University College London

Department of Physics and Astronomy
High Energy Physics Group

First Year Ph.D. Transfer Report

Towards A Distributed Scalable Water Cherenkov Detector

Petr Mánek

supervised by
prof. Jennifer Thomas and Benedikt Bergmann

November 2020

Abstract

The Chips R&D project aims to develop affordable water Cherenkov detectors for large-scale underwater installations. In 2019, a 5 kt prototype detector Chips-5 was deployed in northern Minnesota to study neutrinos generated by the nearby NuMI beam. The presented work is comprised of several discrete tasks. Firstly, an online hit sorting algorithm was developed for Chips DAQ, removing the need for computationally intensive post-processing after run periods. In a benchmark equivalent with Chips-5, the algorithm has shown performance suitable for large-scale application. Secondly, a dedicated low-latency time distribution system was implemented to deliver timing signals from the Fermilab accelerator to the Chips-5 detector with sub-nanosecond precision. In a time-of-flight study, the system has reliably offered a time budget of 610 ± 330 ms for on-site triggering. Finally, a series of software packages was developed for use with the next generation of Chips detector instruments. Among these, a calibration toolkit for photomultiplier gain tuning has achieved a $88\times$ relative speedup compared to previously used programs, advancing towards widespread adoption of the new technology in the current and future Chips deployments.

Contents

Acronyms

ADC Analog-to-digital converter

BSYNC Beam-synchronous clock system

CP Charge-parity (symmetry or violation)

CPU Central processing unit

CW Cockroft-Walton (PMT base)

DAQ Data acquisition

FSM Finite-state machine

GPS Global Positioning System

HV High voltage

CHIPS Cherenkov Detectors in Mine Pits

IRIG-B Inter-range instrumentation group timecode B

MTU Maximum transmission unit

NOνA NuMI Off-Axis ν_{e} Appearance

NuMI Neutrinos at the Main Injector

PCB Printed circuit board

PE Photoelectron

PMT Photomultiplier tube

PWM Pulse-width modulator

RAM Random-access memory

R&D Research and development

SSH Secure shell

TCLK Tevatron Clock

TDS Timing Distribution System

TDU Time distribution unit

UDP User datagram protocol

UTC Coordinated Universal Time

WIPAC Wisconsin IceCube Particle Astrophysics Centre

WR White Rabbit

Introduction

1 Theoretical Background

1.1 Neutrinos & Neutrino Oscillations

The Standard Model of particle physics describes neutrinos as neutral fermions. Predicted in 1930 by W. Pauli [1] and discovered experimentally in 1956 [2], neutrinos are known to only interact at short range through the weak force, exclusively coupling with W and Z bosons. This implies that even though large fluxes of neutrinos may be emitted by human-constructed sources and astronomical phenomena, only a minuscule fraction of particles can be effectively observed.

Neutrinos exist in quantum superposition of three mass and three weak eigenstates. While they freely propagate through space as the former, they are measured as the latter. The mass eigenstates are labeled ν_1, ν_2 and ν_3. The weak eigenstates are the *leptonic flavours* that usually give neutrinos their name: electron neutrinos (ν_e), muon neutrinos (ν_μ) and tau neutrinos (ν_τ). Due to the absence of one-to-one correspondence between the mass and weak eigenstates, mixing between them causes spontaneous flavour changes. This phenomenon, known as *neutrino oscillations*, has been theoretically described by B. Pontecorvo, Z. Maki, M. Nakagawa, and S. Sakata in the 1960s [3, 4, 5]. In their honour, the unitary matrix that associates the mass and weak eigenstates is named the *PMNS matrix* (U):

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \underbrace{\begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix}}_{U} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{1}$$

Even though the PMNS matrix is comprised of nine elements in total, it has been shown that it can be fully determined by only four parameters. These are known as the three mixing angles (θ_{12}, θ_{23} and θ_{13}) and a single phase angle labeled δ_CP. With this parametrization, the matrix U is re-written as

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_\mathrm{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_\mathrm{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_\mathrm{CP}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_\mathrm{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_\mathrm{CP}} & c_{23}c_{13} \end{pmatrix}, \tag{2}$$

where $s_{ij} = \sin\theta_{ij}$ and $c_{ij} = \cos\theta_{ij}$ for brevity. It should be noted that this parametrization asserts neutrinos to be *Dirac* particles in nature, acquiring their mass via a Yukawa coupling. In contrast, if neutrinos are found to be *Majorana* particles, two more complex phase angles need to be introduced, commonly denoted α_{21} and α_{31}. Since these terms, however, lie on the matrix diagonal, they do not affect neutrino oscillations.

1.2 The δ_CP Parameter

While the three mixing angles were experimentally measured with relatively high accuracy, only estimates of δ_CP are available at the time of writing. This phase angle, which is so named because of its relationship with charge-parity (CP) violations, has strong implications for the status of CP symmetry in the lepton sector. Its precise measurement may help explain the matter-antimatter imbalance in the Universe through the hypothetical process of *leptogenesis* [6].

One promising channel for measuring δ_{CP} is by study of neutrino oscillations, in particular the transition from ν_μ to ν_{e} (known as ν_e *appearance*). By theory, its probability in the full three-flavour model is given as [7]

$$P(\nu_\mu \to \nu_{\mathrm{e}}) = \sin^2(\theta_{23}) \sin^2(2\theta_{13}) \frac{\sin^2(\Delta(1-A))}{(1-A)^2} \tag{3}$$

$$+ \alpha \tilde{J} \cos(\Delta \pm \delta_{\mathrm{CP}}) \frac{\sin(\Delta A)}{A} \frac{\sin(\Delta(1-A))}{1-A} \tag{4}$$

$$+ \alpha^2 \cos^2(\theta_{23}) \sin^2(2\theta_{12}) \frac{\sin^2(\Delta A)}{A^2}, \tag{5}$$

with $A = \frac{2\sqrt{2} G_F N_e E}{\Delta m^2_{31}}$, $\tilde{J} = \cos(\theta_{13}) \sin(2\theta_{13}) \sin(2\theta_{12}) \sin(2\theta_{23})$, $\Delta = \frac{\Delta m^2_{31} L}{4E}$ and $\alpha = \frac{\Delta m^2_{21}}{\Delta m^2_{32}}$. In the equations, G_F denotes the Fermi weak coupling constant, N_e is the electron density in matter, Δm^2_{ij} is the mass-square difference between the mass eigenstates i and j, and $\pm$ reduces to $+$ for neutrinos and $-$ for antineutrinos. This motivates the possibility to measure Δm^2_{31} and δ_{CP} from A and term (4) respectively [8].

1.3 Water Cherenkov Detectors

Neutrinos can be observed in water Cherenkov detectors, which exploit the Cherenkov effect. This phenomenon generates electromagnetic radiation, referred to as the *Cherenkov light*, when a charged particle propagates in a dielectric medium at a speed exceeding the speed of light in such medium [9]. Due to the motion of the charged particle, Cherenkov light is emitted at angle θ with particle's velocity, creating the characteristic shape of a cone (shown in Figure 1) reminiscent of a sonic boom wavefront. Furthermore, the emission angle θ satisfies

$$\cos\theta = \frac{1}{n\beta}, \tag{6}$$

where β denotes the particle speed as the ratio of the speed of light, and n is the refractive index of the medium. Thanks to this relationship, the speed of the charged particle may be measured from the angle θ. For instance, this yields a cone with angle 21° for a typical 20 MeV muon in water.

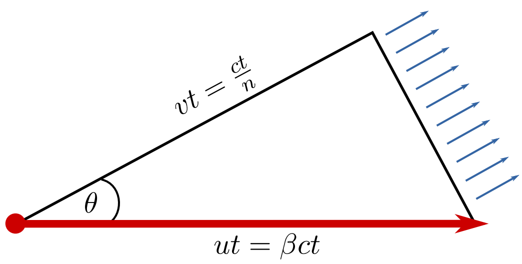

Figure 1: Cherenkov light (blue) emitted with angle θ from a charged particle of velocity u (red).

To observe charged particles such as neutrinos, Cherenkov detectors usually follow common methodology. The sensitive volume is a tank designed to induce the Cherenkov effect in incident particles. This volume is surrounded from multiple angles by a large number of photo-sensitive instruments. When a charged particle enters the volume, emitted cones of Cherenkov light propagate towards the surface of the tank, projecting observable ring-like patterns.

In subsequent analysis of these patterns, various properties of the incident particle can be reconstructed. For instance, the speed of the particle is estimated from the measured cone angle, and its trajectory is given by the cone axis. In some cases, particle identification can even be achieved by more sophisticated inference.

2 CHIPS Experiment

2.1 Design Aims

Due to weak interaction properties of neutrinos, Cherenkov detectors aiming to observe them must often rely on large sensitive volumes in order to collect statistically significant results. This usually implies long construction schedules, enormous instrumentation costs and impractical organizational overhead. The CHIPS R&D Experiment attempts to address these issues by proposing a novel, affordable design for water Cherenkov detectors [10].

The CHIPS detector uses a water tank engineered to be submerged in a large body of water (e.g. a lake or a flooded mining pit). Compared to other alternatives, this concept provides great flexibility as well as variety of cost-saving opportunities. First, due to the surrounding water, the CHIPS detector requires far less structural support underwater than if deployed on dry land. Second, thanks to natural shielding from cosmic radiation provided by the overburden of water above, it is unnecessary to invest in expensive mining operations in order to place the CHIPS detector deep underground. Next, the detection medium can be obtained from a local source of water. Last, in construction, standardised off-the-shelf components are preferred to proprietary hardware.

By design, CHIPS assemblies follow a unified modular architecture, offering a wide range of versatile configurations suitable for a diversity of applications. This implies that CHIPS detectors can be constructed and deployed relatively fast and without much effort. Furthermore, thanks to their interoperability, detector parts can be easily serviced or upgraded throughout their lifetime, reducing possibly prohibitive upfront costs that are frequently associated with conventional water Cherenkov experiments.

At the time of writing, two CHIPS detector assemblies were successively deployed in a flooded mine pit in the north of Minnesota, USA: (1) in 2014 the pilot CHIPS-M project [11, 12, 13], and (2) in late 2019 the 5 kt CHIPS-5 project, which is the subject of this report.

2.2 The CHIPS-5 Detector

The CHIPS-5 detector is the latest-generation CHIPS prototype. Using a 12 m tall cylindrical tank with 12.5 m radius (illustrated in Figure 2), the detector provides 1924 m^2 available surface area, offering the first practical insights into construction, deployment and operation of CHIPS devices at such scale.

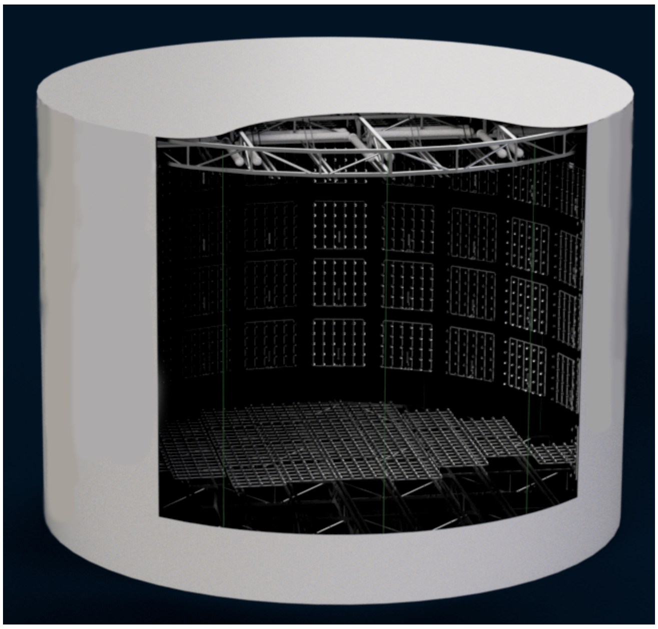

Figure 2: Artist's rendering of the CHIPS-5 detector. Image courtesy of Thomas Dodwell.

The detector was constructed and deployed throughout the summer of 2019. First, its stainless steel frame was built in a shallow part of the partially flooded mining pit. This frame, comprised of two *end caps*, represents the main support structure of the detector tank. Once completed, both end caps were installed with electronics, and joined together by 12 m cables. Later in the year, the pit was flooded and the ready-to-deploy assembly was floated from the shallows to its final location. Securely anchored in place, the detector's bottom cap was released, descending to its target depth due to its mass. Conversely, the top cap remained in its position due to floatation, stretching the cables between the end caps, and giving the tank its intended tall shape. Since the net mass of the entire detector exceeded that of the surrounding water, it further submerged deeper into the pit, gaining natural protection from cosmic rays and harsh weather conditions.

The detector tank is fully insulated from its surroundings by lightproof fibreglass-reinforced plastic liner. In addition to protecting sensitive instruments within the tank from external light sources and wildlife, the liner also serves as a watertight divider. This permits on-shore facilities to continuously pump and purify the water contents of the tank in order to reach desired photon attenuation conditions. The concrete water purification plant used in CHIPS-5 is expected to produce photon attenuation length of 133 ± 2 m after approximately two months of uninterrupted operation [14].

2.3 Detector Planes

Photo-sensitive components of CHIPS detectors are organised in *detector planes*, standardised arrays of photomultipliers (PMTs) that surround the detection volume from multiple angles. This design concept is in part motivated by the notion of distributed systems, with individual planes operated independently of each other, and responsible for pre-processing of their respective data outputs.

For increased affordability, frames of CHIPS planes are usually constructed from commercially available materials such as standard schedule 40 PVC pipes, glue and cement. This allows PMTs to be laid out in simple matrices, sometimes, depending on their intended location, tilted at an angle. As a protection in the event of flooding, detector planes contain water-blocking components at important junctions. This includes DAQ electronics boxes that are attached to each plane.

For the purposes of data analysis, measured outputs of detector planes are represented as sequences of PMT *hits*, discrete events describing light digitised by individual PMTs. Due to precise time synchronization, individual hits are associated with timestamps that allow their clustering. In addition, hits also carry information about light intensity that is derived from the time-over-threshold of the digitised sensor pulse. This data, along with PMT locations within the detector known *a priori*, is considered invaluable for 3D reconstruction of Cherenkov event vertices as well as particle trajectories.

Later in Section 4, implementation of plane *triggering* is discussed. This refers to the task of selective data acquisition with the aim to only measure PMT hits from time periods, when a particular neutrino source is known to be active. Based on their requirements for hardware support, two types of triggering are distinguished: (1) *active* triggering, which relies on planes to measure data only in precisely defined time periods, and (2) *passive* triggering, where data is measured continuously (even outside desired times), and later rejected by software downstream of the DAQ pipeline. Furthermore, depending on scientific objectives, triggers can be operated in two modes: (1) *in-phase* mode, which focuses on the neutrino signal, or (2) *out-of-phase* mode, which focuses on the background.

As shown in Figure 3, the CHIPS-5 detector uses planes of two hardware types: (1) *Nikhef planes*, each comprised of HZC PMTs and KM3NeT DAQ electronics [15, 16], and (2) *Madison planes*, which consist of Hamamatsu PMTs and dedicated DAQ electronics developed in joint effort by CHIPS and the Wisconsin IceCube Particle Astrophysics Centre (WIPAC). Due to their direct relevance to this report, Madison planes are further described in Section 5.1. In total, the CHIPS-5 detector design accommodates up to 226 Nikhef and 30 Madison planes, containing over 6500 PMTs.

In software, all planes are orchestrated by a finite-state machine (FSM) program. This program issues instructions to other component of the CHIPS DAQ system, for instance the DAQonite program, which is responsible for receiving PMT hit data, or the DAQontrol program, which performs plane configuration and run control.

Figure 3: Detector planes installed within the Chips-5 detector. Lower-density Madison planes are seen in the front, while higher-density Nikhef planes are in the back. Image courtesy of Simeon Bash.

2.4 Neutrino Source

The Chips-5 detector site in northern Minnesota has been intentionally selected to lie 7 mrad off the axis of the NuMI neutrino beam [17] (shown in Figure 4). Historically, this beam has been studied in a variety of experiments, and therefore represents a well-understood source of neutrinos.

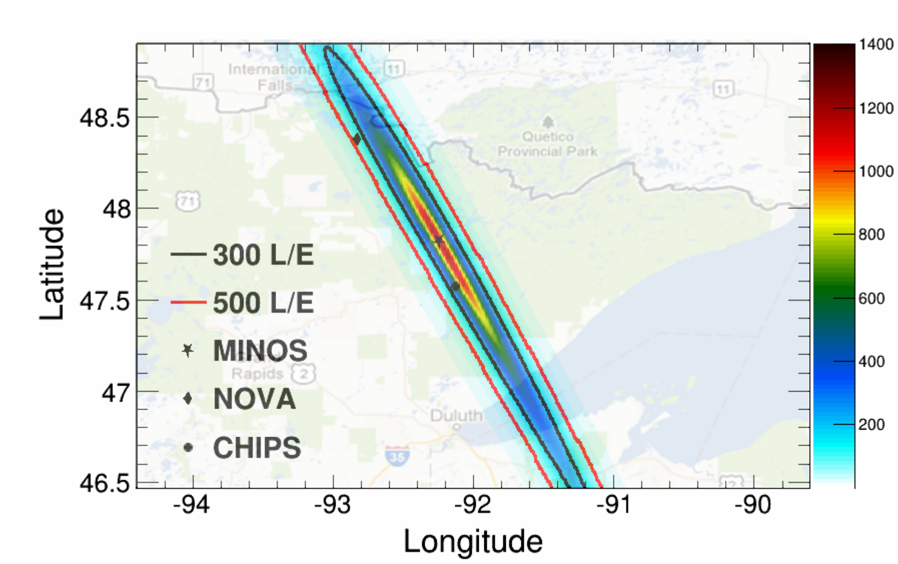

Figure 4: Map showing the Chips-5 detector location superimposed with expected NuMI beam intensity, viewed in terms of expected neutrino flux (assuming no oscillations) [10].

The NuMI beam originates at Fermilab near Chicago, Illinois. Under normal operation, the Fermilab accelerator's Main Injector periodically generates a beam predominantly comprised of muons. These particles decay in-flight into a mixed beam of electron and muon neutrinos, which can be observed to exit the facility in 10 µs spills with 0.75 Hz frequency. The remaining non-neutrino components of the beam are absorbed when it enters the Earth's crust, resurfacing 707 km away in Minnesota due to Earth's curvature. Depending on the accelerator's configuration, the NuMI beam can also be generated with antiparticles instead of particles, i.e. electron and muon antineutrinos.

In addition to NuMI neutrinos, the Chips-5 detector is also sensitive to cosmic muons, which survive its water overburden. In order to reliably detect Cherenkov events produced from such source, Chips-5 top-cap planes use outward-facing PMTs in addition to conventional inward-facing PMTs. The motivation is that these additional instruments can act as a veto signal for cosmic events.

Completed Work

3 Online Hit Sorting

In data analysis, PMT hits from CHIPS detectors are expected to be sorted in a time-ordered sequence. Since CHIPS DAQ is inherently a distributed system with many instruments operating simultaneously, this is by no means guaranteed at runtime. For this reason, offline hit sorting is usually performed after individual runs are completed. Even though this practice is effective, it is considered suboptimal due to severe computational and storage penalties incurred in the process. This section presents an algorithm designed to perform hit sorting *online* during ongoing runs. The proposed method permits significant reduction in memory and storage footprint, and offloads a significant portion of computing load onto peripheral facilities of the DAQ system.

3.1 Background

To better understand the proposed algorithm, let us first examine a simplified case. Consider two finite sequences of PMT hits corresponding to the same time period, which were produced by two detector planes. Furthermore, assume that both these sequences were individually sorted earlier. In order to merge such sequences into a single sorted sequence, a forward pass through both sequences can be performed that trivially interleaves hits based on their timestamps (illustrated in Figure 5). This requires $\mathcal{O}(n)$ computational steps, where n is the input hit count.

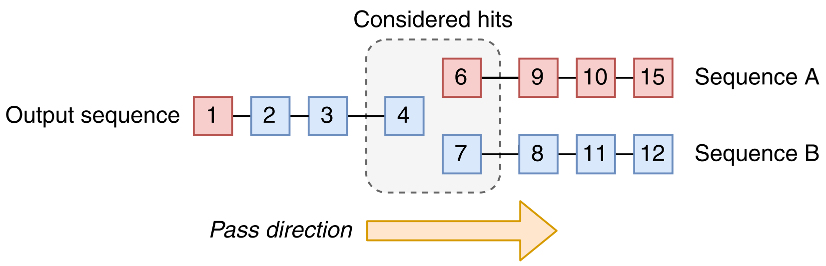

(a) State during the interleaving process, partially merged sequence is seen on the left.

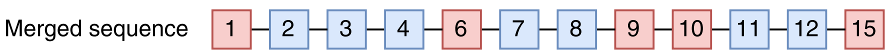

(b) State after the interleaving process, only a fully merged sequence remains.

Figure 5: Interleaving of two sorted PMT hit sequences by a linear pass. Square boxes with numbers represent hits with timestamps. The dashed area location, from which the next output hit is selected.

The interleaving approach can be efficiently generalised for an arbitrary number of input sequences; let us denote that number $k > 2$. Applying recursion to the problem, the k input sequences can first be grouped in pairs and merged individually. Next, the $k/2$ interleaved sequences can be grouped and merged again, and again, until finally only a single sequence remains. This is known as the k-way merge strategy [18].

For complexity analysis, it is convenient to model this algorithm with binary trees. Consider the smallest tree that contains k leaves; such tree has $\mathcal{O}(\log_2 k)$ levels. The algorithm begins by placing the

k input sequences in the leaves at the bottom level. It can be observed that actions of the algorithm correspond with the binary layout of the tree. In particular, connections between individual node pairs and their parent nodes in the level above represent acts of interleaving two sequences into one; this is shown in Figure 6.

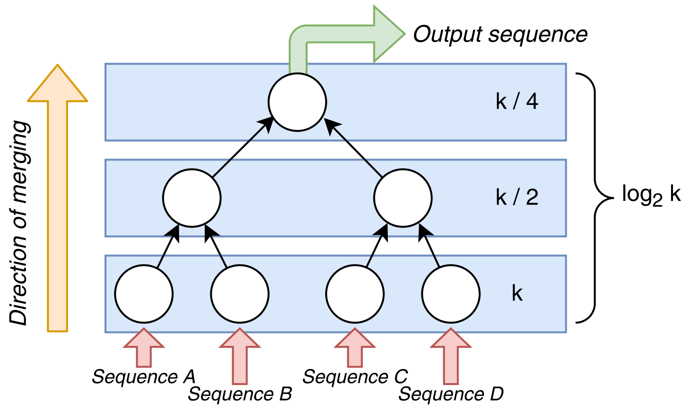

Figure 6: Interleaving of k sorted PMT hit sequences. Input sequences (red) are propagated from the bottom of a binary tree upward. The output sequence is obtained from the root of the tree (green).

The diagram indicates that the algorithm may also be viewed as a series of grouping stages corresponding to individual tree levels, each of which reduces the number of sequences by a half. Counting the number of interleaved sequences (nodes) from the bottom of the tree up, the following geometric series emerges:

$$\frac{k}{2} + \frac{k}{4} + \frac{k}{8} + \cdots + 2 + 1 = \sum_{p=1}^{\lceil \log_2 k \rceil} k \cdot \left(\frac{1}{2}\right)^p. \tag{7}$$

The sum of this series, combined with the known complexity of each interleaving step, gives the algorithm its desirable $\mathcal{O}(n \log_2 k)$ running time.

3.2 Practical Application

In Chips DAQ, k-way merging was implemented by upgrading the DAQonite program, which is responsible for receiving PMT hits from detector planes. In this context, data streams produced by individual planes were viewed as the input sequences processed by the algorithm.

To ensure that plane data streams were well-ordered prior to processing, two steps were taken. First, microcomputers on-board planes were modified to sort hits before their transmission to the rest of the DAQ infrastructure. Since planes consist of multiple PMTs, this effectively implies execution of the same algorithm within each individual plane. Secondly, due to concerns about plane processing power, an automatic failsafe was implemented in DAQonite that allows distributed load balancing in scenarios where planes become unable to cope with excessive computational strain.

The failsafe robustly applies insertion sort to each plane data stream prior to processing [19]. While it may seem redundant to perform additional sorting on a sequence that is already expected to be well-ordered upon reception, the motivation will soon become clear. Due to the desirable properties of insertion sorting, the algorithm's running time on fully sorted sequences is $\mathcal{O}(n)$. This implies that under normal circumstances, where planes are able to sort hit data online, no additional computing overhead is introduced in DAQonite. If at any point planes become overloaded, their software can attempt to alleviate the situation by temporarily suspending in-situ sorting, thereby transmitting data to DAQonite in unordered or partially ordered state. When that happens, the computing load associated with hit sorting is effectively transferred from planes to machines the central DAQ infrastructure, which will observe increase in insertion sorting complexity.

During implementation, the assertion that all presented algorithms operate on PMT hit sequences of finite size had to be reconciled with parameters of Chips runs, which often cannot be technically considered finite until completed. For that reason, DAQonite was configured to divide incoming PMT hit sequences at runtime into finite time windows, which can be processed independently. To maintain tractability, a life cycle policy was introduced for resource management. This ensures that windows are opened, filled with hits, closed, processed with k-way merge, streamed to a permanent storage facility, and recycled in timely manner.

3.3 Testing

Before the presented upgrades were deployed, a series of performance tests was conducted to investigate scalability in real-world conditions. In each test case, n PMT hits were randomly generated from k simulated planes (data streams) and merged while the wall time was tracked.

The benchmark was performed on a laptop[1] for $(n, k) \in \{16, 32, 64, 128, 256, 512, 1024, 2048\} \times \{4096, 8192, 20000, 65536\}$. Its results, shown in Figure 7, indicate that the running time of the algorithm appears to scale linearly with n. Consistently with expectations a hierarchy emerges, where the slope of the complexity characteristic is seen to be directly proportional to k. Overall, the presented algorithm yielded satisfactory performance, processing a realistic batch of 65 thousand hits organised in 256 streams in less than a second.

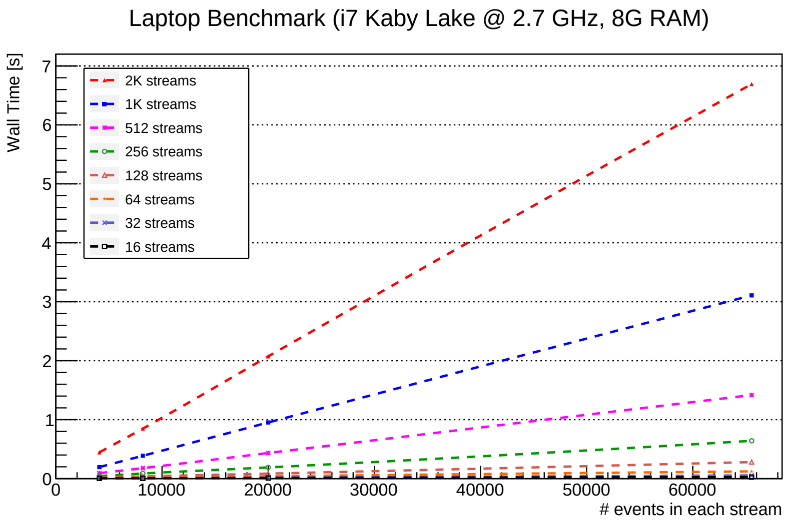

Figure 7: Results of merging scalability benchmark. The number of hits n is given by the X-axis, the number of simulate planes (data streams) k is indicated by various colours (see legend), and the measured wall time is plotted on the Y-axis.

Following integration with CHIPS DAQ, the presented work was further verified in laboratory conditions. In particular, its effects were studied on datasets measured by real planes located in an experimental dark room. Figure 8, which compares the states the hit sorting disabled and enabled, shows that the algorithm successfully merged out-of-order hit bursts into a fully sorted sequence.

4 Low-latency NuMI Trigger

To implement triggering on the NuMI beam, a low-latency Timing Distribution System (TDS) was developed that allows CHIPS DAQ to filter out undesirable data sources at runtime. Based on a similar system used by the NOνA experiment, CHIPS TDS relies on precise time synchronization with UTC and fast signal forwarding between the accelerator at Fermilab and the CHIPS-5 detector. This section describes design and implementation of such system.

[1]Intel® Core™ i7-8565U CPU (Kaby Lake, 2.70 GHz frequency), 8 GB RAM

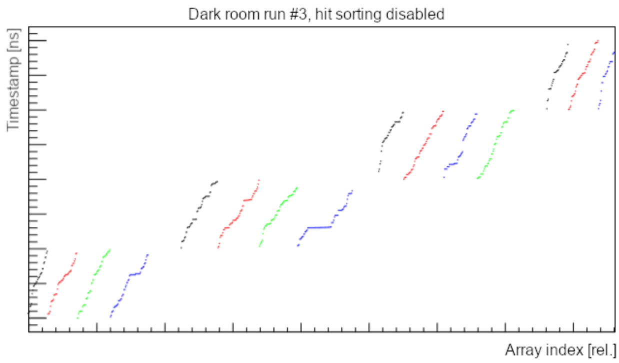

(a) Dataset with hit sorting disabled.

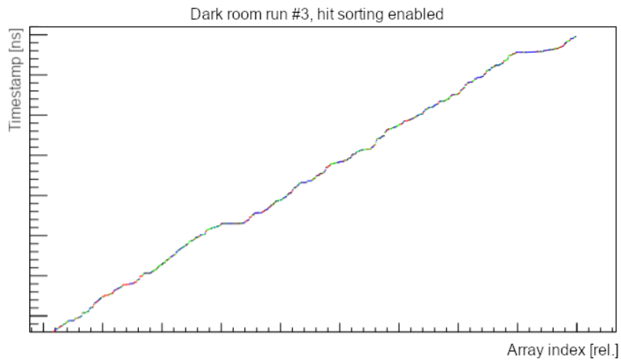

(b) Dataset with hit sorting enabled.

Figure 8: Effects of the presented algorithm on a dark room dataset (colours correspond to streams). PMT hits are plotted as timestamps (Y-axis, ns) vs. their position in the array (X-axis).

4.1 NuMI Beam Spill Cycle

The Fermilab accelerator employs precise timing systems to track operation of its principal components: (1) the *beam-synchronous clock system* (BSYNC), and (2) the *Tevatron Clock* (TCLK) [20]. Under normal conditions, these systems periodically emit characteristic signals that describe various events of interest in the accelerator duty cycle.

Available accelerator signals that are relevant to the NuMI beam are shown in Figure 9. Out of these, in order to be considered viable for triggering, a time signal must satisfy several requirements. Firstly, the time elapsed between the signal and the accelerator neutrino spill must have minimal jitter. Secondly, the signal must be reliably emitted for all spills. And finally, the signal must be available with a sufficient time in advance. With these requirements, two possibilities remain: (1) the TCLK signal \$A5, which marks the reset of the accelerator prior to the start of a new spill cycle, and (2) the BSYNC signal MIB\$74, which is emitted when neutrinos exit the accelerator complex.

4.2 NOνA Time Distribution Units

To obtain access to Fermilab accelerator signals, CHIPS-5 relies on existing facilities used by the NOνA experiment [21]; in particular, its proprietary TDS [22]. Similar in purpose to the CHIPS TDS, the NOνA TDS timestamps accelerator signals with high-precision UTC provided by a commercial GPS satellite receiver. This information is transmitted through a hierarchy of Timing Distribution Units (TDUs, shown in Figure 10) that permit system-wide synchronization to within 7.8 ns across all elements.

Internally, NOνA TDUs contain a PowerPC 8347 computer providing a Linux platform. Usually, this computer hosts NOνA software that is responsible for consumption of the received timing signals. All timestamps handled at that point are compliant with the NOνA time specification that measures time elapsed as UTC ticks from the "NOνA Epoch" defined as 00:00:00 January 1, 2010 GMT.

The source of information for the CHIPS-5 TDS is a prototype NOνA TDU, which was procured and installed at Fermilab in the autumn of 2019. To serve as an adapter between NOνA and CHIPS timing systems, a dedicated program was developed to replace standard NOνA software in the TDU. Among other tasks, this program is responsible for decoding consumed timing signals, converting timestamps from NOνA to UTC specification and forwarding data to the CHIPS-5 detector site through a low-latency communication channel.

4.3 Signal Delivery

The signal delivery chain must overcome a variety of security systems while maintaining relatively short delivery times. For this purpose, a relay computer was installed at Fermilab to facilitate elaborate tunnelling scheme (illustrated in Figure 11). From the CHIPS-5 detector site, a reverse SSH tunnel was initiated to the relay that permits the DAQonite program to expose a server interface. By a similar

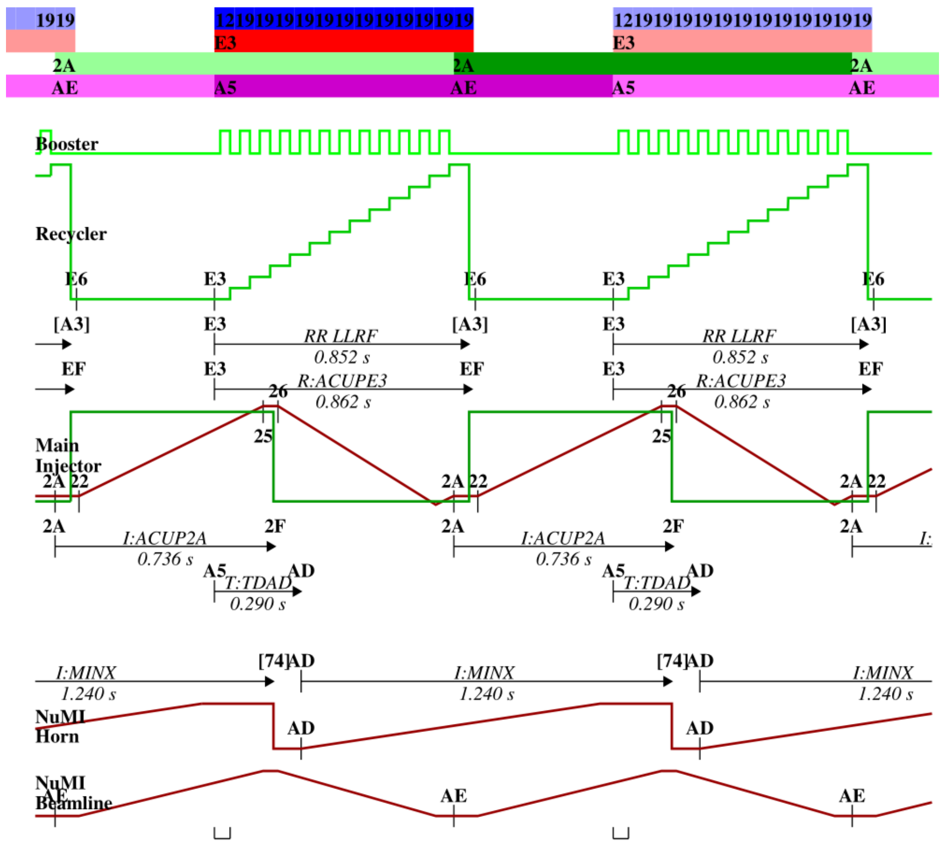

Figure 9: Selected parts of the timing structure of the Main Injector cycle relevant to the NuMI beam. In addition to component names and known fixed delays, time signals are labeled by two-letter identifiers in bold (e.g. AE, 2A or 74). Plot courtesy of Phil Adamson.

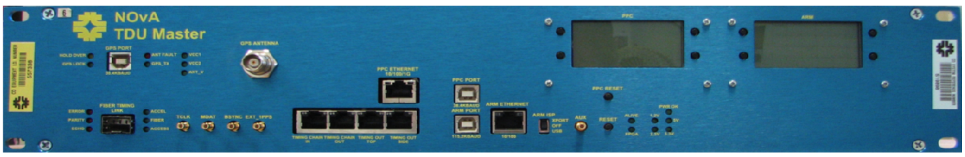

Figure 10: NOνA Timing Distribution Unit (TDU). [22]

mechanism, an identical interface was propagated from the relay to the TDU, completing the delivery chain. Under normal operation, when a timing signal is consumed by the TDU, the information travels through the relay to the DAQonite server backend with standard latency within 10 ms.

For the purposes of monitoring, the relay computer was further upgraded to multiplex transmitted signals, track their frequency, and produce warnings if signals of various types fail to be observed on a programmed schedule. In addition, the device also possesses automatic fault recovery capability that addresses a known issue of occasional spurious TDU reboots. When such a condition is detected, the relay automatically re-configures the TDU and bootstraps the proprietary CHIPS software.

4.4 Trigger Implementation

Once received by DAQonite, timing signals serve a variety of purposes in CHIPS DAQ. First, along with other information retrieved from the Fermilab accelerator complex, all signals are logged in the on-site storage facility, which supplies information to DAQ monitoring displays (shown in Figure 12) and offline analysis jobs. Next, if a data run is ongoing, timing signals can be used for passive triggering. Last, if a data run is ongoing that includes Madison planes, signals may also be utilised for active triggering. Depending on the desired type of an ongoing data run, both triggering mechanisms support in-phase and out-of-phase operation, effectively determining whether PMT hits are recorded during

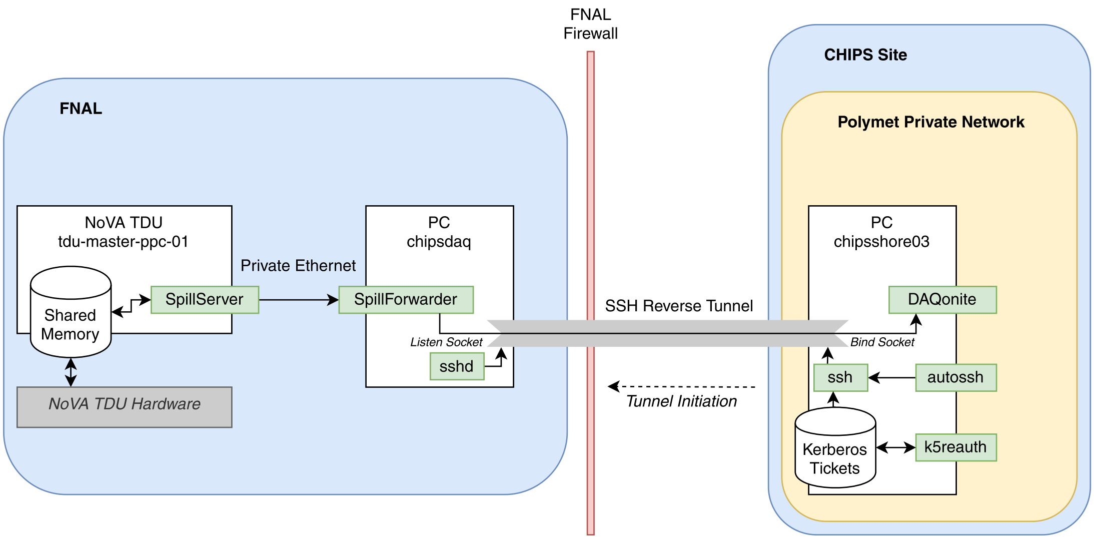

Figure 11: Schematic of the spill signal delivery system. Fermilab accelerator signals are decoded by a NOνA TDU, which forwards them through the "chipsdaq" relay computer to CHIPS detector site.

time periods of NuMI beam activity, or their complement respectively.

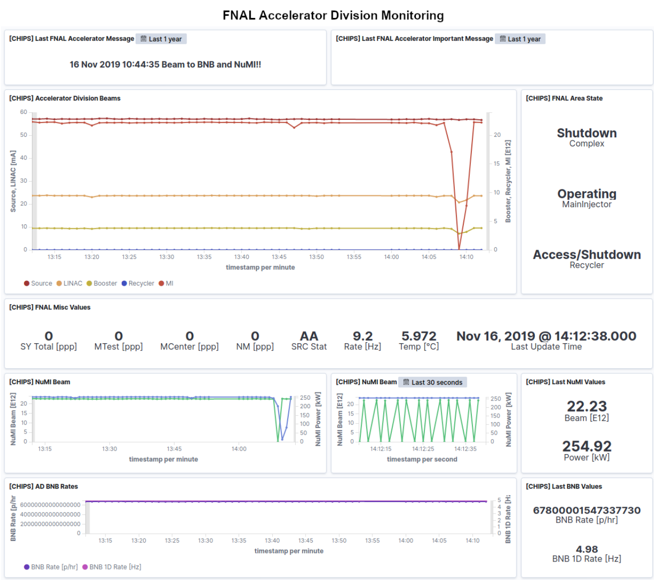

Figure 12: Real-time dashboard that was developed to display information about the state of the Fermilab accelerator to CHIPS DAQ operators. Visualization is performed by the Kibana front-end.

Passive triggering is realised by a continuous period of data acquisition combined with software rejection of undesirable PMT hits downstream in DAQonite. Since such implementation is platform-agnostic, it is compatible with all CHIPS-5 plane types. Conveniently, software components that facilitate passive triggering are also responsible for scheduling time windows for hit merging (first introduced in Section 3.1). For that reason, no additional implementation was required to discard hits outside of trigger periods.

In contrast to passive triggering mode, active triggering does not require continuous data taking, implying considerable reductions to DAQ load. Instead it relies on hardware support of detector planes, which was only implemented in Madison planes at the time of writing. During active triggering runs, planes remain in idle state by default, ready to start data acquisition at short notice. When a timing signal is received, a corresponding time window is calculated and programmed in the detector planes. Since their peripherals are precisely synchronised with UTC, which is conventionally used for PMT hit timestamping, planes control their own acquisition and only measure hits inside of the programmed time window. For that reason, no out-of-window hits are transmitted to, and processed by DAQonite, implying reduced network bandwidth as well as computing load.

4.5 Testing

The presented Chips TDS solution was successfully implemented, and its software packages were deployed to the Fermilab and Chips-5 detector sites. To thoroughly test the new system, real timing signals were observed and recorded in runs, each spanning a period of several days during the NuMI beam operation. Throughout the monitored period, the system has demonstrated resilience to random as well as deliberately introduced failures, and the capability to automatically recover into nominal state once all faults were corrected.

Relative accuracy of delivered timestamps was evaluated by analyzing frequencies of various periodic timing signals and comparing them with the known durations in the Fermilab accelerator duty cycle. Among the examined signals, one that is of particular interest is the \$74 signal, which can be expected to show consistent[2] period of 1.333 s. Aggregating roughly 128,000 signals observed over a 3 day run (shown in Figure 13), the period was experimentally determined to be $1.333\,\mathrm{s} \pm 317.7\,\mu\mathrm{s}$. While this agreement provides independent verification of the system, the \$74 signal appears unsuitable for use in triggering due to occasional discrete variations in its period. Following further examination, it was determined that the \$A5 signal would be used instead.

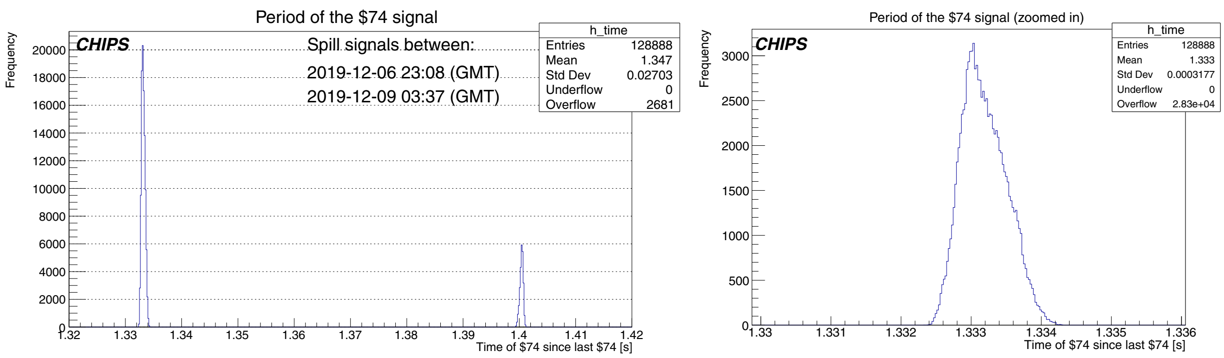

(a) Discrete variations observed in the \$74 signal.

(b) Detail of the peak at 1.333 ms in the left plot.

Figure 13: Time between subsequent emissions of the \$74 signal aggregated in a histogram over the period of roughly 52 hours of NuMI beam operation.

In order to assess trigger signal viability, the amount of time from its emission to the subsequent arrival of neutrinos at the Chips-5 detector must be considered. This period represents the *total time budget*, during which the signal must be delivered between sites in order to remain viable for triggering. The budget is constituted by two components: (1) the time from the signal emission to the neutrino spill at the accelerator site $t_{\$A5\to\mathrm{spill}}$, and (2) the neutrino time of flight t_{travel} that is given by d/c, where $d = 707\,\mathrm{km}$ is the baseline length a c is the speed of light. Since the \$74 signal marks the moment of accelerator neutrino spill, the duration between two subsequent signals $t_{\$A5\to\$74}$ known from the accelerator duty cycle can be used to calculate $t_{\$A5\to\mathrm{spill}}$ in theory.

Alternatively, the same duration can also be measured experimentally (as shown in Figure 14) with the added benefit of considering the jitter $\sigma_{\$A5\to\$74}$ in the calculation. In such case, the estimate is conservatively given by $t_{\$A5\to\$74} - \sigma_{\$A5\to\$74}$. Combining all the listed components, the time budget evaluates as

$$\begin{aligned} t_{\mathrm{budget}} &= t_{\$A5\to\mathrm{spill}} + t_{\mathrm{travel}} && (8)\\ &= t_{\$A5\to\$74} - \sigma_{\$A5\to\$74} + t_{\mathrm{travel}} && (9)\\ &\approx 1.437\,\mathrm{s} - 3.48\,\mu\mathrm{s} + 2.5\,\mathrm{ms} = 1.4395\,\mathrm{s}. && (10) \end{aligned}$$

[2]Even though the observed period is known to vary in discrete increments depending on the active Fermilab accelerator cycle, the given period can be expected to be dominant.

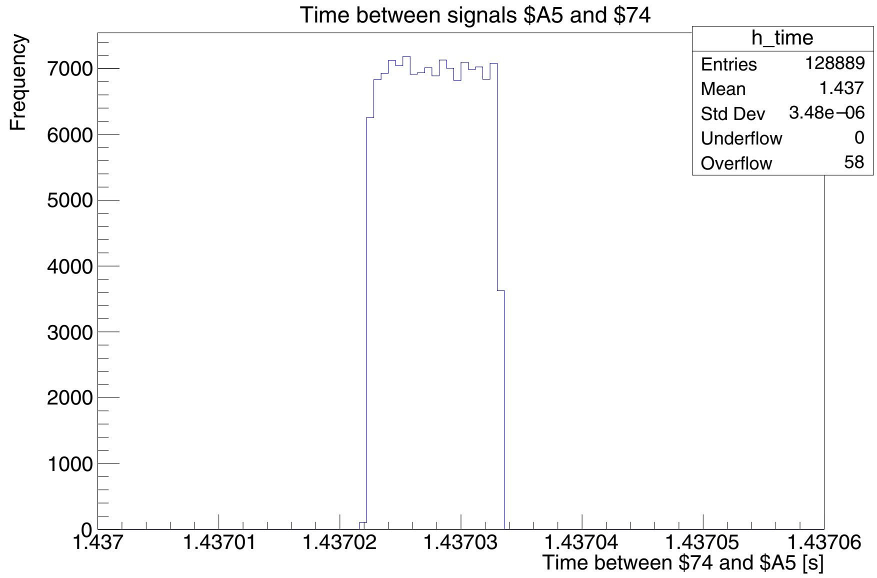

Figure 14: Time between subsequent emissions of the $A5 and $74 signals aggregated in a histogram over the period of roughly 52 hours of NuMI beam operation.

After subtraction of signal delivery time due to network latency, the remaining budget for scheduling time windows at the CHIPS detector site based on this calculation is plotted in Figure 15.

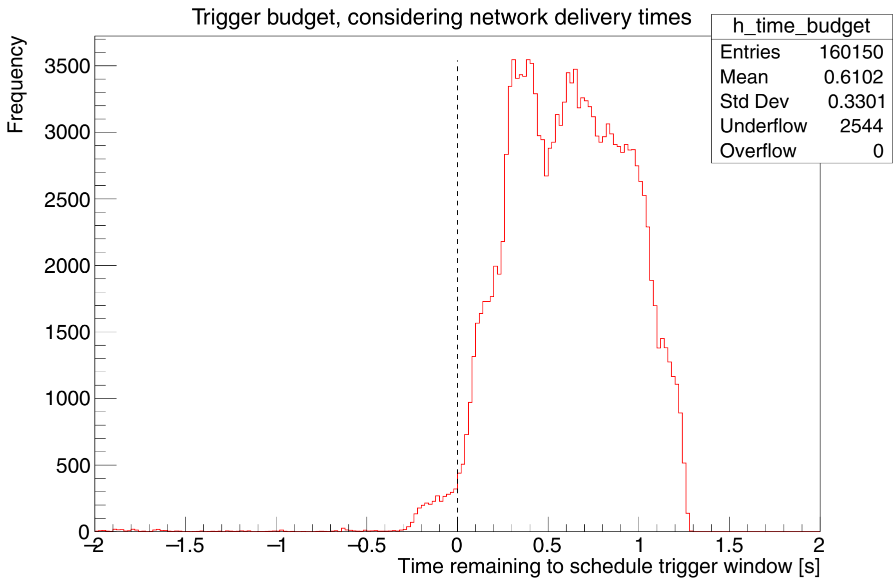

Figure 15: Histogram of the budget remaining to schedule time window in CHIPS DAQ. Negative values are not viable since neutrinos beat the signal to the detector.

According to the analysis, out of roughly 160,000 observed signals, approximately 3.69 % had negative time budget, meaning that in such cases signals were delivered to the detector only *after* it had already encountered NuMI neutrinos. The remaining 96.3 % of accelerator signals in some form preceded their corresponding neutrino spills, opening the possibility for scheduling time windows ahead of the incoming spills. At this point, it should of course be noted that this fraction is still bound to shrink depending on the latency of the final DAQ implementation. However, even if it is conservatively assumed that the on-site scheduling process would take additional 300 ms to complete, more than 80 % of spills would still remain viable.

5 Madison Plane Upgrade

As described in Section 2.3 of the introduction, CHIPS-5 detector planes can be divided into two groups based on their hardware: (1) Nikhef planes, which have already been successfully tested, and (2) Madison planes, which can be viewed as a next-generation PMT readout setup, and are still at a proof-of-concept stage. This section describes various upgrades to the software operating Madison planes to close this gap by improving their integration with CHIPS DAQ systems. Overall, the presented efforts aim to bring Madison planes closer towards a robust and scalable solution that is appropriately equipped to withstand adversarial conditions of a long-term deployment within the CHIPS-5 detector.

5.1 Background

Madison planes are organised in a tree topology (illustrated in Figure 16) that places CHIPS DAQ in the root, and detector PMTs in the leaves [23]. This architecture permits low-latency run control by system-wide broadcasts, while allowing data payloads to be relayed from the bottom of the tree upwards, maintaining favourable bandwidth characteristics.

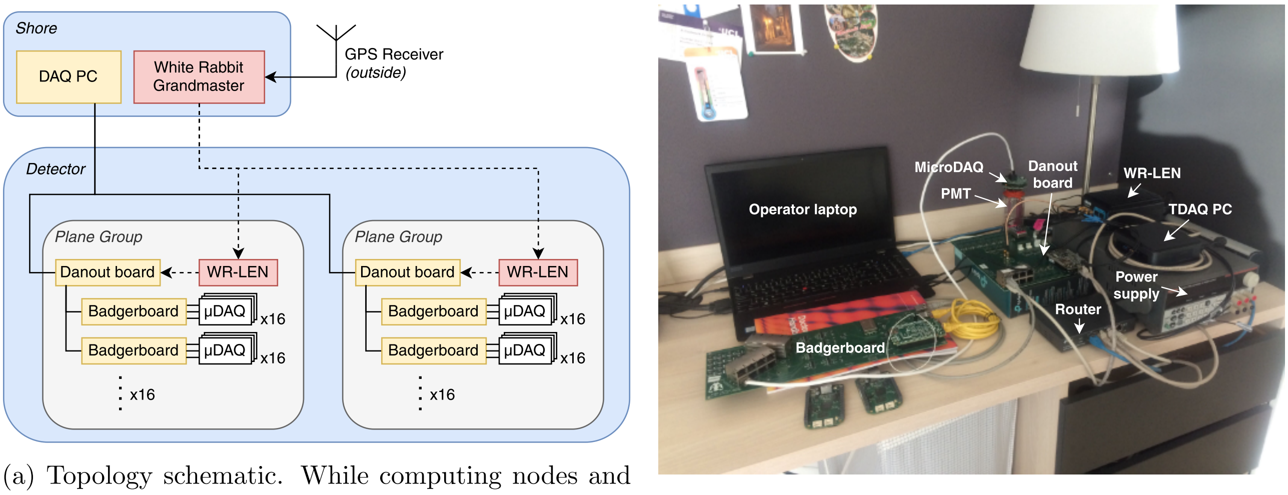

(a) Topology schematic. While computing nodes and data links are drawn in yellow boxes and solid lines, timing nodes and links are drawn in red boxes and dashed lines, respectively.

(b) Minimal testing setup used during the SARS-CoV-2 pandemic, comprised of a single PMT, MicroDAQ, Badgerboard, Danout board and a WR-LEN.

Figure 16: Components of a Madison detector plane.

The smallest building block of a Madison setup is a *MicroDAQ*, a compact microprocessor board with a form factor of the Cockroft-Walton (CW) base that is usually tightly coupled with a single PMT. Developed in the frame of CHIPS R&D, MicroDAQ has been designed to facilitate all activities associated with its PMT. For instance, the board generates oscillator signal that is converted by the CW base into power for high-voltage (HV) operation. In addition, MicroDAQ also reads out PMT signals, handling digitisation of the PMT sensor pulse and timestamping with sub-nanosecond resolution.

A single Madison detector plane consists of 16 PMTs arranged in a 4×4 matrix, with their corresponding MicroDAQs orchestrated by a *Badgerboard*, a proprietary PCB that performs configuration, monitoring, power switching and data multiplexing. The Badgerboard is accessible by a conventional Ethernet link via on-board Beaglebone microcomputer that conveniently provides a Linux platform for high-level DAQ operations. Finally, groups of planes are controlled by a *Danout board*, another PCB that distributes power and accurate UTC time information to up to 16 planes. Analogous to Badgerboards, Danout boards also expose a Linux interface through an on-board Beaglebone.

Precise time information is distributed by the White Rabbit (WR) system, which relies on a hierarchy of proprietary devices synchronised with sub-nanosecond accuracy by extended Ethernet protocol [24]. At the CHIPS-5 detector site, a commercial GPS satellite receiver decodes UTC time, which drives the grandmaster WR node that synchronises all detector planes. From that point, time is relayed through fibre-optic links to WR-LEN nodes [25], which transform it into a reference 10 MHz clock signal, and a IRIG-B timestring. These signals are forwarded through copper links over a relatively short distance through Danout boards and Badgerboards, and finally consumed by MicroDAQs.

5.2 Data Streaming Implementation

Madison planes are operated from the Beaglebones on the Badgerboards by the *BadgerApp*, a program that reads out PMT hit data and saves them in the Beaglebone's local storage. In the scope of the presented work, this software was upgraded to continuously stream PMT hit data to the rest of the CHIPS DAQ infrastructure.

To harmonise Madison setups with their Nikhef counterparts, a dedicated network protocol based on UDP was designed. Under normal operation, this protocol defines communication channels to contain a sequence of variable-size packets (further illustrated in Figure 17) corresponding to time windows with PMT hits. This design has a variety of advantages. For instance, it allows network bandwidth to be conserved by using fewer bits to encode hit timestamps relative to the window start. In addition, at the other end of the data stream, receivers can carry out gap detection by monotonic sequence number comparison.

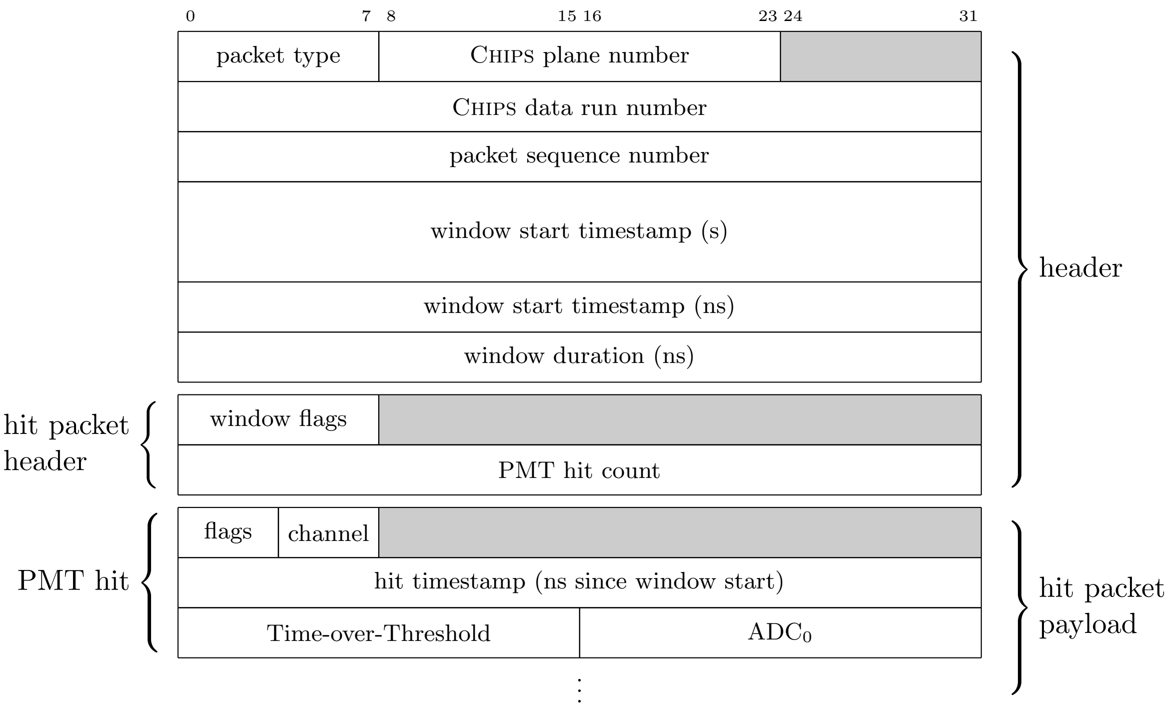

Figure 17: Structure of a UDP time window packet that contains Madison PMT hits.

The proposed protocol was implemented in the BadgerApp and the DAQonite programs. Their new components were successfully tested with randomised as well as adversarial datasets, and cleared for large-scale deployment in CHIPS-5. In a subsequent effort, the BadgerApp was additionally upgraded to allow hit data auditing by emitting monitoring packets, similar in structure and implementation to the PMT hit data stream. This feature proved particularly useful for detection and troubleshooting of network infrastructure issues; for instance, congestion and MTU size tuning. While initial tests on limited setups were successful, works on a wider-scale deployment are ongoing.

5.3 Plane Configuration & Run Control

In addition to data streaming, new network-based software packages were developed to facilitate remote configuration and run control of Madison planes. Specifically, novel "badgerd" and "danoutd" daemons were implemented to control BadgerBoards and Danout boards, respectively. Unlike the

BadgerApp, which is active only during run periods, these programs are designed to operate continuously throughout the entire lifetime of their corresponding Beaglebone hardware, exposing high-level server interface towards Chips DAQ systems. Their features include per-channel and per-plane power switching, and MicroDAQ firmware flashing, which is a convenient maintenance feature. Plane configuration and run control are achieved by manipulating local BadgerApp settings files and controlling the runtime of the BadgerApp program, respectively.

For network communication related to the presented tasks, a minimalistic protocol was devised and implemented. Operating within conventional client-server paradigm, the new daemons expect to receive request packets (illustrated in Figure 18) containing specific instructions to carry out. In addition to such requests, a periodic heartbeat signal is emitted that reports the current state of the controlled hardware as well as the latest information retrieved from on-board sensors; this includes accelerometer, pressure, humidity and temperature data. The solution was integrated with the Chips FSM & DAQontrol program, and successfully tested in laboratory conditions.

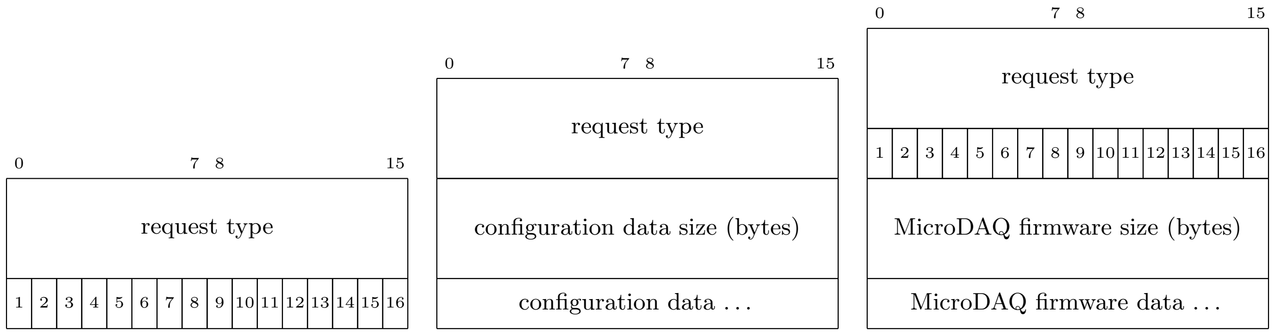

(a) MicroDAQ power control packet. (b) Plane configuration packet. (c) MicroDAQ flashing packet.

Figure 18: Structure of selected Madison control packets. Where applicable, MicroDAQ channel selection is implemented by bitmaps indexed 1-16.

5.4 PMT Calibration Procedure

Prior to and during installation, Chips PMTs undergo elaborate calibration procedures that determine their optimal settings for HV operation in runs. In Madison planes, these activities have been usually performed by Jupyter notebooks that were prototyped alongside hardware. To ensure viability for large-scale deployment, a toolkit of high-performance programs based on the ROOT framework [26] was implemented to replace Python-based tools in the most frequent calibration tasks.

One such task is the adjustment of parameters that determine the gain of the PMT. At runtime, MicroDAQs internally compare the PMT analog pulse with a fixed threshold level using a built-in discriminator. When this level is exceeded, the MicroDAQ *trigger* is activated[3] and a hit is produced from the digitised pulse. Accounting for manufacturing irregularities, PMT thresholds must be tuned individually for each tube, so that all detector instruments are equalised. Inconveniently, MicroDAQ discriminator configuration is fixed, only allowing thresholds to be changed indirectly by manipulating PMT gain. This is achieved by setting the bias voltage generated by the CW base, which is controlled by adjustable oscillation frequency of a pulse-width modulator (PWM) of the MicroDAQ. The calibration procedure therefore aims to map the PWM-threshold characteristic, and find PWM frequency such that a desired threshold is attained. For Chips-5 PMTs, this has been chosen to be 0.3 photoelectrons (PE) [27].

First, multiple dark noise datasets with both random and discriminator-driven triggering are measured at various PWM frequencies. The randomly triggered data bypasses the MicroDAQ discriminator, and therefore the cut-off, which allows a value known as the *pedestal* to be estimated as the mean ADC (shown in Figure 19). Even though this value is calculated for each PWM frequency separately, it is expected to be independent of the PWM setting.

[3]In spite of similar names, MicroDAQ internal trigger is *not* associated with per-plane triggering described in Section 4.

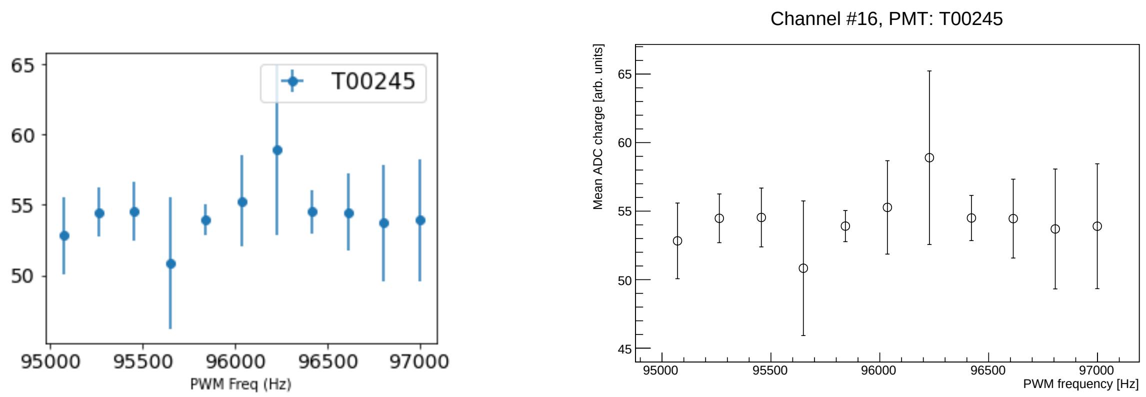

(a) Result produced by Jupyter notebook.

(b) Result produced by the new toolkit.

Figure 19: Pedestals estimated by the old and the new tools. Identical results can be observed.

Next, the discriminator-triggered data are analyzed to locate ADC peaks for individual PWM frequencies (shown in Figure 20). Even though the data may be burdened with backgrounds, it is usually massively dominated by 1 PE values due to dark noise. To avoid instability due to outliers, a robust peak-recognition method is utilised. In addition to detected peaks, cut-off ADC values are determined as the smallest observed ADCs exceeding the pedestal.

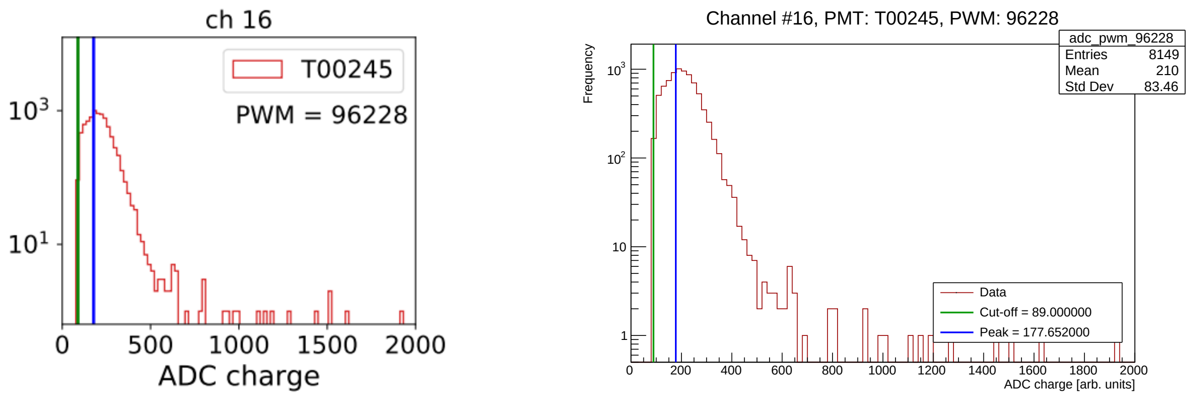

(a) Result produced by Jupyter notebook.

(b) Result produced by the new toolkit.

Figure 20: ADC peak & cut-off found by the old and the new tools. Identical results can be observed.

Using values gathered from randomly triggered measurements, cut-off ADCs are converted to absolute scale in PE as

$$\mathrm{ADC}^{(\mathrm{abs.})}_{\text{cut-off}} = \frac{\mathrm{ADC}_{\text{cut-off}} - \mathrm{ADC}_{\mathrm{ped}}}{\mathrm{ADC}_{\mathrm{peak}} - \mathrm{ADC}_{\mathrm{ped}}}\ [\mathrm{PE}]. \tag{11}$$

With absolute ADC cut-off plotted as a function of set PWM frequency, a continuous exponential is fitted through all collected data points using a least-squares minimiser provided by the MINUIT package [28]. Once this characteristic function is fully determined, it is easily inverted, and a PWM frequency corresponding to 0.3 PE is calculated (shown in Figure 21).

In order to assess the relative speedup delivered by the new calibration toolkit with respect to the older Python-based solution, a performance benchmark was conducted on hardware[4] comparable to workstations used in CHIPS facilities. During the benchmark, the presented calibration procedure was executed on a set of previously measured data files, and CPU wall time was tracked. Furthermore, results produced by both methods were saved and compared. The benchmark results are listed in Table 1. While both evaluated methods yielded identical outputs within the tolerance of the test, the new developed toolkit required a considerably smaller period of computing time. In relative terms, this is equivalent to a speedup of approximately 88× on identical hardware.

[4] Intel® Core™ i7-8565U CPU (Kaby Lake, 2.70 GHz frequency), 8 GB RAM

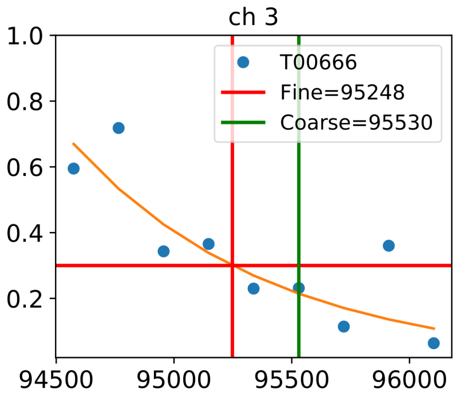

(a) Result produced by Jupyter notebook.

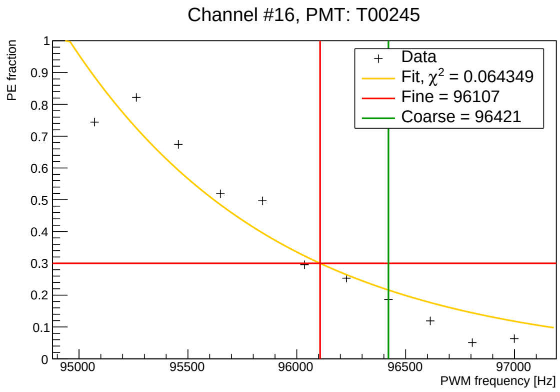

(b) Result produced by the new toolkit.

Figure 21: PWM fine-tuning performed by the old and the new tools on similar datasets. The new result (on the right), calculated by the MINUIT package [28], produces a more robust fit.

Method	Wall time	# channels	Wall time / channel	Speedup factor
Jupyter notebook	01:31.560	16	00:05.688	1.000×
New toolkit	00:00.904	14	00:00.065	88.042×

Table 1: Results of the calibration benchmark.

Conclusion

6 Summary

During the course of the past academic year, the author's activities were focused on a wide variety of discrete tasks related to the ongoing CHIPS-5 project. This report describes the most important of such tasks.

First, thanks to the online sorting algorithm (presented in Section 3), CHIPS DAQ will no longer necessitate computationally intensive post-processing procedures following run completion. Next, the novel CHIPS Timing Distribution System (described in Section 4) increases signal-background separation efficiency in the CHIPS-5 detector. Last, through dedicated software (detailed in Section 5), calibration procedures are accelerated with the factor of 88×, and the integration of Madison planes with CHIPS DAQ systems is improved, advancing towards their widespread adoption in the current and future CHIPS deployments.

In addition to these activities, the author has also dedicated a considerable amount of time and energy to implementation of the CHIPS FSM facility, aiming to efficiently orchestrate increasingly distributed CHIPS hardware. Furthermore, by continuously working to parallelise CHIPS DAQ systems, the author has enhanced their overall scalability and load balancing capability. Finally, throughout the year 2019, the author has spent nearly 3 weeks at the CHIPS-5 site, contributing to joint detector deployment efforts in various capacities.

7 Further Work

In the short term, future work will predominantly focus on finalization, large-scale tests and benchmarks of the developed Madison plane software. Following their completion, Madison plane components will be further upgraded for the purposes of a spin-off CHIPS R&D project.

The new project aims to develop a compact next-generation Cherenkov detector setup, based on the current design of Madison planes. This apparatus will be designed for fully autonomous and fault-tolerant operation over the Internet, without the requirement of conventional CHIPS DAQ backend in close vicinity. This renders it particularly suitable for long-term modular deployments across large distances; e.g. in near and far detectors, or observation of astronomical phenomena.

Bibliography

[1] W. Pauli. "Letter to the radioactive ladies and gentlemen". In: *Letter to the physical society of Tübingen* (1930). URL: http://www.pp.rhul.ac.uk/~ptd/TEACHING/PH2510/pauli-letter.html.

[2] C. Cowan et al. "Detection of the free neutrino: a confirmation". In: *Science* 124.3212 (1956), pp. 103–104. ISSN: 0036-8075. DOI: 10.1126/science.124.3212.103.

[3] Z. Maki, M. Nakagawa, and S. Sakata. "Remarks on the unified model of elementary particles". In: *Progress of Theoretical Physics* 28 (1962), pp. 870–880. DOI: 10.1143/PTP.28.870.

[4] B. Pontecorvo. "Neutrino experiments and the problem of conservation of leptonic charge". In: *Soviet Physics JETP* 26 (1968), pp. 984–988.

[5] V. Gribov and B. Pontecorvo. "Neutrino astronomy and lepton charge". In: *Physics Letters B* 28.7 (Jan. 1969), pp. 493–496. DOI: 10.1016/0370-2693(69)90525-5.

[6] Ko Abe et al. "Constraint on the matter-antimatter symmetry-violating phase in neutrino oscillations". In: *arXiv preprint arXiv:1910.03887* (2019).

[7] Leigh H Whitehead. "Neutrino oscillations with MINOS and MINOS+". In: *Nuclear Physics B* 908 (2016), pp. 130–150.

[8] Carlo Giunti and Chung W Kim. *Fundamentals of neutrino physics and astrophysics*. Oxford university press, 2007.

[9] Pavel A Cherenkov. "Visible emission of clean liquids by action of γ radiation". In: *Dokl. Akad. Nauk SSSR*. Vol. 2. 8. 1934, pp. 451–454.

[10] P. Adamson et al. "CHerenkov detectors In mine PitS (CHIPS) letter of intent to FNAL". In: (2013). arXiv: 1307.5918 [physics.ins-det].

[11] A. Perch. "Construction of the CHIPS-M prototype and simulations of a 10 kiloton module". In: (2015). arXiv: 1505.00042 [physics.ins-det].

[12] M. Pfützner. "Prototype detection unit for the CHIPS experiment". In: *Journal of Physics: Conference Series* 888 (Sept. 2017), p. 012059. DOI: 10.1088/1742-6596/888/1/012059.

[13] M. Pfützner. "Sensitivity study and first prototype tests for the CHIPS neutrino detector R&D program". https://discovery.ucl.ac.uk/id/eprint/10052874/1/Pfutzner_thesis.pdf. PhD thesis. University College London, July 2018.

[14] Medbh Campbell. "Measuring neutrino oscillations in the NOvA and CHIPS Detectors". https://discovery.ucl.ac.uk/id/eprint/10097512/1/Campbell_10097512_thesis_sig-removed.pdf. PhD thesis. University College London, Aug. 2020.

[15] UF Katz, KM3NeT Consortium, et al. "Status of the KM3NeT project". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 602.1 (2009), pp. 40–46. DOI: 10.1016/j.nima.2008.12.215.

[16] A. Martinez et al. "Letter of intent for KM3NeT 2.0". In: *Journal of Physics G: Nuclear and Particle Physics* 43.8 (2016), p. 084001. DOI: 10.1088/0954-3899/43/8/084001.

[17] P. Adamson et al. "The NuMI neutrino beam". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 806 (Jan. 2016), pp. 279–306. ISSN: 0168-9002. DOI: 10.1016/j.nima.2015.08.063.

[18] Thomas H. Cormen. *Introduction to Algorithms, 3rd Edition (The MIT Press)*. The MIT Press, July 2009. ISBN: 0262033844. URL: `https://www.xarg.org/ref/a/0262033844/`.

[19] Donald E. Knuth. *The Art of Computer Programming: Volume 3: Sorting and Searching (2nd Edition)*. Addison-Wesley Professional, May 1998. ISBN: 0201896850. URL: `https://www.xarg.org/ref/a/0201896850/`.

[20] David G Beechy and Robert J Ducar. "Time and data distribution systems at the fermilab accelerator". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 247.1 (1986), pp. 231–238.

[21] D Ayres, NOvA Collaboration, et al. "NOvA proposal to build a 30 kiloton off-axis detector to study neutrino oscillations in the Fermilab NuMI beamline". In: *arXiv preprint hep-ex/0503053* (2005).

[22] A Norman et al. "The NOvA Timing System: A system for synchronizing a long baseline neutrino experiment". In: *Journal of Physics: Conference Series*. Vol. 396. 1. IOP Publishing. 2012, p. 012034.

[23] D. van Eijk. "Electronics and DAQ for the CHIPS experiment". In: (2018). arXiv: `1805.12206 [physics.ins-det]`.

[24] PPM Jansweijer, HZ Peek, and E De Wolf. "White Rabbit: Sub-nanosecond timing over Ethernet". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 725 (2013), pp. 187–190.

[25] *White Rabbit LEN*. `https://web.archive.org/web/20200219023827/http://sevensols.com/index.php/products/wr-len/`. Accessed: 2020-02-19.

[26] Ilka Antcheva et al. "ROOT—A C++ framework for petabyte data storage, statistical analysis and visualization". In: *Computer Physics Communications* 180.12 (2009), pp. 2499–2512.

[27] MA Unland Elorrieta et al. "Characterisation of the Hamamatsu R12199-01 HA MOD photomultiplier tube for low temperature applications". In: *Journal of Instrumentation* 14.03 (2019), P03015.

[28] Fred James and Matts Roos. "MINUIT: a system for function minimization and analysis of the parameter errors and corrections". In: *Comput. Phys. Commun.* 10.CERN-DD-75-20 (1975), pp. 343–367.