem in our lab
- We had started to prepare large scale setups
  - If enough samples are running, we must observe this issue at some point!
- 50 devices running continuously inside a temperature chamber
  - We could control the temperature and try to provoke the failures
- And this happened without even turning on the chamber:

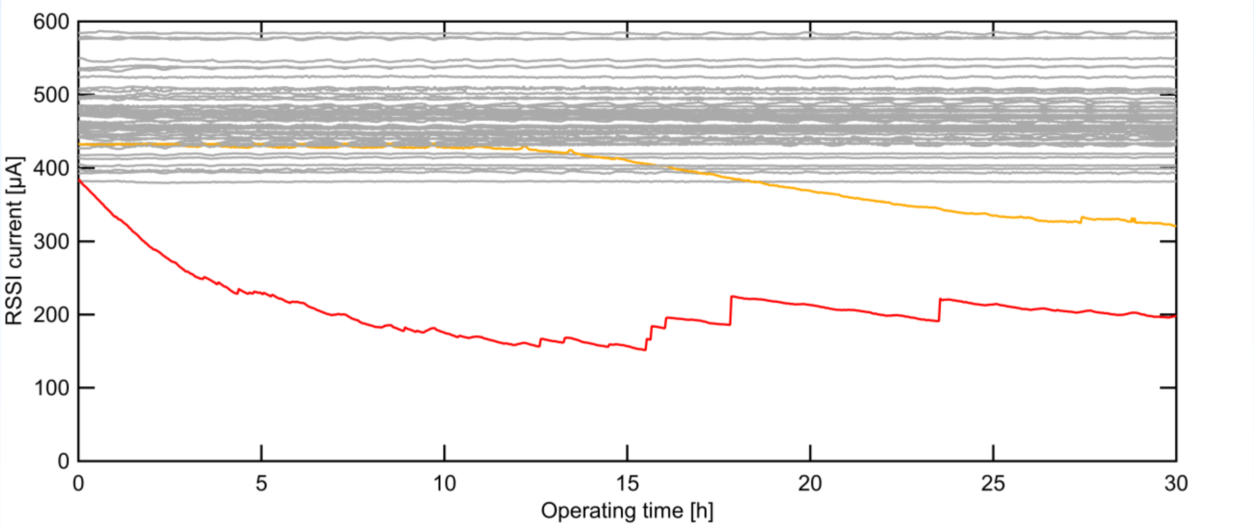

- This setup was left running to collect failure statistic and another similar setup was right away prepared…

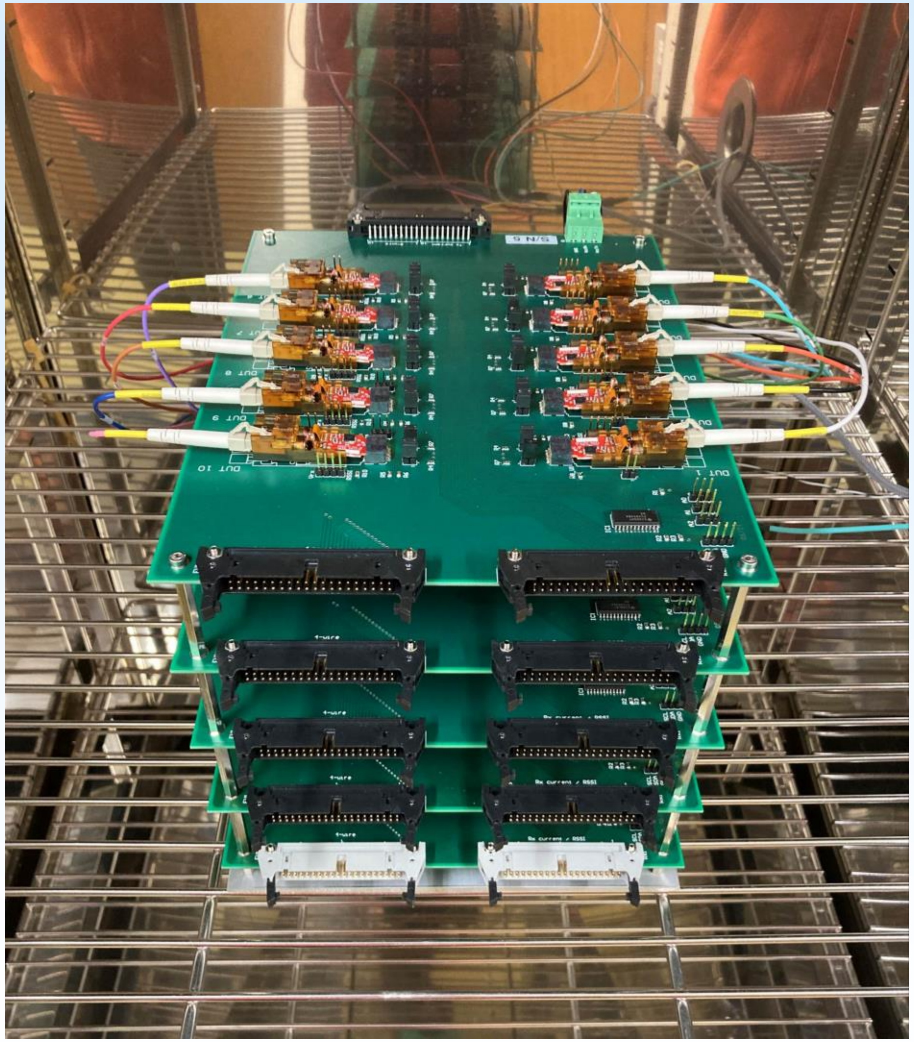

Second large-scale setup: Setup 2

Versatile Link

- Second setup was constructed in a similar way but with an additional high-speed test slot that allowed full functional test on one VTRx
- First drifting module was moved to this slot:

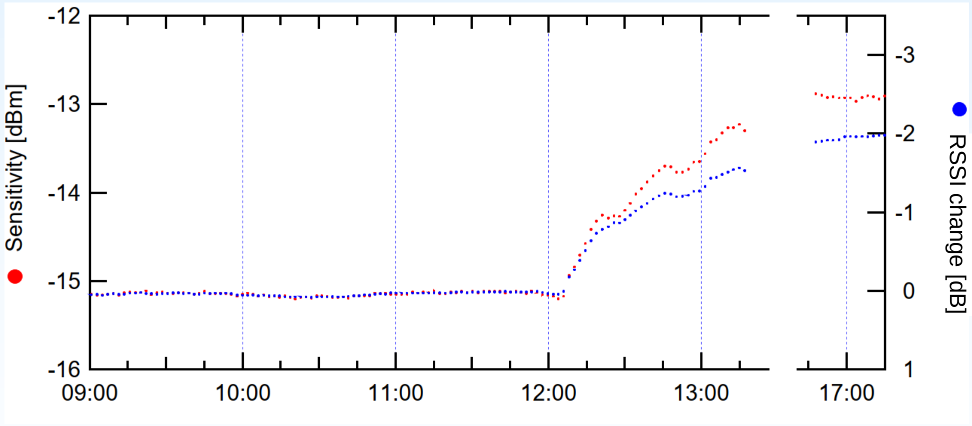

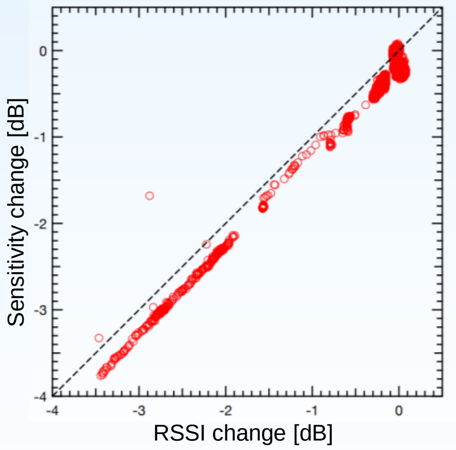

- The RSSI drift shows almost identically in the sensitivity!
- We are losing the light input somehow

While Setup 1 was collecting statistics, that did not look good…

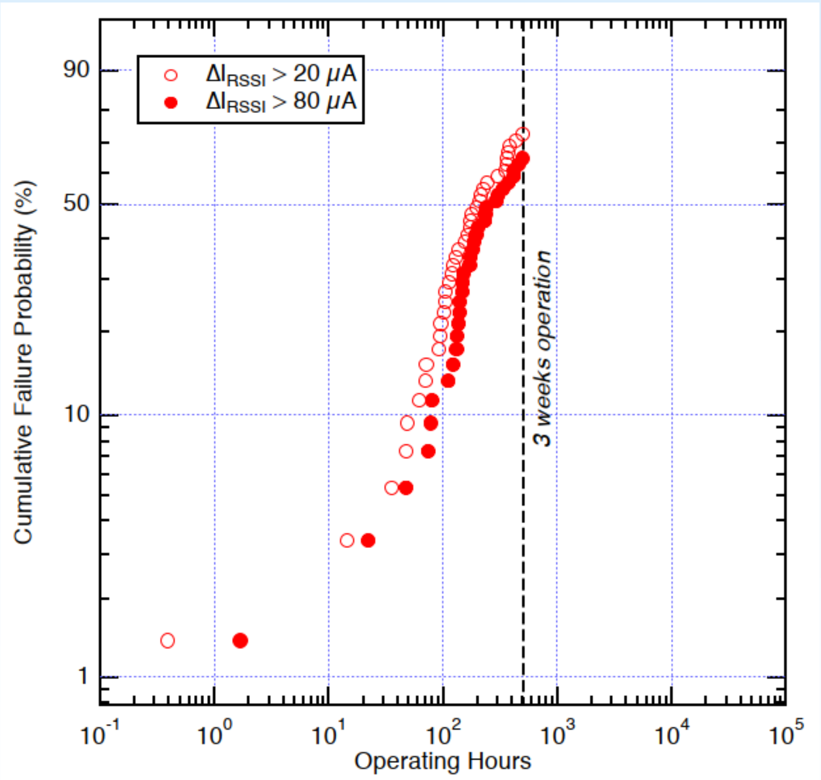

In right conditions it looks like the whole (at least almost) population can be affected!

Why are we losing the light?

- Something on the light path? Misalignment?
- Added mystery: seemingly random recoveries, sometimes partial, sometimes complete!?
- If we have introduced some contamination inside the ROSA during the assembly, maybe it could move when heated, but still how would the contaminant just suddenly disappear?
- Maybe high temperature softens something in the assembly and allows small movements?
  - If you try hard enough, you could imagine a process where external stress causes slow movements and even sudden changes when some kind of relaxation occurs

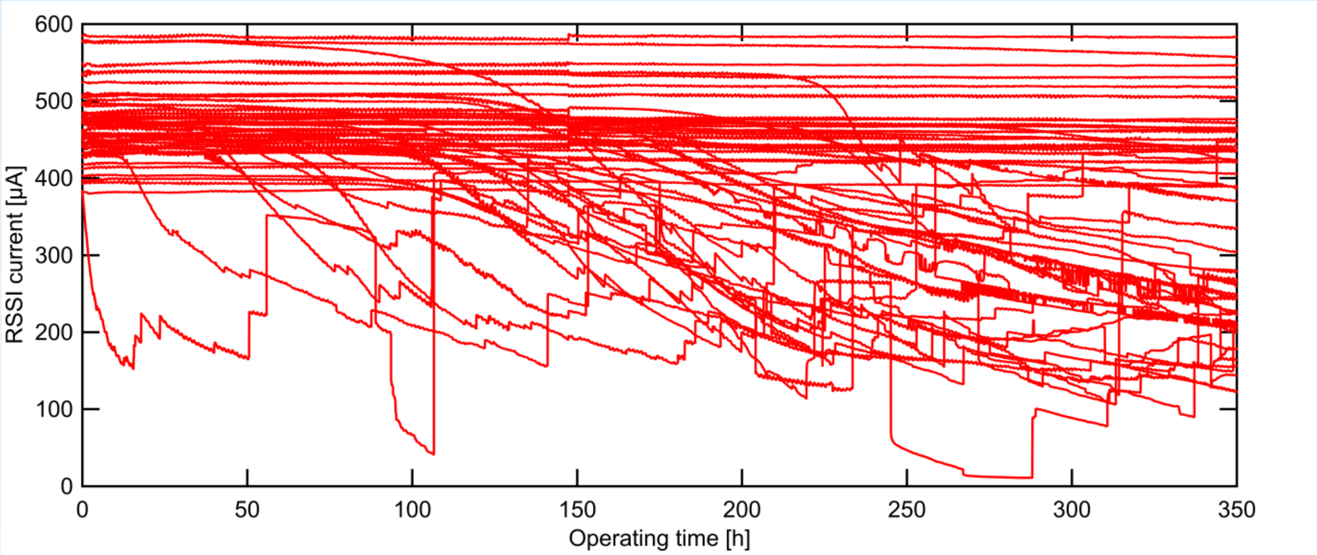

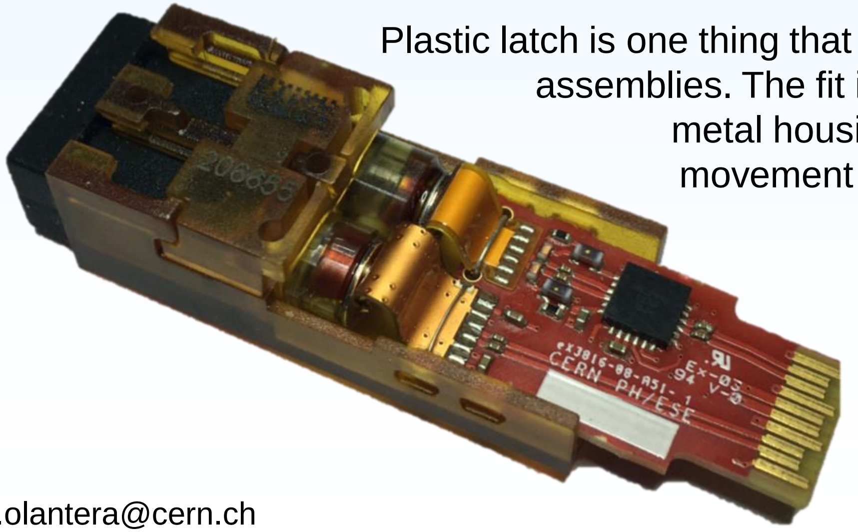

Plastic latch is one thing that differs from a commercial assemblies. The fit is not as tight as in typical metal housings and allows little more movement between the components

Misalignment

- With simple optics simulation one can estimate that some tenths of a degree movement between the plastic receptacle (1) and metal TO-can (2) could cause observed losses
- This was then tested in two ways
  1. Movement is small but could be detected under a microscope: VTRx and a microscope inside insulated box and recording images and RSSI
     - No movement observed even when RSSI jumped almost 50%

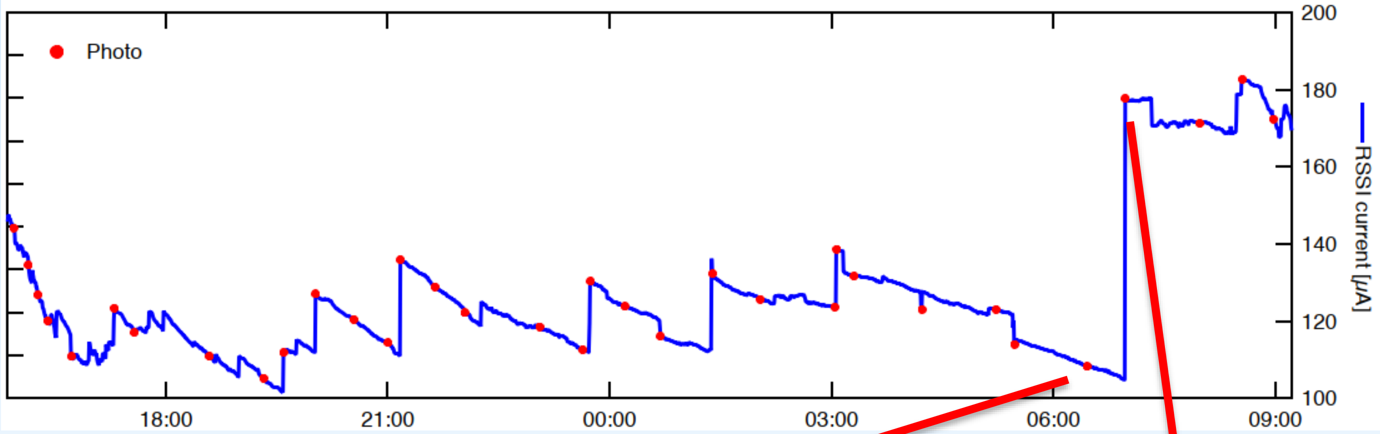

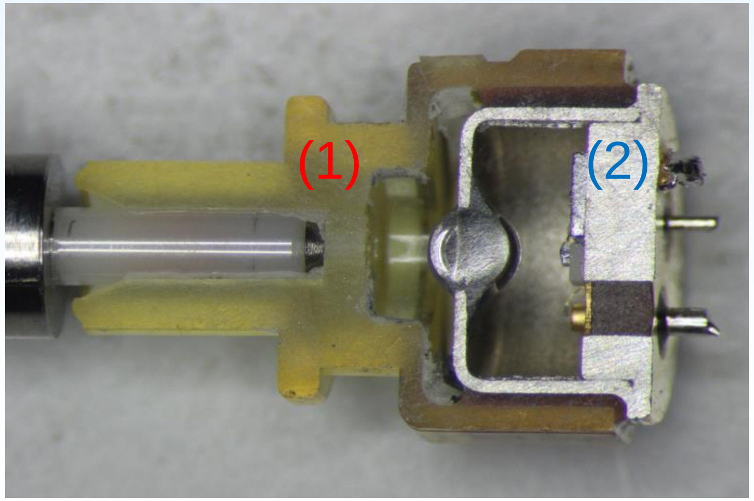

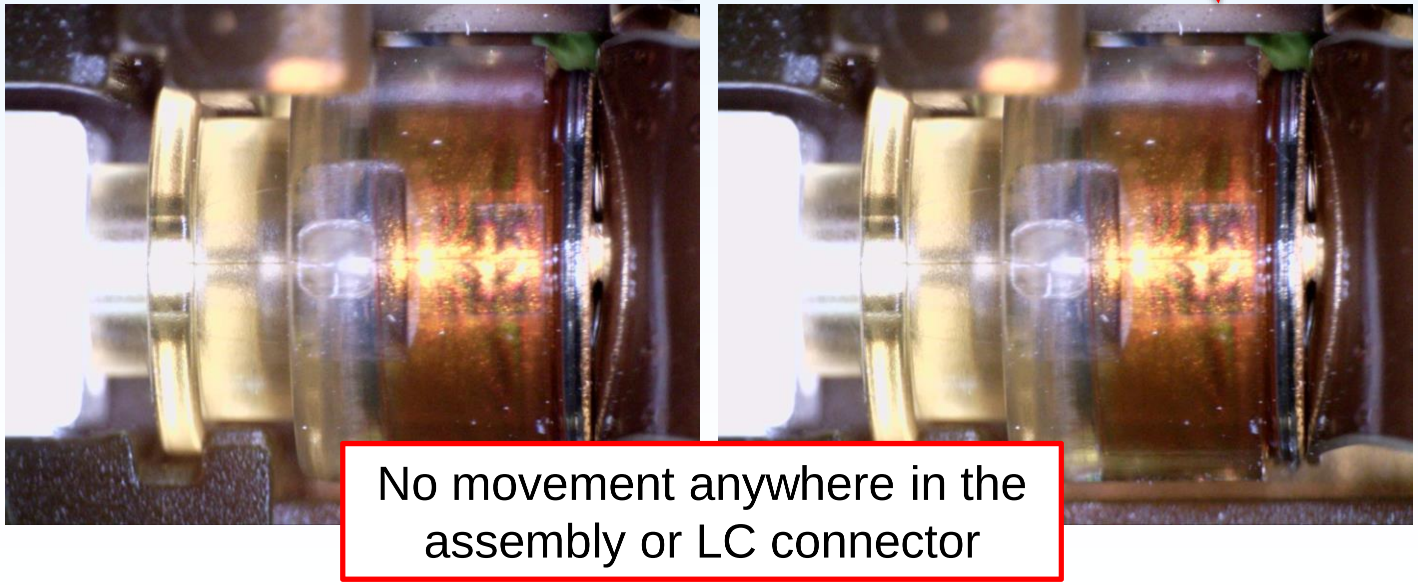

Misalignment

- With simple optics simulation one can estimate that some tenths of a degree movement between the plastic receptacle (1) and metal TO-can (2) could cause observed losses
- This was then tested in two ways
  1. Movement is small but could be detected under a microscope: VTRx and a microscope inside insulated box and recording images and RSSI
     - No movement observed even when RSSI jumped almost 50%
  2. Weights pulling the fiber in different directions either using VTRx or bare ROSA while whole assembly heated in temperature chamber
     - We definitely exceeded realistic mechanical forces that a ROSA would be subject to in a real application
     - And still no changes in RSSI observed

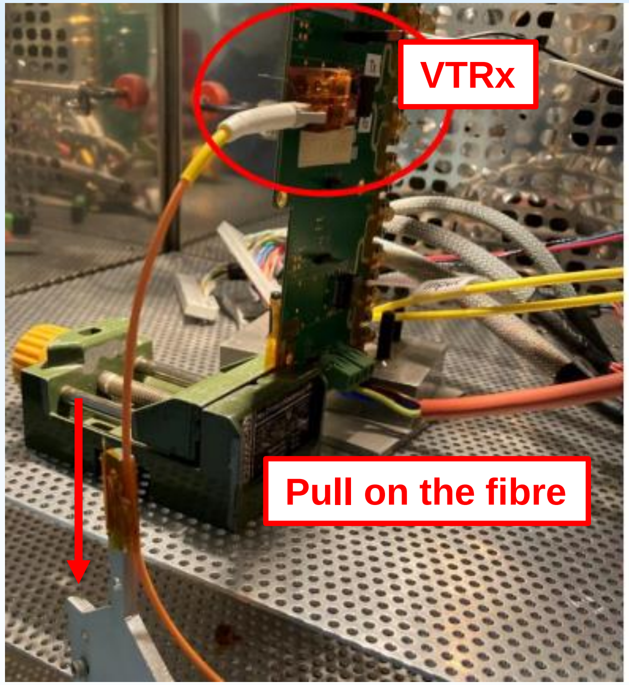

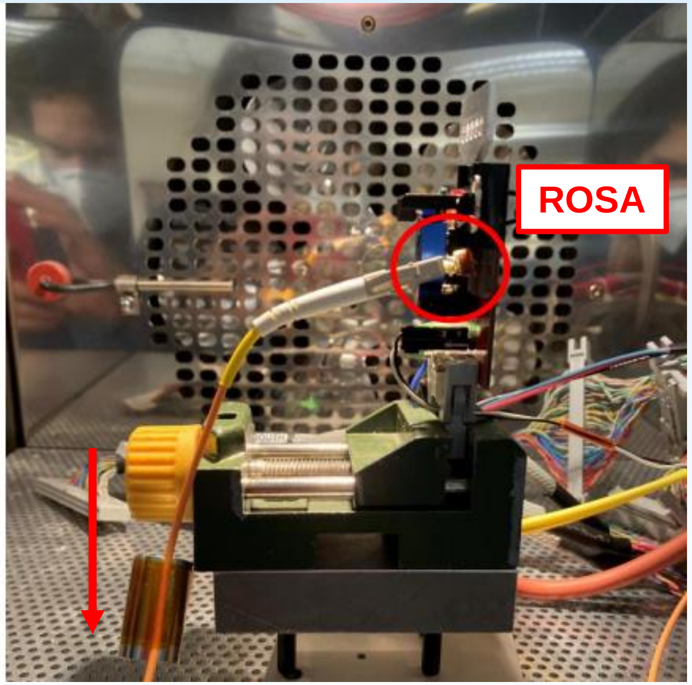

No, misalignment theory does not work

Back to the drawing board, what do we know so far?

- Contamination inside the ROSA on the ball lens: NO
- Flex problem: NO (well, it is an issue in a fraction of the devices, but not the cause of the main problem)
- GBTIA issue: NO
- Misalignment: NO
- **Temperature** plays an important role: YES
  - Requires elevated temperature and added cooling helps
- Let's have a closer look on the temperature performance, and confirm this fact:

Back to the drawing board, what do we know so far?

- Contamination inside the ROSA on the ball lens: NO
- Flex problem: NO (well, it is an issue in a fraction of the devices, but not the cause of the main problem)
- GBTIA issue: NO
- Misalignment: NO
- **Temperature** plays an important role: YES
  - Requires elevated temperature and added cooling helps
- Let's have a closer look on the temperature performance, and confirm this fact:

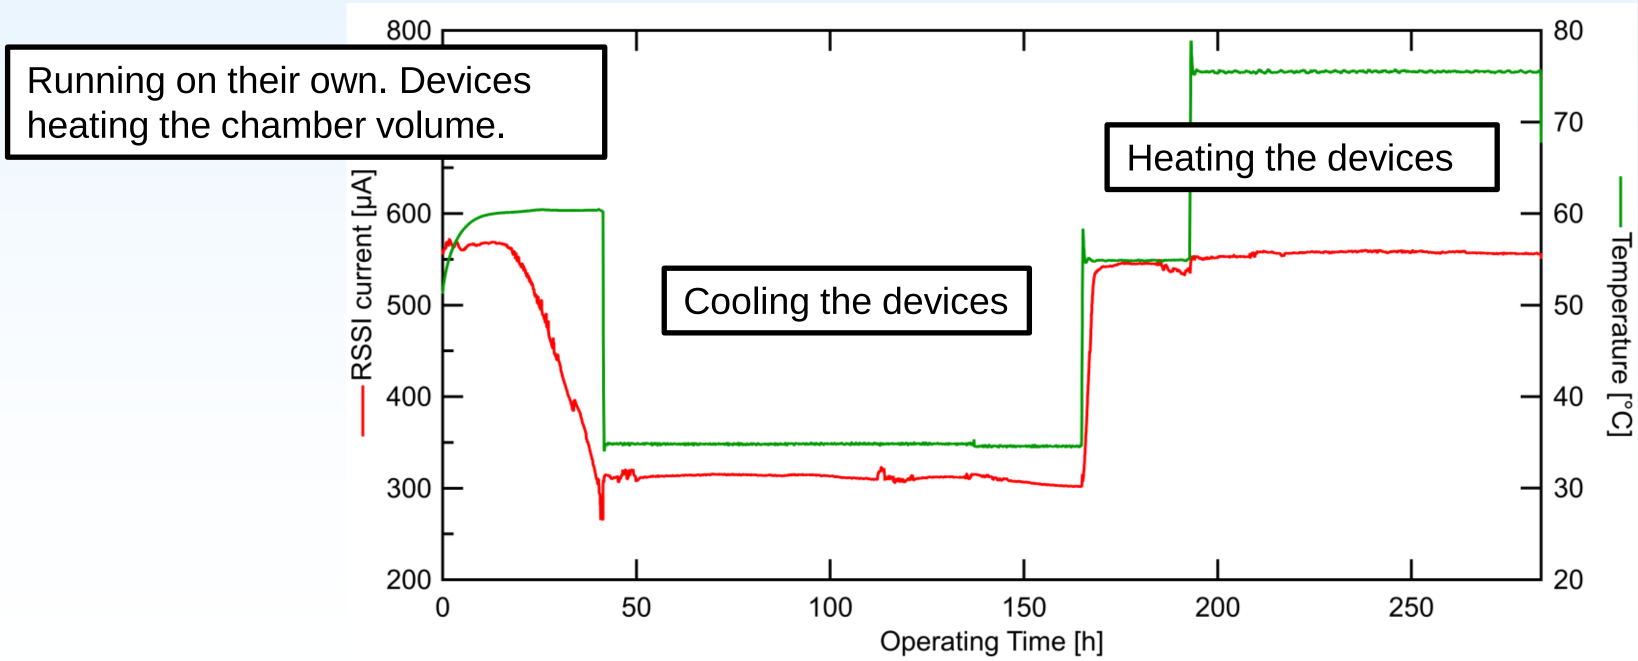

Wait, what? Cooling stops the RSSI drift, but heating makes it even better??

Break in the case!

- A new idea came from CMS HCAL investigations who observed contaminate on a fiber end-face
- This observation was very quick to confirm using the drifting samples in our setups
- Contaminant is volatile and evaporates at room temperature in minutes/hours depending on quantity and conditions

Now we know what causes the optical losses, but what is it? How and where is it coming from?

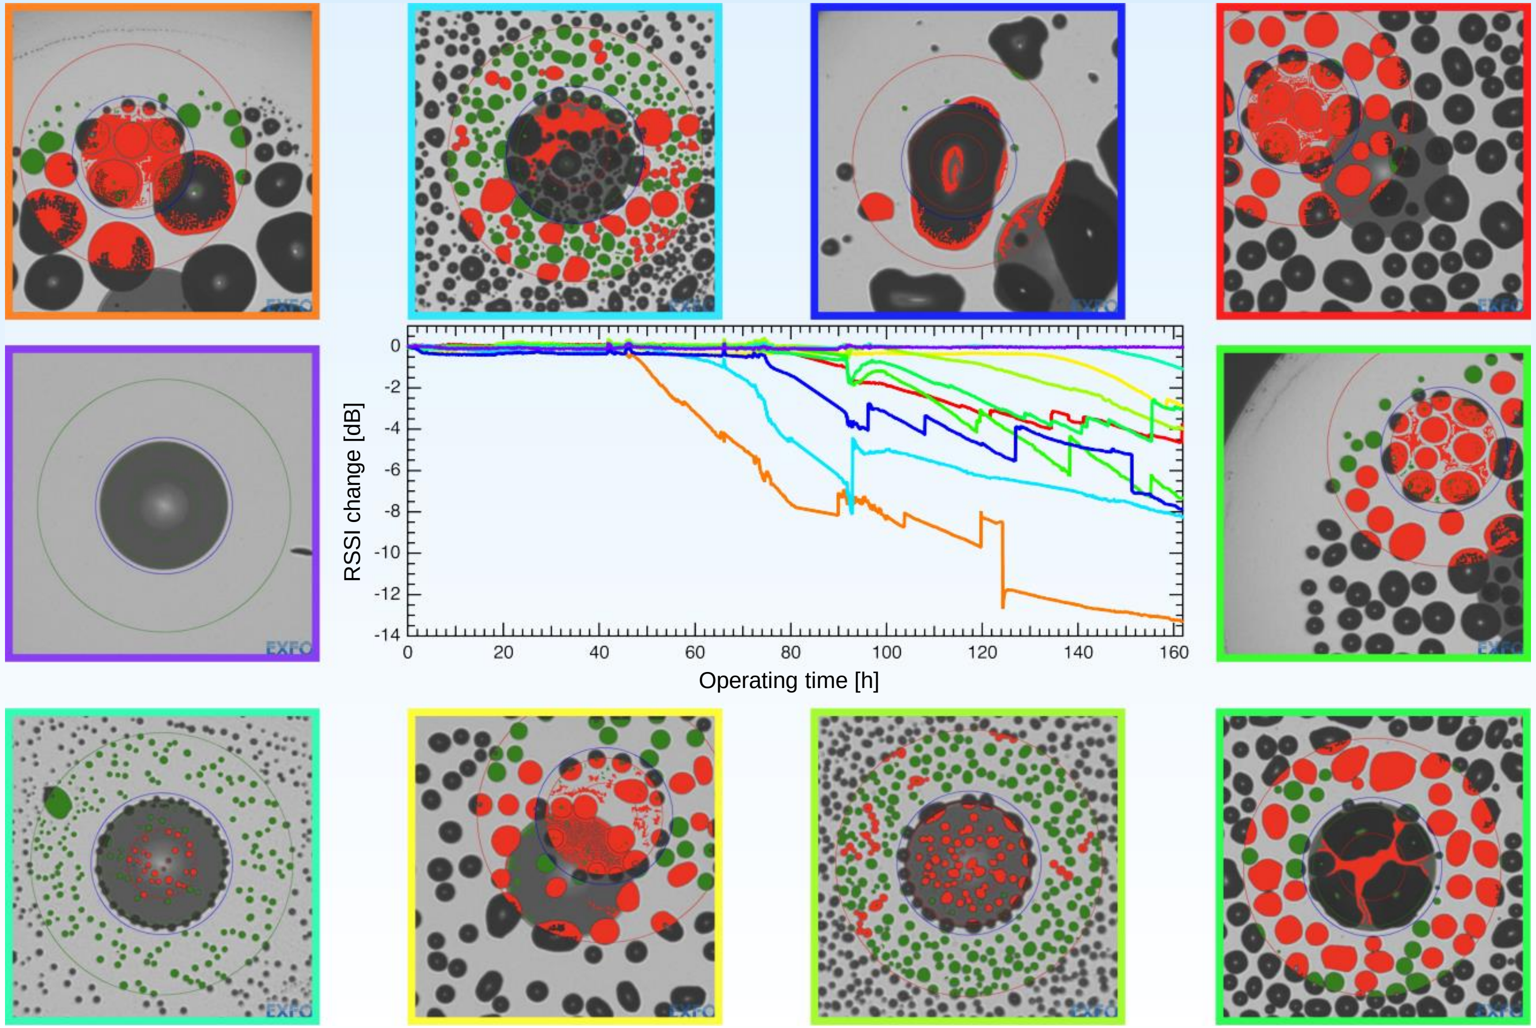

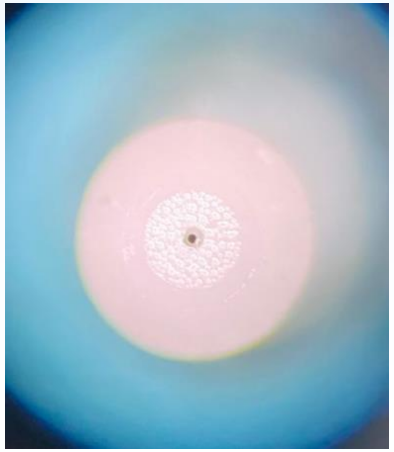

Image from
J. Dittmann CMS HCAL

Something out-gasses and condenses

- Now the earlier observations about the weird temperature behavior, where cooling helped but heating solved the problem started to make sense:
  - In order to out-gas the device must be hot enough
  - In order to condense the fiber end-face must be cool enough
  - **Temperature difference is the key!**
- These conditions can often happen in experiments, where VTRxs are on a motherboard with other warm electronics but the fibers are on the front panel in cooler air
- Not the case in our typical temperature testing, where the temperature differences are created with a temperature test chamber
  - When device is heated, fibers are heated as well

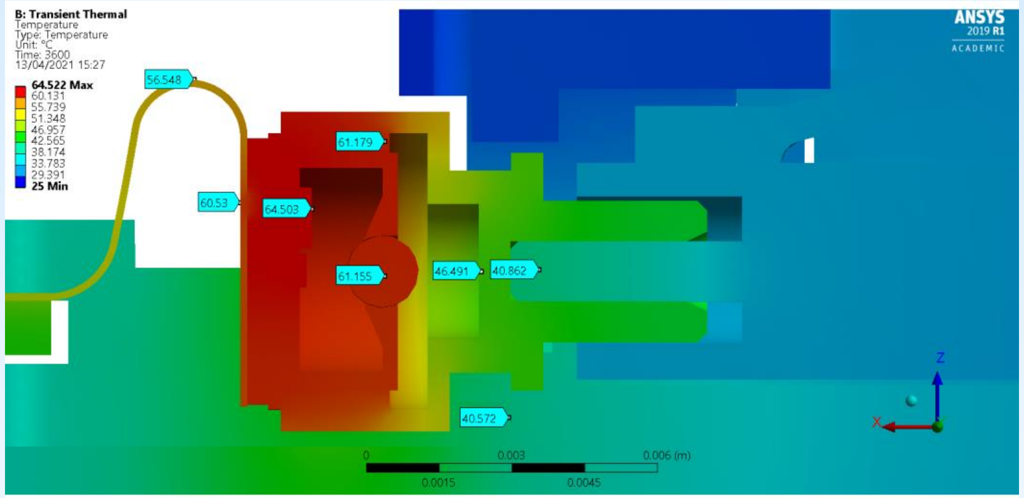

Thanks to CERN EN-MME-EDS

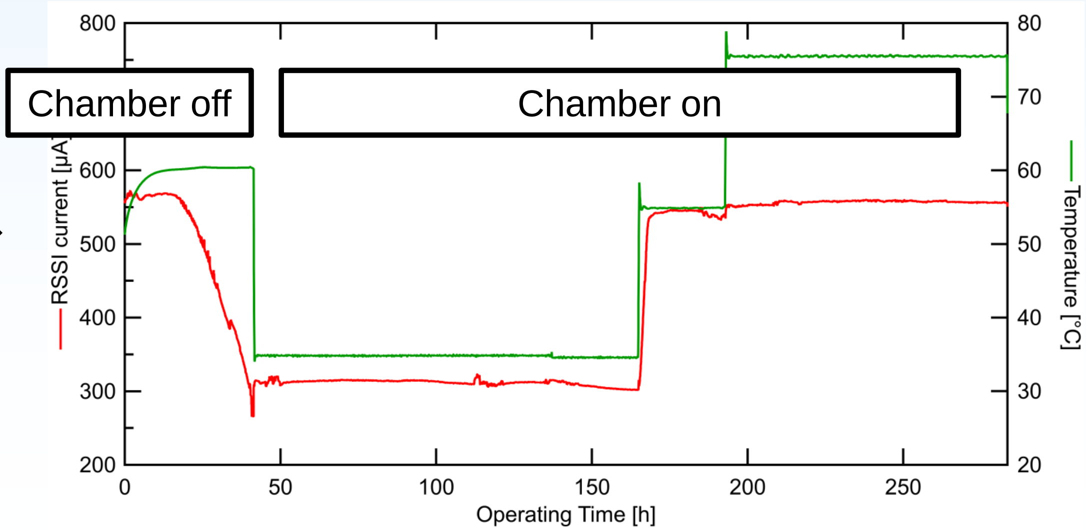

Explains also the sudden changes

Versatile Link

- Contaminant condenses and forms bubbles that grow
- Bubbles are somewhat mobile, especially due to surface tension when smaller bubbles combine and form larger bubbles
- The main reason for optical losses is ***not the attenuation*** caused by the contaminant ***but the refraction!***
  - Bubbles act as small lenses on the fiber-end face

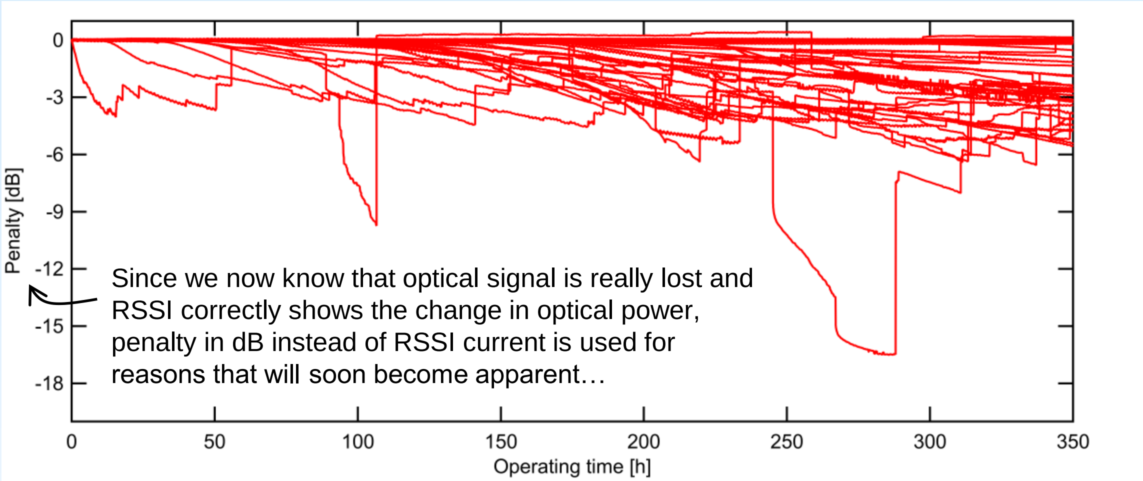

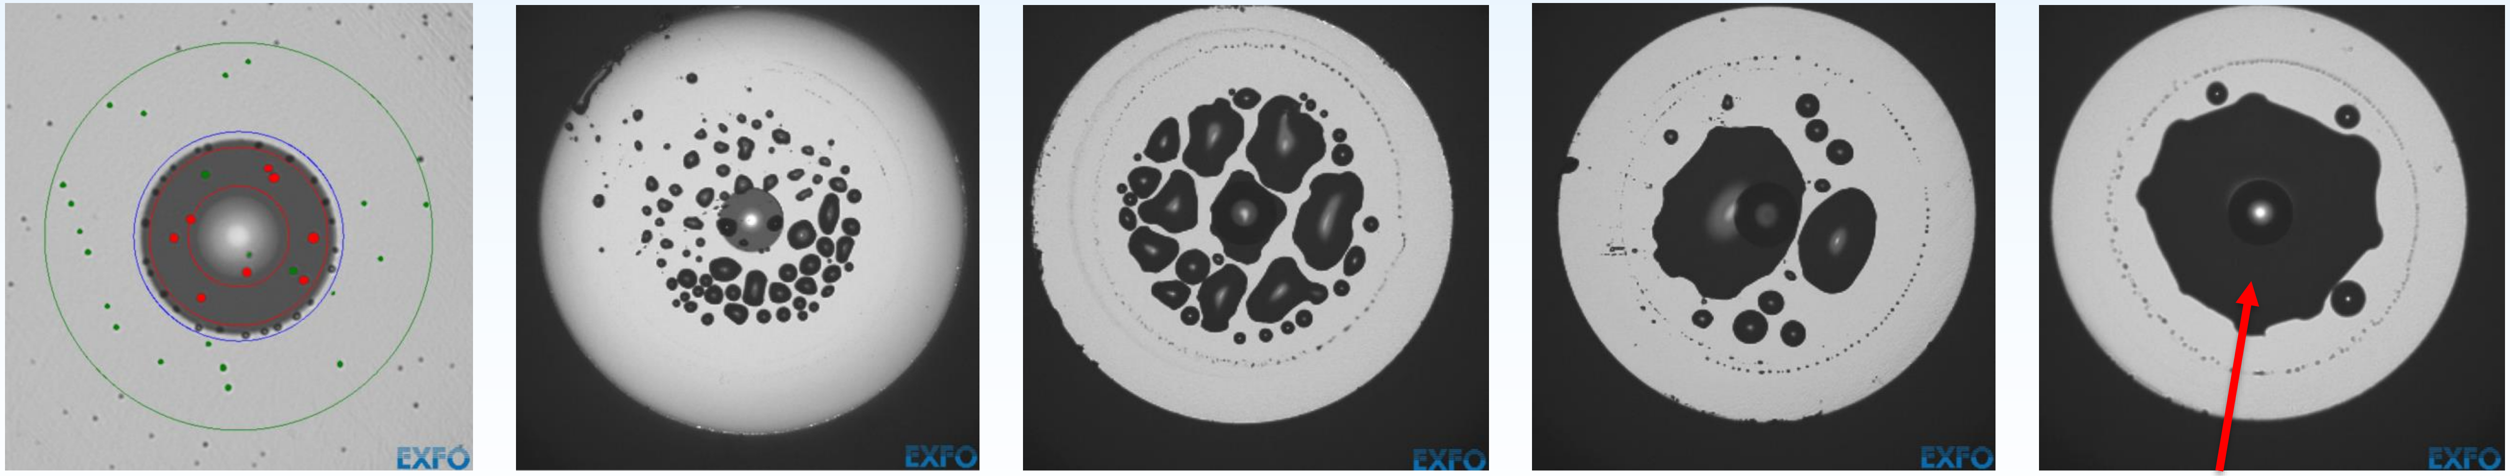

Examples of contamination bubbles at different stages (not the same fiber).

Interesting sidenote: looks pretty bad, right? Worked perfectly: uniform contaminant layer does not refract the light too much…

Quick detour

How can the contaminant even get to the fiber end-face?

- OSA (optical sub-assembly) can be designed in different ways
- V in VTRx stands for Versatile
  - One aspect of the versatility is options for 850 nm and 1310 nm wavelengths
- The same supplier provided both ROSA versions and their design decision was to use a receptacle with a hole for the light path
  - If light travels in air the wavelength does not matter and the same design suits for both versions!

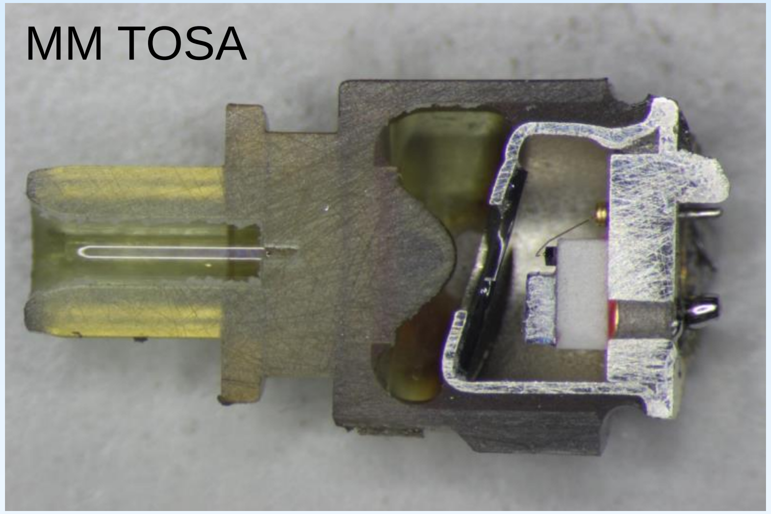

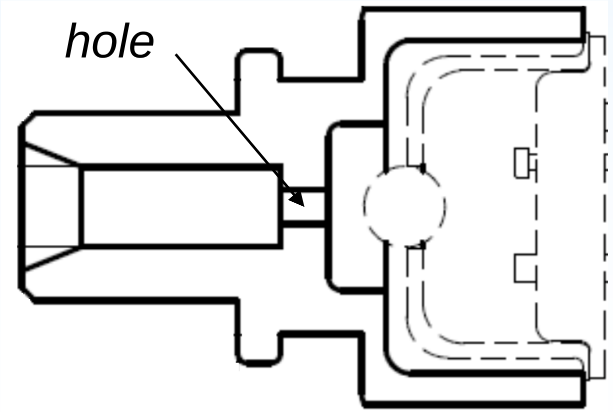

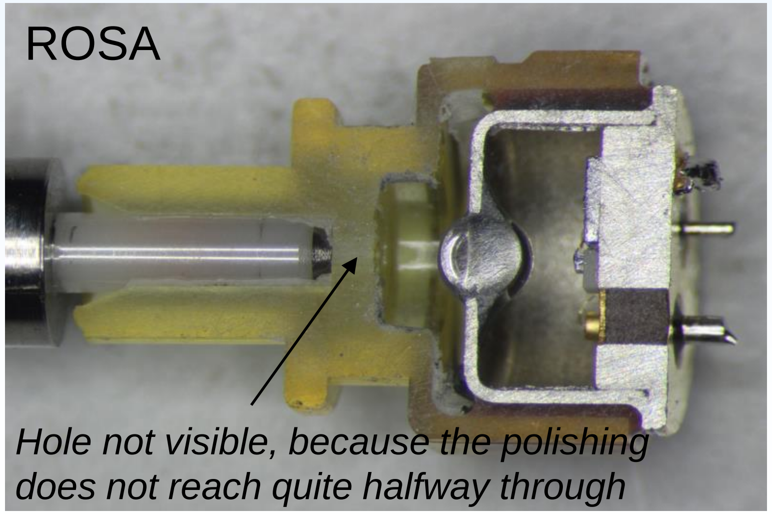

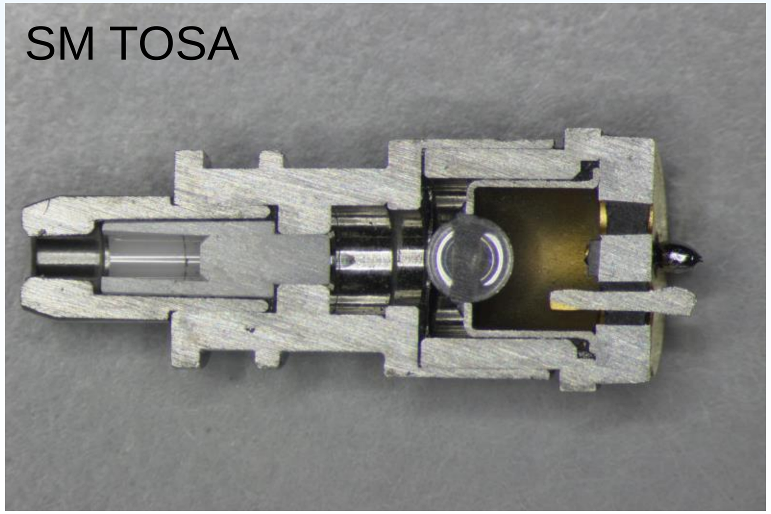

Nature of the failures

Versatile Link

- Due to the dynamic nature of the process a failing link is not permanently damaged
- The percentage of links failing seems to stabilize over time
  - While some links drift to state of unacceptable loss, others recover without any intervention
- Conditions can also change what dominates the process
- Example:
  - Device warms up fast and out-gassing process starts quickly in a device that has drifted already earlier
  - Fiber warms up slower but eventually is too warm to act as an effective condensation surface

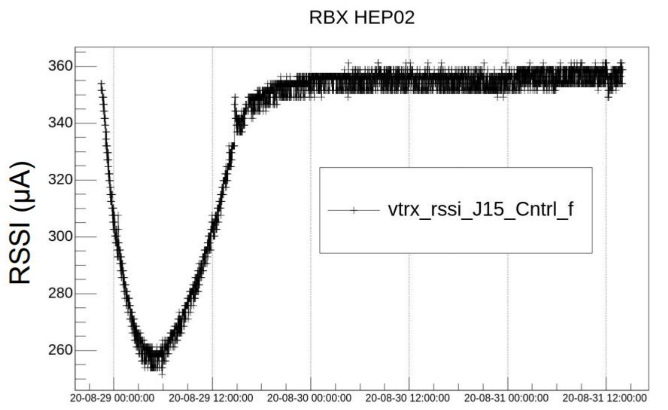

Beautiful real-life example from CMS HCAL

- Obviously, not a great situation, but it is not like wait long enough and all the links are permanently dead
- The percentage of failing links depends on the power margin

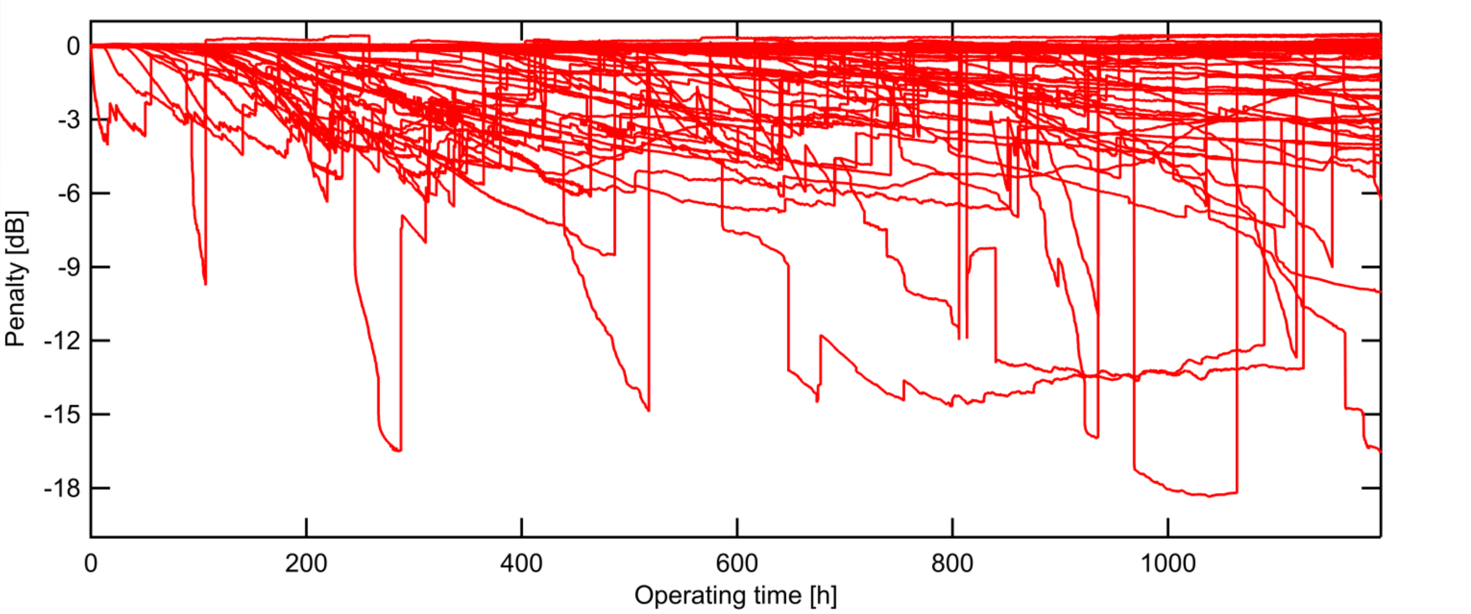

Power budget and power margin

- Let's use the Versatile Link specification (which your links should follow) as a starting point:

	Specification	
Min. Tx OMA	-3.2 dBm	For example, MiniPOD 99%: **> -1.5 dBm** 50%: **> -0.4 dBm**
Max. Rx sensitivity	-13.1 dBm	→ (see chart)
Power budget	**9.9 dB**	
Fiber attenuation	0.6 dB	150 m MM fiber
Insertion loss	1.5 dB	3 × connection 0.5 dB each (MPO mean random mated: 0.2 dB)
Link penalties	1.0 dB	Takes into account other signal properties
Rx radiation penalty	2.5 dB	The specified Rx radiation penalty only if irradiated up to 5e14 n/cm^2, if below 1e14 n/cm^2 penalty < 0.5 dB
Fiber radiation penalty	0.1 dB	
Margin	**4.2 dB**	

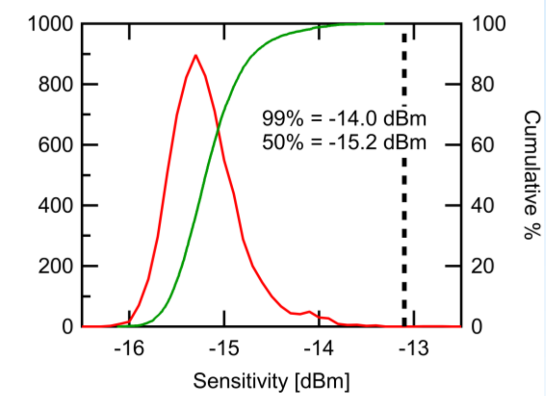

	Realistic worst case	Typical
Min. Tx OMA	-1.5 dBm	-0.4 dBm
Max. Rx sensitivity	-14.0 dBm	-15.2 dBm
Power budget	**12.5 dB**	**14.8 dB**
Fiber attenuation	0.6 dB	0.6 dB
Insertion loss	1.5 dB	0.6 dB
Link penalties	1.0 dB	1.0 dB
Rx radiation penalty	0.5 dB	0.5 dB
Fiber radiation penalty	0.1 dB	0.1 dB
Margin	**8.8 dB**	**12 dB**
Margin w/ GBT coding	**10.8 dB**	**14 dB**

Error coding scheme, for example, the GBT coding can result in in an additional 2-3 dB gain in margin.

What does a drift of X mean in practice?

Versatile Link

- Specification:
  - 5253 / 102130 device hours failing = **5.1% dead time**
- Realistic worst case (all components in the worst 1%):
  - 1791 / 102130 device hours failing = **1.8% dead time**
- Typical (50% devices):
  - 956 / 102130 device hours failing = **0.9% dead time**

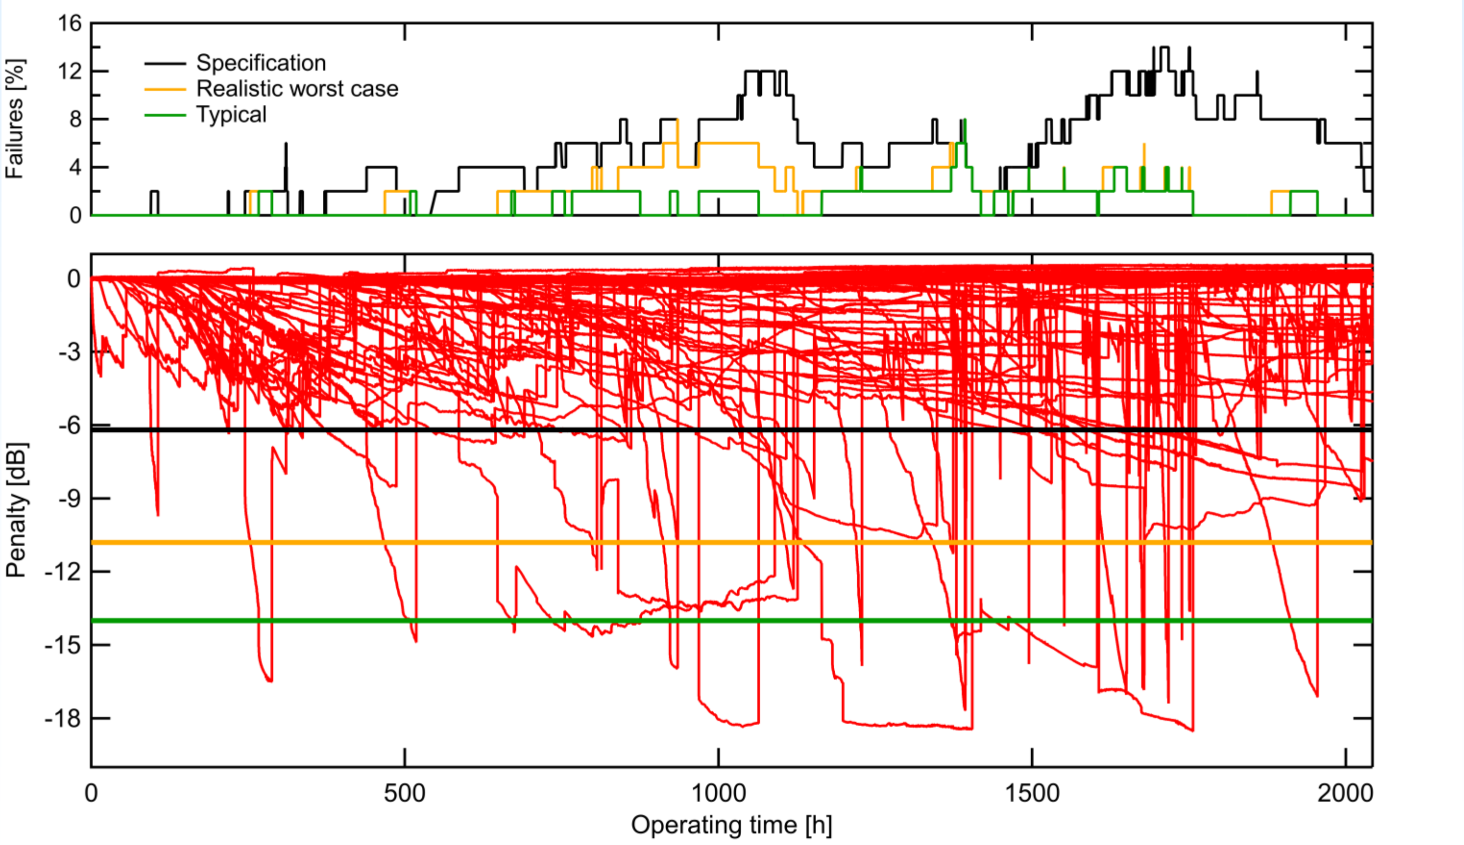

	Specification	Realistic worst case	Typical
Min. Tx OMA	-3.2 dBm	-1.5 dBm	-0.4 dBm
Max. Rx sensitivity	-13.1 dBm	-14.0 dBm	-15.2 dBm
Power budget	**9.9 dB**	**12.5 dB**	**14.8 dB**
Fiber attenuation	0.6 dB	0.6 dB	0.6 dB
Insertion loss	1.5 dB	1.5 dB	0.6 dB
Link penalties	1.0 dB	1.0 dB	1.0 dB
Rx radiation penalty	2.5 dB	0.5 dB	0.5 dB
Fiber radiation penalty	0.1 dB	0.1 dB	0.1 dB
Margin	**4.2 dB**	**8.8 dB**	**12 dB**
Margin w/ GBT coding	**6.2 dB**	**10.8 dB**	**14 dB**

…not done yet

- Two lines of investigations were started immediately:

1. What is this stuff and where does it come from?

2. If the stuff is released from inside, it must run out at some point?

1. What is this stuff and where does it come from?

- In order to determine at which stage of the assembly process the contaminant was introduced different sample sets were tested
  - Samples from different times in the series production
  - Samples that had been in contact with the dust caps and without them
  - Samples with box-fresh ROSAs straight from manufacture’s original packaging

- All above, even the box-fresh ROSAs, showed typical drifting, which meant that the contaminant was present already from the ROSA manufacturer

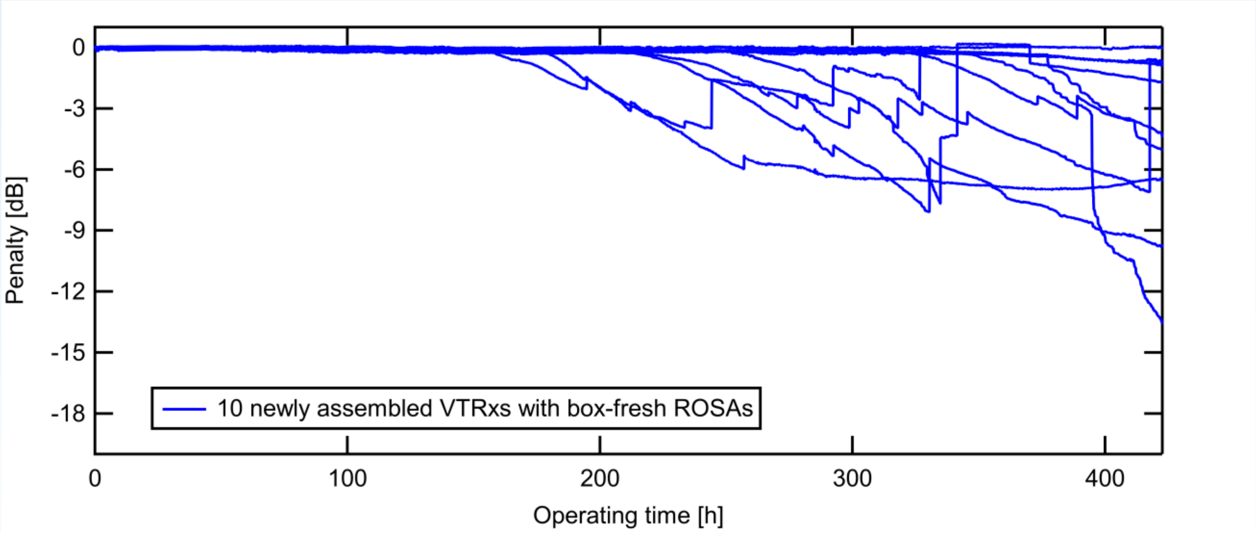

1. What is this stuff and where does it come from?

Versatile Link

- Metal TO-can is actively aligned with the plastic receptacle and tacked in place with two dots of UV glue on opposite sides of the package
- Final structural strength is achieved with epoxy that is applied on all sides of the assembly and then heat cured
- Assembly recipe gave us a list of potential contaminant sources

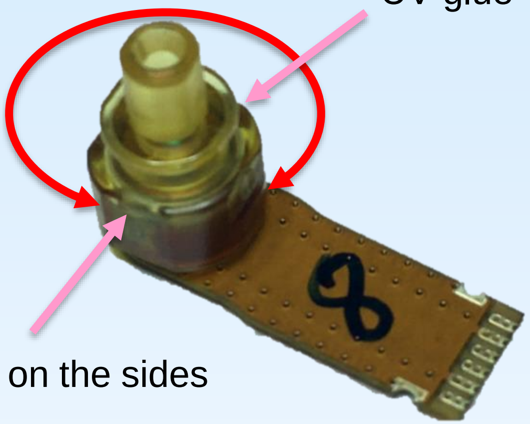

Candidates:

- Structural Adhesive ("epoxy")
- UV-cure Adhesive ("UV-glue")
- Unknown process chemical

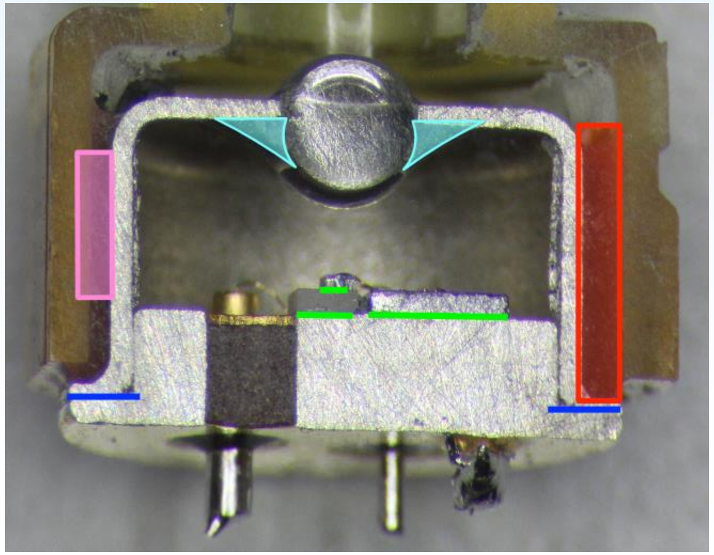

Other parts, not really candidates:

- Conductive Adhesive
- Solder Glass
- Laser Weld

Source of the contaminant

Versatile Link

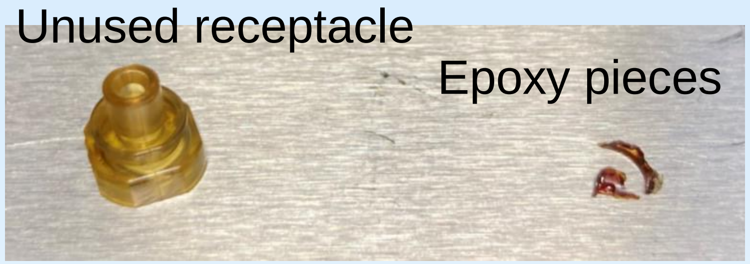

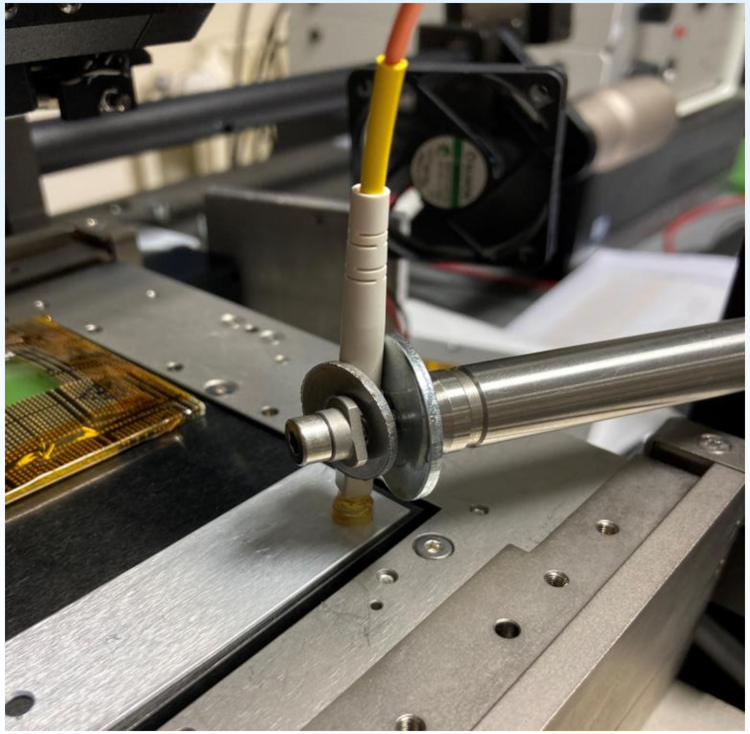

- Epoxy sample from ROSA heated on a hot plate and collected on a clean fiber using a new receptacle
  - Looked like the contaminant observed in our lab and at the experiments
  - Also sticked on the fibre end-face for longer time
  - We were also able to change the epoxy colour with extra curing, which could be a sign of insufficient initial curing

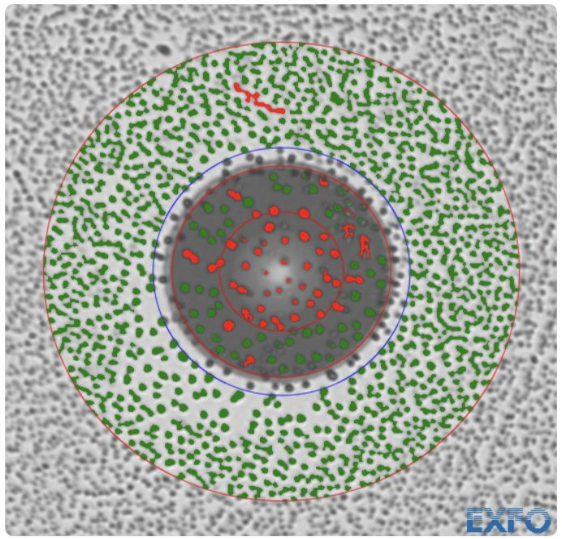

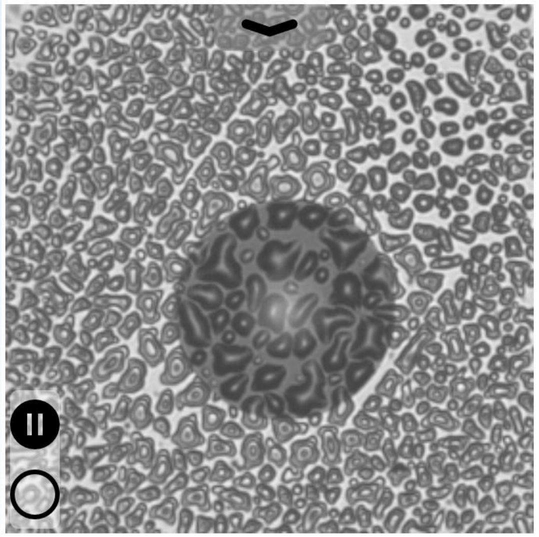

Looks really promising, but from now this point on we needed proper tools and help from real chemists

Contaminant Identification Steps

- Attenuated Total Reflection Fourier Transform Infra-Red (ATR FT-IR) spectroscopy
  - First pass attempt at identification from spectrum
- Differential Scanning Calorimetry (DSC)
  - Cure-level analysis for epoxy
- Headspace Gas Chromatography with Mass Spectroscopy (GC-MS)
  - Highly sensitive and accurate identification of individual contaminants
- Visual light microscopy
  - Detailed visual inspection of UV adhesive within ROSA assembly

WARNING! What follows is a bunch of confusing plots,
which we don't have time to explain in detail.
Organic chemistry lovers, these are included for your pleasure.
The rest of you, don't worry!

FT-IR Spectroscopy

Attenuated Total Reflection Fourier Transform Infra-Red spectroscopy

- FT-IR instrument available in CERN Chemistry lab
  - First attempt at identifying the contaminant
  - Directly wiped contaminant off of the fibre ferrule onto the ATR crystal for analysis
- **Data best match with a diol**
  - FT-IR identification is a bit difficult, especially for mixtures

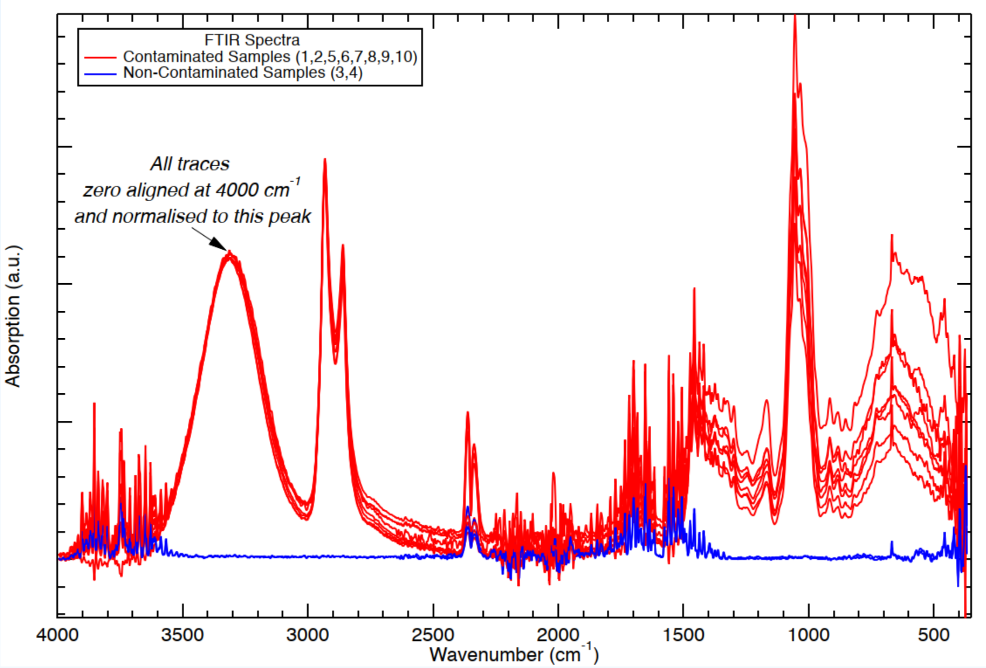

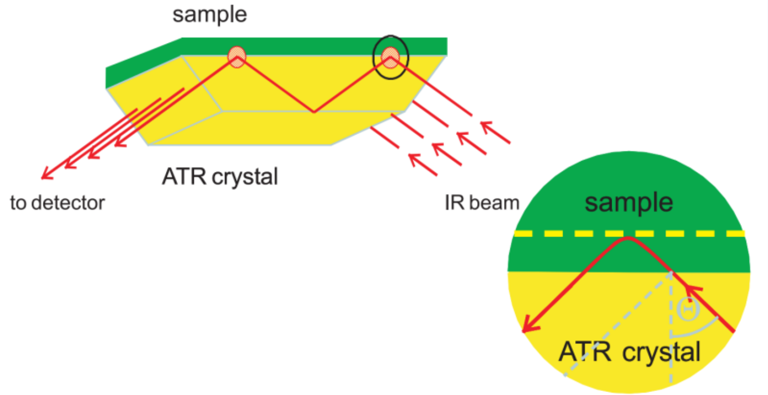

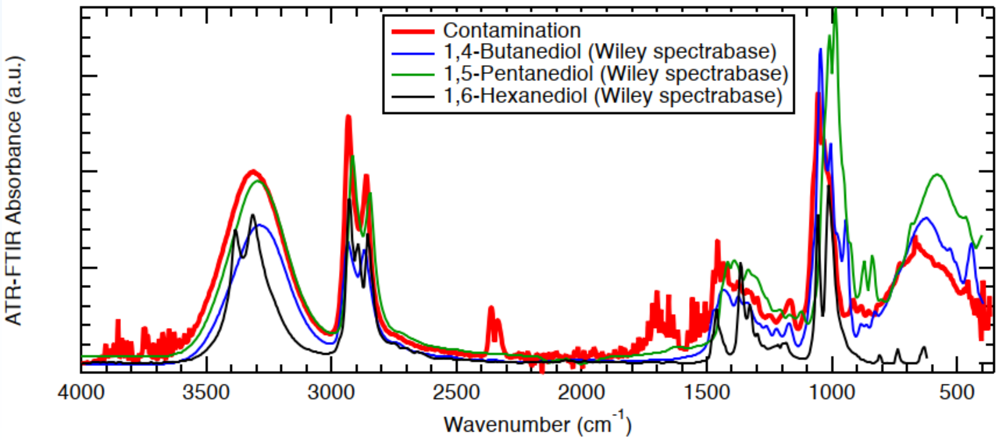

Adhesive cure analysis

Differential Scanning Calorimetry

Versatile Link

- DSC on epoxy from ROSA shows a relatively low Tg (glass transition temperature)
- Epoxy can be further cured using additional application of heat

- DSC on UV adhesive technically challenging
  - The Tg is low, and not well defined
  - Too small samples for other methods

- Why could the UV adhesive be not properly cured?
  - Ultem (receptacle material) absorbs UV very strongly
  - UV-adhesive harvested from ROSAs often found to be still "tacky", as if not entirely cured

- **Signs of insufficient curing of both epoxy and UV-glue**
  - **This could allow unwanted reactions and mixing of the two adhesives**

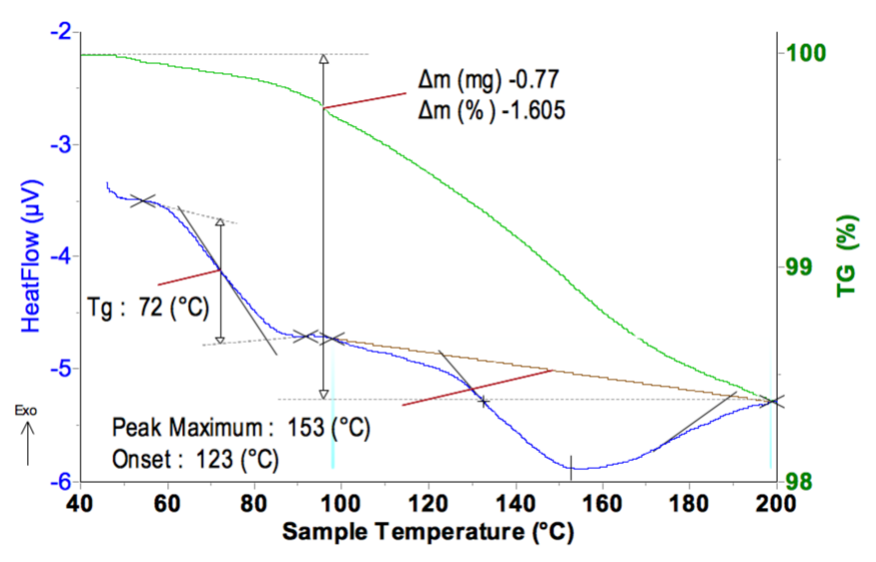

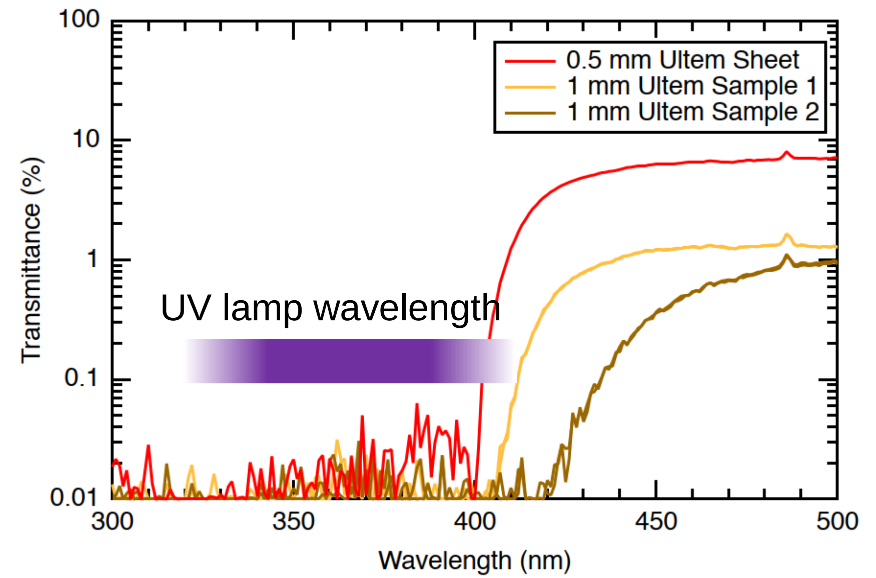

Contaminant id by GC-MS

Headspace Gas Chromatography with Mass Spectroscopy

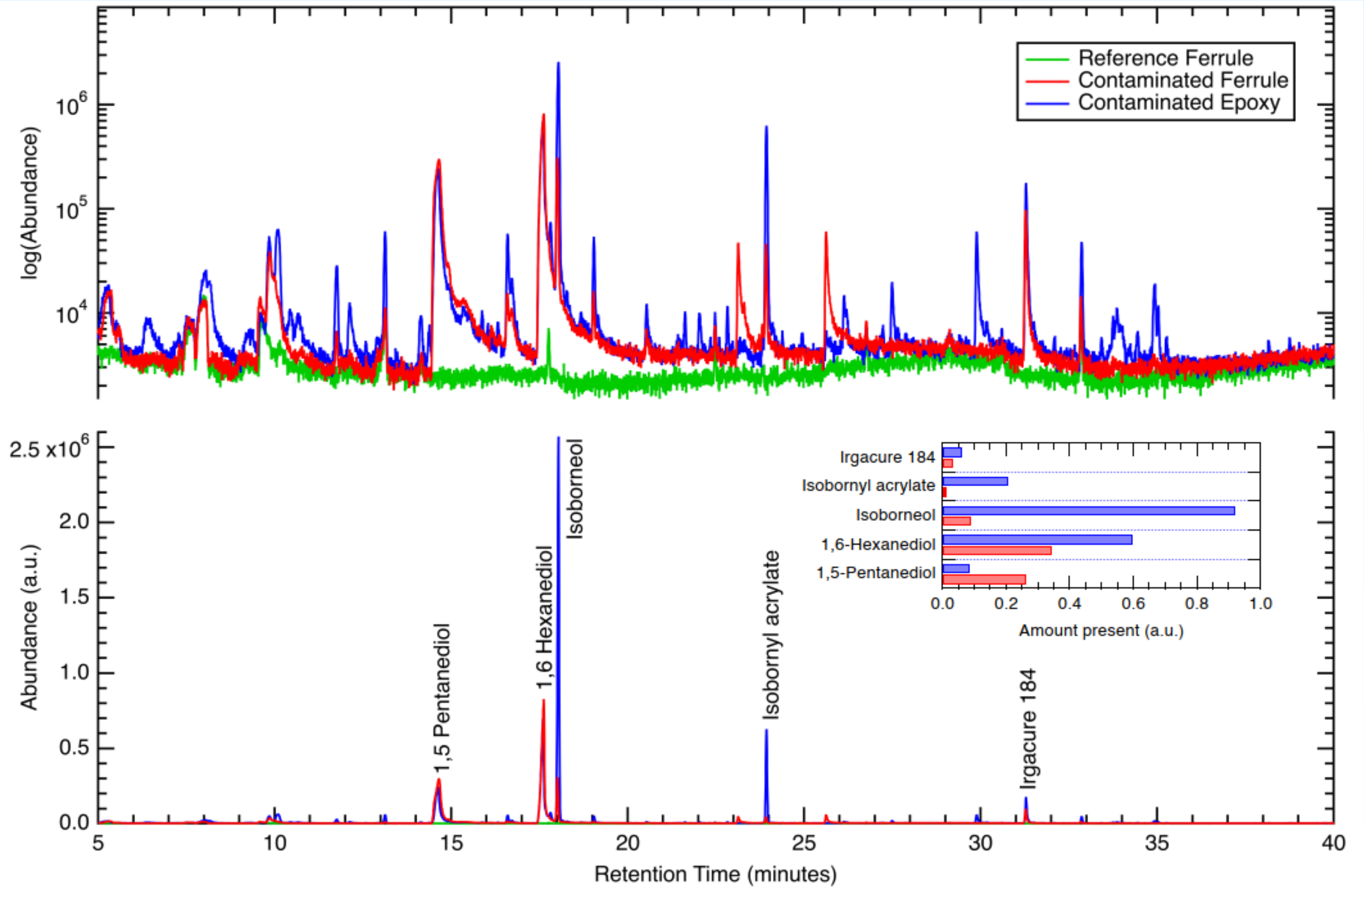

- **The contaminant is clearly a mixture of two diols**
  - Isoborneol, Isoboryl acrylate, Irgacure should not be present in the epoxy, they come from the UV adhesive
- **Compounds in the epoxy that should not be there → mixing of the two adhesives**

Same test with fresh adhesives

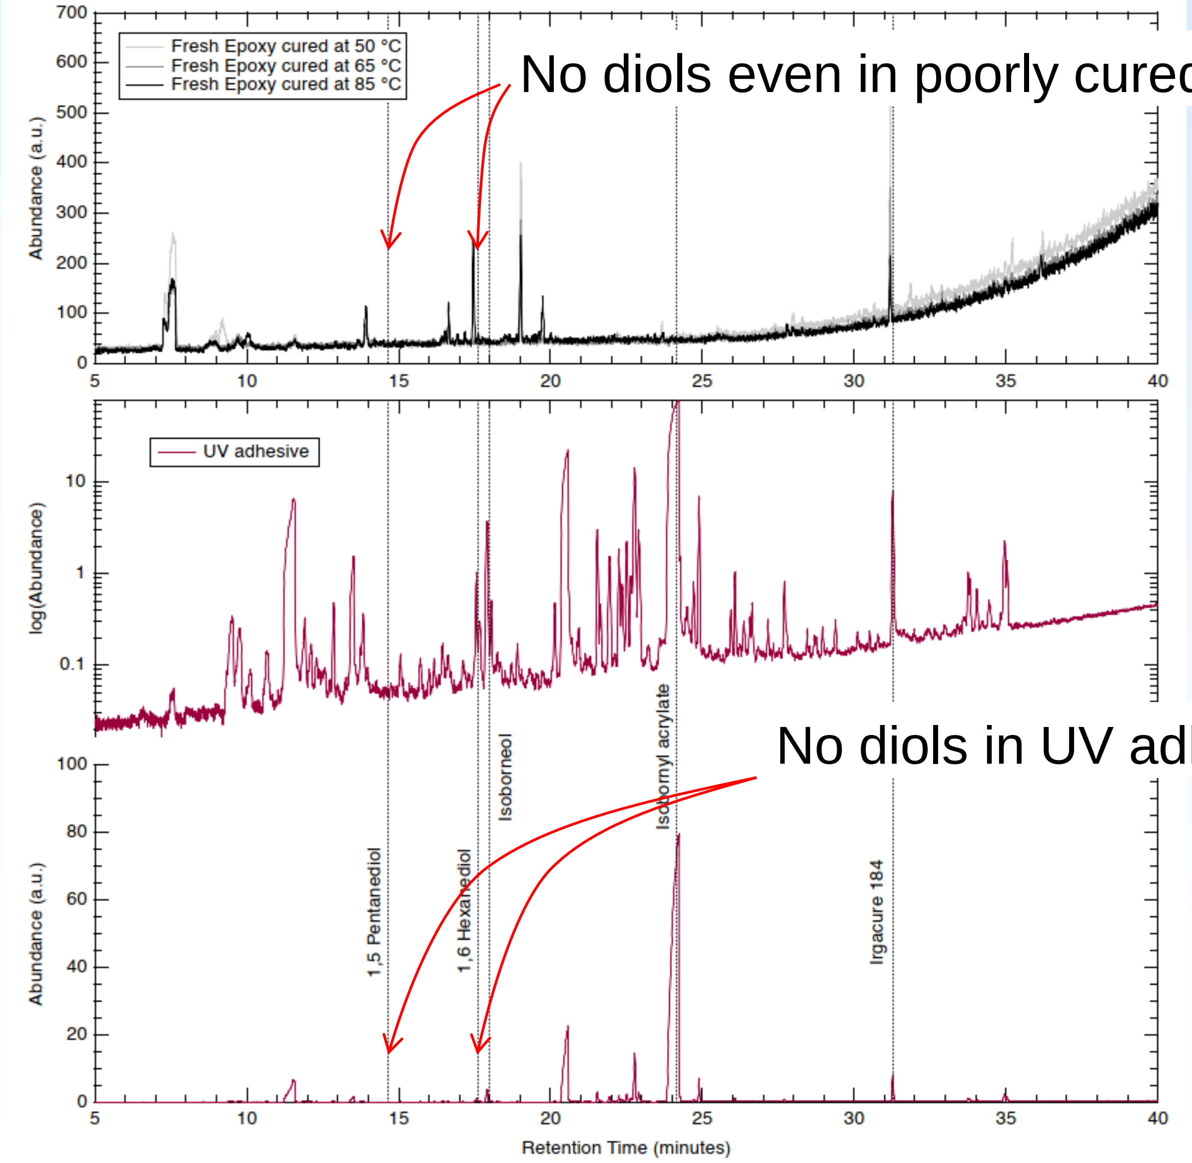

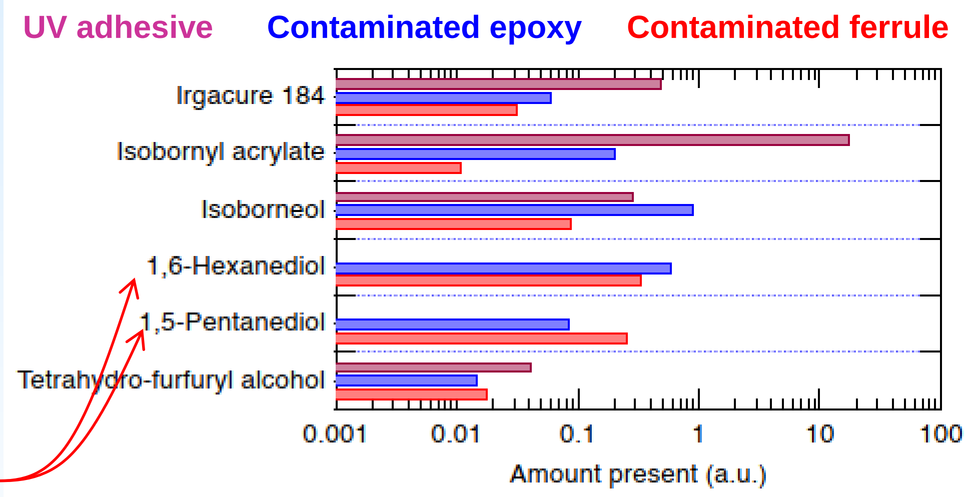

- **No diols in the fresh, “out from the bottle” adhesives**
  - **Some reaction is needed to create the contaminant product**

Microscope images

- UV-glue supplier technical experts gave opinion that multiple regions within the glue clearly indicate incomplete curing
- What happens to unreacted monomers is difficult to predict
  - e.g. mixing with epoxy, reactions catalysed by metal TO-can, …
- **One of the unreacted monomers from the UV-glue could decompose into diols**

- Application of the UV-glue has been rather **poorly controlled process**
  - Quantity, location, apparent cure level varies
  - **Could explains why there are so big differences in how and when the drifting occurs in different samples**

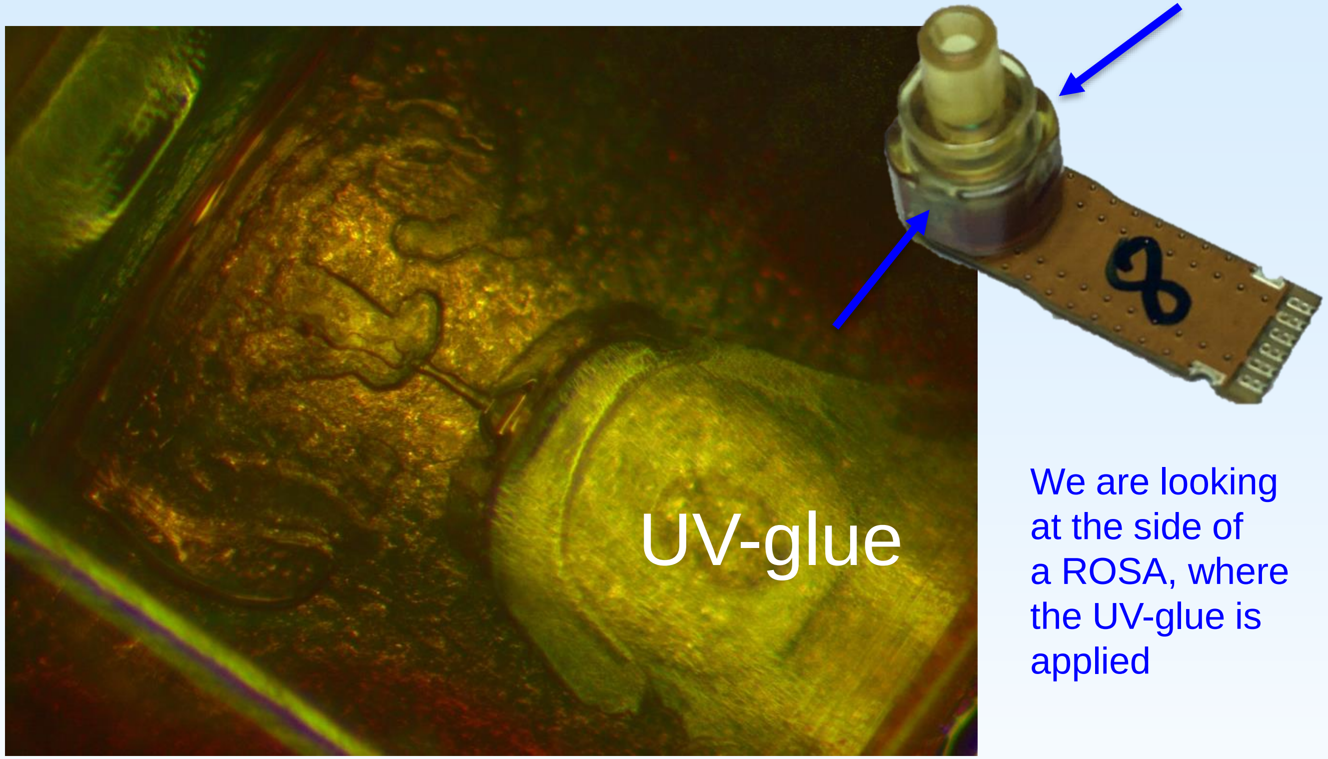

We are looking at the side of a ROSA, where the UV-glue is applied

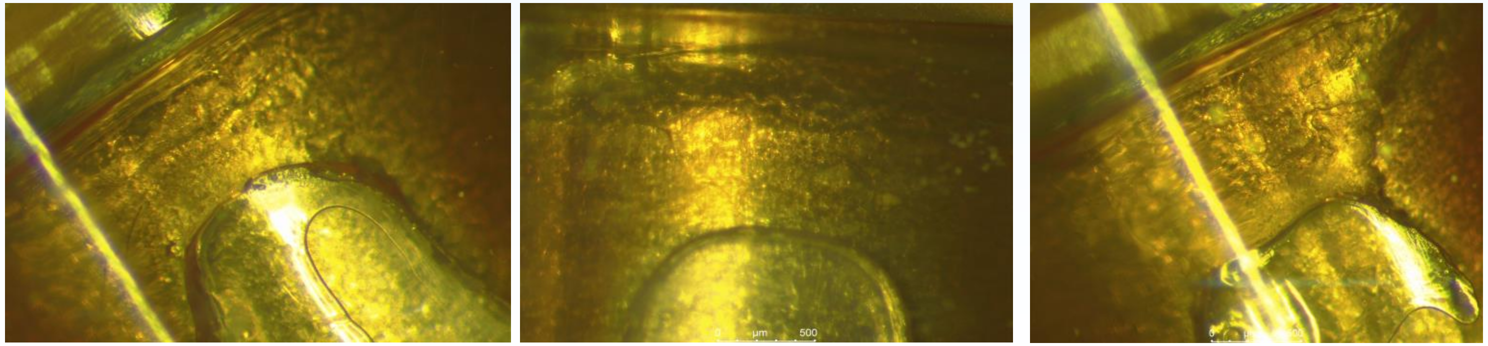

Conclusions for 1. “What is this stuff?”

- Diols outgas within the assembly during operation & condense on the optical fibre ferrule impeding the light path

- Precursor components of the contaminant is in the UV adhesive
  - Requires combination of poor curing and reactions with other compounds in the assembly (likely also rather poorly cured epoxy) to create diols

- Diols can be absorbed into the epoxy and slowly migrate and release when heat is applied

- Other process chemicals were ruled out as none of the chemicals that could have been in contact with ROSAs during the assembly matched the chemical composition of the contaminant

2. If the stuff is released from inside, it must run out?

- It was very quickly shown that the contaminant is not water, and therefore it is not continuously absorbing into the ROSAs from air
  - There must be only a limited quantity of the stuff inside the assembly
- The question was if it can be baked-out and how could it be done without damaging the device and endangering its long-term reliability
  - What is required bake-out temperature?
  - Required duration?
- …or are we so unlucky that it is an integral part of one of the adhesives and it will out-gas as long as there is adhesive left in the assembly

Bake-out

- Good starting point was a sample set from the original QA
  - 10 samples that had been kept at 85°C for 2000 hours
  - So called *High Temperature Storage*, HTS, test
- Results were promising
  - No characteristic drifts after more than 3000 hours
- 85°C is still safe temperature, because devices are not powered
  - No penalty on the long-term reliability
  - Surely lower would always be better
- 2000 hours is a long time
  - Can we make that shorter?

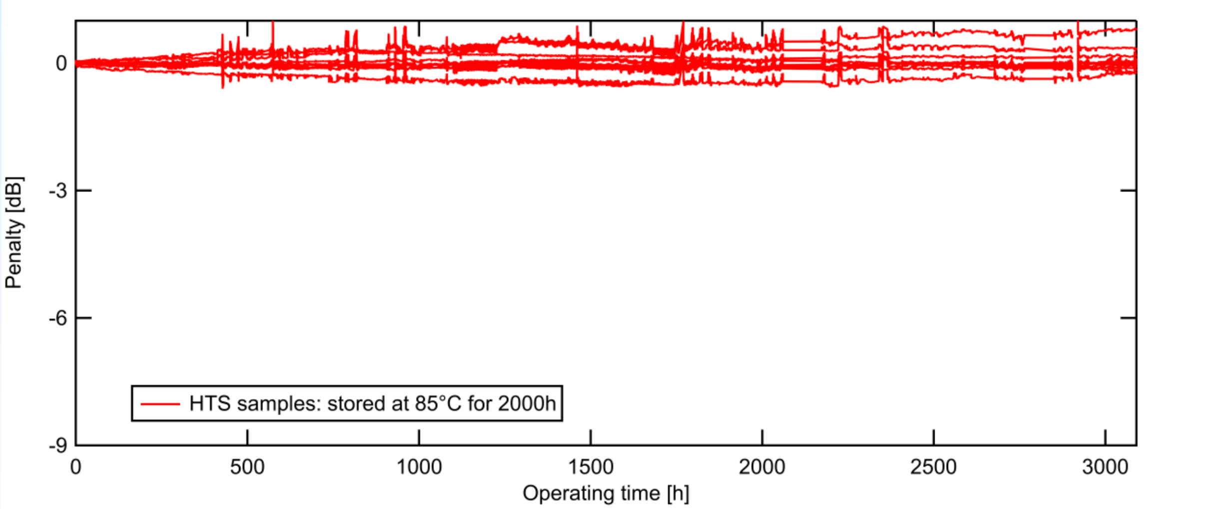

Bake-out optimization

- Several different bake-out profiles and post cures, both heat curing and UV curing, were tested
  - 120°C for 2h (post curing the epoxy)
  - 65°C for 48h
  - 85°C for 50h
  - 85°C for 130h
  - UV-post curing different times

  Signs of drifting

- Finally, 300h at 85°C seemed to have made a significant difference

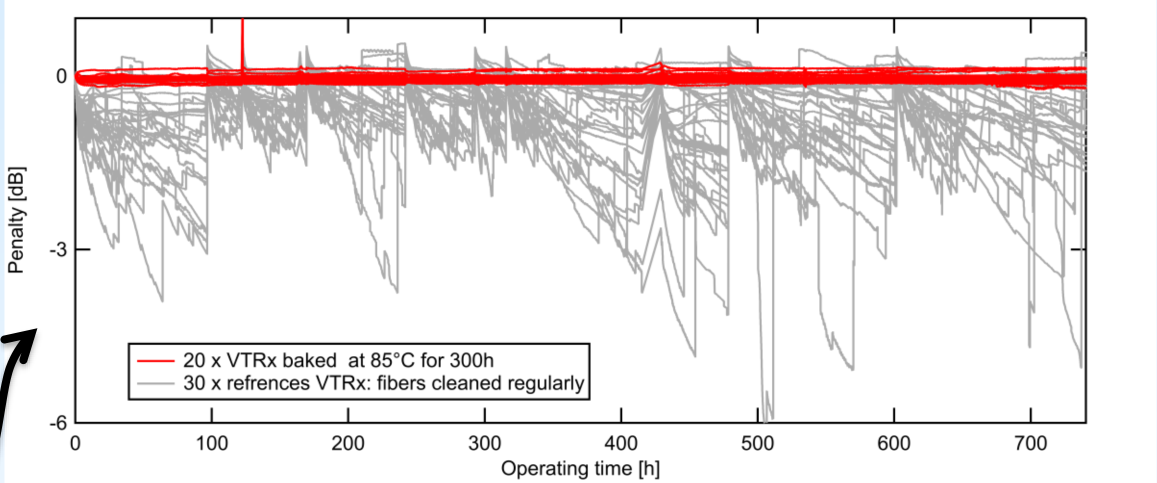

- While we started to prepare bake-out facilities, a larger sample set baked 300h at 85°C was being tested
  - Eventually, and unfortunately, 300h bake did not seem to be a perfect solution
  - However, it definitely made a big difference

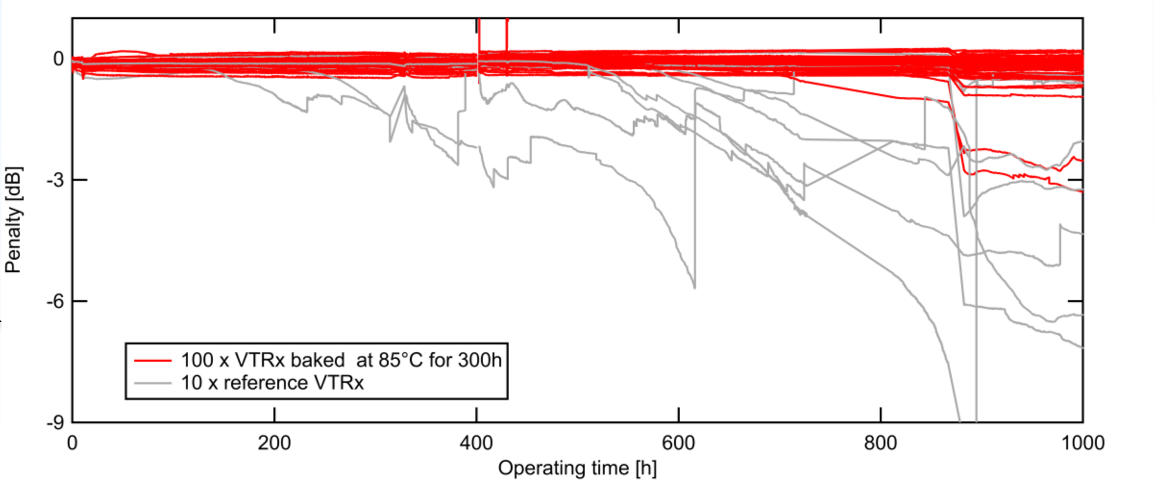

Gaps in the data taking

- Standard bake-out process was chosen to be **500h at 85°C**

Bake-out facilities

- Ovens were ordered and a room, instruments and software prepared to handle the bake-out process
  - Device tracking
  - Bake-out
  - Testing

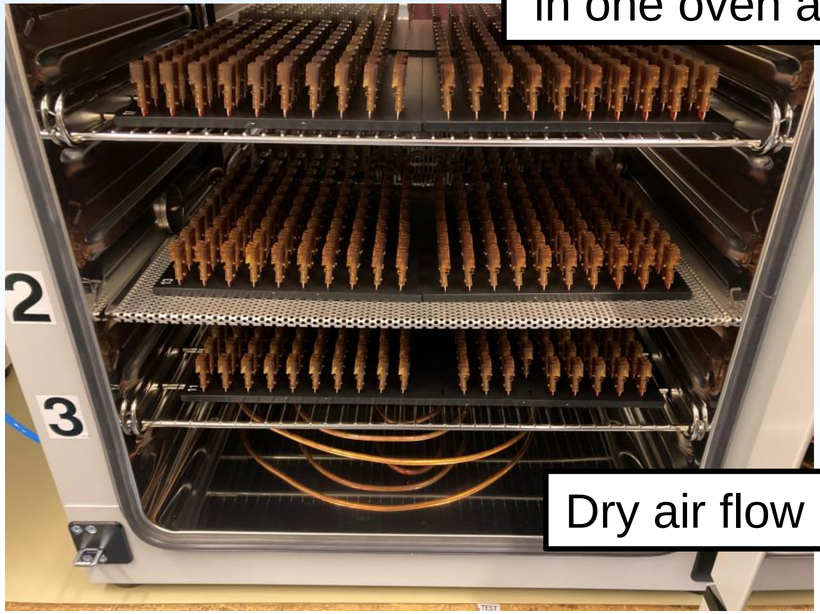

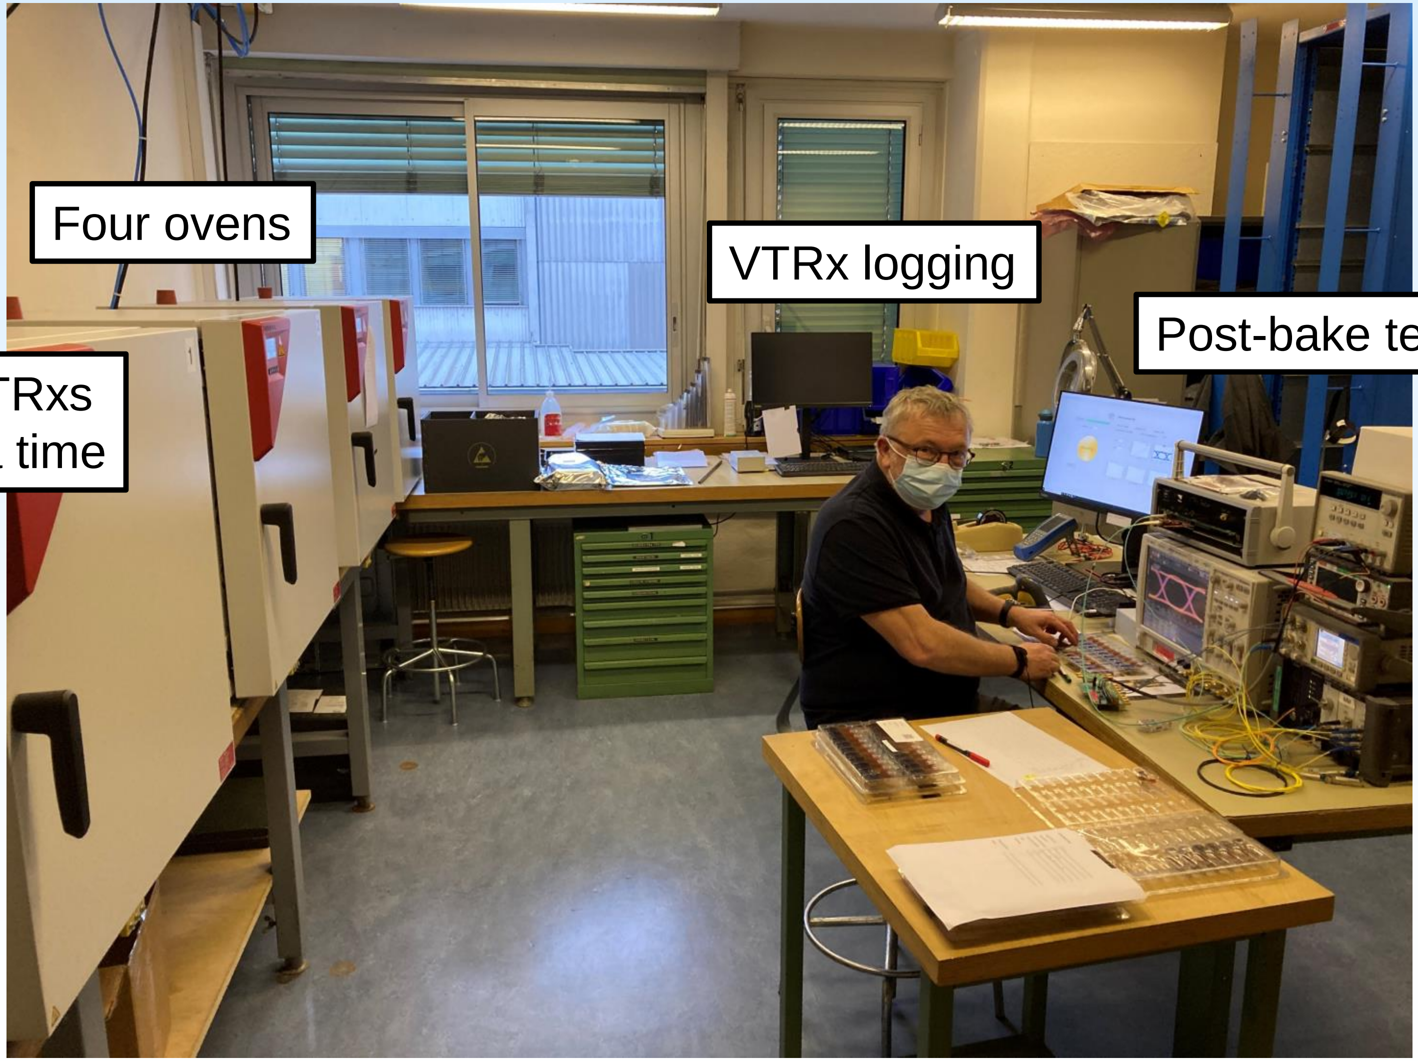

Bake-out statistics

- **4652** VTRxs have gone through the full bake-out process
  - Yield after bake-out: **93.4%**
  - Majority of the 307 failures were caused by issues in the receiver flex
    - Flex replacement solved 270 failures
    - Yield after repair and retest: **99.2%**
- **2618** VTRxs have gone through the bake-out and post bake-out test
  - Yield after bake-out: **95.2%**
  - Repair and retest in progress…
- **7270 / 18800 ≈ 40%**

Running test for 500h baked samples

- We have a long-term test continuously running with 50 samples from the standard 500h at 85°C bake-out process
  - This will be kept running for a long time

- We try to make to conditions bad on purpose
  1. Samples inside insulating boxes
  2. Fibers outside the heated volume in air-conditioned lab
  3. Additional heating with resistors on top of each ROSA

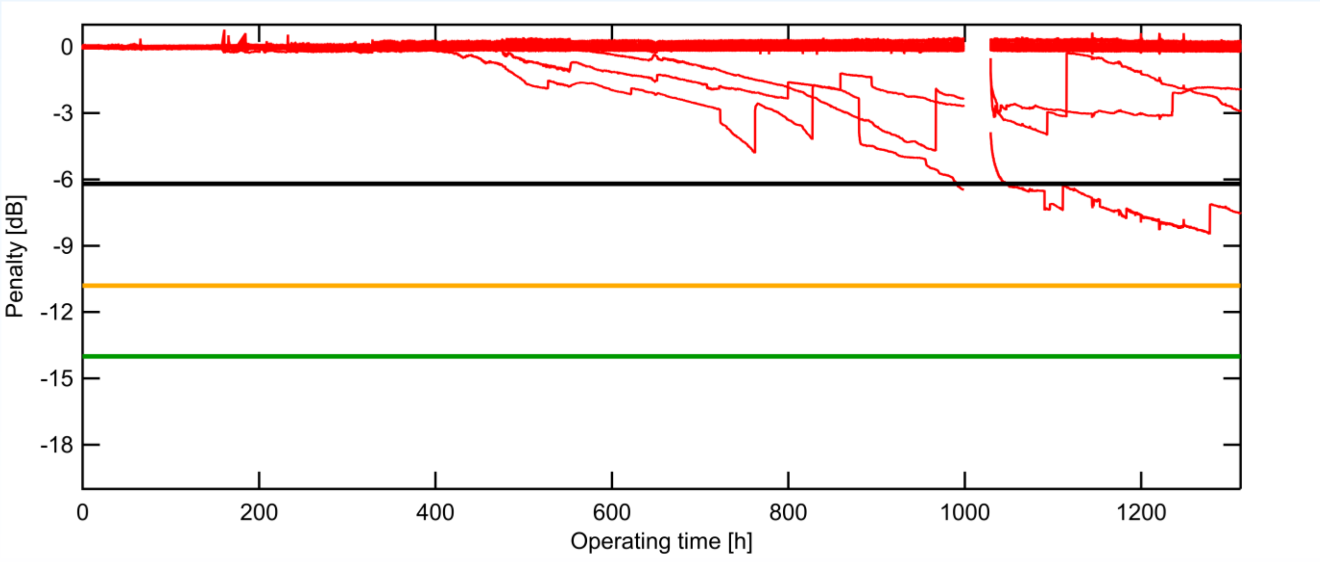

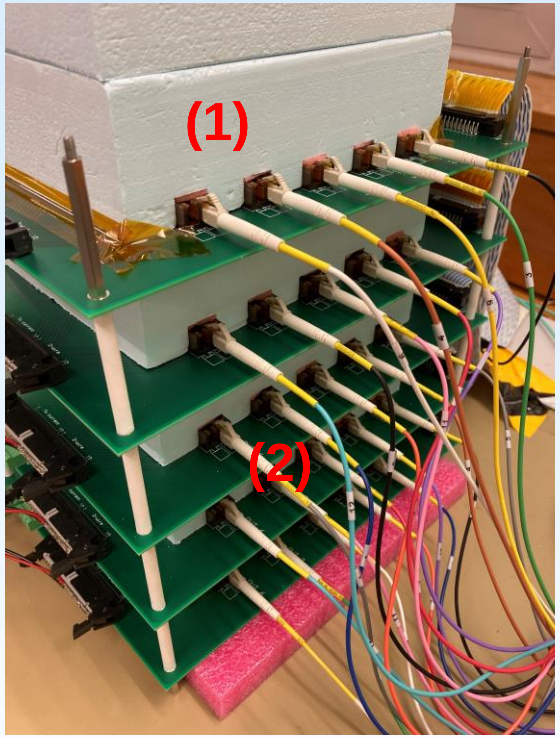

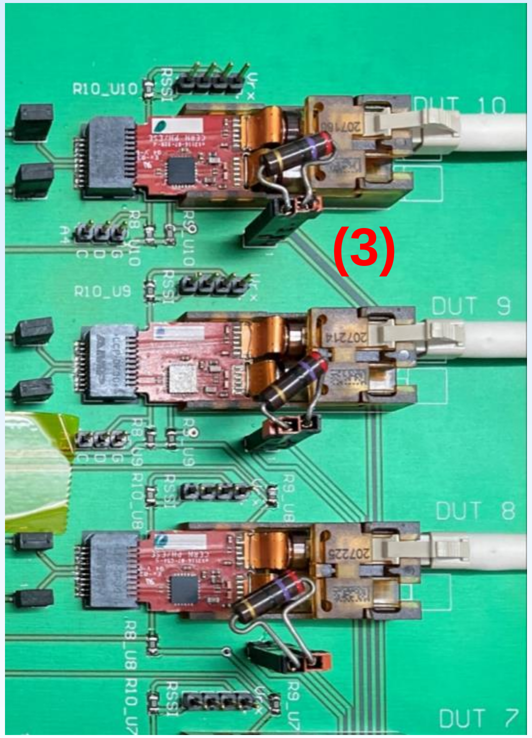

- Specification:
  - 272 / 65650 device hours failing = **0.4 % dead time**
- Realistic worst case (all components in the worst 1%):
  - 0 / 65650 device hours failing = **0.0 % dead time**
- Typical (50% devices):
  - 0 / 65650 device hours failing = **0.0 % dead time**

Cooling

- Cooling is another highly recommended mitigation method
  - If a VTRx runs cool, contaminant is not released **whether the VTRx is baked or not**
- Added cooling also improves long-term reliability
  - Like with all electronics, cooler you operate, longer they last
  - Improve your cooling as much as possible (easier said than done, I know)
- Cooling solutions are highly experiment specific
  - For example, successful VTRx cooling improvements carried out by CMS HCAL and ALICE TPC (using HCAL's model)
  - Not going into details here
    - Contact us or colleagues from CMS HCAL or ALICE TPC, and we are all happy to share experiences and ideas

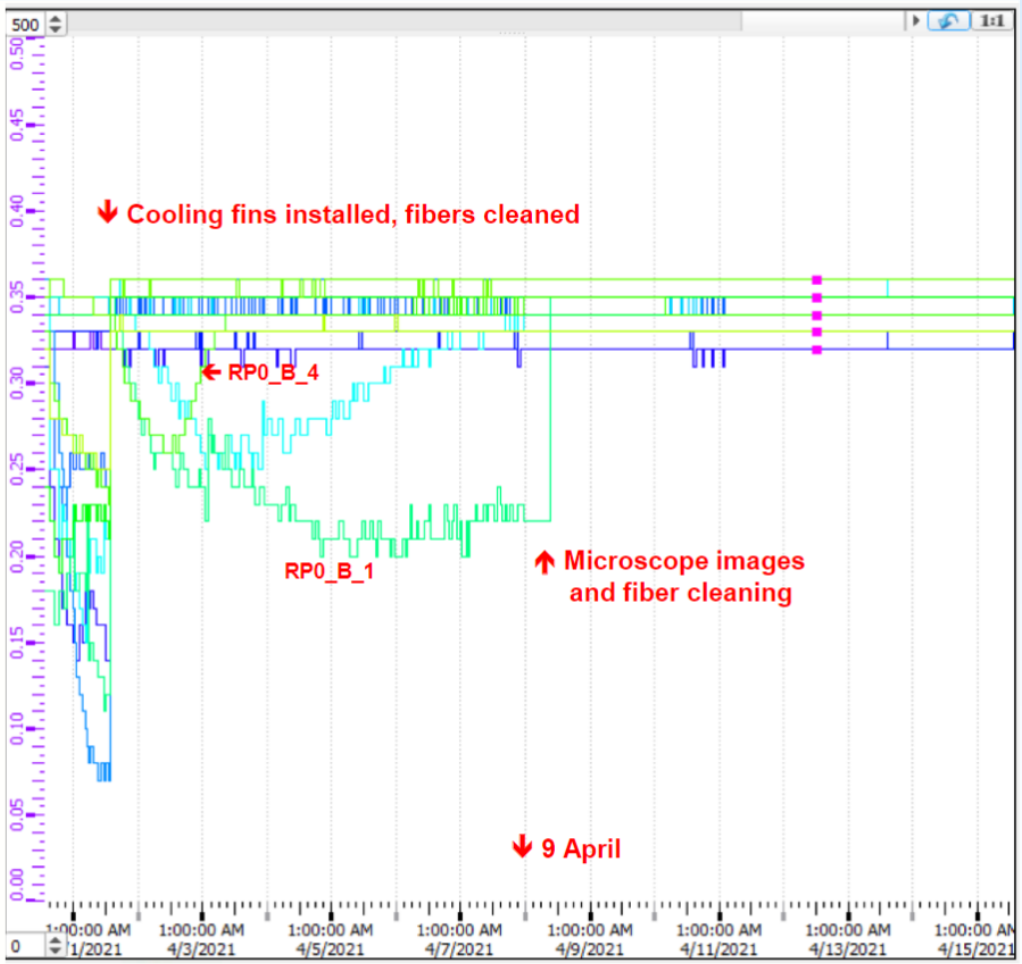

ALICE TPC example

Summary

Versatile Link

- VTRx receiver reliability issue that could potentially cause communication issues in large fraction of the datalinks has been investigated
- The root cause has found to be contaminant that out-gases from the ROSA and condenses on the optical fibre ferrule impending the light path
  - Not permanent damage
- Precursor components of contaminants are from the UV adhesive
  - Likely due to incomplete polymerisation of the UV adhesive prior to heat-curing of epoxy
  - Epoxy is porous material that can absorb contaminant from the UV glue, after which slow migration and release when heated
  - The process can take hundreds of hours before any signs are observable
- Bake-out have been shown to effectively mitigate the issue
  - Most of the contaminant released before devices are placed in the final application
- If the temperature difference between the fibre and ROSA is small out-gassing/condensation does not happen
  - Cooling is effective mitigation method

Some lessons learned

CERN Fermilab SMU Versatile Link

- Failure analysis work can be extremely slow
  - Especially reliability tests are *by definition* very time consuming
  - There is no way to speed up a test that must last for example 2000h (~3 months!) or 500 temperature cycles (easily a couple of months)
  - Also consulting experts can be surprisingly time consuming: finding right people, right services and methods, test labs etc.
- and expensive...

- It is very difficult to know beforehand what is really the most stressful condition for a device
  - We can, and we are doing industry standard reliability testing, but these may not reveal the problem
  - On the other hand, we are not able to test all different conditions devices may be placed in the final application
  - It would be important to have significant number of devices running in conditions that mimic the final conditions as well as possible, as early as possible, so that there could be time to do something before detectors are closed and sealed
    - **This requires help from the users!**

- Industry is filled with NDAs
  - At least slows down, but often completely blocks information sharing
  - If companies are used to operate need-to-know basis and protect their information, it may lead to situations where all details are not propagated down to the last subcontractor, who makes decision based on limited specs you did not manage to write perfectly including absolutely everything
  - Even the best experts make poor choices if based on bad information, and you *assume* that their choices are great because they are the experts

- Don’t assume

Thanks to all people involved in the failure analysis!

Versatile Link

- All collaborating users and colleagues
- Especially the active partners from
  - ALICE: Felix Reidt, Piero Giubilato, Christian Lippman, Alex Kluge and others
  - CMS: Jay Dittman, Grace Cummings, Jon Wilson, Magnus Hansen and others
  - LHCb: Ken Wyllie (who conveyed us updates from surely many people within LHCb collaboration)
- CERN services:
  - Chemistry lab TE-VSC-SCC: Benoit Teissandier
  - Engineering Design and Simulation group EN-MME-EDS: Jorge Guardia Valenzuela
  - Materials and Metrology lab EN-MME-MM: Stephan Pfeiffer
- Experts from the material suppliers and component manufactures
  - to remain nameless due to the NDAs…
- Technicians who have helped enormously with the bake-out and testing
  - Lilian Janvier
  - Paulo Jorge Torres Dos Santos
  - Sergio Cuadrado Calzada