KM3NeT: a proposal design for a detection unit data transmission system based on a copper backbone

F. Ameli [a,*] for the KM3Net Collaboration [b]

[a] *Istituto Nazionale di Fisica Nucleare*
[b] *http://www.km3net.org/home.php*

Abstract

The experience gained in Antares and Nemo experiments suggests to explore new ways of realizing the data transmission at the level of the detection unit (i.e. the string or the tower). The implementation of a copper connection with simple tracts of cable between contiguous storeys could provide an easier scalability of the structure, simple maintenance of the backbone, lower costs because of cheaper connectors, safer integration and transportation. This work is aimed at the presentation of an electronic board prototype designed to test the feasibility of the project.

Key words: Neutrino telescope, Copper backbone, KM3Net
PACS: 01.30.Cc

---

1. Introduction

Lessons learnt from current experiments developing neutrino telescopes, advised us to look for new ways of implementing data transmission at the detection unit level. The main reason to avoid the use of fiber optics in the unit backbone is the high cost of cables and connectors, the integration (because fiber handling is a delicate matter), the difficulty of testing the integrated system, the power required by electro-optical transceivers and, in the case of DWDM systems, the cost of the transceiver itself. Most of these issues are easily addressed using copper links, e.g. handling is much easier, connectors are more common and one order of magnitude cheaper than the optical counterpart, the devices which implement the physical layer of the transmission system are inexpensive. On the other side, the use of copper links requires a sophisticated implementation of the transmission system.

Recent technological progresses, specifically in the audio/video data transmission, made available on the market suitable devices capable to transmit at high data rate (about Gb/s) at long distances (hundreds of meters, depending on the medium). The idea behind this paper is to use such an established technology for the purposes of a submarine experiment.

The basic unit of neutrino telescopes can be defined as the Detection Unit (DU). Each DU consists of a set of sensors (photomultipliers or PMTs) usually grouped in 3 to 6 elements; each set is mechanically placed in a "storey". The total number of sensors per DU ranges from 60 to 120, distributed vertically in tens of storeys in structure as high as 800 m. One of the major effort in the development of this project is to suite the maximum number of possible apparatus topologies, in order to be independent from mechanical choices.

The DU is far from the shore laboratory, typi-

---

[*] Corresponding author. Address: INFN Sez. Roma1. Tel.: +390649913467
Email address: `fabrizio.ameli@roma1.infn.it` (for the KM3Net Collaboration).

cally it is placed at distances of many tens of kilometers, and produces an amount of data on the order of 1 Gb/s; these constraints require a data transmission system based on fiber optics. One or more storeys in the DU will collect all the data produced by a group of sensors providing the link with on-shore. The data transmission inside the DU can be implemented with optical fibers. This solution has already been successfully tested in the Antares and NEMO experiments [1] [2] [3], where DWDM backbones based on Add and Drop modules have been realized.

Because of short distances (maximum 50 m) between storeys and low data rates per storey, we could try a different approach: we could use copper wires to implement the backbone transmission system for the detection unit.

In the next paragraph, 2, the requirements of a copper system will be discussed and a set of possible implementations are presented. In section 3 an implementation of the electronic unit, which the copper transmission is based on, is shown. Finally, in section 4 the feasibility and impact of this design is evaluated.

2. Copper link highlights

The main constraints taken into consideration in the design of a copper backbone are:
- storeys have a distance of maximum 50 m;
- a reasonable data rate on the link is about 1 Gb/s;
- the cable is mechanically manageable;
- storeys are connected by single tracts of cables.

The data rate over the backbone is highly asymmetrical: data flowing from on-shore to the storeys are for control and setup, i.e. the data rate can be as low as few Mb/s per storey. On the opposite direction the data rate must support the transmission of physics data, about 10 Mb/s per 10” PMT. For both direction , we assume the best (ad only choice, if minimization of cable number is required) is to transmit a synchronous bitstream which embeds both clock and data. This choice guarantees the minimum number of logical links. The low speed of the control channel allows the receiver to recover the clock signal guaranteeing the necessary sub-nanosecond precision required by this kind of experiments.

The requirement of a manageable submarine cable implies a low number of wires inside it; moreover, that means the storeys should be connected with independent tracts, without signals extraction from the backbone by means of breakouts, and the connectors can be realized with a low number of electrical pins.

In the next subsections, three different solutions for copper backbone links are proposed.

2.1. *“Waterfall” implementation scheme*

In this scheme the slow control data stream is fed only to the highest storey in the DU, as shown in Fig. 1; this node recovers the clock from the low speed link and uses it to drive the high speed chain which “falls” down toward the DU base. The data rate are asymmetrical: for the up stream 10 Mb/s and for the down stream 1.25 Gb/s.

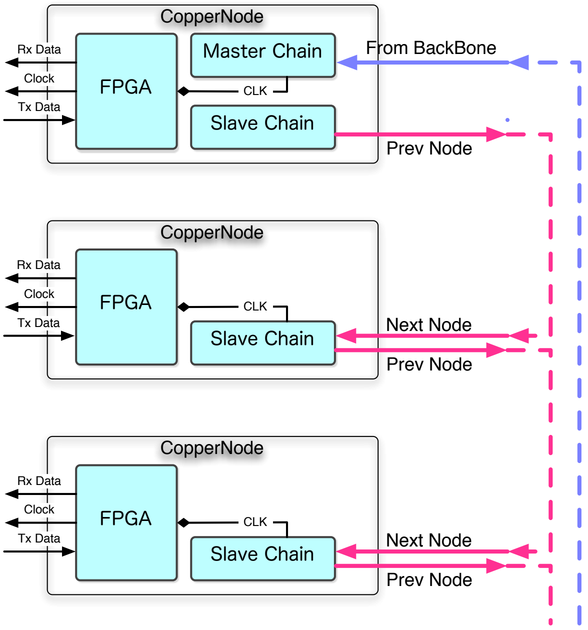

Figure 1. Block diagram of the “Waterfall” scheme.

The high speed link is implemented by means of a daisy chain topology: each node (but the highest) receives a stream from the previous one, extracts the clock, add its own payload to the received data and transmits all the data again to the next node. Each nodes manages the whole link data rate (about 1 Gb/s).

The advantages of this design are the simplicity of the up-link which reaches only one node; all the control data are then retransmitted using the down going stream, reaching each node in the chain. Some power can be saved because only one node needs the slow speed interface. Some problems arise from the need of recovering the clock using the high speed link, which could be more problematic than extracting it from a low speed stream. Moreover, in case of failure either of the control link or of the highest node, all the detection unit would be lost. Finally the communication protocol is quite complicated by the necessity of using the down going data transmission link to transport both physics data and control data. The scheme is not straightforwardly scalable, because of the limitation of the distance between the highest floor and the control link source.

2.2. *Multi-Tap implementation scheme*

The scheme in Fig. 2 is similar to the one described in the previous section with the difference that the slow control channel is received by each node independently: for this reason it is called "multi-tap". The up-going link runs at a low speed, say 20 Mb/s, from which the clock is derived and provided to the storey. The down-going link uses the recovered clock to synchronize the transmission at high speed.

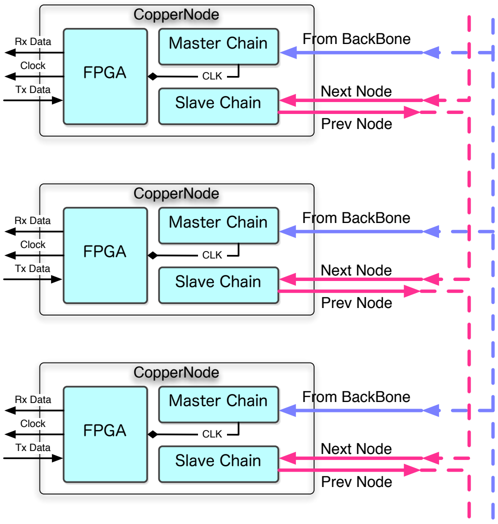

Figure 2. Block diagram of the Multi-Tap scheme.

The control link, in this case, is quite simple and guarantees good clock quality; the protocol must be conceived for a broadcast channel shared between all the nodes. The dynamic allocation of data bandwidth is easily managed by the control data source through a simple communication protocol. The failure of one node can be non-destructive, if the node itself can be properly sectioned out from the backbone. On the other hand, signal extraction from the backbone could be a challenging task: the pick-up should be as non-intrusive as possible to avoid signal deterioration which would worsen clock (and data) quality for upper floors. In this sense, this scheme could be not scalable because of the limited number of nodes allowed.

2.3. *Daisy Chain implementation scheme*

The last implementation scheme, shown in Fig. 3, uses a full daisy chain mechanism to transmit both the up-going and the down-going stream: the only difference is in the speed, about 200 Mb/s for the up-going and about 1 Gb/s for the down-going. Once again the control link is as low as desired to allow the highest quality of the recovered clock. All the nodes have the same structure, but the first and the last in the chain have a different setup with respect to clock managing. Each node handles all the data flowing through the link, also data regarding the other nodes, and "adds and drops" the payload it is interested in.

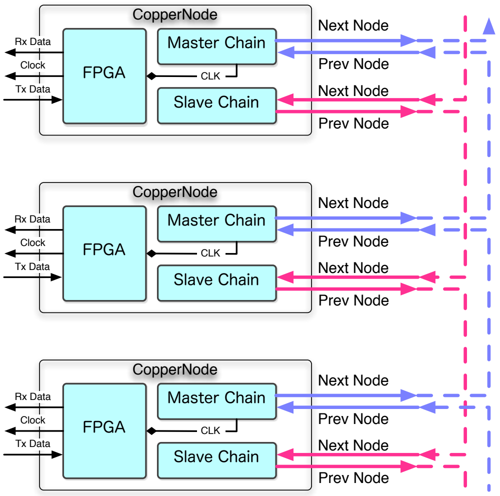

Figure 3. Block diagram of the full Daisy Chain scheme.

Obviously this scheme implies higher power with respect to the previous two, because the node has to implement the two daisy chain interfaces. Adjacent storeys are connected through simple cable tracts; adding another storey means adding another tracts[1] This is again a "single point failure" scheme, because the fault of one node breaks all the

---

[1] Remembering that the first and the last node in the chain must have a different configuration regarding clock use, the best way to add nodes is insertion in the middle of the already built chain.

chain. This argument will be briefly discussed in section 2.4.

2.4. *Reliability issues*

As already stated in previous sections, using a daisy chain mechanism implies a single point failure of the whole DU. To avoid this risk, the first approach is to increase the number of backbones, so that in case of failure only a part of the DU is lost. If we have more than one backbone per DU, the number of PMTs per backbone decreases and the required bandwidth per daisy chain link can be reduced. As a consequence, it is possible to reach longer distances, i.e. the faulty node could be sectioned out. In this way a single fault can be considered recoverable. Moreover, the increased complexity due to the multiple backbones is as expensive (under all points of view, i.e. power, money, complexity) as having a single backbone serving all the PMTs on the DU.

The sectioning of the broken node can be implemented relying on the independent communication link used to manage storey power: in fact, it seems reasonable that the two functions, data taking and apparatus monitoring on one side and power management with some slow control functions related to power on the other side, are separated conceptually and practically. Then the power communication line could be instructed to switch off faulty nodes and to bypass them.

At the DU junction box, a BackBones Concentrator (BBC) must be designed. The task of the BBC is bridging the long haul optical fiber, connecting the DU to the optical network, and the multiple copper backbones in the DU. The interface toward the copper backbones is exactly the daisy chain source node described in the previous implementation schemes: it is labeled as “source” node because the clock of the slow control link is originated there. Hence the BBC acts as a multiplexer-demultiplexer of the backbones data.

Considering current experience, the making of a 2.5 Gb/s link using DWDM commercial electro-optical transceiver and serializer-deserializer (SerDes) is a well-established technology at reasonable cost and power. We will not delve into the details of this design in this paper

3. Copper Node

In this section a brief description of the copper node is given. Fig. 4 shows the block diagram of such a node: at the heart of the system an FPGA is foreseen; the two chains, the *Master* and the *Slave*, are functionally identical but for the use of the clock: in fact the Master Chain recovers the clock from the control up-going stream, cleans it and provides the timing to the whole node including to the Slave Chain. The node is designed as a pluggable mezzanine board which can be hosted by a user board. Exploiting this concept, it is useful to let the mezzanine manage all the burden due to daisy chain handling. The user interface, provided by the FPGA, consists of two parallel buses to transmit and receive data, clock and control signals and, eventually, some slow control interface for asking the status of the node itself.

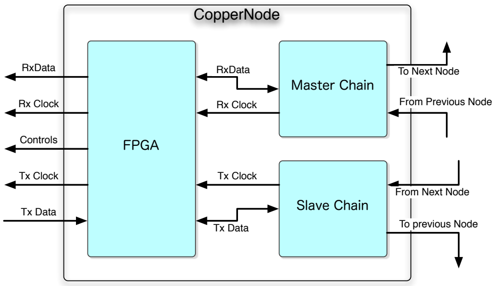

Figure 4. Block diagram of the copper node.

The daisy-chain mechanism is completely transparent for the user, whose interface to the backbone is a simple SerDes-like interface. There is no user intervention in the management of the physical layer. Moreover, the BBC can insert in the control data some fields intended for the node itself for some maintenance tasks like on-the-fly reprogramming, dynamic bandwidth allocation of the whole payload, speed selection of the interface, and more.

In Fig. 5 the block diagram of the Master Chain is shown. The received signal is properly equalized to allow stream recovery; then the clock is extracted by a Clock and Data Recovery (CDR) device and jitter cleaned by a Phase Locked Loop (PLL) device. Such recovered clock is provided to the serializer, to the rest of the electronics in the copper node and also to the host board. Data are transmitted using a driver which mates the equalizer on the receiver side.

The clock for the high speed transmission link implemented in the Slave Chain block can be obtained

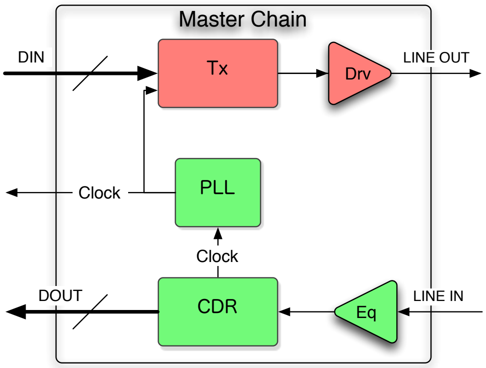

Figure 5. Block diagram of the copper node.

multiplying the clock recovered by the Master Chain block, so to have a fully synchronous system in both directions.

4. Conclusions

The proposal of a coper backbone seems to be feasible thanks to the great evolution of commercial data transmission systems on copper. This architecture could improve some aspects of detection unit testing and integration, cost (because optical fibers could be avoided at the DU level), manageability of assembled system. The need of a single DWDM electro-optical transceiver per DU represents a big saving in money. Moreover, all the data of a DU could be transported by a single pair of fibers, making the overall requirement of colors, i.e. fibers, less demanding than current projects. The connection between storeys can be implemented using small tracts of cable containing only copper wires, allowing the use of inexpensive submarine connectors.

In the design of the copper system we are not taking into consideration the effect of high pressure: this item is under study and the electronic board is designed having in mind that an impedance matching will be necessary to compensate the change of the characteristic impedance due to variations of cable mechanical properties. First goal is to demonstrate that the system is able to work even after many nodes put in cascade.

A procedure for assessing the timing relationship between node is currently under study and will be tested when the first prototypes will be available.

References

[1] Nemo website.
URL `http://nemoweb.lns.infn.it/`

[2] Antares website.
URL `http://antares.in2p3.fr/`

[3] F. Ameli, M. Bonori, C. A. Nicolau, Synchronous data transmission protocol for NEMO experiment, in: E. de Wolf (Ed.), VLVνT Amsterdam, NIKHEF, Amsterdam, 2004, pp. 196–200.