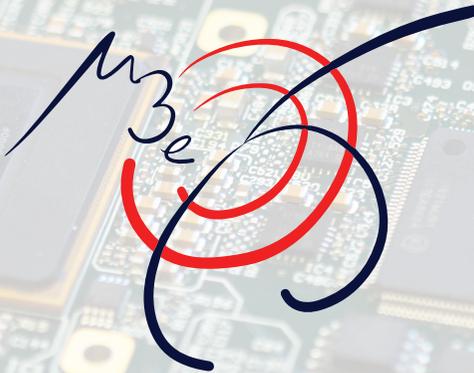


The Mu3e DAQ

A Trigger-less, FPGA-based Readout Approach



Niklaus Berger

Institut für Kernphysik, Johannes-Gutenberg Universität Mainz

Terascale Detector Workshop
March 2023

Overview



Searching for charged lepton flavour violation:

- The Mu3e experiment

100 Gbit/s streaming readout:

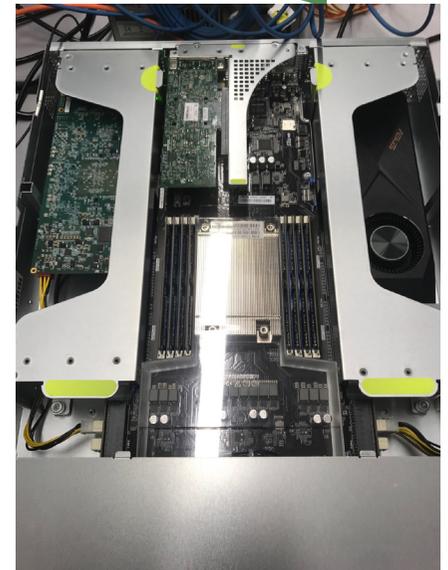
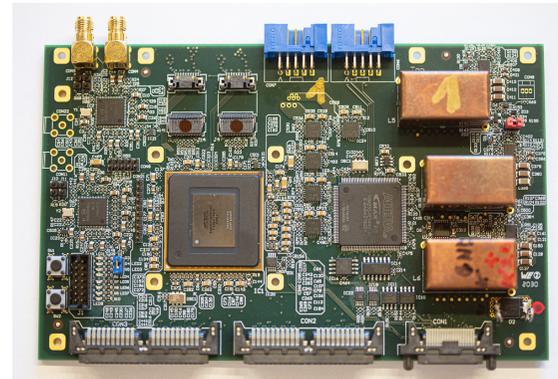
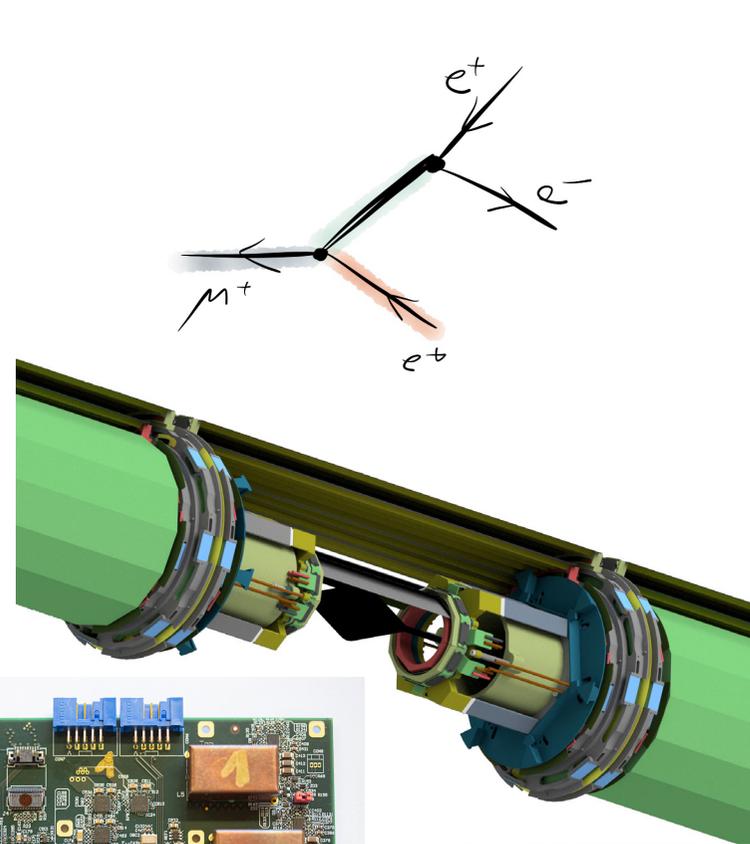
- The Mu3e data acquisition

$> 10^9$ track fits/s on GPUs:

- The Mu3e filter farm

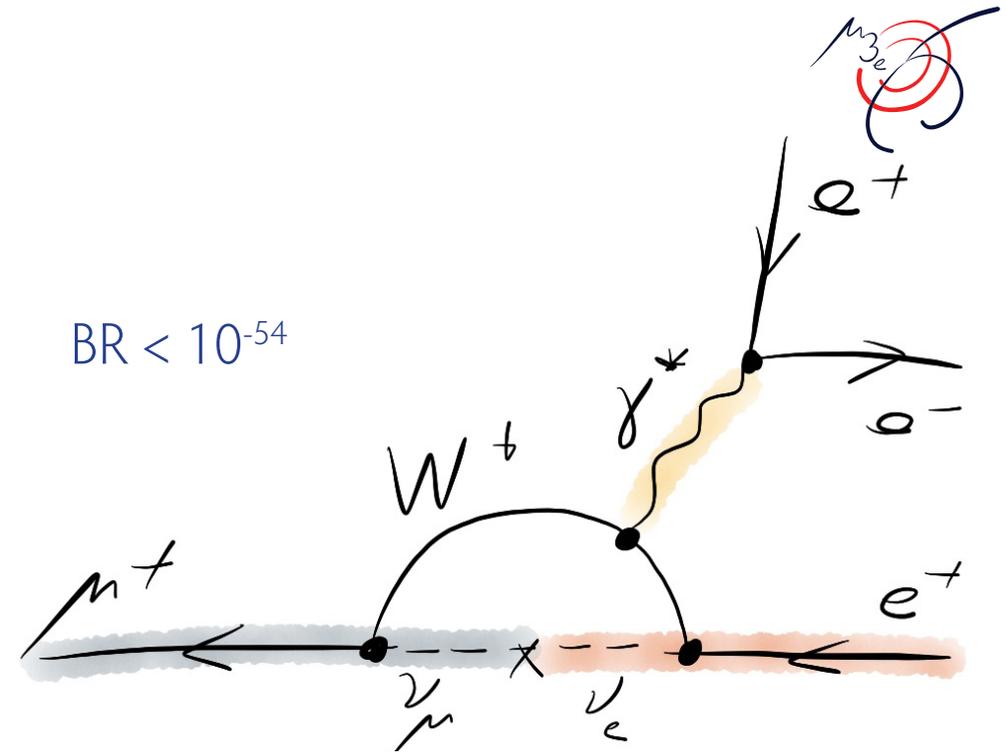
We are not done yet

- Lessons learned so far



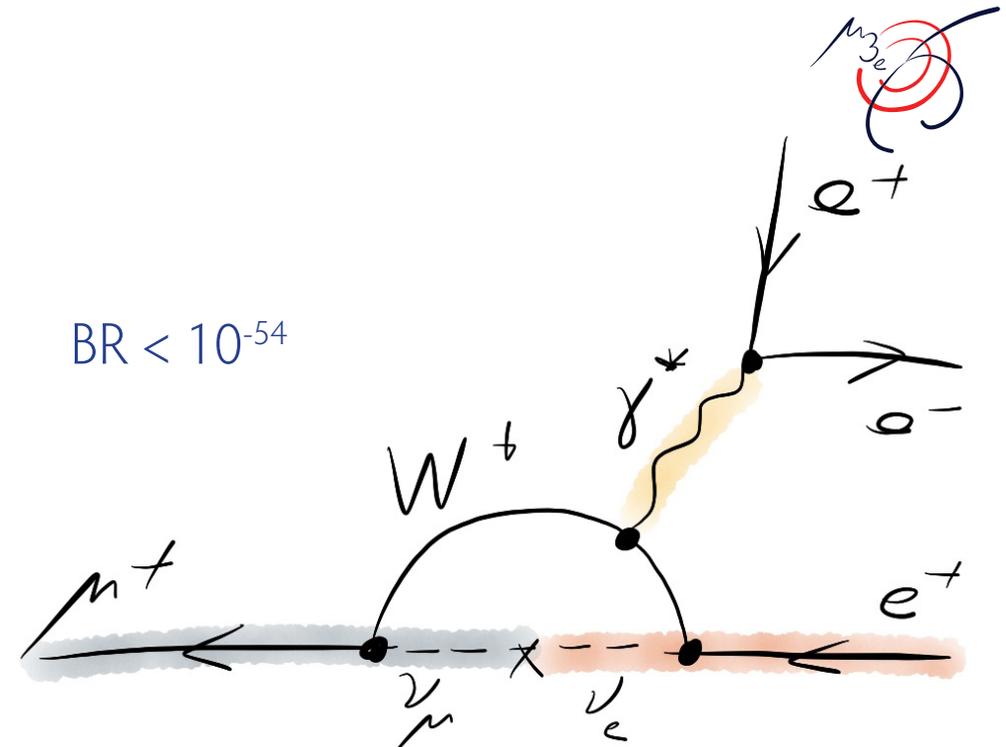
Searching for $\mu^+ \rightarrow e^+e^-e^+$

- Lepton flavour violating muon decays
- Extremely low branching fractions in the Standard Model
- Excellent probes for new physics
- $\text{BR}(\mu^+ \rightarrow e^+e^-e^+) < 10^{-12}$ (SINDRUM, 1988)

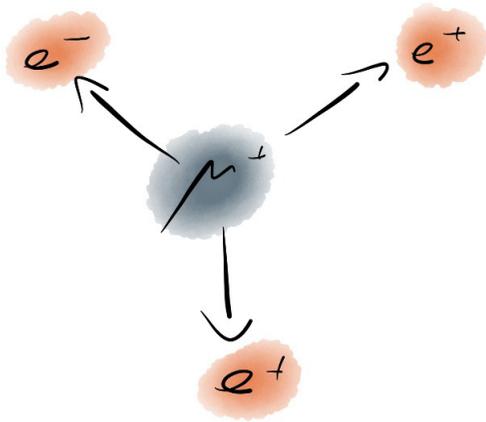


Searching for $\mu^+ \rightarrow e^+e^-e^+$

- Lepton flavour violating muon decays
- Extremely low branching fractions in the Standard Model
- Excellent probes for new physics
- $\text{BR}(\mu^+ \rightarrow e^+e^-e^+) < 10^{-12}$ (SINDRUM, 1988)
- Mu3e aims for a sensitivity of 1 in 10^{16}
- Very intense muon beam: Paul Scherrer Institute (PSI), Villigen, Switzerland
- $2 \cdot 10^{15}$ in a first phase at an existing beam line with 10^8 muons/s - this talk
- Plans for new high-intensity muon beam line (HiMB) with $> 10^9$ muons/s



Signal and Background

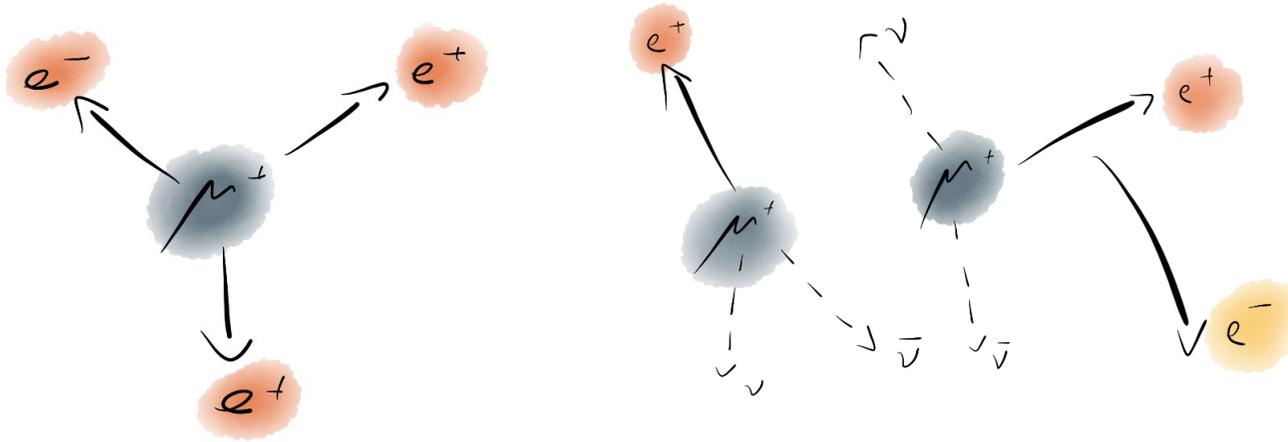


Signal

- $\mu^+ \rightarrow e^+e^-e^+$ at rest
- Two positrons, one electron
- From same vertex
- Same time
- $\sum p_e = m_\mu$
- Maximum momentum:
 $\frac{1}{2} m_\mu = 53 \text{ MeV}/c$



Signal and Background



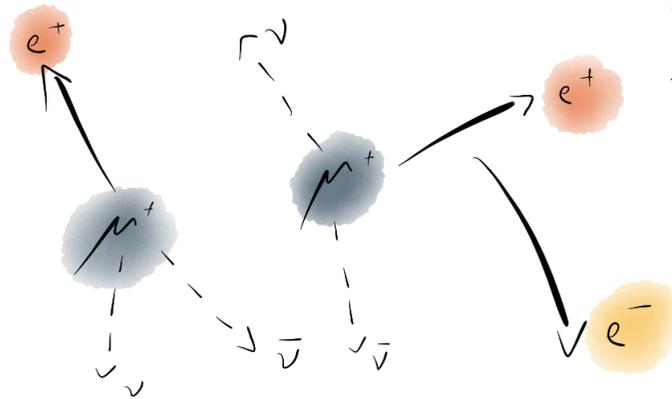
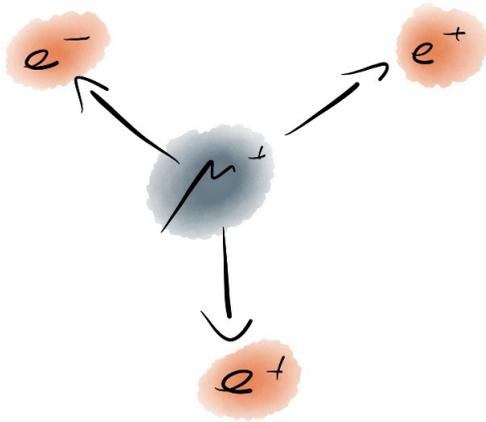
Signal

- $\mu^+ \rightarrow e^+e^-e^+$ at rest
- Two positrons, one electron
- From same vertex
- Same time
- $\sum p_e = m_\mu$
- Maximum momentum:
 $\frac{1}{2} m_\mu = 53 \text{ MeV}/c$

Accidental Background

- Several muon decays
- Plus an electron
- Need good vertexing
- Need good timing

Signal and Background



Signal

- $\mu^+ \rightarrow e^+e^-e^+$ at rest
- Two positrons, one electron
- From same vertex
- Same time
- $\sum p_e = m_\mu$
- Maximum momentum:
 $\frac{1}{2} m_\mu = 53 \text{ MeV}/c$

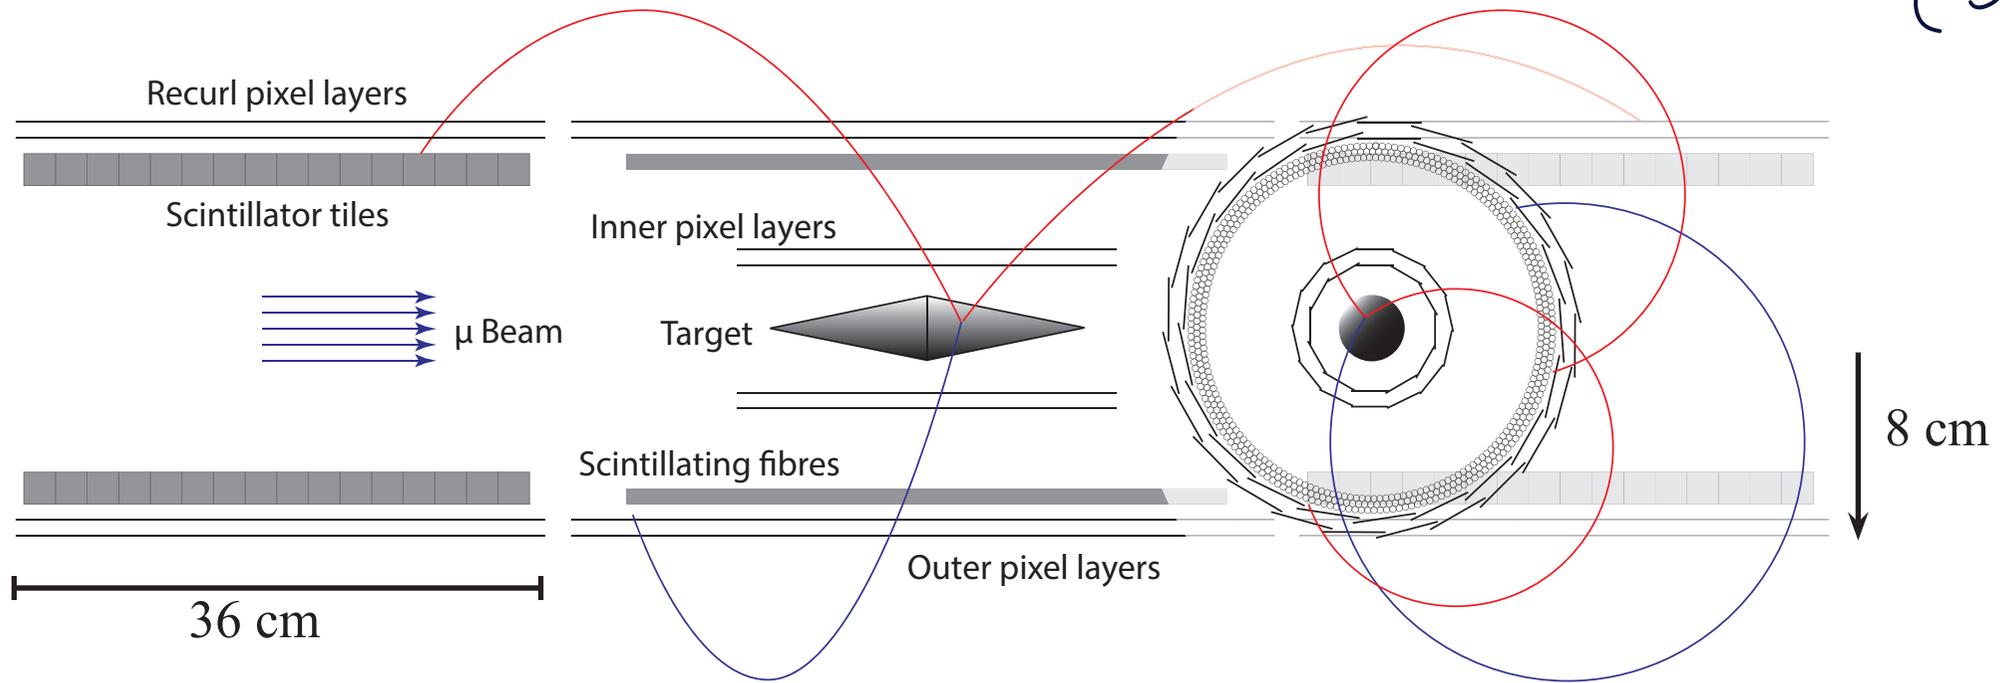
Accidental Background

- Several muon decays
- Plus an electron
- Need good vertexing
- Need good timing

Internal conversion decay

- Allowed rare decay
- $\mu^+ \rightarrow e^+e^-e^+\nu\bar{\nu}$
- Detect missing energy carried by neutrinos
- Need excellent momentum reconstruction

The Mu3e Detector

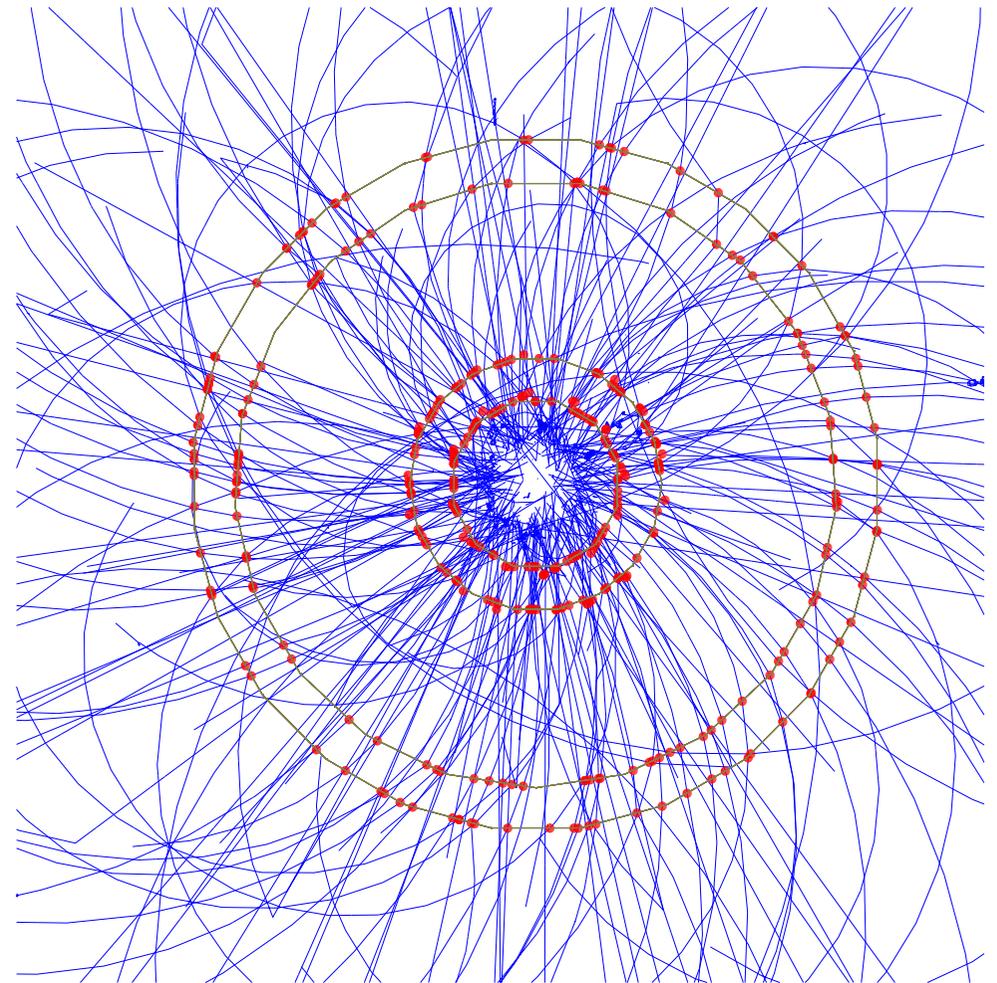


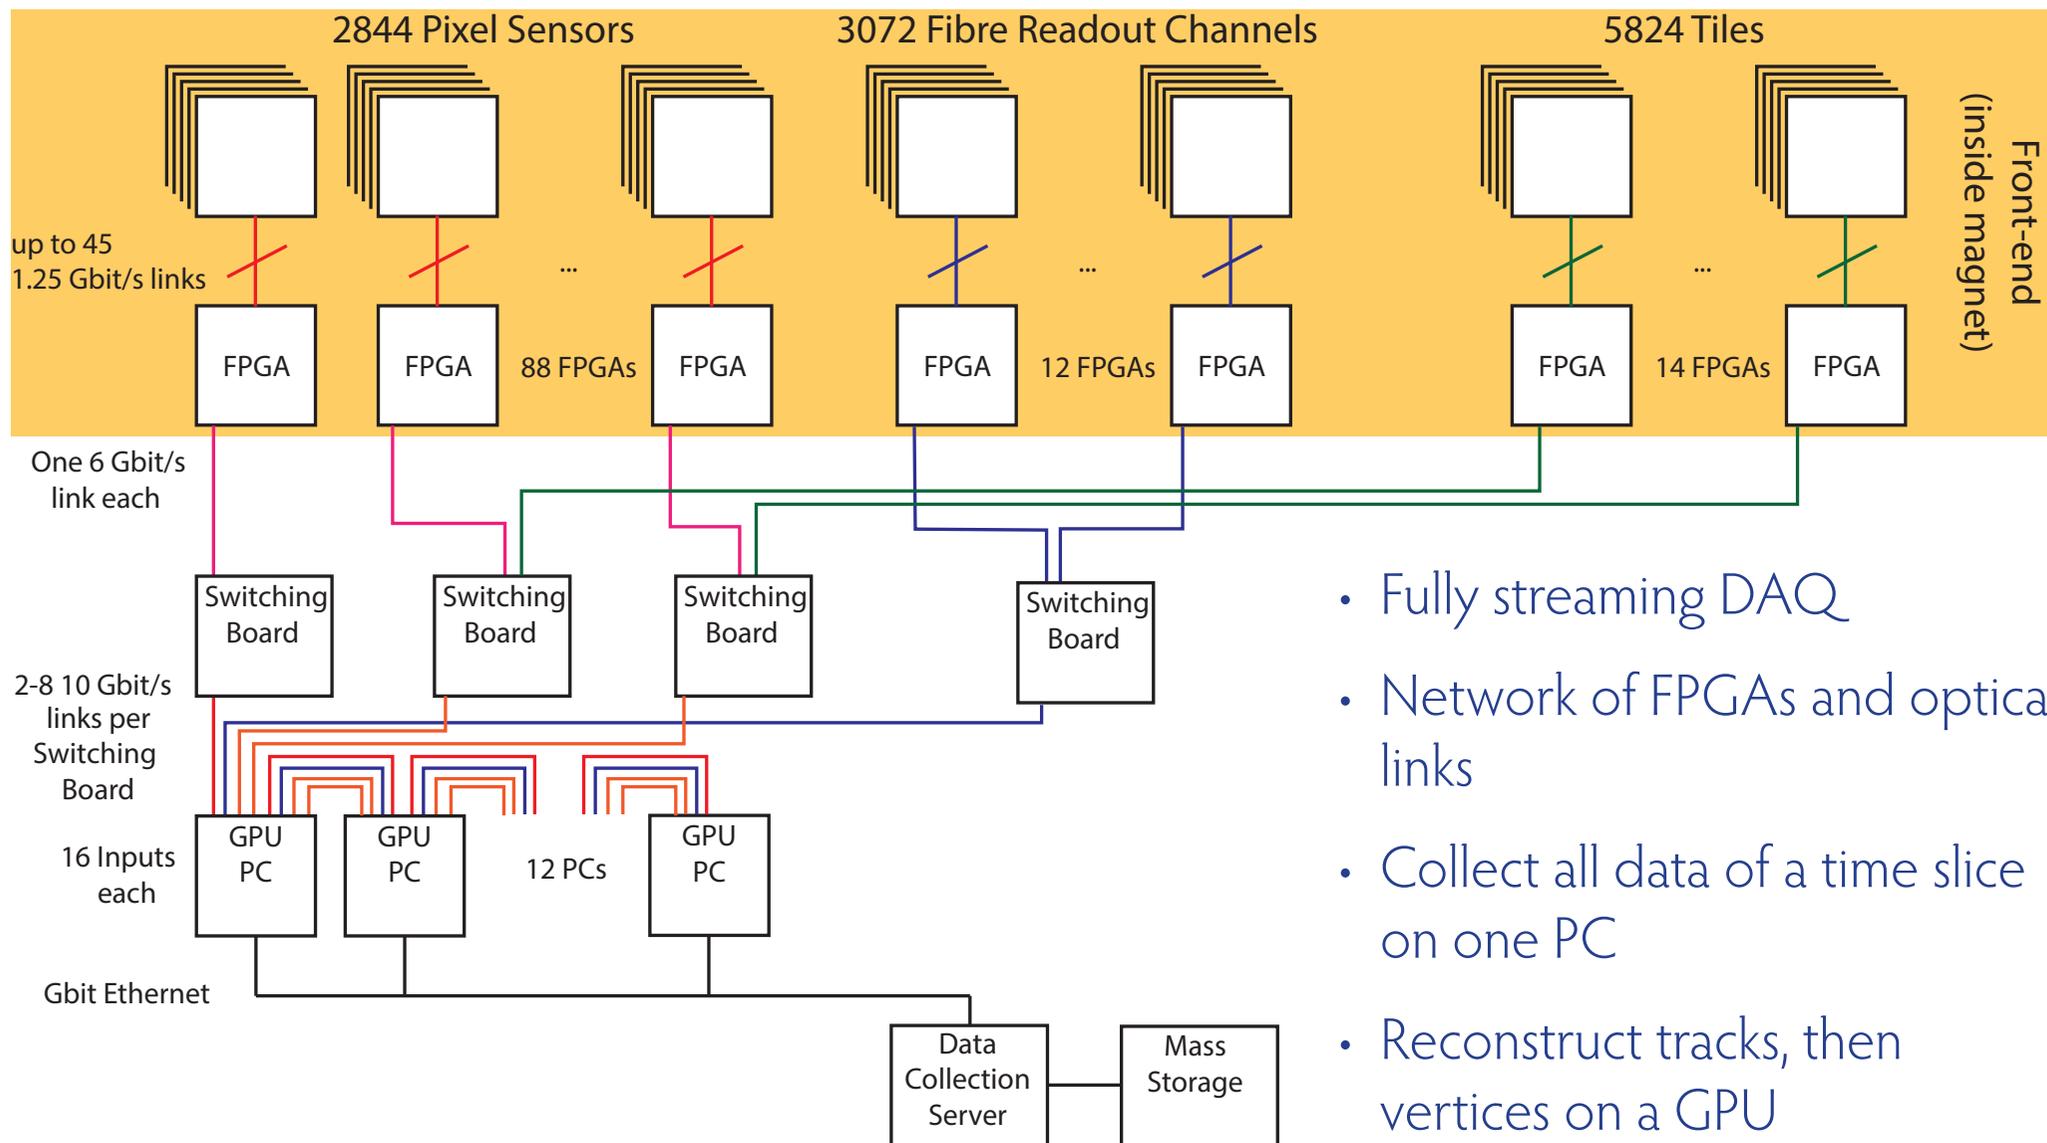
- 1 T solenoid field
- Helium atmosphere to reduce scattering and for cooling
- Minimize material to minimize scattering
- Ultra-thin layers of high-voltage monolithic active pixel sensors (HV-MAPS)
- Scintillating fibres and tiles for improved timing measurements
- Long lever arm of recurling tracks gives precise momentum measurement

Requirements for the data acquisition



- Up to 10^8 muon decays/s
- Highly non-local signal signature
- Rate makes simple three-track coincidences infeasible
- 2844 MuPix sensors with 182 million pixels
- 8896 SiPM readout channels - 278 MuTrig TDC ASICs
- ~ 100 Gbit/s data after zero suppression on ASICs
- Can write about 100 MB/s to mass storage



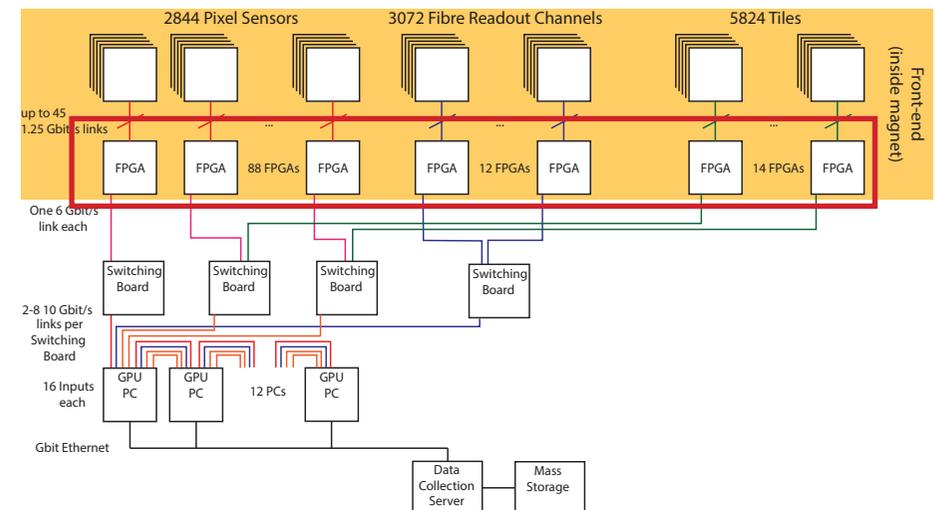
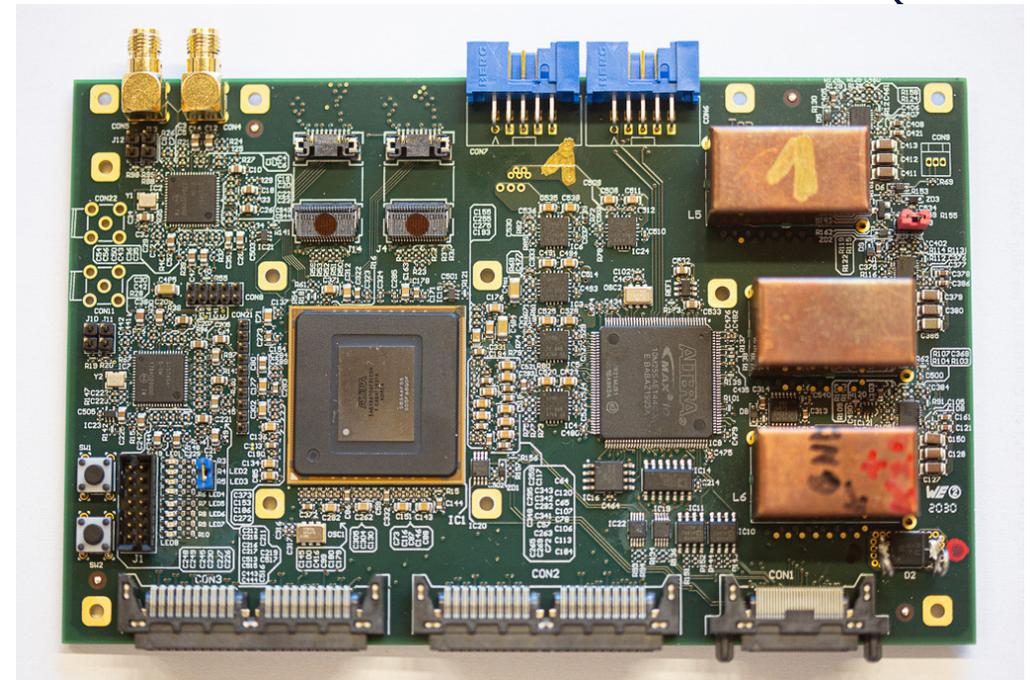


- Fully streaming DAQ
- Network of FPGAs and optical links
- Collect all data of a time slice on one PC
- Reconstruct tracks, then vertices on a GPU
- Write interesting events to disk

Front-end board

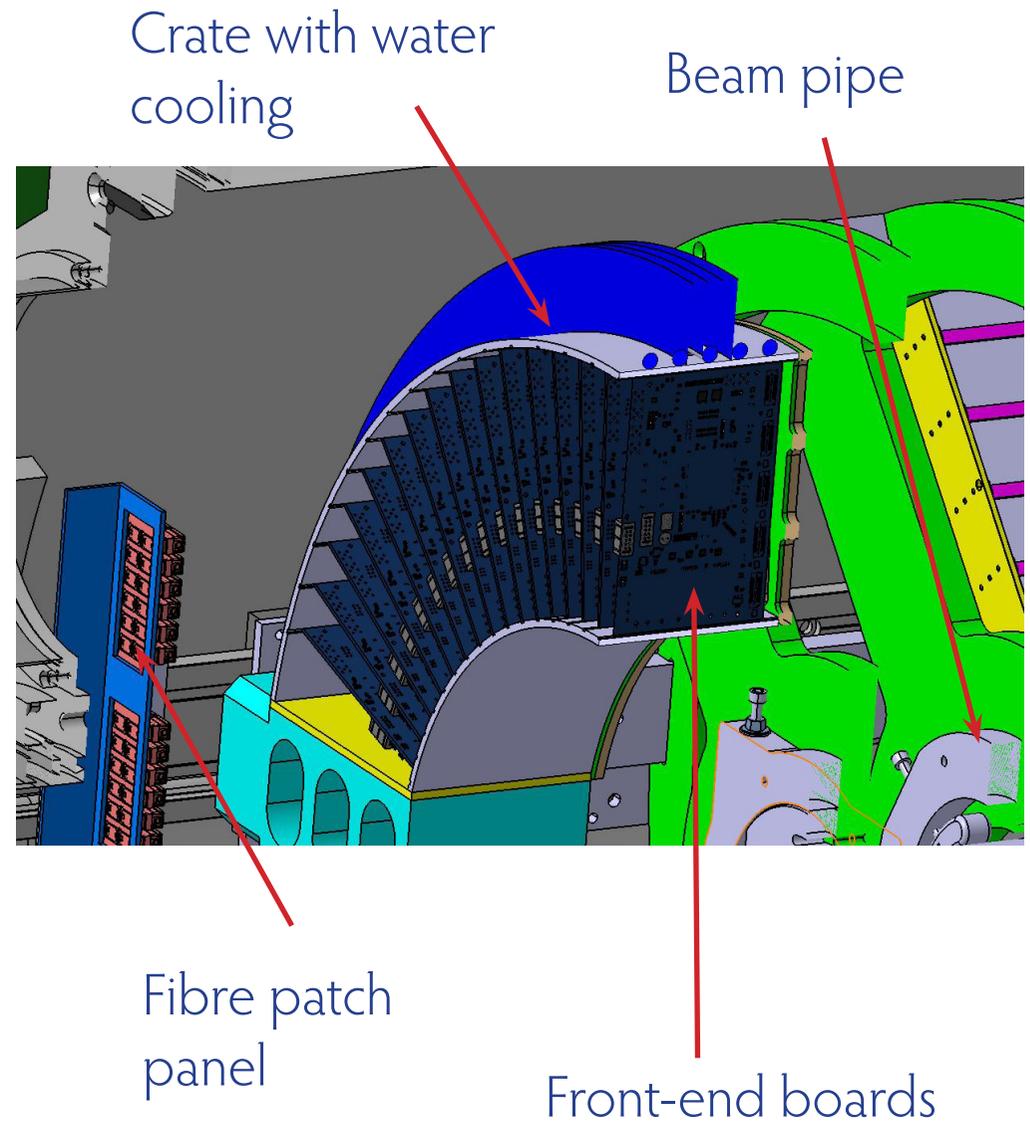


- Operates in magnet and helium atmosphere, space is tight
- Up to 45 1.25 Gbit/s LVDS inputs from detector ASICs
- Intel Arria V A7 FPGA for time-sorting and clustering of hits
- Output to a 6 Gbit/s optical link on a Samtec Firefly Transceiver
- Two SiLabs 5345 jitter cleaners and clock multipliers provide FPGA and detector clocks
- Intel MAX10 FPGA for configuration and monitoring
- Air-coil DC/DC converters for powering



Front-end board

- Mounted in quarter-circular crates inside the 1 m diameter solenoid
- Backplane for control connections and connection to detector
- Adaptors on back of backplane for detector specific cabling
- Aluminium cooling plates connected to water-cooled crate with heat pipes
- ~ 1000 multi-mode optical fibres to the outside world



Lessons learned



- Prototype your main firmware algorithm early (started hit sorting almost 10 years ago)

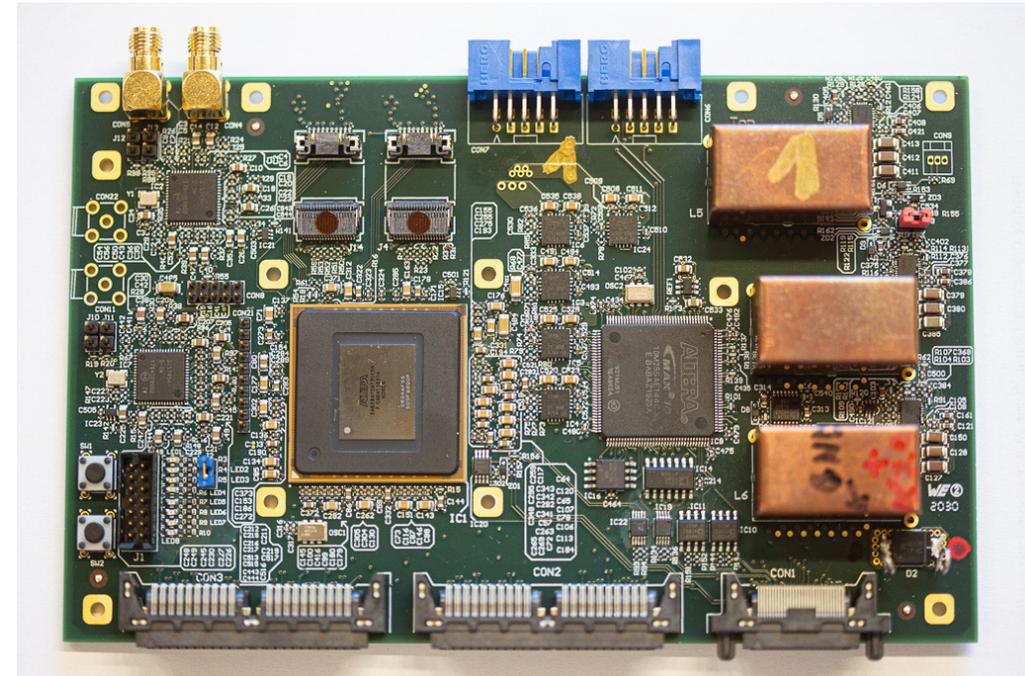
Correct dimensioning of FPGAs

Time for re-writes and debugging

- Think well about programming FPGAs not physically reachable during running (optics works great, *SPI very slow*)

- *Design around available parts (Buy before designing?)*

- Optical links are lovely, electrical ones tricky

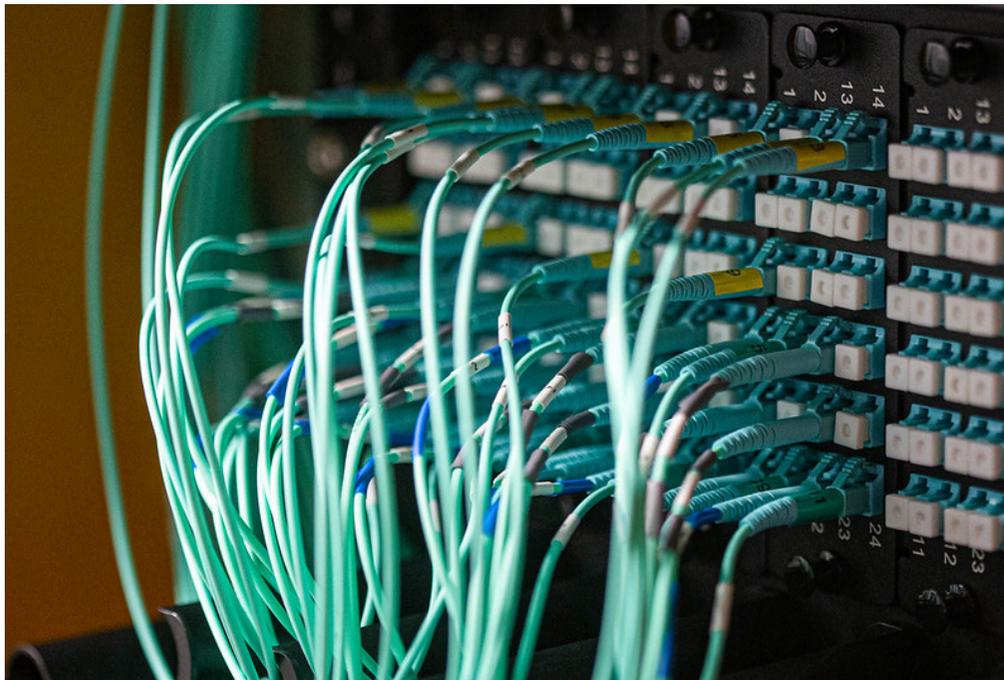


- Having plenty of bandwidth also to the detector is nice - we believe we can configure 300 M pixels in < 4 s
- Generic board for all subdetectors and adaptor boards works well

Optical cabling



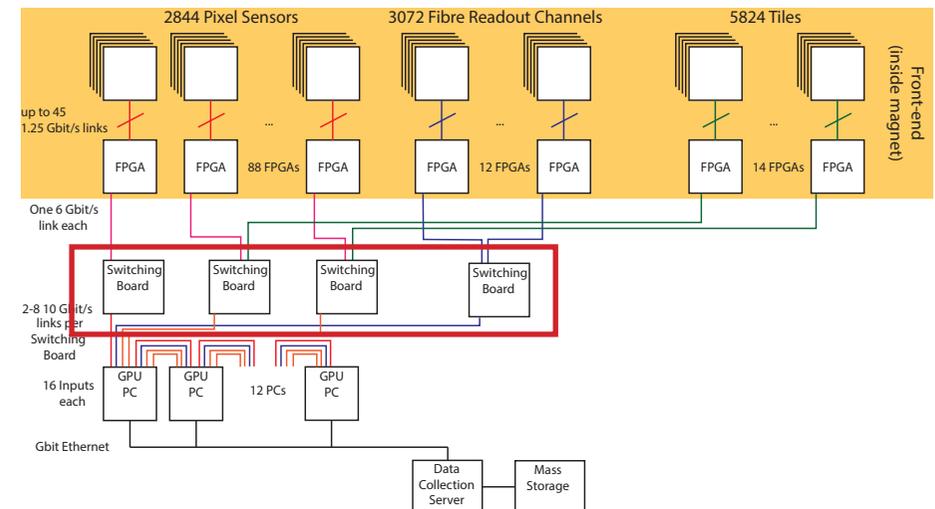
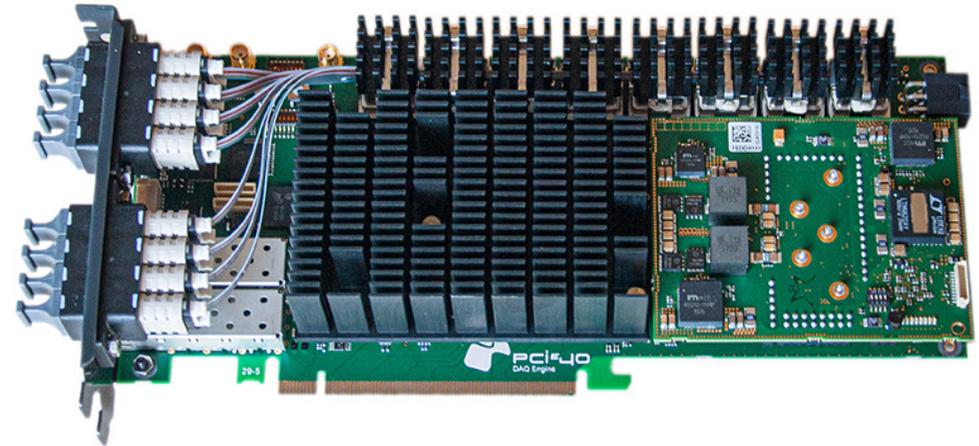
- Nice and compact
- A lot of very convenient and affordable commercial equipment
- 300 Million channel detector read out and controlled with 48 cables (24 fibres each)



Switching board



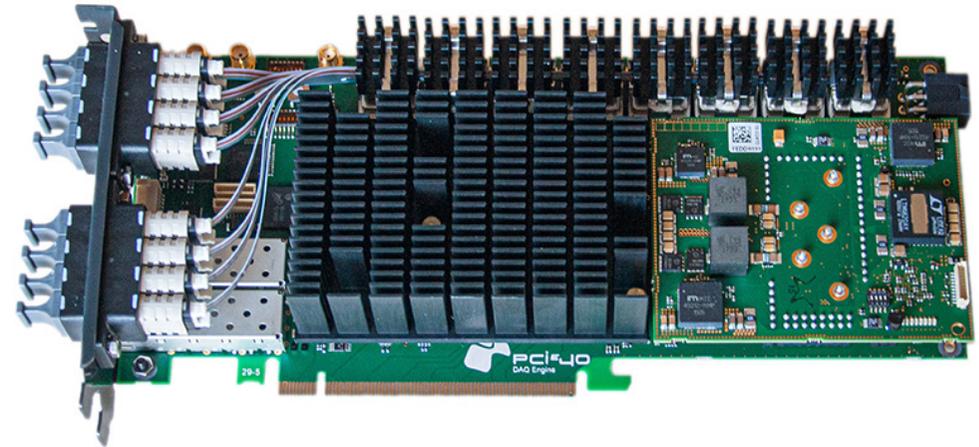
- Operates in a PC case
- Up to 37 front-end board inputs (and control lines)
- Up to eight 10 Gbit/s outputs to filter farm
- Use PCIe40 board developed in Marseilles for LHCb and ALICE upgrades
- Intel Arria 10 - 115 FPGA
- Avago MiniPod Transmitters and Receivers
- Two 8-lane PCIe 3.0 interfaces (used for control and monitoring data)



Lessons learned



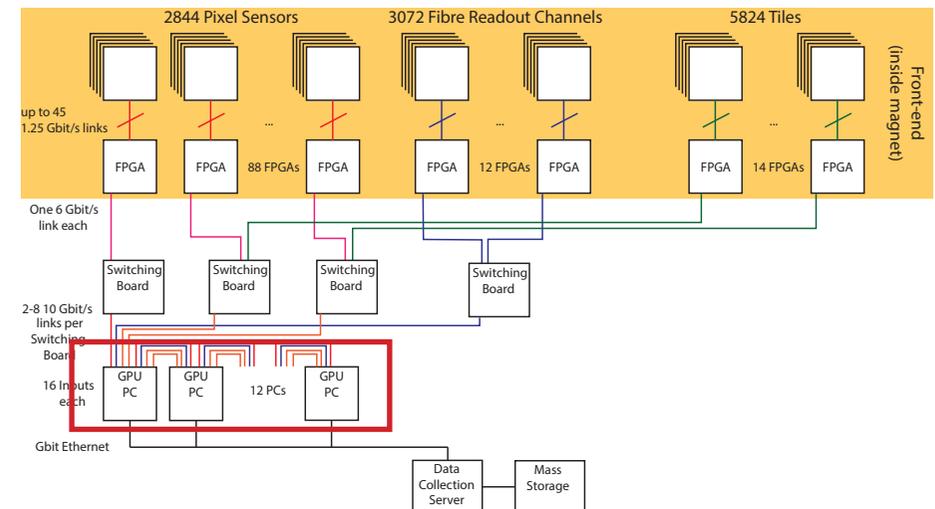
- You do not have to develop everything yourself
- Passively cooling $> 100\text{ W}$ on a card is possible, but loud
- Writing our own generic PCIe interface and driver framework was extremely helpful
(and time consuming)
Be aware of what you get into when one end of your interface is a commercial standard with a few decades of history



Receiving board



- Operates in a PC case, together with a GPU
- 16 10 Gbit/s inputs and outputs (daisy chain)
- Use commercial DE5A NET board from Terasic Inc.
- Intel Arria 10 - 115 FPGA
- DDR 3/4 memory for buffering
- QSFP Transmitters and Receivers
- 8-lane PCIe 3.0 interface



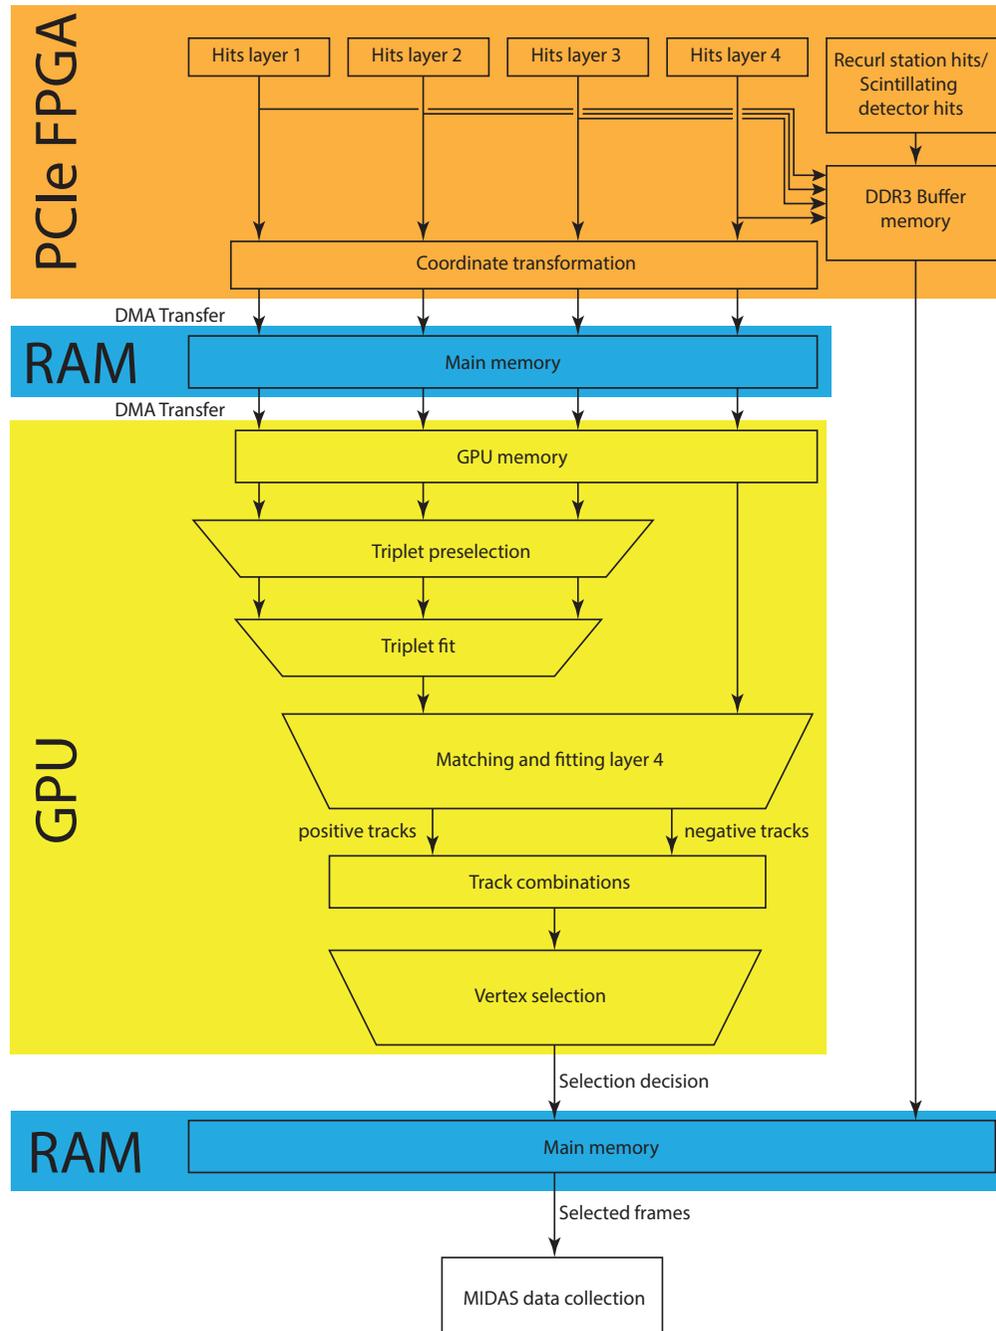
Lessons learned



- You do not have to develop everything yourself
- Some manufacturers have very attractive university programs
- Fast (DDR3/4) memory interfaces are tricky
- PCIe Interface (see above)

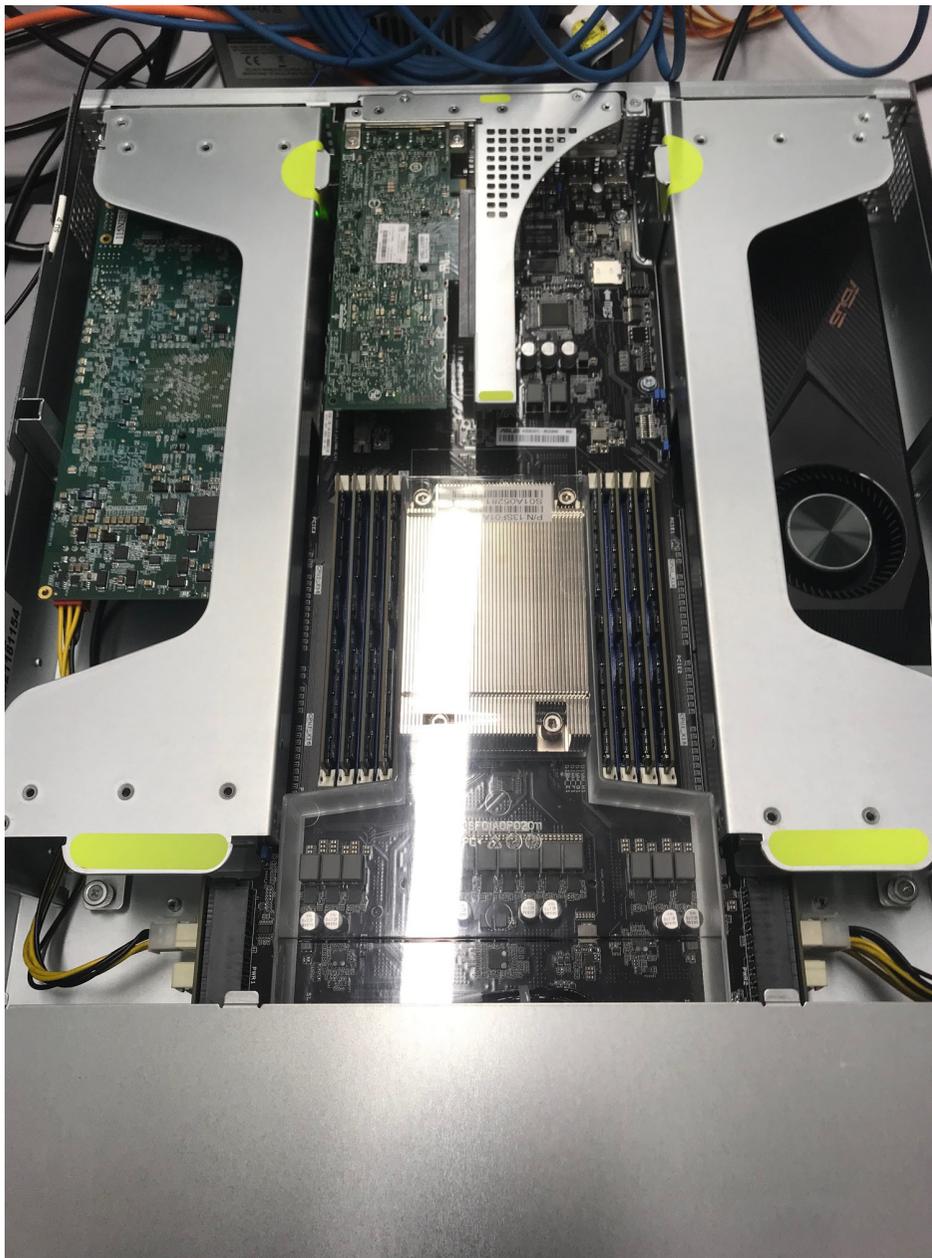


Farm data flow



- Buffer all incoming data in DDR memory
- Use subset from central detector for track and vertex finding on a GPU
- If interesting: Get full data from buffer, send to PC
- Up to 38 Gbit/s PCIe DMA transfers using custom firmware and driver
- After full reconstruction: Send off to mass storage
- Use the MIDAS software for data collection, detector control and monitoring etc.

GPU reconstruction



- GPU reconstruction on gaming cards
- Have achieved $> 10^9$ track fits/s per GPUs (Nvidia GTX 980)
- Twelve GTX 1080Ti are sufficient for dealing with 10^8 muon decays/s
- Excited about the possibilities with the latest cards...

Lessons being learned



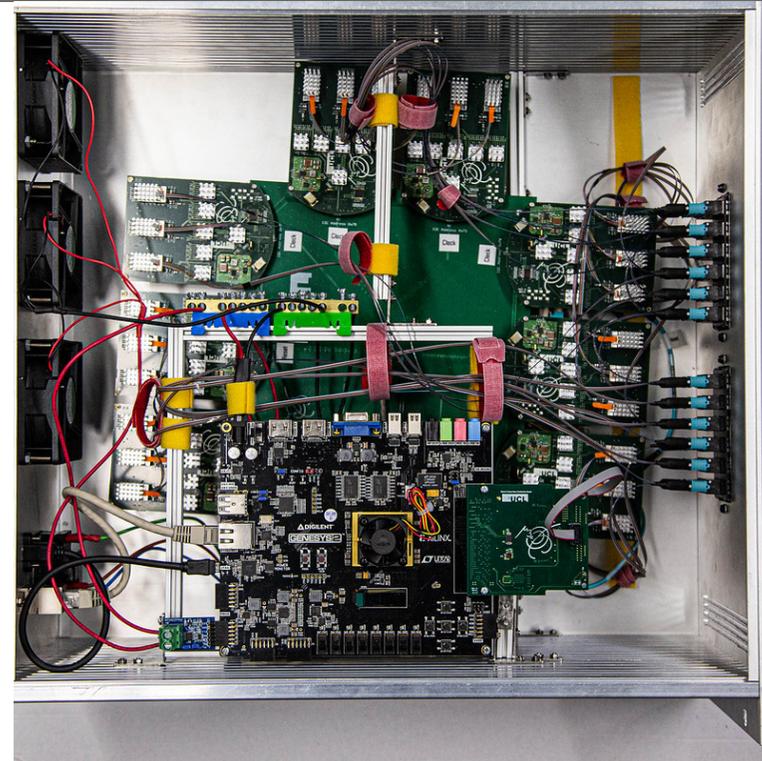
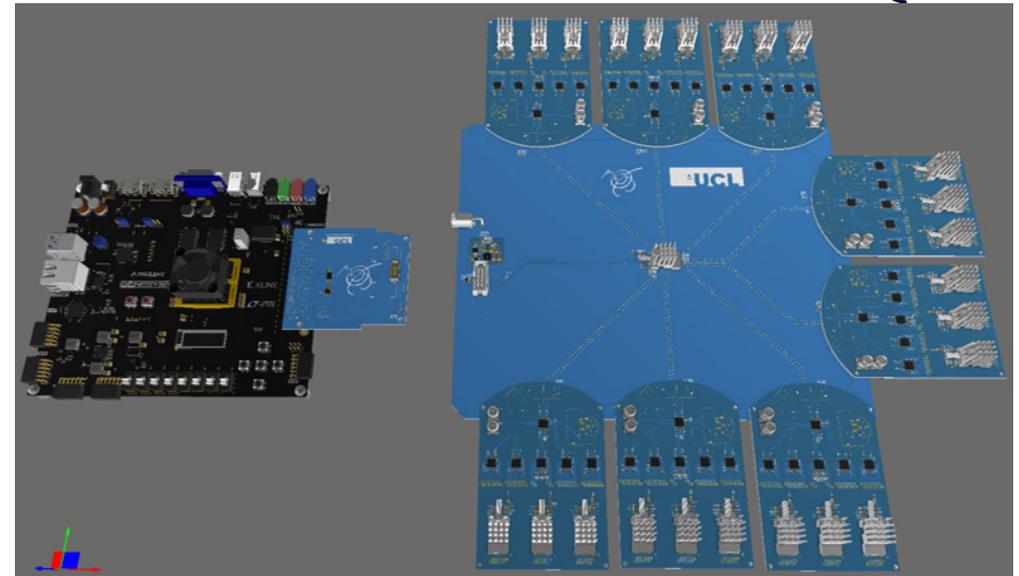
- GPUs get faster and cheaper all the time
- Except if they do not
- If everything else is also a bit late, things improve again
- Optimizing algorithms has to happen for every GPU generation again



System synchronization



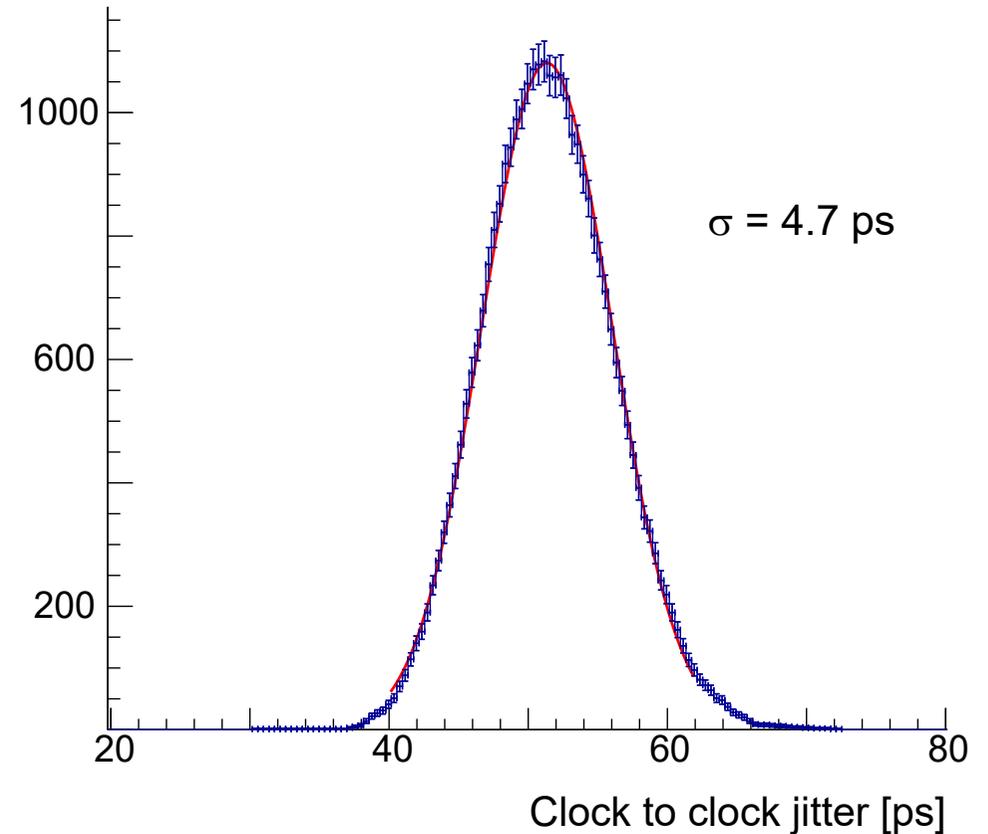
- Produce 144 copies of the 125 MHz system clock
- Produce 144 copies of the 1.25 Gbit/s, 8bit/10bit encoded reset and state transition signal
- Digilent Genesys FPGA board
- Samtec Firefly optical transmitters



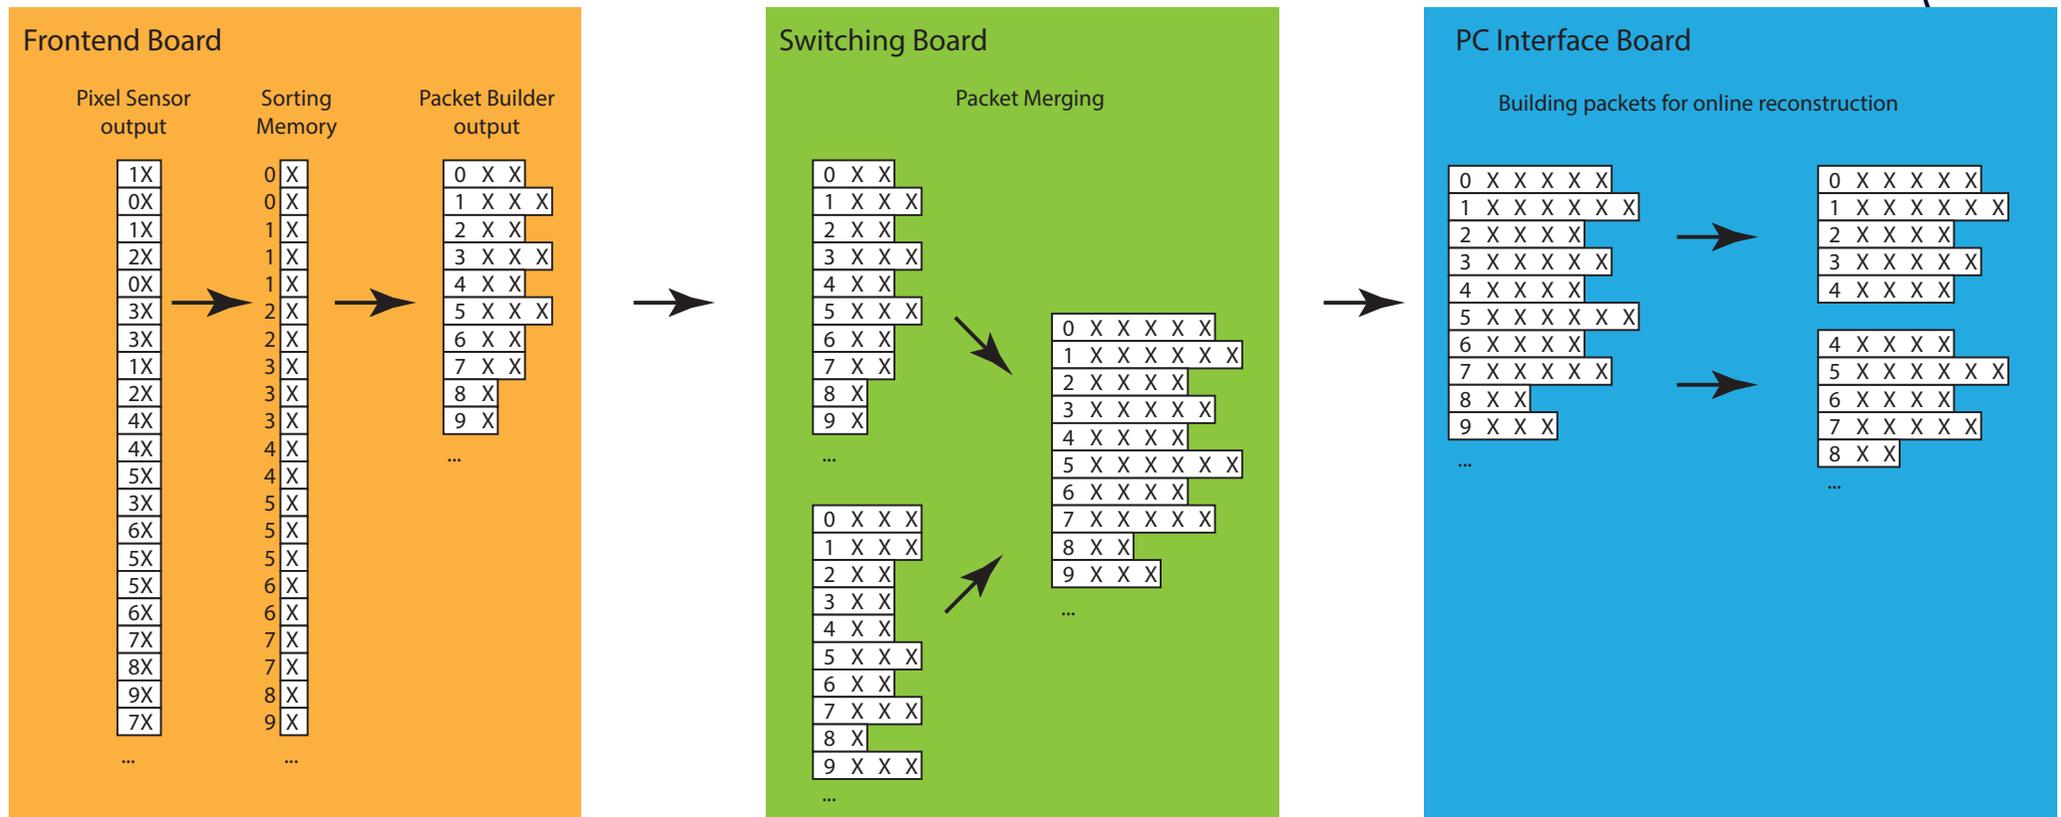
System synchronization



- Produce 144 copies of the 125 MHz system clock
- Produce 144 copies of the 1.25 Gbit/s, 8bit/10bit encoded reset and state transition signal
- Digilent Genesys FPGA board
- Samtec Firefly optical transmitters
- Less than 10 ps clock-to-clock jitter



Data synchronisation



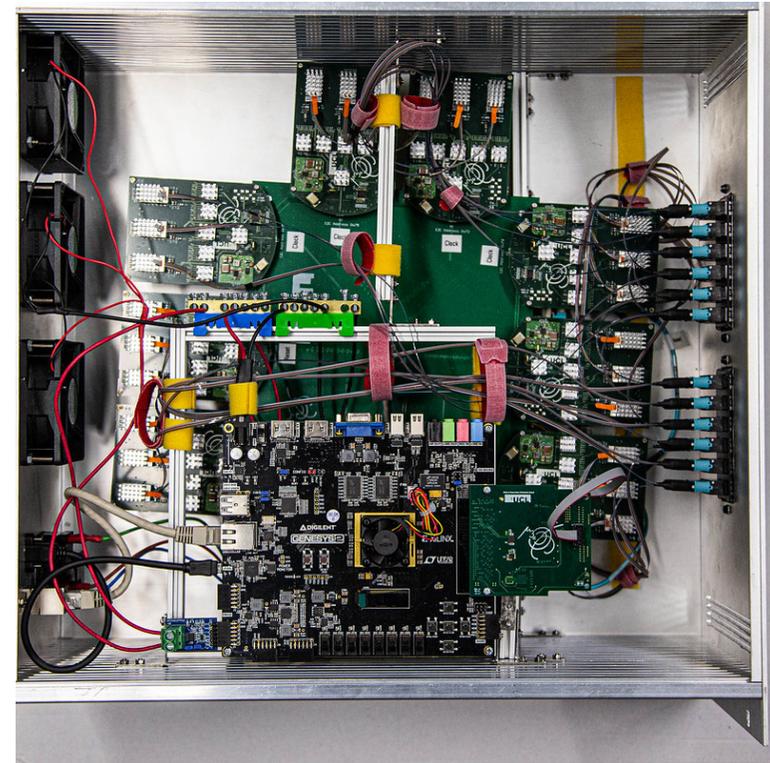
Thinking about streaming readout (vs trigger):

- Latency does matter little
(have big enough FIFOs everywhere)
- Bandwidth either needs to be constant across the system or data need to be rejected in a well controlled way

Lessons learned



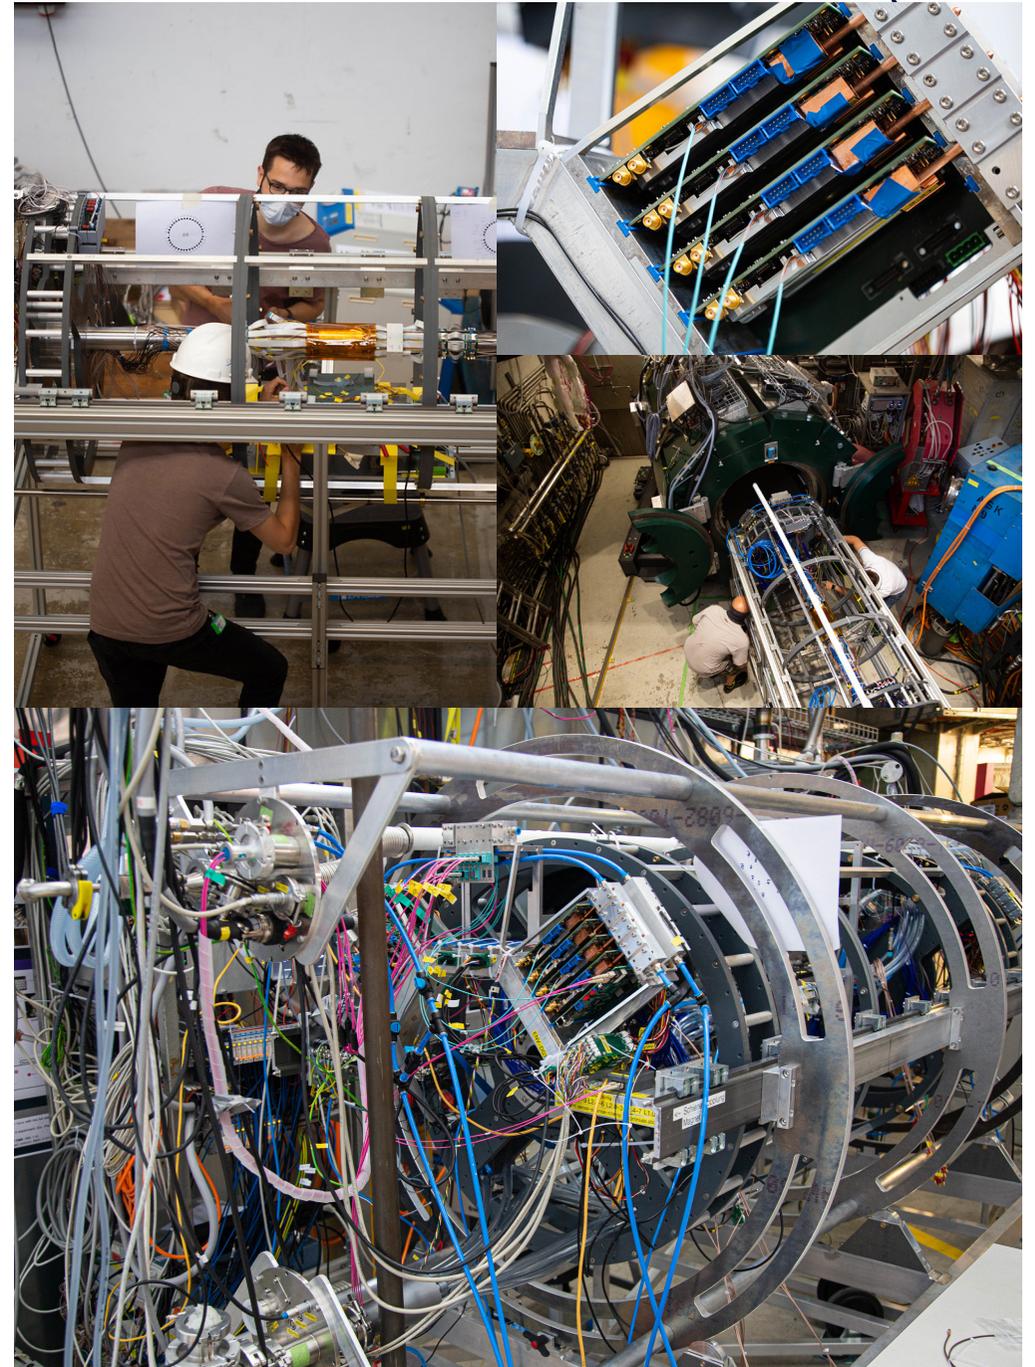
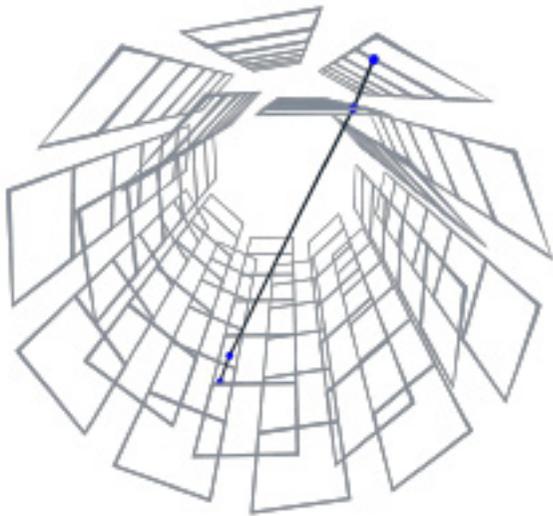
- You can get less than 5 ps jitter with standard commercial components
- Don't use a single Xilinx FPGA in an Intel/Altera system
- Don't use complicated interfaces that are not needed in your context (Here: IPbus) - reduce dependencies
- Make sure you have a simple replacement system if you operate a small slice of the detector



System integration

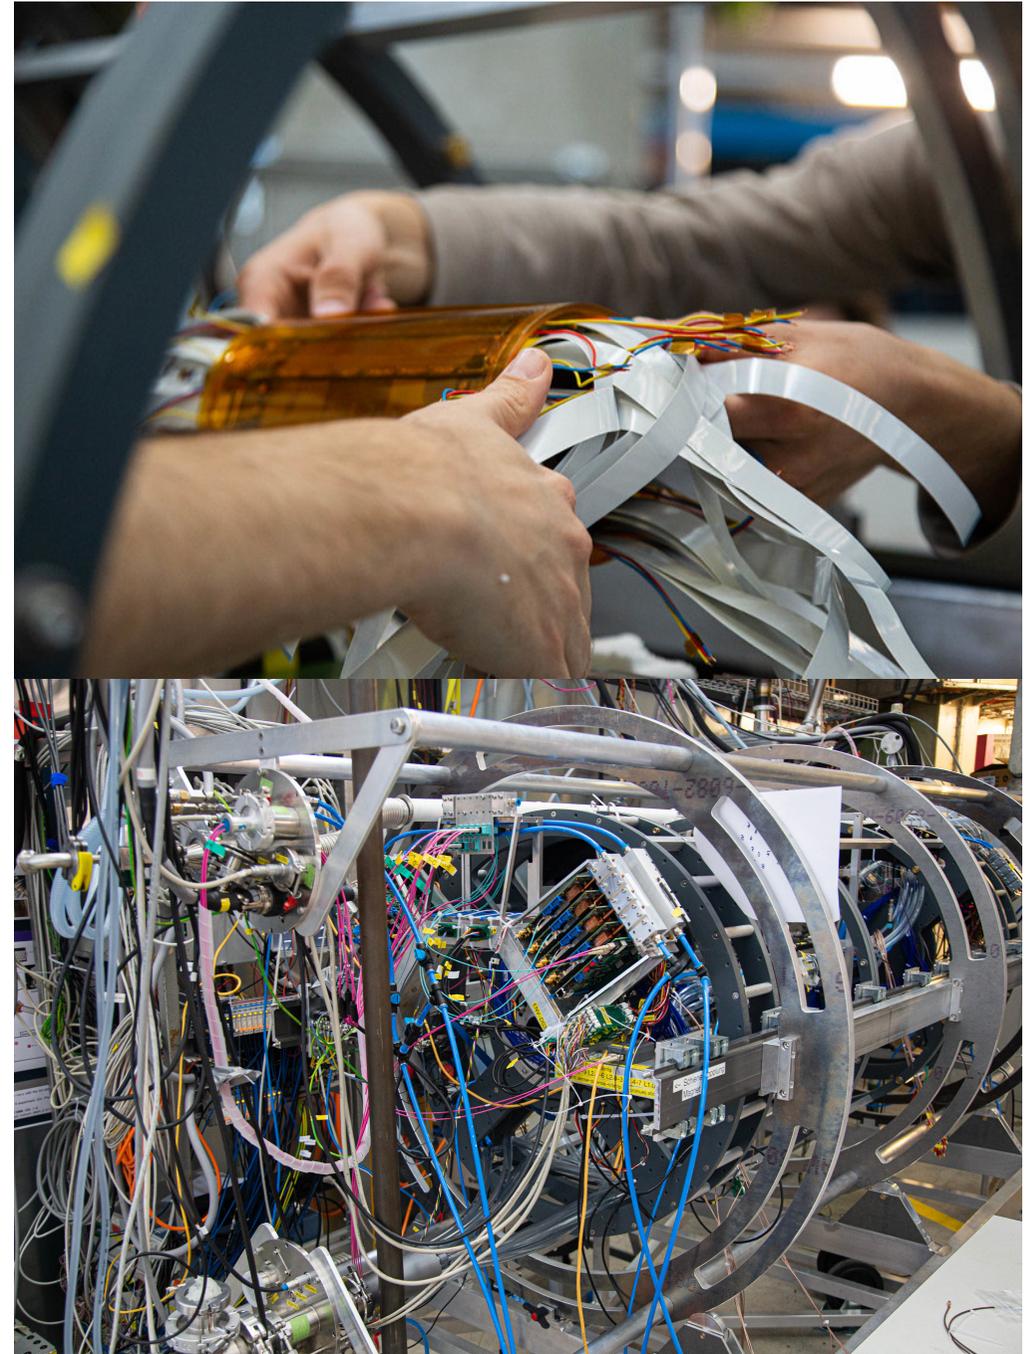


- Used many parts of the system for detector test(-beams)
- Started doing integration tests with sub-detectors very early (DAQ-weeks in Mainz)
- Interrupted by Corona
- Did two very extensive integration runs at PSI 2021 and 2022 (one with beam, one with cosmics)



Lessons learned

- Start early with integrating detector ASICs and DAQ
Needs ASIC and DAQ experts physically at the same place
- Documentation, documentation, documentation - especially of interfaces
- There is never enough preparation time for a run
- There needs to be a run to end preparation time
- Monitoring, monitoring, monitoring - finding out that it works is as much effort as making it work
- At some point you need professional cabling



Lessons learned - Firmware & Software



- Start early
- Good scripting of your firmware compilation is great
- Have tools to make things automatically consistent between firm- and software (e.g. register maps)
- Continuous integration is great (Is someone automatically uploading to FPGAs? connected to “real” detectors?)
- Physicists writing firmware is great (Engineers & Computer Scientists: Also great) - but we need more of them

7 of 7 builds passed

- ✓ A10Farm
A10Farm · 6 days ago
- ✓ A10Switch
A10Switch · 6 days ago
- ✓ CI
CI · 6 days ago
- ✓ Max10
Max10 · 6 days ago
- ✓ MupixFEB
MupixFEB · 6 days ago
- ✓ SciFiFEB
SciFiFEB · 6 days ago
- ✓ SciTileFEB
SciTileFEB · 6 days ago

Summary



- Mu3e is searching for charged lepton flavour violation:
Aiming for a sensitivity for $\mu \rightarrow eee$ of one decay in 10^{16}
- Mu3e Phase I:
Search for $\mu \rightarrow eee$ with a sensitivity of $2 \cdot 10^{-15}$
- 10^8 muons/s and 100 Gbit/s data
- Mu3e DAQ:
Optical links and FPGAs for transporting and sorting data
- Mu3e filter farm:
> 10^9 tracks/s reconstructed on just a dozen GPUs
- Now putting everything together:
the real fun is just starting