Published by IOP Publishing for Sissa Medialab

Received: *July 28, 2025*
Revised: *October 22, 2025*
Accepted: *November 24, 2025*
Published: *December 17, 2025*

Readout electronics for low occupancy High-Pressure Gas TPCs

N. Khan,[a] **Y. Hua,**[a] **I. Xiotidis**[a,*] **T. Alves,**[a] **E. Atkin,**[a] **G. Barker,**[b] **D. Barrow,**[c] **A. Booth,**[d] **J. Borg,**[a] **A. Bross,**[e] **M. F. Cicala,**[f] **L. Cremonesi,**[d] **A. Deisting,**[g] **K. Duffy,**[c] **R. Gran,**[h] **P. Green,**[c] **A. Habig,**[h] **M. Judah,**[l] **T. Junk,**[e] **A. Kaboth,**[i] **A. Klustová,**[a] **H. LeMoine,**[h] **A. D. Marino,**[j] **F. Martínez López,**[d,k] **T. Mohayai,**[k] **D. Naples,**[l] **R. Nichol,**[f] **D. Parker,**[a] **M. Pfaff,**[a] **J. Raaf,**[e] **P. Rubinov,**[e] **P. Singh,**[d] **A. Srivastava,**[g] **A. Waldron,**[a,d] **L. Warsame,**[a,i] **M. Wascko,**[a,c] **A. Wilkinson,**[b,f] **A. Ritchie-Yates**[i] **and P. Dunne**[a]

[a] *Imperial College of London, London, U.K.*
[b] *University of Warwick, Warwick, U.K.*
[c] *University of Oxford, Oxford, U.K.*
[d] *Queen Mary University of London, London, U.K.*
[e] *Fermi National Accelerator Laboratory, Batavia, U.S.A.*
[f] *University College of London, London, U.K.*
[g] *Johannes Gutenberg-Universität Mainz, Mainz, Germany*
[h] *University of Minnesota Duluth, Duluth, U.S.A.*
[i] *Royal Holloway University of London, London, U.K.*
[j] *University of Colorado Boulder, Boulder, U.S.A.*
[k] *Indiana University, Bloomington, U.S.A.*
[l] *University of Pittsburgh, Pittsburgh, U.S.A.*

E-mail: i.xiotidis@imperial.ac.uk

Abstract: High-Pressure Gas Time Projection Chambers (HPgTPCs) have benefits such as low energy thresholds, magnetisability, and 4π acceptance, making them ideal for neutrino experiments such as DUNE. We present the design of an FPGA-based solution optimised for Gaseous Argon Near Detector (ND-GAr), which is part of the Phase-II more capable near detector for DUNE. These electronics reduce the cost significantly compared to using collider readout electronics, which are typically designed for much higher occupancy and therefore, for example, need much larger numbers of FPGAs and power per channel. We demonstrate the performance of our electronics with the Teststand for an Overpressurised Argon Detector (TOAD) at Fermilab in the US at a range of pressures and gas mixtures up to 4.5 barA, reading out ~10 000 channels from a Multi-Wire Proportional Chamber (MWPC). The operation took place between April and July of 2024. We measure the noise characteristics of the system to be sufficiently low, and we identify sources of noise that can be

*Corresponding author.

https://doi.org/10.1088/1748-0221/20/12/P12023

2025 JINST 20 P12023

further mitigated in the next iteration. We also note that the cooling scheme used in the test requires improvement before full-scale deployment. Despite these necessary improvements, we show that the system can fulfil the needs of a HPgTPC for a fraction of the price of collider readout electronics.

Keywords: Data acquisition concepts; Front-end electronics for detector readout; Gaseous detectors; Neutrino detectors

ArXiv ePrint: 2507.17425

2025 JINST 20 P12023

Contents

1 Introduction

Gas detectors offer greatly reduced energy thresholds, magnetisability, and 4π acceptance, which are not currently achieved with other detector technologies [1]. However, gases have much lower density than solid or liquid detectors, reducing the interaction rate for particles passing through. The advent of MW-power neutrino beams means that detectors with gas targets will, for the first time, be practical for long-baseline neutrino experiments. Despite these higher beam powers, a gas Time Projection Chamber (TPC) detector in these beams would still need to be pressurised to accumulate a sufficient number of interactions.

Harnessing the capabilities of such a High-Pressure Gas Time Projection Chamber (HPgTPC) for the Deep Underground Neutrino Experiment (DUNE) would provide important insights into neutrino-argon interactions and enable key measurements of the acceptance corrections of other detectors. As such, a gaseous argon TPC surrounded by an Electromagnetic Calorimeter (ECal) has been recognised as crucial to the experiment's physics program by the recent US P5 process [2], and the DUNE collaboration's Phase II white paper [3].

The detector's position resolution is a critical factor in achieving the required momentum resolution and detection thresholds for complex low-energy interactions with high multiplicity. The resolution obtainable is closely related to the number of readout channels that the detector can accommodate and the sampling frequency of those channels. Whilst the numbers are not final, instrumenting a 5 m diameter TPC sufficiently to achieve the desired low energy thresholds for the DUNE would require $O(100\,\mathrm{k}) - O(1\,\mathrm{M})$ channels with sampling frequency $O(10\,\mathrm{MHz})$ [1]. The development of readout electronics in the particle physics community has previously focused on the design and operation of high-throughput radiation-hard electronics required at the Large Hadron Collider (LHC)-based experiments [4]. Readout electronics from sPHENIX [5] or the TPC of A Large Ion Collider

2025 JINST 20 P12023

Experiment (ALICE) [6] require a high-specification Field Programmable Gate Array (FPGA) to readout groups of ~10 front-end cards. Naively scaling such designs for neutrino experiments would be prohibitively expensive. This is due to the hundreds of thousands of channels required to meet DUNE's resolution needs. Hence, this approach would result in a total cost larger than possible, given budgetary constraints for the DUNE near detector. In fact, the electronics were previously listed as the most expensive component for the Gaseous Argon Near Detector (ND-GAr) for DUNE.

Since neutrino experiments have, by their nature, much lower occupancy and hence lower throughput than LHC experiments, a re-optimisation of the readout system enables significant cost savings whilst retaining physics performance.

In the following sections we present the design of a scalable and highly re-configurable FPGA-based solution optimised to be affordable for the ND-GAr DUNE Phase-II detector [3]. This system delivers significant cost-saving compared to the systems described above while showing no performance degradation in an HPgTPC environment. A prototype has been produced and was then sent for testing to Fermilab as part of the Teststand for an Overpressurised Argon Detector (TOAD). Whilst DUNE has not decided on an amplification stage for its HPgTPC, a Gas Electron Multiplier (GEM)-based solution is the current frontrunner. However, TOAD used a Multi-Wire Proportional Chamber (MWPC) from ALICE, due to its availability at the time of the test.

2 TOAD

The TOAD detector design and operation largely follow that of the U.K. high-pressure platform, described in [7]. The detector was sent to Fermilab, where it was commissioned, and the newly developed readout electronics were installed. These readout electronics interface with an ALICE Outer Read-Out Chamber (OROC), meaning TOAD was a full slice test of a HPgTPC.

2.1 ALICE multi-wire proportional chambers

Given that the ALICE TPC [6, 8] required an upgrade [9], the Read-Out Chambers (ROCs) can be repurposed. The OROC used in TOAD was part of the ALICE 'test' ROCs, rather than a production ROC. The specifications of the OROC used are detailed at length in [7].

The ROCs employ a typical wire plane schematic: an anode-wire grid above the pad plane, a cathode-wire grid and a gating grid to separate the drift volume from the readout. The pad plane and the three wire planes are all spaced 3 mm apart. All wires run in the azimuthal direction (approximately vertically in the orientation we have the chamber) and have different separations in the various planes. The length of the shorter and longer parallel sides of the OROC are 45.8 cm and 86.0 cm, respectively. The perpendicular length of the OROC is 114 cm. Voltages are applied to the anode and gating grid wires, while the cathode wires are held at ground. The pad plane consists of 9984 pads in total, of two sizes, which are summarised in table 1. Here, "inner" refers to the rows closer to the shorter parallel side, whereas "outer" refers to rows closer to the longer side of the OROC. The difference in pad dimensions is due to the trapezoidal shape of the ROC. The readout electronics must interface with this ROC. Therefore, groups of 21 or 22 pads are read out via a socket on the back of the OROC. With the sockets grouped in threes, our Front-End Cards (FECs) connect directly to these to read out a total of 64 channels each.

2025 JINST 20 P12023

Table 1. Dimensionality of the pad plane on the OROC.

	Pad Dimensions (mm^2)	Pad Rows	Number of Pads
Outer Pads	6×15	32	4032
Inner Pads	6×10	64	5952
Total	-	96	9984

2.2 Pressure vessel

TOAD utilises the same pressure vessel described in [7, 10], which allows the testing of the readout electronics with an ALICE OROC at pressures up to 4 barG (~5 barA). The vessel has an inner diameter of 140 cm and an outer diameter of 142 cm, made from 304 L Stainless Steel. The vessel has domed ends, giving a total volume of 1.472 m^3. The vessel can be opened by removing one of the domed faces. An image of the vessel in the test environment is shown in figure 1. Further details of the vessel properties can be found in [7, 10]. The way in which the feedthroughs were utilised for TOAD differs from the original setup detailed in [7]. In addition, the vacuum system was replaced. There are various KF25 and KF40 flanges on the vessel, which are used for the High Voltage (HV), gas, vacuum and readout-electronics systems feedthroughs. The gas and evacuation system is connected to the vessel so that the vessel can be evacuated to $O\left(10^{-5}\,\text{barA}\right)$ to ensure the purity of gas in the detector. The pump system was upgraded to an Edwards Turbo Cart, consisting of a backing pump, a turbo pump and a controller. The gas system allowed the vessel to be filled with up to 5 barA of gas mixtures, with the gas mix controlled by filling to a given partial pressure of gas from up to four different gas cylinders connected to the system. For the run described in this paper, three input gases were mixed: Ar-CH_4 with a molar ratio of 92:8, pure Ar, and pure N_2.

The TPC is the same as that described in [7]. It utilises a HV drift cathode, a field cage, a terminating plate, and the OROC. There are three HV feedthroughs on the vessel, which are used to apply voltages to the drift cathode, anode and gating grid. Other feedthroughs are also used to ensure the equipment inside the vessel is all safely grounded. The cathode and OROC anode were previously tested with voltages up to 16 kV and 3 kV respectively, with the maximum achievable voltage varying with the gas mixture used [7]. The results shown in this paper use various configurations of HV, gas mixtures and pressures. These are summarised in table 2. The measurements outlined in this paper are noise characterisation studies, and so the particle tracking ability of the TPC is not employed.

Table 2. Different operational configurations used for data shown in the Results section.

Mode	Gas Mixture	Pressure (barA)	Anode (V)	Gating Grid (V)	Cathode (V)
1	Air	1.0	0	0	0
2	Ar:CH_4, 96:4	4.5	0	0	0
3	Ar:CH_4, 98:2	3.0	1500	−100	0
4	Ar:CH_4, 98:2	3.0	1500	−100	−16 000

In addition to the connections described above, the readout electronics were powered and streamed data through KF40 feedthroughs on the vessel door. One feedthrough was used to provide power, and three RJ45 feedthroughs are used to control, configure, monitor and stream data to the readout electronics. These were connected directly to a Commercial Off the Shelf (COTS) network switch,

2025 JINST 20 P12023

Figure 1. The TOAD pressure vessel in the test environment at Fermilab. Around the vessel, the hydraulic pressure clamps can be seen. In the centre, there is access to five flanges, which are unused in TOAD but previously used for mounting an optical readout system. The KF40 feedthroughs that are used for data readout and Control Configuration and Monitoring (CCM) via the blue cables are labelled. Various other KF40 and KF25 flanges on the vessel provide access for the different subsystems, but cannot be seen in this image. The gas and evacuation system is to the right of the vessel.

which was, in turn, connected to two DELL PowerEdge servers: one running CCM software, and the other Data Acquisition (DAQ) software — both of which are outlined in detail in section 3.3.

3 Readout setup

A key challenge for future neutrino experiments is the harsh conditions in which the electronics must operate. In contrast to collider experiments, those conditions are not due to extensive radiation but instead factors such as cryogenic temperatures present in Liquid Argon (LAr)-based detectors and high pressure, such as in an HPgTPC based detector. In the case of the high-pressure detectors, an important consideration is the number of feedthroughs required to power and communicate with the electronics, so there is a need to aggregate data inside the vessel. Whilst power dissipation and thermal management present additional challenges, they can be addressed using water cooling through a small number of vessel penetrations. Hence, a readout system will need to optimise the number of cables

2025 JINST 20 P12023

required to stream data from the large number of analogue readout channels. In addition, based on the physics requirements for the DUNE near detector (ND) complex outlined in the Conceptual Design Report [1], the proposed readout system should not constrain or limit the advantages provided by the installation of a magnetised high-pressure gaseous TPC in the complex. Such requirements foresee enhanced tracking capabilities in comparison to the LAr-based detector counterpart for more accurate event classification. A summary of the requirements taken into account can be seen in table 3 with the equivalent design choice and the technology implementation. Obviously, choices like the number of links aggregated into a single FPGA are driven based on the capabilities of the package as well as the cost, making it an optimisation process based on physics versus cost.

Table 3. Physics requirements and expected timing and throughput values for ND-GAr as described in [1] with the corresponding choices in the design and solutions applied to those decisions. The requirement on the track length resolution arises from the ALICE TPC capabilities, which should suffice for ND-GAr but will be refined in the future.

Requirement	**Design choice**	**Implementation**
Detection threshold (O(MeV) protons)	High-Pressure TPC	Reduced feedthroughs
Track separation (1 – 5 cm)	ADC sampling	20 MHz
	Num. of detector channels	O(700k)
Low event rates (1 Hz)	Gigabit readout	Kintex Ultrascale FPGA

The proposed solution outlined in this paper uses a fast-sampling FEC with zero-suppression capability and a custom FPGA-based Power, Aggregation and Timing (PAT) card, both residing in-situ within the high-pressure volume. Data are then streamed over a standard Ethernet-based network to the DUNE DAQ software operating on standard servers, avoiding extensive requirements within the pressurised environment. The block diagram of the proposed readout system is shown in figure 2.

Compared to a collider experiment, the channels in neutrino experiments have very low occupancy (in the similarly sized ALICE, collisions occur at ~ 35 kHz and each generate thousands of tracks, whereas in DUNE, beam spills happen at ~1 Hz and generate only tens of tracks per spill). This difference means far less processing is required as there are far fewer above-threshold hits. Therefore, the system can have fewer FPGAs per channel (1 per 3584 channels with the particular FECs in this design) and use much lower specification FPGAs as we stream out all hits to a software-based DAQ. Comparisons with the ALICE and sPHENIX readout electronics are shown in table 4.

Table 4. Comparison in terms of event rates, sampling rates and channels per readout card among the sPHENIX readout system, the ALICE calorimeter readout and the ND-GAr readout.

System	**Event Rate**	**Sampling Rate**	**Channels/board**
sPHENIX	15 kHz	60 MHz	192
ALICE	35 kHz	10 MHz	Vary per detector region
ND-GAr	1 Hz	20 MHz	3584

Furthermore, the lower data rates from the FEC allow the native transceivers on the ADC Application Specific Integrated Circuits (ASICs) to be used to communicate with the PAT, so we also

2025 JINST 20 P12023

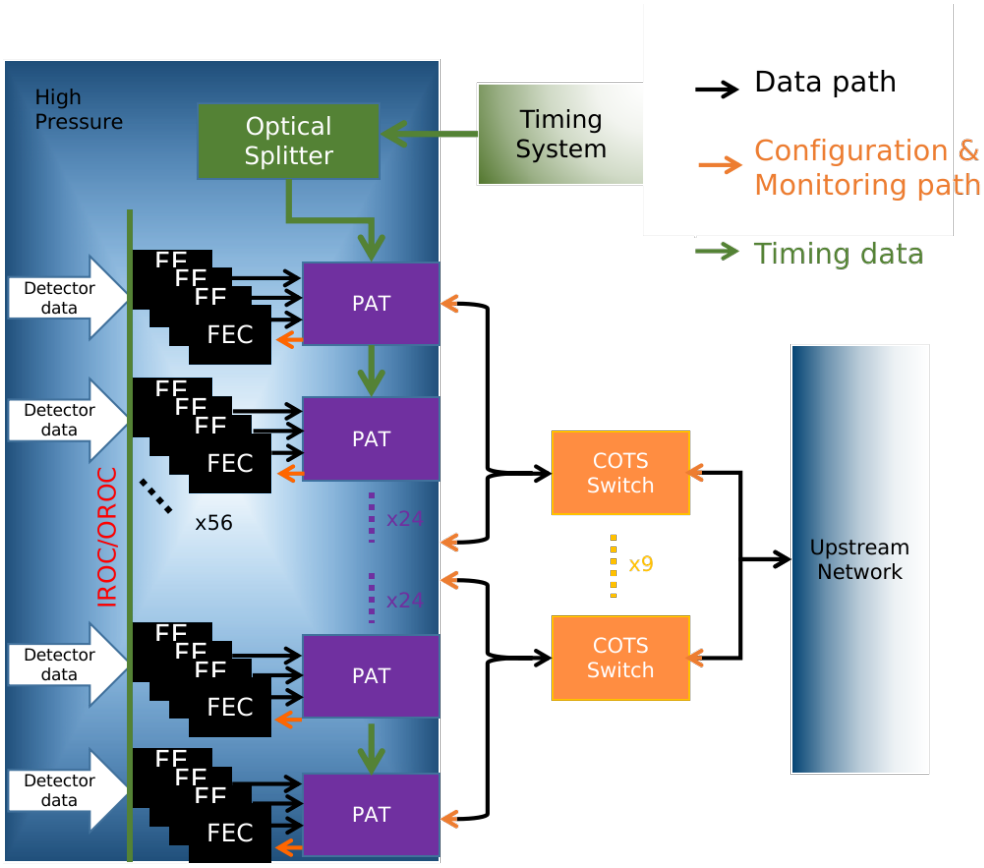

Figure 2. ND-GAr readout scheme with direct power provided to all DAQ cards and explicit timing signal distribution. The vertical, dark green line represents the ALICE Inner/Outer Readout Chamber. The orange boxes represent COTS network switches.

do not need the additional transceiver modules that are required for higher occupancy detectors [9]. These differences result in a much lower cost to instrument the hundreds of thousands of channels necessary for a neutrino HPgTPC.

3.1 Hardware layout

As described, our system requires two types of custom hardware boards. The first type are the FEC boards, which host the digitiser ASICs and connect to the detector itself. The second type are the PAT boards, which utilise an FPGA device for interfacing and aggregating the data for a multitude of FECs as well as providing them with power and distributing clock signals. Both the FECs and the PATs are custom-designed, with the FECs designed for the TOAD test beam detector and its ALICE OROC. The flexibility of the FPGA-based design of the PAT cards means that the PAT can be reused even if the FECs change, so, whilst designed for TOAD, this board could be used with minimal modifications for the final ND-GAr detector or other detectors with similar data volume and occupancy. In designing the hardware, we avoided components with known air gaps, e.g. lidded FPGA packages. However, as there is no standard for electronics operation in the 10 bar regime, a goal of TOAD is to test the operation of components at high pressures.

2025 JINST 20 P12023

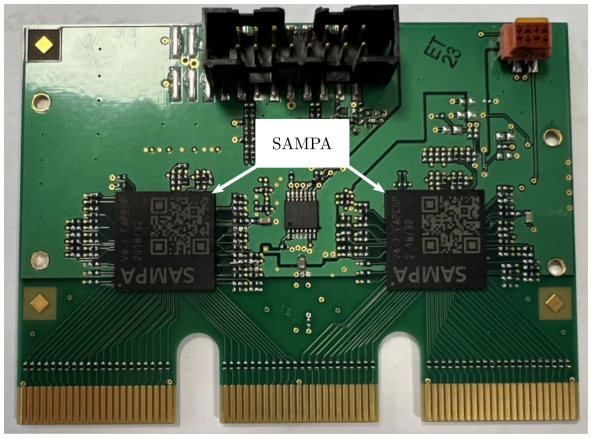

Figure 3. FEC board hosting two daisy-chained SAMPA chips designed for the TOAD OROC based test beam readout.

3.1.1 Front-end card

GEMs are currently the leading candidate for the charge readout for ND-GAr. The SAMPA ASIC used for this readout system was developed by the ALICE collaboration for the GEM upgrade of their TPC [9]. This upgraded detector has similar characteristics to those needed for ND-GAr, so SAMPA is expected to meet the requirements of ND-GAr. The parameters of the SAMPA measured by ALICE can be seen in table 5 [9, 11]. The electronics system described here is, therefore, designed around deploying the SAMPA ASIC [11] used by ALICE, in a high-pressure environment, and at a much lower cost than in the full ALICE system.

Table 5. SAMPA reference parameters measured for the ALICE TPC upgrade with GEMs [9, 11].

Specification	GEM
Supply Voltage	1.25 V
Detector Capacitance	40-80 pF
Peaking Time	300 ns
Shaping Order	4th
Equivalent Noise	<1600 e Cd = 80 pF
Sampling Rate	10 MHz, up to 20 MHz
Total Power (Idle)	< 15 mW/Ch

In addition, using an ASIC that has been proven to work at scale for another particle physics detector minimises risks and allows us to benefit from the extensive studies of their performance and limitations that have already been performed.

Another important reason for the choice of the SAMPA chips is their ability to sample at 20 MHz with fast ADCs. Gas detectors have about an order of magnitude faster electron drift velocities than liquid; hence, this high sample frequency is necessary.

For TOAD, the ALICE MWPC has a group of three connectors per slot, which correspond to a total of 64 channels. Given that the SAMPA ASIC has only 32 channels, the FEC designed for TOAD, which can be seen in figure 3, has three connectors and hosts two SAMPAs to capture all 64 channels.

2025 JINST 20 P12023

The final ND-GAr amplification stages may have a different number of channels per connector; however, the more SAMPAs per FEC, the lower the cost of the system, so designing for two SAMPAs represents a reasonable worst-case cost. To minimise the required readout channels connecting to the PAT board, the two SAMPAs on each FEC are daisy-chained, as the low occupancy of neutrino experiments does not require the full bandwidth of the SAMPA readout link. However, as this is a prototype detector, the layout of the FEC has been designed to allow each SAMPA to be read out individually with minimal changes to the fitted components and no re-design of the Printed Circuit Board (PCB).

For TOAD, the FECs were operated in the configuration shown in table 6. This configuration gives an expected current draw of 800 mA/channel with the FEC operating at full capacity. It is, however, not expected that the FEC will have to operate at full power due to low occupancy and zero suppression. Whilst the design of the FEC is specific to TOAD, its operation at TOAD showcases the expected performance of SAMPA more generally in the type of environment that is required for HPgTPCs, which has not been previously tested.

Table 6. Operation parameters for SAMPA ASIC.

Parameter	Value
Supply Voltage	1.25 V
Reference Clock	160 MHz
Sampling Frequency	20 MHz
Operation Mode	Streaming/Trigger

3.1.2 Power Aggregation and Timing cards

The PAT card's primary function is to interface between the FEC cards and a standard Ethernet network, where data can be streamed to servers for data selection and storage. The prototype, designed at Imperial College London, utilises the advantages of FPGAs as the main data processing unit for the PAT. An image of the prototype board can be seen in figure 4.

Data aggregation is a primary requirement that drives the design of the PAT. The cards are designed to take inputs from 56 FECs each streamed over a differential pair into the FPGA package. The received data payloads from the many FECs are then extracted and propagated through a total of 64 buffers using Round-Robin (RR) logic to aggregate them into a single output stream. This output data stream is then wrapped into an Ethernet-based protocol and sent over the FPGA serial transceivers to the COTS network switch. The full layout can be seen in figure 5. Due to the low occupancy of the detector, the aggregation factor of 56 to 1 can be achieved easily within the FPGA fabric.

As well as data aggregation, the powering and timing roles of the PAT are also crucial. For powering, we investigated two possible solutions, for which a summary of key metrics is provided in table 7. The first solution provides 48 V to each PAT via dedicated terminals on the board. The FPGA DC/DC converter then outputs 12 V which is further fed into the FEC DC/DC converter to be transformed to the 1.8 V needed by the FEC on-board Low-Dropout regulator (LDO). The advantage of this scheme is that each PAT card receives a dedicated power line, minimising potential interference with signal traces on the board. The disadvantage is that ND-GAr will need dedicated power feedthroughs in the pressure vessel, which increases the overall feedthrough count.

2025 JINST 20 P12023

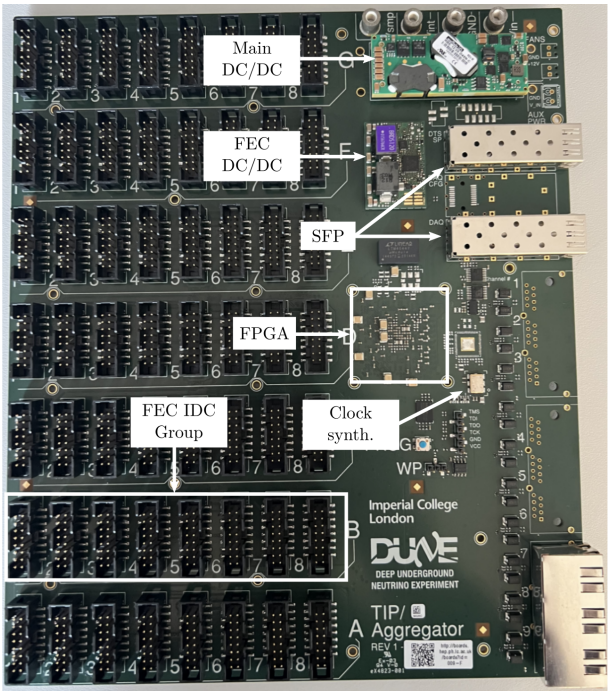

Figure 4. PAT card top-level image, on the top the two Direct Current to Direct Current (DC/DC) converters can be seen, one dedicated for the main voltage translation (Main DC/DC) and the second one for the FEC power (FEC DC/DC). The FPGA position is on the back side of the board at the position indicated, the FEC communication is organised in groups of Insulation-Displacement Contact (IDC) connectors shown with the white box on the left. Finally, there are four placeholders for clock synthesizers, however, only two are mounted, with one of them indicated by the white arrow on the right.

The second 'power over serial' powering scheme is inspired by Power over Ethernet (PoE), and aims to allow PAT cards placed within the pressure vessel to receive power over the incoming network cable. The circuit diagram used can be seen in figure 6. However, some modifications to the normal PoE standard are required due to the intended environment. PoE works by injecting a Direct Current (DC) component on top of the alternating data component using transformers. Copying the same methodology for a magnetised detector wouldn't be viable, as transformer performance may be impacted by the surrounding magnetic field. Therefore, an alternative system was developed using air-cored inductors to extract the differential signal from the bias voltage. Such a scheme provides the following two advantages. Primarily, the number of feedthroughs and cables required for the complete detector readout would be reduced compared to the direct power solution. A second major difference with this option is that a board that is outside of the pressure vessel is required to receive these power plus data signals. This second board would allow additional buffering and processing to be done on the FPGA if necessary. In the explored design, both the PAT boards inside and outside of the pressure vessel were based on the same board design with minor differences in the fitted components.

The implementation on an early prototype for the PAT cards was found not to give sufficient signal stability, so further optimisations are required to achieve a production-level system with this powering option. This work is underway with a new dedicated hardware board allowing modifications to both the power injection and the power decoupling circuits for further analysis. However, as it has not yet converged, we therefore chose the first solution (direct power) for the board design discussed here.

2025 JINST 20 P12023

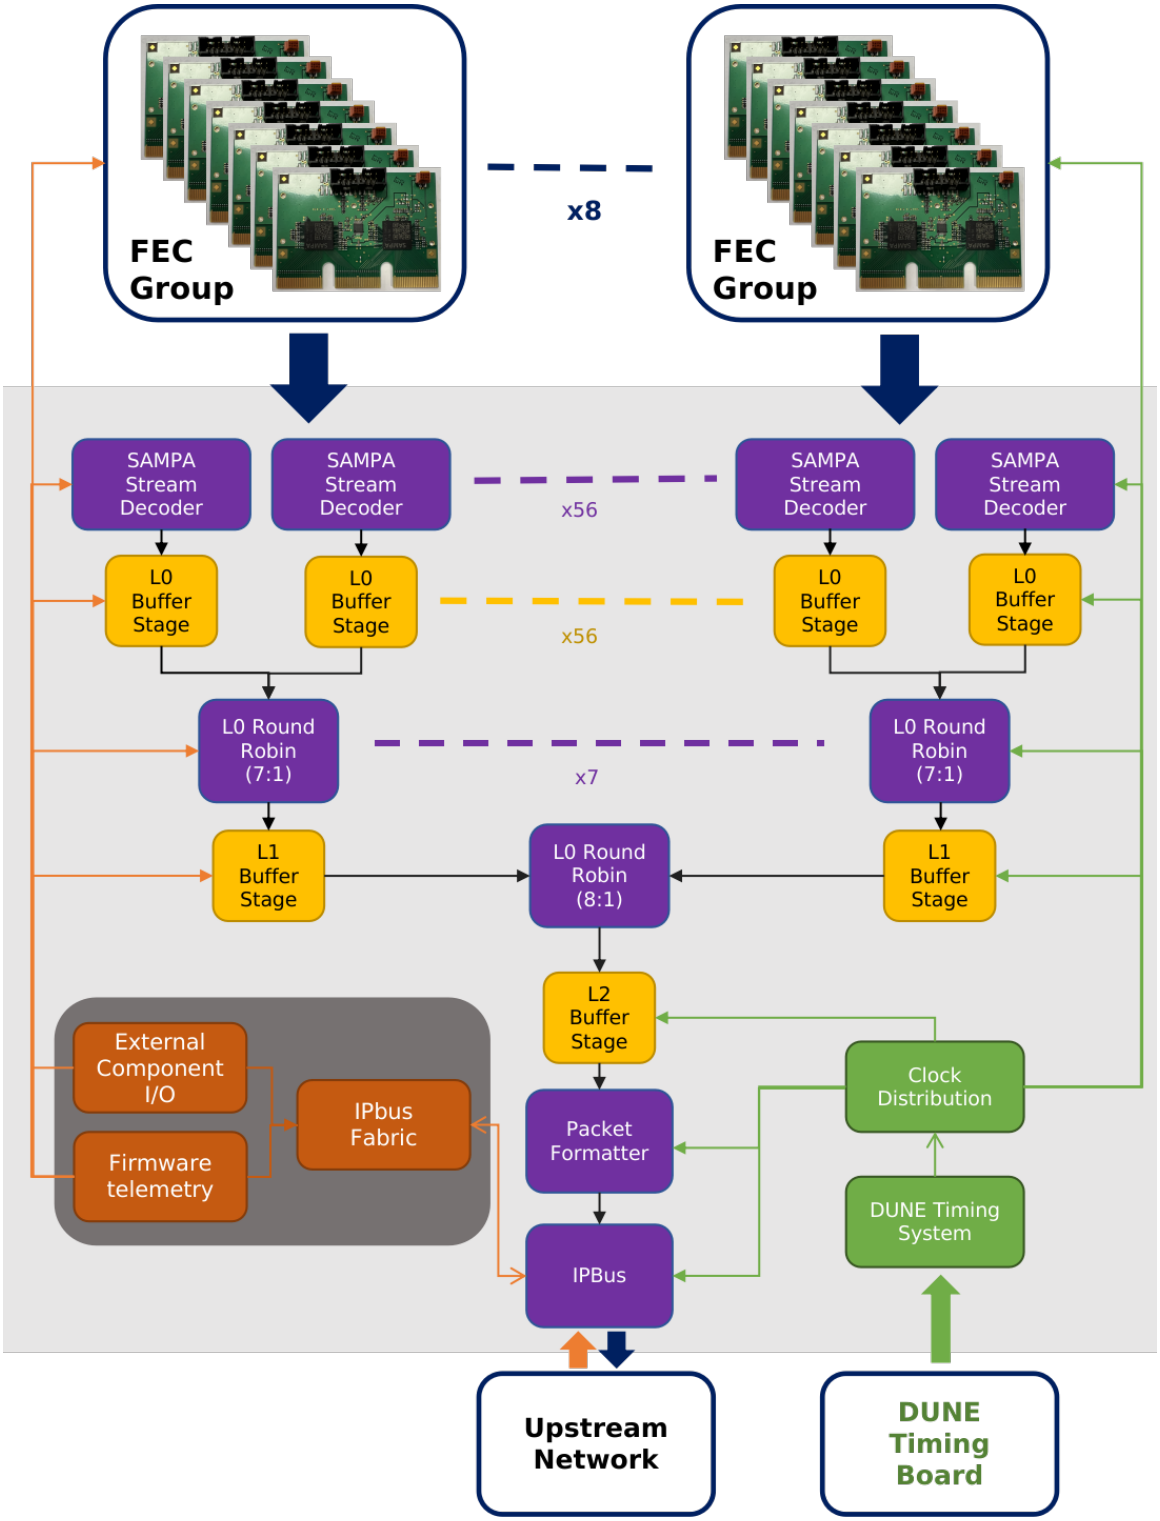

Figure 5. Diagram of the firmware schematic. The black lines show the data path, through successive levels of buffering (L0, L1, L2), format conversions and the RR data flow management scheme. The green lines show the clock distribution (derived from the DUNE Timing System). The orange lines show telemetry signals to/from internal/external components and the CCM system. The thicker arrows denote serial communication to other boards or switches in the system.

2025 JINST 20 P12023

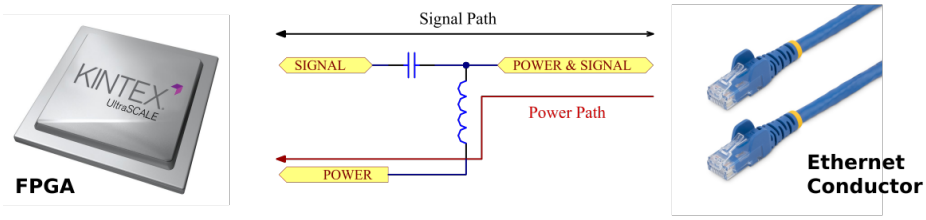

Figure 6. Circuit diagram for an alternate powering scheme where data and power are provided over the same cable. Power is injected by biasing the POWER & SIGNAL line to 48 V and thereby also the subsequent Ethernet Conductor (seen with red arrow: Power Path). At the FPGA, the capacitor causes the SIGNAL pins to only see the alternating component of the voltage and thereby removes the DC bias. Finally, the inductor acts as an open circuit switch, ensuring that signal doesn't leak towards the POWER node.

In addition, a study, summarised in table 7, was done to estimate the number of feedthroughs needed for both options and also a third potential solution which would use optical fibres to stream out the data and a separate power line for each PAT. This option allows several PAT cards to share one output feedthrough as optical fibres are smaller in form factor than the electrical connectors needed.

Table 7. Feedthrough and component count for different readout schemes when using the PAT-based system for an ~700 000 channel detector.

Component	**Power over Serial**	**Direct Power**	
Readout Type	Electrical	Optical	Electrical
# PAT (In-Pressure)	196	196	196
# PAT (Out-Pressure)	22	0	0
# FPGA	218	196	196
# Feedthrough	196	76	206
# Net. Switch Eth. Ports	22	196	196

Decisions on the power distribution of the system also have an impact on the timing. For the Power over Serial method, timing can be provided to the upstream PAT card that is outside the pressure vessel and all the cards within the vessel can then be synchronised by the serial transceiver recovered clock. In the case of Direct Power, the data stream is separate from power and is running through commercial Small Form-Factor Pluggable (SFP) modules. We therefore require a separate dedicated input for clock distribution. The PAT cards are designed to be able to interface with the DUNE timing system [12] over an optical fibre. In order to implement this, extra optical feedthroughs are required to receive all the timing fibres into the pressure vessel. The number of feedthroughs can be mitigated by fanning out the signal inside the vessel, but there are limits to the number of boards that can be supplied from any one fibre. In both cases, the clock received at the PAT cards in the pressure vessel is transmitted from the FPGA to a dedicated clock synthesiser for cleaning and then further transmitted to the FECs.

Furthermore, distribution of the FEC clock, synchronisation, and slow control signals happen, depending on the use case, via fan-out chips or multiplexers. Given that the designed readout is synchronous to the DUNE clock, all the FEC cards are expected to receive the same copy of the clock synchronisation signal. For this reason, an organisation of eight FECs in seven groups has been

2025 JINST 20 P12023

designed. The advantage of this method relies on the layout design of the PCB. The chosen FPGA can drive seven individual groups with the same clock synchronisation signal pulse. Each of these seven individual groups is then copied in phase to all eight FECs connected to that pin. The main reason for deciding to fan-out 1:8 was to ensure that the clock quality provided to all the FECs wasn't impacted by the fan-out ASIC. Equivalently, for the slow control interface, a series of multiplexers is implemented, allowing individual access to all the FECs, providing enhanced telemetry of the FEC operation.

Finally, the PAT cards are mounted at three positions on the OROC, in such a way that when 18 OROCs are mounted into a circle, the PAT cards are shared between neighbouring OROCs. In this way, the cable lengths required between FECs and PAT cards are minimised, and so, these connections are optimised. The PAT cards are mounted on aluminium plates, which, through thermal contact with the ROC, provide some amount of passive cooling. An image of the FECs and PATs connected to the OROC inside the vessel is shown in figure 7.

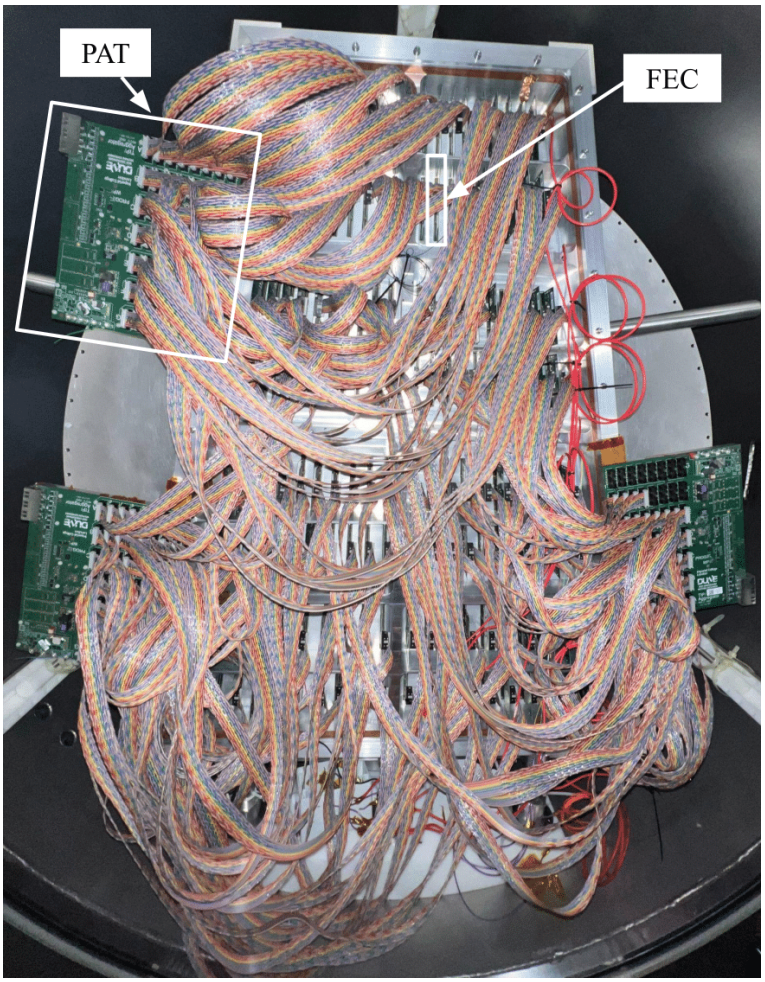

Figure 7. The readout side of the OROC, which has 156 FECs connected to 3 PAT boards via ribbon cables. These cables are either 0.5 m or 1.0 m in length. The PAT boards are mounted on aluminium plates and secured to the OROC frame. Each FEC has 3 connectors to the OROC, where each connector reads out groups of 21 or 22 pads. In this way, each FEC reads out the charge signal on 64 channels/pads, and so all 9984 channels of the OROC can be read out. A PAT and FEC have been highlighted in this image.

2025 JINST 20 P12023

3.2 Firmware design

The PAT card's main processing unit is a Kintex Ultrascale FPGA package. This device was selected as it provides the required number of serial transceivers to accommodate both readout schemes for ND-GAr whilst satisfying the remaining requirements. The power consumption (based on the AMD power estimate tool) is within the nominal limits for operation in gases at atmospheric pressure or above when actively cooled. Finally, the availability of this package at low cost through CERN and the extensive experience of using it that has been built up at CERN make this FPGA a favourable choice given the cost-driven nature of the design.

In order to take advantage of the FPGA, dedicated firmware had to be developed that allowed data to be reliably and efficiently streamed. Here, we discuss the firmware solution for direct powering, as this was the firmware that was eventually tested with the TOAD test beam. The design has two primary functions. First, the data flow, i.e. how FEC data are reliably received into the software. The second is Control Configuration and Monitoring (CCM) operations. A schematic representation of the firmware design can be seen in figure 5.

3.2.1 Data flow

The main objective of the data flow firmware is to reliably stream the FEC data into a server running the DUNE DAQ software [13]. The input to this firmware is the data stream from the daisy-chained SAMPAs, which is fed into the FPGA pins via a Low-Voltage Differential Signalling (LVDS) stream. The SAMPAs stream first a 50b header which includes important information, like the number of 10b samples that will be streamed, the time of arrival information, the fired channel number, etc. After the header, the payload follows which is made up of N 10b words, where N is defined in the header.

The data stream is then decoded using dedicated decoder firmware based on a Finite State Machine (FSM) that can handle all the different operational modes of the SAMPA. The extracted payload, along with some of the header information (channel number, timestamp, and number of clusters), is converted into a 32b stream that is subsequently sent through a RR based buffering system. The output Ethernet sample additionally contains a single 32b header with some of the header information received from the SAMPA header, as well as the 64b DUNE timestamp for data alignment.

The firmware implements three options for the streaming protocol used to do this. The first option, which is the one we used for all data shown here, uses the standard IPbus library over User Datagram Protocol (UDP) [14]. This option provides a stream of a given number of words on request from the DAQ software.

The second option interfaces the SiTCP suite from BeeBeans [15] with IPbus. SiTCP is a lightweight interface that provides streaming capabilities over Transmission Control Protocol (TCP). We developed an interface between the standard pull-request based IPbus firmware and the SiTCP IP core, thereby providing a TCP-based interface for the IPbus suite. While SiTCP includes its own UDP-based interface for CCM, it does not have enough address space for our application.

The third option was used for the power over serial powering mode investigations. The number of serial links on the PAT only allows for a single hardware interface between the in-pressure and out-of-pressure boards for both data and CCM and hence the SiTCP solution can be used in the out-of-pressure electronics. The link between the in-pressure and out-of-pressure electronics still cannot use the SiTCP solution, so it is implemented as an Aurora-based stream interface. In this operation mode, IPbus is not used at all, and the received data are converted into a TCP stream interface using SiTCP.

2025 JINST 20 P12023

3.2.2 Control, configuration and monitoring

The firmware includes many registers containing information allowing full telemetry of the hardware from software. The CCM aspect of the firmware's job is to monitor and set these registers. The design of the PAT cards allows the FPGA to also monitor and change the behaviour of several other components on the board using standardised slow control interfaces (e.g. Inter-Integrated Circuit (I2C), Serial Peripheral Interface (SPI), Power-Management Bus (PMbus)). The list of components can be found in table 8 along with a short explanation of their function.

Table 8. List of PAT card components that can be addressed through slow control.

Component	Interface	Application
FEC DC/DC Converter	PMbus	Power converter for FEC
SAMPA slow control	I2C	Configuration of SAMPA device
SFP Connectors	I2C	Ethernet I/O
Clock Synthesizer	I2C	Clock cleaning
Flash Memory	SPI	Remote firmware programming
EEPROM	I2C	Design information
I2C Switches	I2C	Slow control routing
Port-Expanders	I2C	Power switches for FEC

Each detector operation mode requires a specific configuration of these components. For example, with the inclusion of the I2C port-expanders, individual FECs can be turned off and removed from the run if there is a problem with their operation.

In addition to the external hardware components, the CCM part of the firmware implements a variety of standard IPbus registers and memories dedicated to monitoring the data flow or configuring operational modes. A list of these registers can be seen in table 9.

Table 9. IPbus configuration and monitoring register list.

Reg. Name	Control/Monitor	Number of Registers
Clock control	Enable clock transmission to fanout chips	1/FEC group
Clock rate control	Control over internal clock multiplexers for (320 MHz, 160 MHz or 80 MHz)	1/PAT
Serial links control	Ethernet serial link monitor and configuration	1/Serial Link
Sync. control	Send sync. signal to all FEC sampling windows	1/FEC group
Memory content	Live data flow memory content	1/FEC
RR control	Enable, Veto and Reset of RR logic	1/RR block
Logic Reset	Reset signals for all logic components	1/Firmware block in the design

2025 JINST 20 P12023

3.3 Software development

Operating the PAT-based readout requires software as well as the custom firmware. The PAT cards have three main functions: data flow, CCM, and timing. Each function has a software counterpart that has been developed for the TOAD test beam and designed so as to minimise modifications needed for integration in the final ND-GAr. All three sets of software are designed to run on standard commercial servers.

The data flow application is based on the DUNE DAQ software suite. For CCM, an IPbus-based framework has been developed, allowing full access to all the in-pressure electronics' configuration registers. Finally, for timing, the software is based on the framework from the DUNE timing system, which we have wrapped in the bespoke PAT CCM framework since the timing system uses IPbus for configuration as well.

3.3.1 DUNE DAQ integration

Once the data has been streamed out, it is received by the DUNE DAQ software. While this software already existed prior to this work, significant effort was required to integrate TOAD's DAQ into the larger DUNE software stack. A schematic of the general data flow within DUNE DAQ can be found in [16]. The integration of the TOAD readout into this framework involved adding the Ethernet packet structure described in section 3.2.1 into the upstream DAQ. Additionally, the zero-suppression functionality of the FECs allows for the PAT boards to only output data above the threshold we wish to be considered for offline processing. This means that TOAD also demonstrated that the DUNE DAQ software framework worked in untriggered continuous readout mode.

A further difference between TOAD and the other systems using DUNE DAQ is that the output stream data format allows each hit that is streamed out to contain a variable number of ADC samples, while all other systems had fixed numbers of samples per transmitted data packet. This implementation expands the capabilities of the DUNE DAQ software and allows it to handle variable-length data packets [13].

3.3.2 Control, configuration and monitoring software suite

The data from TOAD is received by standard DUNE software [13]. However, the CCM software suite is a new, scalable and modular software framework. The framework has two key base classes to talk to the two classes of objects on the PAT boards: firmware-implemented registers and external devices. The framework is summarised in figure 8. Each base class contains general methods needed by all registers of that type. Communication to the firmware registers is via IPBus, so the base class contains IPBus-specific methods, while the inherited classes for each specific register contain their own specific bit maps that relate their 32-bits to specific functions. This scheme allows scalability to all electronic devices that use IPBus and similar firmware implementation.

The external device base class contains methods that generate instructions for the FPGA to access registers on external devices. For a device that uses the I2C interface, the inherited class will contain an additional I2C address and a list of register addresses for the different registers inside the device. Additional slow control protocols, such as SPI or PMbus, are integrated in a similar manner.

The CCM suite also provides various low-level scripts that act as a Command Line Interface (CLI) for the user to manage the hardware. By combining multiple low-level scripts, one can create high-level

2025 JINST 20 P12023

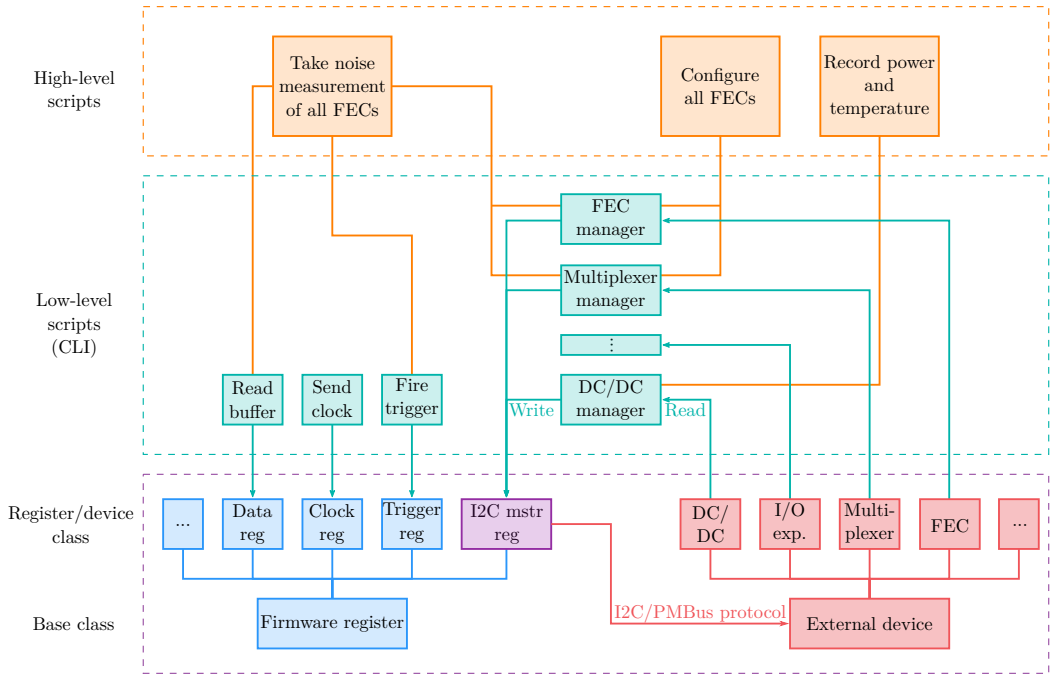

Figure 8. Block diagram for the CCM suite. Orange blocks at the top represent high-level scripts that interact with various low-level scripts to fulfil certain tasks. Low-level scripts are shown as green blocks, they provide CLI for users to interact directly with individual components on the PAT boards. Individual firmware registers and external devices are implemented as separate classes that inherit from the corresponding base class. Firmware registers are shown in blue and external devices are shown in red. The I2C master register in purple facilitates the I2C communication to external devices.

scripts that are tailored towards specific use cases. Examples include the noise measurements outlined in section 4.1 and continuous temperature and power monitoring used in section 4.2.

4 Results

4.1 Noise characterisation

Understanding the noise profile is crucial in verifying that we have a sufficient signal-to-noise ratio to identify the signals we are looking for. We characterise noise as the Root Mean Square (RMS) error of the waveform in the absence of any signal. Prior to measuring the noise and in order to ensure that everything is correctly set up, the PAT cards have been tested for synchronisation. This has been achieved on the table top setup, where the clock skew between the different FEC links has been measured and found to be in the sub-ps regime, indicating that all the FECs were clocked at the same time. The same test with the same outcome has been performed for the trigger signal, which is crucial for collecting noise data.

Noise data were collected by sending trigger signals to the FECs using the high-level script shown in figure 8. The noise is measured by taking the RMS of a 25 µs waveform, which corresponds to 500 samples at 20 MHz, for each channel in each FEC. Figure 9 shows the spatial distribution of the RMS noise of all FECs in the OROC, measured under 4.5 barA pressure. Channels with RMS noise less than

2025 JINST 20 P12023

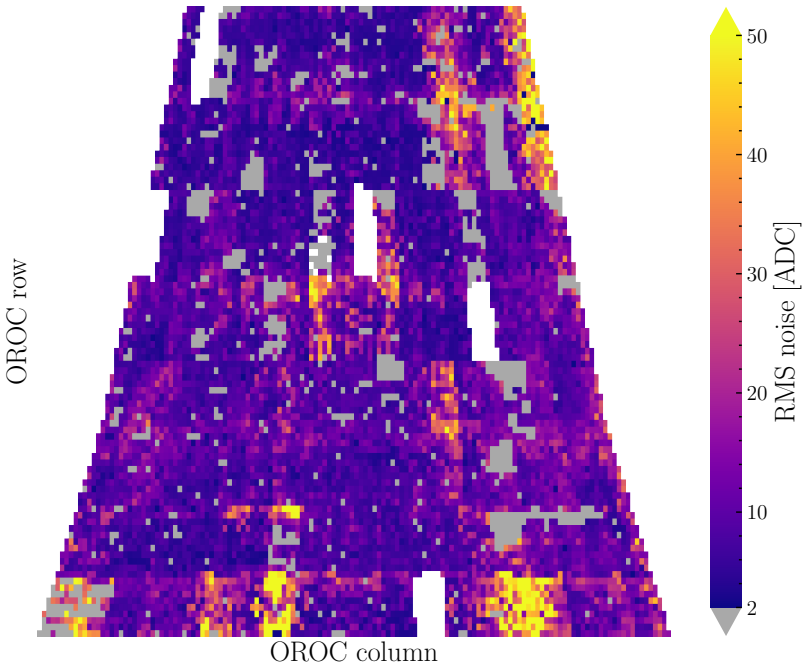

Figure 9. A noise map of the OROC pads in 4.5 barA pressure. The noise calculated from each channel on each FEC is mapped to the spatial position within the OROC and the RMS noise in ADC counts is plotted. The white regions are caused by non-responding FECs.

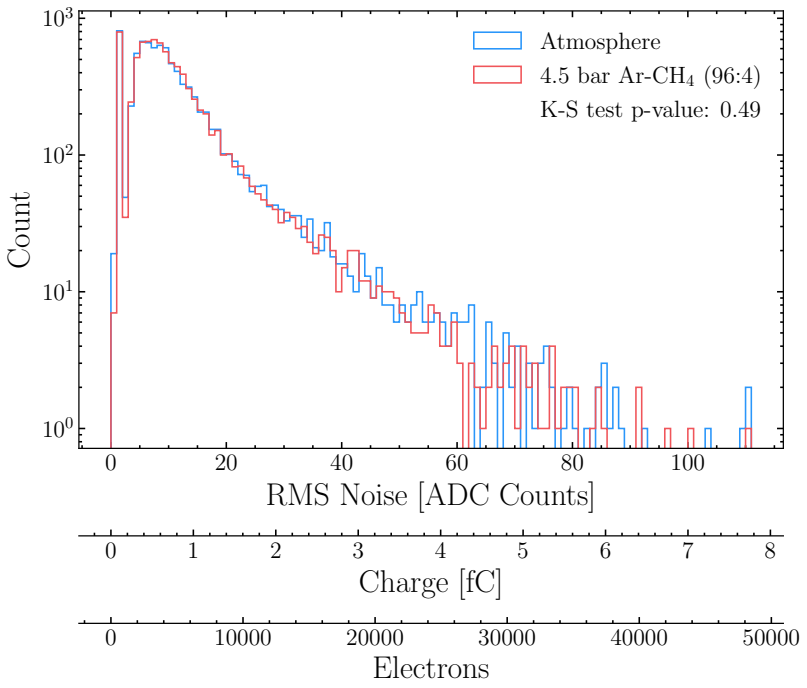

Figure 10. Distributions of RMS noise for all valid channels across all FECs in mode 1 (blue) and 2 (red) of table 2.

2025 JINST 20 P12023

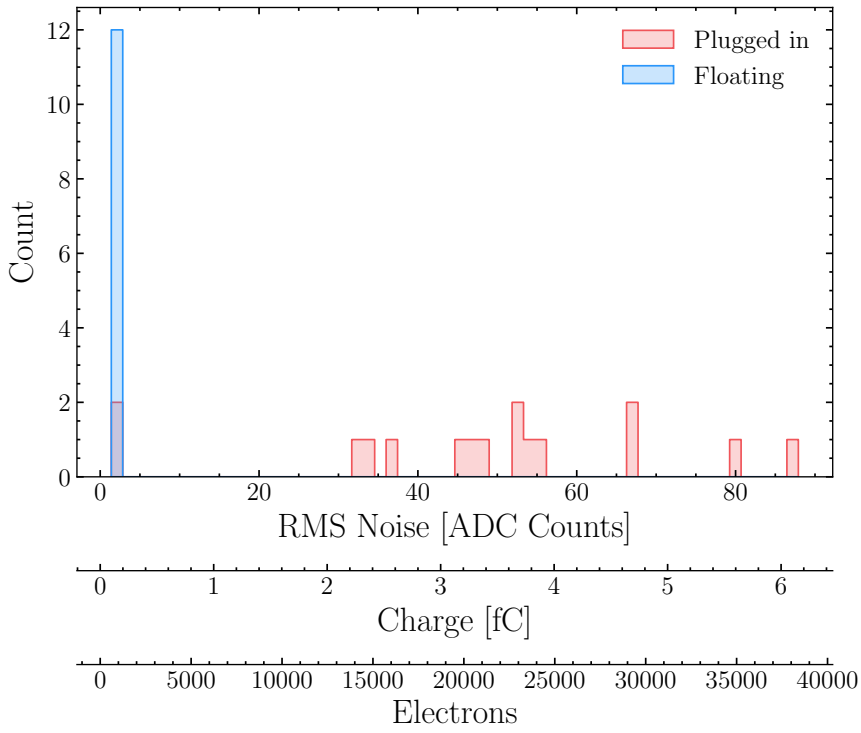

Figure 11. Distribution of RMS noise on a subset of 12 channels of a noisy FEC before (red) and after (blue) it is removed from the OROC.

2 ADC are coloured grey, while channels with noise above 50 ADC are shown with the highest colour in the scale. The blank regions are caused by FECs that did not respond to any I2C communications.

Figure 10 plots the distribution of the RMS noise. Noise from a total of 9600 (9656) channels was measured for the atmospheric pressure (4.5 barA) case, which corresponds to ~96% of the total number of OROC pads as described in table 1. We require the waveform baseline to be greater than 10 ADC counts in figure 10. This cut is placed to remove waveforms that have noise cut off at 0 ADC due to the low baseline, which makes the waveform appear to have low noise. This cut removes 10% of the data, and results in 8621 (8645) channels for the atmosphere (4.5 barA) case. Correction of the baseline can be applied to individual SAMPA channels, which would mitigate the issue of low baseline and therefore recover these channels. This calibration was not performed during this operation period of TOAD due to time constraints.

The noise distribution is bimodal with a single-bin peak at just above 0 and another peak at ~10 ADC. The first peak is caused by FECs that are not properly connected to the OROC. This peak is spatially shown as the grey patches in figure 9. This is demonstrated by measuring the RMS noise from the FECs after they are removed from the OROC and left floating. Figure 11 shows a comparison between the RMS noise for the same subset of 12 channels on one FEC, before and after the card is unplugged from the OROC. When the FEC is floating, all of the measured channels have RMS noise of < 5 ADC. This level of noise corresponds to the intrinsic noise level of the SAMPA (~1600 electrons), as shown in table 5, which is equivalent to 0.26 fC and 3.6 ADC. A gain of 14 ADC/fC is used for the conversion [11]. For most of the FECs that are well-connected to the OROC, their noise distribution is characterised by the second peak.

2025 JINST 20 P12023

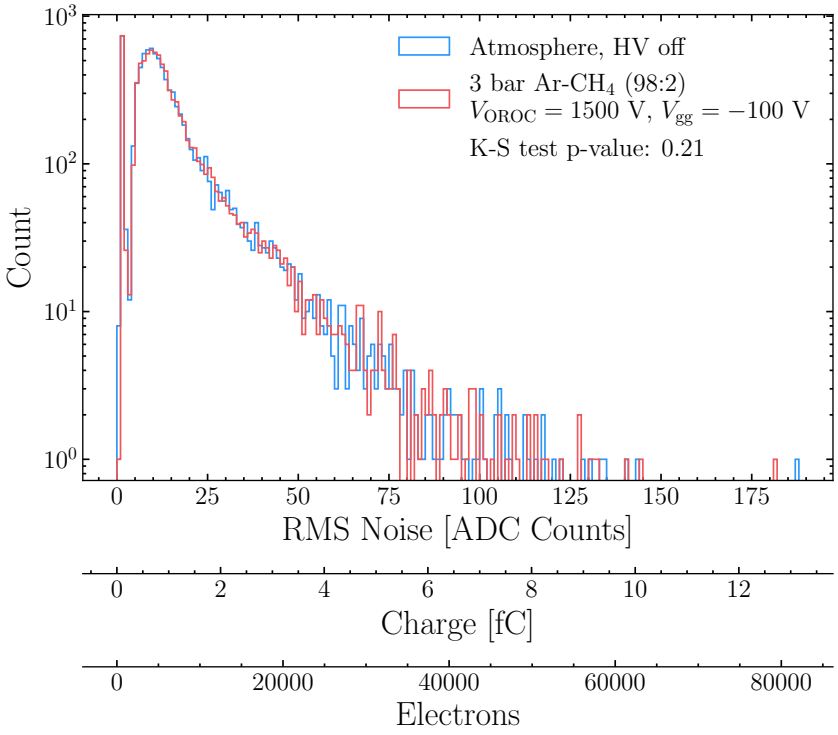

Figure 12. Distributions of RMS noise for all valid channels across all FECs in mode 1 (blue) and 3 (red) of table 2. Data for this plot were taken in a different condition and before some noise improvements were made compared to figure 10, so no direct comparison could be made.

The large tail that extends past 100 ADC in figure 10 is shown as distinct noisy regions in figure 9. This is due to a partial connection between the FEC and the OROC, which results in a floating ground, as this was not observed in benchtop tests with direct connection to the FECs. The difficulty in connecting the FEC to the OROC, for both the case of improper connection, which left the FEC afloat, and partial connection, which resulted in a floating ground, is due to the legacy OROC connectors used in TOAD. It is observed that the same FEC can gain good performance when moved to a different slot on the backplane. It is also worth noting that the ALICE ROCs are now end of life, so a future HPgTPC would be designed so as to achieve a more reliable connection to the front-end electronics.

From figure 10, no noticeable difference can be seen in the noise measurements taken in atmosphere (mode 1 of table 2) and 4.5 barA 96:4 Ar-CH_4 (mode 2). A Kolmogorov-Smirnov (K-S) p-value of 0.49 is obtained from the unbinned data, which suggests the same underlying distribution.

Additionally, we see that powering of the HV devices does not have a noticeable effect on the noise. This can be seen in figure 12 which compares the noise when the HV devices are off in atmosphere (mode 1 of table 2) and when $V_{\mathrm{OROC}} = 1500\,\mathrm{V}, V_{\mathrm{gg}} = -100\,\mathrm{V}$ in 98:2 Ar-CH_4 (mode 3). A K-S p-value of 0.21 is obtained and again suggests the same underlying distribution. The data for this plot were taken in a different condition and before some noise improvements were made to address the known issue of poor grounding in the test beam facility. As such, no direct comparison could be made with figure 10. Nevertheless, the bimodal distribution and the positions of the peaks are still the same between the two sets of data. We also tested mode 4 (same OROC and gating grid voltages and −16 kV cathode voltage) for a small subset of the channels, shown in figure 13. The

2025 JINST 20 P12023

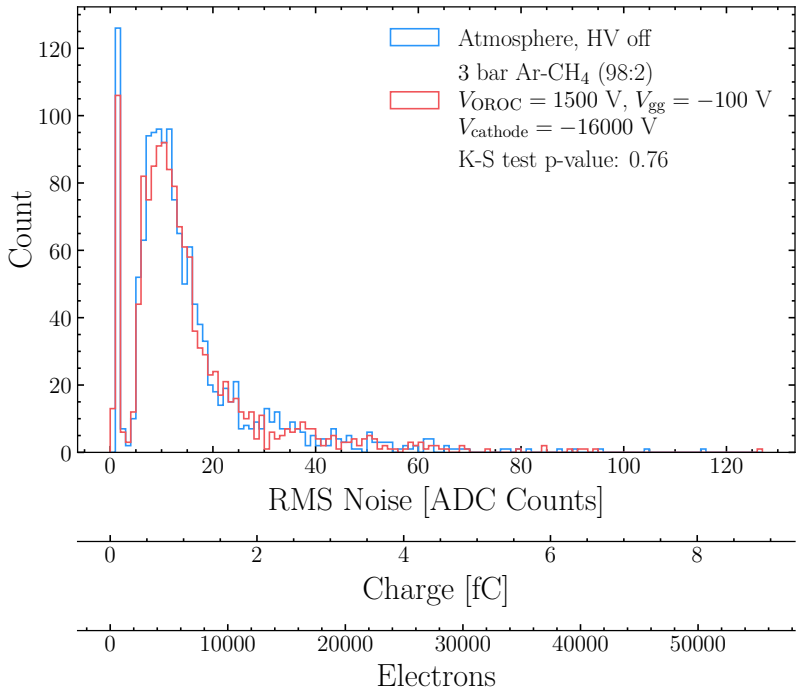

Figure 13. Distributions of RMS noise for a subset of channels across all FECs in mode 1 (blue) and 4 (red) of table 2. A linear scale is used due to the smaller sample size.

same data for mode 1 (blue) was used for both figure 12 and figure 13, but for the latter figure, only the subset of channels matching those collected for mode 4 were shown. A K-S p-value of 0.76 is obtained and the main features in the noise distribution remained identical.

The contribution of noise from different frequencies can be studied through Discrete Fourier Transformation (DFT) of the waveforms. Figure 14 shows the DFTs from a single channel on a FEC. To avoid random noise, 1000 waveforms are collected and the average of the DFT of 1000 individual waveforms is plotted. Large peaks at 400 kHz and integer multiples of it can be noticed in figure 14. This corresponds to the switching frequency of the FEC DC/DC converter. Less prominently, the smaller peaks at multiples of 200 kHz correspond to the switching frequency of the main DC/DC converter.

To summarise, the DC/DC components on the PAT boards are major contributors to the noise measured by the FECs. However, since this iteration of the board is not the one expected to be deployed in ND-GAr in a future revision, more care will be taken with regard to the choice of the DC/DC converter to not limit the physics capability of the detector. In addition, FECs that were not conductively connected to the OROC would appear to have extremely low noise. High noise channels in the OROC are seen in small regions, usually corresponding to a single FEC, and could be due to improper grounding of the OROC connectors.

To assess whether the level of noise measured in TOAD would be sufficient for a neutrino experiment, we estimate the size of the signal on a FEC that could be generated by a minimally ionising particle. The expected signal-to-noise ratio can be calculated using stopping power 1.5 MeV cm^2/g, atmospheric argon gas density $1.66 \times 10^{-3}\,\mathrm{g/cm}^3$, and ion pair production energy 26 eV, giving $95\,\mathrm{e}^-/\mathrm{cm}$ at atmosphere [17]. At 10 bar, using a projected gas gain of 10 000 from [7] and accounting for the change in gas density, one obtains $9\,500\,000\,\mathrm{e}^-/\mathrm{cm} = 1520\,\mathrm{fC/cm}$. We take 1 cm as an

2025 JINST 20 P12023

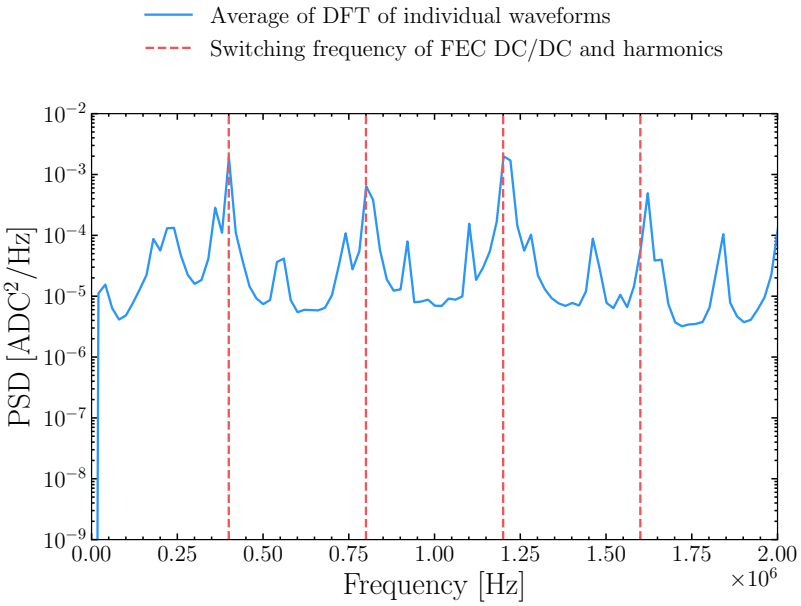

Figure 14. Power spectral density of a single channel on a FEC. The average of all 1000 DFT of the waveforms is shown in blue. Dashed lines show the switching frequency of the FEC DC/DC and its harmonics.

approximate track length across an OROC pad, i.e., a channel in a FEC. Simulations have been done to show that the electrons spread across multiple pads, with only 22% lie within 0.5 cm of the centre. Furthermore, it takes ~7 samples for the charge to be fully deposited. So using 14 ADC/fC from [11], on average, a minimally ionising particle is expected to deposit 668 ADC per channel per sample. The 95 percentile of figure 10 gives a conservative estimate for noise of 32 ADC, which leads to an estimated signal-to-noise ratio of ~21.

4.2 Thermal profile

To achieve a reasonable gain in a gas TPC, a stable pressure must be maintained. The thermal profile of the system thus needs to be well-understood as in a closed vessel the temperature directly affects the pressure. Cooling is currently achieved in TOAD through thermal contact between the aluminium mounting plate of the PAT board and the OROC. The FECs are cooled through contact with the gas in the vessel. The temperature and the power output from each PAT board are recorded via the FEC DC/DC converter on each board, and the average values in 4.5 barA are plotted in figure 15. The idle power consumption, with no power delivered to the FECs, for the PAT boards all connected to the same power bus, ranges between 0.89-3.60 W. Additionally, the idle temperature in equilibrium with the environment has been found to be 40 °C, which is well within the operational margin provided by AMD [18].

The FECs constitute the majority of the power consumption. With all available FECs in full operation, the power consumption for each PAT board ranged between [60.7, 64.8]W, [74.0, 75.4]W, and [79.7, 81.6]W for the PAT board with 43, 53, and 55 FECs connected, respectively. Assuming the power consumption scales linearly with the number of FECs, an average power consumption of 1.4 W per FEC can be calculated. This corresponds to a maximum power consumption of ~85 W per PAT board.

2025 JINST 20 P12023

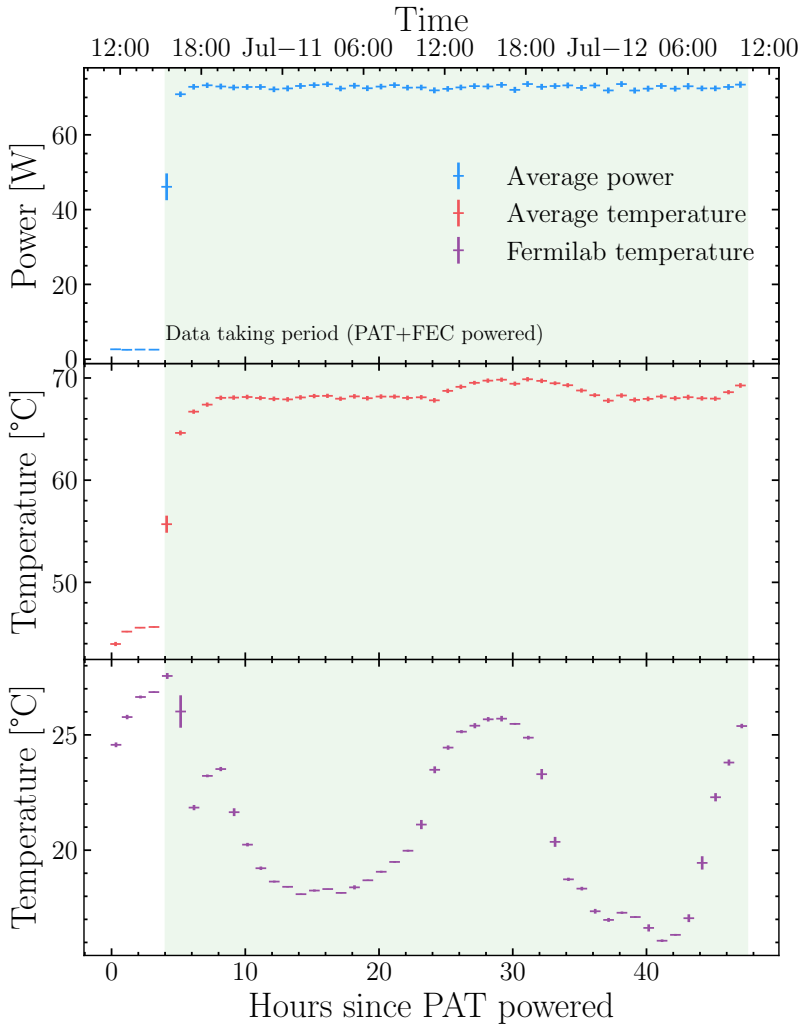

Figure 15. Power output (blue) and temperature (red) of the FEC DC/DC converter since the PAT boards were powered in 4.5 barA pressure, averaged across the three PAT boards. The error bar is mainly due to the difference across the three boards. The temperature measured from a Fermilab weather station is plotted in purple. The slight increases in PAT temperature at noon on both 11 Jul and 12 Jul correlate with high environmental temperature. The green region is the data taking period where both the PAT boards and FECs are in operation.

Relying solely on the thermal coupling between the PAT boards and the OROC, an average change of temperature of 23 °C on the DC/DC converter was observed as a result of the operation of the FECs. While the final temperature reached is within the operation threshold of the electronics, a large temperature change and use of the gas as the cooling medium can lead to a long stabilisation time for the pressure in the vessel as well as uncertain gas drift properties. As shown in figure 16, by fitting an exponential to the pressure profile, it takes ~6 hours for the pressure to have a rate of change of pressure $\dot{P} < 1$ mbar/hr. It is also worth noting that the heavy pressure vessel is in contact with the gas, increasing the time to heat up. Furthermore, the OROC is also in thermal contact with the TOAD vessel and the gas, further increasing the pressure stabilisation time and also exposing it

2025 JINST 20 P12023

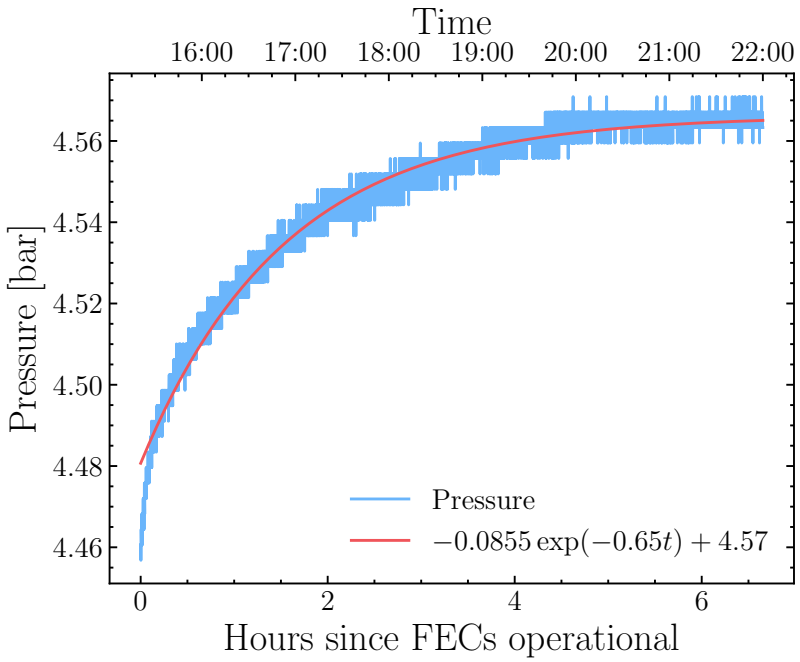

Figure 16. Pressure in barA since the FECs have been powered (blue), fitted with an exponential function (red). The pressure measurement is limited by the manometer resolution of 0.0038 barG.

to the high equilibrium temperature. Such a long stabilisation time prolongs the time taken to go from the detector being off to stable operations, increasing downtime.

Each PAT board is contributing a maximum of only ~85 W of power which is well within the range that can be mitigated by the addition of active cooling or a better heat sink design. This investigation of the thermal profile is a necessary component of future research.

Stable performance of the electronics was observed during the pressure operation. During the time-constrained test-beam period, we observed stability in pressure for the duration of multiple days, with no significant variation in the electronics performance observed. Similar stability was observed in atmosphere during the test-beam period, as well as throughout benchtop testing.

5 Conclusion

Many readout systems developed for particle physics experiments target the high-occupancy environment of a hadron collider experiment. Scaling systems such as readout electronics from sPHENIX or ALICE would be unnecessarily costly. Similarly, technology such as LArPix would require overclocking to meet the resolution requirements for ND-GAr. Commercial solutions, such as CAEN, would also be considerably costly. We have, instead, described a custom system designed for much lower occupancy detectors, such as those in neutrino physics experiments. The system leverages the lower buffering requirements to use far fewer FPGAs per channel and thereby significantly reduces the cost of the system. Even though it is a prototype system manufactured at a small scale, the readout for TOAD (including FECs, PAT cards and all the hardware to connect and power them) costs less than $20 000 for ~10 000 channels. Scaling this up to the full DUNE Phase-II ND, ignoring likely economies of scale, gives a conservative cost estimate of ~$1.4 million.

2025 JINST 20 P12023

The ND-GAr detector that this system was designed for requires operation in a magnetic field and at high pressure (up to 10 barA). The high-pressure requirement also necessitates a limited number of cables exiting the pressure vessel. We have described the designed system and demonstrated its operation at a range of pressures and gas mixtures up to 4.5 barA. These tests were carried out through the TOAD experiment and involved the readout of ~10 000 channels of an MWPC from the ALICE TPC.

The system performed well, with stable operation seen throughout the test period, which included multiple days of running at atmosphere and pressure as well as benchtop testing. However, the passive cooling used in the test has been identified as an area that must be improved before final deployment. Nevertheless, the total heat load is well within the range that can be dealt with using active cooling. The noise characteristics of the system were measured and found to be within the range that will allow us to read out the intended signals for almost all channels. We have identified poor connection due to the mounting of the FECs in the legacy connectors of the OROC as the primary item contributing to channels that are not readable. This is specific to the OROC used in TOAD, and is not expected to be a problem for future HPgTPC, e.g. ND-GAr, which will use different connectors. Notwithstanding these necessary minor modifications, we have shown that this system is able to fulfill the needs of a high-pressure gas TPC detector for a fraction of the price of systems used for LHC type detectors.

Acknowledgments

We would like to thank the STFC, Royal Society and UKRI, U.K.; the DOE, U.S.A.; and the Imperial College President's Scholarship scheme for funding this work. Particularly, for Fermilab, this document was prepared using the resources of the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, Office of High Energy Physics HEP User Facility. Fermilab is managed by Fermi Forward Discovery Group, LLC, acting under Contract No. 89243024CSC000002.

We would also like to acknowledge the invaluable assistance of the mechanical and electrical workshops at Royal Holloway University of London and Imperial College London, as well as the staff at the Fermilab test beam facility.

References

[1] DUNE collaboration, *Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report*, *Instruments* **5** (2021) 31 [arXiv:2103.13910].

[2] P5 collaboration, *Exploring the Quantum Universe: Pathways to Innovation and Discovery in Particle Physics*, arXiv:2407.19176 [DOI:10.2172/2368847].

[3] DUNE collaboration, *DUNE Phase II: scientific opportunities, detector concepts, technological solutions*, 2024 *JINST* **19** P12005 [arXiv:2408.12725].

[4] ALICE collaboration, *Upgrade of the ALICE Time Projection Chamber*, CERN-LHCC-2013-020 (2014).

[5] H. Klest, *Overview and design of the sPHENIX TPC*, *J. Phys. Conf. Ser.* **1498** (2020) 012025 [arXiv:2110.02082].

[6] ALICE collaboration, *The ALICE experiment at the CERN LHC*, 2008 *JINST* **3** S08002.

[7] A. Ritchie-Yates et al., *First operation of an ALICE OROC operated in high pressure Ar-CO_2 and Ar-CH_4*, *Eur. Phys. J. C* **83** (2023) 1139 [arXiv:2305.08822].

2025 JINST 20 P12023

[8] ALICE collaboration, *ALICE: Physics Performance Report*, *J. Phys. G* **32** (2006) 1295.

[9] ALICE TPC collaboration, *The upgrade of the ALICE TPC with GEMs and continuous readout*, 2021 *JINST* **16** P03022 [arXiv:2012.09518].

[10] A. Deisting et al., *Commissioning of a High Pressure Time Projection Chamber with Optical Readout*, *Instruments* **5** (2021) 22 [arXiv:2102.06643].

[11] J. Adolfsson et al., *SAMPA Chip: the New 32 Channels ASIC for the ALICE TPC and MCH Upgrades*, 2017 *JINST* **12** C04008.

[12] A. Barcock et al., *Timing and synchronization of the DUNE neutrino detector*, 2023 *JINST* **18** C01067 [arXiv:2210.15517].

[13] A. Abed Abud, *Trigger and Data AcQuisition (TDAQ) System Design*, CERN Engineering & Equipment Data Management Service (2023), https://edms.cern.ch/document/2812882.

[14] C. Ghabrous Larrea et al., *IPbus: a flexible Ethernet-based control system for xTCA hardware*, 2015 *JINST* **10** C02019.

[15] T. Uchida, *SiTCP*, Electronics system group, IPNS, KEK (2021), https://www.sitcp.net/doc/SiTCP_eng.pdf.

[16] DUNE collaboration, *Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume IV: Far Detector Single-phase Technology*, 2020 *JINST* **15** T08010 [arXiv:2002.03010].

[17] Particle Data Group, *Review of Particle Physics*, *PTEP* **2022** (2022) 083C01.

[18] AMD, *Kintex Ultrascale FPGAs Data Sheet: DC and AC Switching Characteristics*, 2025.

2025 JINST 20 P12023