ITALIAN SUMMER PROGRAM AT FERMILAB

UNIVERSITA' DEGLI STUDI DI PISA

Analysis and design of HW/SW electronic system for Mu2e data acquisition.

Student: Marco Ferrini

Supervisors: Ryan Rivera, Micol Rigatti

October 7, 2022

CONTENTS

1 ABSTRACT

The purpose of the Mu2e experiment at Fermi National Accelerator Laboratory (Fermilab) is the research for the neutrino-less coherent conversion of the muon into an electron, in the field of an aluminum nucleus. The observation of this physics process would unambiguously demonstrate the existence of physics beyond the Standard Model. Although in the past there's always been a huge amount of experimental research activity, all the previous research for this process have given null results. The experimental technique employed by Mu2e experiment has been designed to improve the sensitivity by four orders of magnitude with respect to similar experiments.

Mu2e is a complex experimental apparatus composed by a high intensity pulsed muon beamline and several independent particles detectors; including a straw-tracker and a crystal-based electromagnetic calorimeter. The calorimeter has been designed and will be built by the collaboration among the Italian National Institute of Nuclear Physics (INFN), the California Institute of Technology (Caltech) and the Fermi National Accelerator Laboratory (FNAL or Fermilab).

The Mu2e Tracker will precisely measure momentum of charged particles that traverse it. This is critical to distinguish the well-known momentum of the signal electrons from background particles that have different momenta. The momentum measurement can be made because a charged particle will trace a helical path through the uniform magnetic field of the Detector Solenoid and the radius of this helix is directly proportional to its momentum. The Mu2e calorimeter is vitally important in reducing backgrounds. Its primary purpose is to provide a set of measurements that complement the information from the Tracker and enable us to reject backgrounds due to reconstruction errors and cosmic ray interactions not vetoed by the cosmic ray veto. For real tracks, activity in the Tracker and in the calorimeter will be correlated in time. The combination of these two-timing measurements provides a time-of-flight system that could be capable of providing particle identification information.

Both Tracker and calorimeter are challenging detector, operating in a hostile environment of 1 T magnetic field, a harsh radiation level and 10^-4 Torr vacuum. Moreover, the detector will be installed inside an evacuated cryostat and will be accessible for maintenance only for an extremely limited number of weeks per year.

Operation in vacuum has an important impact on the detector design: all the components have to be vacuum-compatible, and a dedicated cooling system is necessary to maintain electronic components within a range of temperature suited for safe long-term operation. Electronics needs to take into account also the expected high radiation levels, using radiation-hard components.

2 THE EXPERIMENT

The Mu2e experiment at Fermilab will be 10,000 times more sensitive than previous experiments looking for muon-to-electron conversion [1]. This precise and complex experimental apparatus will produce 200 million billion muons per year. The accelerator complex repurposes elements of the infrastructure that produced high energy proton and anti-protons beams for the Tevatron experiments to produce the high-intensity muon beams necessary for Mu2e and the Muon (g-2) experiments.

The Fermilab Booster will accelerate protons to the 8 GeV needed to produce the intense muon beam employed by Mu2e.

The protons will travel from the Booster to the Recycler where they will be stacked, bunched, and extracted to the Delivery ring. The Delivery ring is located in the repurposed Debuncher. Once in the Delivery ring the protons will be slow extracted and delivered to the Mu2e apparatus. A system of three superconducting solenoidal magnets (Production Solenoid, Transport Solenoid and Detector Solenoid) will transport the intense low-energy muon beam to the experimental area where Mu2e is located.

The 8 GeV protons will arrive in bunches from the Delivery ring and enter the Mu2e Production Solenoid at a slight angle to its axis and strike a tungsten production target about the size of a pencil. These collisions will create a cascade of particles, including pions that decay into muons. The magnetic field of the Production Solenoid will capture some of the muons and spiral them into the Transport Solenoid. Only about 1 in 300 protons that collide with the production target will generate a muon that moves into the Transport Solenoid. Throughout the experiment's projected three-year running period, roughly 10 billion muons per second will be stopped.

Muons in the Transport Solenoid will travel inside an evacuated vessel towards the Mu2e aluminum target. The Mu2e detector is the particle physics detector embedded inside the Detector Solenoid that provides a magnetic field in the detector region that allows the momentum of the conversion electrons to be accurately determined. The Mu2e detectors consist of two main parts: the magnetic spectrometer to measure particles momentum, and the electromagnetic calorimeter to measure particles energy and time of impact.

Improvements to the accelerator could extend the initial Mu2e sensitivity by a factor of ten or more. This would provide a valuable tool for physics research whether or not Mu2e discovers muon-to-electron conversion during its first, lower-intensity phase. If Mu2e does observe charged lepton conversion, an upgraded accelerator would enable Mu2e to study in depth the details of the conversion by providing more data. If Mu2e does not observe the conversion, the collaboration could continue the search with a wider net and also search for signs of never-before-seen physics in rare processes that have previously been out of reach of physics machines.

3 THE MU2E DETECTORS

The most relevant Mu2e detectors are the straw tracker, the electromagnetic calorimeter and the cosmic ray veto. The Mu2e straw tracker has been designed to perform an accurate measurement of charged particles momentum. This measurement is critical to distinguish signal conversion electrons, that have the momentum equal to the muon rest mass, from background electrons produced in Michel decays, that have a broad momentum spectrum with the muon rest mass as endpoint. Charged particles momentum can be measured because they follow helical trajectories through the uniform magnetic field of the Detector Solenoid and the helix radius is directly proportional to the particle momentum. The tracker has been designed to intercept the helical path in many points and allow an accurate trajectory reconstruction.

The Mu2e calorimeter provides information that complements the particles trajectory reconstruction performed by the tracker and helps to reject backgrounds due to reconstruction errors and cosmic ray interactions undetected by the cosmic ray veto. Although the calorimeter energy resolution is not competitive with the tracker, even a coarse confirmation of the particle energy helps to reject backgrounds due to spurious combinations of hits generated by low energy particles. Moreover, for real particles traversing the entire detector, the activities in the tracker and calorimeter are correlated in time. This timing information provides a time-of-flight measurement extremely useful to reduce the background level in the detector.

3.1 THE TRACKER

The tracker [2] must accurately and efficiently identify and measure 105 MeV/c electrons while rejecting backgrounds and it must provide this functionality in a relatively unique environment. The tracker resides in the warm bore of a superconducting solenoid providing a uniform magnetic field of 1 Tesla; the bore is evacuated to 10^-4 Torr. A key feature of Mu2e is the use of a pulsed beam that allows to reject prompt backgrounds by looking only at tracks that arrive several hundred nanoseconds after the proton pulse. The tracker must survive a large flux of particles during the early burst of “beam ash" particles that result from the proton pulse striking the production target, but it does not need to take data during this time. The Mu2e signal window is defined as 700 < t <1695 ns, where t = 0 is the arrival of the peak of the beam pulse at the stopping target. However, in order to study backgrounds such as radiative pion capture, the tracker must be fully efficient during the interval 500 < t < 1700 ns; we take this as the tracker's live window. To calibrate using positrons from $\pi^+ \rightarrow \nu\, e^+$ decays the tracker must also be able to collect data, during special runs with reduced beam intensity, for 300 < t < 1700 ns.

The Detector Solenoid provides a uniform 1 Tesla field in the region occupied by the tracker. To have good acceptance for signal electrons without being overwhelmed by DIO electrons (including electrons scattered into the active region), the active area of the tracker extends from about 40 < r < 70 cm (where the radius r is measured from center of the muon beam). Mechanical support, readout electronics, etc. are to be placed at r > 70 cm, out of the way of both signal and DIO electrons. These dimensions depend on the size and geometry of the muon stopping target, the size of the muon beam, and the magnetic field properties; they have been optimized to maximize the acceptance to conversion electrons while minimizing the number of low energy electrons that intersect the tracker. The momentum resolution requirement is based on background rejection: the signal is sharply peaked, whereas backgrounds are broad (cosmic rays, radiative pion capture) or steeply falling (DIO electrons). For a Gaussian error distribution, the requirement is that σ < 180 keV/c. This is simply a convenient reference point; the actual resolution is not Gaussian and may be

asymmetric. Furthermore, scattering and straggling in material upstream of the tracker are significant contributors to the final resolution.

3.1.1 Font End Electronics

To minimize penetrations through the cryostat, digitizers and zero-suppression logic are located directly on the detector [3]. Signals from the straws need to be amplified, digitized and trasmitted to the TDAQ. Front end electronics is defined as all electronics residing on the detector. Digitization is performed in the front end, then, digital data is transmitted to the TDAQ over optical fibre. The bulk of the electronics is dedicated to amplifying, digitizing, and transmitting signals from the ~ 25 K straws, but in addition there are sensors for monitoring detector parameters; data from these must also be transmitted to the TDAQ system. Each straw is read out from both ends.

The front end electronics is divided in the following categories:

- Infrastructure: includes power and cooling.
- Preamplifiers: there is one at the end of each straw to amplify and send an analog signal via a PCB transmission line to digitizers
- Digitizers: the digitizer receives and digitizes the signals from both ends of each straw and transmits that data to the Readout Controllers. The digitized signals for each “hit" consist of two-timing measurements, one at each straw end, and one amplitude.
- Readout Controller (ROC): interposes a link between the TDAQ and all other FEE components.

The time difference between the two ends is used as a measure of the hit position along the straw. The average time is used in the conventional fashion to measure the drift distance.

The communication between the digitizer and ROC proceeds via LVDS signals (Low-Voltage Differential Signalling).

3.2 THE CALORIMETER

The Mu2e detectors have been designed to reject backgrounds to a level consistent with a single event sensitivity for the $\mu N \rightarrow eN$ coherent conversion of the order of 10^-17. The electromagnetic calorimeter [4] is a vital link in the chain of background defences. Since the quality of track reconstruction is fundamental for background rejection, a particular concern is due to false tracks arising from pattern recognition errors, due to high rate hits on the detector. These errors are frequently due to accidental noise and lower energy electrons' hits. They may create a trajectory similar to a higher energy electron and mimic the muon conversion signal. One of the primary functions of the calorimeter is to provide a redundant set of measurements to complement the tracker data and minimize the level of backgrounds.

The electrons produced in muons' decay follow helical trajectories in the solenoidal magnetic field. Then, they hit the front face of the calorimeter crystals with a maximum energy in the 100 MeV range. In this energy regime a total absorption calorimeter employing a homogeneous continuous medium is required to meet the Mu2e energy and satisfy the time resolution requirements. The sensitive material could either be a liquid, such as xenon (Xe), or a scintillating crystal.

The Mu2e collaboration chose the scintillating crystals technology. Several types of crystals were considered, including barium fluoride (BaF2) and cesium iodide (CsI). The baseline design uses an array of relatively cheap undoped CsI crystals, arranged in two annular disks. Every crystal is read out by two large-area solid-state photodetectors (SiPMs), preferred to standard photomultipliers because of the high magnetic field employed by Mu2e.

While the Front End Electronics (FEE), servicing the SiPMs is mounted on the rear side of every disk, the voltage distribution, slow control and data acquisition boards are hosted in 10 crates mounted on the external lateral surface.

A laser flasher system provides light to each crystal through a network of optical fibers, for relative calibration and monitoring purposes. A circulating radioactive liquid source system, housed in the front plate of the disks, provides absolute calibration and allows to determine the absolute energy scale.

3.2.1 Calorimeter electronics

The entire calorimeter electronics can be divided in two subsystems with different functions and locations: the front-end electronics (FEE) and the digital acquisition electronics (DAQ). The first subsystem is composed by the SiPMs and the front-end boards. It is placed in the backplate, facing the rear side of the crystals. The second subsystem is composed by the data acquisition boards, which perform the digitization of the analog signals received from the front-end boards. They also provide power and monitor the front-end electronics status.

3.2.1.1 The Front End Electronics

The front end unit of the calorimeter is composed by the parts (Figure 3.1):

- one CsI crystal: the electrons impinging on the crystal's frontal surface penetrate the material and generate an electromagnetic shower. The photons produced by scintillation diffuse through the crystal volume and get transported to the rear side of the crystal.
- two SiPMs: they face the rear side of the crystal and convert the light produced by the scintillation into an electric signal.
- two front end boards: electrically connected to both the SiPMs. They provide power to the SiPMs and amplify their signals. The amplification has to be performed in the immediate proximity of the weak signal source, otherwise would be lost in electric noise.
- one mechanical support: made of copper to support the SiPMs and the FEE boards. It also acts as a “bridge" for heat transfer.

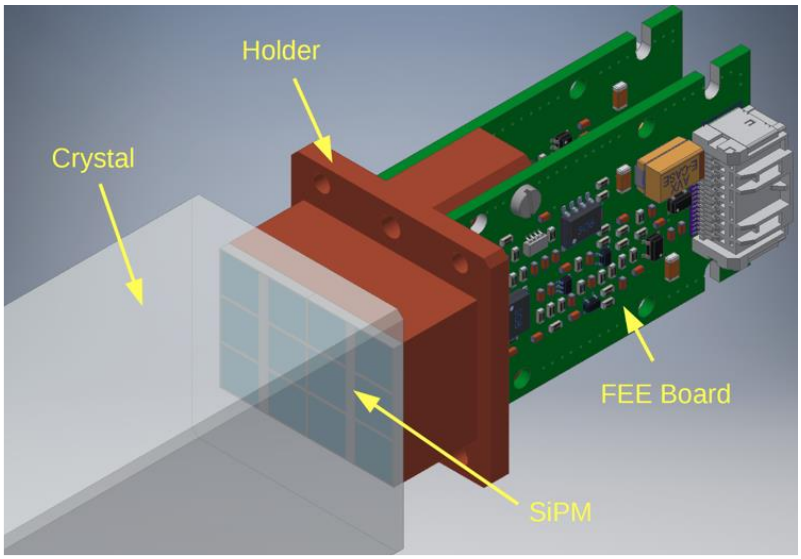

Figure 3.1: CAD model of one front-end unit

The front end electronics for the calorimeter readout consists of two discrete and independent chips (Amp-HV) placed on one unique front end board electrically connected to the back of the photo-sensor pins. The chips provide both the amplification and the local linear regulation to the photo-sensor bias voltage.

Groups of 16 Amp-HV chips are controlled by one dedicated ARM controller placed on one interface board located in the DAQ crate. This board distributes low voltage and high voltage reference values, sets and reads back the locally regulated voltages. The Amp-HV is a multilayer double-sided discrete component board that performs out the two tasks of amplifying the signal and providing a locally regulated bias voltage, significantly reducing the noise loop-area.

3.2.1.2 DAQ electronics

The analog signals produced by the front end electronics are transmitted to the data acquisition boards, hosted in the DAQ crates. Since the main function of the data acquisition boards is to digitize and transmit the analog signals to the global Mu2e data acquisition, these boards are named waveform digitizers or DIRAC (Digitizer ReAdout Controller). Additional boards are necessary to provide and distribute power to the front end boards, monitor photosensors and front-end electronics performance. These boards are called interface boards. In the current design, there are 10 DAQ crates per disk, and each crate hosts 8 DIRAC and 8 interface boards (coupled together to form 8 “bigger boards"). Every crate also hosts one further board to provide the clock distribution.

4 TDAQ – TRIGGER AND DATA ACQUISITION

The Mu2e Trigger and Data Acquisition (TDAQ) subsystem [2] provides necessary components to collect digitized data from the Tracker, Calorimeter, Cosmic Ray Veto and Beam Monitoring systems (Stopping Target Monitor and Extinction Monitor) and deliver that data to the online and offline processors for further analysis. It is also responsible for detector synchronization, control, monitoring, and operator interfaces. The Mu2e TDAQ is based on a “streaming" readout. This means that Tracker and Calorimeter detector data is digitized, zero-suppressed in front-end electronics, and then transmitted off the detector to the TDAQ system. While this approach results in a higher off-detector data rate, it also provides greater flexibility in data analysis and filtering, as well as a simplified architecture. The Mu2e TDAQ architecture is further simplified by the integration of all off-detector components in a “TDAQ Server" which functions as a centralized controller, data collector and data processor. A single TDAQ Server can be used as a complete standalone data acquisition/processing system or multiple TDAQ Servers can be connected together to form a highly scalable system.

The TDAQ monitors, selects, and validates physics and calibration data from the Mu2e detector for final stewardship by the offline computing systems. The TDAQ combines information from about 450 detector data sources and applies filters to reduce the average data volume by a factor of at least 100 before it can be transferred to offline storage. The TDAQ also provides a timing and control network for precise synchronization and control of the data sources and readout, along with a Detector Control System (DCS) for operational control and monitoring of all Mu2e subsystems. Figure 4.1 shows a full view with focus on the interfaces connected to TDAQ: Tracker and Calorimeter ROCs, Detector Hall, WH Control Room.

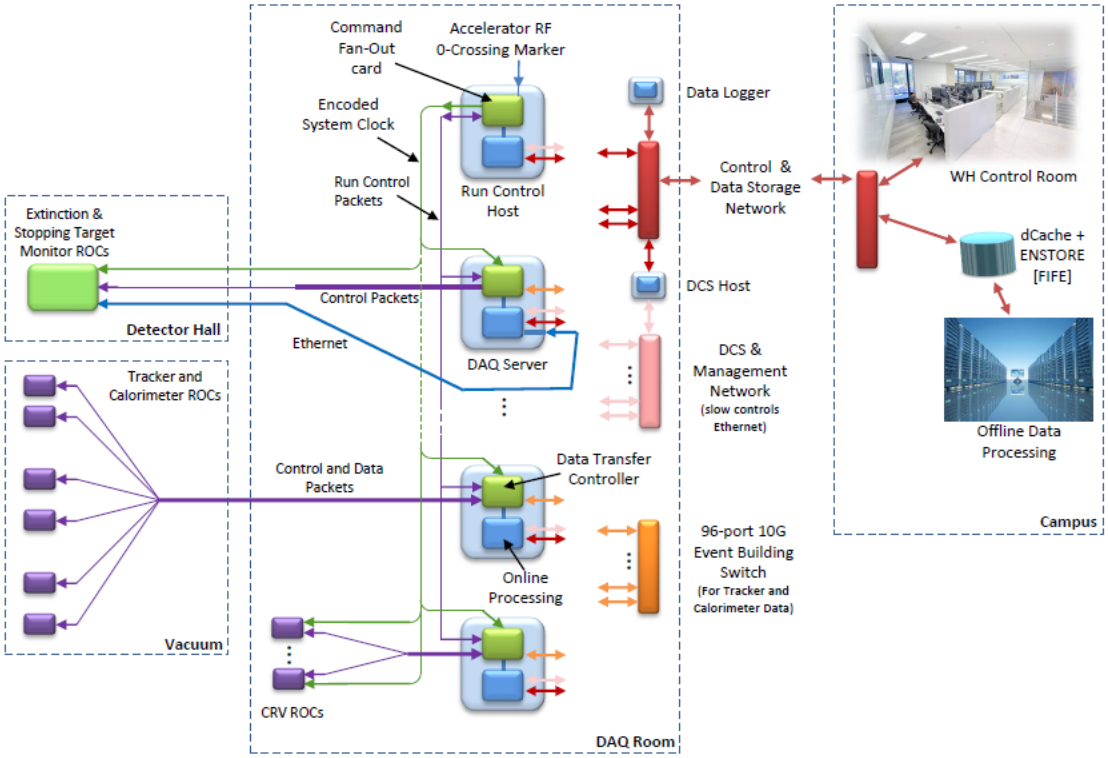

Figure 4.1: Full view of TDAQ and its interface with subsystems

4.1 ARCHITECTURE

Readout Controllers digitize and zero-suppress data at the detector. The data is then transmitted over optical links to TDAQ Servers in the surface level electronics room. Control information is sent from the TDAQ Servers to the Readout Controllers over the same bidirectional optical links. Data is exchanged between TDAQ Servers (via the Event Building Network) to form complete events.

The TDAQ Servers filter these events and forward a small subset of them to offline storage.

4.1.1 Readout Controllers

Readout Controllers (ROCs) are not part of the TDAQ system, but rather are included separately in each detector subsystem. The number of Readout Controllers and the estimated data rate for each subdetector are listed in the following table 4.1:

DETECTOR	NUMBER OF ROCS	AVERAGE RATE PER ROCS MB/S	TOTAL DATA RATE GB/S	NUMBER OF OPTICAL LINKS	NUMBER OF TDAQ SERVER
TRACKER	216	83	18	108	18
CALORIMETER	192	42	8	72	12
CRV	14	214	3	14	3

Table 4.1

Readout Controllers have the main purpose of data collection, buffer management and processing. They are based on an FPGA architecture. This FPGA provides the high-speed serial transceivers (SERDES) for the optical links and manages all kind of communications: both data transfer then Detector Control System (DCS) “slow control" operations.

Readout Controllers in or near the detector will be exposed to a high neutron flux. SRAM based FPGAs are sensitive to radiation induced single-event upset (SEU) in the configuration and application memory. Mu2e Readout Controllers in higher radiation areas will use Microsemi PolarFire series FPGAs [5] which provide on-chip microcontroller and SERDES and a number of features to mitigate SEU, including ash based configuration and ECC protected memory and registers. Commercial integrated circuits can typically tolerate total dose of at least 100 Gy without significant degradation. In the region where the Tracker and Calorimeter ROCs are located, total dose is estimated at 10 Gy/yr.

A block diagram for a digitizer/ROC is shown in figure 4.1. As a basic description, ROCs receive and phase aligns as necessary a “punched" or “encoded" System Clock with frequency +/- 10% of 40 MHz. A clock generator multiplies the recovered System Clock to drive the digitizer sample clocks (typically 50-100 MHz). The marker encoded on the System Clock is the time-zero reference for an event window (i.e. 1695 ns uBunch during beam ON). A local timestamp counter (driven by the sample clock) measures the time offset within the event window the timestamp data. The microcontroller and FPGA interface logic operates from a local oscillator, independent of the System Clock. The ROC receives a Heartbeat Request packet for each event window. This packet contains event window readout control information, along with the System

Timestamp. Data from the digitizers is zero-suppressed, formatted and written to the ROC Data Buffer during the beam spill. Data packets are read from the Data Buffer and transmitted on the optical link during the full accelerator supercycle. The buffer is large enough to hold at least 1 second of ROC output data, and uses ECC memory for SEU mitigation.

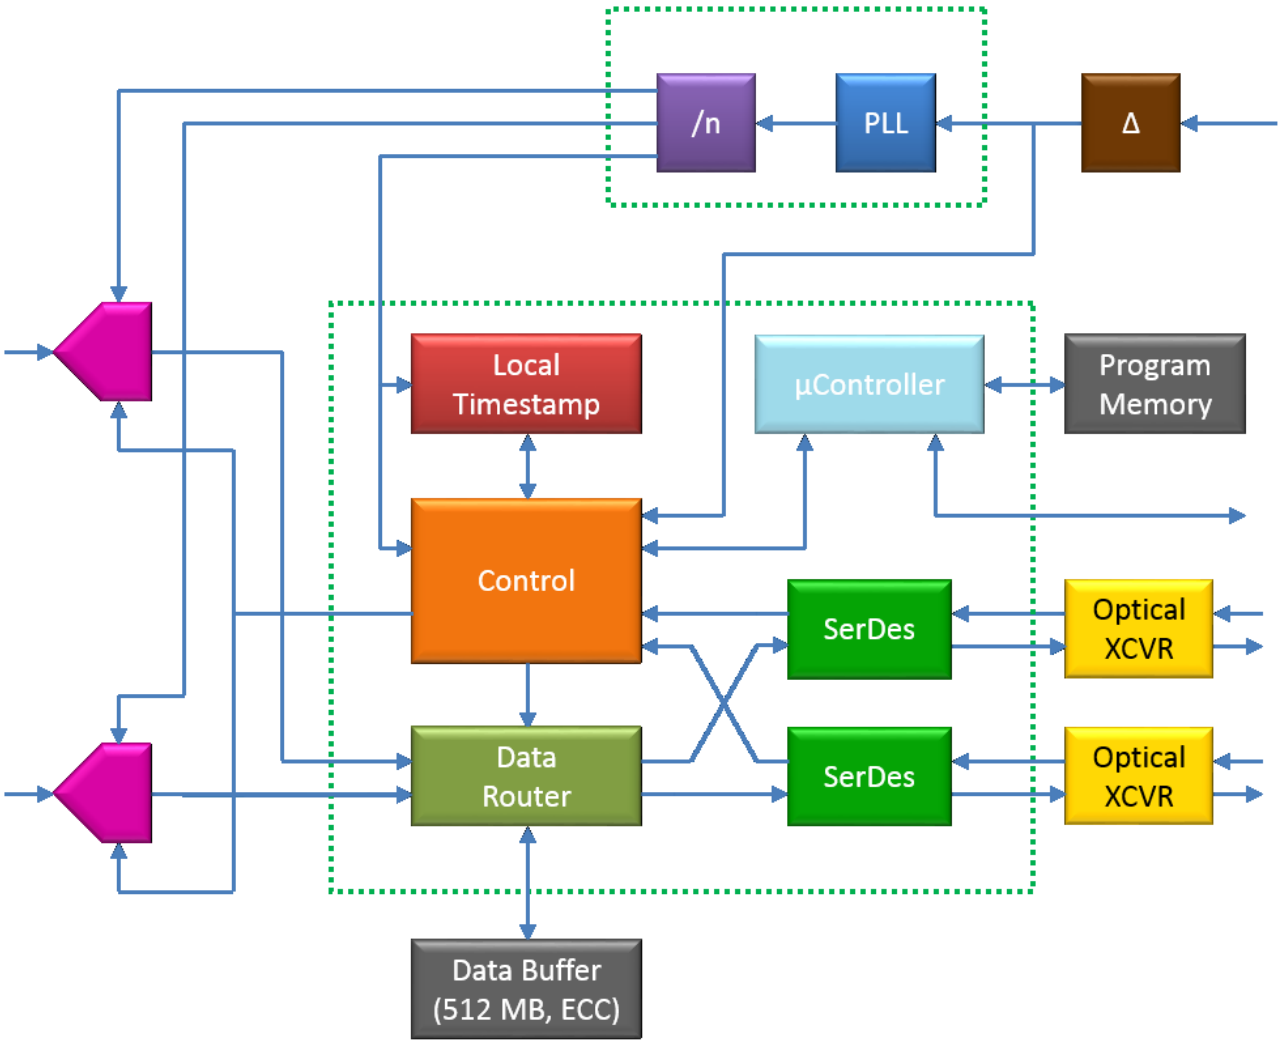

Figure 4.1: Basic Digitizer/Readout Controller Architecture

4.1.2 Data Transfer Controller

The Mu2e Data Transfer Controller (DTC) collects data from multiple detector Readout Controllers, optionally performing event building and data pre- processing. The DTC module provides an interface between the Mu2e Readout Controller (ROC) modules, and the Trigger and Data Acquisition (TDAQ) servers running the TDAQ online software framework. For Mu2e, the DTC is implemented using a commercial PCIe (Peripheral Component Interconnect Express) card located in the TDAQ Server. It is based on the HiTech Global Kintex-7 (HTG-K700) PCI Express expansion card. This card features an eight lane Gen 2 PCI Express interface, a DDR3 SODIMM socket, and a 400 pin FMC connector, all wired to a Xilinx K325T Kintex-7 FPGA. The FMC connector allows the installation of an FMC card with the optical fiber interface. This

provides for multi-gigabit serial links for up to six ROC Links, a port for data exchange for hardware event building, and a port for the Command Fan-Out (CFO) interface. Firmware for the DTC's FPGA is based on a modified reference design provided by Xilinx.

High-speed serial ports are provided by an adapter module which plugs into the FMC (FPGA Mezzanine Card) connector on the PCIe card. This adapter has eight bidirectional SFP+ (enhanced Small Form-factor Pluggable) ports and can be used with optical or copper cabling. Six of the ports are used to connect to Readout Controller rings optical links. One port can be used to connect to the Event Building Network to exchange data between DTCs. The last port is used to communicate with the Run Control Host computer.

The DTC receives Heartbeat packets from the Run Control Host. These packets are forwarded on each attached ROC ring link. Data packets from the Readout Controllers are returned on the same links. The DTC multiplexes data from six links into one timeslice which is then transferred to the Server over PCIe, or to other DTCs via the Event Building Network.

4.1.3 Run Control Host

The Run Control Host receives beam status and timing information from the Accelerator Controls network, and operator commands from the remote-control room. The Command Fanout (CFO) module in the Run Control Host is responsible for generating and synchronizing Heartbeat packets.

The DTCs receive control and timing information from the Run Control Hoston the DTC Control rings as shown in figure 4.2. These rings operate at 2.5 Gbps (3.125 Gbps 8b/10b encoded). The Command Fanout (CFO) card/optical link adapter in the Run Control Host is physically identical to that used for the DTCs.

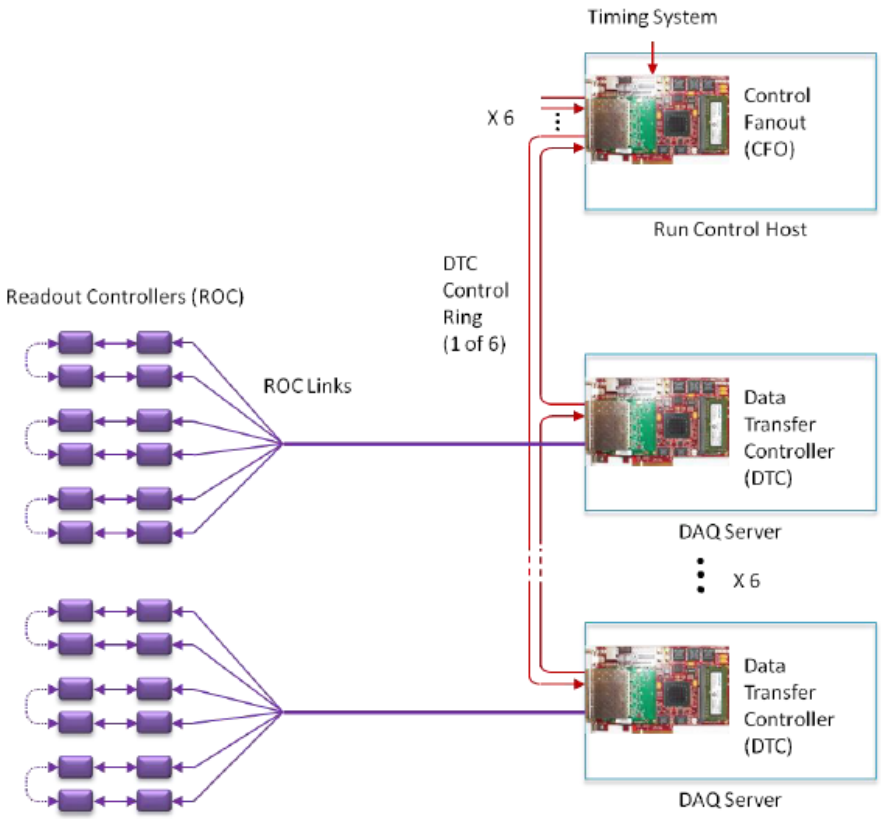

Figure 4.2 Control Hierarchy

4.1.4 Command Fan-Out

The Command Fan-Out Module (CFO Module) [6] provides an interface between the CFO Host and the DTCs. The CFO is based on the HiTech Global Kintex-7 (HTG-K700) PCI Express expansion card. This card features an eight lane Gen 2 PCI Express interface, a DDR3 SODIMM socket, and a 400 pin FMC connector, all wired to a Xilinx K325T Kintex-7 FPGA. The FMC connector allows the installation of an FMC card with eight SFP+ slots. This provides for multi-gigabit serial links for up to eight DTC Links, and ports for the System Clock and Super Cycle Start inputs. Firmware for the CFO's FPGA is based on a modified version of the DTC code.

Once configured, the CFO will issue a Heartbeat Packet for each System Clock cycle (uBunch) of a Super Cycle. A System Clock period is 1695 ns.

4.1.5 Event Building

Each Readout Controller collects data from a small subset of the detector. The Event Building (EVB) function combines these subsets to form a complete detector data set for analysis by an online processor. Event building is typically done in a switching network.

The event building function can be performed by software in the Server or by firmware in the DTC FPGA:

- Server option (“software" event building) - the event building network is a commercial 10Gbase-T Ethernet or Infiniband switch (Ethernet is more widely used, but Infiniband switches are less expensive with less software overhead). Event fragments are copied from the DTC to the Server over PCIe. EVB input and output buffers are in Server memory, and the processor handles all buffer management. The advantage of the Server option is that the event building software is part of the existing artdaq framework and is therefore easier to maintain. The disadvantage is additional processing load for event building and management of the network interface. If performance and scaling are satisfactory, then this is the preferred method. Tests of artdaq on a small, four server system with Infiniband networking have demonstrated a throughput of approximately 900 Mbytes/sec per server.

- DTC option (“hardware" event building) - the event building network is a commercial 10G SFP+ Ethernet switch. A port on each DTC card is connected to the switch via a direct-attach SFP+ copper cable. Input and output buffers are in DTC memory, and FPGA firmware handles the buffer management. Complete events are then copied from the DTC to the server over PCIe. No IP stack is necessary. The switch is programmed with a static MAC address table (one entry for each TDAQ server). The EVB network operates synchronously. Each server loops through its input buffers sending up to 2 KBytes packets. The advantages of the DTC option are that it offloads the processor in the TDAQ Server and allows use of the DTC FPGA for pre-processing or triggering on fully assembled events if needed. The disadvantage is that the FPGA firmware for event building is more difficult to develop and maintain than the artdaq event building software.

4.2 TIMING SYSTEM

The TDAQ will generate a continuous Mu2e System Clock [7] with frequency of 40 MHz at the Run Control Host Clock Fanout module (CFO). The CFO also receives the RF-cavity 0-crossing marker from the Accelerator. Note that this marker signal is synchronous with the arrival of proton pulses every 1695 ns to Mu2e and is only active through a ~43 ms spill. There will be at least 100 us of markers without proton pulses to start each spill. The time between spills is arbitrary, minimum is on order ~5 ms.

The CFO outputs a “punched" or “encoded" clock indicating the start of the Mu2e event window, which is synchronous with the Accelerator marker during beam ON to ~10 ns accuracy - the “punch" or “marker" is a change of the duty cycle of two cycles of the clock to either 25%/75% duty cycle, or 75%/25% duty cycle as shown in Figure 4.3. These two encodings are alternated so that a lost marker can be identified by the ROCs.

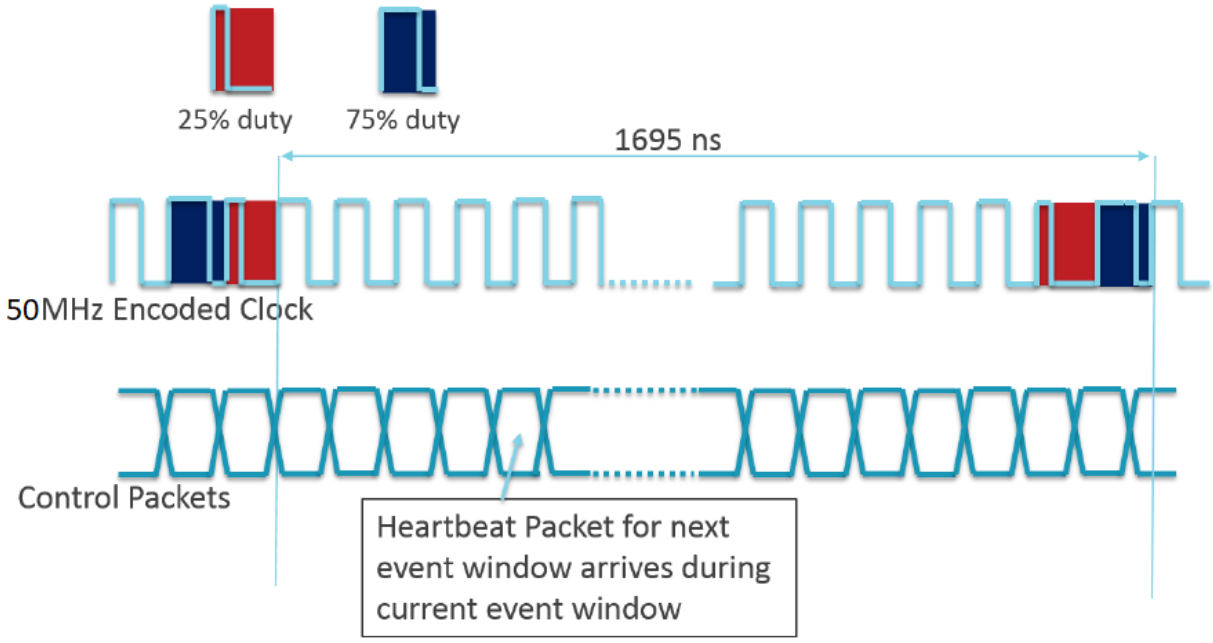

Figure 4.3 Encoded clock fanned out by the CFO

This encoded clock will be fanned out and distributed to outside the cryostat in the detector hall. Because of the grounding requirements, the TDAQ will distribute the signal optically from the TDAQ room to the Detector Hall.

The System Clock is not guaranteed to arrive to all detector ROCs phase aligned. Phase alignment at individual ROCs is accomplished by an adjustable delay at the ROC clock input. Additional alignment is provided by software or firmware calibration and may result in internal timestamp synchronization offsets. Each ROC generates its own internal high speed digitization clocks, phase locked to the System Clock. More specifically, to line up Event Windows the approach is to delay each front end to match front end with longest latency. Delay is determined calculating the Loopback signal. It consists in determine the round-trip time by sending a marker and returning it many times and taking the average time. Once the timing path latency is known for each front end, it is possible to calculate each front end's delay offset define as “longest timing path latency" minus “individual timing path". At the start of each run, the front ends can be configured to apply their own delta delay.

4.3 TDAQ SOFTWARE: *ARTDAQ*

The software architecture is based on artdaq [8]. This software runs on TDAQ servers and on dedicated control and monitoring computers. artdaq is a toolkit of C++ 2011 libraries and programs for use in the construction of TDAQ systems.

It provides functionality that includes the following:

- management of the readout and configuration of the TDAQ hardware. This makes use of experiment-supplied software components.
- routing of data between threads within a process, between different processes, and between different machines, and for assembling complete events from these data.
- encapsulation of the data being routed, and support for experiment-specific raw data formats to provide type-safe data access.
- event analysis and filtering using the art event-processing framework.
- basic control and monitoring applications
- infrastructure for distributing configuration data to TDAQ processes

The artdaq data acquisition toolkit is used to build the Mu2e TDAQ software system. artdaq provides software applications for managing the data flow as well as libraries and applications for encapsulating the data, analyzing the data, and performing other basic data acquisition functions.

The toolkit is designed to provide core functionality while allowing experiments to customize the hardware readout and event analysis as needed.

4.3.1 Otsdaq

otsdaq is the Ready-to-Use data-acquisition (DAQ) solution aimed at test-beam, detector development, and other rapid-deployment scenarios. Otsdaq uses the artdaq DAQ framework under-the-hood, providing flexibility and scalability to meet evolving DAQ needs and provides a library of supported front end boards and firmware modules which implement a custom UDP protocol. Additionally, an integrated Run Control GUI and readout software are provided, preconfigured to communicate with otsdaq firmware.

Otsdaq comes as a web page where it presents:

- State Machine: allows to easily configure and run DTCs
- Macro Maker: a tool that allows the user to execute front-end interface writes and reads and build sequences of writes and reads. Macros can be saved per user or made public. Macro can be exported to C++ or directly to a target plugin as a FE Macro.
- FE Macro: C++ member functions.
- Configuration trees: defines the hierarchical relationship between all entities in the online DAQ system, and all their parameters. When otsdaq is launched the executables that starts are the ones enabled in the configuration tree for that node. Then later, when the state machine transitions to the Configured state, the children of the executables are instantiated based on the parameters defined by the chosen configuration alias.
- Code editor: a tool that allow editing and viewing of source code and text files

- Console: web app that allows the user to exist remotely, the core functionalities are built on artdaq message facility.

Data processing is the primary responsibility of the online DAQ. Mu2e's event window data will be processed through artdaq modules. However, ots allows for data processor plugins in general.

5 OTSDAQ: CONFIGURE()

The configure() is the function called by otsdaq to configure the daq electronics and is accessible via web-interface using the “state machine”.

The configuration process, that right now involves only a chain of 10 DTCs on the calorimeters, was taking 5-10 minutes. This was more of a problem for the testing of the system, where configurations are repeated several times, than for the actual operating phase in which the electronics, once configured, remain that way ideally forever.

Due to the lack of documentation, the first step was to study the configuration process by printing out address, data and timestamp of every writing access to the DTCs registers. By doing so, it was noticed that the configure() consisted of 5 steps:

- Step1: DTC reset
- Step2: Jitter Attenuator input select
- Step3: Configure Jitter Attenuator
- Step4: DTC registers setup
    a. CFO marker Enable
    b. Link Enable
    c. DMA Timeout Enable
- Step5: Reset RX CDR unlock count register

By comparing the printouts with a script used to configure the DAQ electronics manually, it was noticed the configure() steps were all necessary except for the jitter attenuator configuration needed only after a power cycle.

The configuration process was taking so long because the steps were performed sequentially between different DTCs, this means that first step 1 was carried out for DTC1 then for DTC2 and so on. This was discovered by looking at the printouts timestamp.

The process was sped up by creating two new configuration profiles, one, to be run only after power cycles, in which you had the jitter attenuator configuration and one without. In both the configuration steps were parallelized by adding more threads to the process itself. Now the Configure() takes more or less 1 minute and since the configuration happens in parallel it shouldn’t be affected by the future addition of more DTCs chains.

6 DTC Loss Of Lock

When the experiment starts the CFO will provide control signals on up to 8 chains of 9 DTCs maximum. Currently the test for DTC synchronization is carried out on a chain of 10 DTCs i.e. a worst case scenario.

At first, to save on system costs, the clock distribution from the CFO to the DTCs was performed serially, this caused the jitter, at the end of the chain, to be out of the system specifications. So, to switch to a parallel distribution, ad RTF (Rj45 Timing Fanout Module) was added which receives the clock from the CFO and fans it out to the DTCs.

An RTF module has the basic function of combining the inputs and distributing multiple outputs, using, both as input and output, an RJ45 interface. The one used in the DAQ room works as a multiplexer.

After the installation of the RTF another problem was noticed, the DTC couldn't hold the lock on the encoded signal received. It was first thought that the issue still connected to the clock jitter, so both RTF and DTC firmware were updated to lessen it.

The RTF firmware was modified so the clock in was put through a PLL working as a jitter attenuator. While in the DTC, figure 6.1, in the first configuration the clock was retrieved from the received data stream by comparison with an internal reference clock and then this was passed through a jitter attenuator in which the clock from the RTF was used as a reference. In the current configuration, in all DTCs, the clock received from the RTF is sent both to the Rx and Tx.

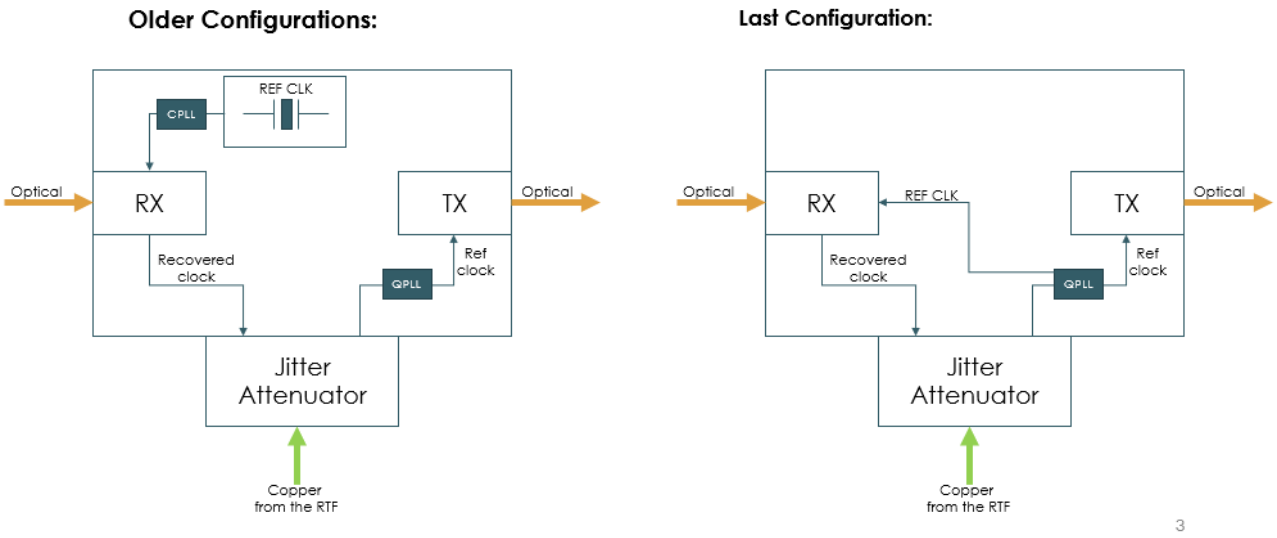

Figure 6.1 DTC clock distribution configuration

Currently has been reached a semi-stable configuration in which the DTCs are able to hold the lock for a maximum of 5 days.

The next step is to complete a jitter measurement, with an oscilloscope, on the system to eliminate it as a possible cause of the loss of lock.

7 EVENT BUILDER – CONNECTIONS

Each DTC collect data from several ROCs related to a sensor subsystem. The EVB function combines these subsets to form a complete detector set for analysis by an online processor. As already said, this function can be implemented via software, the Event Builder are the TDAQ servers, or via hardware where the Event builder are the DTCs with the help of a commercial 10G SFP+ Ethernet switch (Figure 7.1).

The HW solution is the one adopted in the DAQ room because it offloads the TDAQ servers using the DTC FPGA for preprocessing.

The first 10 port of the Ethernet switch has been connected to the DTCs related to the trackers (from 01 to 05) using 10 duplex multimode optic fiber cables (50/125MM) of 2mm of diameter.

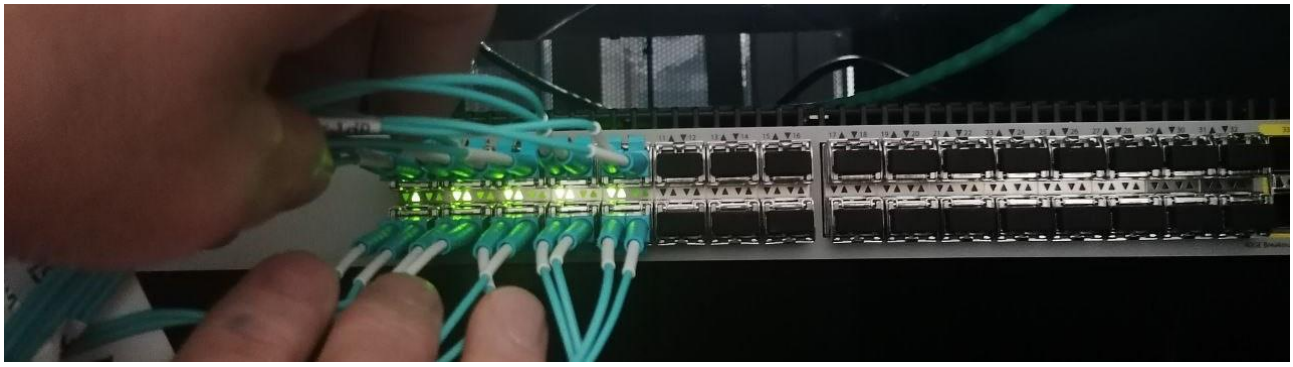

Figure 7.1: EVB Ethernet switch output ports

Special attention was paid to the design of the design of the cable labels. Each fiber has 2 labels in each end in which the position of the end itself is specified and to do so, the proposed hierarchical organization of the DAQ room has been followed:

The DAQ is divided into 10 racks, 5 on the east side, 5 in the west side; each rack is identified with a number from 1 to 5 and a letter to specify the side of the room.

In the west side of the room, the racks have one or two blocks, one at the top and one at the bottom, consisting of 4 or 5 nodes and a chassie but nevertheless the enumeration of the nodes and of the chassie goes from bottom to top depending on the functionality of the node (Figure 7.2).

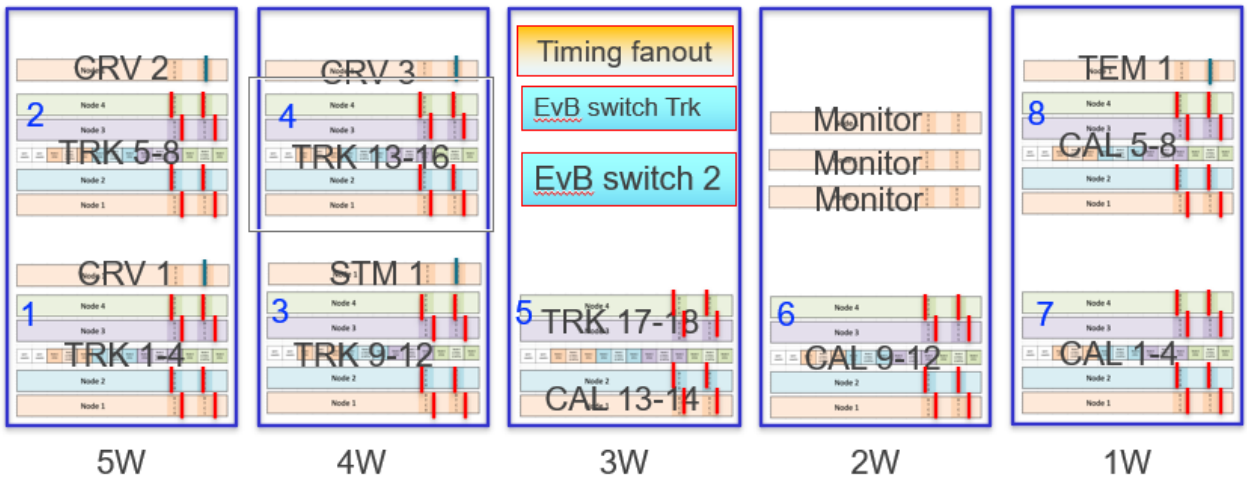

Figure 7.2: Proposed Racks organization on the west side of the DAQ room, the red bars in each node are the DTCs

In a chassie there are slots for 14 cassettes and those are enumerated from left to right. In each chassie there are 8 cassettes for the DTCs of the block, 4 ROC cassettes and two spare cassettes. Each ROC cassette “talks” to 12 ROCs while in each DTC cassette a DTC has 12 2x MTP connections.

In the each chassie we group 2 DTC cassette with 1 ROC cassette using a D-R-D scheme(Figure 7.3).

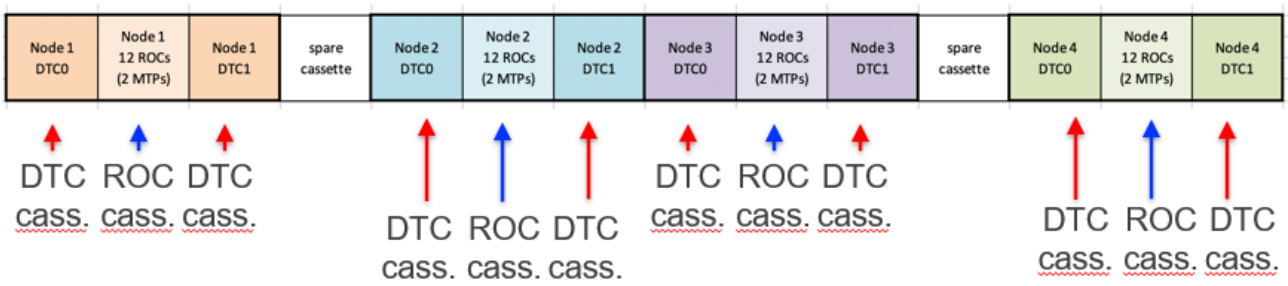

Figure 7.3 Proposed chassie organization

So in the labels, the position of each fiber end is identified by specifying Rack, chassie, cassette, pin.

8 RTF FIRMWARE

The serial distribution of the clock from the CFO to the DTCs didn't match the jitter requirements of the system so an RTF was needed to have a parallel distribution.

Right now, the CFO is connected to a chain of 10 DTCs, placed in 5 different TDAQ server (calorimeters). At the end, there will be 36 TDAQ server occupying 4 different racks of the DAQ room so there will be 72 DTC overall. To have parallel distribution the RTF needs to fan out the clock to all the DTCs.

The RTF firmware has been updated so that it fans out the clock to all its output ports, the new output ports were defined in the entity, and then routed on the correct pin of the RTF FPGA. The firmware has been tested on the hardware, to do so the corresponding bitfile was generated and uploaded on the RTF then the presence of the clock on all output ports was verified 10 at a time by configuring the DTCs via otsdaq, verifying that they hold the lock and then moving them to the next 10 ports.

Since all the output worked, the firmware was saved in the memory of the FPGA so it will survive power restart.

9 BIBLIOGRAPHY

[1] United States Department of Energy Fermilab., «Mu2e Experiment. 2020.,» [Online]. Available: https://mu2e.fnal.gov/.

[2] For the United States Department of Energy Fermi Research Alliance FRA, «Mu2e Technical Design Report,» 2014.

[3] V. Rusu, «Tracker Electronics,» Fermilab, United States Department of Energy, 2018.

[4] INFN, «The Mu2e Calorimeter Final Technical Design Report,» 2014.

[5] Microsemi, «Polarfire FPGAs,» 2020. [Online]. Available: https://www.microsemi.com/product-directory/fpgas/3854-polarfire-fpgas#documentation.

[6] R. A. Rivera e R. Kwarciany, «Mu2e Command Fan-Out Module Hardware User Guide, Mu2e-doc-5619,» Fermilab, United states Department of Energy, 2018.

[7] R. Rivera e R. Kwarciany, «TDAQ Timing Distribution Specification Mu2e-doc-7048,» Fermilab, United States Department of Energy., 2018.

[8] R. Kwarciany e G. H.-S. E. Flumerfelt, «Trigger and DAQ Architecture and Status,» Fermilab, United States Department of Energy, 2019.

[9] United States Department of Energy Fermilab., «Mu2e Experiment.,» [Online]. Available: https://mu2e.fnal.gov/.