Data Flow in the Mu3e DAQ

Marius Köppel[1] on behalf of the Mu3e collaboration

Abstract—The Mu3e experiment at the Paul Scherrer Institute (PSI) searches for the charged lepton flavour violating decay $\mu^+ \to e^+e^+e^-$. The experiment aims for an ultimate sensitivity of one in 10^{16} μ decays. The first phase of the experiment, currently under construction, will reach a branching ratio sensitivity of $2 \cdot 10^{-15}$ by observing 10^8 μ decays per second over a year of data taking. The highly granular detector based on thin high-voltage monolithic active pixel sensors (HV-MAPS) and scintillating timing detectors will produce about 100 Gbit/s of data at these particle rates. The Field Programmable Gate Array (FPGA) based Mu3e Data Acquisition System (DAQ) will read out the different detector parts. The trigger-less readout system is used to sort, time align and analyse the data while running. A farm of PCs equipped with powerful graphics processing units (GPUs) will perform the event reconstruction and data reduction. This paper presents the ongoing integration of the sub detectors into the DAQ system, especially focusing on the time aligning and the data flow inside the FPGAs of the filter farm. It discusses the DAQ system used in the Mu3e Integration Runs performed in spring 2021 and 2022, while providing results of the synchronisation of the different detectors.

Index Terms—Data acquisition, Field programmable gate arrays, Physics

I. THE MU3E EXPERIMENT

THE Mu3e experiment [1] at the Paul Scherrer Institute (PSI) searches for the charged lepton flavour violating decay $\mu^+ \to e^+e^+e^-$. The experiment aims for an ultimate sensitivity of one in 10^{16} μ decays. This would improve the current limit of $\mathcal{B}(\mu^+ \to e^+e^+e^-) < 1.0 \times 10^{-12}$, set by the SINDRUM experiment [2], by four orders of magnitude. The first phase of the experiment, currently under construction, will reach a branching ratio sensitivity of $2 \cdot 10^{-15}$ by observing $10^8\mu$ decays per second over a year of data taking. The planned second phase depends upon a new High intensity Muon Beam-line (HiMB) [3] which will provide $10^9 - 10^{10}$ μ per second.

The 590 MeV proton beam at PSI hits a graphite target, producing secondary pions, which decay into muons. The resulting 29.8 MeV/c muon beam is further guided into the
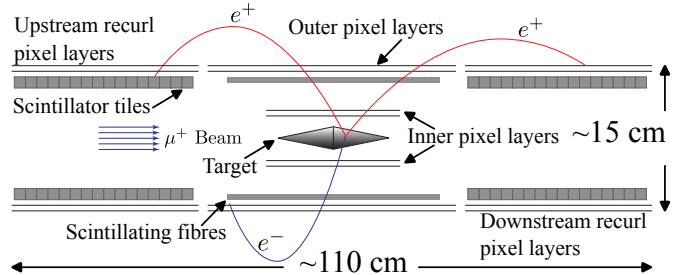
Fig. 1. A sketch of the phase I Mu3e detector concept is shown. The 110 cm long, cylindrical detector is cut along the beam axis. The incoming muons are stopped on a target and decay at rest. The target is surrounded by a two-layer vertex detector built out of MuPix pixel sensors. The second detector station holds scintillating fibres used for the timing measurement. Further out, a second double layer of pixel sensors is mounted. Additionally, two upstream and downstream recurl stations holding a third double layer of pixel sensors and scintillating tiles are used to improve the momentum and timing measurements for particles curling back in the 1 T magnetic field.

Mu3e detector, sitting inside a 1 T solenoid magnet, and stopped by a hollow double cone target made out of Mylar foil. These stopped muons decay at rest and trajectories of their decay products are measured using a silicon pixel tracking detector.

In Figure 1, a sketch of the Mu3e detector is shown. The detector is built of three parts, one inner central station and two recurl stations, one upstream and one downstream. The recurl stations improve the momentum resolution by providing additional track points when the particles curl back towards the beam axis in the magnetic field.

Multiple scattering is the largest contribution to the momentum resolution. Therefore, it is crucial to reduce the material budget. Hence, a new type of thin high-voltage monolithic active pixel sensors (HV-MAPS), called MuPix, is used, which can be produced as thin as 50 μm [4]–[10]. To have precise time-of-flight measurement and to suppress the background from random accidental coincidental decays, a precise timing measurement is needed. Therefore, a scintillating fibre detector [11], placed in the central station of the detector, and scintillating tile detectors [12], located in the two recurl stations, are used. These detectors are read out via silicon photomultipliers and a custom-designed Application-Specific Integrated Circuit (ASIC), called MuTRiG [13].

To be able to handle the 100 Gbit/s of data, which are produced by the pixel and scintillating timing detectors, an online selection needs to be performed. A farm of PCs equipped with powerful graphics processing units (GPUs) will perform the event reconstruction and data reduction.
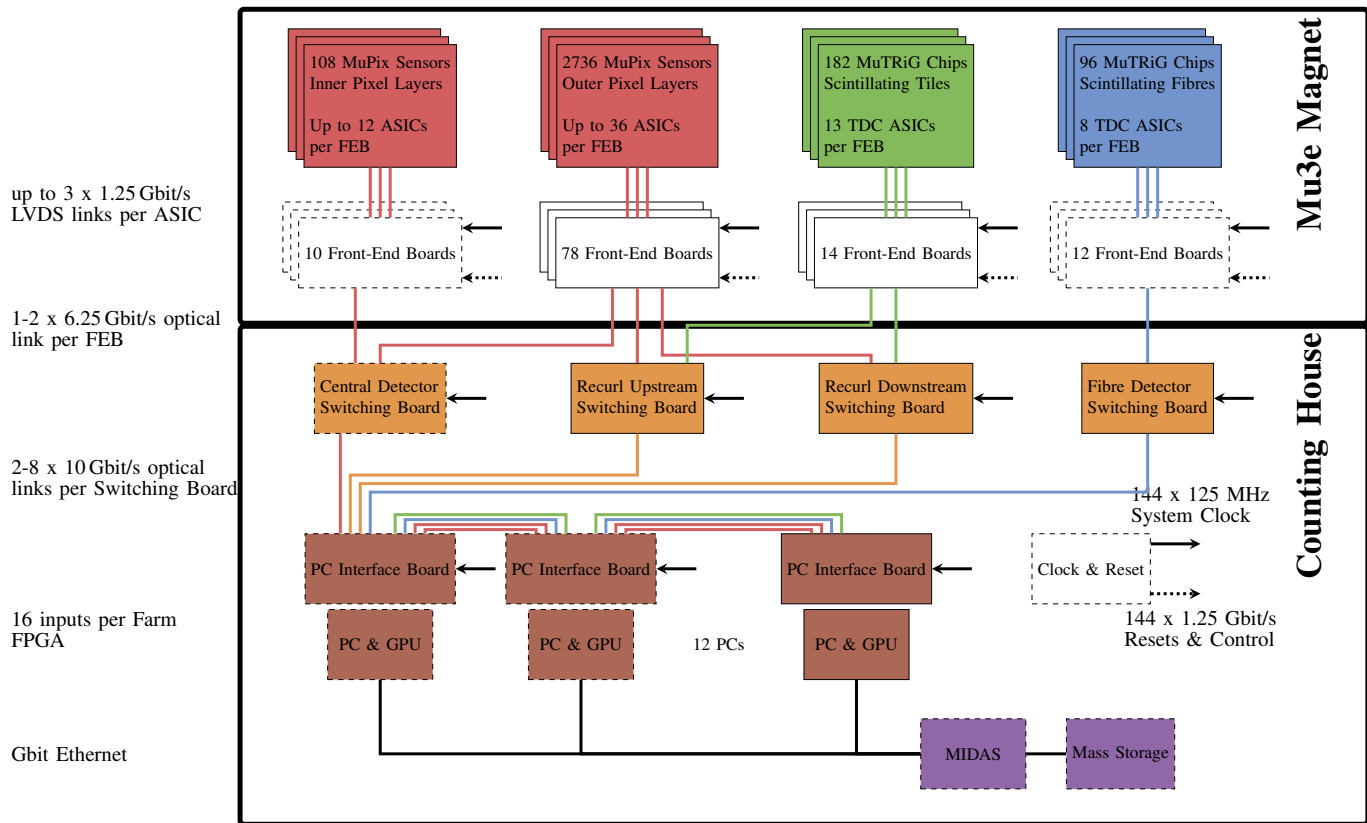
Fig. 2. A sketch of the final Mu3e data acquisition system is shown. For the Integration Run 2021 and Cosmic Run 2022, the dashed parts of the system were used. The whole system is synchronised to a global 125 MHz system clock using a clock and reset system holding the Genesys-2 FPGA board. This system is producing a 1.25 Gbit/s reset and control line, which is used for run transitions. The inner vertex detector is read out via ten FEBs while the fibre detector is read out via two FEBs. The switching layer of the DAQ system holds one PCIe40 board (Switching Board) which is receiving the data from the FEBs, synchronises them and forwards it off to a minimal farm holding two Terasic DE5 boards (PC Interface Board).

II. Mu3e Data Acquisition System

In Figure 2, the final Mu3e DAQ system is shown. The total number of the MuPix sensors used for the phase I detector is 2844 while the scintillating fibres and scintillating tiles are read out by 278 MuTRiG ASICs. The DAQ is designed to be able to handle the expected data rate of 100 Gbit/s for the full detector. Both ASICs (MuPix and MuTRiG) send zero-suppressed and time unsorted hit data over up to three links for the MuPix and one link for the MuTRiG with a bandwidth of 1.25 Gbit/s per link.

The MuPix chips from the inner pixel layer transmit up to 30 Mhits/s over three links, while the outer layers are connected via one link. The hit timestamp of the MuPix has a resolution of 8 ns.

The MuTRiG chip generates a 625 MHz coarse and fine counter with a bin size of 50 ps, to ensure the high time resolution of less than 500 ps for the fibre and less than 100 ps for the tile detector [1]. Both counters are needed for the required timing resolution, but are generated in different ways on the MuTRiG chip. In Section III-A, the special treatment of the coarse counter on the FEB is explained.

The actual DAQ system is built up of three layers of FPGA boards. For reading out the different detectors, the custom developed Front-end boards (FEBs) are used. This board is placed inside the magnetic field and is connected via 1.25 Gbit/s Low Voltage Differential Signaling (LVDS) links to the different sub-detectors. The task of the board is to read out and configure the different detector ASICs, and sort the received hits in time. The second layer of the DAQ system consists of the PCIe40 Board [14] (in the paper the board will be called Switching Board), which was developed for the A Large Ion Collider Experiment (ALICE) and Large Hadron Collider beauty (LHCb) upgrades. These boards are placed outside the magnet and are connected via 6.25 Gbit/s optical links to the FEBs. They perform the time alignment of the different data streams, and they send the detector configuration via optical connections to the FEBs. The last layer of the system consists of the commercial Terasic DE5a-Net-DDR4 boards [15] sitting in a PC equipped with a GPU (in the following, one of this units will be referred to as the farm PC). The Terasic DE5a-Net-DDR4 boards hold an Arria10 FPGA and an onboard Double Data Rate Synchronous Dynamic Random-Access Memory (DDR SDRAM); in the following, this will be referred to as DDR4-RAM. The DDR4-RAM is employed to buffer the data of the full detector, while the hits of the four layers of the central pixel detector are used to perform an online event selection on the GPU. The total filter farm of the final system contains twelve units of farm
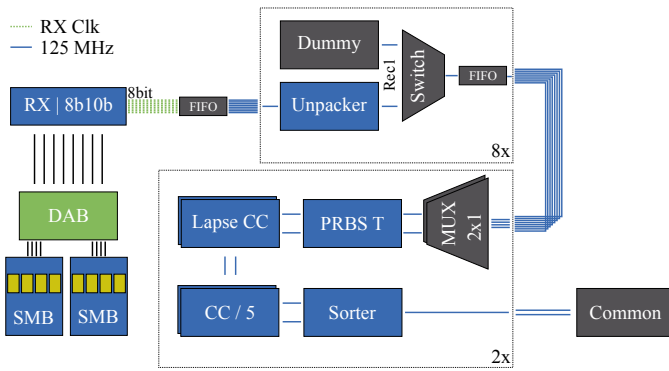
Fig. 3. Sketch of the datapath of the scintillating fibre detector. One scintillating fibre (SciFi) module board (SMB) holds four MuTRiG ASICs, and is connected to the FEB via a detector adapter board (DAB). The receiver (RX) on the FEB is byte-aligning the incoming serial LVDS links and 8b/10b [22] decoding them. Afterwards, the data is buffered in First In First Out (FIFO) queues and unpacked, multiplexed, and time-corrected. After this pre-processing, the hit data can be sorted using onboard memory, and transmitted to the common part of the FEB.

PCs, which are daisy-chained to process the expected data rate of 100 Gbit/s from the detector. To synchronise the different detectors, a dedicated clock and reset system was built. The system is centralised around the Genesys-2 FPGA board [16], which generates a 125 MHz clock and synchronised reset signals for run starts and stops. Additional electronics are used to create 144 copies of the clock and reset signals, which are connected to each part of the DAQ system. A comprehensive description of the final DAQ of the Mu3e detector is given in [17].

Since the search for $\mu^+ \rightarrow e^+ e^+ e^-$ is a three body decay at rest, the DAQ system needs to read out the full detector in order to be able to select physics events. This is in contrast to trigger-based DAQ systems, which only read out data if a certain detector response is over a specific threshold. Examples of trigger based DAQ systems are ATLAS [18] and Compact Muon Solenoid Experiment (CMS) [19], the two largest experiments at the Large Hadron Collider. However, the concept of full online reconstruction of the whole detector data is also being implemented at the LHCb [20] upgrade, and for parts of the ALICE detectors [21].

The main contributions of this work, in contrast to the previous description of the Mu3e DAQ system [17], are:

1) Detailed insights in the developed firmware parts of the Mu3e DAQ system, which are used to process the data from the different detectors to the filter farm.
2) Description of the time alignment firmware used on the switching board (SWB) to synchronise the different detector systems.
3) Overview of the Mu3e Cosmic Run 2022 and first timing studies of the different detectors using a vertical slice of the Mu3e DAQ system.

III. DATA FLOW IN THE MU3E DAQ

A. Front-end Board

As an example for the developed data processing firmware on the FEB, the data path of the scintillating fibre detector

is shown in Figure 3. Two scintillating fibre (SciFi) module boards (SMBs) holding eight MuTRiG ASICs are connected, via a specific detector adapter board (DAB), to a FEB. After receiving the data on the Arria V FPGA, the hit data is unpacked and grouped into a detector-specific record type (Rec1). The data lines from the different ASICs are merged into groups of two for further processing.

To be able to correlate the hits from the different detectors, time sorting is needed. Unfortunately, the two ASICs create the hit time with two different frequencies. The timestamp generation of the MuTRiG chip runs at 625 MHz while the MuPix chip generates the timestamp with 125 MHz. The faster clock for the MuTRiG is needed for the required time resolution of the two timing detectors. To encounter this problem the firmware needs to provide a common time base for both detectors. Therefore, the firmware takes the timestamp from the MuTRiG hits and divides them into a 125 MHz and a 625 MHz part. Afterwards the firmware sorts the hits according to the 125 MHz part. Since the MuPix chips are running already at 125 MHz only the sorting is needed.

In detail the algorithm for sorting the MuTRiG hits into 125 MHz is a bit more involved. In order to run at 625 MHz, the coarse counter timestamp is implemented as an optimised fifteen-stage linear feedback shift register, which runs through a deterministic order of $2^{15} - 1$ stages. The feedback loop of the shift register is implemented using an exclusive-or (XOR). These stages can appear to be pseudo random values (e.g. 1. stage is 0x1234, 2. stage 0x5678, 3. stage 0x4242), but can be translated back to binary values (e.g. 0x1234 = 0, 0x5678 = 1, 0x4242 = 2) using a dual-port lookup RAM on the FEB (in Figure 3, this is indicated with the PRBS T block). Dual-port RAMs are used to be able to process the two data streams at the same time.

Since the timestamp of the MuPix chip counts up to 2^{15} and the coarse counter of the MuTRiG chip counts up to $2^{15} - 1$, the counting of the MuTRiG chip needs to be corrected. As shown in 2 the whole system (FPGA-Boards and ASICs) is running in sync to a global 125 MHz clock. Both the 125 MHz clock on the FEB and the 625 MHz clock on the MuTRiG are synchronised to the global 125 MHz clock. Because of this synchronisation the lapse correction block (Lapse CC) can mimic the counting of the MuTRiG ASIC. It can then estimate the number of counter overflows and subtracted them from the coarse counter to correct it for counting up to 2^{15}.

After this correction of the counter value, the 625 MHz counter needs to be divided by five to have a part which follows the 125 MHz clock and a remainder part which holds the 625 MHz value. After all the corrections, the various hits can be sorted and transmitted to the SWB via up to two 6.25 Gbit/s optical links (for the pixel and tile detectors only one optical link is needed).

B. Switching Board

One task of the SWB is to time-align the different data streams. The time alignment of the hit data is needed to be able to perform the full track reconstruction on the GPU. In Figure 4, a sketch of this time alignment firmware of the SWB
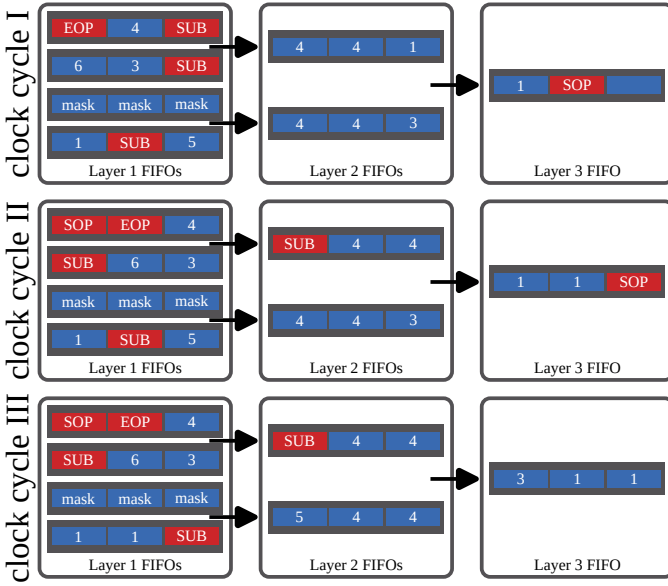
Fig. 4. Example of a four-to-one time alignment firmware on the SWB using pairwise comparison in a tree. The numbers on the left indicate three example steps (clock cycles) of the algorithm. The markers SOP and EOP indicate the start and the end of a data package sent from the FEBs. The marker SUB shows the so-called sub header which contains more timestamp information. The values between show the hit time of hits inside a sub header. The alignment firmware combines the various packages into one common stream while keeping the hits' time sorted.

is shown. Each data package from the different FEBs contains hit data sorted into 8 ns bins. A new data package is sent every 16 µs. Each hit contains space, time-over-threshold (ToT) of the comparator, and time information. The data package for each detector is designed to hold the whole hit information using 32 bit. To be able to keep the size of 32 bits, we only have four bits for the time information (in Figure 4 the time information of the 32 bit hit word is the white number of the blue rectangles). Therefore, we have to send 128 sub headers (SUB) which contain the upper seven bits of the global timestamp (in total one package as 2^{11} possible timestamps) when the four bits of the hit time overflow.

The time alignment firmware is using pairwise comparison of the hit timestamps in a tree-based architecture. In Figure 4, the start of each package is marked with SOP and the end of each package is marked with EOP. Each layer of the tree contains two input streams and one output stream. The streams are buffered in FIFO queues, the example in Figure 4 uses a queue size of three (in the final system the queue size is of the order of 2^{10}). In the first layer, the input stream is running with 125 MHz and the output of the FIFO is running at 250 MHz. At each layer of the tree, both streams are buffered until each has a SOP, then the hit with the lowest timestamp is sent to the next layer. If a stream holds a SUB, the hits from the other stream are forwarded until both streams hold the same SUB. Since the FEB always sends all 128 SUBs for each package - even in the case where there are no hits for these timestamps - the firmware only needs to compare the four bits of the two individual hits. These are reproduced over the full tree until the EOP marker is reached. Figure 4 shows, as an example, three steps (clock cycles) of a four-to-one tree.

The actual implementation is done with an eight-to-one tree following the same principle. After the synchronisation, the merged stream is sent via a 10 Gbit/s link to the first Terasic DE5 board of the filter farm.

As described above, the firmware combines eight input streams into one output stream. Since there is only one clock frequency change from 125 MHz to 250 MHz, the time alignment firmware has a bottleneck of effectively $4 - 1$. Using simulated data, the average rate for the whole pixel detector was estimated to be 56 Gbit/s, including 8 bit/10 bit encoding and 75 % protocol efficiency [17]. The total number of input links for the central detector SWB from the FEBs is 36, while there is a total of eight output links to the Terasic DE5 board. Therefore, the possible throughput of the time alignment firmware for the central detector SWB is of the order of 64 Gbit/s. This reduces the bottleneck of the time alignment firmware to around $2 - 1$ for the central pixel detector. However, this bottleneck is only hypothetical and only if all 36 input links of the central pixel SWB are 100% saturated which will never be the case in the actual detector. The expected rate of the whole pixel detector is 56 Gbit/s, including the two recurl stations and the central station. Since the central SWB is only connected to the central pixel detector, the time alignment firmware does not have to process the full 56 Gbit/s. To be able to use backpressure and to handle possible rate bursts, the FIFO queue size for the first layer of the time alignment firmware was set to hold one data package (32 kB). For the two recurl SWBs, two 10 Gbit/s links to the Terasic DE5 board is intended. The recurl pixel stations produce only a fraction of the data from the central station, leading to no significant bottleneck. For the other two detectors, the overall situation is relaxed as well. The fibre detector will produce an overall rate of 28 Gbit/s and one SWB is processing this data. For the connection from the fibre detector SWB to the Terasic DE5 board four 10 Gbit/s links are provided. The two stations of the tile detector producing an overall rate of 17 Gbit/s, while each recurl SWB has a 10 Gbit/s link for the tile detector to the Terasic DE5 board.

C. Filter Farm

The first Terasic DE5 board of the filter farm receives the whole 100 Gbit/s detector data from all four SWB. Because all data is synchronised and sorted to 8 ns, the board receives a stream of time-continuous data. This data is buffered on one of the two onboard memory interfaces. If the buffer gets full, the data is forwarded to the next Terasic DE5 board which is daisy-chained. Since the full farm contains twelve farm PCs, each of them needs to process only $1/12$ of the overall 100 Gbit/s rate. To perform the online track reconstruction, the data of the central pixel detector needs to be transferred to the GPU which is hosted in the same PC. The central pixel data is packed into 2 MB packages and transmitted via the PC RAM to the GPU using direct memory access (DMA) [23], [24]. A reference of the selected events are sent back to the Terasic DE5 board using Peripheral Component Interconnect Express (PCIe) registers and stored in a request FIFO to trigger the readout of the data buffer. By using two DDR4 RAMs, one
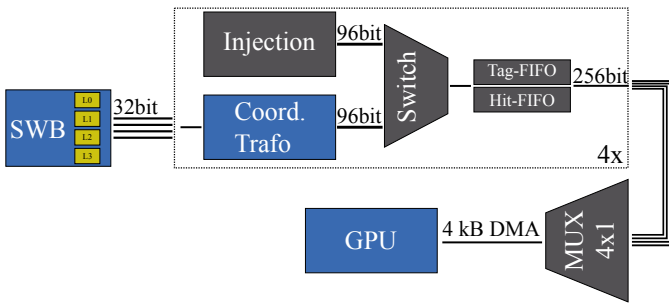
Fig. 5. The central switching board (SWB) is connected to the first farm PC, sending the pixel hits of the four layers (L0-L3). The incoming hit positions are converted to 32 bit floating-point values of the global x, y, z position. For debugging and eventual blinding proceedings, an injection entity is implemented. Using a Tag-FIFO the absolute number of hits and total number of 256 bit words for all hits inside a time-frame is stored. By multiplexing between the four layers, the required GPU events are contracted and sent via DMA to CPU RAM, where it is forwarded to the GPU.
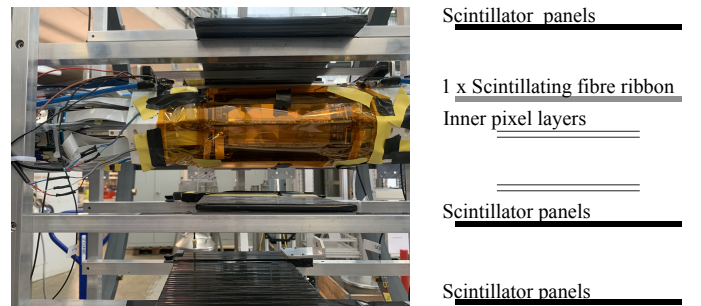

Fig. 6. **Left:** Picture of the Mu3e Cosmic Run 2022 detector prototypes. The inner vertex detector is enclosed with Kapton foil to distribute the helium flow. Directly over the pixel detector, a layer of scintillating fibre ribbons is mounted. For further noise and cosmic studies, three additional scintillating panels are mounted around the detectors. The coincidence of the panels are used as a cosmic ray trigger. The three panels are not perfectly aligned, which results in a non optimal coverage of the detector acceptance. **Right:** Sketch of the detector setup of the Mu3e Cosmic Run 2022.

can be used to buffered the data while the other one can be read out via a second DMA engine and transferred to the Maximum Integration Data Acquisition System (MIDAS) [25], [26] which is used as data acquisition software.

To be able to efficiently perform the tracking on the GPU, the data needs to be prepared on the FPGA. In Figure 5, the data flow of the central pixel detector on the Terasic DE5 board is shown. Each layer sends its individual hit position and timing information using a total of 32 bits. In the first step, the position is converted into global 32 bit floating point values of x, y, z using 96 bits in total. For this, a conversion from the local coordinate system on the chip to the global coordinate system of the detector needs to be performed. This is done by using the chip ID as address to lookup the global corner value (row=0, column=0) of the corresponding pixel chip in a RAM-based lookup-table. Knowing the corner value of each chip and the column (col) and row information the global position x, y, z in the whole pixel detector of the ith chip can be calculated using:

$$\vec{h_i} = \vec{s_i} + \text{col}_i \cdot \vec{c_i} + \text{row}_i \cdot \vec{r_i}, \qquad (1)$$

where $\vec{h_i}$ is the hit position vector (x, y, z), $\vec{s_i}$ is the global corner position, $\vec{c_i}$ and $\vec{r_i}$ are the column and row directions. Following that, the hit data is packaged into frames of 8 ns and packed into 256 bits to be processed by the DMA engine running at 256 bits times 250 MHz. For each GPU frame, the hits are sorted in time for each of the four layers. The hits of the four layers are stored in 0.5 MB sub-packages per layer inside the overall 2 MB package. At the end of each sub-package, references to the 8 ns borders of the hits are transmitted. With this packaging, the GPU is able to dynamically create time frames in multiple of 8 ns. Since one has the reference on the 8 ns level, it is possible to include hits from neighbouring frames ($\pm$ n x 8 ns) without copying them twice. These overlaps are needed to perform the tracking using all four layers [23], [24].
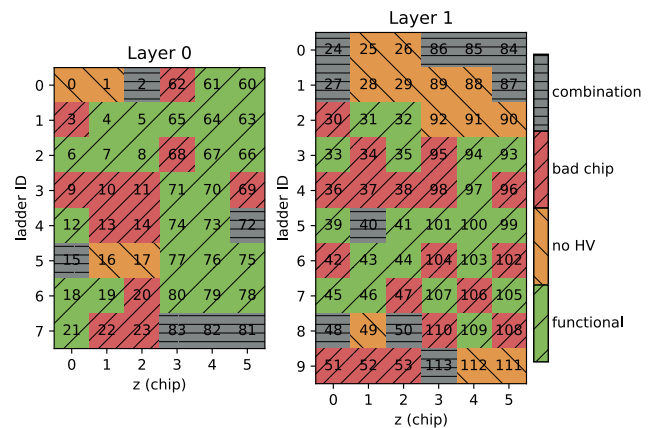

Fig. 7. Overview of the chips mounted on the two layers of the inner vertex detector. All chips without problems are marked in green (with a wide bottom left to top right hatching), while the chips in orange (with a wide top left to bottom right hatching) run into thermal runaway when applying high voltage. The chips marked in red (with a fine bottom left to top right hatching) had problems in terms of broken links, non-working LVDS connections or noisy pixels. The grey chips (with a fine left to right hatching) had some combination of the two problems above.

IV. THE MU3E COSMIC RUN 2022

During the integration of the different detector systems into the final DAQ system, two intensive system tests were performed. The first one, the Mu3e Integration Run 2021, was performed to integrate the inner vertex detector into the DAQ system and is explained in detail in [29]. The second test run, the Mu3e Cosmic Run 2022, was performed to finalise the integration of the inner vertex detector and to integrate the scintillating fibre detector using cosmic rays. In contrast to the Integration Run 2021, the detectors were not operated inside a magnetic field. In Figure 6, the Cosmic Run setup is shown. One intention of the run was to operate the detectors under conditions as close as possible to the final experiment and study their noise behaviour. This includes cooling of the MuPix detector with gaseous helium.
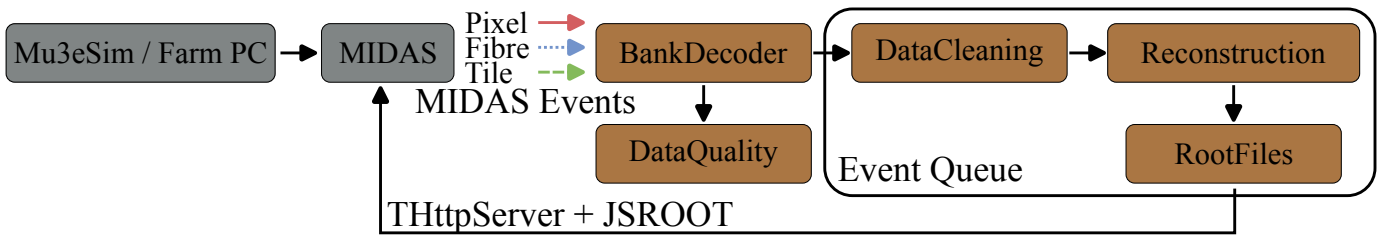
Fig. 8. Sketch of the online analyzer framework used at the Mu3e Cosmic Run 2022. The design of the system is completely integrated into MIDAS and receives all the different detector events. In the first step, the data banks inside the events are decoded and sent to a queue where different cleaning and clustering steps are performed for each detector type. Using the Mu3e Reconstruction software, online tracking can be performed. The final tracks are written to a ROOT [27] file. At the same time, the raw data is buffered using a second queue to perform data quality plots (DataQuality) of the system. Using the THttpServer class of ROOT combined with JavaScript ROOT, provides interactive ROOT-like graphics in the web browsers. For debugging and performance tests, the whole setup can be fed with simulation data from the Mu3e software [28].
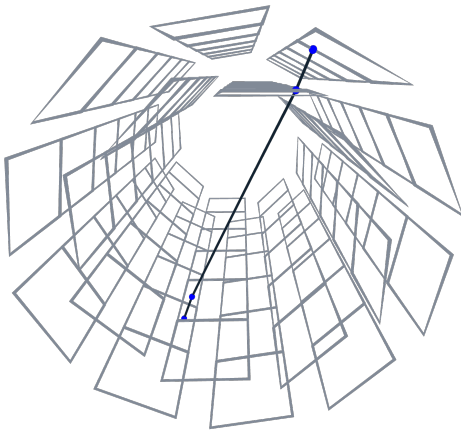

Fig. 9. Sketch of the Mu3e Event Display (PhD thesis in progress: B. Gayther, Preparations for Phase I of the Mu3e experiment). The display is showing a track which was reconstructed using the Mu3e Online Analyzer. The track was fitted to four hits in a certain time frame using only the data from the inner vertex detector.

As for the final DAQ system, the Mu3e Cosmic Run 2022 contained ten FEBs for reading out the inner vertex detector. On the other side, only two FEBs for the fibre detector were used compared to the final twelve. Furthermore, there were no tiles and none of the other pixel detectors present. All twelve FEBs are read out with only one SWB, which contains a third of the final links for the central vertex detector and a small fraction for the fibre detector. The farm contained only two Terasic DE5 boards and no GPUs, in contrast to the twelve Terasic DE5 boards used in the final system. The dashed blocks in Figure 2 show the components used in the Mu3e Cosmic Run 2022. Furthermore, three layers of scintillator panels are mounted around the scintillating fibre and inner vertex detectors. Using Nuclear Instrumentation Module (NIM) logic, the three panels are used to create a cosmic ray coincidence trigger. This trigger signal was digitised via an additional FEB and sent to the SWB. The geometry of the scintillator panels was not optimal, and the panels did not cover the whole detector acceptance.

By tuning the vertex detector to a minimal noise level, it was possible to read out the full detector data. Due to a broken readout board, only one side of the fibre detector was operating. In addition, not all pixel chips were able to

run under optimal conditions. In Figure 7, all chips without problems are marked in green, while the chips in orange run into thermal runaway when applying high voltage. The chips marked in red had problems in terms of broken links, non-working LVDS connections or noisy pixels. The grey chips had some combination of the two problems above. Overall, only 10 V could be applied to the chips since there are unforeseen problems in the prototype of the MuPix chip which are resolved in a newer version. Having all these problems even the best-working chips were not operated at design conditions, leading to a worse timing resolution. A detailed description of the inner vertex detector used in the test runs can be found in [30].

A. Online Analyzer

The Mus3 Online Analyzer software was implemented to analyse the data quality and perform tracking. This software component is based on the MIDAS-compatible manalyzer framework [31]. A parallel readout pipeline was used to clean the data and perform the needed quality checks. Figure 8 demonstrates the various steps of the analyzer. In the first part, the system connects to MIDAS and checks if detector events are present. The different detector events for the fibre and the vertex detector are decoded and sent to a queue for further processing. Even with the threshold tuning and masking pixels, a few chips still had noisy pixels left which could not be turned off in hardware. This required to filter them out of the raw data in software. Since the fibre data path had no hardware sorter running, the data were sorted in time to be able to correlate with the pixel data. The additional scintillator panels were considered as pixel hits and were filtered out at the pixel readout chain. After the cleaning, a full track reconstruction was performed using an adapted version of the Mu3e triplet reconstruction software [28] to reconstruct cosmic tracks in the pixel detector. The reconstructed hits were forwarded to an event display to have a live view on specific events (see Figure 9). For general time correlations, the scintillator panels were used as a trigger to correlate the fibre and pixel hits in a time-window around the trigger. Simple quality plots such as hit rate, hitmaps, timestamp plots, etc., were created for the different detectors.
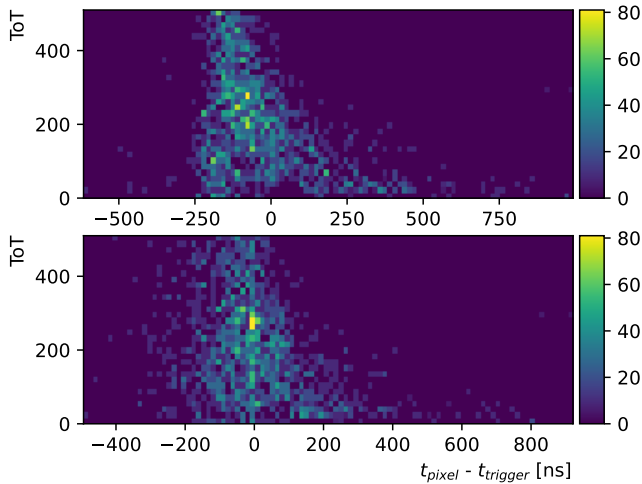
Fig. 10. **Top:** ToT versus time difference of the inner vertex detector and the scintillating panels. **Bottom:** ToT versus time difference of the inner vertex detector and the scintillating panels with time walk corrected pixel hit time.
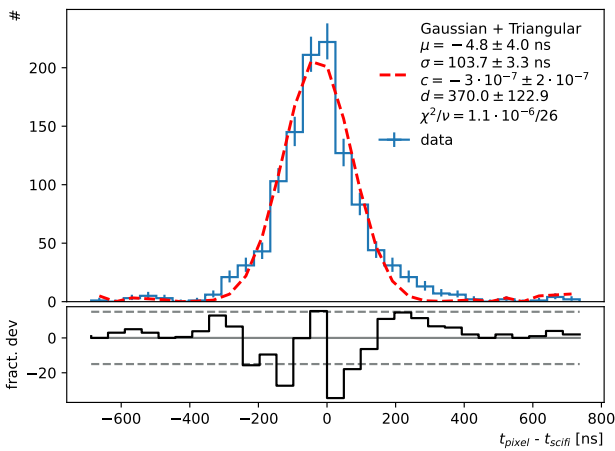

Fig. 11. The time correlation between the inner vertex detector and the scintillating fibre detector requires a coincidence of the scintillator panels.

B. Results Of The Cosmic Run 2022

To simplify analyses of the time correlation of the inner vertex detector and the scintillating fibre detector, only the hits which met a coincidence with the three scintillator panels are used. The top plot of Figure 10 shows the ToT versus the time difference of the inner vertex detector and the scintillator panels. The clear asymmetric shape of the distribution is caused by the time walk effect. Since the comparator of each pixel cell has an absolute threshold on the rising edge of the analog signal to determine the hit time, signals with higher amplitude have an earlier time, and signals with lower amplitude have a later time. The difference of the two times is referred to as time walk. To overcome this problem, the MuPix chips can be tuned to use a second timestamp, sampled on the falling edge of the signal, to calculate the ToT and correct the hit time. During the Cosmic Run this feature was not turned

on, and an offline correction, using the scintillator panels as a reference, needs to be performed. To correct this, the mean of each ToT value was calculated and subtracted from the sampled hit time. The bottom plot of Figure 10 shows the corrected time difference. During the test run, the setting for the readout speed of the second timestamp was not tuned for each individual pixel chip. This leads to an imperfect sampling of the second timestamp, which can be seen in the top of Figure 10 for hits on the lower left part of the distribution ($t_{pixel} - t_{trigger} < 0$ and ToT< 200). There, the ToT was not calculated correctly and the variance of the correction for these ToT values (see bottom of Figure 10) is higher. In contrast to the inner vertex detector, the scintillating fibre detector was not configured to send the second timestamp. Therefore, a time walk correction was not possible for this detector.

In Figure 11, the difference between the hit time of the inner vertex detector after correcting the time walk and the scintillating fibre detector is shown. Since all hits were correlated with each other, an unintended triangular-shaped background was created. Therefore, the used Gaussian fit function was extended by $f_{bg}(x) = c(d - |x|)$. The performed fit showed a total time resolution of 103.7(33) ns.

The systematic errors are dominated by the pure timing resolution of the vertex detector caused by the untuned chips, which is in the order of 100 ns. The sampling of the scintillator panels was done with 8 ns on the FEB, which is the main source of systematic uncertainty for this detector. The scintillating timing detector should have a smaller uncertainty on its timing resolution [11] than the systematics of the scintillator panels. Nevertheless, despite all these limitations on the detector side, the presented results prove that the DAQ system is able to process different detector data, perform online time synchronisation to provide the required hit data to study time resolutions between the different detectors and track reconstruction.

V. SUMMARY AND OUTLOOK

The paper presented detailed insights in the developed firmware parts of the Mu3e DAQ system, which are used to process the data from the different detectors to the filter farm. Especially, the time alignment firmware used on the SWB to synchronise the different detector systems was shown. Furthermore, a detailed description of the Mu3e Cosmic Run 2022 and first timing and tracking studies of the different detectors using a vertical slice of the Mu3e DAQ system are given.

Showing first physical plausible results for the detected cosmic rays and being able to have a consistent readout chain are first steps in commissioning the final DAQ system of the Mu3e experiment.

During the end of 2022, the commissioning of the filter farm will start. The final detector will be constructed in the beginning of 2023 and first data taking is planned for the end of 2023.

REFERENCES

[1] K. Arndt *et al.*, "Technical design of the phase I Mu3e experiment," *Nucl. Instr. Meth. A*, vol. 1014, p. 165679, 2021.

[2] U. Bellgardt *et al.*, "Search for the Decay $\mu^+ \rightarrow e^+e^+e^-$," *Nucl. Phys.*, vol. B299, p. 1, 1988.

[3] M. Aiba *et al.*, "Science Case for the new High-Intensity Muon Beams HIMB at PSI," 2021. [Online]. Available: https://arxiv.org/abs/2111.05788 Accessed: 2022-12-12.

[4] H. Augustin *et al.*, "The MuPix high voltage monolithic active pixel sensor for the Mu3e experiment," *JINST*, vol. 10, no. 03, p. C03044, 2015.

[5] ——, "MuPix7 - A fast monolithic HV-CMOS pixel chip for Mu3e," *JINST*, vol. 11, no. 11, p. C11029, 2016.

[6] ——, "The MuPix System-on-Chip for the Mu3e Experiment," *Nucl. Instrum. Meth.*, vol. A845, pp. 194–198, 2017.

[7] ——, "Efficiency and timing performance of the MuPix7 high-voltage monolithic active pixel sensor," *Nucl. Instr. Meth. A*, vol. 902, p. 158, 2018.

[8] ——, "MuPix8 — Large area monolithic HVCMOS pixel detector for the Mu3e experiment," *Nucl. Instrum. Meth. A*, vol. 936, pp. 681–683, 2019.

[9] ——, "Performance of the large scale HV-CMOS pixel sensor MuPix8," *JINST*, vol. 14, no. 10, p. C10011, 2019.

[10] ——, "MuPix10: First Results from the Final Design," *Proceedings of VERTEX2020*, 12 2020.

[11] S. Bravar *et al.*, "Scintillating fibre detector for the Mu3e experiment," *JINST*, vol. 12, no. 07, p. C07011, 2017.

[12] H. Klingenmeyer *et al.*, "Measurements with the technical prototype for the Mu3e tile detector," *Nucl. Instr. Meth.*, vol. A958, p. 162852, 2019.

[13] H. Chen *et al.*, "MuTRiG: a mixed signal Silicon Photomultiplier readout ASIC with high timing resolution and gigabit data link," *JINST*, vol. 12, no. 01, p. C01043, 2017.

[14] J. P. Cachemiche, P. Y. Duval, F. Hachon, R. Le Gac, and F. Réthoré, "The PCIe-based readout system for the LHCb experiment," *JINST*, vol. 11, no. 02, p. P02013, 2016.

[15] "DE5a-NET FPGA Development Kit User Manual," Terasic Inc, Tech. Rep., 2019.

[16] Digilent, "The Digilent Genesys 2 Kintex-7 FPGA development board." [Online]. Available: https://reference.digilentinc.com/reference/programmable-logic/genesys-2/ Accessed: 2022-12-12.

[17] H. Augustin *et al.*, "The mu3e data acquisition," *IEEE Transactions on Nuclear Science*, vol. 68, pp. 1833–1840, 10 2020.

[18] W. P. Vazquez *et al.*, "The ATLAS Data Acquisition System: from Run 1 to Run 2," *Nuclear and Particle Physics Proceedings*, pp. 273–275, 2016.

[19] CMS collaboration, "CMS Technical Design Report for the Level-1 Trigger Upgrade," 2013. [Online]. Available: http://cds.cern.ch/record/1556311 Accessed: 2022-12-12.

[20] R. Aaij *et al.*, "Allen: A High-Level Trigger on GPUs for LHCb," *Computing and Software for Big Scienc*, vol. 4, 2020.

[21] ALICE collaboration, "Real-time data processing in the ALICE High Level Trigger at the LHC," *Computer Physics Communications*, vol. 242, pp. 25–48, 2019.

[22] A. X. Widmer and P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," *IBM Journal of Research and Development*, vol. 27, p. 440, 1983.

[23] D. vom Bruch, "Pixel Sensor Evaluation and Online Event Selection for the Mu3e Experiment," Ph.D. dissertation, Heidelberg University, 2017.

[24] V. Henkys, B. Schmidt, and N. Berger, "Online Event Selection for Mu3e using GPUs," 2022. [Online]. Available: https://arxiv.org/abs/2206.11535 Accessed: 2022-12-12.

[25] S. Ritt, P. Amaudruz, and K. Olchanski, "The MIDAS data acquisition system," *Proc. IEEE 10th Real Time Conf.*, pp. 309–312, 1997.

[26] ——, "Maximum Integration Data Acquisition System," 2001. [Online]. Available: https://midas.triumf.ca/ Accessed: 2022-12-12.

[27] R. Brun and F. Rademakers, "ROOT - An Object Oriented Data Analysis Framework," *Nucl. Inst. & Meth. in Phys. Res. A*, vol. 389, pp. 81–86, 1997.

[28] N. Berger *et al.*, "A New Three-Dimensional Track Fit with Multiple Scattering," *Nucl. Instr. Meth A.*, vol. 844, p. 135–140, 2017.

[29] M. Köppel, "Mu3e Integration Run 2021," *Proc. PoS NuFact2021*, p. 233, 2022.

[30] T. Rudzki, "The Mu3e vertex detector - construction, cooling, and first prototype operation," Ph.D. dissertation, Heidelberg University, 2022.

[31] K. Olchanski, "MIDAS analyzer." [Online]. Available: https://bitbucket.org/tmidas/manalyzer/ Accessed: 2022-12-12.