

15TH TOPICAL SEMINAR ON INNOVATIVE PARTICLE AND RADIATION DETECTORS
14–17 OCTOBER 2019
SIENA, ITALY

KM3NeT acquisition: the new version of the Central Logic Board and its related Power Board, with highlights and evolution of the Control Unit

D. Real,^{a,1} C. Bozza,^{f,g} D. Calvo,^a P. Musico,^b P. Jansweijer,^c S. Colonges,^d
V. van Beveren,^c F. Versari,^{e,2} T. Chiarusi,^e F. Carriò,^a G. Pellegrini^e and A.F. Díaz^{f,g}
on behalf of the KM3NeT collaboration

^aIFIC — Instituto de Física Corpuscular (CSIC — Universitat de València),
c/Catedrático José Beltrán, 2, 46980 Paterna, Valencia, Spain

^bINFN, Sezione di Genova, Via Dodecaneso 33, Genova, 16146, Italy

^cNikhef, National Institute for Subatomic Physics,
PO Box 41882, Amsterdam, 1009 DB, Netherlands

^dAPC, Université Paris Diderot, CNRS/IN2P3, CEA/IRFU, Observatoire de Paris,
Sorbonne Paris Cité, 75205 Paris, France

^eINFN, Sezione di Bologna, v.le C. Berti-Pichat, 6/2, Bologna, 40127, Italy

^fUniversità di Salerno e INFN Gruppo Collegato di Salerno, Dipartimento di Fisica,
Via Giovanni Paolo II 132, Fisciano, 84084, Italy

^gUniversity of Granada, Dept. of Computer Architecture and Technology/CITIC,
18071 Granada, Spain

E-mail: real@ific.uv.es

ABSTRACT: The KM3NeT collaboration is currently building two deep sea neutrino telescopes at the bottom of the Mediterranean sea. The acquisition electronics for the first phase of the telescopes has been produced and several Detection Units have already been deployed. For subsequent phases, an improved version of the acquisition electronics has been designed with the goal of reducing the power consumption and improving the long term reliability of the boards. The control software suite, named Control Unit, is also being upgraded, in particular to overcome hardware failures. In this article, we present the last versions of the Central Logic Board and its associated Power Board, together with the evolution of the Control Unit.

KEYWORDS: Digital electronic circuits; Neutrino detectors

¹Corresponding author.

²Speaker.

Contents

1	Introduction	1
2	Acquisition hardware improvements	2
2.1	Central Logic Board version 4	2
2.1.1	Central Logic Board clock generation system	3
2.2	Power Board version 3	4
3	Control Unit highlights and evolution	5
3.1	The Detector Manager	6
3.2	The TriDAS Manager	7
3.3	The Master Control Program	7
3.4	The Data Base Interface	8
3.5	The Local Authentication Provider	9
4	Conclusions	9

1 Introduction

The KM3NeT Collaboration is building two neutrino telescopes at the bottom of the Mediterranean Sea [1]. One, called Oscillation Research with Cosmics in the Abyss (ORCA), 40 km away from Toulon and the second one, called Astroparticle Research with Cosmics in the Abyss (ARCA), 100 km away from the southern tip of Sicily. The telescopes, designed to detect neutrinos by means of the Cherenkov photons induced by relativistic charged particles while traveling through the detector, consist in three-dimensional arrays of PhotoMultiplier Tubes (PMTs). Each Detection Unit (DU) is a string, anchored on the sea bed and tightened by a buoy, bearing 18 Digital Optical Modules (DOMs) [2]. Each DOM houses 31 PMTs together with its acquisition electronics [3], which main acquisition board is the Central Logic Board (CLB), supplied by its ancillary board, the Power Board (PB). The PMT signals are digitized at the PMT base boards. The PMT base board outputs a Low Voltage Differential Signal (LVDS) with a duration equal to the time over a pre-configured threshold, called Time over Threshold (ToT). The LVDS signals are conducted to the CLB by an aggregation board called Signal Collection Board (SCB). Then, the arrival time and the ToT are determined, with one nanosecond resolution, by the Time to Digital Converters (TDCs) implemented in the Field Programmable Gate Array (FPGA) of the CLB. In order to properly determine the trajectory of the Cherenkov photons it is needed to synchronize the different DOMs with a resolution better than one nanosecond. The synchronization of the DOM is performed in the FPGA of the CLB by means of the White Rabbit (WR) protocol [4], which allows the data communication and the synchronization using the optical link available. The shape of the detector changes continuously because of the sea currents. In order to obtain high directional accuracy and time resolution, the position and

orientation of each DOM needs to be measured. Acoustic beacons provide reference signals whose arrival times are recorded by hydrophones and sent via the CLB to the shore station. While the main acquisition tasks are performed at the DOM level, specifically at the CLB, the control of the acquisition is performed by the Control Unit (CU), a software suite that allows to handle the DOM functions. It is a highly modular system, designed for scalability and high availability, which orchestrates the operation of the CLB as well as the triggering and data processing system. Section 2 is dedicated to a new version of the acquisition electronics. Section 3 briefly describes the CU and its latest evolution about hardware failure management. Conclusions are presented in section 4.

2 Acquisition hardware improvements

For the first phase of the project there has been produced enough CLB and PB to manufacture DOMs for 31 DUs, some of which have been already deployed and are currently taking data. Taking into account the return of the experience and in order to increase the performances and the reliability of the boards, new versions of both boards are being developed. The improvement of reliability is of paramount importance as, once deployed the DU, it will be impossible to perform any kind of maintenance. The main PB improvements include a modification of some of the DC/DC converters used to operate the PB more efficiently. This increase in efficiency results in decreasing the power consumption, thermal losses and the temperature inside the DOM and therefore increasing the reliability of the electronics.

2.1 Central Logic Board version 4

Several changes have been introduced in this new version of the CLB in order to improve the mechanics coupling with the rest of the DOM mechanics. Also some additional sensors have been introduced, as a pressure sensor, and the clock system has been improved to reduce the phase noise of the CLB clocks which has a great importance to improve the synchronization of the detector. In addition, the layout has been modified to use a new type of optical transceiver with higher reliability and a watchdog system has been introduced in order to further decrease the risk of losing access to the CLB due to an error in the reconfiguration of the acquisition firmware. An extensive list of the main improvements (figure 1) are here reported:

- Two flash memories. In the previous version there was only one Serial Port Interface (SPI) flash memory, while in this version it has been added an extra SPI flash memory. The main objective is to reduce the cycles of write/read in the flash memory storing the image in order to increase the operation life of this memory. The second memory is dedicated to an internal log and its end of life is not critical. The memories wear out of use. Separating the log in a different memory increase the life of the main flash memory.
- A new optical transceiver is used. This transceiver, from Glenair, has a higher reliability and a different layout with respect to the previous one. Therefore, the CLB has been modified to accommodate this new model of transceiver.
- New sensors have been included to the CLB. A pressure sensor has been added, and the compass, accelerometer and gyroscopes are moved from the previous daughter board to the CLB itself.

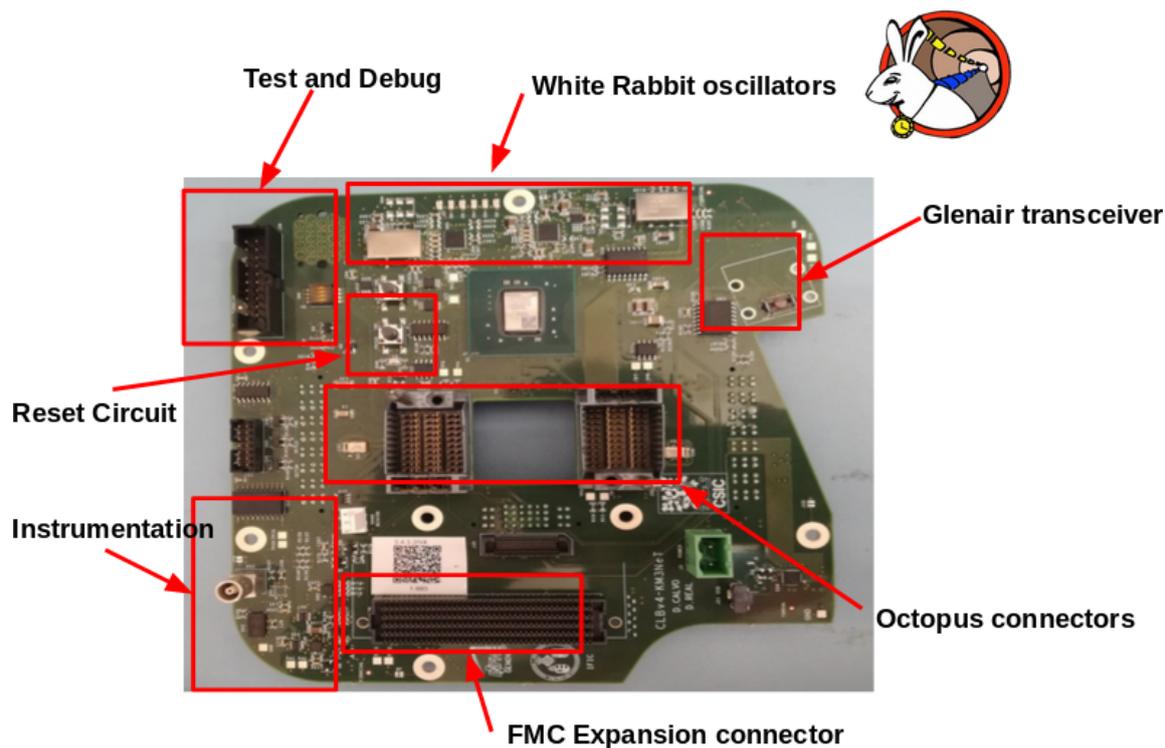


Figure 1. Prototype of CLBv4 with its main elements that have been improved with respect to the previous version.

- It has been included a hardware watchdog and the reset system has been updated.
- Two different clock schemes have been included in order to evaluate the best solution for WR from the point of view of stability and phase noise.

2.1.1 Central Logic Board clock generation system

The clock scheme of the CLBv4 is presented in figure 2. In the first system implemented, a clock of 25 MHz is generated by a quartz oscillator, and, in order to generate the needed 125 MHz, a low jitter clock generator is used to increase the frequency to the 125 MHz needed. The same applies for the generation of the 124.992 MHz. Instead, in the second system, the 125 MHz and the 124.99 MHz are generated directly by a quartz oscillator manufactured to oscillate at that frequency. As the quartz creates the frequency exactly it improves the quality of the clock signal. Both frequencies are needed to measure the phase using the Dual Mixer Time Difference (DMTD) technique.

Four prototypes of the CLBv4 have been produced, where the new feature including the new clock generation system has been tested. The tests performed (see figure 3), where a CLBv4 connected to a standard WR switch has been turned off and on several times, show the stability of the clock generated when used by WR. The skew of the Pulse Per Second (PPS) presents very low jitter, with values below ± 25 ps standard deviation, a value well below the jitter requirements, of 100 ps. These values include also the effect of the reset and power down of the board as it is in this operation when usually higher differences are obtained. The jitter has been measured

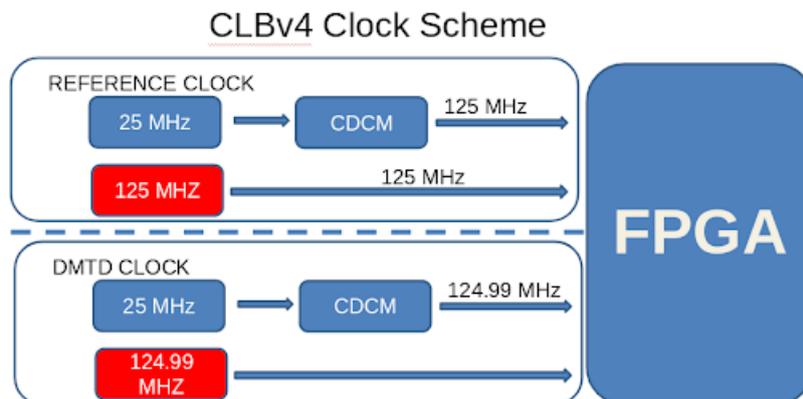


Figure 2. Clock generator scheme for the WR system of the CLBv4.

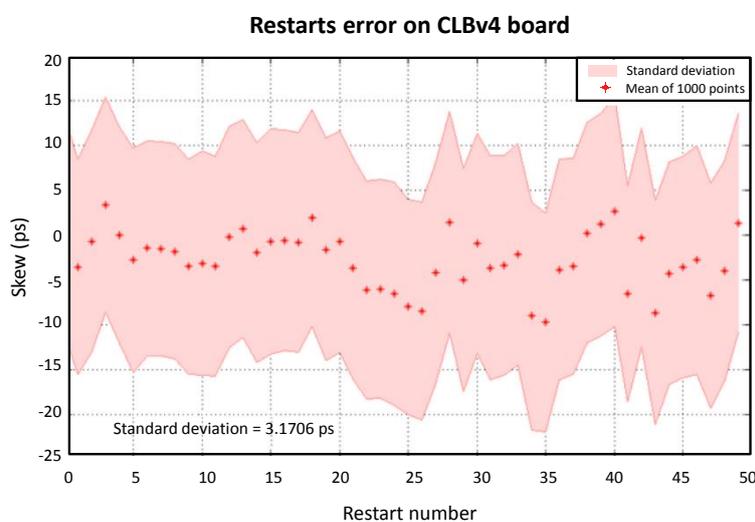


Figure 3. Tests performed on the CLBv4. The skew of the PPS signal has been measured after the restart of the board. As it can be seen the restart of the board almost does not add skew.

with conditions close to the final ones, including the use of fiber with the same length and it is not expected to deteriorate once deployed. The overall results are very promising for their use in KM3NeT.

2.2 Power Board version 3

The new version of the PB modifies several of the DC/DC converters in order to improve the global efficiency of the board. In particular, there has been changed the DC/DC converters of the rails of 2.5 V, 3.3 V and 5 V, obtaining an overall decrease in the consumption of 1 Watt when the PB supplies all the DOM acquisition electronics. (See table 1).

The reliability of the board has also been improved. The FIDES analysis has been performed and, when compared with the previous version of the board, the Failure In Time (FIT) (given in failures in 10^9 hours) value has decreased to 783 (while the FIT of the previous version of the PB was 947).

Table 1. DOM Voltages, currents and efficiency of the PB broken down by power rail. The efficiencies of the previous version of PB are given as well for comparison.

Volt (V)	Current(A)	Efficiency PB v2.3 (%)	Efficiency PB v3.0 (%)
1	0.13	80	80
1.8	0.33	80	80
2.5	0.33	60	78
3.3	0.81	65	90
3.3 PMTs	0.46	90	90
5	0.10	60	90

3 Control Unit highlights and evolution

CLBs are driven and operated in a coordinated fashion by higher level software, named the CUs. The same software is used in detectors deployed in the deep sea as well as in test benches for integration and qualification of PMTs, DOMs and DUs. Figure 4 shows the structure of the CU and how its services relate to the detector, the remote central database and the software used for triggering and data acquisition (TriDAS).

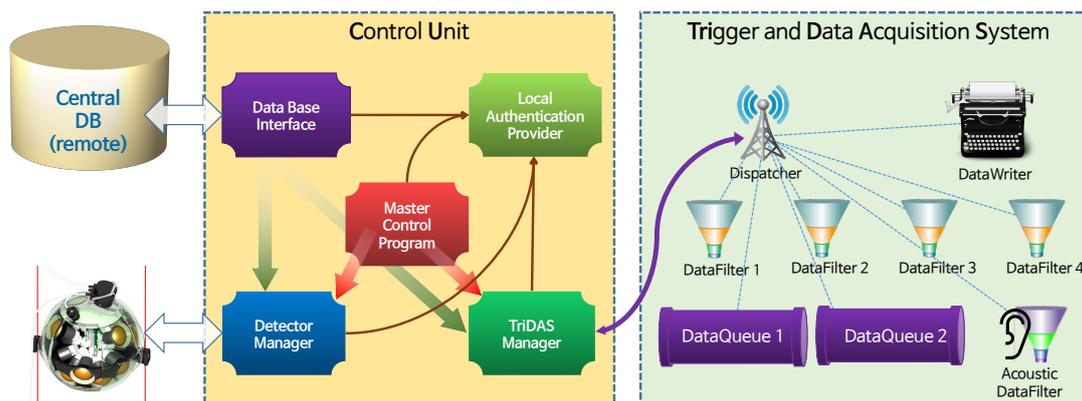


Figure 4. Functional overview of the Control Unit and its relationships with physical objects and software components. Only the flow of control information is shown.

The CU is built as a suite of software services. Each one has a certain degree of independence from the others, in the sense that it can work without communicating with other services for some time (minutes or even hours in some cases). Such modular approach has several advantages:

- Improved fault tolerance, because a single service getting stopped cannot cause data taking to stop.
- Increased separation of tasks and easier software development and maintenance.
- Ability to upgrade services one by one with the acquisition running, thus minimising the downtime.

Each CU service hosts its own web server that provides a graphical user interface relying on an HTTP layer; the latter is also used for a disconnected remote procedure call protocol, to communicate with each other service. The CU can run all on a single physical machine or can be distributed on several machines. In the following subsections, the various CU components will be described.

3.1 The Detector Manager

The Detector Manager (DM) is the CU service that directly controls the CLBs, through a UDP-based network protocol. Each CLB is controlled individually with a closed feedback loop that is asynchronous with respect to all others. This approach allows quick recovery of operation if a CLB loses contact or hangs and must be restarted. During the execution of recovery actions data acquisition can continue with the other CLBs.

The DM sets all operational parameters of the CLBs and continuously reads out monitoring data, including temperature of several devices, humidity, acceleration, compass, power levels and network link strength. The sampling frequency can be tuned in the range 0.1–10 Hz. All such parameters are exploited in two ways:

- They are stored into the so-called *datalog* files, staged to be later uploaded to the remote central database.
- They are exposed in a *Virtual Directory* which is accessible via HTTP. This is used by the graphical user interface (figure 5) and can be directly read by external programs or scripts to provide real-time monitoring and reporting.

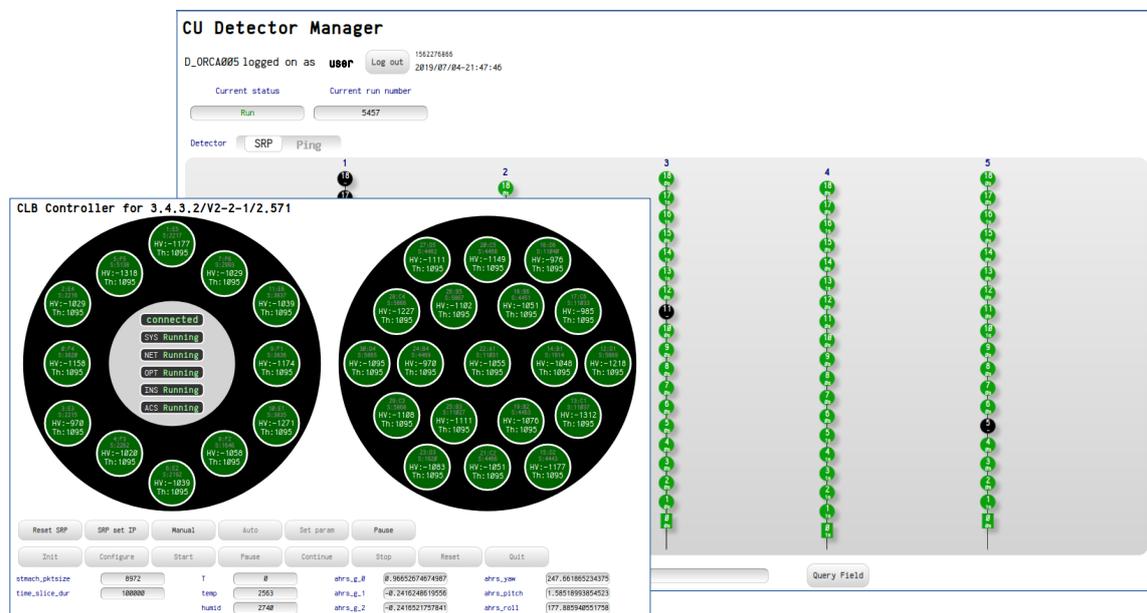


Figure 5. Screenshot of the DM graphical user interface with the monitor of a single DOM superimposed to the overall monitor of five DUs.

For each CLB, the DM sets up a mirror Finite State Machine (FSM), shown in figure 6, that is used to drive the CLB state machine, using the aforementioned UDP protocol. The state machine is conceptually identical for all components, both hardware and software, of the KM3NeT data acquisition and triggering, providing a single operational paradigm for the whole detector. Some states of the FSM are transient, whereas others correspond to operational targets of the CU:

- *Off*: PMTs are not powered.
- *On*: PMTs are powered but neither optical nor acoustic data are generated.
- *Run*: PMTs are powered, optical and acoustic data are generated.

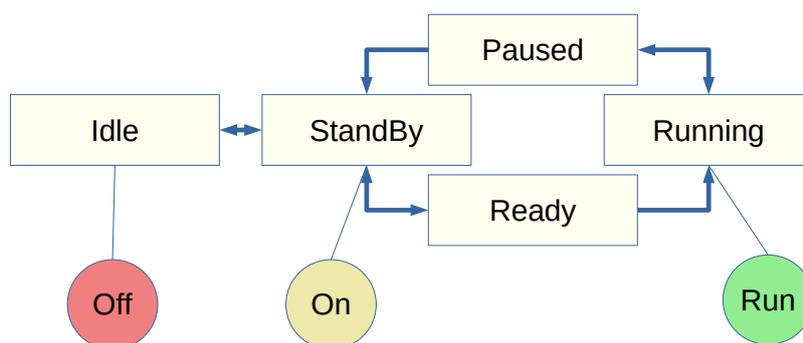


Figure 6. The FSM (boxes) that defines the states of each component of the data acquisition system and the operational targets (circles) of the Control Unit.

3.2 The TriDAS Manager

The TriDAS Manager (TM) controls processes running on a local computer farm that performs *triggering* and online *data acquisition and processing system* (TriDAS). They are built on top of the same Finite State Machine as the CLBs. Data from the CLBs are sent to the Data Queues which rearrange them for the Optical and Acoustic Data Filters, so that each Data Filter can apply the triggering algorithms on the full detector. Triggered events are sent to the Data Writer(s) to be saved on permanent storage. Communication between the TM and the TriDAS processes occurs through a TCP-based network protocol (*ControlHost*) and a dispatcher that provides bidirectional time-ordered streams of messages. Like the DM, the TM fills *datalog* files that are later written to the remote central database, noting all events, actions performed and running parameters of the TriDAS processes. The same data are also exposed in a Virtual Directory for real-time access by the graphical user interface and possible other programs.

3.3 The Master Control Program

The Master Control Program (MCP) coordinates all other CU services. The MCP is the single authority that is responsible for defining:

- The current configuration of the detector (KM3NeT detectors are built incrementally and addition of new DUs causes a change in the detector configuration).

- The current set of operational parameters (defined as *runsetup* in KM3NeT terminology).
- The current run, i.e. a well defined timespan during which the runsetup does not change. For practical reasons of data management, it is convenient to break a run every few hours, hence forcing a run change even without a runsetup change.
- The current operational target.
- A list of scheduled jobs with a priority system to orderly manage the active time of the detector including routine duty tasks and special tasks such as calibrations or tests for detector development. At each time, the job with the highest priority is executed. A change of the current job causes a run change. A lower priority job may be preempted by a higher priority one and when the latter ends the MCP returns to the previous job if not expired yet, but now with a different run number. The logic is sketched in figure 7. For the convenience of detector operators, jobs can be automatically generated and enqueued in advance.

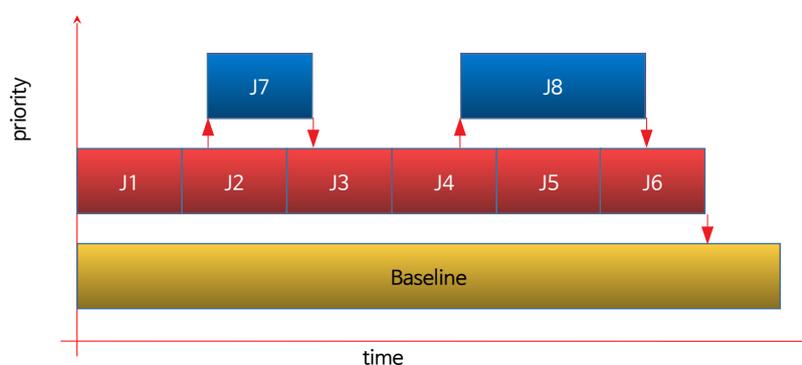


Figure 7. MCP jobs on three priority levels. The run number changes every time there is a switch from a job to another. A baseline job stands as a *fallback* to define the detector state in case nothing else is scheduled. Normal duty data-taking is performed by jobs of equal duration at the priority level above the baseline. Special tasks happen at a higher priority level.

Like other CU services, the MCP also produces *datalogs* and exposes its state in a Virtual Directory.

3.4 The Data Base Interface

The CU services use information provided by the remote central database [5] and use it to book-keep the data acquisition activities and monitoring parameters. A Data Base Interface service is devoted to facilitate the interaction. The DBI polls the central database for user authentication information and user privileges, also considering data taking shifts done by human operators, which are centrally managed. The DBI downloads the detector definition, structure and calibration parameters that are to be applied at a certain time. On the other hand, run book-keeping information, jobs, *datalog* files and processed acoustic data (*Time of Arrival* of signals to the hydrophones) are uploaded to the central database when several MBs are available to optimise the transfer. The central database may not be reachable through the Internet because of several reasons, including network malfunctions or software/hardware failure on either end of the communication channel. Acquisition should not

be stopped just because the database is temporarily unreachable. The Data Base Interface manages a two-way local cache to make vital information always available locally and to optimise the upload of data to the database. In addition, the DBI frees the local CU services from knowing the inner structure of the database and the SQL queries. If the database schema changes, the DBI is the single point where changes have to be reflected, all other services being left unperturbed. CU services communicate data with the DBI in the form of data files in XML format. The DBI is in charge of performing the translation to/from SQL. This also allows for evolution of the database schema or possible technology change, provided data can still be represented in a way that is convenient to the CU services. The DBI also exposes its status in a Virtual Directory accessible via HTTP.

3.5 The Local Authentication Provider

The CU has been designed to work in a protected environment, where security is provided by proper network access rules. Nevertheless, users have to be identified and proper permission sets have to be generated and applied to prevent possibly harmful mistakes. The Local Authentication Provider (LAP) acts as the local authority for authentication and identity management. In addition to this, the LAP used to hold the link between each CU service and the machine where it runs. Recently the tasks of the LAP have been upgraded to provide a full Dynamic Resource Provisioning and Failover. The computational power needed to run the TriDAS largely overwhelms the requirements of the CU, but a hardware failure in the CU can paralyse the acquisition for hours or days if an experienced administrator is not ready on call. This is particularly important in case a failure occurs during a transient neutrino emission, which is too short to wait for administrative intervention to migrate services from the failed machine to a healthy one. In order to make the system resilient and capable of automatic self-reconfiguration, CU services are configured to run on several machines and kept standing by except for one instance that is running. TriDAS services are also handled by the Dynamic Provisioning and normally occupy all available resources. If a failure is detected, the resources allocate to the TriDAS are revoked and used to power the CU services that automatically migrate to an unaffected server. In this working fashion, there is one LAP per server running a *Health Checker* that periodically diagnoses the ability of the machine to comply to its tasks. If one LAP reports that the machine has a problem, or fails to answer polling by other LAPs, the corresponding server is blacklisted until it comes up again and all the TriDAS and CU services hosted there are migrated elsewhere or run in downgraded mode (e.g. the number of Optical Data Filters is reduced). A complete reconfiguration takes 5–10 s and requires no action from human operators.

4 Conclusions

This work presents the last version of the main acquisition boards of KM3NeT, the Central Logic Board together with its Power Board, which improve its reliability as well as its functionality. Improvements in reliability and fault-tolerance have been made for the enhanced version of the Control Unit.

Acknowledgments

The authors acknowledge the financial support of the funding agencies: Agence Nationale de la Recherche (contract ANR-15-CE31-0020), Centre National de la Recherche Scientifique (CNRS), Commission Européenne (FEDER fund and Marie Curie Program), Institut Universitaire de France (IUF), IdEx program and UnivEarthS Labex program at Sorbonne Paris Cité (ANR-10-LABX-0023 and ANR-11-IDEX-0005-02), Paris Île-de-France Region, France; Shota Rustaveli National Science Foundation of Georgia (SRNSFG, FR-18-1268), Georgia; Deutsche Forschungsgemeinschaft (DFG), Germany; The General Secretariat of Research and Technology (GSRT), Greece; Istituto Nazionale di Fisica Nucleare (INFN), Ministero dell’Istruzione, dell’Università e della Ricerca (MIUR), PRIN 2017 program (Grant NAT-NET 2017W4HA7S) Italy; Ministry of Higher Education, Scientific Research and Professional Training, Morocco; Nederlandse organisatie voor Wetenschappelijk Onderzoek (NWO), the Netherlands; The National Science Centre, Poland (2015/18/E/ST2/00758); National Authority for Scientific Research (ANCS), Romania; Plan Estatal de Investigación (refs. FPA2015-65150-C3-1-P, -2-P and -3-P, (MINECO/FEDER)), Severo Ochoa Centre of Excellence program (MINECO), Red Consolider MultiDark, (ref. FPA2017-90566-REDC, MINECO) and Prometeo and Grisolia programs (Generalitat Valenciana), “la Caixa” Foundation (ID 100010434) through the fellowship LCF/BQ/IN17/11620019, and the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 713673, Spain.

References

- [1] KM3NET collaboration, *Letter of intent for KM3NeT 2.0*, *J. Phys. G* **43** (2016) 084001 [[arXiv:1601.07459](https://arxiv.org/abs/1601.07459)].
- [2] E. Leonora and V. Giordano, *Design and production of the digital optical module of the KM3NeT project*, *EPJ Web Conf.* **136** (2017) 04008.
- [3] D. Real and D. Calvo, *Digital optical module electronics of KM3NeT*, *Phys. Part. Nucl.* **47** (2016) 918.
- [4] M. Lipinski et al., *White Rabbit: a PTP application for robust sub-nanosecond synchronization*, *Proceeding to IEEE International Symposium on Precision Clock Synchronization for Measurement, Control and Communication*, Munich, Germany (2011) [[DOI](https://doi.org/10.1109/ISPCS.2011.6044444)].
- [5] A. Arnauld and C. Bozza, *The relational database system of KM3NeT*, *EPJ Web Conf.* **116** (2016) 07004.