PROCEEDINGS OF SCIENCE

The data acquisition system of the KM3NeT detector

Simone Biagi[*†]
INFN, Laboratori Nazionali del Sud, Catania, Italy
E-mail: biagi@bo.infn.it

Tommaso Chiarusi
INFN, Sezione di Bologna, Bologna, Italy
E-mail: chiarusi@bo.infn.it

Paolo Piattelli
INFN, Laboratori Nazionali del Sud, Catania, Italy
E-mail: piattelli@lns.infn.it

Diego Real
IFIC, Instituto de Física Corpuscular (CSIC-Universitat de València), València, Spain
E-mail: real@ific.uv.es

The KM3NeT neutrino telescope is part of a deep-sea research infrastructure being constructed in the Mediterranean Sea. The basic element of the detector is the Detection Unit, a 700 meter long vertical structure hosting 18 Digital Optical Modules (DOMs). The DOM comprises 31 3" photomultiplier tubes (PMTs), various instruments to monitor environmental parameters, and the electronic boards for the digitization of the PMT signals and the management of data acquisition. Dedicated readout electronics have been developed and are installed inside each DOM, allowing to measure the time of arrival and the duration of photon hits, on each of the 31 photomultiplier tubes, with a time resolution of 1 ns. Moreover, the data transmission system supports the maximum throughput from a DOM of 200 Mbps, which correspond to a photon-hit rate of $\sim$130 kHz per each 3" PMT. Due to the extreme operation conditions of the abyssal site, the all-data-to-shore concept is used in order to minimize the complexity of the offshore detector. The processing of the data transmitted to onshore is performed by the Trigger and Data Acquisition System (TriDAS). The networking infrastructure and computing resources are conceived to be modular and scalable in order to manage the full data rate from the final cubic-kilometer scale telescope. The electronics and the DAQ system described in the poster are currently under test in the first Detection Unit to be deployed offshore Toulon.

The 34th International Cosmic Ray Conference,
30 July- 6 August, 2015
The Hague, The Netherlands

[*]Speaker.
[†]On behalf of the KM3NeT Collaboration.

PoS(ICRC2015)1172

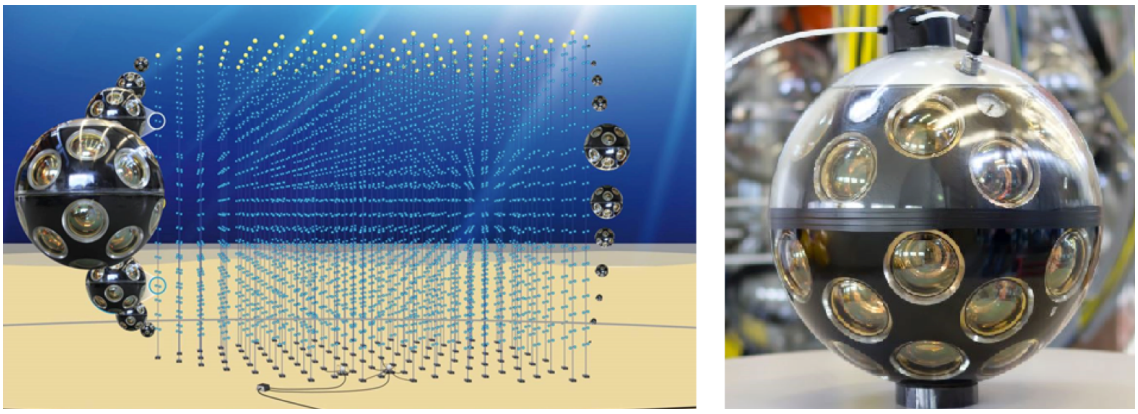

Figure 1: Left: A sketch of the KM3NeT detector in its final configuration. Right: A Digital Optical Module.

1. Introduction

KM3NeT (KM^3-scale Neutrino Telescope) is a future European research facility in the Mediterranean Sea which will house a neutrino telescope of cubic kilometer scale [1]. Cherenkov light from neutrino induced secondary particles will be detected by an array of optical modules consisting in high pressure resistant glass vessels with photomultipliers inside. This vessel is called Digital Optical Module (DOM) and it is composed of 31 3" Photomultiplier Tubes (PMT) distributed around the glass sphere, which collect the Cherenkov light and transform it into electronic signals [2]. 18 DOMs are arranged on string-like structures anchored on the sea bed and kept vertical by a buoyancy system, the so-called Detection Units. The Detection Units are connected with submarine Junction Boxes and through them to shore for power feed and data transmission. Fig. 1 shows an overview of the KM3NeT detector and an assembled DOM.

The all-data-to-shore data taking approach follows the choice done by ANTARES [3]. In this contest, no trigger is done underwater, but all signals from PMTs arrive onshore where they are triggered and processed by farm of computers. The Trigger and Data Acquisition System (TriDAS) must be able to handle an extremely large throughput of data arriving from off-shore; the DAQ modularity and scalability are needed properties for a detector made of hundreds of Detection Units that could be added or recovered without interfering with the data acquisition.

2. Readout electronic components

Inside the DOM, the 31 PMTs are suspended in a foam support structure: 19 PMTs in the lower hemisphere and the remaining 12 in the upper hemisphere. Each PMT has its own adjustable high voltage supply integrated in the PMT base. The PMTs collect the Cherenkov light and convert it into electronic signals. In order to translate these signals into the arrival time of the photons, they are processed by Time to Digital Converters (TDC) core embedded in the Field Programmable Gate Array (FPGA) of the Central Logic Board (CLB) [4]. The CLB integrates the White Rabbit Protocol [5], a fully deterministic Ethernet-based network for general purpose data transfer and synchronization, that allows to synchronize all the KM3NeT DOMs with 1 ns resolution. The data

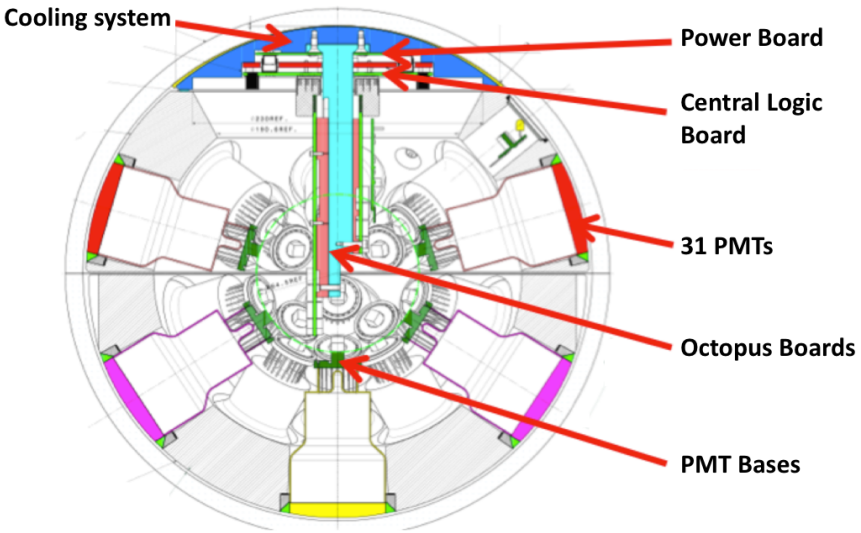

Figure 2: The Digital Optical Module in details. The Power Board, the Central Logic Board, the Octopus Board and the PMT Bases are indicated by arrows.

provided by the PMT bases is collected and distributed to the CLB by means of two boards (one for each hemisphere), the so-called Octopus Boards. It also contains the electronic components for an optical link to the shore. All necessary DC power is provided by the Power Board (PB) [6]. An aluminium structure provides heat conduction between the electronics inside and the exterior of the sphere. In Fig. 2 a schematic view of the DOM is shown. In this section, all the electronic boards contained in the DOM are reviewed.

2.1 Photomultiplier Base

The PMT Base is in charge of discretizing the signal read by the PMT and to provide the High Voltage (HV) for the PMT. The PCB contains a pre-amplifier, a comparator (Time over Threshold) and an I^2C decoder. Connection to the PMT is done by flying leads with a PCB that has a diameter of 38 mm. Every PMT must give the same output signal when it is hit by a single photon. The gain of a PMT depends on the supplied high voltage. The HV for each PMT is individually adjustable from 800 to 1400 V. Consequently each PMT gets its own HV circuit board. I^2C protocol is used to be able to program the PMT Base and to change the HV. The power consumption of each PMT Base is around 4.5 mW.

An additional function of the PMT Base is the digitization of the analogue output signal of the PMT. The output signal is converted from a charge signal to a voltage signal, followed by a conversion to a digital level by a comparator, resulting in a Time over Threshold (ToT) signal. The comparator can be adjusted to the required ToT value using I^2C protocol. The ToT signal is transferred to the DOM logic by a LVDS connection. To identify the PMT Base an ID circuit is added. The analogue and digital signal conversion functionality of the PMT Base is performed in an ASIC. The PMT Base diagram is shown in Fig. 3.

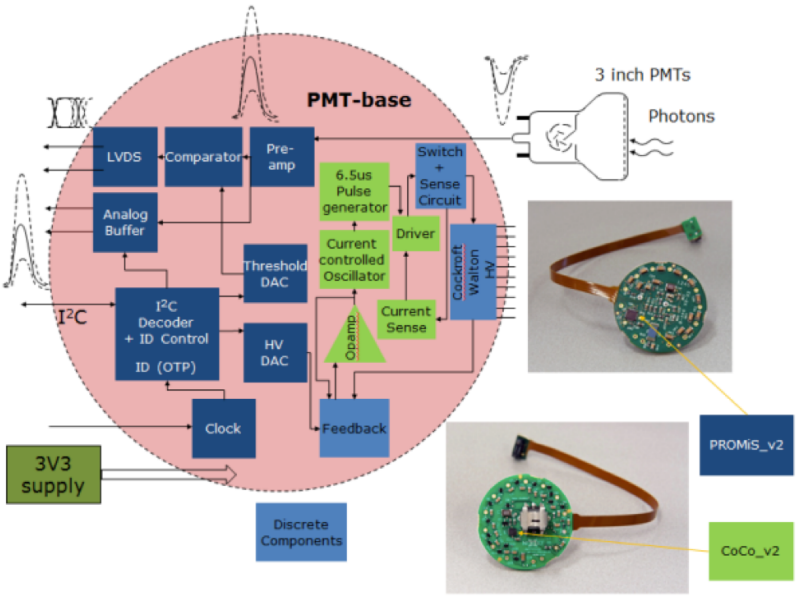

Figure 3: The diagram and 2 pictures (front/rear) of the PMT base.

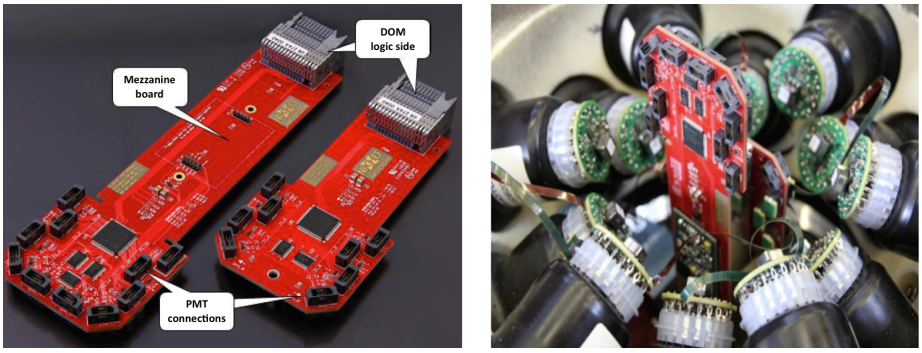

Figure 4: Left: Large and Small Octopus Boards. Right: Close up of the Octopus Board mounted in the DOM.

2.2 The Octopus Board

Within each hemisphere, the LDVS signals from the PMT bases are collected on a custom electronics board (Fig. 4), referred to as the Octopus Board and transferred to the CLB of the DOM. The Octopus Boards also provide connection for the electrical power to the PMT bases and the I^2C communication control. For each PMT, the electrical power can be switched on/off individually by the slow control and in case of overload of a PMT the power will be switched off automatically. This can be monitored by the Fault Flag (FFLG). A clock enabled signal to the PMTs for the I^2C communication avoids digital interference.

2.3 The Power conversion Board

For an efficient transfer of the electrical power the voltage level must be high and the current low, because of the power loss depends on $I^2 \cdot R$ of the cable. In addition, the different electronic components of the DOM require many different voltage levels for their performance. Therefore,

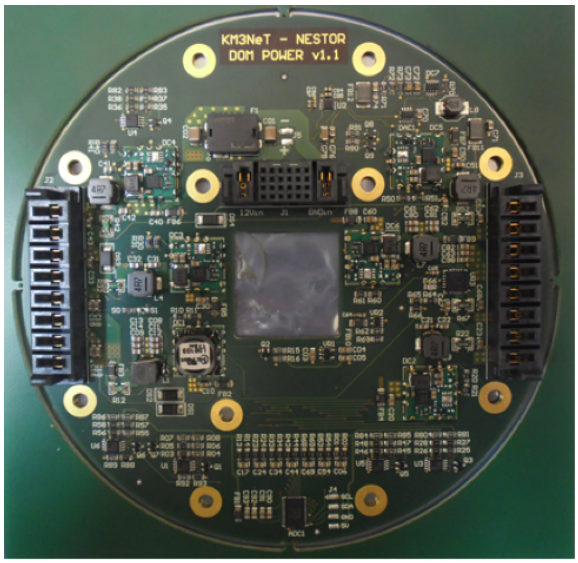

Figure 5: Picture of the first prototype of the DOM Power Board.

the Power conversion Board (PB) inside the DOM derives all different client voltage levels from an input voltage of 12 V. Modern converters at high frequency are used to obtain a high efficient power conversion. To protect the other electronics inside the DOM against possible high frequency noise interference, the converter board is located in a shielded part of the cooling system. In addition, the cooling of the board is more efficient at this location. The PB is shown in Fig. 5.

2.4 The Central Logic Board

The DOM Central Logic Board (CLB) is the main electronic board in the readout chain of KM3NeT. The LDVS signals generated by the PMT Bases and collected and distributed by the Octopus Boards arrive to the CLB where they are discretized by means of 1 ns resolution TDCs. The TDC data are sent onshore after being organized and timestamped by the CLB. The CLB takes care also of the readout of several instruments, as it is the case of the compass, tilt-meter and temperature sensor, all of them integrated on the same PCB, an acoustic piezo sensor, an LED nano-beacon and a hydrophone. In order to synchronize the DOMs in KM3NeT, the CLB integrates the White Rabbit protocol [5]. It provides sub-nanosecond accuracy and picoseconds precision of synchronization for large distributed systems (more than 1000 nodes over optical fibres and copper wires with a length of up to 100 km). Using the White Rabbit, we are able to achieve precise time-tag measured data and trigger data taking in large installations and, contemporarily, the same network can be used for data transmission. The key technologies used are the Synchronous Ethernet (SyncE) and the Precision Time Protocol (PTP).

The main component of the CLB is a Kintex7 FPGA [7]. This device also allows the reconfiguration of the firmware of the CLB. It is feasible to store up to four FPGA images in an SPI memory, three of them reconfigurable. The non-reconfigurable image provides a safe start for the FPGA in case of corruption of the 3 reconfigurable images, being possible to choose to boot the FPGA with any of the other three. Fig. 6 shows the CLB board containing the FPGA where the readout and communication systems are implemented.

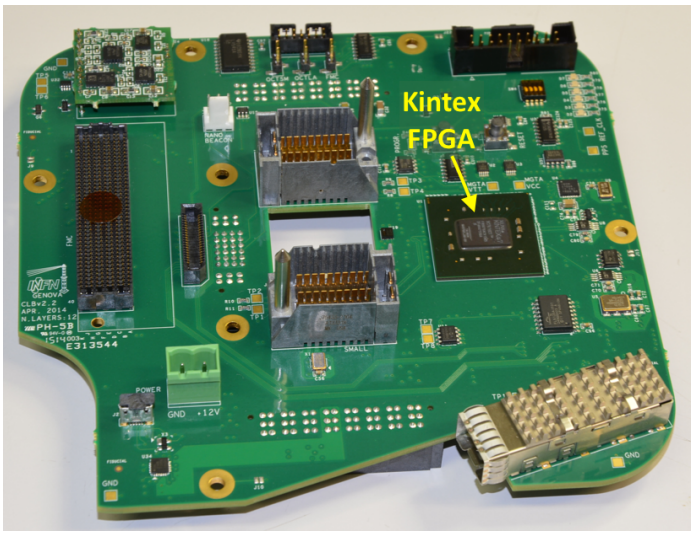

Figure 6: The Central Logic Board.

The Kintex7 FPGA is an FPGA family providing very low power consumption allowing a total CLB power consumption below 4 watts. The control of the CLB is achieved through embedded software running in an LM32 [8], an open source firmware microprocessor from Lattice. No operative system is used in order to reduce power consumption. The CLB firmware implements two LM32 CPUs. One of them resides in the WR core which is dedicated to handle the PTP traffic and controlling the Phase Locked Loops (PLL) that are part of the timing system. The other CPU takes care of the slow control communication with the shore station.

2.5 The CLB readout firmware

TDC channels were designed in a Xilinx Kintex7. Because of the programmable characteristic of an FPGA, a TDC readout implemented in this device has flexible characteristics. Here, a simplified TDC is designed to verify the idea of deserializing the raw data by means of dedicated input/output blocks of the FPGA. The CLB board includes a 25 MHz crystal oscillator, the clock signal is first transferred from a clock pin to a buffer in the centre of the FPGA, and then fanned out to the inner PLLs to provide two high frequency clocks of 250 MHz and 90° phase shifted. A 4-oversampling method increases the sampling frequency up to 1 GHz achieving the desired accuracy of 1 ns. The samples produced by deserializers are sent to specific blocks called Data Recovery Units (DRUs) where the data are reorganized and the digital pulse information is computed. The system readout generates an output of 48 bits where the 8 most significant bits are used to encode the PMT identifier, the next 32 bits indicate the time-stamp and the 8 less significant bits are used to digitize the length of the pulse.

Data are chopped in periods of time called frames or timeslices. The data-stream split into timeslices grants the scalability of the DAQ system; each timeslice is self-consistent and the probability of having events which are cut in two parts across two subsequent timeslice is negligible. Frames start at regular (programmable) intervals. Each frame is assigned a UTC time stamp and it is therefore uniquely identified. The CLB timing reference clock with a 1 ns phase precision and

the UTC time are supplied by the White Rabbit PTP Core (WRPC) [5] that is implemented in the firmware. TDC data is transmitted to shore via the endpoint of WRPC. This endpoint is basically a normal Ethernet MAC that uses the PTP system; a sub-nanosecond timing precision is achieved. The MAC is connected to an IP/UDP packet buffer stream selector (IPMUX). This IPMUX splits the data connecting to the MAC into separate streams, based on the UDP port number. In the first place, data need to be shipped from the DOM to the shore station, so the IPMUX in mainly used as an UDP transmitted. Data received from the TDCs are stored in the Front-End FIFO memories and consequently passed to the IPMUX.

More functions have been integrated into the TDC firmware or are under development, such as the high-rate veto function to avoid overloading the communication bandwidth and the multi-hit function to process pulse widths larger than 255 ns. 31 TDC channel have been implemented, but the appropriate number of channels can be chosen according to the requirements of each DOM. The current implementation also offers a wide variety of interface options like an enable interface, which allows enabling or disabling remotely the TDC channels using the embedded software based on the LM32 microprocessor.

3. The Trigger and Data Acquisition System

The ensemble of hardware infrastructures and software collections in charge of the data acquisition, aggregation, filtering and save-to-disk is called TriDAS (Trigger and Data Acquisition System). All the TriDAS elements are connected within a local 10Gbps network which comprises also the underwater detector. No hardware triggers are implemented at the underwater detector; all the measured signals are sent to the shore station where they are filtered and recorded.

The first step of the data aggregation onshore is the Data Queue process (DQ, Fig. 7), that distributes unfiltered data arriving from a sector of the detector to the computer farm that takes care of the trigger, the DataFilter (DF). There are two different types of DF: the acoustic DataFilter is responsible of the online analysis of data from acoustic sensors for the positioning system; the optical DataFilter is the software devoted to the trigger of PMT data (Fig. 7). Each optical DF receives from all the active DQs a block of data related to a specific time window (the timeslice) with a typical length of 100 ms. The optical DF handles the data of the entire apparatus with respect to the assigned timeslice. The filtered data are finally sent to a DataWriter, that writes on disk a ROOT file [9].

Each acoustic DF receives all data arriving from a DQ in a continuous stream, calculates the sound arrival time produced by the underwater positioning system and sends the result to a DataBase. The ControlUnit (CU, Fig. 7), which is the user interface of the detector, coordinates all the operations of the TriDAS and operates the DOMs using a dedicated Slow Control (SC) protocol. The number of necessary DFs scales with the detector throughput and with the algorithms complexity without changing the DAQ design. The data reduction is more than a factor 10^3 with respect to the throughput from the detector.

The optical background measured by the 3-inch PMTs, due to the environmental sources such as the ^{40}K decay and the bioluminescence, is $\sim$7 kHz and causes a large throughput from the detector. The acoustic stream from the positioning system is also relevant (from $\sim$1/4-1/2 of the

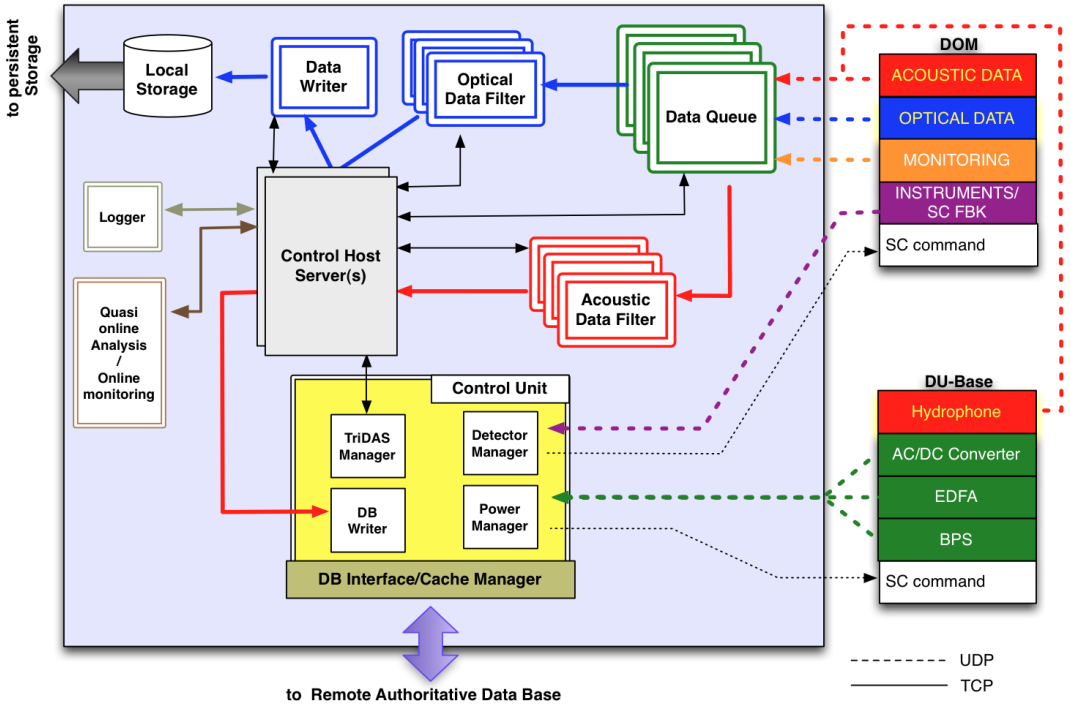

Figure 7: General architecture of the TriDAS for KM3NeT according to the all data to shore approach. The shore network infrastructure (gray filled parts) is based on the 10 Gb Ethernet.

optical stream). A ∼120 Gbps throughput is expected from the full detector composed of 690 Detection Units, assuming a photon-hit rate of 6 kHz for each 3" PMT.

4. Conclusions

The data acquisition system of KM3NeT has been presented. The electronics inside every DOM reads and digitizes the signal from the 31 3-inch PMTs. Data are processed in parallel by an FPGA which implements 31 TDCs with 1 ns resolution and the White Rabbit protocol. Data are sent onshore packed in timeslices of a fixed time duration with a unique timestamp. Onshore data are filtered and triggered by the TriDAS, and finally stored on disk for analysis. The first Detection Unit of KM3NeT that uses all the technologies described here will be installed soon offshore the French coasts near Toulon [10].

References

[1] `http://www.km3net.org`

[2] S. Adrián-Martínez et al. (KM3NeT Collaboration), *Expansion cone for the 3-inch PMTs of the KM3NeT Optical Modules*, *JINST* **8** (2013) T03006.

[3] M. Ageron et al. (ANTARES Collaboration), *ANTARES: the first undersea neutrino telescope*, *Nucl. Instr. Meth. A* **656** (2011) 11.

[4] S. Biagi and A. Orzelli (KM3NeT Collaboration), *The Central Logic Board and its auxiliary boards for the optical module of the KM3NeT detector*, *JINST* **9** (2014) C12033.

[5] http://www.ohwr.org/projects/wr-cores/wiki/wrpc_core

[6] E. Anassontzis et al. (KM3NeT Collaboration), *Design and development of the Power Converter Board within the Digital Optical Module in KM3NeT*, PoS(TIPP2014)188 (2014).

[7] http://www.xilinx.com/products/silicon-devices/fpga/kintex-7.html

[8] https://en.wikipedia.org/wiki/LatticeMico32

[9] http://root.cern.ch

[10] A. Creusot (KM3NeT Collaboration), *Calibration, performances and tests of the first detection unit of the KM3NeT neutrino telescope*, in proceedings of *34th ICRC, The Hague, The Netherlands*, PoS(ICRC2015)1154 (2015).