

The Trigger and Data Acquisition System for the KM3NeT-Italia towers

M. Favaro^{1,2,4,a}, T. Chiarusi^{1,b}, F. Giacomini^{2,c}, M. Manzali^{2,4,d}, A. Margiotta^{1,e}, and C. Pellegrino^{1,3,f}

¹ INFN BO, Bologna, Italy

² INFN CNAF, Bologna, Italy

³ Dipartimento di Fisica e Astronomia, Università degli Studi di Bologna, Bologna, Italy

⁴ Università degli Studi di Ferrara, Ferrara, Italy

Abstract. KM3NeT-Italia is an INFN project supported with Italian PON fundings for building the core of the Italian node of the KM3NeT neutrino telescope. The detector, made of 700 10'' Optical Modules (OMs) lodged along 8 vertical structures called *towers*, will be deployed starting from fall 2015 at the KM3NeT-Italia site, about 80 km off Capo Passero, Italy, 3500 m deep. The *all data to shore* approach is used to reduce the complexity of the submarine detector, demanding for an on-line trigger integrated in the data acquisition system running in the shore station, called *TriDAS*. Due to the large optical background in the sea from ⁴⁰K decays and bioluminescence, the throughput from the underwater detector can range up to 30 Gbps. This puts strong constraints on the design and performances of the TriDAS and of the related network infrastructure. In this contribution the technology behind the implementation of the TriDAS infrastructure is reviewed, focusing on the relationship between the various components and their performances. The modular design of the TriDAS, which allows for its scalability up to a larger detector than the 8-tower configuration is also discussed.

1. Introduction

The INFN's project KM3NeT-Italia [1], supported with Italian PON fundings, consists of 8 vertical structures, called Towers, instrumented with a total number of 700 Optical Modules (OMs) and will be deployed 3500 m deep in the Ionian Sea, at ~ 80 km from the Sicilian coast [2, 3]. A Tower is made of 14 horizontal bars, piled up one by one with 90° heading difference. Each bar hosts 6 OMs. Each OM contains a 10'' PMT and the readout electronics. The detection principle exploits the measurement of the Cherenkov light from relativistic particles outgoing high-energy neutrino interaction within a fiducial volume around the telescope. In order to reduce the complexity of the underwater detector, the *all data to shore* approach is assumed, demanding for a Trigger and Data Acquisition System

^a e-mail: matteo.favaro@cnafe.infn.it

^b e-mail: tommaso.chiarusi@bo.infn.it

^c e-mail: francesco.giacomini@cnafe.infn.it

^d e-mail: matteo.manzali@cnafe.infn.it

^e e-mail: annarita.margiotta@unibo.it

^f e-mail: carmelo.pellegrino@bo.infn.it

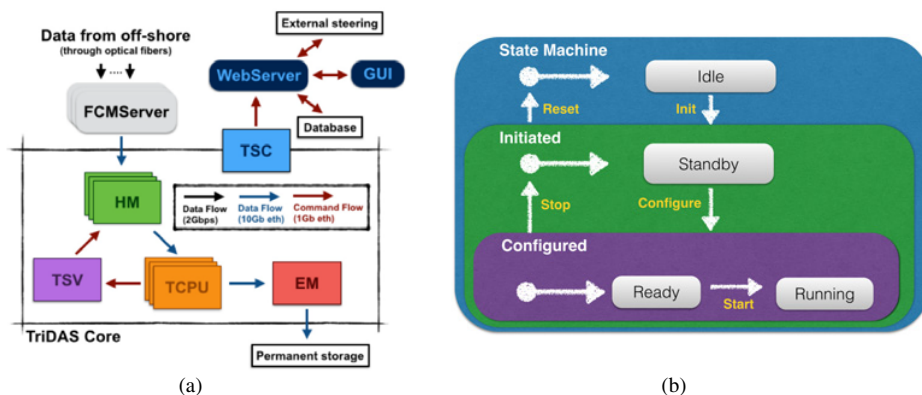


Figure 1. The Fig. 1a (left) represents the TriDAS block diagram with the various managed data flows. The Fig. 1b (right) is the finite state machine adopted for steering the TriDAS components.

(TriDAS) [4] running at the shore station. The collected data stream from all the Towers is largely affected by the optical background in the sea [5], mainly due to the ^{40}K decays and bioluminescence bursts. Ranging up to 30 Gbps, such a large throughput puts strong constraints on the required TriDAS performances and the related networking architecture. In the following sections there are the description of the final implementation of the physics-data handling (TriDAS Core), the user and management interfaces (TriDAS Control) and the large-band network infrastructure.

2. The TriDAS software

2.1 TriDAS Core

The TriDAS Core [6] (Fig. 1a) is formed by the HitManager (HM), the Trigger CPU (TCPU), the TriDAS-SuperVisor (TSV) and the Event Manager (EM). The data to the TriDAS Core are provided from the FCMServer (FCMS) units that read the optical data from the detector and send them to the HitManager. The FCMServers are designed to be, onshore, the gate for all the kind of data streams (slow control, optical and acoustic), going to and coming from the offshore detector [7]. One single FCMServer can handle the optical data coming from 4 floors of a Tower. With 8 Towers, the total number of FCMServers is 32. The FCMServers forward the data coming from the OMs to the first layer of the TriDAS, the HitManagers. Every single HitManager process runs in a dedicated server and it is linked to a fixed number of FCMServers, which correspond to a portion of the detector, called Sector. Test proved that a HitManager process is capable to withstand the data from an entire Tower. All the HitManagers share the same time line, originated from a common timestamp, which is quantized in subsequent intervals of time of the same duration, called *TimeSlices*. In this way, a full set of optical data occurred during a particular TimeSlice are asynchronously managed by all the HitManagers, which organize their own fraction of data in a special data-structures called the “*Sector Time Slices*” (STSs). The role of the TriDAS-SuperVisor is then to steer all the HitManagers sending the STSs belonging to the same TimeSlice to the first available TriggerCPU, according to a free-token-scheduler mechanism. On its turn, one TCPU collects all the STSs of a TimeSlice into the so called “*Telescope Time Slice*” structure (TTS), then process it according with the trigger algorithms [4]. Many TriggerCPUs process different TTS from different TimeSlices at the same time. The fraction of data which fulfill the trigger selection criteria is sent to the EventManager, which records the filtered data in binary files on the local storage. Offline, the written post-trigger files are transferred from the Shore station infrastructure of

Portopalo to the storage facility at LNS via a dedicated 10 Gbps connection. The design for the TriDAS is modular and scalable with the number of deployed Towers. The required number of TriggerCPUs processes depends on the complexity of the trigger algorithms and increases with the number of OMs. It is currently under study the determination of the necessary computing resources as a function of the detector dimension.

2.2 TriDAS Control

The *TriDAS Control* (TSC) is the software component that orchestrates all the TriDAS processes running on the data acquisition farm. The TriDAS Control implements a simple hierarchical state machine with four states, as shown in Fig. 1b:

Idle. This is the initial state of the overall TriDAS state machine, where no processes are running. An *init* transition, which takes a *run setup* identifier as a parameter, executes an action that retrieves a run datacard corresponding to the given run setup. The datacard describes the geometry of the detector and the configuration of the TriDAS system (such as the role of each node) for this run. If the action is successful, the state machine moves into the *Initiated* sub-state machine.

Standby. This is the initial state of the *Initiated* sub-state machine. Here the TriDAS Control is aware of the configuration of the TriDAS system but no processes are running yet. A *configure* transition executes an action that retrieves the *run number* and starts the Trigger CPU, HitManager and Event Manager processes on the corresponding nodes. If all the processes start successfully the state machine moves into the *Configured* sub-state machine.

Ready. This is the initial state of the *Configured* sub-state machine. Here the Trigger CPUs, the HitManagers and the Event Manager are ready to acquire, filter and save physics data coming from the FCMServers. The *start* transition executes an action that computes the start time of the run and starts the TriDAS-SuperVisor. The TriDAS-SuperVisor's role is to schedule which Trigger CPU process will compute a given TTS. The scheduling follows a credit-based mechanism to balance the load among the Trigger CPUs. If the TriDAS-SuperVisor starts successfully the data acquisition starts and the state machine moves into the *Running* state.

Running. In this state the data acquisition is running.

Transitions exist to move the system back to the *Idle* and *Standby* states.

If any error occurs during a transition, the transition is aborted. Depending on the severity of the error, the system may stay in the current state or even shut-down completely.

The communication with the TriDAS Control, for example to trigger the transitions described above or to query the state of the system, is stateless and happens over a UNIX socket. Only one client can use the socket at any one time.

2.3 WebServer

The *WebServer* is the only entry point for controlling the DAQ. This component provides a set of RESTful API which allows to inquire the TriDAS Control. Therefore, it provides user authentication based on hierarchical configuration, which is implemented via different privileged groups. The privileged groups present are: administration, DAQ control and monitor. A user can have a combination of privileges, belonging to different groups. The TriDAS Control has a unique way of communication as described in 2.2. The TriDAS Control is a local-single-client program. The WebServer can be contacted from several different concurrent people at a time. The system allows controlling the DAQ via an escalation procedure only one user at a time, that permits users to acquire the privilege for controlling the TriDAS Control.

The use of WebSocket allows to implement a real time feedback system to the users. The WebServer can communicate instantly feedback and alarms during an acquisition phase or escalation.

3. Conclusion

The TriDAS has been improved to sustain the foreseen 8 towers detector. Its performances and scalability are under intense test, with long duration runs and varying the incoming throughput using either real FCMServer and simulation programs, using a test bench that reproduces the real farm in Portopalo. New trigger algorithms are under development serving different kinds of physics analysis, e.g. multi-messenger external alerts, high energy neutrino induced showers and astrophysical source detection.

The preliminary test phase demonstrates that the system is stable and the users are able to control it properly. Moreover, in November 2015 the installation and functionality test of the farm in Portopalo has been completed. Extended tests of the TriDAS will be also realized in the Portopalo infrastructure, in advance with respect to the first deployment of the Towers. This will make possible tests in the real situation, exploiting the actual facility available in the shore station. In addition to that, increasing the computing resources with respect to what available in the test-bench, more realistic results will be achieved.

References

- [1] *Km3net web site*, <http://www.km3net.org/home.php> (2015), [Online; accessed 11-Dec-2015]
- [2] S. Aiello et al., *Astroparticle Physics* **66**, 1 (2015), ISSN 0927-6505, <http://www.sciencedirect.com/science/article/pii/S0927650514001960>
- [3] M.S. T. Chiarusi, *European Physical Journal* **65**(649), 3 (2010)
- [4] C. Pellegrino, et al. (Nemo), *AIP Conference Proceedings* **1630**(158), 7 (2014)
- [5] Pellegriti, M. G., et al., *AIP Conference Proceedings* **1630**(1), 94 (2014)
- [6] M. Manzali, T. Chiarusi, M. Favaro, F. Giacomini, A. Margiotta, C. Pellegrino, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (2015), ISSN 0168-9002, <http://www.sciencedirect.com/science/article/pii/S0168900215014266>
- [7] A.L. et al., *Journal of Instrumentation* **10**(04), C04011 (2015), <http://stacks.iop.org/1748-0221/10/i=04/a=C04011>