Dissertation
submitted to the
Combined Faculty of Mathematics, Engineering and Natural Sciences
of Heidelberg University, Germany
for the degree of
Doctor of Natural Sciences

Put forward by
Lorenzo Sanfilippo
born in: Desio, Italy
Oral examination: 29.10.2025

Testing the high-precision timing distribution for the ATLAS experiment at the HL-LHC

Referees: Prof. Dr. Hans-Christian Schultz-Coulon

Prof. Dr. Peter Fischer

Abstract

The High-Luminosity LHC (HL-LHC) at CERN will significantly increase the collision rate at CERN, presenting major challenges for detector operation, particle reconstruction, triggering and analysis. This thesis addresses the critical role of high-precision 40 MHz timing distribution in the ATLAS experiment, focusing on the development, implementation, and validation of the Local Trigger Interface (LTI) module—a key component of the Phase-2 Trigger and Data Acquisition (TDAQ) system upgrade.

We detail the architecture and operation of the LTI module, including its FPGA-based design, optical and electrical I/O interfaces, clocking system, and integration with ATLAS subsystems. Extensive studies were conducted to evaluate phase stability and mitigate timing uncertainties introduced by temperature variations, transceiver resets, and hardware-induced jitter. Using tools such as Digital Dual Mixer Time Difference (DDMTD) and Transmitter Phase Interpolator (TxPI), we demonstrate picosecond-level timing resolution under realistic operating conditions.

In addition to component-level validation under laboratory conditions, the LTI module was tested in integration setups with the ATLAS Liquid Argon Calorimeter and High-Granularity Timing Detector. Our results confirm the LTI's capability to provide deterministic, low-jitter timing across large-scale distributed systems such as the ATLAS detector. This work supports the broader effort to enable accurate event reconstruction and reliable triggering in the HL-LHC era.

Zusammenfassung

Der High-Luminosity LHC (HL-LHC) am CERN wird die Kollisionsrate drastisch erhöhen und stellt damit erhebliche Anforderungen an den Detektorbetrieb, die Teilchenrekonstruktion, das Triggersystem und die Analyse. Diese Doktorarbeit befasst sich mit der entscheidenden Rolle der hochpräzisen Verteilung des 40 MHz-Taktes im ATLAS-Experiment, mit Fokus auf die Entwicklung, Implementierung und Validierung des Local Trigger Interface (LTI) Moduls - einer zentralen Komponente des Phase-2 Upgrades des Trigger- und Datenerfassungssystems (TDAQ).

Die Architektur und Funktionsweise des LTI-Moduls werden im Detail beschrieben,

einschließlich des FPGA-basierten Designs, der optischen und elektrischen Schnittstellen, des Taktungssystems, sowie der Integration mit den ATLAS-Subdetektoren. Zur Bewertung der Phasenstabilität und zur Minderung von Timing-Unsicherheiten, die durch Temperaturänderungen, Transceiver-Resets und Hardware-Jitter verursacht werden, wurden umfangreiche Studien durchgeführt. Mit Werkzeugen wie der Digital Dual Mixer Time Difference (DDMTD) und dem Transmitter Phase Interpolator (TxPI) wird eine Zeitauflösung im Pikosekundenbereich unter realistischen Betriebsbedingungen demonstriert.

Neben der Validierung auf Komponentenebene in Laborumgebung wurde das LTI-Module auch in Integrationsumgebungen mit dem ATLAS-Flüssigargon-Kalorimeter und dem High-Granularity Timing Detector getestet. Die Ergebnisse bestätigen die Fähigkeit des LTI-Moduls, deterministischen und jitterarmen Takt in einem großen und verteilten System wie dem ATLAS-Detektor bereitzustellen. Diese Arbeit leistet einen wichtigen Beitrag zur präzisen Ereignisrekonstruktion und zu einem zuverlässigen Betrieb des Triggersystems am HL-LHC.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The European Organization for Nuclear Research (CERN) is an international research organization based in Geneva, Switzerland, dedicated to the study of fundamental particles and their interactions. Its primary scientific objective is to test and extend the Standard Model of particle physics, which currently provides the most comprehensive description of known elementary particles and three of the four fundamental forces: electromagnetic, weak, and strong interactions. The Standard Model does not incorporate gravity and is known to be incomplete [1].

At the center of CERN's experimental program is the Large Hadron Collider (LHC), a circular superconducting collider with a circumference of 27 km, located approximately 100 m underground along the France–Switzerland border. The LHC is capable of colliding proton–proton, proton–ion, and ion–ion beams[1] at the highest energies ever achieved in a laboratory setting. These collisions are studied with large-scale detectors to explore the properties of known particles and probe for potential signs of physics beyond the Standard Model [1].

There are four main interaction points along the LHC ring, each hosting a major experiment: ATLAS, CMS, ALICE, and LHC*b*. ATLAS and CMS are general-purpose detectors optimized for a wide range of physics goals, including precision measurements, Higgs boson studies, and searches for Physics Beyond the Standard Model (BSM). ALICE is designed to investigate the properties of strongly interacting matter under extreme conditions, such as those created in heavy-ion collisions, where a Quark–Gluon Plasma (QGP) is expected to form. LHC*b* focuses on flavor physics, particularly the study of heavy-flavor hadrons and CP violation in the quark sector.

The LHC ring is divided into eight sectors, numbered sequentially around the circumference, with the junctions between sectors referred to as "Points" (P1–P8). Four of these points host the main experiments: P1 (ATLAS), P2 (ALICE), P5 (CMS), and P8

[1]Heavy-ion runs at the LHC have primarily used fully stripped lead nuclei ($^{208}\mathrm{Pb}^{82+}$), while recent short runs have also included oxygen ($^{16}\mathrm{O}^{8+}$) and neon ($^{20}\mathrm{Ne}^{10+}$) beams for lighter-ion studies.

(LHC*b*), while others contain major accelerator systems and facilities.

The LHC began operations in 2010 and has delivered collisions in successive multi-year periods known as "runs". To sustain and improve detector performance in the face of increasing luminosity, and to meet the physics goals of the data-taking, a series of hardware and software improvements have been implemented. Major improvements are usually implemented during Long Shutdown periods, following every run, during which both the accelerator and the experiments undergo maintenance and upgrades. The most recent such program, known as the Phase-1 upgrades, was carried out during Long Shutdown 2 (2019–2022) and prepared all LHC experiments for the higher luminosities and more challenging conditions of Run-3. At the time of writing, the LHC is operating in Run-3, which is scheduled to conclude in mid-2026 and be followed by the Long Shutdown 3 [2].

The LHC is the final stage in CERN's accelerator complex, a series of interconnected machines that progressively accelerate and prepare the particle beams. The acceleration chain includes Linac4, the Proton Synchrotron Booster (PSB), the Proton Synchrotron (PS), and the Super Proton Synchrotron (SPS), which together deliver protons at 450 GeV to the LHC. A schematic overview of this complex is shown in Figure 1.1.

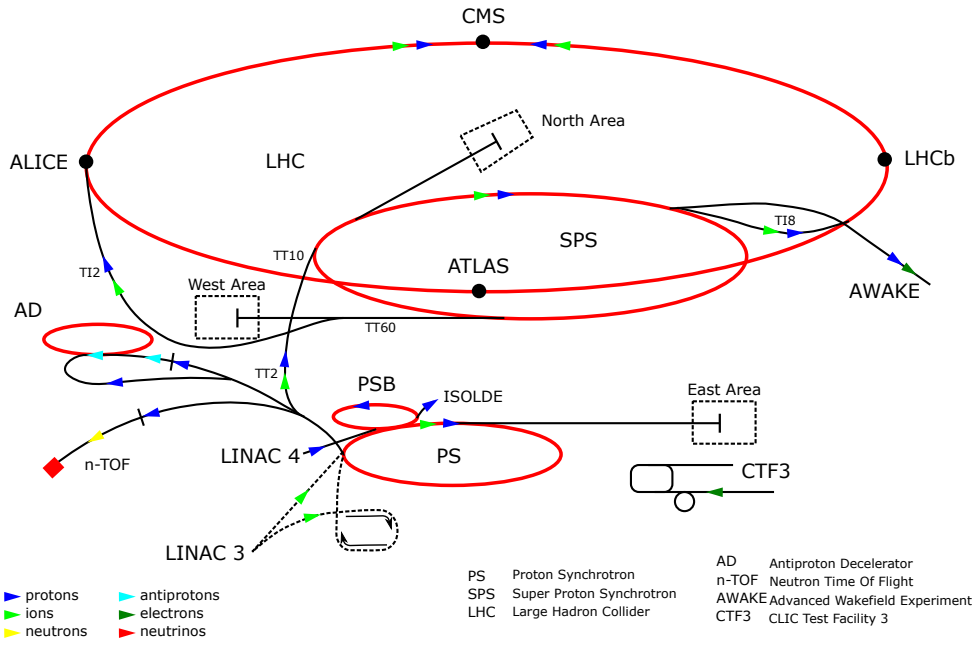

Figure 1.1: CERN's accelerator complex [3].

Inside the LHC, two counter-rotating proton beams circulate in ultra-high-vacuum beam pipes, guided by superconducting dipole magnets and focused in the transverse plane by quadrupole magnets. The longitudinal beam structure is composed of discrete proton packets known as bunches. Each bunch contains approximately 1.1×10^{11} protons, whose organization is governed by the superconducting radio-frequency (RF) system located at P4.

The LHC RF system operates at the master frequency of 400 MHz, which defines discrete RF buckets separated by 2.5 ns. In principle, each RF bucket can hold a bunch of protons. Even so, in nominal operating conditions the bunches can occupy every tenth bucket. This frequency division defines the 40 MHz **bunch clock**, and consecutively 3564 bunch slots spaced 25 ns. In addition to the bunch clock, the RF system also produces the once-per-turn **Orbit** signal, which has a period of 89 µs and a pulse width of 1 µs.

The bunch clock is the heartbeat of the LHC and its experiments. Almost everything is synchronous to the bunch clock: collisions, detector readout, front-end digitization, and data selection pipelines. This means that at any time the bunch clock (and the Orbit) has to be distributed everywhere with the best possible quality.

The bunch clock and Orbit signals for each beam are distributed through the Trigger, Timing and Control (TTC) backbone, a CERN-wide optical network that delivers the timing signals to the CERN Control Center (CCC), and to all LHC experiments. The backbone employs standard VME[2]-format RF optical transmitters and receivers, ensuring low-jitter transmission over long distances — with a typical output jitter of about 10 ps RMS at the experiment side [5]. The "jitter" refers to small, rapid variations in the timing of a signal or event. It is the deviation from the expected, regular timing of a signal or process. Figure 1.2 shows the TTC backbone distribution from P4 to the main experiments, passing from the CCC.

In the ATLAS experiment, the TTC backbone signals are received in the counting room, which hosts the off-detector electronics next to the main experimental cavern. An electronics module recovers the clocks and Orbit markers with a jitter RMS of about 11 ps [6] before distributing them via the ATLAS-specific TTC network to all front-end systems. More information on the connection between global-TTC and ATLAS-TTC will be provided in Section 2.2.7.

The synchronization of ATLAS with the bunch clock is affected by the specific filling scheme of the LHC beams. The filling scheme determines which RF buckets are populated and which are left empty. The filling scheme represents a trade-off between maximizing collision rates, machine-protection requirements and overlapping interactions within the detector. Figure 1.3 shows a typical filling scheme.

After injection at 450 GeV, the beams are accelerated to a top energy of 6.8 TeV per beam, corresponding to a center-of-mass energy $\sqrt{s} = 13.6$ TeV; at this point the beams are focused and brought into collision. In general the detectors are kept recording as many physics events as possible, however only the events corresponding to LHC stable beams are used for physics analyses. With the beam configuration in place and stable collisions achieved, attention turns to the physics processes that these collisions can produce and how they are recorded.

[2]Versa Module Eurocard [4]

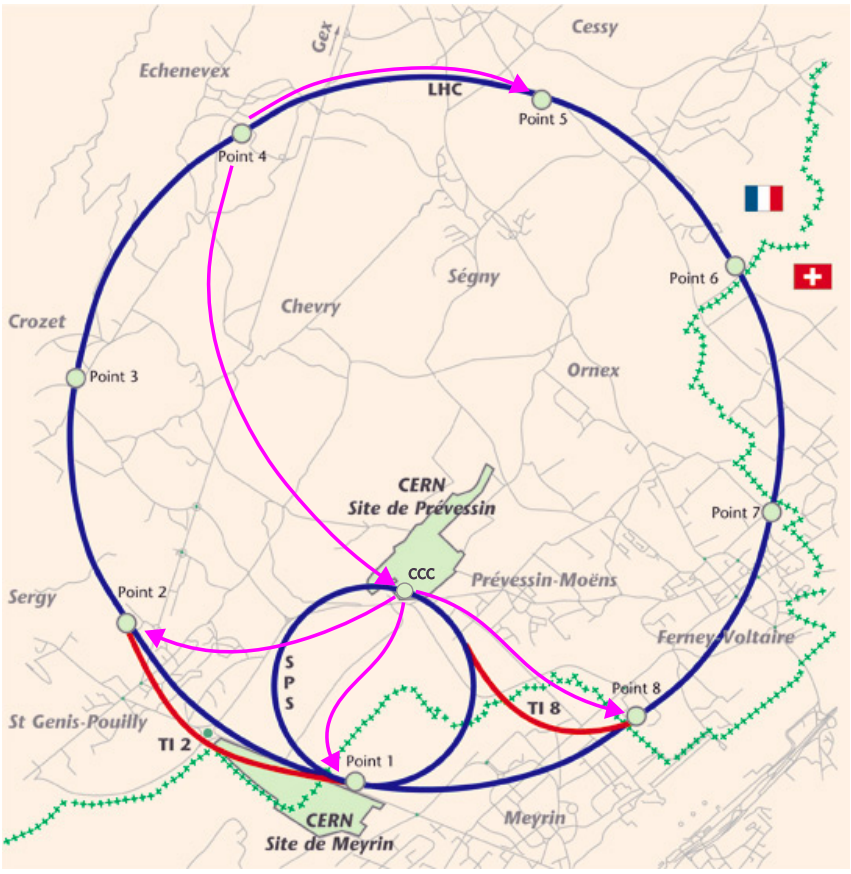

Figure 1.2: TTC backbone distribution (highlighted in magenta). Starting from the LHC RF at P4, passing through the CCC, and to the main experiments. Only the CMS experiment at P5 receives timing signals directly from P4.

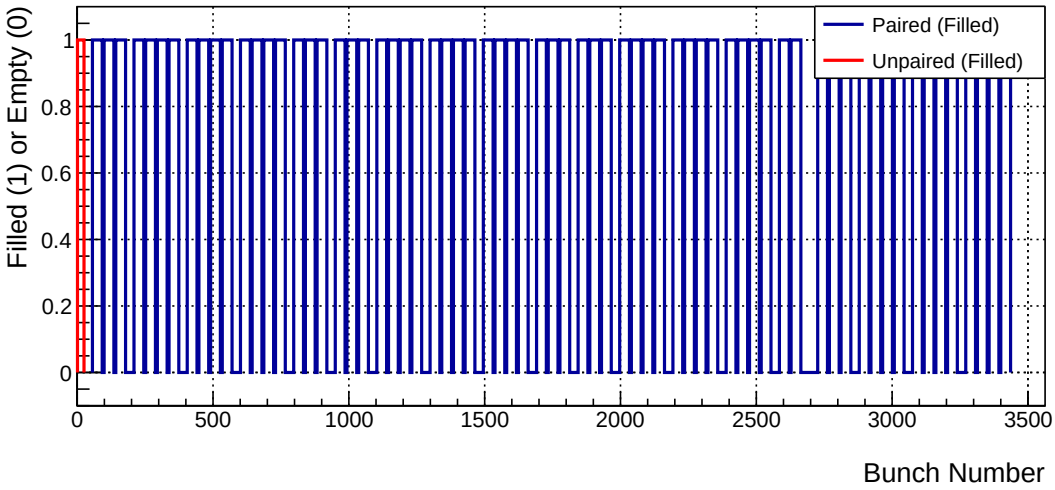

Figure 1.3: LHC filling scheme for a typical Run-3 configuration. The horizontal axis shows the bunch slot index, while the vertical axis indicates whether the slot is filled.

Since only a small fraction of the millions of collisions occurring every second can be recorded, experiments at the LHC rely on fast and efficient data selection systems. In the case of ATLAS, a two-level trigger architecture is used: the hardware-based Level-1 (L1) Trigger reduces the event rate from 40 MHz, corresponding to a 25 ns bunch spacing. Front-end electronics are required to sample detector signals at precise intervals, with timing alignment typically accurate to within a few nanoseconds. Even small deviations can result in signals being assigned to the wrong bunch crossing, degrading both triggering efficiency and offline reconstruction performance.

The interpretation of detector signals in a collider experiment like ATLAS relies on accurate temporal alignment with the LHC's bunch structure. Sub-Detectors such as calorimeters and muon chambers produce analog signals with characteristic pulse shapes and time profiles. To reconstruct these signals correctly, front-end electronics must sample them at specific, phase-locked intervals relative to the 40 MHz bunch clock. Even nanosecond-level misalignments can distort the signal shape, lead to incorrect energy reconstruction, or cause hits to be assigned to the wrong bunch crossing.

The Timing, Trigger, and Control (TTC) system is responsible for distributing this clock, along with synchronous control signals, to thousands of front-end modules throughout the detector. It ensures that all channels operate within a coherent global timing framework, enabling accurate signal digitization, consistent trigger decisions, and proper event reconstruction. Low-jitter, phase-stable timing distribution is therefore essential to maintaining detector performance and data integrity, even under standard LHC running conditions.

This requirement is especially stringent for sub-detectors with long signal integration times. In calorimeters, for example, analog pulses can span multiple bunch crossings. The ATLAS liquid-argon (LAr) calorimeter uses bipolar shaping, with signals that peak approximately 50 ns after the interaction and extend beyond 400 ns. Sampling must occur at precisely calibrated phases relative to the bunch clock. A timing offset on the order of 1 ns can significantly affect the measured energy and cause energy deposits to be assigned to the wrong vertex in high pile-up conditions [7]. This is illustrated in Figure 1.4, which shows the shaped LAr pulse and sampling points.

Muon systems are similarly sensitive. Muons typically traverse several meters from the interaction point to the outermost detector layers, and their arrival time depends on both geometry and momentum. Timing calibration ensures that hits in drift tubes or resistive plate chambers are correctly assigned to their bunch crossing. Offsets of just a few nanoseconds can result in misassigned or missed hits. In drift-based detectors, such misalignment affects the conversion from drift time to spatial coordinate, which degrades track and momentum reconstruction.

The role of precise timing is particularly evident in the performance of ATLAS muon

drift tube (MDT) chambers. These detectors measure the distance of a muon from the wire based on the time required for ionization electrons to drift across the gas volume. This drift time must be calibrated against the global TTC-distributed clock. In Run-2, after calibrating individual tube offsets, an average single-hit spatial resolution of 81.7 ± 2.2 µm was achieved across the MDT system [8]. Given a typical drift velocity of 30 µm/ns, this level of spatial resolution corresponds to timing precision at the 1–2 ns level. Maintaining such synchronization is essential for accurate hit reconstruction and reliable muon tracking. Figure 1.5 shows the spatial resolution as a function of drift radius.

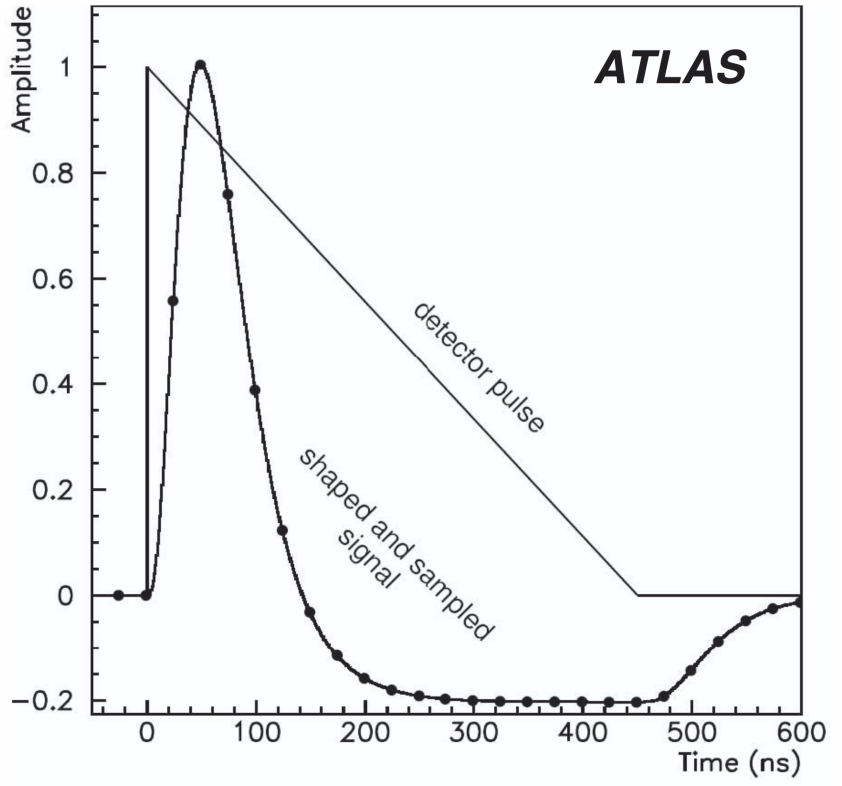

Figure 1.4: Shaped bipolar signal from the ATLAS liquid-argon calorimeter [7].

These challenges become more severe at high luminosity, when many overlapping interactions occur within the same bunch crossing. Accurate timing is not only required to assign signals to the correct bunch crossing, but also to suppress pile-up. Maintaining precise timing across all detector systems is therefore not a technical detail, but it is essential to achieving high-quality reconstruction and efficient triggering. Upcoming detector upgrades aim to enhance both in-time and out-of-time pile-up rejection by incorporating high-resolution timing layers with precision below 50 ps, enabling four-dimensional vertex reconstruction.

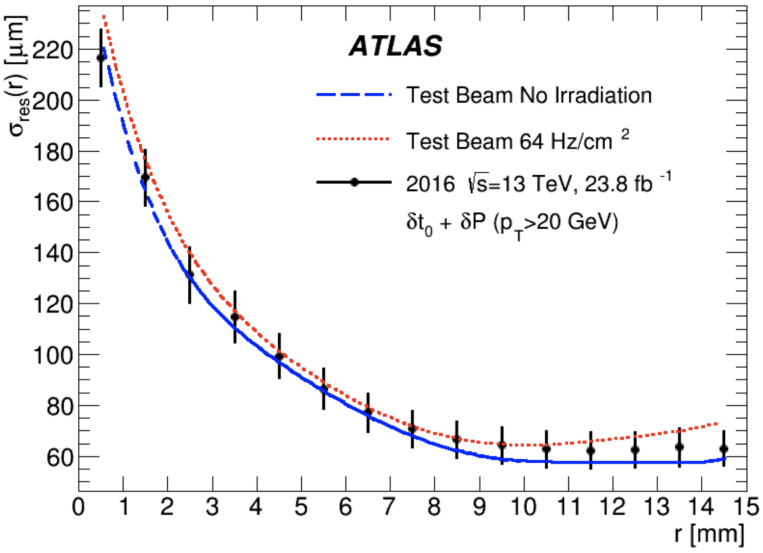

Figure 1.5: Single-hit spatial resolution of ATLAS MDT chambers as a function of drift radius, based on 2016 Run-2 data [8].

1.1 High-Luminosity LHC

The end of LHC Run-3, scheduled for mid-2026, will be followed by the Long Shutdown 3 (LS3), extending until mid-2030. During this period the Phase-2 upgrade program will be implemented in preparation for the High-Luminosity LHC (HL-LHC) phase. The HL-LHC aims to increase the LHC's instantaneous luminosity by a factor of five to seven, reaching up to $7.5 \times 10^{34}\text{cm}^{-2}\text{s}^{-1}$, allowing for the collection of a much larger dataset and improving sensitivity to rare physics processes. However, this increase in performance also presents major challenges for the experiments in terms of data volume, radiation tolerance, and the need for tighter timing and trigger control. The increased luminosity will lead to an average of $\langle\mu\rangle = 200$ simultaneous inelastic collisions per bunch crossing, which significantly complicates event reconstruction and analysis.

At these pile-up levels, vertex reconstruction and track-to-vertex association become increasingly challenging, particularly in the forward region. While the upgraded Inner Tracker (ITk) provides excellent performance in the central region, its longitudinal impact parameter resolution degrades at high pseudorapidity, exceeding 200 µm at $|\eta| > 2.5$ [9]. To maintain efficient pile-up rejection in the forward region, the High-Granularity Timing Detector (HGTD) will be installed in the range $2.4 < |\eta| < 4.0$, providing a per-track time resolution of approximately 30 ps [10]. This enables four-dimensional vertexing[3], improving the association of tracks and calorimeter deposits to the correct primary in-

[3]Four-dimensional vertexing refers to the process of reconstructing collision vertices not just in space, but also in time.

teraction. Achieving this timing precision requires a low-jitter clock distribution system. The Local Trigger Interface (LTI) will be the new ATLAS module distributing the experiment's master clock signal. Its function is essential to ensure that high-resolution timing detectors such as the HGTD receive a phase-stable, low-jitter reference signal, as any fluctuations in the distributed clock would directly degrade timing performance and reduce the efficacy of 4D vertexing in the high pile-up environment of the HL-LHC.

More details about the ATLAS upgrades in view of the HL-LHC will be given in Chapter 3.

1.2 Author's Contributions

The work presented in this thesis was carried out within the ATLAS Collaboration during Run-3 and in preparation for the ATLAS HL-LHC upgrade.

The author played a central role in the development of the Local Trigger Interface (LTI), a key component responsible for distributing timing and trigger signals to ATLAS sub-detectors in the HL-LHC era. A central focus of this work was the study of clock phase stability: evaluation kits and prototypes were tested to identify potential sources of instability, quantify their impact on timing performance, and investigate mitigation strategies for long-term drifts caused by temperature variations within the LTI and its environment. In addition, the author was responsible for the configuration and implementation of the LTI control software and actively defined its operational logic and integration strategies.

Beyond upgrade activities, the author was responsible for the operation of the Run-3 ATLAS Central Trigger system, a core element of ATLAS data-taking. This work included on-call expert support, maintenance of the timing infrastructure, and the development and upkeep of monitoring tools. In particular, the author developed and maintained a Central Trigger monitoring package that ensured the stability and reliability of data-taking conditions.

The author's ATLAS Qualification Task (a prerequisite for full authorship within the collaboration) was also carried out in the Central Trigger system. It focused specifically on monitoring and data quality, and its successful completion granted the author full ATLAS authorship. Together, these efforts span both upgrade development and Run-3 operations, supporting the technical infrastructure, reliability, and evolution of the ATLAS trigger and timing systems in preparation for the HL-LHC.

As a member of the ATLAS Collaboration, the author has been included in the author list of all official ATLAS physics publications since January 2024. In addition, the author was co-author of selected ATLAS internal and public notes within the Level-1 Central Trigger group. While the author did not participate directly to the writing of these

documents, technical and operational work described in this thesis was incorporated into the results they present. The most relevant notes are:

- *Integration and commissioning of the ATLAS Muon-to-Central-Trigger-Processor Interface for Run-3* [11], ATL-COM-DAQ-2021-083, CERN (2021). This note presented work the author carried out as a CERN Technical Student. Although the publication preceded the start of the doctoral program, the author continued related integration and operational activities on this system throughout the doctoral research.

- *Phase-II Upgrade of the ATLAS L1 Central Trigger* [12], ATL-DAQ-PROC-2024-007, CERN (2024). While the note itself describes the high-level design and status of the Phase-2 upgrade for the Central Trigger, the author's contributions were at the technical and operational level supporting this broader activity.

- *High-Precision Timing Distribution in the ATLAS Phase-2 upgrade* [13], CERN (2024). During the poster session of the October 2024 ATLAS week the author presented the status of my studies on clock phase stability and mitigation procedures with LTI evaluation kits. This poster presentation marked the first internal ATLAS publication on this topic.

Chapter 2

The ATLAS experiment

This chapter provides an overview of the ATLAS Experiment. The first part describes briefly the main ATLAS sub-detectors, while the second part introduces the Run-3 ATLAS Trigger and Data Acquisition (TDAQ) system.

2.1 The ATLAS detector

ATLAS is a general-purpose particle detector designed to investigate a wide range of physics phenomena resulting from high-energy proton–proton, ion–ion or proton–ion collisions. It is built with a forward–backward symmetric cylindrical geometry and nearly complete coverage in solid angle. The detector is composed of several concentric subsystems, each optimized to measure specific properties of particles. For each selected bunch crossing, the full event from the detector readout is reconstructed. This includes all the collisions that occurred within that 25 ns interval. Typically, only one of these collisions, referred to as the primary hard-scatter interaction, is of interest for physics analyses. The goal is to assign tracks to the primary vertex so that reconstructed objects originate from the hard scattering. For this reason pile-up interactions are also reconstructed and must be taken into account [14].

Surrounding the collision point from the inside out, the sub-detectors are:

- Inner detector for measurements of charge, trajectories and momentum of charged particles. The inner detector is enclosed by a solenoidal magnet.

- Electromagnetic and hadronic calorimeters for energy measurements of electrons, taus, photons and hadrons.

- Muon spectrometer for identification and momentum determination of muons. The muon spectrometer operates within a large air-core toroidal magnet system.

The detector measures approximately 46 m in length and 25 m in height, with a total mass of about 7×10^6 kg [14]. A schematic longitudinal cut-away view of the detector is shown in Figure 2.1.

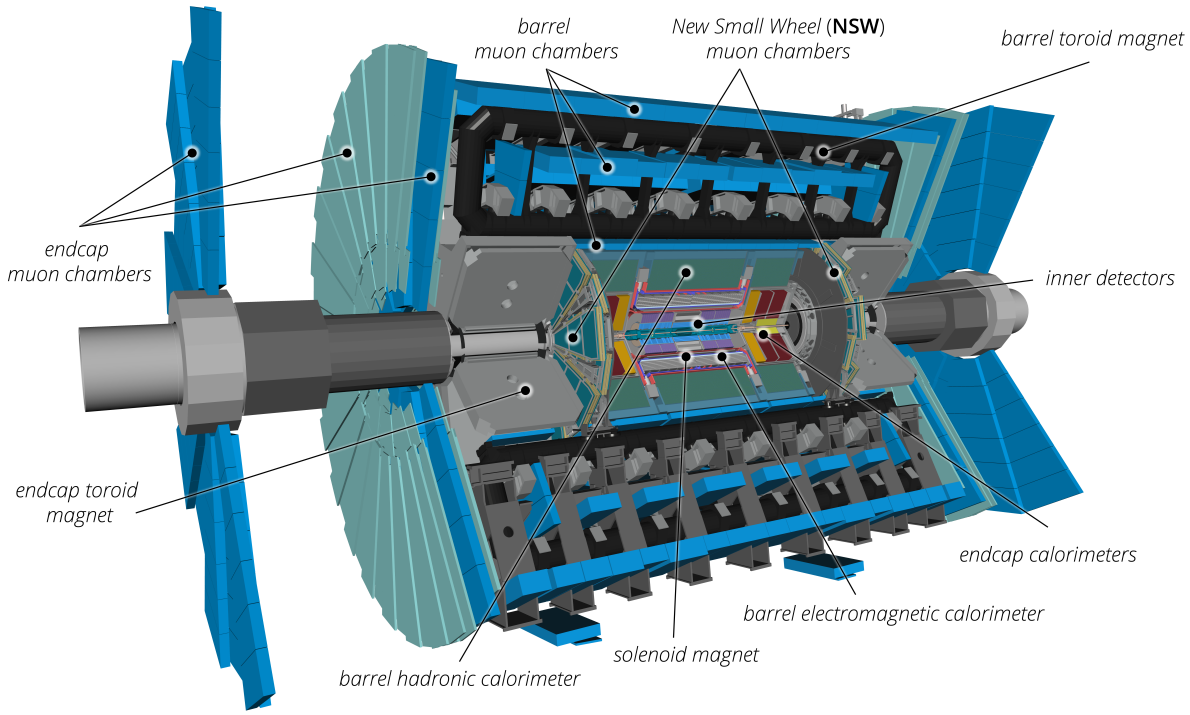

Figure 2.1: Cut-away view of the ATLAS detector, illustrating the main sub-detectors arranged in a layered structure around the interaction point [15].

In addition to its diverse sub-detectors, ATLAS relies on a Trigger and Data Acquisition system to decide, in real time, which of the bunch crossings are stored for analysis. The TDAQ operates in two stages: a fixed-latency hardware-based Level-1 (L1) trigger and a software-based High-Level Trigger (HLT). The L1 trigger must deliver a decision within a maximum latency of 2.5 µs, reducing the event rate from 40 MHz to about 100 kHz. Events accepted by the L1 Trigger are fully read out and processed by the HLT, which applies refined reconstruction algorithms to bring the output rate down to around 3 kHz for permanent storage.

The diverse range of physics goals pursued at the LHC imposes several stringent design requirements on the ATLAS detector [14]:

- Fast and radiation-hard electronics to withstand high radiation levels and allow real-time signal processing.

- High spatial and temporal granularity to handle intense particle fluxes and mitigate event pile-up.

- Precision muon momentum measurements across a broad range of momenta, along with reliable charge determination.

- Excellent electromagnetic calorimetry for identifying and measuring photons and electrons, along with robust hadronic calorimetry for accurate jet reconstruction and missing transverse energy (MET) measurements.

- Efficient tracking and vertex reconstruction for charged particles, even under high-luminosity conditions.

- Wide pseudorapidity acceptance and full azimuthal coverage to maximize detection efficiency and transverse energy reconstruction.

2.1.1 Coordinate system

The coordinate system used in ATLAS is a right-handed Cartesian system with its origin at the center of the detector, located at the nominal interaction point [14]. The positive x-axis points towards the center of the LHC ring, while the z-axis is aligned along the beam direction. The positive y-axis points upwards, though it is slightly tilted with respect to the vertical due to the overall inclination of the LHC tunnel.

In the transverse plane, cylindrical coordinates (r, ϕ) are used, where $\phi \in [-\pi, \pi]$ is the azimuthal angle measured around the beam axis from the x-axis, and r is the radial distance from the interaction point [14]. The polar angle $\theta \in [0, \pi)$ is measured from the positive z-axis, but it is more commonly expressed in terms of the pseudorapidity η, which is defined as:

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right]. \tag{2.1}$$

Regions of the detector close to the beam axis are often referred to as "forward" (for $|\eta| \gtrsim 2.5$). Additionally, the side of the detector where $z > 0$ is called side A, and the side where $z < 0$ is referred to as side C.

For distance measurements in the (η, ϕ) space, the separation variable ΔR is defined as:

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}. \tag{2.2}$$

This variable is dimensionless and represents the angular distance between two objects in the detector, combining differences in pseudorapidity and azimuth.

2.1.2 Magnets system

Achieving precise reconstruction of charged particle momenta in ATLAS, especially at the multi-TeV scale, requires a strong and well-shaped magnetic field over large detector volumes. The magnetic field across the ATLAS detector is produced by four systems of superconducting magnets: one solenoid, one barrel toroid, and two endcap toroids [14]. The cryogenic infrastructure supplies both liquid and gaseous helium at temperatures of 4.5 K and 60 K respectively, to cool the superconducting coils and to maintain thermal shielding [14]. The configuration of the magnet systems is shown in Figure 2.2.

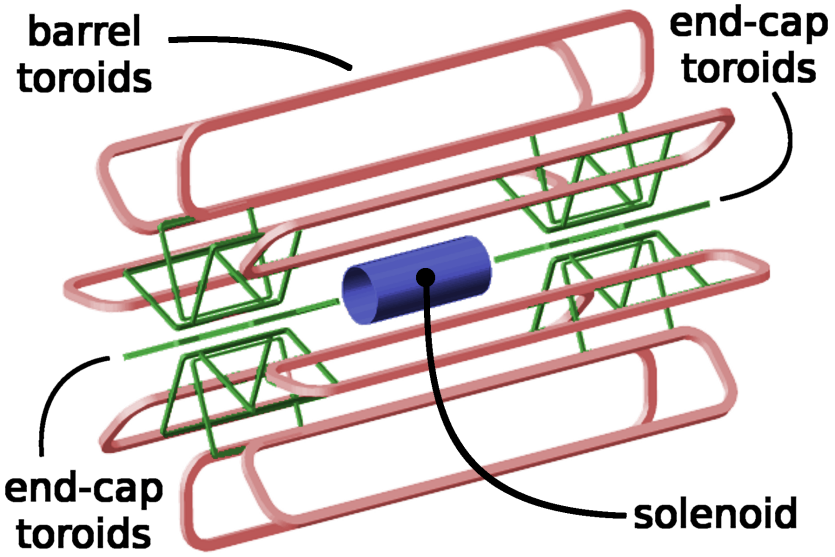

Figure 2.2: Layout of ATLAS superconducting magnet systems [16].

The choice of magnet technologies and geometry was driven by the need for large active volumes and strong magnetic fields. These parameters are critical for achieving good momentum resolution. To first order, the transverse momentum resolution scales as $\sigma_{p_T}/p_T \propto p_T/(BL^2)$, where B is the magnetic field strength and L is the tracking length [17].

The solenoid magnet is aligned coaxially with the beam line and surrounds the Inner Detector. Located immediately outside the Inner Detector volume and inside the cryostat of the electromagnetic calorimeter barrel, it generates an axial magnetic field of approximately 2 T used for inner tracking. Because the solenoid is located in front of the calorimeter, its design had to minimize the amount of dead material which degrades the calorimeter energy resolution. For this reason, it was constructed using Al-stabilized NbTi with a thickness of only 4.5 cm, corresponding to about 0.66 radiation lengths [16].

The endcap toroid magnets generate magnetic fields of up to 3.5 T in the regions surrounding the endcap calorimeters. Their distinctive gear-like structure allows them to interlock with the barrel toroid at an angle of 22.5°, improving the field coverage and

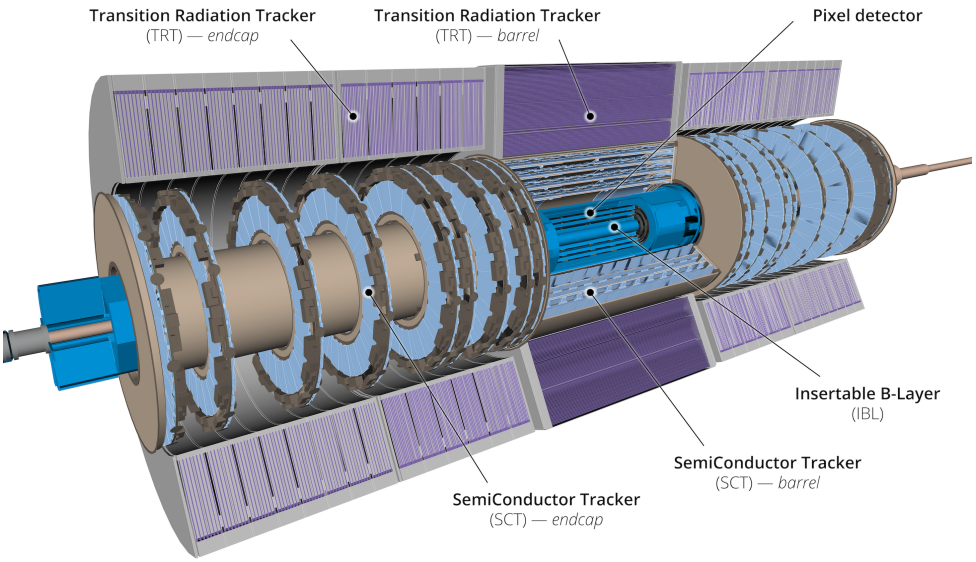

Figure 2.3: Cut-away view of the ATLAS inner detector [15].

providing a smooth transition between the barrel and endcap regions [14].

The typical bending power of the magnetic system is around 3 T m in the barrel and up to 6 T m in the endcap regions [16]. Due to the finite number of coils, the magnetic field deviates from an ideal toroidal shape and exhibits a periodic ripple pattern in azimuth. As a result, the bending power experienced by a particle depends on its initial position, direction, charge, and momentum. These local variations degrade the momentum resolution, particularly in the barrel–endcap transition region, where a sizable radial field component is present [14].

2.1.3 Inner Detector

The inner detector (ID) is the innermost tracking system of ATLAS, surrounding the beamline in a high-radiation environment. It operates in the axial magnetic field provided by the central solenoid (see section 2.1.2). The ID provides tracking coverage up to $|\eta| < 2.5$ and is used to reconstruct primary and secondary vertices [15].

It consists of three sub-detectors arranged in concentric cylindrical layers: the pixel detector (PIX), the semiconductor tracker (SCT), and the transition radiation tracker (TRT). The full system is contained in a cylindrical volume 2.1 m in diameter and 6.2 m in length [15]. The inner detector is shown in Figure 2.3.

The inner detector information is not used in the L1 trigger. Real-time track reconstruction at L1 would require processing the full inner detector hit pattern within the 2.5 µs latency window. The inner detector tracks are used in the high-level trigger and offline reconstruction.

Pixel Detector (PIX): The PIX consists of four barrel layers of silicon pixel sensors with a total of about 92 million readout channels. The innermost layer, the insertable B-layer (IBL), was installed during Long Shutdown 1 after Run-1 to maintain performance under higher pile-up conditions [18]. Pixel pitch in the IBL is 50 µm × 250 µm in the $r\phi$ and z directions, respectively. In the other three layers the pitch is 50 µm × 400 µm. The spatial resolution is approximately 8 µm in $r\phi$ and 40 µm in z for the IBL, and 10 µm in $r\phi$ and 115 µm in z for the other layers.

Semiconductor Tracker (SCT): The SCT is located outside the PIX and consists of four barrel layers and nine disks in each endcap. Each module contains two single-sided silicon microstrip sensors mounted back-to-back at a stereo angle of about 40 mrad. The strips on both sensors are nearly parallel to z, but the stereo tilt causes them to shift sideways in $r\phi$ as z increases. A track crossing the module therefore produces hits on different strip numbers in the two sensors, and the offset is proportional to the z-position of the hit [19]. This geometry allows the reconstruction of both the $r\phi$ coordinate and the second coordinate (z in the barrel, r in the endcaps) from a single module. In total, the SCT contains about 4000 modules with roughly six million strips. The spatial resolution is about 16 µm in $r\phi$ and 580 µm in z. Two different tracks can be resolved if they are distant more than 200 µm from each other [19].

Transition Radiation Tracker (TRT): The TRT is the outermost ID sub-detector and consists of straw drift tubes arranged in the barrel and endcap regions. Each straw is a 4 mm diameter cylindrical drift chamber with a central gold-plated tungsten anode wire, operated at a cathode potential of −1.5 kV [14]. A charged particle crossing a straw ionizes the gas, and the electron drift time to the anode is used to determine the track's radial distance from the wire. Since Run 3, the TRT straw tubes have been operated with an argon-based gas mixture (70% Ar, 27% CO_2, 3% O_2) [15], replacing the previous xenon-based mixture due to its high cost and limited availability [15]. The TRT provides about 30 tracking measurements per charged particle, with an intrinsic resolution of 120 µm in $r\phi$. In addition, the TRT read-out measures the time-over-threshold for each straw signal, which is proportional to the energy deposited in the gas. Combining this information across many straws yields an estimate of the track's average energy loss (dE/dx). Layers of polypropylene radiators are interleaved between straw planes to increase the production of transition-radiation photons from highly relativistic particles. Transition-radiation deposits are typically 8–10 keV per hit for electrons and about 2 keV for pions, enabling electron identification [20].

2.1.4 Calorimeter System

Calorimeters are detectors that measure the energy of particles by causing them to interact with a dense material, which initiates a cascade of secondary particles known as a particle shower. The total energy of the primary particle is then inferred by measuring the energy deposited by the shower within an active medium. This is in contrast to tracking detectors, which measure particle momentum by observing their trajectories. The composition and shape of the particle shower depend on the type and energy of the primary particle ($e^{\pm}$, γ, or hadrons). This distinction motivates two main applications of calorimeters: (i) electromagnetic calorimeters and (ii) hadronic calorimeters.

Electromagnetic calorimeters are optimized to measure electrons and photons, which develop showers dominated by bremsstrahlung and pair production. Hadronic calorimeters are designed to measure hadrons, such as protons and neutrons, whose showers are initiated primarily through strong interactions. In practice, hadronic showers also contain a sizable electromagnetic component, arising mostly from the decay of neutral pions ($\pi^0 \to \gamma\gamma$). Consequently, a fraction of the hadronic shower energy is deposited in the electromagnetic calorimeter, while the remainder is absorbed in the hadronic calorimeter. Compared to electromagnetic showers, hadronic showers are typically more complex and spatially larger, requiring greater material depth for full containment. Their development involves a broad range of particles and processes, including the production of slow neutrons, which may travel long distances before depositing their energy.

Beyond their physics application, calorimeters can also be classified by their structure. They can be homogeneous—constructed from a single material that both initiates and detects showers—or sampling calorimeters, which consist of alternating absorber and active layers. The absorber, often made of high-Z materials such as lead, tungsten, or steel, initiates repeated electromagnetic or hadronic interactions that produce cascades of secondary particles. The active medium converts the passage of shower particles into a measurable signal: in liquid argon calorimeters this is done via ionization of the liquid argon (LAr), producing a drift current, while in scintillator calorimeters it is done via scintillation light collected by optical fibers.

The ATLAS calorimeter system consists of two main concentric structures with nearly full azimuthal coverage around the beam axis: the Liquid Argon (LAr) calorimeter and the Tile calorimeter, as shown in Figure 2.4.

Spatial segmentation in calorimeters allows measurement of the direction and position of showers. Calorimeter measurements also provide the basis for calculating MET, which is essential for identifying non-interacting particles such as neutrinos.

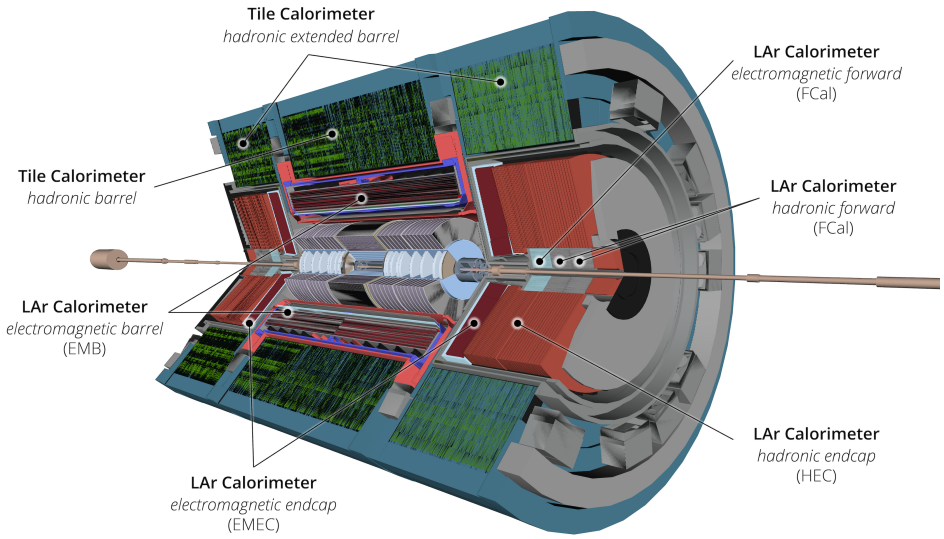

Figure 2.4: Cut-away view of the ATLAS calorimeter system [15].

Liquid Argon (LAr) Calorimeters

The ATLAS liquid argon (LAr) calorimeters are sampling calorimeters that use liquid argon as the active medium and dense metals as absorbers. Liquid argon is chosen for its excellent radiation tolerance and stable, linear response. Operation requires cryogenic cooling to about 87 K inside hermetically sealed cryostats to ensure signal quality and avoid contamination [7].

The system includes four main sections:

- Electromagnetic Barrel: covers the central region up to $|\eta| < 1.475$.
- Electromagnetic Endcaps: extend coverage to $1.375 < |\eta| < 3.2$.
- Hadronic Endcaps: positioned behind the electromagnetic endcaps and covering the same η range.
- Forward Calorimeter: extends coverage to the forward region $3.1 < |\eta| < 4.9$.

The electromagnetic barrel and endcap calorimeters use an accordion geometry to provide continuous azimuthal coverage. Copper electrodes are embedded between lead absorber plates, dividing the argon gap into two half-gaps for bipolar signal readout. The accordion geometry also facilitates a fast signal readout by providing a short signal path to the front-end electronics located at the detector's inner and outer radii. The electromagnetic sections have a total thickness exceeding 22 radiation lengths in the barrel and 24 radiation lengths in the endcaps, ensuring full containment of electromagnetic showers.

The hadronic endcap calorimeters consist of two independent wheels per endcap, segmented longitudinally by copper plates of 25 mm (inner wheel) and 50 mm (outer wheel) thickness. Combined with the electromagnetic calorimeter in front, the total depth corresponds to about 10 nuclear interaction lengths, which is sufficient to contain most high-energy hadronic showers in the endcap region for ATLAS physics requirements [7].

The forward calorimeters operate in the high-radiation forward region. They provide an almost full η coverage as needed to measure E_T^{miss}, and help to shield the muon spectrometer. The absorbers are in copper for the electromagnetic section and in tungsten for the hadronic section. Each forward calorimeter module contains three longitudinal segments: the innermost copper section optimized for electromagnetic energy measurements, and two tungsten sections optimized for hadronic energy.

The bipolar ionization signals from the LAr calorimeters are shaped and amplified in the front-end boards. Shaping times of approximately 400 ns optimize noise performance while limiting pile-up sensitivity. The shaped signals are sampled every 25 ns and stored in pipeline memories deep enough to accommodate the 2.5 µs L1 trigger latency.

Starting in Run-3, the front-end boards on the trigger path perform real-time digital processing to form trigger primitives, namely a summary of energy collected over a specific, predefined region of the calorimeter. The granularity of these primitives was improved in the barrel electromagnetic layers, with layer-specific energy readout enabling improved shower-shape discrimination [15]. These digital sums provide the input to the calorimeter-based trigger selection system, described later in Section 2.2.1.

Tile Calorimeter

The Tile calorimeter, located immediately outside the electromagnetic barrel, covers the central region up to $|\eta| < 1.7$. It consists of a barrel section and two extended barrels, segmented azimuthally into 64 modules. Radially, the calorimeter is divided into three layers. The system provides a total depth of approximately 9.7 λ at $\eta = 0$, ensuring sufficient energy containment for hadronic showers.

The Tile calorimeter is also a sampling calorimeter, using plastic scintillating tiles as the active medium and iron plates as absorbers. Scintillation light produced in the tiles is collected via wavelength-shifting fibers and directed to photomultiplier tubes.

Photomultiplier signals are shaped with a fast bipolar shaper ($\tau \approx 50$ ns) to suppress out-of-time pile-up, then digitized every 25 ns and stored in front-end pipeline memory pending the trigger decision.

In parallel, a low-latency analog summation path combines photomultiplier signals into fixed $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ trigger-tower sums. These analog sums are transmitted and digitized in preprocessor modules for use in the calorimeter trigger system.

2.1.5 Muon Spectrometer

The muon spectrometer is the outermost sub-detector of ATLAS, designed to identify muons and measure their momentum with high precision. It provides standalone momentum measurements in the pseudorapidity range $|\eta| < 2.7$, and contributes to the L1 muon trigger up to $|\eta| < 2.4$. The system achieves a momentum resolution of about 3% at $p_T = 100\,\text{GeV}$ and ∼10% at $p_T = 1\,\text{TeV}$ [21].

Muon momentum is determined from the sagitta s of the curved trajectory in the magnetic field generated by the barrel and endcap toroid magnets (see subsection 2.1.2). The sagitta is the maximum perpendicular distance between the reconstructed track and the straight-line chord connecting its entry and exit points in the magnetic field. The toroid field strength is typically sufficient to produce measurable curvature even for high-p_T muons (e.g. sagitta of about 0.5 mm for p_T =1 TeV compared to a ∼50 µm measurement precision) [21, 14].

The spectrometer combines precision tracking chambers with fast trigger detectors:

- Precision tracking: Monitored Drift Tubes (MDTs) throughout most of the spectrometer, complemented by Micromegas (MM) in the forward region.
- Trigger detectors: Resistive Plate Chambers (RPCs) in the barrel, Thin Gap Chambers (TGCs) in the endcaps, and small-strip TGCs (sTGCs) in the upgraded forward region.

The barrel region contains three concentric cylindrical layers of detectors with an eight-fold ϕ symmetry, split into 16 sectors (eight large, eight small). In the endcaps, detectors are arranged into the "Small Wheel" (innermost) and "Big Wheel" (outermost) structures. Each set of layers with similar geometry is referred to as a station. Starting from Run-3 the New Small Wheel replaces the innermost forward muon station in the region $1.3 < |\eta| < 2.7$ with detectors capable of both precision tracking and fast triggering under high background conditions. It integrates MM for precision tracking and sTGCs for triggering [15]. This combined system allows the New Small Wheel to deliver triggers with improved spatial and timing resolution, essential for suppressing beam-related background and maintaining high trigger efficiency in the forward region. Figure 2.5 shows the ATLAS muon spectrometer system.

Precision Tracking Detectors

- **Monitored Drift Tubes** – MDTs provide high-precision tracking across most of the spectrometer. Each drift tube is a 30 mm diameter aluminum cylinder filled with 93% Ar and 7% CO_2 at 3 bar pressure [15]. A 50 µm gold-plated tungsten–rhenium wire runs along the central axis as the anode, with the tube wall serving as the

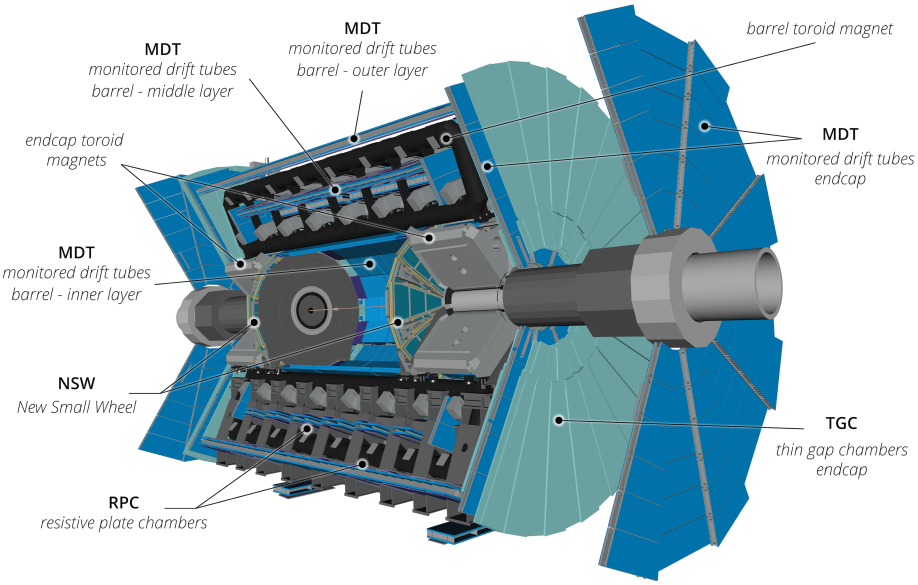

Figure 2.5: Cut-away view of the ATLAS muon spectrometer [15].

cathode [21]. Electrons from ionization drift toward the wire, with a maximum drift time of about 700 ns [22]. MDTs are not used in the L1 trigger due to this relatively long drift time, but are essential for offline momentum reconstruction and for refining muon candidates at the HLT. The Phase-I upgrade introduced small MDTs (sMDTs) with 15 mm tube diameter, halving drift times and increasing granularity [15].

- **Micromegas** – Micromegas detectors are parallel-plate gaseous chambers with separate drift and amplification regions. They achieve typical spatial resolutions of 100 µm and can handle hit rates up to 15 kHz/cm^2 [23]. The drift gap (5 mm) is operated at 600 V/cm, while the amplification gap (128 µm) operates at 40 to 50 kV/cm [23]. They provide a timing resolution of 5 ns to 10 ns, sufficient to reject of out-of-time backgrounds and to ensure time-coincidence with the TGC trigger signal.

  The micromesh between the drift and amplification gaps blocks the majority of avalanche ions, resulting in an ion backflow fraction of only 1–2% at nominal operating fields [23]. This suppression prevents space-charge build-up and preserves spatial and timing performance even at maximum hit rates.

Trigger Detectors

The trigger detectors provide fast bunch crossing identification and coarse position measurements to seed muon reconstruction. They are designed for nanosecond-level timing resolution, ensuring correct bunch crossing assignment.

- **Resistive Plate Chambers** – RPCs are parallel-plate gaseous detectors used in

the barrel region ($|\eta| < 1.05$). Each chamber consists of two resistive bakelite plates separated by a 2 mm gas gap, filled with a mixture of 94.7% $C_2H_2F_4$, 5% C_4H_{10}, and 0.3% SF_6. High voltage of about 9.6 kV is applied across the plates, producing fast avalanches with signal rise times of $\mathcal{O}(1\,\text{ns})$. Copper readout strips on the external surfaces of the plates provide measurements in both the η and ϕ views. On-detector front-end cards amplify and discriminate the signals, which are then combined in coincidence matrices to select patterns compatible with muon trajectories. Typical performance includes a time resolution of about 2.5 ns and spatial resolution of 2 cm in η [21].

- **Thin Gap Chambers** – TGCs are multiwire proportional chambers used in the endcap region ($1.05 < |\eta| < 2.4$). Each chamber has closely spaced anode wires (1.8 mm pitch) stretched between cathode planes, operated in a gas mixture of 55% CO_2 and 45% n-pentane. This mixture provides high gas gain while suppressing secondary discharges. Readout strips measure the ϕ coordinate, while anode wires provide the η coordinate and precise timing. Signals are processed by on-detector electronics for amplification, shaping, and discrimination before entering pattern-recognition logic. The typical time resolution is 4 ns, sufficient for unambiguous bunch-crossing identification [21]. Spatial resolution is in the order of 4 mm [14].

- **Small Thin Gap Chambers** – The sTGCs are similar in operating principle to standard TGCs but with finer strip segmentation, achieving spatial resolutions of a few hundred µm. They are arranged in quadruplets, with pad electrodes providing coarse hit position information and strips refining the coordinate measurement. On-detector electronics perform fast amplification, shaping, and discrimination; pad signals are processed in dedicated pad trigger ASICs running at 40 MHz for bunch-crossing identification. Strip data are buffered locally for matching with pad coincidences.

2.1.6 Specialized forward detectors

In addition to the main sub-detectors, several small forward systems are used primarily for timing, luminosity monitoring, and targeting forward physics signatures.

- **ATLAS Forward Proton Detectors (AFP)** – The AFP system is located symmetrically on both sides of the ATLAS interaction point, with two "NEAR" stations at $z \approx \pm 205$ m and two "FAR" stations at $z \approx \pm 217$ m [24]. Each station houses four planes of 3D silicon pixel tracking detectors, tilted to improve spatial resolution to about 6 μm in the horizontal direction. The FAR stations additionally host time-of-flight detectors, providing timing resolutions of $\mathcal{O}(20$ ps) in optimal conditions.

The AFP measures intact protons scattered at very small angles in diffractive and photon-induced processes, enabling precise reconstruction of the proton fractional energy loss[1] and vertex position for exclusive event selection.

- **Beam Conditions Monitor (BCM)** – The BCM employs diamond sensors with a high voltage applied across them to detect the passage of charged particles. It is positioned near the beam pipe at $z = \pm 1.84$ m. The BCM provides bunch-by-bunch luminosity measurements and beam loss detection with sub-nanosecond timing resolution [14, 15].

- **Beam Pick-up Timing (BPTX)** – The BPTX system consists of electrostatic button pick-up detectors located at $z = \pm 175$ m from the ATLAS interaction point [25]. The signals are used both in the ATLAS Level-1 trigger to provide a bunch-passage reference and in a dedicated monitoring system to measure the phase between the LHC clock and the beam with better than 100 ps precision. The system also provides bunch-by-bunch intensity and longitudinal structure measurements, and is crucial for ensuring correct synchronization of detector readout with collisions.

- **Beam Loss Monitors (BLM)** – The BLM system protects LHC magnets and other sensitive components by detecting beam losses with ionization chambers and secondary emission monitors placed at critical locations around the ring [26]. In ATLAS, forward BLM devices close to the experiment are used to monitor losses near the interaction region, aiding both machine protection and background diagnostics for physics analyses. They have a large dynamic range, fast response (down to a single LHC turn, 89 µs), and are integrated into the LHC interlock system to request a beam dump if thresholds are exceeded. The forward BLM devices are located at various positions from a few meters to about 250 m from the interaction point in both beam directions [26].

- **Minimum-Bias Trigger Scintillators (MBTS)** – The MBTS is a scintillator detector located on the inner face of the endcap cryostats at $z = \pm 3.6$ m. The MBTS are used for triggering events in minimum-bias and forward physics, as well as for measuring luminosity. Thanks to their position, they detect low-energy particles produced at the edges or in the peripheral regions of a collision, helping to select events that are important for studying background processes and the beam's conditions [14].

[1]The fraction of the original LHC beam energy that a scattered proton has lost after the interaction at the ATLAS IP

- **LUminosity measurement using a Cherenkov Integrating Detector (LUCID)** – LUCID is a pair of Cherenkov detectors positioned around the beam pipe at about $z = \pm 17$ m, LUCID serves as the primary online and offline luminosity monitor. By detecting the Cherenkov radiation emitted by high-energy charged particles, LUCID provides precise luminosity measurements, which are crucial for real-time beam monitoring and integrated luminosity estimation during both data-taking and analysis [15].

- **Zero Degree Calorimeters (ZDC)** – The ZDC are Tungsten/quartz calorimeters located at $z = \pm 140$ m from the interaction point, the ZDCs are designed to detect neutral particles emitted at zero degrees (along the beam axis). These calorimeters are primarily used for heavy-ion physics, studying diffractive events, and triggering on ultra-peripheral collisions, where the nuclei interact via electromagnetic forces without participating in a full hadronic collision [15].

2.2 The ATLAS Trigger and Data Acquisition system

The ATLAS Trigger and Data Acquisition (TDAQ) system, as anticipated in Section 2.1, is responsible for the real-time selection, readout, and processing of data from the detector. The trigger function is to reduce the 40 MHz bunch-crossing rate to an output rate suitable for permanent storage. This is achieved through a two-stage architecture comprising the L1 fixed-latency hardware-based trigger and the software-based HLT, along with a distributed data acquisition system that handles the readout, event building, and data flow to storage.

The bunch clock is used by the detector front-end electronics for the digitization of the analog signals. The sampled data is stored locally in pipeline buffers, while waiting for the trigger decision. When one interesting event triggers, the readout system extracts data from the correct location in the buffers using its bunch crossing identification. Eventually, if no trigger fires, data is overwritten.

The L1 trigger is implemented using custom hardware to cope with high data rates and latency constraints. It identifies potential physics signatures, such as high-p_T leptons, energetic jets, and large MET, using coarse detector information from the calorimeters, the muon spectrometer and forward detectors.

The L1 trigger system comprises:

- L1 Calorimeter Trigger (L1Calo)

- L1 Muon Trigger (L1Muon)

- Topological Processor (L1Topo)

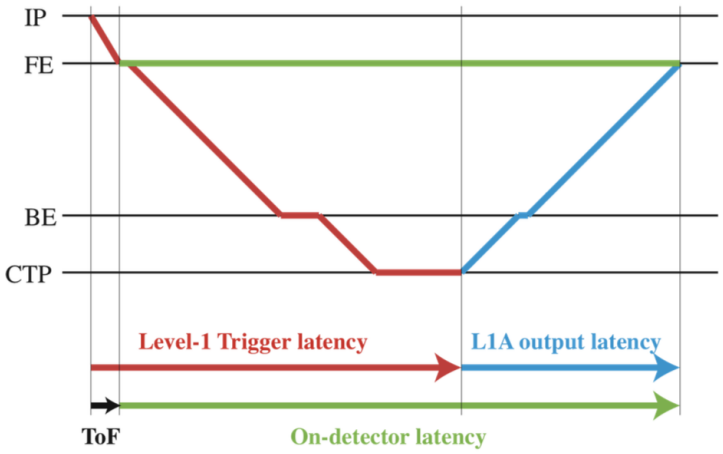

Figure 2.6: Schematic representation of latency definitions in the ATLAS Level-1 trigger system [27].

- Central Trigger (CT)
  - Muon-to-Central Trigger Processor Interface (MUCTPI)
  - Central Trigger Processor (CTP)
  - Trigger and Timing Control (TTC)

The Level-1 triggering conditions are defined as requirements on the multiplicity of objects above programmable p_T/E_T thresholds — e.g., ≥ 2 muons with $p_T > 20\,\text{GeV}$ or ≥ 1 EM cluster with $E_T > 20\,\text{GeV}$. In contrast, conditions on ΣE_T and E_T^{miss} are based solely on total energy thresholds and do not require an object multiplicity condition.

The Central Trigger system collects all trigger signals and makes the Level-1 Accept (L1A) decision. In Run-3, the total L1 trigger latency — measured from the Interaction Point to the CTP output — is fixed and must remain below the maximum allowed value of 2.5 µs (about 84 BC). Once formed, a positive L1A is distributed to all sub-detectors via the Trigger and Timing Control system, initiating data readout from the pipeline memories. Figure 2.6 illustrates the definitions of the different latency components in the L1 trigger system. The Level-1 trigger latency (red) is measured from the interaction point (IP) to the output of the CTP. The L1A output latency (blue) covers the signal propagation from the CTP to the sub-detector front-end electronics, while the on-detector latency (green) accounts for the full round trip from the front-end through the back-end and back to the front-end.

The second stage of selection is the HLT, which runs on a scalable computing farm. It refines L1-selected events using full detector granularity and precision reconstruction algorithms. Events passed to the HLT include associated Region-of-Interest (ROI) information, which are localized coordinates in η and ϕ, as well as trigger object types and threshold levels.

A schematic representation of the TDAQ system used during Run-3 is shown in Figure 2.7.

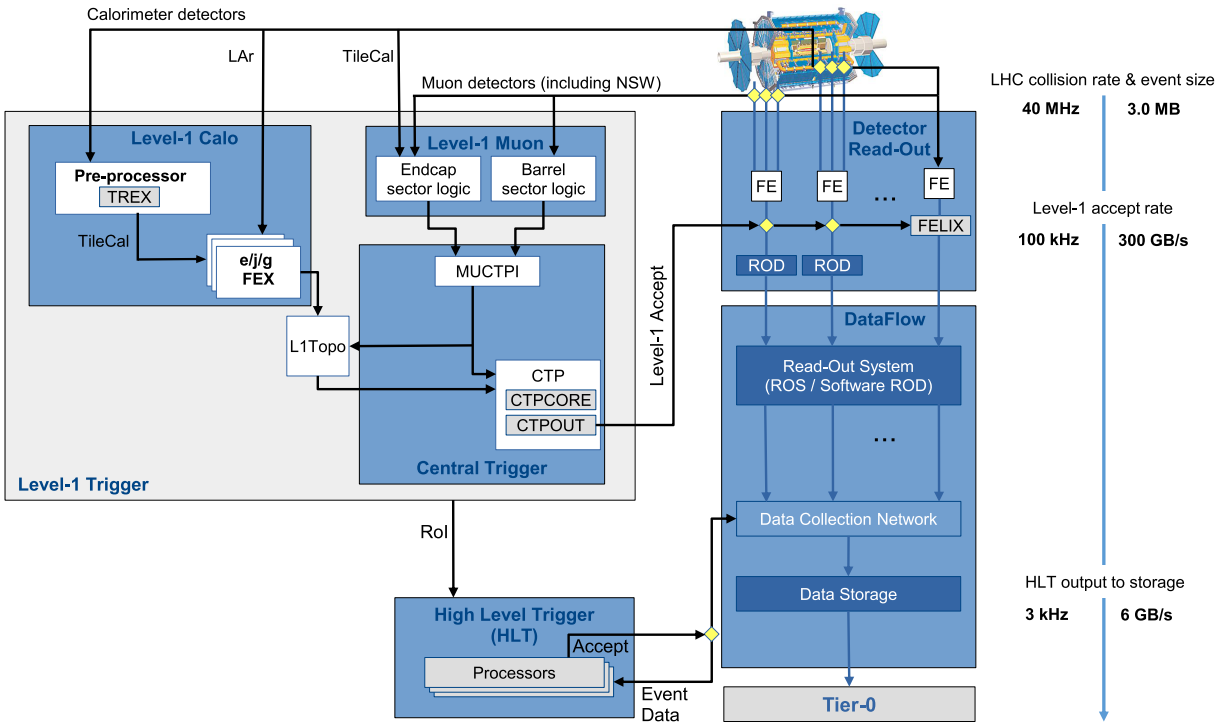

Figure 2.7: Schematic overview of the TDAQ system after the Phase-I upgrade. Adapted from [28].

2.2.1 L1Calo Trigger

The L1Calo Trigger selects events containing energetic electrons, photons, τ-leptons candidates, jets, large transverse energy sums, or MET using calorimeter data for every bunch crossing.

Trigger towers and SuperCells

The first stage of the calorimeter trigger is the formation of fixed granularity trigger towers, each covering $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ in most of the detector. In the Tile calorimeter, these towers are formed by analog summation of the relevant photomultiplier signals at the front-end and transmitted to the trigger system. In the Liquid Argon calorimeter, the Phase-1 upgrade replaced the previous tower sums with SuperCells. Each tower is now subdivided into up to ten SuperCells, adding longitudinal layer information and finer η segmentation (as fine as $\Delta\eta \times \Delta\phi = 0.025 \times 0.1$ in the barrel EM layers). This additional granularity improves shower-shape discrimination and pile-up suppression in the trigger. Figure 2.8 illustrates an example showing the L1Calo improved resolution.

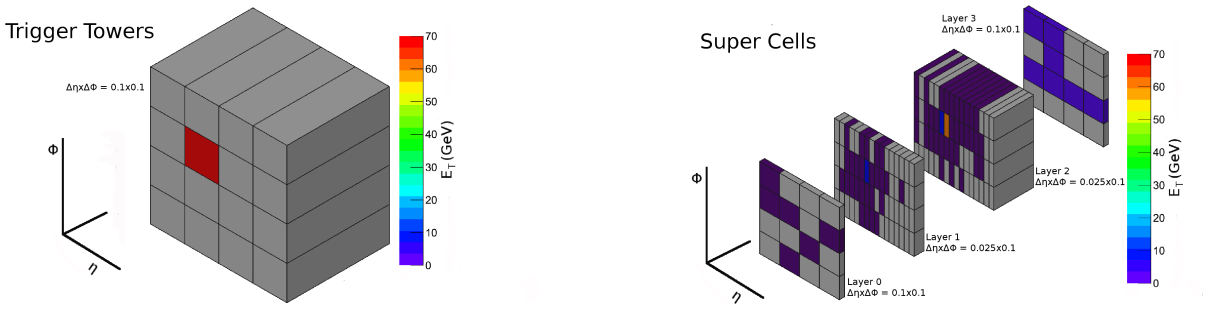

(a) Run-2 L1Calo trigger using legacy trigger towers with granularity $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$.

(b) Run-3 Phase-I upgrade using SuperCells with finer granularity and layer-dependent readout.

Figure 2.8: Simulated energy deposit of a $E_T = 70\,\text{GeV}$ electron in the L1Calo trigger system. Figures adapted from [29].

Feature extraction

The trigger-tower and SuperCells data are processed in real time by dedicated Feature Extractor (FEX) processors:

- **eFEX** (electron Feature EXtractor) for electrons, photons, and τ-leptons candidates, exploiting shower-shape and isolation variables.
- **jFEX** (jet Feature EXtractor) for jets, large-radius τ objects, and MET.
- **gFEX** (global Feature EXtractor) for event-wide quantities such as total transverse energy ΣE_T and global MET.

The eFEX and jFEX modules operate with multiple boards in parallel, processing specific regions of the calorimeter, while the gFEX processes the full coverage within a single board.

Output

The outputs from the FEX processors are reduced to a compact set of trigger objects and global quantities like the **Trigger Objects (TOBs)**. The TOBs contain object-level information (position, E_T) for identified electromagnetic clusters, jets, τ candidates, energy sums, MET, and scalar sum of jet transverse momenta (H_T). All this information is transmitted to L1Topo.

2.2.2 L1Muon Trigger

The L1Muon Trigger identifies muon candidates using fast trigger chambers: RPCs in the barrel, TGCs and New Small Wheel in the endcaps. Additional input is provided by the Tile extended barrel regions. Signals from these detectors are processed for every

bunch crossing to find hit patterns compatible with muons originating from the interaction point [14].

Hit coincidences

In the barrel, each trigger sector contains three layers of RPCs arranged in concentric cylinders around the beam line. Low-p_T candidates are identified by requiring coincident hits in the two innermost layers, while high-p_T candidates must also produce a hit in the outermost layer. This geometric coincidence ensures that selected tracks have sufficient curvature in the toroidal field to exceed a programmed p_T threshold.

In the endcaps, the trigger uses three layers of TGCs in the "Big Wheel" stations to form muon candidates. Run-3 operation benefits from additional matching to hits in the New Small Wheel sTGCs or MMs, to reject non-collision backgrounds. The input from the Tile calorimeter provides additional coincidence. This matching suppresses low-p_T and beam-related fake triggers by up to 45% in high pile-up conditions [23].

Figure 2.9 shows an example of coincidences in both barrel and endcap region.

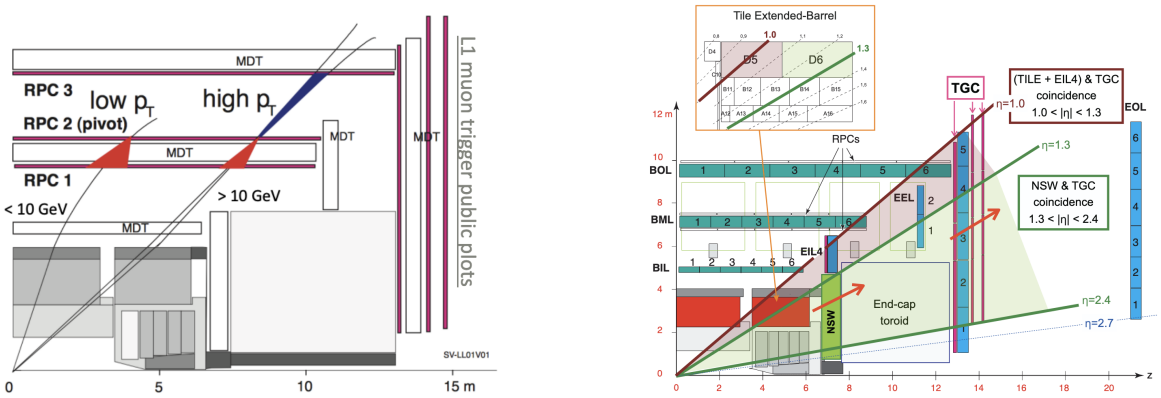

(a) L1Muon trigger coincidences in the barrel region.

(b) L1Muon trigger coincidences in the endcap region.

Figure 2.9: Examples of hit coincidences in the muon spectrometer trigger sectors [30].

Sector Logic

Each muon trigger sector has dedicated Sector Logic (SL) electronics that receive the pattern-recognition results from the on-detector trigger chambers. The SL compares the hit patterns to a set of predefined coincidence maps stored in look-up tables, which correspond to different p_T thresholds. These maps are defined using reference trajectories: the straight line of an infinite-momentum muon provides the baseline, while tracks of finite p_T traced through the toroidal field and detector geometry determine the allowed hit combinations.

For each bunch crossing, the SL identifies the highest-p_T candidate per threshold within its sector and encodes the candidate's p_T code, η–ϕ coordinates, and quality bits into a ROI word. The ROI words are transmitted over optical links to the Muon-to-Central Trigger Processor MUCTPI.

Output

Within each trigger sector, up to 2 (or 4) of the RPC (or TGC) p_T candidates for each programmed threshold are selected by the SL electronics. Each candidate is assigned a ROI coordinate in η–ϕ, together with its p_T threshold. The ROI words and multiplicity counts for all thresholds are sent to the MUCTPI.

2.2.3 L1 Topological Processor

The L1Topo system consists of three Advanced Communications Computing Architecture (ATCA) boards. It is the dedicated processor stage that refines the trigger decision using geometric and kinematic relationships between trigger objects from the calorimeter and muon systems. In Run-3, the L1Calo and muon TOBs are sent to L1Topo.

Input

L1Topo receives TOBs from L1Calo and from the MUCTPI for every bunch crossing. Each TOB encodes the object type (e.g. electron, jet, muon), η–ϕ coordinates, transverse energy or momentum, and quality bits. For each object type, the number of TOBs is capped at a configurable maximum (highest-p_T first), a limit set by the available bandwidth and processing resources, including the need to stay within the Level-1 latency budget. Global quantities such as ΣE_T, E_T^{miss}, and H_T are also provided as TOBs.

Topological and multiplicity algorithms.

L1Topo implements topological algorithms requiring correlations between objects. Examples include:

- Minimum angular separation between two jets (ΔR or $\Delta\phi$ cuts).
- Invariant mass of an e^+e^- or $\mu^+\mu^-$ pair above a threshold, or in a mass window (e.g., used in B-physics).
- Back-to-back topology between a photon and a jet.
- Large azimuthal separation between E_T^{miss} and leading jets.

Output

The resulting multiplicity counts and algorithm decision bits are sent to the CTP, which evaluates them against a selection of trigger conditions to decide whether to issue a L1A.

2.2.4 Muon-to-Central Trigger Processor

The Muon-to-Central Trigger Processor Interface collects muon trigger candidates from all muon trigger sectors, resolves overlaps between sectors, calculates global muon multiplicities, and forwards both multiplicities and detailed muon candidate information to the CTP and the Level-1 Topological Processor (L1Topo).

The Phase-1 MUCTPI, introduced for Run-3, replaced the legacy system of 18 VME 9U boards with a single ATCA blade. This compact redesign significantly increases bandwidth, and simplifies maintenance. The board hosts three large processing FPGAs: two Muon Sector Processor (MSP) FPGAs for candidate merging and TOB formation, and one Trigger Readout Processor (TRP) FPGA for multiplicity calculation and readout formatting. A 64-bit ARM-based System-on-Chip (SoC) running embedded Linux provides slow-control, configuration, and monitoring. The main MUCTPI components are shown in Figure 2.10.

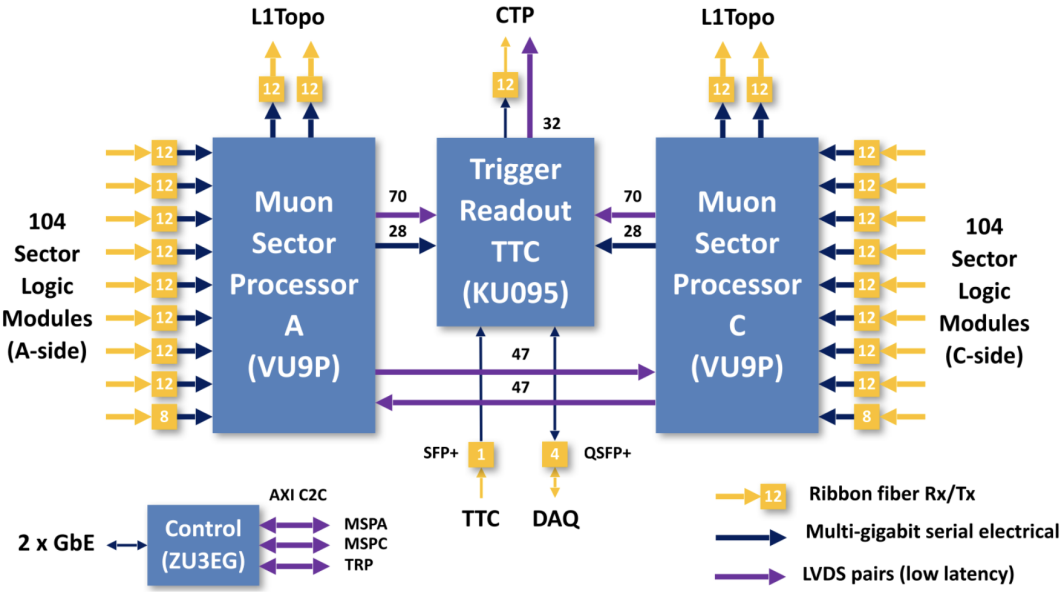

Figure 2.10: Block diagram of the Phase-1 MUCTPI (v3).

Input

For each bunch crossing, the MUCTPI receives 208 optical input links from the barrel and endcap Sector Logic boards, each carrying 2 or 4 (RPC or TGC respectively) of the

highest-p_T candidates for a set of programmable p_T thresholds, together with coarse η–ϕ position and quality bits.

Processing

The MSP FPGAs resolve overlaps between candidates from neighboring trigger sectors to avoid double counting muons in overlapping sectors. Surviving candidates are sorted by p_T and encoded into TOBs containing p_T code, η–ϕ coordinates, quality flags and the origin of the candidates. The TRP FPGA counts the number of candidates above each p_T threshold, producing multiplicities used in the Level-1 decision.

Output

The TOBs are transmitted via high-speed optical links to the L1Topo processors, enabling topological selections based on muon–muon or muon–calorimeter correlations. In parallel, multiplicity counts are sent to the CTP for combination with other Level-1 trigger inputs. For each L1A, the MUCTPI readout path delivers detailed muon candidate information, multiplicities, and TOBs to the Data Acquisition system, where they are available to the High-Level Trigger for further event selection.

MUCTPI monitoring

The MUCTPI includes comprehensive monitoring capabilities to ensure correct operation and to allow detailed diagnostics of the muon trigger path. Internally, each MSP FPGA maintains per-threshold candidate counters, while the TRP FPGA monitors the muon multiplicities for programmable p_T thresholds. All those monitoring capabilities exist also per-bunch.

In addition to the monitoring counters, the MUCTPI includes spy memories that can capture raw data, providing low-level diagnostic snapshots for firmware debugging without disturbing normal operation.

Figure 2.11 shows a histogram of MUCTPI muon candidates per-bunch per-sector for the ATLAS run 503778. The bins are well aligned, indicating good timing from L1Muon.

2.2.5 Central Trigger Processor

The CTP is the core of the ATLAS L1-1 trigger system. It collects trigger information from all L1 subsystems, evaluates the active triggering conditions, and issues the final L1A signal. The CTP also interfaces to the general TTC network for timing synchronization, and is the primary source of trigger and timing commands to the local TTC system, and ultimately, to the detector.

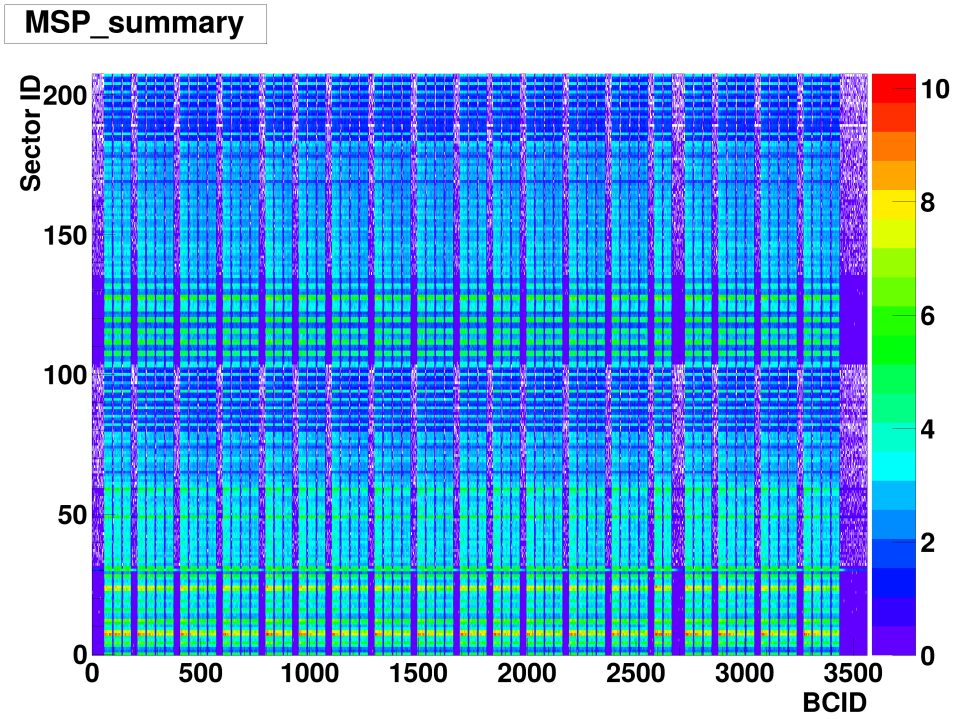

Figure 2.11: A MUCTPI monitoring histogram showing the per-bunch rates of muon candidates per-sector. The unit of the temperature gradient is kHz.

The **trigger menu** defines the complete set of Level-1 trigger conditions evaluated by the CTP. Each trigger item specifies one or more physics objects (e.g. muons, electrons/photons, jets, hadronic taus, MET, ...) together with multiplicity and threshold requirements. Up to 512 trigger items can be configured simultaneously. Each item is assigned an integer, the prescale factor, which reduces its effective acceptance rate by only allowing one event in N that fulfills the item to pass. The menu also specifies additional constraints such as bunch-group masks, which selects the triggering bunches in the filling scheme. Different trigger menus are defined depending on the type of data-taking, e.g., normal proton-proton, ion-ion, cosmic-ray, calibration sequences, or technical runs.

During operation, the CTP receives trigger signals from L1Topo, MUCTPI, and selected forward detector inputs. These are aligned to the bunch clock before being compared against the configured trigger menu conditions. Items firing from the wrong bunch group are immediately rejected. Surviving items are processed by their individual prescale counters, and those still active are subject to deadtime restrictions to limit bursts or sustained high rates. The remaining active items are then combined in a logical OR to form the L1A signal distributed to all sub-detectors.

Partitioning

The CTP supports multiple independent partitions, each with its own trigger menu, bunch group configuration, and prescale settings. Partitions allow different sets of sub-detectors to be operated in parallel, for example to run calibration triggers for one subsystem while others collect physics data, or to take cosmic-ray data during LHC fills. Inputs to each partition are masked independently, and the L1A signals generated by each partition are merged before distribution. Partitioning thus provides flexible, simultaneous operation modes without interfering with the main physics data-taking stream.

Busy mechanism

The CTP implements a global busy-veto system to prevent buffer overflows and maintain synchronization between the trigger and sub-detector readout systems [31]. Busy signals can originate from multiple sources:

- Sub-detectors, through the TTC module.
- CTP timing interface, with two external burst veto inputs from RPC and TCG.
- CTP processing logic, asserting a busy signal if its internal readout or monitoring buffers approach capacity.

All busy sources are OR'ed and continuously monitored. The CTP inhibits trigger generation whenever busy is active. Busy duration counters in each module record the integrated veto time, and programmable alarms can be issued if a source exceeds a threshold occupancy. This mechanism operates in parallel with deadtime generation, but unlike deadtime, busy vetoes are asserted dynamically in response to instantaneous readout or sub-detector status.

Deadtime generation

The CTP applies programmable deadtime to prevent buffer overflow and loss of synchronization in sub-detector front-end electronics. Two categories are implemented. The **simple deadtime** algorithm enforces a fixed minimum interval, expressed in bunch crossings, between consecutive L1As. The **complex deadtime** category includes algorithms that regulate the average L1A rate over a period of time.

The main complex algorithm in use is the **leaky-bucket** model [32]: a counter ("bucket") is incremented on each L1A and decremented at a fixed rate ("leakage"), allowing short bursts of triggers while limiting the sustained L1A rate to a programmable threshold. In the current configuration, four independent leaky-buckets operate in parallel, each with parameters tuned to the constraints of different groups of sub-detectors.

In addition, a **sliding-window** algorithm is implemented, which monitors the number of L1As within a moving time window and inhibits further triggers if the count exceeds a programmable maximum.

CTP architecture and modules

The CTP is implemented as a 9 U VME crate hosting 14 custom boards of 8 different types, interconnected via the backplane. The CTP modules types are:

- **CTPMI** (Machine Interface) – bridge the LHC timing system to the CTP.

  In the ATLAS counting room the RF2TTC module[2] receives two bunch clocks (BC1, BC2) from the CCC and a BCRef from a local clock generator. It generates a Main bunch clock and Main Orbit signal, which are electrically sent to the CTPMI. In addition to the main clocking, the CTPMI produces the Event Counter Reset (ECR), that is a synchronous command that resets all L1ID counters (see item 2.2.5, CTPCORE), typically issued at the start of a run or during re-synchronization procedures. The CTPMI distributes the timing signals to the other modules through the backplane.

- **CTPIN** – receives trigger inputs via LVDS electrical links. It synchronizes and align them in time. The CTPIN maps the inputs on a dedicate PIT "Pattern-In-Time" backplane bus, making them available to the CTPMON and CTPCORE modules.

- **CTPCORE** – FPGA-based processing engine that implements the Level-1 decision logic at fixed latency [33]. For each bunch crossing, CTPCORE trigger information via three physical paths:

  - **Direct low-latency electrical inputs:** Connected by direct electrical links to the CTPCORE FPGA
    - L1Topo topological selections, and direct calorimeter- or muon-based thresholds from the eFEX, jFEX and gFEX processors.
    - Selected special-purpose triggers where minimum latency is required.
  - **Direct long-latency optical inputs:** Connected by direct optical links to the CTPCORE FPGA
    - L1Topo multiplicities.
    - Muon multiplicities from the MUCTPI.
  - **CTPIN inputs via the PIT backplane:** Triggers mapped onto the Pattern-In-Time (PIT) backplane for CTPCORE to read.

[2]The RF2TTC is part of the global-TTC distribution network.

The fixed-latency decision pipeline inside CTPCORE, shown in Fig. 2.12, processes the inputs in the following stages:

1. **LUT and CAM item formation:** The incoming trigger signals are first mapped onto a set of intermediate *threshold bits* by programmable Look-Up Tables (LUTs). Each LUT implements a truth table that can combine multiple raw inputs to encode object multiplicities or threshold conditions. The complete set of threshold bits is then passed to fast ternary Content-Addressable Memories (CAMs), which are specialized memory devices that operate in reverse to a normal RAM: instead of providing an address to read stored data, the current input pattern is compared in parallel against all stored entries. Each CAM entry specifies a combination of threshold bits (with support for "don't care" states) corresponding to one trigger item in the Level-1 menu. Items whose conditions are met produce a Trigger Before Prescale (TBP) bit.
2. **Bunch-group masking:** TBP bits are enabled only for bunch crossings belonging to their assigned LHC bunch group — predefined sets of BCIDs (Bunch Crossing Identifier, a counter labeling bunch crossings within one orbit) corresponding to specific beam conditions (e.g. colliding, unpaired, or empty bunches). This produces the BGRP output.
3. **Prescales:** Surviving items are reduced according to their individual prescale counters, producing Trigger After Prescale (TAP) bits. The CTPCORE implements random prescaling [34]. In this mode, the decision to accept or reject a trigger item is based on the output of a hardware pseudo-random number generator implemented as a hardware circuit. This generator produces a repeatable but statistically uniform sequence of bits[3]. On each bunch crossing, the generator's output is compared to a programmable threshold that sets the desired average acceptance fraction. Twenty-two discrete threshold values are available, corresponding to different mean acceptance probabilities, from nearly 100% down to a fraction of a hertz. Because the decision is based on a pseudo-random bit sequence rather than a fixed counting pattern, the actual acceptance of successive events is non-deterministic.
4. **Deadtime and busy veto:** TAP bits are subject to deadtime restrictions and to the global busy veto mechanism, yielding Trigger After Veto (TAV) bits.
5. **L1A formation:** If at least one enabled TAV bit is set, an L1A is issued via CTPOUT to all sub-detectors.

For each accepted event, the CTPCORE assigns and manages identifiers and synchronous commands:

[3]No bias toward certain values being produced more often than others

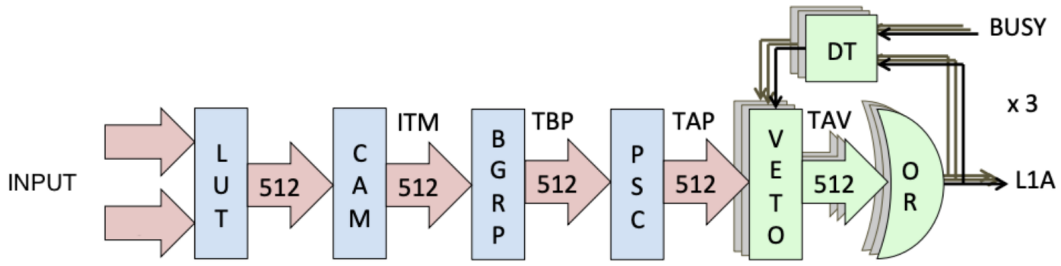

Figure 2.12: Functional diagram of the Level-1 Accept formation in the Central Trigger Processor during Run 3.

  - **BCID** – a 12-bit counter identifying the specific bunch crossing within the current LHC Orbit. It is generated using the Orbit signal as reference.
  - **Level-1 Event Number (L1ID)** – a 24-bit sequential counter incremented for each accepted event, used to uniquely identify events at the Level-1 stage within a run.
  - **Trigger Type Word (TTYP)** – a code included in each accepted event that specifies the origin or purpose of the trigger (e.g. physics, calibration, monitoring).

  The BCID, L1ID, and TTYP are assigned in the CTPCORE. All are passed to the CTPOUT module for distribution via the TTC network.

- **CTPOUT** – distributes the L1A and other TTC timing signals from the CTP to the ATLAS sub-detectors via the TTC distribution, see next Section 2.2.7. It receives from CTPCORE the accepted event information — including the bunch crossing identifier, the L1ID, and the trigger type word — and formats these into TTC broadcast messages, ensuring synchronous delivery with respect to the LHC bunch clock. The CTPOUT also receives sub-detector calibration request (CALREQ) signals and drives them onto the dedicated calibration bus (CAL) on the CTP backplane [35]. In addition, the CTPOUT receives and back-propagates busy signals from the sub-detectors via the TTC modules ATLAS Local Trigger Interface (ALTI), masks and monitors them, and places the resulting busy status onto the common busy line of the COM backplane for use by the CTPCORE in trigger inhibition [35].

- **CTPCAL** – receives calibration request signals from the CAL bus. The CTPCAL implements a time-sharing scheme for calibration triggers by assigning specific LHC turns to individual sub-detectors. A programmable turn counter, incremented by the Orbit signal and wrapping after a configurable maximum (up to 20), identifies the current turn number. A look-up table maps each turn number to one group of three calibration request lines from a single sub-detector. During a given turn,

only the CALREQ lines from the mapped sub-detector are forwarded to the CTPIN module; all others are ignored. This ensures that calibration triggers from different sub-detectors are separated in time, avoiding deadtime conflicts while maintaining synchronous distribution through the CTP [36].

In addition, CTPCAL can inject auxiliary trigger signals from its NIM-level front-panel. All selection, mapping, and timing parameters are configurable via the VME interface, and the module includes monitoring counters [36].

- **CTPMON** – captures detailed trigger-decision data for monitoring and diagnostics of the PIT trigger items. It implements rate and per-bunch counters. The per-bunch counters are particularly useful to monitor if triggers are firing in the correct BCID. If this is not the case in the CTPIN and CTPCORE it is possible to perform small phase adjustments, or triggering edge changes to address the issue.

- **NIM2LVDS** – converts NIM-level inputs to LVDS logic levels compatible with the CTPIN.

2.2.6 Beam pickup based timing system

The Beam Pick-up Timing (BPTX) detectors provide a precise timing reference for the L1 Trigger system and are essential for synchronization of the ATLAS detector to within a fraction of the LHC bunch-crossing period [37, 25].

Each BPTX station is located approximately 175 m upstream and downstream of the ATLAS interaction point, one per incoming beam. A station consists of four electrostatic button pick-ups arranged symmetrically around the beam pipe. The signals from the four buttons are summed to produce a single signal per station with high signal-to-noise ratio [37].

The BPTX signals are transmitted via about 200 m of low-loss coaxial cable from the detectors to the counting room, where they are used for two main purposes:

1. **Bunch pattern identification:** After discrimination, the waveform acquired from the oscilloscope is processed to identify filled bunches, detect empty bunch slots in the LHC pattern, and provide the BCID.

2. **Clock phase monitoring and adjustment:** The raw analog signals are digitized by a dedicated readout system based on commercial sampling oscilloscopes, allowing the phase between the LHC clock and the passage of each bunch to be measured with a precision of 20 ps [37]. These measurements enable continuous monitoring of clock stability, detection of phase drifts, and adjustment of the clock phase to compensate for variations in the length of the optical fiber link between the CCC and

ATLAS—primarily due to temperature changes. To maintain precise timing stability over long-term operation, ATLAS dynamically compensates for such phase shifts introduced by temperature variations in the timing distribution system. As fibers and electronic components undergo thermal expansion and contraction through seasonal and diurnal changes, the effective propagation delays can drift by several tens of picoseconds. A dedicated board in the RF2TTC crate (CORDE board) adjusts the timing offsets continuously, ensuring that the bunch crossing clock and Orbit signals remain aligned with the actual arrival times of LHC bunches as observed by the BPTX detectors. This active compensation maintains synchronization across the trigger system despite ambient temperature fluctuations. Satellite bunches[4] can also be identified in this way.

2.2.7 Trigger, Timing and Control System

The ATLAS Trigger, Timing and Control (TTC) system distributes the Level-1 Accept (L1A) and synchronous timing signals to all sub-detectors, ensuring that front-end electronics operate synchronously with the 40 MHz bunch clock. It is based on the standard TTC architecture developed for all LHC experiments, which only provides downstream signal propagation using dedicated TTC optical links.

Architecture and distribution chain

In Run-3, the interface between the CTP and the TTC distribution network is provided by the **ATLAS Local Trigger Interface** (ALTI) module. The ALTI replaces legacy TTC modules used in Run-1 and Run-2, and was developed primarily to support the Phase-I upgrade systems due to the scarcity of legacy modules.

The ALTI performs the encoding of trigger and timing information into the TTC data format and transmits it via an optical link. It also receives the Orbit signal from the CTP and generates the Bunch Counter Reset (BCR) signal, which instructs the front-end electronics to reset their BCID counters, ensuring synchronisation of bunch counters in both front- and back-end electronics. The TTC stream, including the BCR, is distributed using a passive optical fan-out tree to each sub-detector partition. At the receiving end, each partition uses one or more radiation-hard **TTCrx** ASICs, which interface the TTC optical signal to the front-end electronics.

ALTI operation

The ALTI module can operate in two modes:

[4] Unintended bunches of protons that occupy the wrong RF bucket relative to the main bunch.

1. **Standalone mode** – the ALTI acts as a local trigger source, generating timing and L1A patterns internally. This "MiniCTP" mode is used for commissioning or sub-detector standalone runs.
2. **Global mode** – the ALTI relays signals received from the CTP, enabling global synchronization across the ATLAS detector.

TTCrx functions

The TTCrx ASIC recovers the bunch clock and data from the incoming optical TTC stream [38]. The signal uses bi-phase mark (BPM) encoding and carries two time-division-multiplexed channels at 80 Mb/s:

- **Channel A** – used exclusively for transmitting L1A decisions (one bit per bunch crossing).
- **Channel B** – carries broadcast commands (e.g. BCR, ECR), addressed control commands, and configuration data.

The TTCrx includes programmable fine (104 ps steps) and coarse (25 ns steps) delay adjustments to compensate for signal path variations across front-end electronics. It outputs recovered clocks, trigger strobes, and decodes BCID and event counter. Configuration is performed via the I^2C interface.

All data frames in the TTC protocol are protected with Hamming encoding for single-bit error correction and double-bit error detection. The TTCrx maintains internal counters for error statistics and single event upsets (SEUs), which can be accessed via the I^2C interface.

2.2.8 Readout System

When a L1A is issued, data stored in the front-end pipelines of each sub-detector is transferred to **ReadOut Drivers (RODs)**. These off-detector modules format the raw channel data into sub-detector event fragments. In the legacy architecture, these fragments are transmitted over **ReadOut Links (ROLs)** to **ReadOut Buffers (ROBs)**, where they await retrieval by the High-Level Trigger (HLT) [15].

To prevent buffer overflows, each sub-detector can assert a **Busy** signal when its RODs or FELIX-based systems are temporarily unable to accept new L1As. This signal is propagated electrically to the **ALTI** module, which relays it to the Central Trigger Processor (CTP) via a dedicated parallel upstream link. The assertion of Busy inhibits the generation of new L1As until the affected system clears the condition, providing synchronous flow control across the data acquisition system.

For Run-3, the traditional ROD/ROL/ROB chain remains in use for some subsystems, but many upgraded detectors now use the **Front-End Link eXchange (FELIX)** architecture. FELIX is a PCIe interface card hosted in a server, which connects directly to front-end FPGAs and ASICs via high-speed optical links. Each FELIX link can operate in:

- GBT mode (GigaBit Transceiver): 4.8 Gb/s per link, typically used for timing, control, and low-rate data paths.
- FULL mode: 9.6 Gb/s per link, typically used for high-rate data acquisition.

Within the FELIX host, the **Software ReadOut Driver (SWROD)** application aggregates event fragments from multiple E-links (serial sub-links) into ROB-format fragments compatible with the ATLAS Data Acquisition request/response model. This enables FELIX-based systems to integrate with the existing readout and HLT infrastructure, while supporting significantly higher throughput.

2.2.9 High-Level Trigger (HLT)

The second stage of the ATLAS event selection is the **High-Level Trigger (HLT)**, implemented entirely in software and running on a scalable farm of commodity servers [15]. The HLT refines the selections made by the Level-1 trigger using the full detector granularity and offline-quality reconstruction algorithms.

To minimize processing time, the HLT exploits the **Region-of-Interest (ROI)** mechanism: the Level-1 trigger provides the η–ϕ coordinates and type of physics object candidates (e.g. electron, muon, jet, MET) together with the L1A. The HLT can then request only the relevant event fragments for reconstruction in these regions, significantly reducing input data volume and CPU usage.

Run-3 HLT algorithms are implemented within the multithreaded **Athena** framework, which supports concurrent execution of multiple event reconstruction tasks. This improves CPU utilization and throughput. The use of Athena also enables unified development between offline and online reconstruction software.

The HLT processes events at the full Level-1 accept rate of up to 100 kHz. After selection, the rate is reduced to approximately 3 kHz for permanent storage, corresponding to an average output bandwidth of about 6 GB/s. These accepted events are then passed to the Data Acquisition system for transfer to CERN's offline storage and subsequent physics analysis.

2.2.10 Data Acquisition (DAQ)

The ATLAS DAQ system collects event fragments from all sub-detectors following an L1A, assembles them into complete events, and delivers them to the HLT for processing [15]. In Run-3, it operates in a hybrid mode:

- **Legacy path:** ROD → ROL → ROB (in **ReadOut System (ROS)** servers) → Sub-Farm Input (SFI) nodes for event building.
- **FELIX/SWROD path:** Front-end → FELIX → SWROD buffer → DAQ network event building.

The Run-3 DAQ must sustain the full L1A rate of 100 kHz plus contingency, with an average event size of about 2.1 MB at $\langle\mu\rangle \approx 60$. This corresponds to a maximum throughput of about 8 GB/s to storage. The system includes at least 1.4 PB of effective local storage at P1 to buffer one day of data in case of CERN mass-storage downtime.

2.2.11 TDAQ Software Framework

The TDAQ software provides a central remote control framework to coordinate control, configuration, monitoring, and data-taking [15]. It consists of modular processes running on a distributed computing farm.

A central operator interface provides run control by starting, stopping, and configuring groups of applications through a single command stream. Operations are organized around a finite set of well-defined states—`INITIALIZED`, `CONFIGURED`, `RUNNING`, and `STOPPED`—with transitions between these states triggering the corresponding hardware and software actions. In general:

- `INITIALIZED` resets hardware and software to a known baseline.
- `CONFIGURED` applies the selected operational settings and starts monitoring.
- `RUNNING` starts data-taking.
- `STOPPED` halts activity while preserving readiness for further operations.

Each ATLAS subsystem implements its own transition logic within this global framework, ensuring coordinated behavior across the experiment (examples are given in later chapters).

Configuration follows a layered model. High-level settings describe roles and operating modes, while low-level descriptors capture hardware and firmware details. This separation enables reuse of templates across environments (from test benches to full operations) and supports versioning for reproducibility. Configuration data are stored in structured,

persistent text-based files[5], which can be edited through graphical tools or programmatic interfaces. In large setups, these files may be organized into hierarchical databases to maintain system-wide consistency and track historical versions.

To coordinate a large number of independent applications, the framework provides a middleware layer for information exchange based on publish–subscribe patterns[6]. Components can publish status, rates, counters, and environmental data; other components subscribe to the specific items they need. This decouples producers and consumers, reduces point-to-point dependencies, and allows operator panels and automated procedures to react promptly to changes.

Online monitoring is integrated into the same infrastructure. Applications produce quality and performance metrics—commonly as histograms and counters—that are collected, indexed, and redistributed to graphical displays and expert tools. In routine operation, interaction with the TDAQ system is through an integrated graphical user interface (IGUI) [41] that presents the global system status, visualizes error messages, and provides controls for moving the system between run states. Specialised panels can be loaded for different sub-detectors, and hierarchical views of hardware and software resources give operators a coherent, real-time view of the experiment.

Errors are handled by a uniform reporting facility. Messages carry severity and category metadata, can be correlated across applications, and are routed both to operator consoles and to automated recovery logic. Aggregation and throttling prevent message floods, and persistent logs support post-mortem analysis.

MUCTPI Monitoring

As part of the TDAQ slow control and online monitoring infrastructure, the author developed the **MuctpiMonitoring** package, a set of applications dedicated to tracking the operational status and performance of the MUCTPI. These applications acquire real-time metrics from firmware registers via low-level software interfaces and publish them through the middleware information-exchange layer.

A dedicated software tool developed at CERN, the **hwcompiler**, takes the defined map of FPGA registers, memories, and FIFOs and automatically generates the matching firmware logic and software access functions [42]. Those access functions are the basis used by the MuctpiMonitoring to operate on the MUCTPI.

The monitored quantities include:

- Input rates from all Sector Logic inputs, including per-bunch distributions.
- Muon multiplicity rates transmitted to the CTP, also with per-bunch detail.

[5]The Object Kernel Support (OKS) [39].
[6]The Information Service (IS) [40].

- Busy state sources and their fractional contributions.

An auxiliary application, running on a host PC in the ATLAS counting room, publishes histograms derived from per-bunch data for additional diagnostic insight.

Collected data are visualized through web-based dashboards implemented in Grafana web pages. This enables ATLAS control room operators and subsystem experts to perform immediate health checks as well as long-term performance analysis (Figure 2.13).

Together, these services—state-based run control, layered configuration, publish–subscribe information exchange, integrated online monitoring, structured error reporting, and dedicated subsystem monitoring applications—provide a robust software environment that scales from local tests to full detector operations, while keeping behavior predictable, observable, and reproducible.

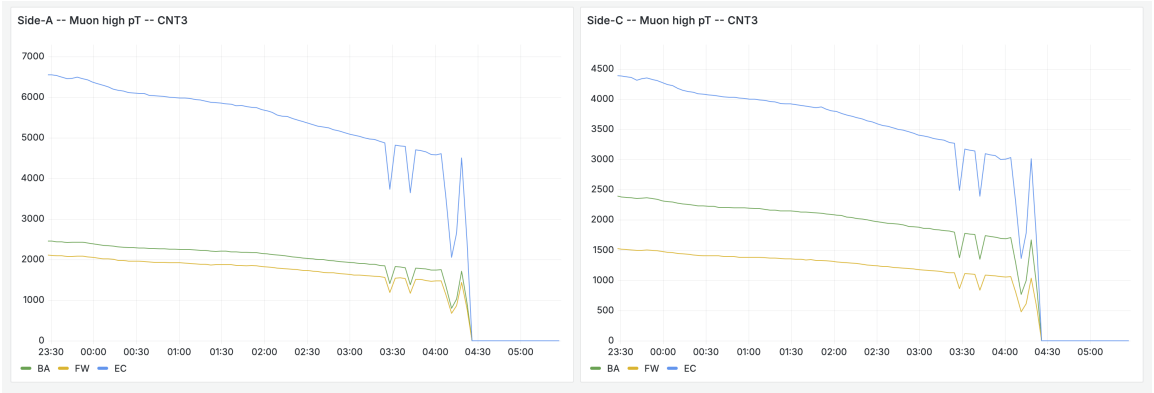

Figure 2.13: Grafana dashboard of MUCTPI Sector Logic Input rates, showing data for high-p_T muons ($\geq 12\,\text{GeV}$) per sector type on Side-A.

Offline Data Quality Monitoring for the MUCTPI

In the context of the author's ATLAS qualification task, Data Quality Monitoring (DQM) tools were developed for the MUCTPI. Although this work is performed within the `Athena` offline framework and is not part of the HLT algorithms described above, it plays a complementary role in ensuring the reliability of the trigger system.

For this task, the content of the histograms was retrieved directly from the MUCTPI raw data readout, including:

- Muon candidates
- Muon TOBs
- Timing information

Figure 2.14 shows two example histograms: the distribution of muon candidates per event and a TOB hitmap for the barrel–A side.

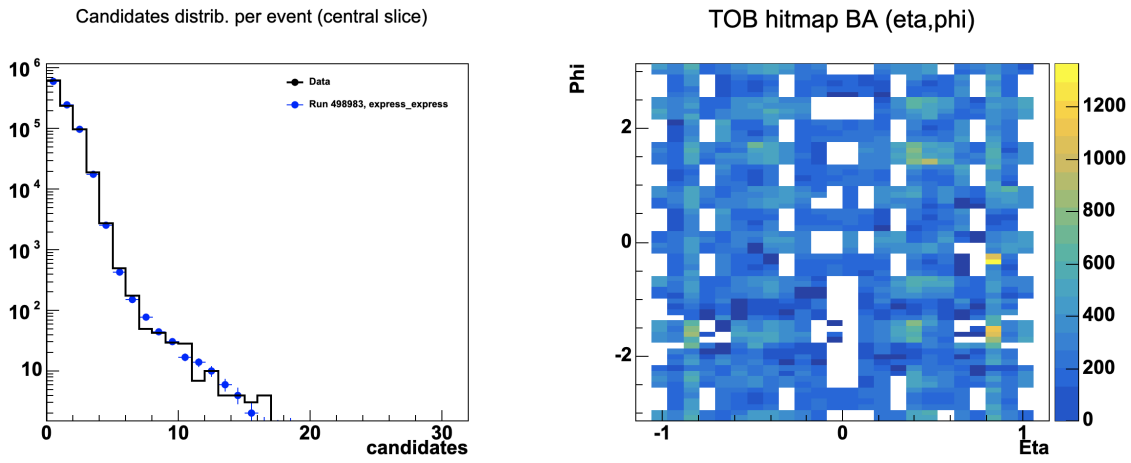

Figure 2.14: Offline MUCTPI DQM histograms from ATLAS run 505244: (left) muon candidate multiplicity per event; (right) TOB hitmap in η–ϕ coordinates for the barrel sectors.

Chapter 3

The ATLAS Phase-2 upgrades

3.1 Overview of the ATLAS Phase-2 upgrade

The High-Luminosity Large Hadron Collider (HL-LHC) will push experimental particle physics into a new era, targeting an instantaneous luminosity baseline of $5 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, up to an ultimate value of $7.5 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The corresponding pile-up will be about 140 in the baseline configuration and up to 200 in the ultimate scenario. During the HL-LHC data-taking period, ATLAS aims to collect about 4000/fb of integrated luminosity [28].

The installation of the Phase-2 upgrades will take place during Long Shutdown 3, scheduled from mid-2026 until the end of 2029. This will be followed by Run 4, marking the start of HL-LHC operation in 2030. Figure 3.1 shows the current LHC schedule from Run 3 through Run 5, highlighting the LS3 period and the runs in which the upgraded ATLAS detector and TDAQ system will operate.

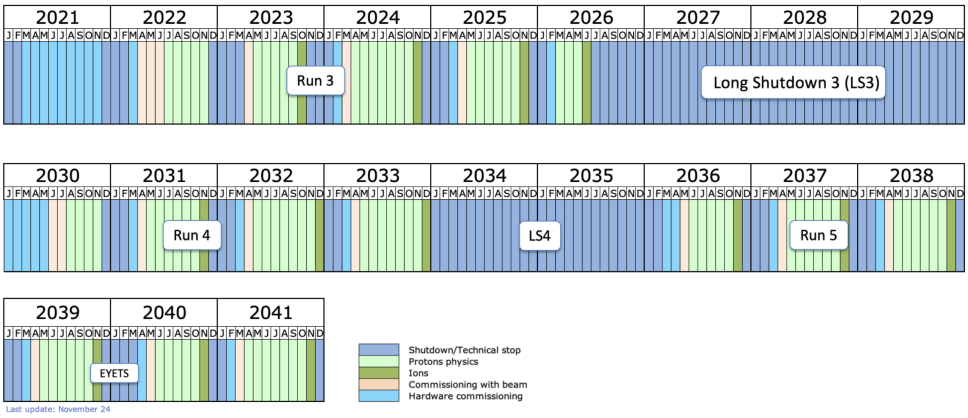

Figure 3.1: Planned LHC operational schedule from Run-3 through Run 5, showing Long Shutdown 3 during which the Phase-2 upgrades will be installed. Run-4, beginning in 2030, marks the start of HL-LHC operation [2].

The HL-LHC will enable an ambitious and wide-ranging ATLAS physics program, spanning precision Higgs boson measurements, searches for new physics, high-precision

Standard Model studies, and heavy-ion and flavor physics. The primary physics goals can be grouped into the following categories:

- **Precision Higgs boson studies** – Measure Higgs boson couplings to Standard Model particles with improved precision, determine its self-coupling via rare processes such as double Higgs production, and study differential cross-sections and decay properties [28].
- **Searches for new physics** – Explore scenarios beyond the Standard Model, including heavy resonances, supersymmetric particles, exotic decays, long-lived particles, and dark matter candidates, covering both conventional and unconventional final states [28].
- **Precision Standard Model measurements** – Perform high-precision measurements of electroweak and QCD processes, vector boson scattering, and top quark properties, requiring high efficiency and minimal trigger bias [28].
- **Heavy-ion physics and flavor studies** – Investigate the quark–gluon plasma, parton energy loss mechanisms, and collect large datasets for flavor physics to study rare decays, CP violation, and mixing phenomena [28].

These physics goals impose stringent performance requirements on all ATLAS subsystems. The most significant upgrades include:

- **Inner Tracker (ITk)** – The ITk is a full replacement of the current ATLAS Inner Detector with an all-silicon tracking system optimized for the high-radiation and high-pile-up environment of the HL-LHC. It consists of a central pixel detector and a surrounding strip tracker, covering a total active area of approximately $180\,\mathrm{m}^2$ and comprising over five billion readout channels [43]. The ITk extends the tracking acceptance up to $|\eta| < 4$ and significantly improves spatial granularity achieving a better separation of closely spaced tracks in dense environments. The detector is engineered to tolerate particle fluences up to 2×10^{16} $\mathrm{n_{eq}/cm^2}$ and total ionizing doses up to 10 MGy, ensuring optimal performance throughout the full HL-LHC program [44].
- **High-Granularity Timing Detector (HGTD)** – A new sub-detector, designed to improve vertex identification and pile-up suppression in the forward region ($2.4 < |\eta| < 4.0$). It consists of two double-sided disks instrumented with silicon Low-Gain Avalanche Detectors (LGADs), achieving a time resolution better than 30–50 ps per track throughout its lifetime [45]. This enables time-based separation of collisions within the same bunch crossing, enhancing forward lepton isolation and vertex assignment. The HGTD includes over 3.6 million channels, with fine granularity

(1.3×1.3 mm^2 pads) and modular design for staged replacement in high-radiation areas [45]. The foreseen timing-resolution of 30 ps to 50 ps imposes an important constraint on the expected performance of the clock distribution, that the bunch clock can be distributed to the HGTD front-end with a phase stability that is small compared to the expected resolution.

- **Calorimeters** – For the Liquid Argon (LAr) calorimeters, all front-end and low-voltage power electronics will be replaced to withstand increased radiation and ensure high precision under elevated occupancy. The new readout chain features the front-end boards with custom ASICs for amplification, shaping, and digitization at 40 MHz, delivering per-channel digital signals to off-detector processors for energy and timing reconstruction [46]. These upgrades enable real-time energy reconstruction and refined trigger decisions with minimal latency. Meanwhile, the Tile Calorimeter will undergo a full replacement of its on- and off-detector electronics. Each cell will be digitized at 40 MHz and data will be transmitted via high-speed optical links to new preprocessors, improving the granularity and timing precision of the calorimeter input to the trigger system. The upgraded design improves radiation tolerance, while maintaining backwards compatibility with existing infrastructure during commissioning [47].

- **Muon spectrometer** – In the barrel region, the trigger coverage will be improved by installing new Resistive Plate Chambers in the inner layer. These will be complemented by compact small-diameter Monitored Drift Tubes, which are better suited to the limited radial space and provide precision tracking capability. In the endcaps, the New Small Wheels will continue to provide precision tracking and trigger signals using Micromegas and small-strip Thin Gap Chambers. The main reason behind these upgrades is to fully integrate precision tracking information from both Monitored Drift Tubes and New Small Wheel systems into the Level-0 trigger. This requires a new trigger and readout architecture capable of processing high-granularity muon data within the latency and trigger rate limits. FPGA-based off-detector processing will replace previous custom logic, enabling real-time reconstruction of track segments with minimal fake rate [48].

- **Trigger and Data Acquisition** – While the underlying hardware and bandwidth capabilities are upgraded, the architecture retains a functional split between a fast hardware trigger stage and a large-scale software-based event filtering farm [28]. As the TDAQ system is central to enabling the HL-LHC physics program, and includes the LTI that is the focus of this thesis, the remainder of this chapter details its Phase-2 upgrades.

3.1.1 Physics-driven requirements for the Phase-2 TDAQ system

The HL-LHC physics goals and the challenging running conditions translate into specific performance requirements for the Trigger and Data Acquisition system. While some of these capabilities already exist in the Run-3 system, they must be significantly improved for Phase-2. These requirements apply across the physics program and are essential to preserve sensitivity to rare processes, minimize trigger bias in precision measurements, and capture unconventional signatures.

- Real-time access to full-granularity calorimeter and muon detector data at Level-0.
- Programmable topological trigger logic, including correlation and invariant mass calculations.
- Integration of precision timing information from the HGTD for pile-up suppression.
- Increased trigger items from 512 to 1024, for a broader and more flexible trigger menu.
- Capability to sustain a Level-0 Accept (L0A) rate of 1 MHz and an output rate to storage of 10 kHz.
- Low-threshold and inclusive triggers to minimize bias in Standard Model measurements.
- Level-0 latency budget (10 µs) set by front-end buffer depth and fully exploited to implement algorithms for efficient selection of low-threshold and topologically correlated objects.

3.1.2 Limitations of the Phase-1 TDAQ system

The Run-3 ATLAS TDAQ system is fundamentally inadequate for HL-LHC conditions. The hardware-based Level-1 trigger is constrained to a maximum accept rate of 100 kHz and a maximum latency of 2.5 µs, limiting the complexity and depth of real-time selection algorithms. As pile-up increases, this leads to poor background rejection and necessitates high thresholds and simplified logic, which reduce acceptance for low-momentum or complex signatures. Additionally, the bandwidth limitations bound adaptation to changing detector or physics conditions.

3.2 The Phase-2 TDAQ System

This section provides an overview of the TDAQ upgrades for ATLAS Phase-2, with particular emphasis on the Central Trigger. Since the work presented in this thesis focuses

on the Phase-2 Local Trigger Interface (LTI), the architecture and requirements of the Central Trigger are central to the discussion. The upgraded architecture is organized as follows:

1. **Level-0 Trigger (L0):** A fully hardware-based stage operating at the full LHC bunch-crossing rate of 40 MHz. It performs real-time selections using calorimeter, muon, and forward-detector trigger primitives, reducing the event rate to at most 1 MHz. The L0A signal initiates the readout of all sub-detectors.
2. **Event Processing Farm (Event Filter, EF):** A software-based system that processes L0-accepted events. It performs complete event building from detector fragments, near-offline-quality reconstruction, and advanced event selection, reducing the final event output to approximately 10 kHz for offline storage.

Figure 3.2 shows a schematic of the upgraded TDAQ system, organized to highlight both the TTC path and the data path. At the center of the TTC chain is the CTP, which receives inputs from the Global Trigger and the MUCTPI. The CTP applies the configured trigger menu, implements deadtime management, and issues the L0A signal within the 10 µs latency budget.

On the TTC path, the L0A, bunch clock, and Orbit signals are distributed from the CTP through the new optical TTC network to the LTIs. Each LTI serves one or more FELIX systems, providing partitioning, phase alignment, and fan-out of TTC signals to the sub-detector front-end electronics. Busy and veto signals from the front-end systems are returned to the CTP via the LTIs for deadtime control and buffer protection.

On the data path, the L0A triggers the readout of detector front-ends via FELIX, which formats and transmits data fragments over high-speed optical links to the Data Handler. The Data Handler performs detector-specific processing and buffering, ensuring that fragments are correctly packaged for the Dataflow framework. The Dataflow system then assembles complete events from all sub-detectors, buffers them, and forwards them to the Event Filter for software-based selection before final storage [49].

3.3 Hardware trigger

3.3.1 Level-0 muon trigger

The Level-0 Muon (L0Muon) trigger processors are responsible for the real-time identification and reconstruction of muon trigger candidates. They process input data from the Resistive Plate Chambers in the barrel, Thin Gap Chambers in the endcap, and the New Small Wheel detector covering the forward region, with a combination of Micromegas and small-strip Thin Gap Chambers. To improve momentum resolution and reduce fake rates,

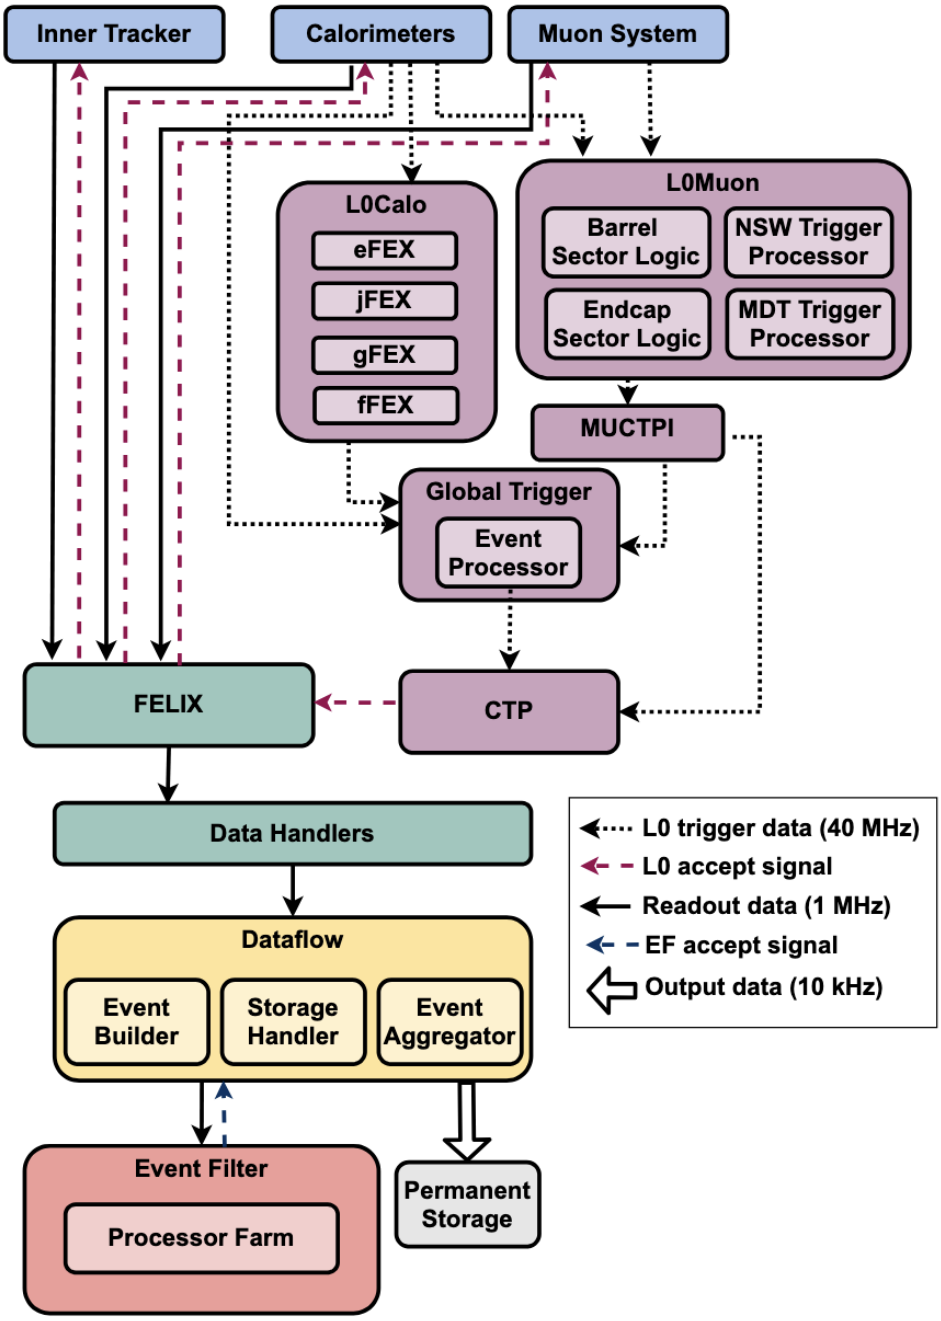

Figure 3.2: Overview of the Phase-2 ATLAS Trigger and Data Acquisition (TDAQ) architecture.

the system optionally incorporates precision information from the Monitored Drift Tubes for matching and quality refinement. Pattern recognition algorithms are implemented in FPGAs and ASICs to extract spatial and temporal coincidences, allowing for coarse p_T estimation and assignment of quality flags.

The L0Muon system is partitioned by detector region, with dedicated sector logic boards handling candidate reconstruction in parallel. Identified muon objects are categorized by p_T thresholds and spatial regions, and are then sent to the Muon-to-Central Trigger Processor Interface for further processing.

3.3.2 Level-0 calorimeter trigger

The Level-0 Calorimeter Trigger (L0Calo) processes ATLAS calorimeter data at the full bunch-crossing rate of 40 MHz, generating real-time trigger objects based on coarse grannularity information [28]. It consists of four FPGA-based feature extractor subsystems:

- **eFEX** (electron/photon/tau Feature Extractor) – identifies electromagnetic clusters and tau candidates ($|\eta| < 2.5$), applying isolation and shape selection.
- **jFEX** (jet Feature Extractor) – reconstructs jets ($|\eta| < 4.9$), applying sliding-window algorithms and pile-up subtraction.
- **gFEX** (global Feature Extractor) – computes event-level quantities such as total transverse energy, missing transverse energy, and identifies large-radius jets.
- **fFEX** (forward Feature Extractor) – processes the forward calorimeter region ($3.2 < |\eta| < 4.9$), extracting forward trigger candidates.

All calorimeter trigger objects from these subsystems are transmitted to the Global Trigger Processor, which applies additional selection logic and correlations before sending trigger decisions onward to the Central Trigger Processor.

3.3.3 Global Trigger Processor

The Global Trigger Processor (L0Global) is a key component of the Level-0 system. It receives and processes trigger objects from the calorimeter feature extractors, the MUCTPI, and zero-suppressed full-granularity calorimeter cell information each bunch crossing. Real-time algorithms exploit this input to perform event selections based on object correlations, invariant-mass and angular requirements, isolation, and refined calorimeter clustering. L0Global encodes the results into a 1024-bit Trigger Input (TIP) word per bunch crossing, which is transmitted to the CTP via optical links [28].

3.3.4 Muon-to-Central Trigger Processor Interface

The MUCTPI aggregates muon trigger candidates generated by the L0Muon processors across the barrel and endcap detector regions. It consolidates this information into muon multiplicities categorized by discrete transverse momentum (p_T) thresholds and segmented by detector region. These multiplicities are transmitted directly to the Central Trigger Processor, where they contribute to the final Level-0 Accept decision.

Hardware implementation

The MUCTPI implementation will be based on two new ATCA boards, sharing a common design with the CTP.

Functionalities

The MUCTPI functions as the central concentrator and formatter for muon trigger data. Its main task is to collect muon candidates from the sector logic boards, remove double-counted muons, and produce a set of p_T-based multiplicities. These multiplicities are sent to the CTP. In addition, the MUCTPI sends muon candidate data to the Global Trigger for topological algorithms. The MUCTPI interfaces on the trigger path are implemented with optical links.

3.3.5 Central Trigger Processor

The Central Trigger Processor integrates inputs from multiple Level-0 subsystems to produce the final L0A decision within the 10 µs latency constraint. Those inputs are received from:

- Combined calorimeter- and muon-based trigger information, including topological combinations, from the L0Global.
- Muon multiplicity counts from the MUCTPIs.
- Trigger signals from forward and specialized detectors (e.g., HGTD, ZDC, LUCID, BCM).

The CTP receives eight optical input links from L0Global and four from the MUCTPIs, and it also provides 40 electrical inputs plus 12 spare optical inputs.

Table 3.1 shows performance and hardware difference between Phase-1 and Phase-2 CTP.

Table 3.1: Comparison of CTP and other parameters between Phase-1 and Phase-2.

Parameter	Phase-1	Phase-2
Max output rate	100 kHz	1 MHz
L1/0 max latency	2.5 µs	10 µs
Architecture	14 × VME 9U	3 × ATCA
Trigger Items	512	1024

Hardware implementation

The CTP will be implemented with three custom ATCA blades: (i) CTPMI, (ii) CTPIN, and (iii) CTPCORE. In order to minimize latency, the system is based on a single Xilinx UltraScale+ FPGA, handling all the latency-critical functions (the full trigger decision path).

Trigger inputs are received primarily via high-speed optical links using Samtec FireFly modules, capable of link rates up to 14 Gb/s. However, the actual TTC links operate at fixed rates of 9.6 Gb/s from CTP to FELIX and 4.8 Gb/s from FELIX to CTP, in accordance with the TTC protocol specifications [50]. The optical transceivers thus operate well below their maximum capacity.

Functionalities

The CTP is the final decision unit in the Level-0 trigger pipeline. Its primary role is to evaluate trigger inputs from upstream processors in real time and decide, for each bunch crossing, whether an event should be read out. It does so by forming the Level-0 Accept signal, which is distributed downstream via the TTC system to initiate detector readout. It also generates additional information including the bunch crossing ID, a unique L0A counter, the Orbit signal, and trigger-type metadata, which are encoded into the TTC stream [50].

Trigger formation logic

The CTP supports up to 1024 logical trigger inputs, which can be combined to define up to 1024 trigger items. Each item corresponds to a programmable logic condition, combining trigger flags from upstream systems via coincidence, anti-coincidence, and threshold logic. These items are prescaled individually and evaluated per bunch crossing. Multiple items can fire simultaneously; their combination, subject to masking and deadtime, forms the basis for issuing the L0A. Programmable bunch masks allow selection of active bunch slots for specific trigger items [50].

Partitioning capabilities

As in the current Run-3 system, the CTP supports operation in multiple concurrent trigger partitions, allowing different sub-detector groups to run independently, each with its own set of inputs, trigger items, prescales, and deadtime configuration. In Phase-2, this capability is extended into the TTC optical distribution layer via the LTI modules, which implement partitioning at the hardware distribution stage. This enables finer granularity — down to individual TTC destinations such as specific FELIX endpoints — and higher link speeds compared to the Run-3 fan-out. The system will continue to support up to three active partitions issuing independent L0As, with busy-signal and veto handling remaining partition-aware [50].

Interfaces

Trigger inputs arrive at the CTP primarily via high-speed optical links, typically from the Global Trigger and the MUCTPI. A limited number of electrical inputs are available via front-panel LEMO connectors, supporting NIM-standard logic levels. These are used for forward detectors or auxiliary systems, calibration inputs, or signals from beam condition monitors. While most Phase-2 inputs are handled via high-speed optical links, the electrical interface ensures compatibility with systems not migrated to serial protocols [50].

The output is distributed via optical links to the Local Trigger Interfaces (LTIs), which in turn fan out the L0A and associated timing signals to FELIX boards and sub-detector front-ends. The busy signals are aggregated from FELIX via LTIs and sent back to the CTP over bi-directional optical links [50].

Deadtime handling

The CTP will implement simple deadtime, complex deadtime and sliding-window programmable facilities[50]. At the time of writing, the requirements from the front-end buffering systems for normal physics data-taking are satisfied with three sliding-window algorithms, shown in Table 3.2. Simple and complex deadtime are foreseen only for special runs [51].

Table 3.2: Sliding window parameters for HL-LHC physics data-taking.

Time window [ns]	Max number of L0As
125	4
500	8
9×10^4	128

Diagnostics and monitoring

The CTP includes extensive online monitoring and test capabilities. All trigger item decisions, prescale counters, and deadtime states can be read out per event and tracked over time. Pattern generators and snapshot memories enable injection and capture of synthetic inputs for debugging and commissioning. The system also maintains per-bunch trigger rate counters, busy statistics, and internal state summaries for diagnostic use. These features are essential for understanding performance under real data-taking conditions and for validating the behavior of complex trigger menus [50].

White Rabbit backbone

For HL-LHC the LHC machine timing is being migrated to a White Rabbit (WR)–based backbone that distributes the reference from P4 to the experiments, including ATLAS. White Rabbit is an open-source, Ethernet-based timing project that extends IEEE 1588 PTP with Synchronous Ethernet and hardware phase tracking to distribute phase-aligned clocks and absolute time with sub-nanosecond accuracy over kilometer-scale fiber networks [52, 53, 54, 55].

The 400 MHz RF, together with the derived bunch clock and Orbit signal, are generated at P4 and transmitted to the experiments over the WR network [53]. In the WR architecture, at the experiment end, WR endpoints reconstruct beam-synchronous clocks for users [52, 54]. In the ATLAS TTC chain, the CTP Machine Interface is clocked from the beam-synchronous 40 MHz bunch clock and Orbit delivered via the experiment's WR endpoint.

The WR backbone specification for the experiments requires a random jitter of $<$ 5 ps RMS at the experiment endpoints, and lab measurements of the WR-based backbone confirm jitter values compatible with this requirement (i.e. $<$5 ps RMS), with some degradation at very low offset frequencies in cascaded-switch topologies [53].

Local Trigger Interface

The Local Trigger Interface (LTI) is the distribution node in the TTC system and plays a central role in this thesis. Each LTI receives the TTC stream from the CTP and redistributes bunch clock, Orbit, and L0A signals to the associated FELIX boards in its subdetector partition. In addition to acting as a TTC fan-out, the LTI performs monitoring, phase alignment, and partition-specific TTC encoding. The hardware is implemented on a custom ATCA board, supporting bi-directional optical links, a Kintex Ultrascale+ FPGA for TTC processing, and a Xilinx Kria K26 for configuration and monitoring. A detailed description of the LTI hardware, its operation, and its importance in the Phase-2 ATLAS TDAQ upgrade is provided in Chapter 4.

Figure 3.3 shows a possible layout of the TTC network in the ATLAS counting room. Every LTI is represented by a circle. The number inside every circle is the number of FELIX boards served by that LTI

3.4 Readout and event processing

After a L0A is issued by the CTP, data from sub-detectors front-end electronics are transmitted to FELIX, assembled into events by the Dataflow system, filtered by the Event Filter, and stored for offline analysis. See Figure 3.2 for an overview of this end-to-end data path.

3.4.1 Front-End Link eXchange readout system

The Front-End Link eXchange system forms the hardware interface between the detector front-end electronics and the commodity server infrastructure of the DAQ system. FELIX receives TTC signals — including the bunch clock, Orbit signal, and L0A — from the Local Trigger Interface and distributes them to the front-end. In response to a L0A, detector front-end electronics transmit event fragments over high-speed optical links to FELIX. The FELIX FPGA transfers these fragments via a PCIe Direct Memory Access (DMA) engine into pre-allocated buffers in the host server's memory. Host-side data handler software then processes, classifies, and forwards the fragments to the Dataflow system for event building [28].

3.4.2 Data handlers and event routing

After the PCIe DMA engine on the FELIX card transfers fragments into the host, the data handlers read these buffers, categorize fragments according to BCID, L0ID, and source identifier, and apply basic integrity checks and any required format conversion to the standard ATLAS fragment format [28].

Processed fragments are then transmitted over the high-bandwidth DAQ network to the Dataflow event builder nodes. This software-mediated separation between hardware readout and network distribution enables flexible load balancing and failure containment, allowing multiple FELIX boards to be deployed across different servers and for their output to be routed independently within the DAQ system.

3.4.3 Dataflow system

The Dataflow system is responsible for reconstructing complete events from fragments received from FELIX host servers and managing their delivery to the Event Filter processing

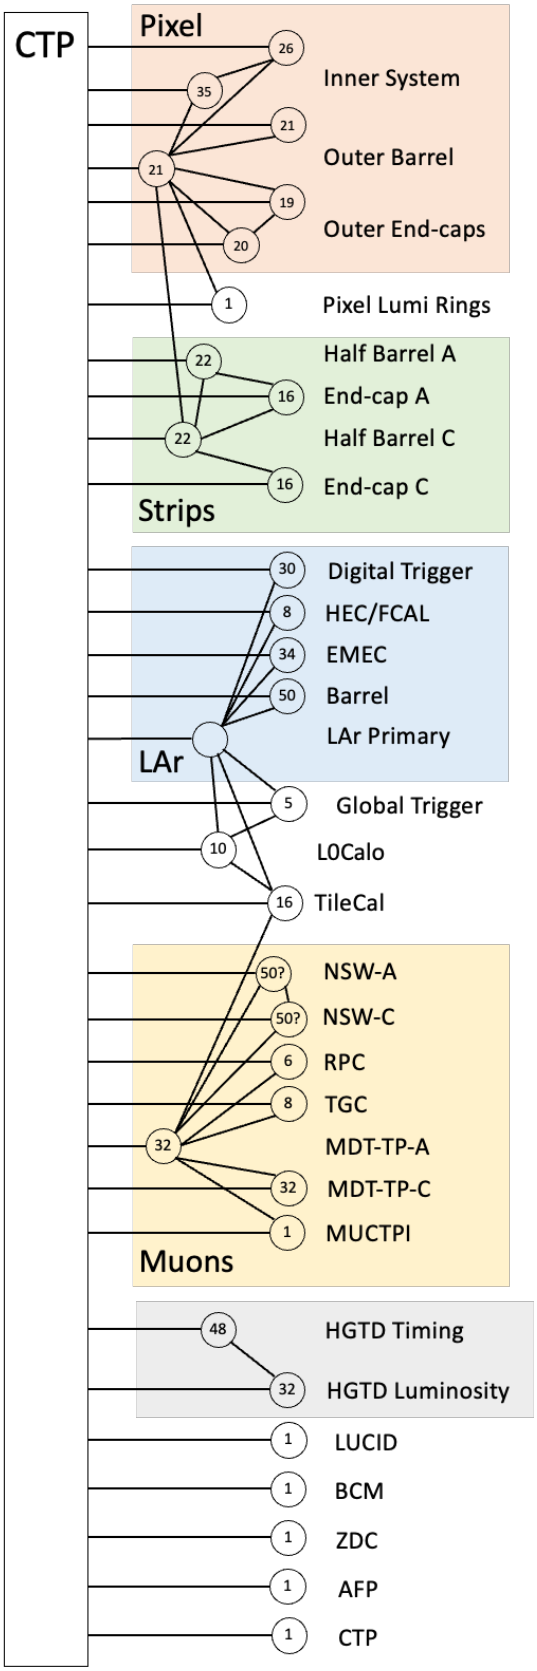

Figure 3.3: TTC network layout showing the CTP-LTI and LTI-LTI connections [56].

farm [28]. It consists of three major components:

- **Event Builder**: assembles event fragments from different sub-detectors corresponding to the same L0A, matching them using BCID, L0ID, and Source Identifier metadata from fragment headers.
- **Storage Handler**: optionally buffers fully built events prior to processing, providing resilience against transient back-pressure and decoupling the timing of readout from selection.
- **Event Aggregator**: groups complete events into output streams and forwards them over the DAQ network to the Event Filter farm.

The system operates at up to 1 MHz L0A input rate and is designed to scale horizontally with the number of processing nodes. Fragment synchronization is based on metadata carried within the readout stream (L0ID, BCID), ensuring deterministic reconstruction of events [28].

3.4.4 Event filter and final selection

The Event Filter (EF) is implemented as a large-scale software farm that performs full event reconstruction and high-level selection [28].

In Run-3, the High-Level Trigger (HLT) software farm operates in two stages: an initial fast-rejection step using simplified algorithms, followed by the Event Filter stage, which performs near-offline-quality reconstruction on the surviving events.

In the Phase-2 system, this structure is replaced by a single unified EF layer capable of executing offline-quality algorithms under low-latency conditions.

The EF receives fully built events from the Dataflow system and applies physics reconstruction algorithms such as tracking, calorimeter clustering, and vertex finding. Based on these results, it makes the final accept/reject decision. Accepted events are sent to permanent storage at a sustained rate of approximately 10 kHz [28].

The EF framework supports multi-threaded processing, heterogeneous computing (including GPU acceleration), and tight integration with offline reconstruction software to ensure consistency. This design ensures high throughput and selection efficiency under HL-LHC pile-up conditions.

Chapter 4

The Local Trigger Interface

This chapter gives an overview of the Local Trigger Interface (LTI). It introduces the main functionalities, hardware specifications, operating modes, and integration within the TDAQ system. This discussion will be preparatory to the next chapter, presenting the experimental work on LTI clock phase stability.

4.1 Overview and scope

The LTI is the crucial node connecting the CTP with the FELIX modules responsible for interfacing with sub-detector readout electronics. It distributes TTC signals using bidirectional serial optical links. In order to minimize the number of fibers, the bunch clock is embedded within the data stream at 9.6 Gb/s in FPGA hardware. This choice requires clock recovery at the receiver, which may introduce phase-stability issues. As previously discussed in Chapter 3, timing signals must maintain rigorous phase stability across the distribution chain; any deviation or drift can significantly degrade event reconstruction accuracy and therefore compromise physics performance [56].

For this reason, the LTI hardware is designed to incorporate advanced techniques to ensure deterministic and stable clock propagation under operational conditions, involving potential system resets and temperature variations [56]. This includes the use of FPGA-based phase measurement and correction methodologies, such as the Digital Dual Mixer Time Difference (DDMTD) and Transmitter Phase Interpolator (TxPI). These mechanisms allow the LTI to dynamically monitor and mitigate phase instabilities.

Moreover, the LTI design includes comprehensive monitoring and diagnostic functionalities, explicitly specified to detect and flag any potential loss of clock lock or timing anomalies [57]. These features enable real-time assessment and correction of timing stability issues.

4.2 Trigger and Timing Control signals

The LTI handles a variety of TTC signals, which originate primarily from the CTP, but some are injected locally or returned from the front-end electronics. The links going from CTP to the detector are called **downlinks**; those in the other direction are called **uplinks**. The exact origin of certain signals depends on the LTI's operating mode—either receiving from an upstream master or generating locally—which will be described in Section 4.4. Table 4.1 summarizes the main signals, their origin, direction of propagation, and relevant notes [56].

Signal	**Origin**	**Direction**	**Notes**
bunch clock (BC)	CTP/LTI or internal	Downlink	40 MHz beam-synchronous clock distributed with low jitter and fixed phase.
Orbit	CTP/LTI or internal	Downlink	Marker for each LHC revolution; used to reset the bunch crossing ID (BCID).
Level-0 Accept (L0A)	CTP/LTI or internal	Downlink	Trigger-accept signal based on global logic inputs.
L0A Counter	CTP/LTI or internal	Downlink	Incremented on each accepted trigger and sent with L0A.
BCID	CTP/LTI or internal	Downlink	bunch crossing identifier for each LHC bunch.
Trigger Type	CTP/LTI or internal	Downlink	Encodes the category of the accepted trigger.
BUSY	FELIX or internal	Uplink	Indicates readout backpressure; used to inhibit further triggers.
Lumi Block (LB)	CTP/LTI or internal	Downlink	Counter identifying a fixed-length time period, used to organize and monitor data-taking intervals.
Calibration Request	CTP/LTI or internal	Uplink	Flag indicating a request for special calibration triggers, such as pulses.
Custom Commands	CTP/LTI or internal	Downlink	Locally injected commands and signals. They can be synchronous with the BC or asynchronous.

Table 4.1: Summary of TTC signals handled by the LTI. Source: [56].

4.3 TTC interfaces

The LTI board implements two distinct types of bidirectional optical TTC interfaces:

- **LTI-in** – Four optical input links (one from the CTP and up to three from other LTIs) that receive TTC data streams. The clock is recovered from the incoming serial stream using the GT transceivers. LTI-in links are phase-monitored to establish the timing reference for downstream distribution.

- **LTI-out** – Sixty optical output links in total, comprising 52 outputs to FELIX modules and 8 outputs to other LTIs. These transmit TTC data streams including the 40 MHz bunch clock, Orbit, L0A signals, and fast commands. Each fiber operates independently.

All TTC optical links use 8b/10b encoding[1]. For the downstream TTC direction, the serial line rate is 9.6 Gb/s. With a 40 MHz bunch clock, each bunch crossing corresponds to:

$$9.6 \text{ Gb/s} \times 25 \text{ ns} = 240 \text{ serial bits.}$$

After 8b/10b decoding, only 80% of these are payload, giving:

$$240 \times 0.8 = 192 \text{ payload bits per bunch crossing.}$$

In the transceivers used on the LTI, a common choice is a 40-bit parallel datapath into the FPGA fabric. At 9.6 Gb/s this yields a user clock of:

$$\frac{9.6 \text{ Gb/s}}{40} = 240 \text{ MHz},$$

which is an integer multiple of the 40 MHz bunch clock and forms the internal processing clock domain for TTC data.

In typical ATLAS configurations, one of the LTI-in links is connected to the CTP, while the LTI-out links fan out identical TTC information to several FELIX modules. The exact link usage depends on the operational mode and partitioning.

In addition to the primary optical TTC interfaces, the LTI also supports auxiliary I/O that can be used for TTC-related tasks in special configurations or standalone setups:

- **LEMO I/O (electrical)** – Front-panel LEMO connectors for local inputs and outputs, supporting standalone lab testing or external synchronization with other sub-detector-specific hardware. They provide direct electrical access to TTC signals such as the bunch clock, Orbit signal, or L0 trigger accept.

- **Gigabit Ethernet (GbE)** – The SoM provides a GbE link to the ATLAS network infrastructure. Although not part of the TTC protocol itself, this path is essential

[1]The 8b/10b encoding is a line code that maps each 8-bit byte to a 10-bit symbol to maintain DC balance and sufficient transition density; receivers align using the unique comma character, after which they decode back to bytes.

for control and monitoring operations, such as the interaction with the ATLAS central remote control.

- **SMA Clock Output** – A SMA connector on the front panel provides a direct copy of the recovered TTC clock. It is intended for high-precision phase and jitter measurements with laboratory instruments such as oscilloscopes.
- **USB–JTAG** – A front-panel USB–JTAG port allows direct programming and debugging of the TTC FPGA. This complements the remote configuration option via the Xilinx Virtual Cable (XVC) server running on the SoM.

Figure 4.1 shows an outline diagram of the LTI interfaces

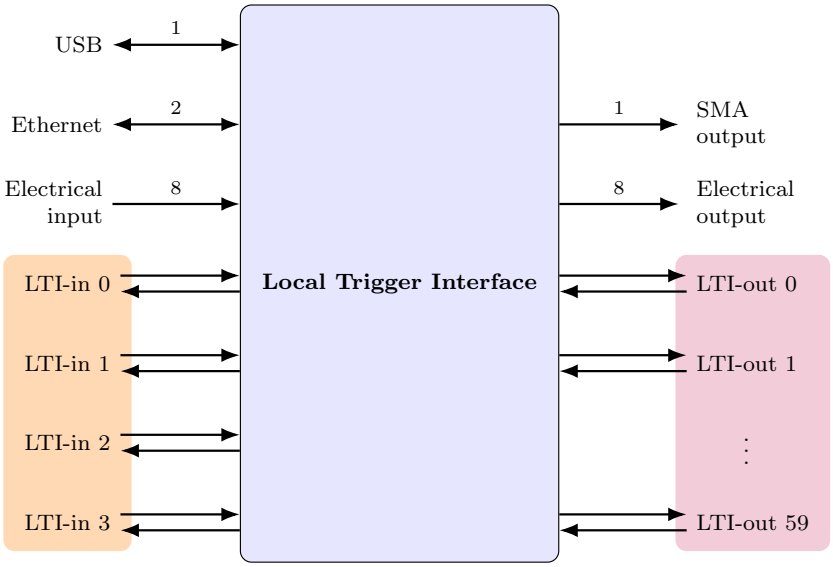

Figure 4.1: LTI interfaces.

4.4 Partitioning and LTI operation modes

The TTC system in ATLAS supports **partitioning**, i.e. dividing the detector's control and readout into independently operable units [56]. The main partition topologies are:

Physics partition: A single global TTC domain with the CTP as master for all sub-detectors.

CTP-combined partition: Multiple independent partitions, all clocked and triggered by the CTP but each with its own trigger menu and set of sub-detectors.

Standalone partition: A LTI acts as local TTC master for a partition, generating bunch clock, Orbit, and triggers without input from the CTP.

Mixed/dual-partition setups: A single LTI can serve two logical partitions, for example one CTP-synchronized and one standalone. Both partitions share the same bunch clock and Orbit.

Figure 4.2 illustrates examples of TTC partition configurations in ATLAS. Panel (a) shows the static cabling between the CTP and sub-detector LTIs, including LTI–LTI links. Panels (b)–(d) highlight different operational partition types:

(b) a single physics partition with the CTP as global TTC master;

(c) a CTP-combined configuration with multiple independent partitions clocked from the CTP;

(d) a standalone partition where one LTI acts as TTC master for its own sub-detector and for another, without input from the CTP.

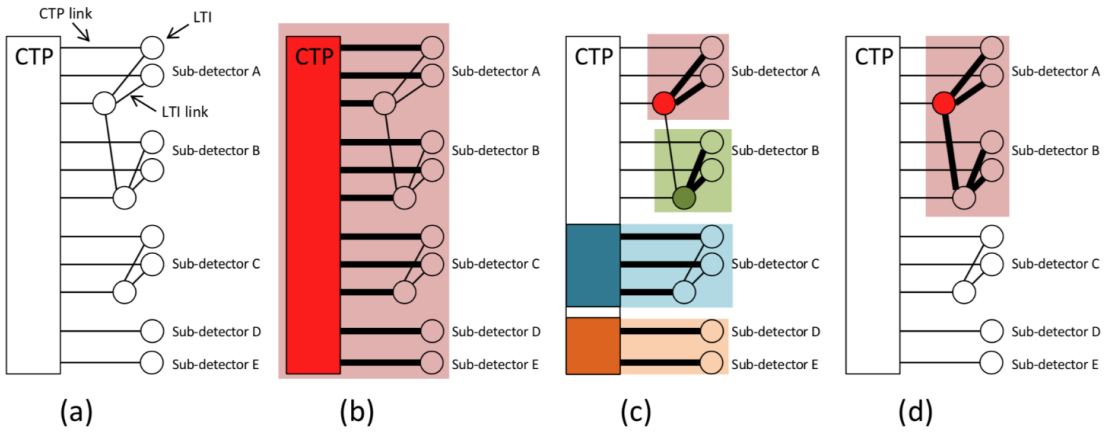

Figure 4.2: Examples of TTC partition configurations in ATLAS. (a) Static optical cabling between the CTP and LTIs, including LTI–LTI links. (b) Physics partition with the CTP as global master. (c) CTP-combined mode with multiple independent CTP-driven partitions. (d) Standalone mode with one LTI acting as TTC master for multiple sub-detectors.

Within a given partition topology, each LTI can operate in one of the following modes:

Slave mode: The LTI locks to the TTC stream from the CTP or an upstream master LTI via one of its LTI-in links, recovering the clock and commands and fanning them out to FELIX modules.

Master mode: The LTI generates TTC signals internally (*mini-CTP* functionality), acting as the clock and trigger source for downstream FELIX modules or other LTIs in slave mode.

Dual-domain mode: The LTI is split into two independent TTC domains sharing the same bunch clock and Orbit signal, but with their own set of outputs and potentially different modes (one master, one slave, etc.).

4.5 LTI hardware architecture

The Local Trigger Interface (LTI) is implemented as a single ATCA blade, offering flexibility and integration to support multiple operational modes, ranging from physics data-taking to standalone subsystem testing [58]. Figure 4.3 provides a high-level overview of the LTI hardware architecture.

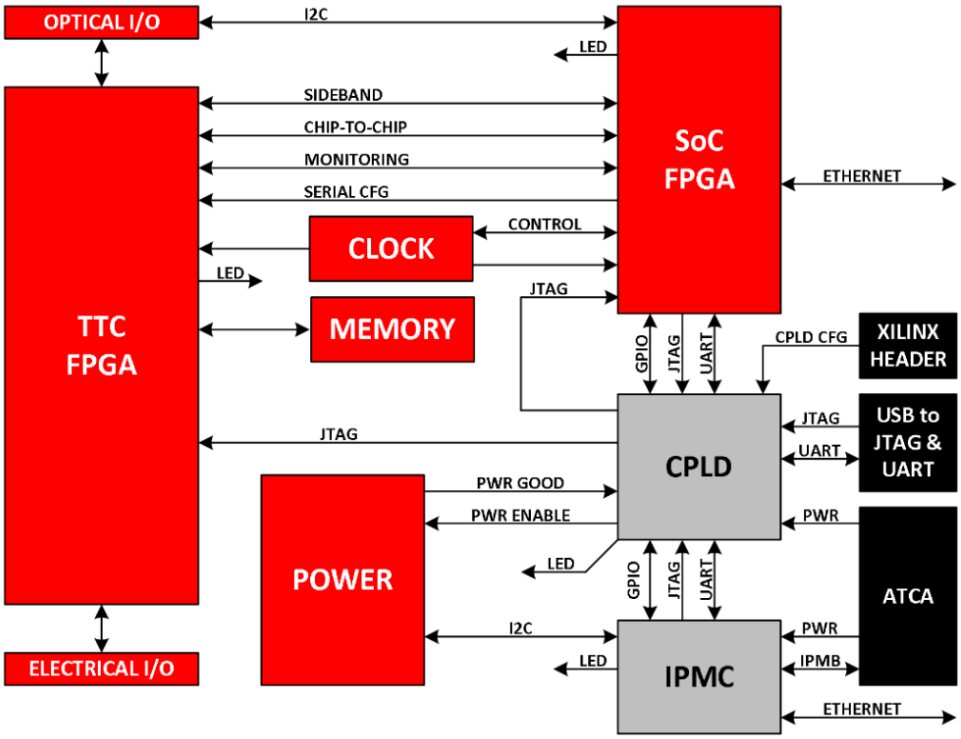

Figure 4.3: LTI board block diagram [58].

An overview of the primary hardware components includes:

- **TTC processing FPGA** – Xilinx Kintex UltraScale+ KU15P FPGA.
- **System-on-Module (SoM)** – Xilinx Kria K26 for configuration and monitoring.
- **Optical I/O** – Samtec FireFly transceivers and MPO connectors.
- **Electrical I/O** – Front-panel SMA and LEMO connectors.
- **Clocking system** – Si5395A-based clock generator and jitter cleaner.
- **Memory** – DDR4 for diagnostics, and pattern replay.

- **Power and cooling** – ATCA-compliant power management.
- **Modularity and monitoring** — Intelligent Platform Management Controller (IPMC) and Complex Programmable Logic Device (CPLD) for low-level board control tasks.

In addition to the main data and control paths, several auxiliary interfaces connect the SoM, the TTC FPGA, and other board-level components [58].

- **C2C (Chip-to-Chip)** — A parallel, memory-mapped, bidirectional bus between the SoM and the FPGA, used for configuration, register access, and runtime control.
- **I^{2}C (Inter-Integrated Circuit)** — A low-speed serial bus used for board-level control and monitoring. The SoM and IPMC use I^{2}C to access temperature sensors, EEPROMs, power monitors, and optical transceiver diagnostics.
- **UART** — A serial console interface used for Linux shell access, boot diagnostics, and debug output.
- **JTAG** — A test and debug interface providing access to internal FPGA signals. It can be accessed via front-panel USB or remotely using XVC over Ethernet.
- **GPIO (General-Purpose I/O)** — A set of configurable digital lines for status, handshaking, and auxiliary control.
- **Sideband signals** — Additional dedicated signals between the SoM and FPGA used for synchronization, interrupts, or hardware events not mapped to the C2C protocol.

4.5.1 FPGA-based processing

The Xilinx Kintex UltraScale+ KU15P is the TTC FPGA handling all the time-critical signal data processing. It manages TTC signal distribution and busy signal collection from the sub-detectors. In the rest of the document we will refer to the TTC FPGA simply as FPGA.

The KU15P includes a total of 76 high-speed multi-gigabit transceivers (MGTs): 44 GTHe4 and 32 GTYe4. The suffix “e4” identifies the fourth generation, which will be omitted in the rest of this document. Earlier TTC distribution studies have shown that the type of transceiver can affect the characteristics of the recovered clock — in particular, its jitter performance and the determinism of its phase after a reset. These aspects are critical for the LTI, where the recovered clock serves as the timing reference for all downstream links. The mix of GTH and GTY transceivers in the KU15P therefore allows

us to compare and optimize link configurations for phase stability at the operational rate of 9.6 Gb/s [58].

These transceivers are embedded hard blocks within the FPGA, optimized for high-speed serial communication. Each includes a serializer/deserializer (SerDes), a clock data recovery unit[2] (CDR), and programmable phase-locked loops (PLLs) to support deterministic link timing.

The configuration bitstream during normal operations is provided by the System-on-Module at startup. Nevertheless, it is possible to configure the TTC FPGA using the USB-JTAG front-panel port.

4.5.2 System-on-Module (SoM)

A Xilinx Kria K26 SoM (which is based on the Zynq UltraScale+ MPSoC) provides the means of control, monitoring, and configuration using the Processing System (PS), the portion of the device containing general-purpose processors, memory controllers, and standard I/O interfaces. It features a 64-bit ARM Cortex-A53 running the 64-bit AlmaLinux operating system. The Programming Logic (PL) has 4 GTH transceivers and offers a moderate amount of logic resources comparable to a mid-range stand-alone FPGA [58].

The collection of software components running on the LTI SoM is built using PetaLinux, a Xilinx-provided toolchain that automates the construction of a full bootable Linux image tailored to the target hardware [58]. This collection includes everything from early-stage bootloaders to the Linux kernel and user-space applications, enabling the initialization and operation of the control-plane functionality of the LTI.

The following sequence outlines the steps performed during a typical power-up event of the SoM, beginning from the hardware reset and culminating in the launch of runtime services:

1. **Power-on and Boot ROM Execution:** After power-on reset, the internal Boot ROM code in the PS is triggered. It reads the boot mode configuration and selects Quad Serial Peripheral Interface (QSPI) flash as the boot source.
2. **Bootloader Execution:** The bootloader process begins with a custom Secondary Initial Program Loader (SIPL), which performs low-level hardware initialization. It then transfers control to U-Boot, a standard embedded bootloader that completes system setup, loads the Linux kernel and device tree into DDR memory, and starts the operating system.
3. **Runtime Services:** Once Linux is fully booted, network services, configuration daemons, and custom tools for managing the main FPGA over the C2C link are

[2]Clock and Data Recovery reconstructs a sampling clock from the incoming serial transitions; after lock the absolute phase is typically not deterministic across resets.

launched.

The Kria SoM interfaces with the main FPGA via a C2C communication scheme, which provides a low-latency, memory-mapped parallel bus. In this architecture, the SoM acts as the master and the FPGA as the slave, enabling the SoM to read and write FPGA internal registers as if they were local memory. This link is used both for configuration at boot time and for runtime monitoring and control. The SoM also manages communication with the ATLAS network infrastructure via Gigabit Ethernet, and provides additional control interfaces such as UART, I^2C, and JTAG via XVC [58].

4.5.3 Optical interfaces

All the LTI optical links are implemented using Samtec FireFly 12-way transceiver modules [58]. Each FireFly optical transceiver module supports 12 TX and 12 RX lanes, with 24 fibers routed through a single 24-way MPO (Multi-fiber Push-On) connector. The FireFly modules are rated 14 Gb/s, which abundantly covers the requirements of the FPGA transceivers.

Figure 4.4 shows an example of assembly with two Firefly optical transceivers linked to one MPO connector.

Figure 4.4: Two Firefly modules connected to one MPO output. Source: [59].

The LTI-out links require five MPO connectors, using ten FireFly modules. The GTH transceivers are distributed over 3 Firefly pairs, and the GTY over 2 Firefly pairs. No Firefly pair serves a mix of transceivers.

The LTI-in functionality requires four links, whereas the available FireFly modules provide twelve. The additional capacity is used to implement three receiver connection schemes for mitigation of phase-instability effects.

1. **Normal mode** - Direct connection from optics to the MGT. The clock is recovered from the serial stream.
2. **Fanout mode** - A fanout chip distributes optics data to the MGT and an FPGA input for phase monitoring.
3. **Clock fiber mode** - The clock is received through a dedicated fiber and routed to the jitter cleaner. This mode requires a separate fiber per link, reducing the maximum number of FELIX cards per LTI by a factor of 2.

These operational modes reflect the flexibility foreseen in the TTC receiver specifications [56], allowing the LTI to adapt to varying FELIX link conditions and clock distribution topologies while preserving timing quality.

4.5.4 Clocking system

Timing precision is ensured using a low-jitter clock generator and a jitter cleaner. They are implemented using two Si5395A chips running in different modes, and configured by the SoM via SPI [58]. The clock generator produces a stable clock signal when the LTI operates standalone without an external reference. The jitter cleaner outputs a low-jitter version (sub-100 fs RMS phase jitter) of the recovered clock from the transceiver. High-speed multiplexers route clocks inside the LTI board without adding significant jitter.

4.5.5 Memory

A single DDR4 memory bank of 4 GB is included on the LTI, connected to the TTC FPGA [58]. It is used for storing large data patterns, captured TTC traffic, and diagnostic snapshots, as well as for replay functions. Typical applications include:

- **Pattern generation** — storing and streaming TTC frames to test downstream FELIX or front-end electronics.
- **Snapshot capture** — recording TTC data over extended periods for later offline analysis.
- **Replay** — reproducing previously captured sequences to reproduce specific timing or trigger conditions during debugging.

This local storage enables fully self-contained standalone testing without requiring continuous upstream TTC input [58].

4.5.6 Power and cooling

The dissipation limit for ATLAS ATCA blades is 350 W, offering a large margin over the LTI's worst-case estimates of 120 W [58].

4.5.7 Modularity and monitoring

The LTI includes an IPMC and hardware monitoring for integration with the ATLAS interface for slow control, i.e. the basic hardware control and monitoring of the detector

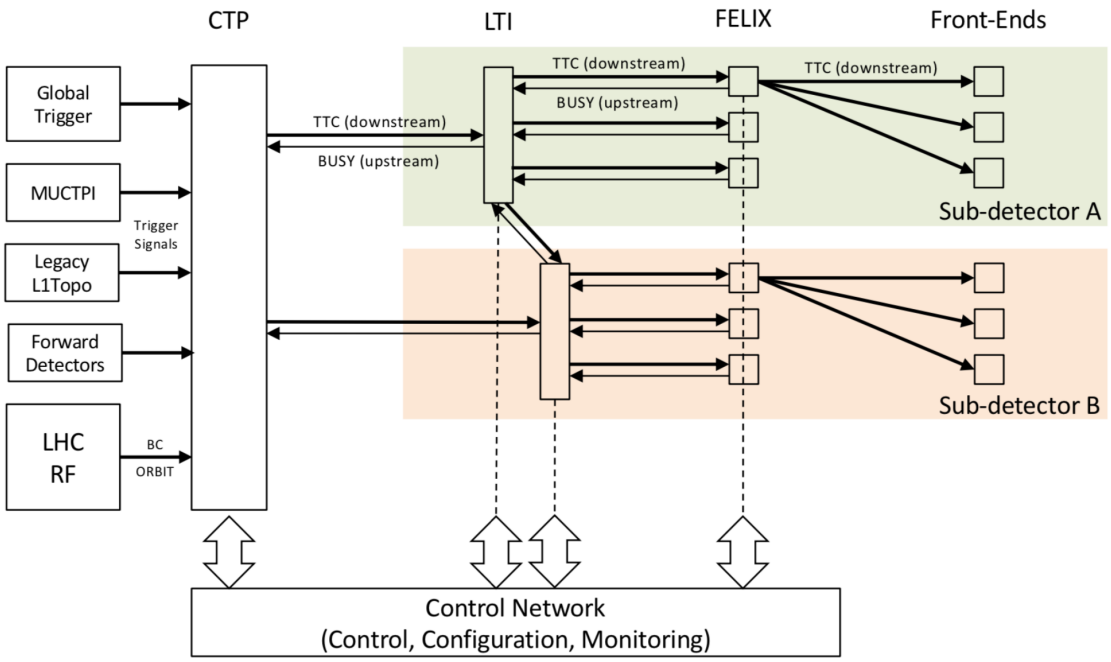

Figure 4.5: TTC context diagram [56].

infrastructure. The IPMC enables remote control of the LTI power state, supports hot-swap operations, and provides real-time monitoring of voltages, temperatures, and fan tray status [58].

Furthermore, the LTI includes a CPLD, a logic device with non-volatile configuration memory, to manage early power-up tasks before the SoM and FPGA are initialized. It handles power sequencing, voltage-level translation, JTAG/UART routing, and mode selection, ensuring correct startup and providing essential low-level control [58].

4.6 Integration in the TTC distribution

The TTC distribution from the CTP to the sub-detectors is shown in Figure 4.5. In this context, a "hop" indicates the connection between two consecutive stages.

Considering the case of the HGTD, the entire TTC distribution chain is illustrated in Figure 4.6, with off-detector components shown in blue and on-detector components in red. Each hop introduces its own contribution to the total timing jitter. The system design must therefore ensure that the combined jitter from all stages remains below the 30 ps requirement, as discussed in Section 3.1.

The **lpGBT** is a radiation-tolerant high-speed transceiver ASIC developed at CERN. It multiplexes data, clock, and control signals between front-end electronics and off-detector systems over optical links, while providing deterministic latency and robust signal integrity [60, 9].

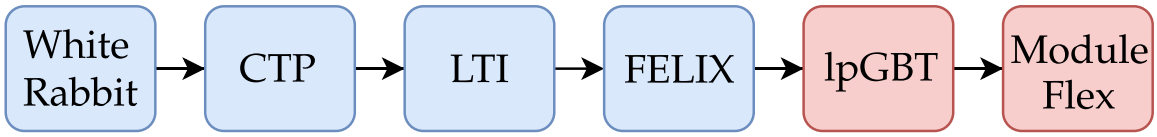

Figure 4.6: TTC distribution for the HGTD, from the White Rabbit to the front-end electronics.

The **Module Flex** is a flexible printed circuit board mounted directly on detector modules. It hosts the lpGBT and associated front-end components, serving as the interface between the sensor and the optical TTC readout path [9].

4.7 Integration with the TDAQ system

The LTI integrates into the ATLAS trigger and data acquisition system through a modular software stack running on its embedded System-on-Module (SoM). This stack supports configuration, control, monitoring, and diagnostics, enabling the participation of the LTI in global run control workflows.

4.7.1 Modular software architecture

The control software for the LTI is structured around a component known as the **LtiModule**. This component abstracts the low-level interactions with the FPGA firmware, providing a clean and modular interface for tasks such as link management, clock configuration, and diagnostic access. The development of the LtiModule package took advantage of hwcompiler, as illustrated for the MUCTPI in Section 2.2.11. The author was heavily involved in the implementation of the LtiModule and its testing.

Each instance of LtiModule can be configured at the start of a run by software programming to operate in different TTC roles—for example, as a slave recovering a clock from upstream, or as a local master generating its own TTC signals. These roles are selected by reading configuration descriptors from a text-based configuration file, which permits defining every aspect of the LTI: from the partitioning of links, to how clocks are sourced, and whether the trigger logic is internal or externally driven.

On the software side, this was realized through a C++ configuration model in which each aspect of the LTI was encapsulated in its own class, with members representing the relevant parameters. To handle optional or context-dependent fields, the `std::optional` construct was adopted, allowing descriptors to remain flexible without resorting to hard-coded defaults. This approach ensured that the same configuration files could be reused across different roles and environments while preserving type safety and simplifying validation. The implementation of this configuration software was carried out by the author,

who was directly responsible for its design and integration.

4.7.2 Configuration abstraction and reuse

Rather than requiring users to manage detailed hardware-level parameters directly, the system supports a layered configuration model. High-level configuration objects define operational roles (e.g., TTC master or slave), while referring to low-level descriptors that capture the precise firmware and clocking setup. This abstraction allows reuse of configuration templates across different environments—such as test benches, standalone subsystems, or full detector operations—without requiring firmware-level expertise from the user.

The configuration objects are stored in structured, persistent text-based configuration files and can be edited or selected through graphical interfaces or scripting tools. In complex environments, these files may be organized into hierarchical databases to support system-wide consistency and versioning.

4.7.3 LTI Monitoring Integration

The LTI participates in centralized remote control through a state-machine-driven interface. As described in Section 2.2.11, the TDAQ system defines a finite set of operational states (such as `INITIALIZED`, `CONFIGURED`, `RUNNING`, and `STOPPED`). Any LTI monitoring or control application integrated in the TDAQ environment would use the `LtiModule`, since it implements the logic and the corresponding firmware actions—such as enabling TTC receivers, resetting internal counters, or read/write a value to a specific register. For example, it is expected that all the LTI links are up in `CONFIGURED` state.

This architecture ensures that the LTI behaves predictably and reproducibly within coordinated data-taking workflows, while still allowing for standalone use in isolated setups where global remote control is not present.

4.8 Diagnostic, monitoring and control

The LTI includes comprehensive infrastructure for diagnostics, runtime monitoring, and board-level control. These features are essential both during detector operations and in standalone environments such as test benches or calibration setups. The goal is to provide full visibility into timing, link health, and configuration state.

4.8.1 Runtime monitoring and status reporting

The LTI continuously monitors the health of its transceivers, clock domains, and TTC protocol state. Key status metrics include:

- Lock status of each TTC receiver and transmitter.
- Detected loss-of-lock (LOS) or clock slips.
- Phase offset between TTC recovered and reference clocks (via DDMTD).
- FIFO occupancy, frame counters, and frame check errors.
- BUSY status returned from front-end electronics.

These quantities are made available through memory-mapped registers accessible from the LtiModule API.

4.8.2 Board-level control and reset handling

Through the System-on-Module, the LTI supports full board management functions:

- Selective transceiver resets (master/slave transmitter/receiver independently or per link group).
- Clock source switching (external vs. internal).
- Firmware reconfiguration and soft reboot of individual components.
- Snapshots and replays of TTC frames for testing.

These operations are exposed through command-line interfaces and APIs, enabling both scripted automation and interactive control. They are also tied into the run control FSM, allowing transitions like `CONFIGURE` or `RESET` to trigger defined control sequences.

4.8.3 Environmental monitoring and protection

The LTI hardware includes sensors and controllers for real-time environmental monitoring:

- Voltage and current levels on critical power rails.
- Temperature measurements from FPGA, SoM, and power modules.
- Optical transceiver diagnostics (e.g., received power, laser bias).

Alarms can be raised when thresholds are violated, and the board management firmware can react by disabling subsystems or requesting operator intervention. These protections ensure reliable long-term operation in both high-availability environments and lab development setups.

4.8.4 Debugging and traceability

For advanced diagnostics, the LTI supports multiple debug interfaces:

- JTAG access to the FPGA via USB or over Ethernet (virtual cable).
- UART console access to the embedded Linux system.
- Snapshot and playback of TTC frames, using the on-board DDR4 memory.

These features allow developers to exercise TTC logic without requiring a live input, debug low-level firmware, and correlate anomalies with environmental or configuration states.

4.9 In-situ phase monitoring and control

In the broader context of phase-stabilization strategies for high-speed links, the LTI design principles incorporate timing-distribution concepts from the TCLink (Timing Compensated Link) system developed at CERN [61]. TCLink is a lightweight, protocol-agnostic FPGA IP core designed for systems requiring high long-term phase stability. It combines in-situ phase measurement via DDMTD and active correction through a phase interpolator, forming a closed-loop control system. Although the LTI firmware does not instantiate the TCLink core directly, its monitoring and correction infrastructure adheres to the same architectural principles: a bidirectional synchronous link, roundtrip phase observation, and programmable control over the correction loop. In the LTI, this control logic is implemented in software rather than in dedicated hardware. Figure 4.7 shows a block diagram of TCLink.

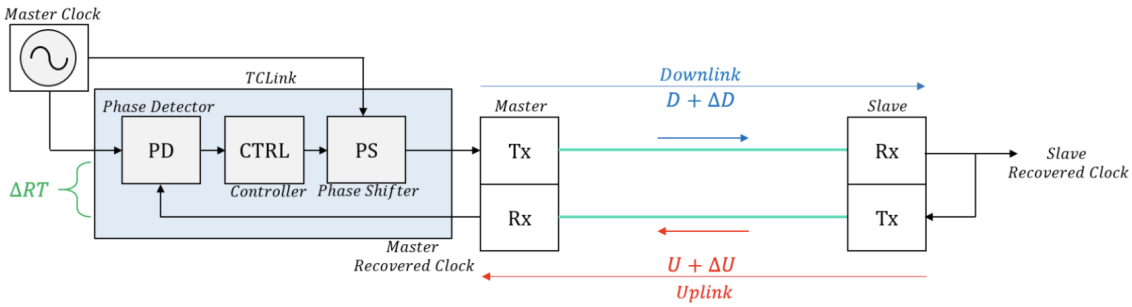

Figure 4.7: TCLink simplified block diagram. On the figure the downlink and uplink have phase variations ΔD and ΔU respectively, while the round-trip phase variation at the phase detector is ΔRT. It is assumed that at the first order ΔD is proportional to ΔRT [61].

In normal operation, the LTI requires a phase measurement method that is available in-situ, low-cost, and deployable on every board without relying on bulky laboratory

equipment. Traditional high-precision tools, such as oscilloscopes, are impractical for continuous use in the detector—especially since each LTI features 60 independent LTI-out links, making individual monitoring infeasible.

To meet these constraints, the LTI implements dedicated firmware logic in the FPGA for direct, real-time phase measurements. This approach enables continuous monitoring during data-taking without external hardware. In this thesis, the FPGA-based method is thoroughly tested and characterized in the laboratory using high performance oscilloscopes.

The remainder of this section introduces the two key FPGA-based tools used for in-situ phase monitoring and control: the DDMTD, which performs high-resolution phase measurements, and the TxPI, which applies fine-grained phase adjustments to correct for phase drifts.

4.9.1 Digital Dual Mixer Time Difference

The DDMTD is a core component of the White Rabbit project and a digital adaptation of the traditional Dual Mixer Time Difference (DMTD) technique [62]. It enables phase difference measurements between two isofrequency signals with picosecond-level accuracy [63].

This method is used for precise phase measurements within the LTI module. Figure 4.8 shows the internal structure of the DDMTD circuit.

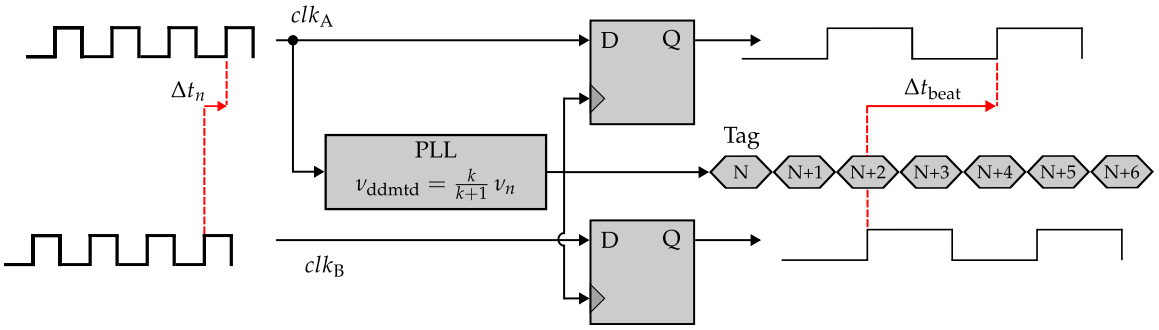

Figure 4.8: DDMTD Diagram, adapted from [64].

The DDMTD operates on two clock signals, A and B, at the same nominal frequency ν_n and period T_n, but with a phase offset ϕ. In the LTI, A is the recovered clock from the LTI-out RX[3] and B is the reference clock from the jitter cleaner. Both run nominally at 240 MHz ($T_n \approx 4.158$ ns).

A helper PLL generates a third clock, ν_{ddmtd}, slightly offset from ν_n as:

$$\nu_{\text{ddmtd}} = \frac{k}{k+1}\nu_n, \tag{4.1}$$

[3]After a round trip: LTI-out TX → LTI-in RX → LTI-in TX → LTI-out RX.

where k is a design-specific integer. For the LTI, $k = 16383$ gives $\nu_{\text{ddmtd}} \approx 240.459$ MHz. This produces a beat frequency

$$\nu_{\text{beat}} = \nu_n - \nu_{\text{ddmtd}} = \frac{\nu_n}{k+1} \approx 14.68 \text{ kHz}, \tag{4.2}$$

slowing the relative phase drift by a factor $M = k + 1 = 16384$. The smallest distinguishable step in the nominal domain is therefore $\frac{T_n}{M} = \frac{4.158\,\text{ns}}{16384} = 0.254\,\text{ps}$, which is the equivalent time resolution achieved by the system.

Signals A and B are digitally mixed with ν_{ddmtd} using D-type flip-flops, yielding outputs modulated at ν_{beat}. This allows phase differences to be measured at low frequency with simple digital logic. Each detected transition in the beat domain is assigned a tag field, which encodes the fine position of the edge within a beat-clock cycle. Combined with a coarse counter, this tag provides sub-cycle accuracy for Δt_{beat}.

The beat-domain offset is related to the nominal-domain offset through the frequency ratio:

$$\Delta t_{\text{n}} \nu_n = \Delta t_{\text{beat}} \nu_{\text{beat}}, \tag{4.3}$$

which directly leads to the expression for the phase offset:

$$\phi_{\text{offset}} = 2\pi \Delta t_{\text{n}} \nu_n = 2\pi \Delta t_{\text{beat}} \nu_{\text{beat}}. \tag{4.4}$$

For $k = 16383$, this scaling yields sub-picosecond equivalent resolution while maintaining an update rate of Equation (4.2) . The magnification effect is analogous to recording a fast-moving object with a high-speed camera: by slowing the apparent motion, fine details that would otherwise be blurred become measurable, allowing picosecond-scale phase differences to be resolved with standard digital logic. Figure 4.9 illustrates this process: the original high-frequency clocks are sampled with the slightly offset ν_{ddmtd}, producing slow beat signals whose relative timing directly encodes the phase difference at the nominal frequency.

DDMTD measurement procedure

The beat signals are accumulated using a digital phase accumulator implemented in the FPGA. This phase accumulator captures the averaged phase offset in a 28-bit register, accompanied by a valid bit and a tag field to ensure readout uniqueness. The DDMTD firmware can be configured to use a 12-bit integer to average many consecutive DDMTD measurements in the firmware. The steps to perform a DDMTD measurement can be summarized as follow:

1. Clear the `valid_tag` bit to initialize a new measurement cycle.

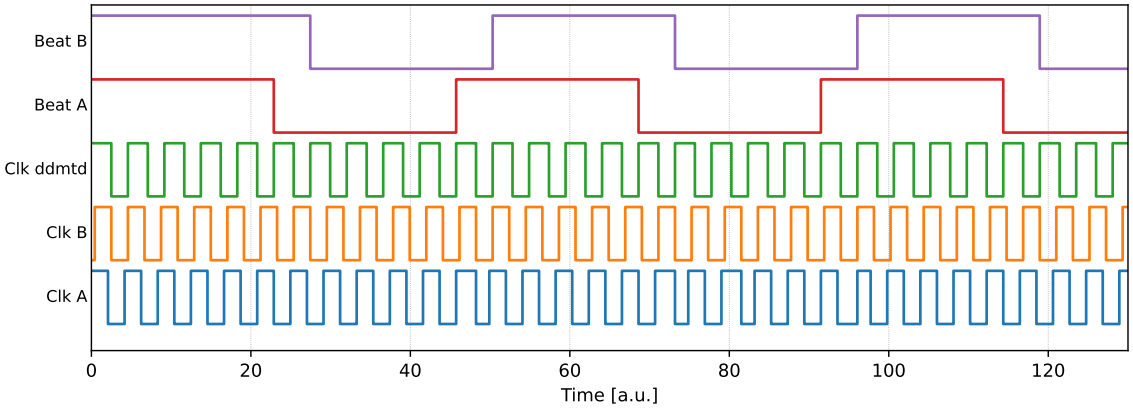

Figure 4.9: Example of DDMTD operation. The two input clocks (Clk A and Clk B) are close in frequency, with Clk ddmtd providing the slight offset. Sampling each input with Clk ddmtd produces the beat signals (Beat A and Beat B), whose slow relative transitions reflect the original phase difference with a magnification factor of $M = k + 1$. Values are scaled for readability; time is in arbitrary units.

2. Wait until the `valid_tag` bit is asserted (logic high).
3. Read the phase-accumulator register.
4. Verify the tag field to confirm measurement uniqueness, then clear the `valid_tag` bit.

Deglitch filtering

The DDMTD firmware includes a deglitch stage after the beat signal sampling. This stage rejects edges that occur within fewer than a configurable number of DDMTD clock cycles after the previous accepted transition. The goal is to suppress glitches due to metastability, or noise, which would otherwise distort the phase-accumulator value. The deglitch threshold is set as an integer count of DDMTD cycles; transitions below this threshold are ignored. In the standard White Rabbit implementation, the default value is 1966, but the parameter can be adjusted in firmware. The impact of this parameter on measurement resolution is evaluated in Section 5.3.2.

4.9.2 Transmitter phase interpolator

The TxPI is an AMD transceiver-integrated block that enables fine-grained adjustment of the transmitted data phase for timing calibration. While the DDMTD is a passive phase-measurement tool, the TxPI functions as an active phase-control mechanism. To leverage this capability, the High-Precision Timing Distribution (HPTD) group at CERN developed a dedicated TxPI controller, which programs the TxPI to apply controlled phase shifts at the transmitter output.

The TxPI is implemented within the transmitter side of each transceiver and allows step-wise shifts of the output phase, typically in increments of a few picoseconds. The nominal step size s is given by [65]:

$$s = \frac{\text{UI [ps]}}{64} = \frac{104.17 \text{ ps}}{64} \simeq 1.63 \text{ ps} \tag{4.5}$$

Here, the unit interval (UI) denotes one period of the serial data rate, expressed in picoseconds; for a line rate of 9.6 Gb/s it corresponds to 104.17 ps.

The phase adjustment is applied to the serial data stream, altering the recovered clock at the receiver. This makes it a useful tool for system-wide phase alignment, especially when combined with local DDMTD-based measurements.

In Chapter 5, the TxPI step size is measured in situ using the DDMTD method and cross-checked with oscilloscope measurements.

4.10 Transceiver operation and link initialization

This section presents an overview of the low-level hardware of the serial transceiver architecture and the link bring-up process mechanism.

4.10.1 Conventions for high-speed serial links

Link (duplex) and directional links In this document a link is a bidirectional (duplex) association between two endpoints, composed of two directional links: the downlink (LTI $\rightarrow$ endpoint) and the uplink (endpoint $\rightarrow$ LTI). Each directional link is a self-contained serial stream with its own CDR, alignment, and error counters.

Whenever behavior is direction-specific (e.g., CDR lock, comma align), we name the direction explicitly as downlink RX or uplink RX.

Lane and physical composition A lane is a single serialization path in an MGT. One directional link occupies exactly one MGT lane and one fiber; therefore, a duplex link uses two MGT lanes and two fibers. The optical module (Samtec FireFly) provides multiple independent TX and RX lanes, aggregated with the MPO connector.

Transceiver (MGT) An MGT is an FPGA transceiver block that performs serialization, deserialization, and clock/data recovery. MGTs are typically organized in quads that share a low-jitter reference and phase-locked loops for line-rate generation. In this document, "MGT" refers specifically to the FPGA transceiver block, while "optical module" denotes the external electro-optical device.

Transmitter (TX) The TX block of an MGT accepts parallel words, applies line encoding (e.g., 8b/10b) if enabled, serializes the data to the configured line rate, and drives the electrical interface toward the optical module. Some devices also include a TxPI for fine phase adjustment for system alignment.

Receiver (RX) The RX block recovers the sampling clock from the incoming serial stream using a CDR, deserializes the data into parallel words, performs word alignment (e.g., via comma detection), decodes 8b/10b symbols if present, and reports status signals such as loss-of-lock or disparity errors.

Optical module The optical module is the board-level device that converts between the MGT electrical interface and optical fibers. Each module exposes multiple TX and RX lanes via a multi-fiber connector.

4.10.2 Transceiver architecture and serial communication

As introduced in Section 4.5, the FPGA includes high-speed GTH and GTY transceivers that handle all TTC serial communication over optical links. These transceivers are physically embedded hard blocks and are grouped into quads—sets of four transceivers sharing a common Quad PLL (QPLL). The QPLL generates the serial line clock from an external low-jitter reference provided by the onboard clocking system, ensuring phase consistency across all lanes within the quad.

Each TTC stream is transmitted at 9.6 Gb/s using 8b/10b encoding, which ensures DC balance and provides sufficient transition density for robust Clock and Data Recovery. On reception, the serial bitstream is deserialized into 40-bit parallel words clocked at 240 MHz, each containing four 10-bit symbols. These are decoded into 8-bit TTC command bytes, with frame alignment achieved through a K28.5 comma character placed in a reserved data frame position. The recovered clock is divided by six to produce the 40 MHz bunch clock, and then routed through a jitter-cleaning PLL.

4.10.3 Link bring-up and synchronization

The TTC link bring-up process is the sequence of events that establishes a reliable and phase-stable serial communication channel between a transmitter and receiver. This occurs after power-up, FPGA reconfiguration, or any transceiver reset. The process ensures both data integrity and timing synchronization. It builds upon the transceiver architecture described in Section 4.10.2.

The typical link initialization sequence proceeds as follows:

1. **Master LTI-out transmitter PLL Lock and Data Transmission** — The master LTI-out transmitter's QPLL locks to a low-jitter reference clock, enabling generation of a stable serial stream. TTC frames are serialized and transmitted at 9.6 Gb/s using 8b/10b encoding.
2. **Slave LTI-in receiver Clock Data Recovery** — The slave LTI-in receiver's internal PLL performs clock recovery from the incoming bitstream. This step determines the sampling phase for deserialization.
3. **Deserialization and Word Alignment** — The incoming serial data is deserialized into 40-bit parallel words at 240 MHz. The receiver searches for the K28.5 comma character to establish word alignment and frame boundaries.
4. **Clock Cleanup and Downscaling** — The recovered 240 MHz clock is divided by six to give a 40 MHz bunch clock phase-locked to the TTC stream. It is then passed through a jitter-cleaning PLL to suppress high-frequency noise.
5. **Frame Validation and Decoding** — Each TTC frame is checked using its CRC[4] field. Valid frames are decoded into trigger commands and distributed internally on the 40 MHz BC domain.

This procedure is typically invisible to end-users but plays a critical role in ensuring timing reproducibility. In particular, the recovered clock phase after CDR lock is not guaranteed to match its previous value, and any shift may affect downstream timing alignment. These effects are at the core of the phase stability studies presented in Chapter 5.

[4]Frames carry an n-bit cyclic redundancy check computed over the header and payload with generator polynomial $G(x) = x^{16} + x^{15} + x^2 + 1$ (width $n = 16$ bits); the receiver accepts a frame only if the division remainder is zero [56].

Chapter 5

Clock phase stability studies

This chapter presents new experimental measurements and systematic investigation of clock phase stability for the LTI. These studies were conducted with two distinct hardware platforms. First with standalone LTI evaluation kits, then with full-format LTI ATCA prototypes, to establish a baseline and uncover any platform-specific behaviors. From there, we quantify the two main sources of instability: the indeterministic phase shifts that occur every time a transceiver link is reset, and the slow phase drift induced by temperature changes in both the fiber and FPGA transceiver circuitry. Alongside those measurements, we characterize the two primary timing tools-the DDMTD and the Tx Phase Interpolator (TxPI)—to understand their intrinsic resolution and noise floors. Finally, all of these strands feed into the development and validation of a software-based phase compensation algorithm using real-time temperature readings to keep the recovered clock phase locked within values compatible with the HGTD sub-30 ps budget, see Sections 1.1 and 3.1. In addition to these measurements, a dedicated framework for systematic testing and data analysis was developed by the author, building on initial exploratory hardware and firmware tests performed by other team members.

In the new ATLAS Phase-2 TTC distribution architecture, the LTI must recover a 40 MHz Bunch Clock from a 9.6 Gb/s serial stream, then distribute it with minimal added jitter.

However, two practical effects threaten this precision in a real-world experiment:

- **Phase indeterminism after resets** – Any re-initialization of the LTI-in transceiver PLL can yield a different recovered-clock phase, requiring a software re-alignment mitigation procedure.

- **Thermal drift** – Changes in ambient or die temperature shift both fiber delay and FPGA-internal delays, causing phase drifts that, if uncorrected, can accumulate to tens of picoseconds over a single LHC fill.

Left unchecked, these instabilities would degrade the timing resolution of sub-detectors

such as the HGTD, and reduce the effectiveness of pile-up suppression. A detailed understanding of both the magnitude and behavior of these effects, together with the development of in situ correction techniques, is therefore essential for reliable operation of the LTI within the ATLAS Phase-2 TTC network.

5.1 Experimental setup and instrumentation

This section describes the hardware platforms, measurement infrastructure, software architecture, and instrumentation used during the LTI clock phase stability studies. The test setup was designed to operate LTI boards in standalone mode under controlled environmental conditions, enabling precise and reproducible measurements of phase behavior. Both LTI kits and full LTI prototypes were tested using a shared measurement and control framework. The setup supported automated operation across temperature-controlled chambers, with phase monitoring performed via both DDMTD and oscilloscope instrumentation.

5.1.1 LTI hardware platforms

The experimental setup included two categories of hardware: evaluation kits and prototypes. The evaluation kits were selected specifically for their transceiver families of interest, although their FPGAs differ from the one used in the final LTI prototype. The AMD development kits employed were:

- AMD Zynq UltraScale+ MPSoC ZCU102 with GTH transceivers.
- AMD Virtex UltraScale+ VCU118 with GTY transceivers.

Each LTI kit was equipped with an FPGA Mezzanine Card (FMC) providing Samtec FireFly optical interfaces, jitter cleaners, clock generators, and SMA outputs. This configuration enabled TTC data transmission, recovery of the 40 MHz clock, and its extraction for external phase measurements. Unless otherwise specified, all kit-level measurements were performed using two ZCU102. The only exception is the GTY transceiver characterization, for which one VCU118 and one ZCU102 were used. Figure 5.1a and Figure 5.1b show the two different LTI kits used during the initial test campaigns.

The production of the first LTI prototype was completed in Q4 2024. By Q2 2025, the prototypes had reached a firmware and software state suitable for testing. As detailed in Chapter 4, the current LTI board differs from the hardware design foreseen for Phase-2 production. Insights gained during the initial testing campaign have motivated several hardware modifications, which will be implemented in a subsequent prototype iteration prior to finalizing the production design. Figure 5.2 displays one of the earliest prototypes.

(a) ZCU102 board with FMC.

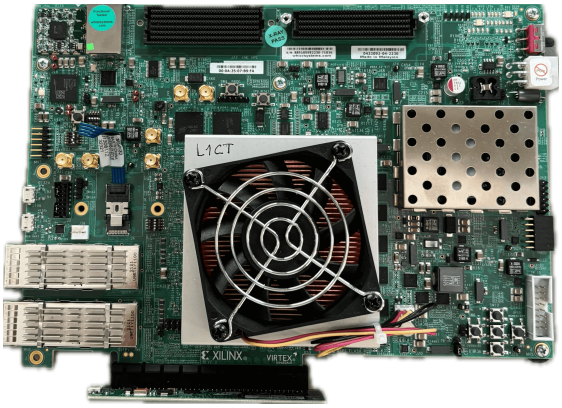

(b) VCU118 board.

Figure 5.1: LTI evaluation kits used for phase stability studies.

Figure 5.2: Fully populated LTI prototype #3.

The prototypes were installed in 3U, 2-slot ATCA shelves. Each shelf included a shelf manager board, which monitored hardware health and environmental status. The shelf manager was connected to the LTI's System-on-Module (SoM) via Ethernet and was accessible through ssh connection [66].

5.1.2 LTI boards configuration

Every test presented in this chapter was performed with two LTI boards, configured to emulate a single hop in the TTC distribution (either CTP↔LTI or LTI↔LTI).

- The **master** board generated the reference clock and transmitted data.
- The **slave** board received the serial stream and recovered the clock.

Throughout this chapter, we will refer to the transmitting board as LTI master and to the receiving board as LTI slave.

Figure 5.3 shows the block-level diagram of this setup. The diagram highlights the subsystems involved in clock generation, transmission, recovery, and measurement. The dashed box labeled "LTI FMC" marks the elements implemented on the mezzanine card in the case of the evaluation kits. The recovered clock from the slave is compared to the master's reference clock using two complementary instruments:

1. Oscilloscope — measures the downlink only.
2. DDMTD — measures the round-trip phase across both downlink and uplink.

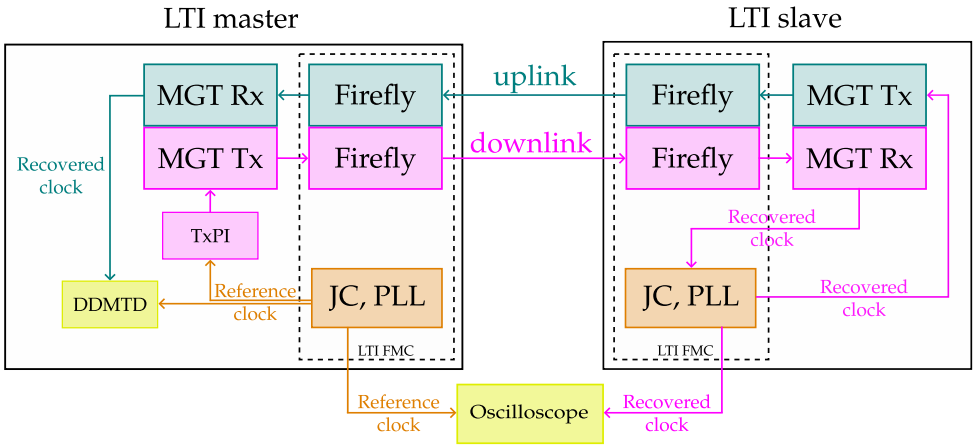

Figure 5.3: Block-level diagram of the setup for the recovered clock phase measurement.

5.1.3 Phase measurement instrumentation

A Keysight MXR254A digital oscilloscope was used to benchmark and cross-validate the phase measurements. This instrument offers a 2.5 GHz analog bandwidth on each channel—more than sufficient for our 40 MHz clock[1]—and a 16 GSa/s sampling rate (62.5 ps between raw samples). By employing its built-in interpolation and fitting algorithms, the MXR254A achieves effective timing resolution down to the picosecond level.

For each test scenario, the oscilloscope was configured with a display window of 10 µs, continuously acquiring and analyzing consecutive waveforms over an acquisition period of 3 s to 7 s. Within each acquisition period, the full waveforms of both the reference and recovered clocks were captured and processed using the scope's built-in interpolation and statistical measurement functions. At the end of each period, the measurement statistics were reset. The following quantity was computed for every acquisition period:

- **Inter-channel skew**: the time offset between the master reference clock and the slave recovered clock, measured on corresponding edges. This parameter quantifies the relative phase alignment between channels.

Figure 5.4 shows a representative oscilloscope screenshot, with the mean and standard deviation readouts for the measurement mentioned above. These were automatically averaged and reported by the oscilloscope, and yielded typical standard deviation spreads of 1.2 to 2 ps under thermally stable conditions.

Typical link jitter

A run of 1500 consecutive recovered-clock phase measurements was taken between two LTI kits under stable environmental conditions to probe the best-case jitter achievable with the oscilloscope-based method. The resulting distribution (Fig. 5.5) has a sample standard deviation $s = \sigma_{\text{obs}} = 0.52\,\text{ps}$.

This spread is not purely the intrinsic link-induced phase jitter but the convolution of

1. the true link-induced relative phase jitter, σ_{link}, and
2. the oscilloscope's internal timing resolution/jitter, σ_{scope}.

Assuming these two contributions are independent and Gaussian so that variances add,

$$\sigma_{\text{obs}}^2 \approx \sigma_{\text{link}}^2 + \sigma_{\text{scope}}^2. \tag{5.1}$$

According to the Keysight MXR254A datasheet [67], the intrinsic scope jitter can be as low as 118 fs under specified conditions. This value is used here only as a lower bound

[1]The rule of thumb says that the oscilloscope bandwidth shall be at least 1.5 times bigger of the signal to be measured.

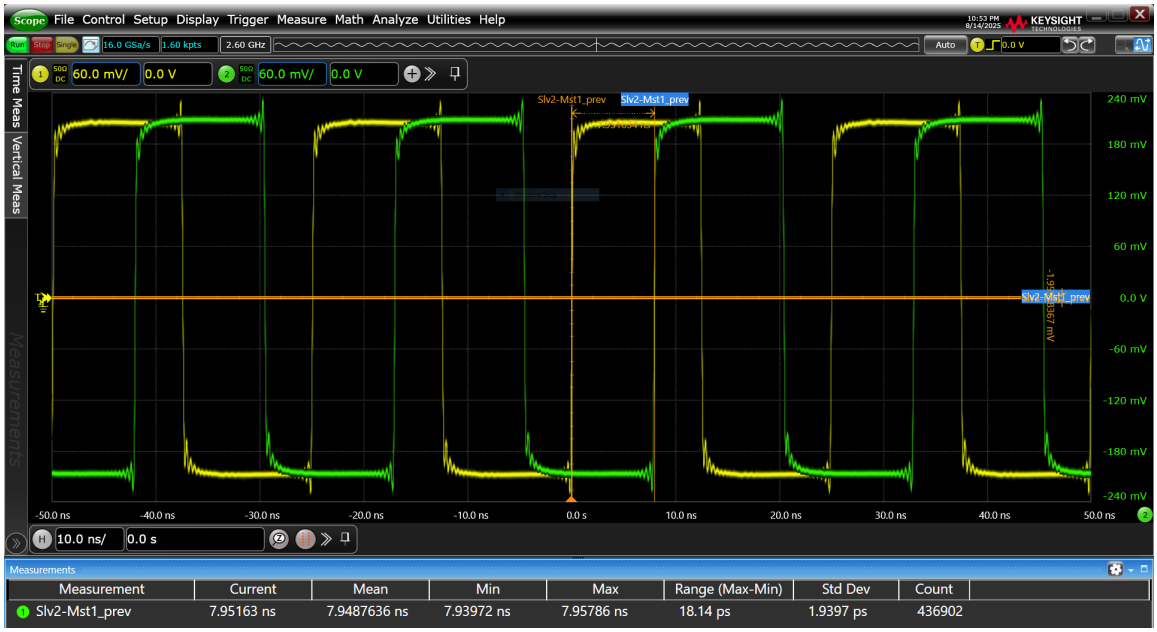

Figure 5.4: Example Keysight MXR254A measurement screen. The display window was adjusted for better readability.

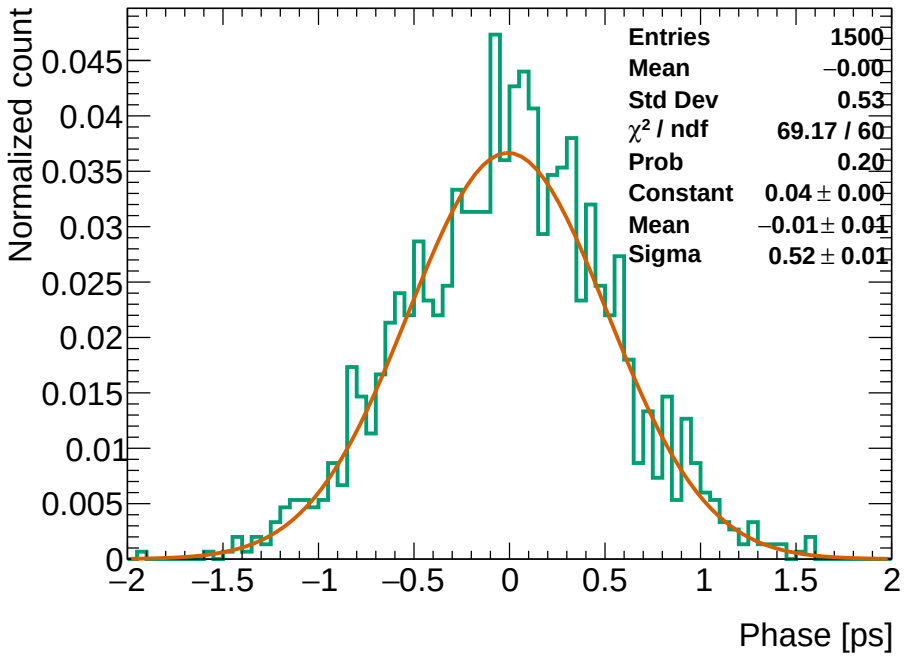

Figure 5.5: Distribution of averaged phase measurements over 1500 oscilloscope acquisitions, without transceiver resets.

for σ_{scope}, yielding an upper bound on the link jitter:

$$\sigma_{\text{link}} \lesssim \sqrt{0.52^2 - 0.118^2}\ \text{ps} \approx 0.51\ \text{ps}. \tag{5.2}$$

On this basis, the scope accounts for $\sigma^2_{\text{scope}}/\sigma^2_{\text{obs}} \approx 5\%$ of the variance; correspondingly, $\approx 95\%$ of the variance is attributable to the link. While small, the scope contribution is not strictly zero. Because the datasheet value is treated as a lower bound, the resulting link jitter estimates should be considered upper bounds.

A follow-up test with the LTI prototype (August 2025, see Section 5.6.2) yielded $\sigma_{\text{obs}} = 0.36\,\text{ps}$, implying

$$\sigma_{\text{link}} \lesssim \sqrt{0.36^2 - 0.118^2}\ \text{ps} \approx 0.34\ \text{ps}. \tag{5.3}$$

Here the scope would contribute about 11% of the variance, still secondary but less negligible than in the 0.52 ps case. This later result is consistent with the method resolving sub-0.5 ps rms jitter, with the measurement floor set largely by the link/clock rather than the instrument's intrinsic jitter.

Under the same stable conditions of Figure 5.5, a concurrent DDMTD acquisition produced the distribution in Figure 5.6 with a fitted width of $\sigma_{\text{obs,DDMTD}} = 0.45$ ps. Every entry in the histogram is a DDMTD measurement with an averaging factor $N_{\text{avg}} = 5000$. The fact that σ_{obs} and $\sigma_{\text{obs,DDMTD}}$ differ by only 14% indicates that both methods are resolving the same underlying link-induced jitter. This supports the conclusion that the oscilloscope's contribution is negligible and validates the measurement fidelity of the DDMTD setup under best-case conditions. The comparable precision also suggests that the DDMTD method is a suitable candidate for continuous phase monitoring during operation.

5.1.4 Temperature control

Every study presented in this chapter involves precise temperature control. For this reason, three climatic chambers were used[2]. The climatic chambers operated in a dedicated special room "13-S-011" or in group's electronics laboratory "14-2-024"; in both cases, the surrounding environment was maintained using air conditioning. Table 5.1 lists the climatic chambers used during the LTI tests together with their location, volume, and range of operating temperatures.

[2]Climatic chambers in general also offer relative humidity control, but this parameter was not considered in the studies presented here.

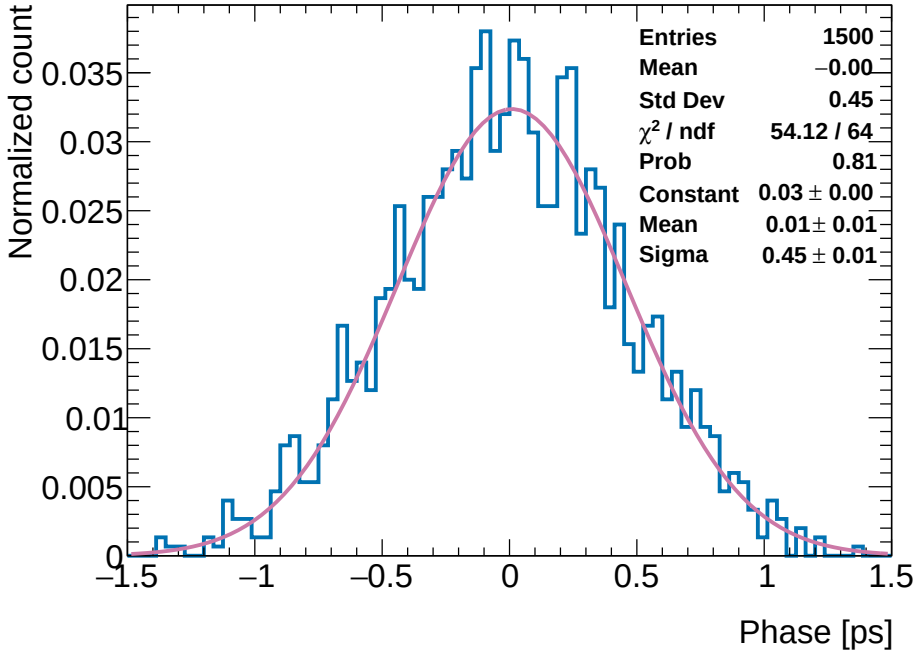

Figure 5.6: Distribution of averaged phase measurements over 1500 DDMTD acquisitions, without transceiver resets.

Table 5.1: Climatic chambers used for LTI tests.

Model	Location	Volume [dm^3]	Temperature (min, max) °C
Climats EXCAL 1423-HA [68]	13-S-011	138	(-70, 180)
CTS T-40/500 [69]	13-S-011	500	(-40, 180)
Binder KB260 [70]	14-2-024	262	(0, 70)

The EXCAL was controlled via the integrated PC, while the CTS and the Binder via custom Python APIs. During the tests with the LTI kits it was possible to drive two independent climatic chambers, one for the master and the other for the slave. With the LTI prototype this was not possible, and only one chamber was used. The reason is logistic: only the CTS and the Binder have enough volume to host the ATCA shelf, but these two chambers are installed in two different rooms and they cannot be moved. Figure 5.7 shows a configuration of LTI prototypes, where one is placed inside the KB260 climatic chamber.

The room and chamber temperatures were measured with a Pt100 Class-B probe (accuracy ±0.1 °C) connected to a Yoctopuce USB module [71, 72]. The chamber temperature variations matched the Pt100 readings within 0.2 °C. However, the absolute temperature values between the internal chamber sensor and the Pt100 often differed by about 1 °C, which was considered unimportant for the purposes of these tests.

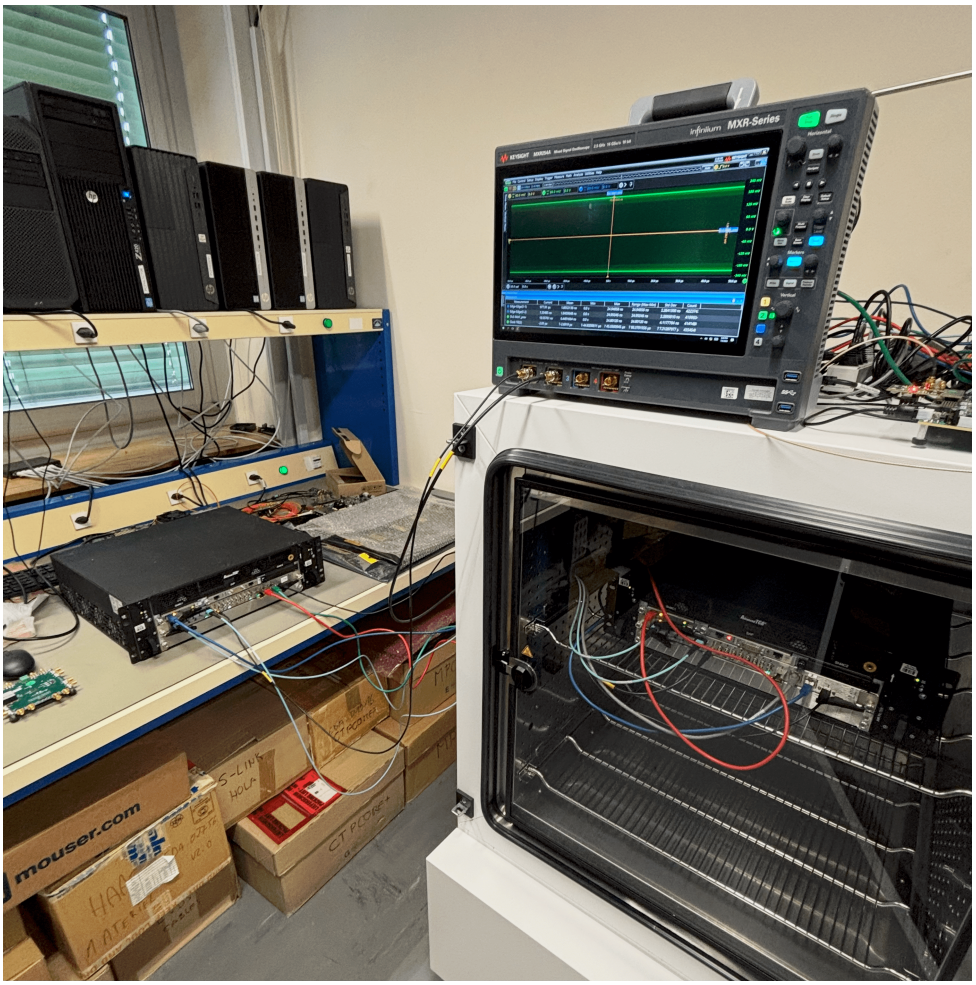

Figure 5.7: Two LTI prototypes configured for recovered clock phase stability tests.

On the LTIs the on-die FPGA junction temperature was measured via the I^2C interface and polled over Ethernet. A single raw temperature measurement is represented by a 10-bit ADC code M, which is converted into a temperature T using the formula [73]:

$$T = \frac{M \times 502.9098\,^{\circ}\mathrm{C}}{2^{10}} - 273.8195\,^{\circ}\mathrm{C}\,. \tag{5.4}$$

The least significant bit corresponds to a temperature increment

$$\Delta T = \frac{502.9098\,^{\circ}\mathrm{C}}{1024} \approx 0.491\,^{\circ}\mathrm{C}\,. \tag{5.5}$$

Thus, for uniformly distributed quantization error, the standard deviation of a single measurement is

$$\sigma_T = \frac{\Delta T}{\sqrt{12}} \approx 0.142\ ^{\circ}\mathrm{C}.$$

Assuming that successive N readings have uncorrelated quantization errors, the standard error on their average decreases as $1/\sqrt{N}$. In the tests presented in this chapter, the FPGA temperature measurements were averaged over the entire oscilloscope acquisition

time (3 s to 7 s) at a sampling rate of 4 Hz, giving N in the range 12–28. This corresponds to a statistical uncertainty from quantization alone in the range of 0.027 °C to 0.041 °C. This value is a lower bound on the total measurement uncertainty, as it does not account for systematic contributions such as sensor calibration accuracy or thermal lag.

The assumption of uncorrelated quantization error was verified on a representative dataset acquired with internal averaging disabled (averaging = 1). For $N = 300$ consecutive FPGA die-temperature measurements, the lag-1 autocorrelation[3] was about 0.032, indicating that successive readings were nearly independent and that the $1/\sqrt{N}$ scaling is applicable.

5.1.5 Software and control architecture

The software and control architecture of the test setup was required to be modular, with centralized data collection and full automated test runs. The infrastructure adopted a hierarchical software architecture, with a central host PC coordinating multiple LTI boards. Communication between the host PC and each board's System-on-Module (SoM) was established over Ethernet using Pyro5 remote procedure calls (RPC). The RPC is a paradigm that permits to call a function or method in another address space (process, machine), as if it were local [75]. In addition to configuring and coordinating the LTI boards, the host PC was also responsible for controlling the climatic chambers via Ethernet or RS232, and acquiring oscilloscope data through PyVISA queries.

Each SoM ran a Pyro5 server that exposed the underlying C++ control framework[4] through Python bindings. This interface provided access to configuration and monitoring of the FPGA via memory-mapped I/O, including control over clocks, transceivers, TxPI, and DDMTD modules. The choice of using Python bindings for the tests was motivated by their flexibility and rapid prototyping capabilities. In particular, they allowed on-the-fly modifications to test scripts and control logic without requiring the time-consuming recompilation of the underlying C++ framework. This significantly streamlined debugging and iterative development during the test campaign.

Figure 5.8 illustrates how the host PC coordinates the LTI boards, climatic chamber, and oscilloscope through the adopted control architecture.

5.1.6 Test scenarios

The test campaign included the following scenarios:

- Characterization of timing tools (DDMTD and TxPI) – Section 5.3

[3] The lag-1 autocorrelation is the Pearson correlation coefficient between successive samples x_t and x_{t+1} [74].

[4] The LTI Module.

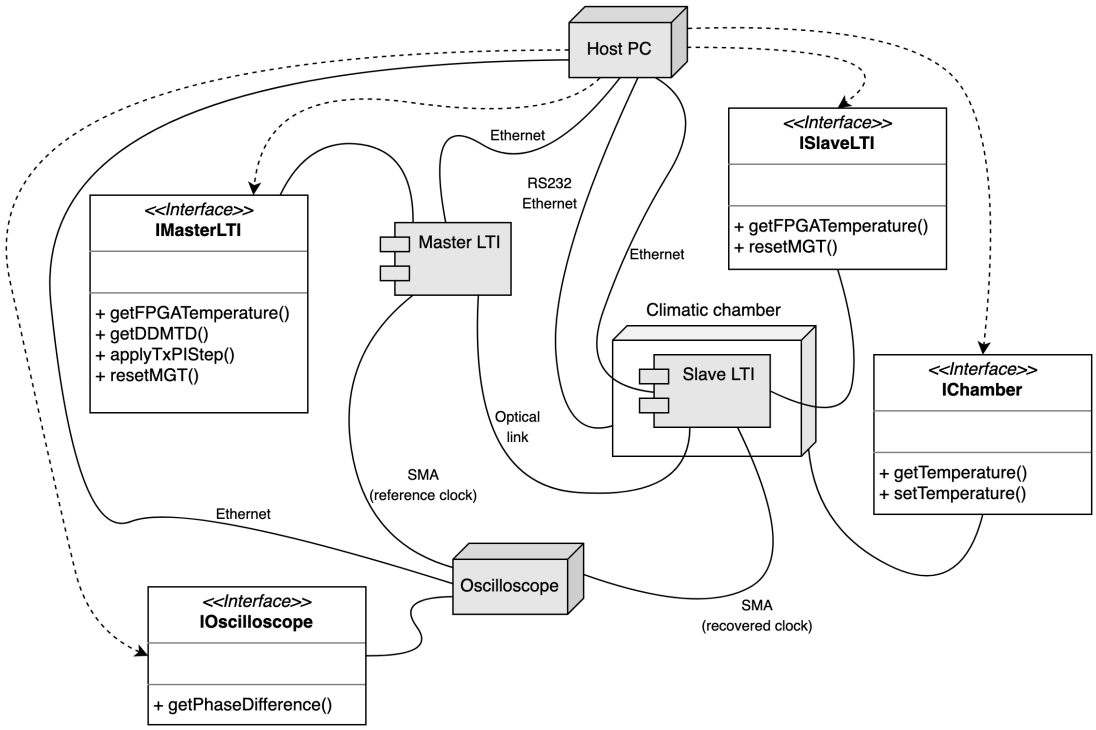

Figure 5.8: Software and control architecture of the test setup, showing device interconnections.

- Phase determinism after transceiver resets – Section 5.4
- Temperature-induced phase drifts with controlled thermal sweeps – Section 5.5
- Mitigation strategies for phase indeterminism and phase drifts – Section 5.6

The tests were repeated across both GTH and GTY transceivers, as well as LTI kits and prototypes, comparing behavior across platforms and environmental conditions.

5.2 Operational temperature analysis from ATLAS 2024 data

The data used in this analysis consist of MUCTPI and ATLAS counting cavern monitoring datasets collected during the 2024 ATLAS operations, covering the period between March and November included, and archived within the CERN Detector Control System[5]

[5]As a vital component in the operation of LHC experiments, the DCS is responsible for monitoring and controlling the detectors' environmental and hardware conditions. It continuously supervises parameters such as temperature, voltage, current, humidity, and gas flow to ensure safe and stable operation. The DCS also provides automated protection mechanisms and issues alarms when predefined thresholds are exceeded. In addition to real-time control, the system archives monitoring data in dedicated databases, allowing long-term storage.

(DCS). The reasoning behind this selection is that the MUCTPI, like the LTI, is a single ATCA blade equipped with three processing FPGAs of comparable size to the LTI FPGA. During the ATLAS 2024 operations the MUCTPI crate in the counting cavern housed two MUCTPI boards: the MUCTPI v3 used for data-taking and the MUCTPI v4 used as a spare. Two ATCA boards running in the same crate recreate a condition similar to what it is expected for the LTIs, where multiple boards will be housed in the same crate. The ATLAS 2023 operations were excluded from this analysis because the MUCTPI crate was housing only the MUCTPI v3. Additionally, we assume that the temperature variations in the counting cavern are sufficiently slow to consider the optical fibers in thermal equilibrium with their environment.

The counting cavern temperature data archived in the DCS consists of measurements from side A and side C on the first floor (L1), as well as from the second floor (L2). Temperature values are averaged over 15-minute intervals, based on the assumption that environmental conditions change slowly and that raw measurements may contain noise, fluctuations, or outliers. No details about the specific sensors used were available. Figure 5.9 shows the distribution of maximum temperature variations per run across all the counting cavern zones.

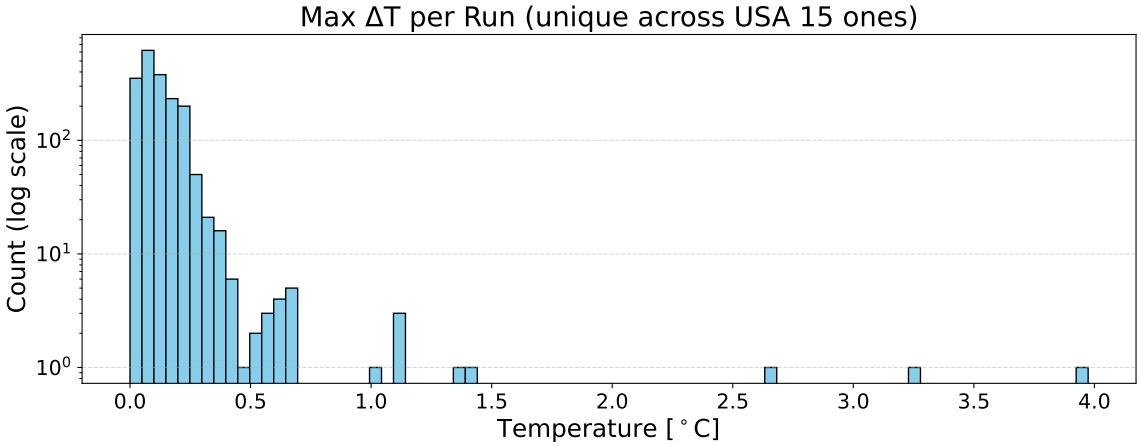

Figure 5.9: Histogram of maximum temperature excursions during 2024 ATLAS operations. For each run only the counting cavern region with the largest value is selected.

The largest recorded temperature variation was 3.97 °C, observed on the C-side of the first floor. This event, however, was exceptional as it was caused by a major incident involving a loss of chilled water in the counting cavern, resulting in a complete power shutdown of all racks. Considering both typical and extreme variations, a conservative temperature variation of 0.5 °C per run is considered[6]. In tests involving optical fibers the temperature range explored with the climatic chamber 18 °C to 22 °C, which covers every condition.

[6]Corresponding to the 98.8$^{\text{th}}$ percentile.

All three MUCTPI FPGA silicon die temperatures have been analyzed similarly, run by run and over the same period of time. Figure 5.10 shows the distribution of average MUCTPI FPGA silicon die temperatures per run. In light of the LTI conditions—specifically, having only one FPGA per board, but multiple boards per ATCA crate—and based on the MUCTPI data, this document operates under the assumption that values between 40 °C and 48 °C are likely.

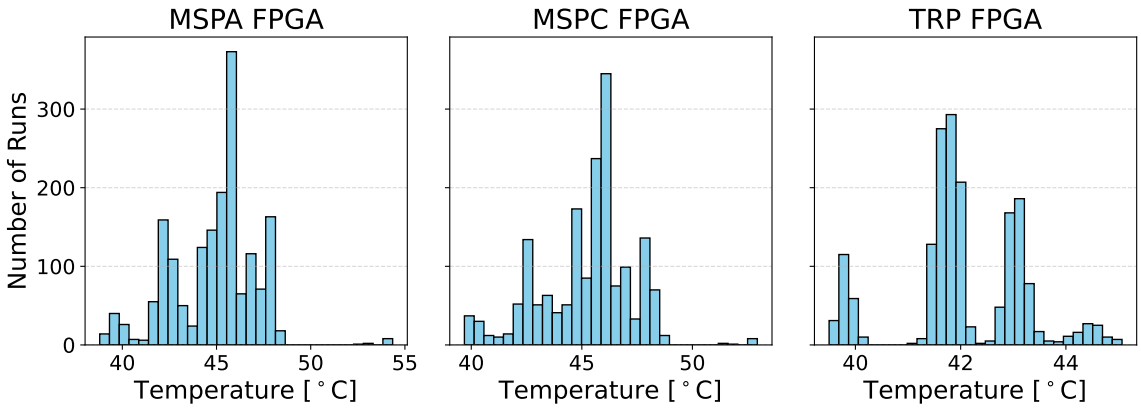

Figure 5.10: Histogram of average MUCTPI FPGAs silicon die temperatures across all the ATLAS 2024 run.

The largest variation of the MUCTPI FPGA silicon die operating temperatures within the ATLAS 2024 runs is 4.9 °C, with mean values across FPGAs and runs ranging from 2 °C to 4 °C. The timeline of this run is shown in Figure 5.11, which exhibits narrow, quasi-periodic excursions superimposed on a slowly varying baseline. Their recurrence at 15-minute intervals and small amplitude are consistent with step responses of the ATCA crate thermal control: when a board-level temperature reaches a pre-defined threshold, the shelf manager increases fan speed, producing a short-lived die-temperature adjustment before a new equilibrium is reached.

The operational scenario of the MUCTPI during the 2024 run shares several features with the anticipated LTI deployment: both involve ATCA crates housing multiple high-power FPGAs in a controlled yet slowly varying thermal environment. Each MUCTPI board contains three FPGAs, whereas each LTI board contains a single FPGA. In particular, two MUCTPI boards running simultaneously in the same crate (six FPGAs in total) approximate the overall crate-level heat load and airflow patterns expected when more than one LTI board operates in the same crate. However, differences remain — the LTI FPGAs may have distinct power profiles, heat dissipation characteristics, and cooling configurations compared to the MUCTPI, and future ATLAS running conditions could introduce different cavern temperature fluctuations or crate-level ventilation efficiency. Furthermore, the thermal coupling between optical fibers and their surroundings, while assumed to be near equilibrium for the present study, may differ for the LTI installation

depending on routing and fiber packaging.

Taking these similarities and possible differences into account, we adopt a conservative approach for test conditions. For the counting cavern environment, a representative operational variation of 4 °C is chosen, while for the FPGA silicon die temperatures a variation of 8 °C is selected — exceeding the largest variation observed in MUCTPI operation (4.9 °C) to ensure margin. This ensures that the thermal excursions tested in the laboratory include both typical and exceptional operational scenarios, providing headroom against unforeseen changes in the final system deployment.

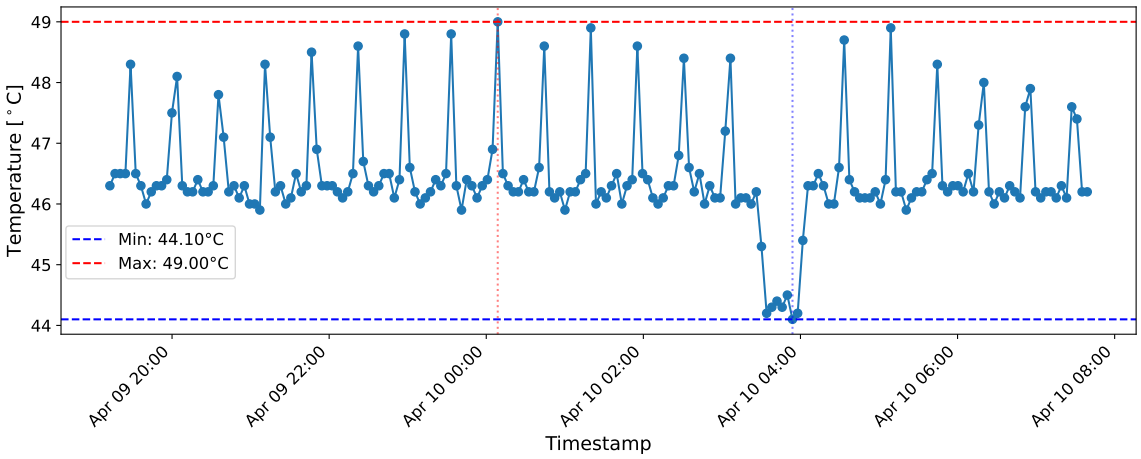

Figure 5.11: MUCTPI FPGA temperature measurements timeline for the ATLAS run 472879.

5.3 Characterization of timing tools

This section presents the characterization results for the DDMTD and TxPI blocks, obtained in a thermally controlled environment using the test configuration described in Section 5.1. The primary objectives were (i) to determine the best trade-off between achieving sub-picosecond precision in DDMTD phase measurements and minimizing acquisition time, so as to limit the monitoring overhead when scaling to up to the 60 LTI-out, and (ii) to validate the deterministic behavior of the TxPI for picosecond-level phase shifts.

To characterize the DDMTD phase measurement system, three parameters were available: the averaging factor in the software (N_{sw}), the averaging factor in the firmware (N_{fw}), and the deglitch threshold. The total effective averaging depth is defined as $N_{\text{avg}} = N_{\text{sw}} \times N_{\text{fw}}$. In this setup, each DDMTD result reported by the test program corresponds to the arithmetic average of N_{sw} phase readings, where each individual phase reading is already internally averaged in the FPGA over N_{fw} samples. This two-level averaging scheme allows flexible tuning of acquisition time and resolution.

5.3.1 Measurement duration and dependence on averaging depth

The measurement duration of the DDMTD system depended on the averaging configuration. Tables 5.2 and 5.3 report the average duration per DDMTD measurement and the corresponding effective time per internal sample for firmware-only and software-only averaging, respectively, as measured on the LTI prototype. The maximum firmware averaging depth was limited by the 12-bit width of the configuration register.

For small N_{fw} the fixed call overhead dominated, so the time per effective sample became large. As N_{fw} grew, the per-sample cost approached a constant determined by the DDMTD beat period $T_{\text{beat}} \approx 68.1\,\mu\text{s}$. In practice, the **software-only** path approached an asymptote of $\sim 136\,\mu\text{s/sample}$ ($\approx 2\,T_{\text{beat}}$), whereas the **firmware-averaging** path approached an asymptote of $\sim 272\,\mu\text{s/sample}$ ($\approx 4\,T_{\text{beat}}$). This factor-of-two difference between the two asymptotes could not be explained by call overhead at large N, which was negligible, but instead reflected implementation details of the FPGA DDMTD averaging logic. Firmware averaging remained advantageous at small to moderate averaging depths, because it reduced host I/O overhead, while software averaging achieved the lower per-sample floor when very large N values were used.

Table 5.2: Measurement duration scaling with firmware averaging ($N_{\text{sw}} = 1$).

N_{sw}	N_{fw}	N_{avg}	Avg. Duration [µs]	Duration/sample [µs]
1	1	1	1222	1222.0
1	10	10	2440	244.0
1	100	100	27240	272.4
1	1000	1000	272300	272.3
1	4095	4095	1115659	272.4

Table 5.3: Measurement duration scaling with software averaging ($N_{\text{fw}} = 1$).

N_{sw}	N_{fw}	N_{avg}	Avg. Duration [µs]	Duration/sample [µs]
10	1	10	2448	244.8
100	1	100	14700	147.0
1000	1	1000	137450	137.5
4095	1	4095	559200	136.6
10000	1	10000	1366686	136.7

Figure 5.12 visualizes these scaling trends. It shows the software-only path converging to ~ 136 µs/sample, while the firmware-only path converged to ~ 272 µs/sample.

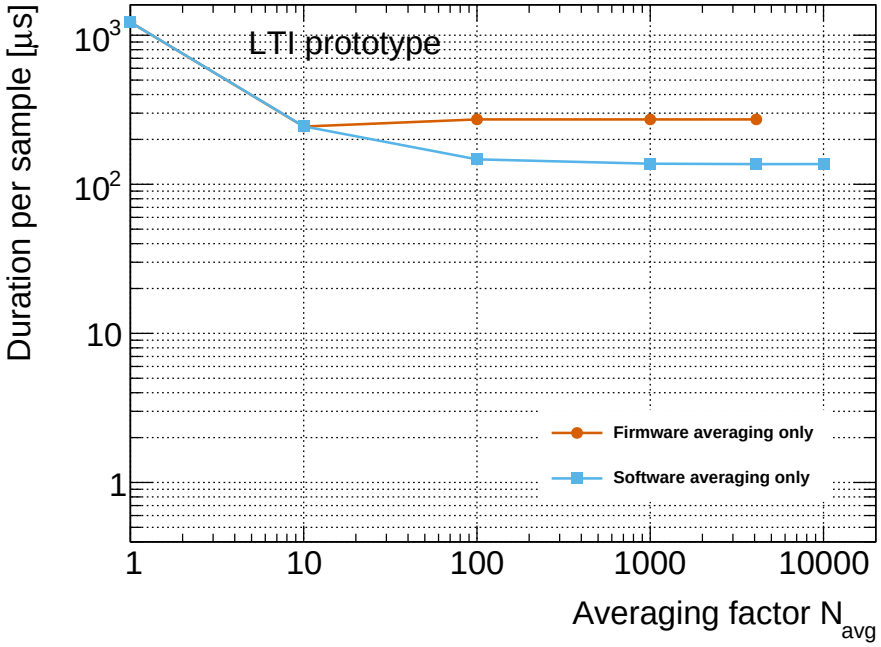

Figure 5.12: Comparison of duration per sample vs. total averaging depth N_{avg} for firmware-only and software-only averaging configurations.

5.3.2 Deglitch threshold selection

In addition to firmware and software averaging, the DDMTD block includes a configurable deglitch threshold parameter. This filter suppresses short-lived transitions in the sampled beat signal, rejecting events likely caused by metastability, or other noise sources. The parameter is expressed in units of the DDMTD clock period $T_{\text{ddmtd}} \approx 4.158\,\text{ns}$ and thus corresponds to a minimum time interval between two accepted beat edges:

$$t_{\text{dg}} = N_{\text{dg}} \times T_{\text{ddmtd}}. \tag{5.6}$$

For the White Rabbit default $N_{\text{dg}} = 1966$, this yields $t_{\text{dg}} \approx 8.18\,\mu\text{s}$. Any transition occurring within this time after the last accepted edge is discarded as a glitch. Since the beat period for the LTI settings is $T_{\text{beat}} \approx 68.1\,\mu\text{s}$, excessively large thresholds approach this value and begin to reject legitimate beat edges.

The impact of the deglitch threshold was characterized on the LTI prototype by computing the average standard deviation from 100 repeated measurements for each N_{dg} value. Error bars in Fig. 5.13 represent the standard error of the mean. The results show only modest fluctuations in resolution without a clear monotonic trend. The lowest average standard deviation was observed near $N_{\text{dg}} = 2000$, while both small thresholds (below 1000) and large thresholds (above 5000) yielded higher variability.

The scan was limited to $N_{\text{dg}} = 8000$ (33.3 µs), roughly half the beat period, as higher

thresholds would reject nearly all valid beat edges and yield no meaningful measurement.

This behavior can be explained by the balance between noise suppression and valid-edge retention. At low N_{dg}, spurious edges are not effectively rejected, introducing additional jitter into the phase accumulator and degrading resolution. At high N_{dg}, the filter begins to discard valid beat edges, reducing the number of samples available for averaging and thus increasing statistical uncertainty. The deglitch threshold does not affect the measurement duration, as the acquisition time is fixed by the beat frequency. Based on this analysis, the White Rabbit default of $N_{\mathrm{dg}} = 1966$ was retained for all subsequent measurements.

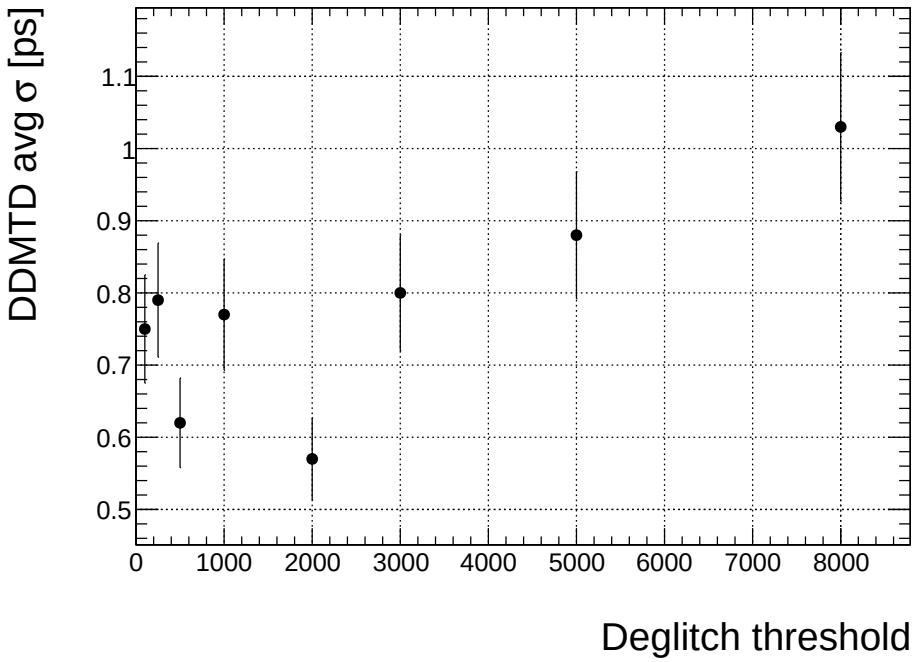

Figure 5.13: Average DDMTD standard deviation σ as a function of the deglitch threshold, measured on the LTI prototype.

5.3.3 DDMTD resolution

The configuration providing the best DDMTD resolution was identified by measuring the average standard deviation from 100 repeated measurements across different parameter combinations. During these tests, the explored parameters were N_{sw} and N_{fw}, while the deglitch setting remained constant. To minimize the impact of temperature drifts, firmware averaging was preferred whenever feasible, reducing I/O overhead.

Figure 5.14 shows the resolution scaling behavior of the DDMTD for both the LTI kit and the LTI prototype. The $\sigma \approx 0.50$ ps point for the kit corresponds to the configuration used for the large acquisition shown in Fig. 5.6.

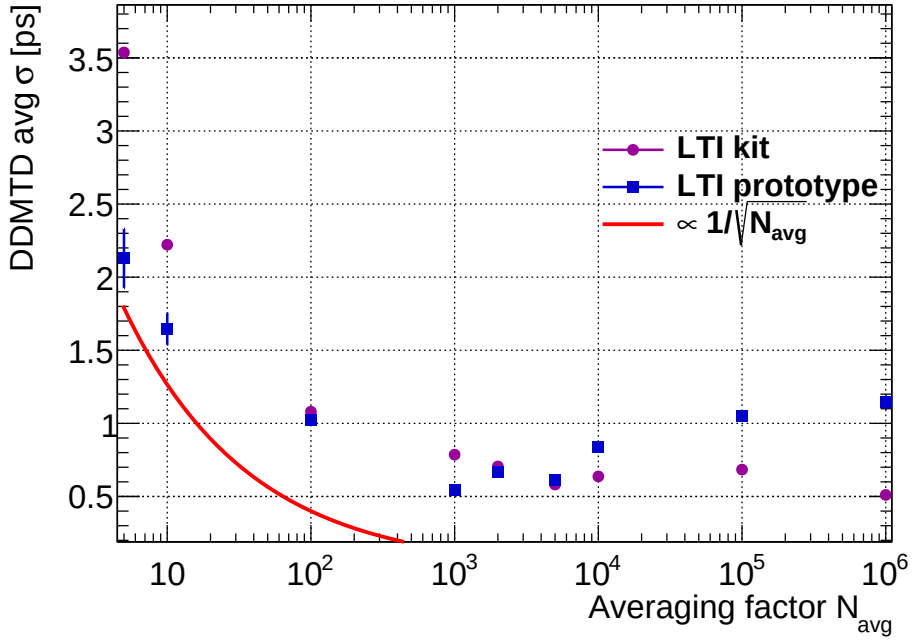

Figure 5.14: Standard deviation of 100 DDMTD phase measurements versus total averaging depth $N_{\text{avg}} = N_{\text{sw}} \times N_{\text{fw}}$, comparing results from the LTI kit and the LTI prototype.

The averaging-depth scan highlights the different resolution behaviors of the LTI kit and the LTI prototype. For the kit, the DDMTD resolution remains close to a floor of $\sim$ 0.45–0.50 ps over the explored range of N_{avg}, confirming that the measurement is instrument-limited under stable conditions. In this regime, further averaging provides negligible resolution improvement and may instead expose the measurement to thermal fluctuations.

The LTI prototype exhibits a different trend, with an earlier plateau or slight degradation in resolution beyond $N_{\text{avg}} \gtrsim 10^3$, indicating a stronger influence of residual correlated noise sources such as temperature drift. This limitation is specific to the current firmware implementation and hardware integration used in the prototype, and is not intrinsic to the DDMTD method itself. Firmware and system-level optimizations, including improved isolation and layout to mitigate crosstalk, are foreseen in future iterations to reduce susceptibility to correlated noise and to bring the prototype performance closer to the instrument-limited floor observed on the kit.

5.3.4 TxPI resolution characterization

The TxPI step size was measured using both the oscilloscope and the DDMTD system on the LTI prototype. The measurement procedure consisted of:

1. Measuring the DDMTD and oscilloscope phase.

2. Incrementing the phase by one TxPI step.
3. Repeat (1) and (2) for 1000 times

The phase data was then fitted as a function of the TxPI step index. Figures 5.15a and 5.15b show the results obtained with the DDMTD and with the oscilloscope, respectively. In both cases the fitted slope agrees with the expected TxPI step size of Equation (4.5).

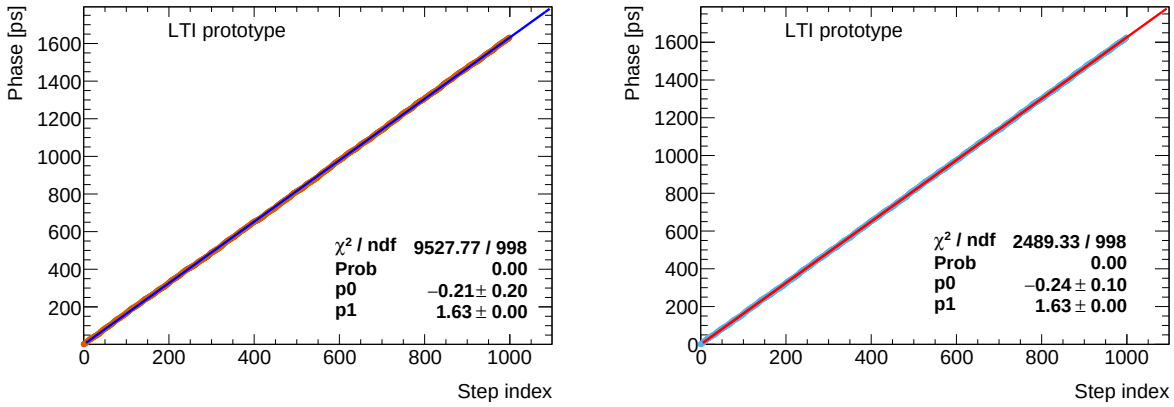

(a) TxPI step size measured on the LTI prototype using the DDMTD.

(b) TxPI step size measured on the LTI prototype using the oscilloscope.

Figure 5.15: Oscilloscope recovered clock phase distributions after resetting the GTH transceivers at constant temperature for the LTI evaluation kit.

The TxPI tests were repeated at 25 °C, 30 °C, and 35 °C. Across all temperatures, the per–step shift is consistent with Equation (4.5), confirming stable TxPI performance.

5.4 Phase determinism after transceiver reset

During ATLAS physics data-taking, or any test run, it may happen that an optical link is lost. Ideally, re-running the link bring-up procedure should result in a recovered clock phase similar to the previous one. In a perfect system, the phase of the recovered clock should remain unchanged after a reset. In practical applications this does not happen due to indeterminism in the CDR procedure, and the phase uncertainty shall be as small as possible. If the distribution of the recovered phase after resets becomes too broad, a mitigation strategy must be implemented.

This section presents a characterization of the recovered clock phase after transceiver resets, for both GTY and GTH transceivers, tested on both the LTI evaluation kit and the LTI prototype. The early results obtained using the evaluation kit were essential to understanding the behavior of the AMD transceivers “e4” generation.

Resetting transceivers in a link involves executing the same sequence used during initial link bring-up, as presented in Chapter 4. Without loss of generality, we assume the link is established between two LTI boards. The master and slave transceivers can be reset independently, nevertheless when resetting a TX the corresponding RX loses the lock and would need to reset as well to re-establish the link.

A full transceiver reset affects both downlink and uplink paths. The correct sequence of actions for a complete reset is as follows:

1. Reset the LTI-out master TX
2. Reset the LTI-in slave RX
3. Reset the LTI-in slave TX
4. Reset the LTI-out master RX

The procedure used for the data-taking sessions illustrated in this section include the following actions:

1. Set the climatic chamber to the temperature set-point.
2. Poll the climatic chamber current temperature every 1 s, and wait until the difference of current temperature and the set-point is $\leq 0.1\,^{\circ}\mathrm{C}$ for at least five minutes. At the end of the five minutes, the temperature was considered stable.
3. Reset the transceivers of interest.
4. Wait 1 s to let the system stabilize.
5. Measure the Oscilloscope for an acquisition window of 7 s. At the same time measure the DDMTD phase.
6. Repeat (3) to (5) 1000 times.

The setup is based on the configuration described in Section 5.1.2. In this arrangement, the oscilloscope measures the downlink only, while the DDMTD monitors both downlink and uplink. High-precision phase stability is only required for the clock-distribution path (the downlink); the uplink does not have critical timing requirements.

5.4.1 GTH Transceivers

The behavior of GTH transceivers after reset using the LTI kit is shown in Figure 5.16. The measurements comprised 1000 resets in two configurations:

- Master LTI-out TX + slave LTI-in RX.
- Slave LTI-in RX.

In both configurations, the distributions are consistent with a Gaussian fit. A comparison of the fitted standard deviations shows that the RX contribution is about 74% of the combined RX+TX reset indeterminism, in agreement with earlier studies [61]. The same tests repeated on the LTI prototype, shown in Figure 5.17, delivered consistent results, confirming that the GTH transceivers meet picosecond-level phase stability requirements.

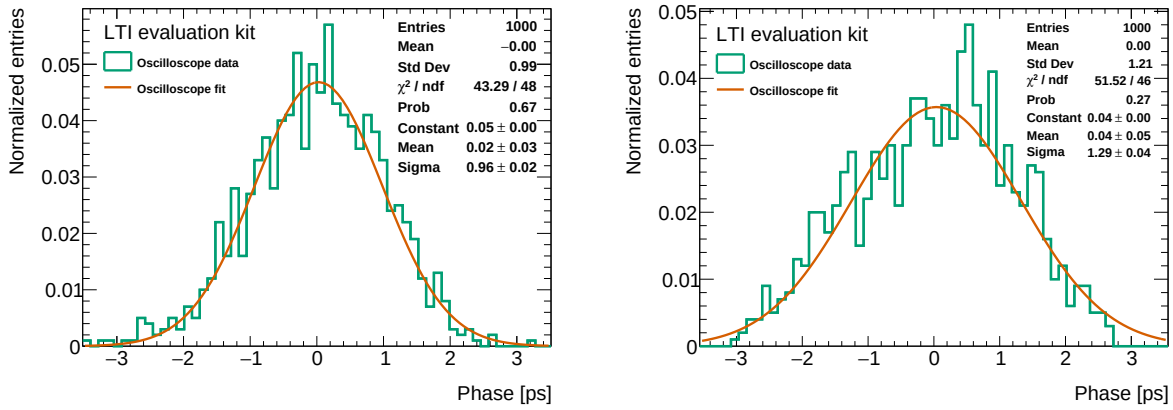

(a) Reset of LTI-in RX.

(b) Reset of LTI-out TX and LTI-in RX.

Figure 5.16: Oscilloscope recovered clock phase distributions after resetting the GTH transceivers at constant temperature for the LTI evaluation kit.

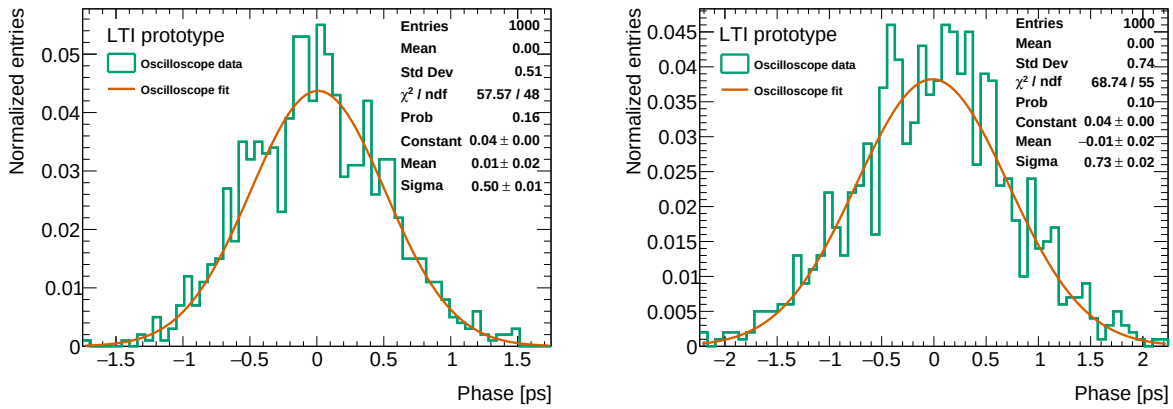

(a) Reset of LTI-in RX.

(b) Reset of LTI-out TX and LTI-in RX.

Figure 5.17: Oscilloscope recovered clock phase distributions after resetting the GTH transceivers at constant temperature for the LTI prototype.

5.4.2 GTY transceivers

The behavior of the GTY transceivers has been studied in a similar way. During the early stages of the LTI development the assignment of LTI-in and LTI-out links to transceiver families was not defined, therefore the need of studying the distribution of the recovered

clock also on GTY transceivers. Figure 5.18 shows the results obtained with the LTI evaluation kit. The Subfigure 5.18a illustrates that the recovered clock follows a distribution with two well-distinct Gaussian profiles, with a range over 10 ps.

In the scenario of GTY transceivers used as LTI-in, this bimodal behavior would compromise timing. After the initialization (or reset) of the link, the recovered clock phase lands displaced in one of the two peaks, eroding on average 4 ps from the timing uncertainty budget. For this reason, the LTI firmware was revised to include a dedicated mitigating procedure, explained in Section 5.6.1.

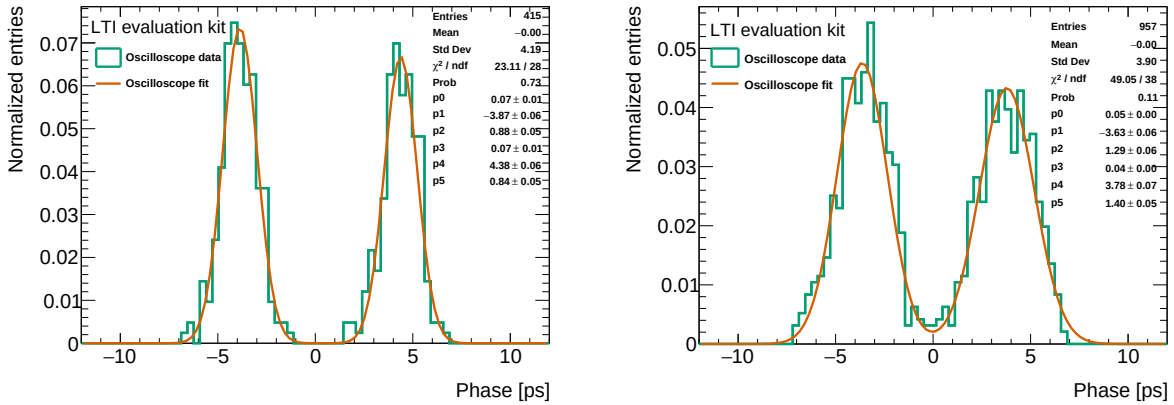

(a) Reset of LTI-in RX.

(b) Reset of LTI-out TX and LTI-in RX.

Figure 5.18: Oscilloscope recovered clock phase distributions after resetting the GTY transceivers at constant temperature for the LTI evaluation kit.

During the production of the first LTI prototype the firmware engineers revised the transceiver family assigned to LTI-in and LTI-out, converging on the use of only GTH for LTI-in.

5.5 Temperature-induced phase drifts with controlled thermal sweeps

The drift coefficients presented in this section measure the sensitivity to phase shifts per °C in the FPGA silicon die temperature. They are a good figure of merit for the estimation of timing budget, and also fundamental for the implementation of a phase correction algorithm. The initial tests with the LTI evaluation kits covered a 20 °C sweep, whereas for the LTI prototype the range was reduced to 8 °C, following the considerations in Section 5.2. The measurement procedure was standardized across all runs and consisted of the following steps:

1. Set the climatic chamber to the temperature set-point, which is the initial value of the temperature ramp.

2. Poll the climatic chamber current temperature every 1 s, and wait until the difference of current temperature and the set-point is $\leq 0.1\,^{\circ}\mathrm{C}$ for at least five minutes. At the end of the five minutes, the temperature was considered stable.

3. Start phase monitoring on the oscilloscope and the DDMTD. As a reminder, the DDMTD data shows the complete round-trip contribution to the phase drift, while the oscilloscope on the downlink only.

4. Increase the climatic chamber set-point by $0.2\,^{\circ}\mathrm{C}$, and wait 180 s before the next set-point increase. At maximum temperature the ramp changes sign and returns to the initial value at the same rate.

Figure 5.19b shows a linear response of both the DDMTD and the oscilloscope. In contrast, Figure 5.19a reveals that the DDMTD exhibits a non-linear response when the FPGA die temperature is varied. o the best of the author's knowledge, this represents the first documented observation of non-linear DDMTD behavior under FPGA thermal sweeps.

For the LTI prototype, the results are shown in Figures 5.20a–5.20b. These measurements confirm the key outcome already established with the evaluation kits: the DDMTD does not exhibit a linear response when the temperature of the hosting FPGA is varied. Figure 5.20c also shows the thermal response of a 40 m optical fiber. In this case, the factor of two observed between the DDMTD and oscilloscope slopes arises because the DDMTD measures phase shifts in both the downlink and uplink paths, effectively doubling the apparent thermal coefficient. This property, and in general the different slopes between oscilloscope and DDMTD characteristics, were later taken into account during the phase-compensation tests described in Section 5.6.3.

While confirming this outcome, the prototype measurements further reveal structured, non-linear oscillations superimposed on the underlying thermal drift. These additional features may be attributable to differences in clock routing in the prototype FPGA, which belongs to a different family and has a larger device size than the evaluation-kit FPGA. In the prototype, certain routing paths of clock signals may be more sensitive to thermal fluctuations, potentially distorting an otherwise linear response.

5.6 Mitigation strategies for phase indeterminism and phase drifts

This section explains the procedures put in place to mitigate GTY phase indeterminism after reset and phase instabilities due to FPGA temperature variations.

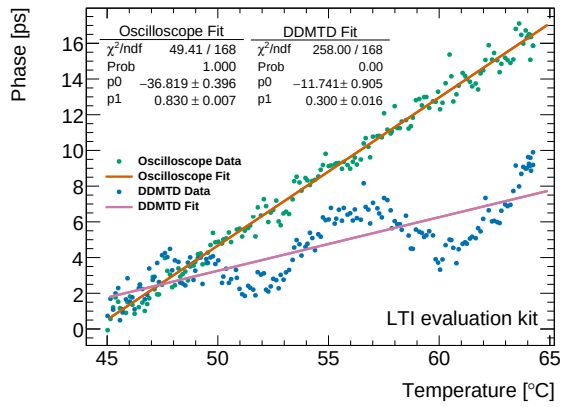

(a) LTI master in the climatic chamber.

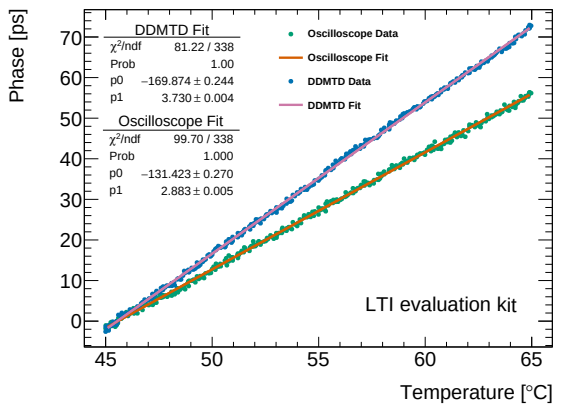

(b) LTI slave in the climatic chamber.

Figure 5.19: LTI evaluation kit phase monitoring data with either LTI master or slave are in the climatic chamber

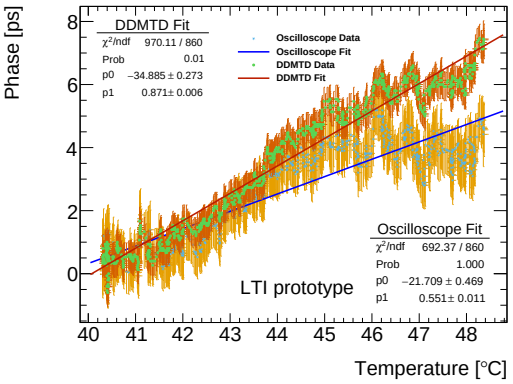

(a) LTI slave in the climatic chamber.

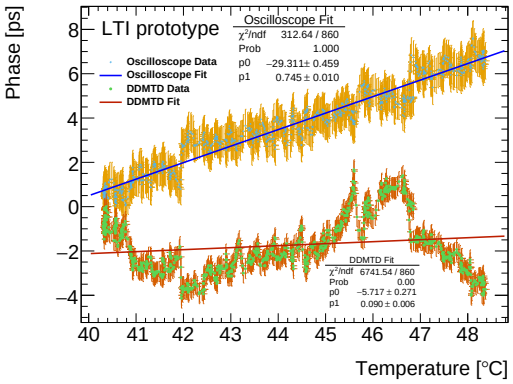

(b) LTI master in the climatic chamber.

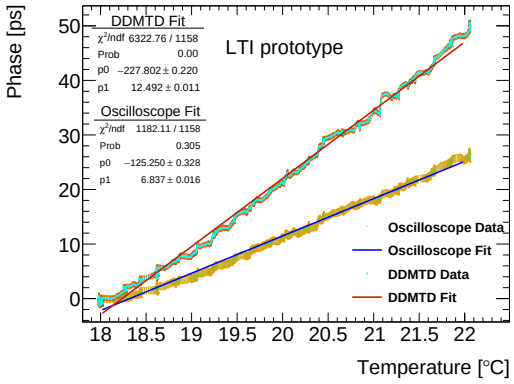

(c) Optical fibers in the climatic chamber.

Figure 5.20: LTI prototype phase monitoring data with either LTI slave (a), LTI master (b), or 40 m of optical fibers (c) are in the climatic chamber.

5.6.1 Rx-fanout reset procedure

The RX-fanout reset procedure (see Section 4.5.3) is a workaround implemented by the LTI firmware developers to address the bimodal distribution of the recovered clock phase after a reset of GTY transceiver. The implementation of this procedure required a hardware modification, with the addition of a high-speed fanout chip at the LTI-in RX.

The method described here successfully rectified a two peaks distribution into a single peak distribution centered in zero, similar to GTH behavior.

1. On the LTI slave the 8b/10b transmission is bypassed, switching to a 240 MHz clock-like pattern.
2. The fanout chip routes the 240 MHz clock (20 UI symbols[7] high, followed by 20 UI symbols low) to both the GTY transceiver, and the FPGA.
3. A clock reference is now available to the DDMTD. At the same time the GTY transceiver remains locked on data, thanks to the "run length" firmware parameter, $\in [0, 256]$, which defines the maximum numbers of consecutive UI symbols before the receiver loses lock.
4. The DDMTD measures in which peak the phase landed.
5. A new reset of the receiver can be issued. This consistently maps the double-peak distribution into a single-peak distribution GTH-like.
6. Switch back to 8b/10b TTC stream on the Tx side.

Figure 5.21 shows the effectiveness of the procedure, with a two-peaks distribution rectified into a single-peak distribution.

5.6.2 Clock fiber

The clock fiber mode (see Section 4.5.3) concerns a clock distribution scheme where every FELIX receives the clock from the LTI over a separate fiber. This mode was tested using two LTI prototypes, with the LTI slave placed in the climatic chamber and the LTI master in equilibrium with electronics lab's air conditioning. The result of a clock stability assessment with this distribution scheme is shown in Figure 5.22. The fitted Gaussian distribution shows that the oscilloscope resolution is at least as good as 0.34 ps, improving the figure of merit shown in Section 5.1.3 by 33%. The dataset used in this experiment includes less than 1000 entries because few of them were affected by large temperature variations in the electronics lab, that in turn induced variations in the FPGAs temperature.

[7]UI symbol (high or low) refers to the logical level (1 or 0) of a transmitted bit held constant over a single Unit Interval. Each symbol occupies exactly one UI.

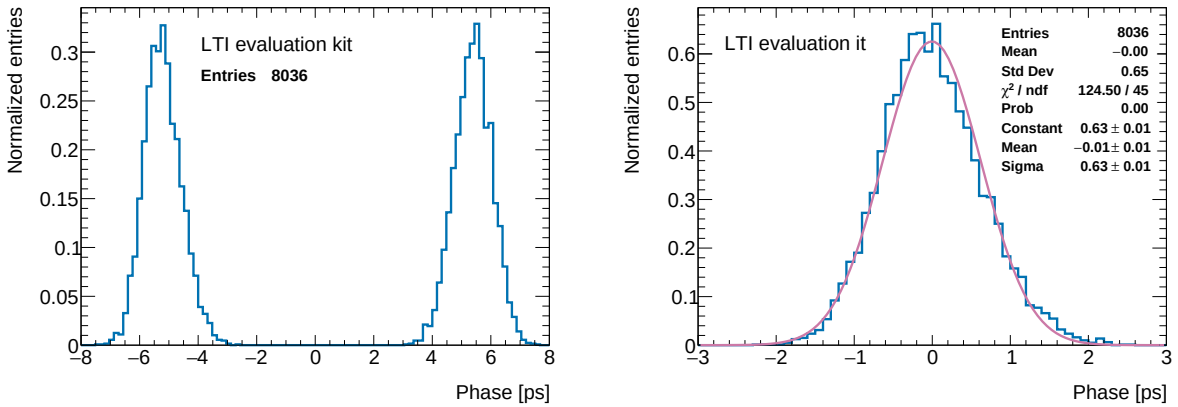

(a) DDMTD measurements before RX-fanout correction.

(b) DDMTD measurements after RX-fanout correction.

Figure 5.21: DDMTD RX-fanout mitigation procedure.

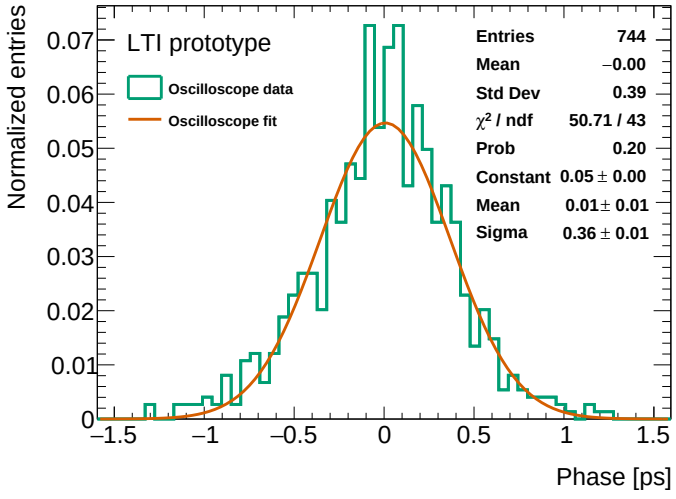

Figure 5.22: LTI prototype recovered clock phase distribution with the "clock fiber" configuration.

5.6.3 Compensation algorithm

The purpose of the phase compensation algorithm is to maintain the recovered clock phase as constant as possible, regardless of any instability. During the future HL-LHC ATLAS operations it will not be possible to have an oscilloscope for each link and for each LTI. Therefore the requirement is to implement a model that would be able to predict the phase drift. The only handles available to feed such a model are the DDMTD and the FPGA silicon die temperatures. Regardless of the implementation of the algorithm, the testing procedure followed the steps:

1. Stabilize the climatic chamber temperature as seen in previous procedures.
2. The script containing the algorithm logic measures a reference initial state.
3. The script controlling the climatic chamber starts the temperature ramp.
4. In a loop, the compensation algorithm evaluates the drift from a reference state.
5. If the drift is comparable to a TxPI step, apply a correction.
6. If needed, update the reference state[8].

In the first implementation, the compensation algorithm relied solely on DDMTD measurements. The DDMTD phase was used as an error signal and mapped onto the true recovered-clock phase using the calibration factor obtained from oscilloscope data. A corrective TxPI action was triggered whenever the DDMTD drift exceeded the equivalent of one TxPI step with respect to the reference. This approach was effective as long as the DDMTD response remained linearly correlated with the oscilloscope, but it became unreliable for temperature excursions of several degrees, where the DDMTD characteristic deviated from linearity and in some cases even inverted its correlation with the true phase.

The DDMTD-based algorithm was tested with LTI prototypes by alternately placing the master or the slave in a climatic chamber and sweeping the temperature by 8 °C. An additional test was performed with 40 m of optical fiber inside the chamber swept by 4 °C. The corresponding results are shown in Figure 5.23, where the y-axis represents the recovered clock phase, measured with both oscilloscope and DDMTD, after applying the DDMTD-based compensation. The x-axis corresponds to the FPGA temperature in panels (a) and (b), and to the chamber temperature in panel (c). For the optical fiber test, the DDMTD signal does not remain constant even though the recovered clock phase appears stable. This originates from the fact that the DDMTD measures the relative phase of both downlink and uplink paths, whereas the oscilloscope observes only the downlink recovered clock. Since the compensation target is the stability of the downlink,

[8]This step depends by the specific algorithm implementation.

it is the oscilloscope measurement that reflects the relevant quantity, while the DDMTD is affected by additional fiber-induced variations outside the scope of the correction.

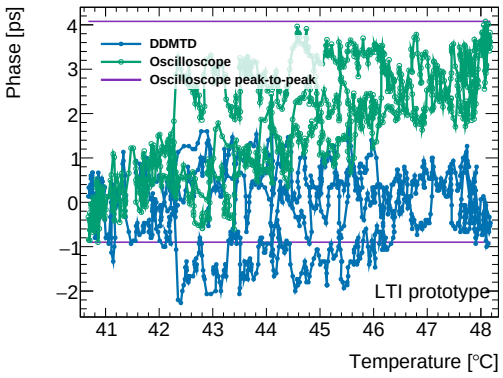

(a) DDMTD compensation of LTI slave.

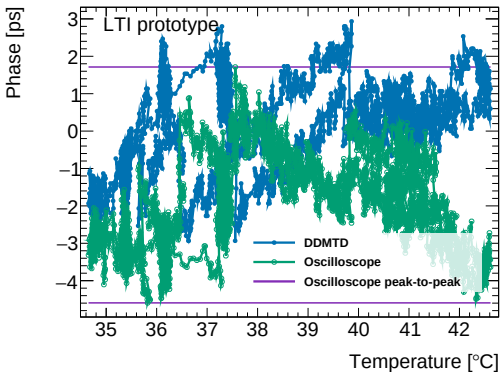

(b) DDMTD compensation of LTI master.

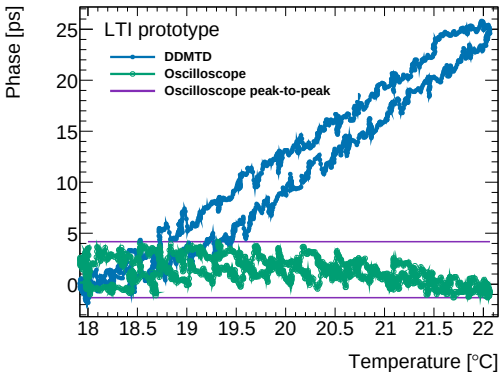

(c) DDMTD compensation of optical fibers.

Figure 5.23: DDMTD phase compensation of either LTI slave (a), LTI master (b) or 40 m of optical fibers (c).

Although the DDMTD-based phase monitoring provides high-resolution tracking of relative phase differences, its performance degrades with temperature variations of 5 °C to 8 °C. Under such conditions, the DDMTD reference becomes unreliable. To overcome this limitation, a compensation algorithm based on local temperature measurements was developed. The approach relies on a linear model that predicts phase shifts as a function of the temperatures at both the LTI master and slave:

$$\phi = \alpha T_{\mathrm{M}} + \beta T_{\mathrm{S}}\,, \tag{5.7}$$

where T_{M} and T_{F} denote the temperatures of the master and slave, respectively. The regression model was trained on data sets where either the master or slave board was placed in a climatic chamber and subjected to controlled thermal sweeps. For each data set, both temperature and phase measurements were transformed into relative displace-

ments with respect to their initial values. The regression was performed without an intercept term to avoid introducing an artificial offset.

Figure 5.24 shows the measured phase drift as a function of the prediction from the temperature-based regression model. The data points represent bin averages, with vertical error bars given by the standard error of the mean in each bin. The Ordinary Least Squares [76] (OLS) regression and the robust Huber estimator [77] yielded consistent coefficients within 5%. The master and slave temperatures contribute $\alpha = (0.87 \pm 0.02)\,\mathrm{ps/°C}$ and $\beta = (0.50 \pm 0.01)\,\mathrm{ps/°C}$, respectively, to the phase drift. The quoted uncertainties are purely statistical, obtained from a weighted regression with per-bin standard errors of the mean. Consequently, the vertical error bars in Figure 5.24 are much smaller than the point-to-point scatter, so they are not visible at the scale of the figure. In contrast, the deviation of the individual points around the regression line is significantly larger (standard deviation $\sim$0.8 ps), reflecting the intrinsic noise of the phase measurement as well as possible systematic effects such as sensor calibration offsets, correlations between the master and slave temperatures, or deviations from linearity. The consistency between OLS and robust regression ($\leq 5\%$ difference in coefficients) suggests that such systematic contributions are limited to the few-percent level. A global goodness-of-fit test gives a reduced $\chi^2 \approx 1.0$ with $p \approx 0.48$, showing that the residuals are consistent with statistical noise and that no significant model misfit is present.

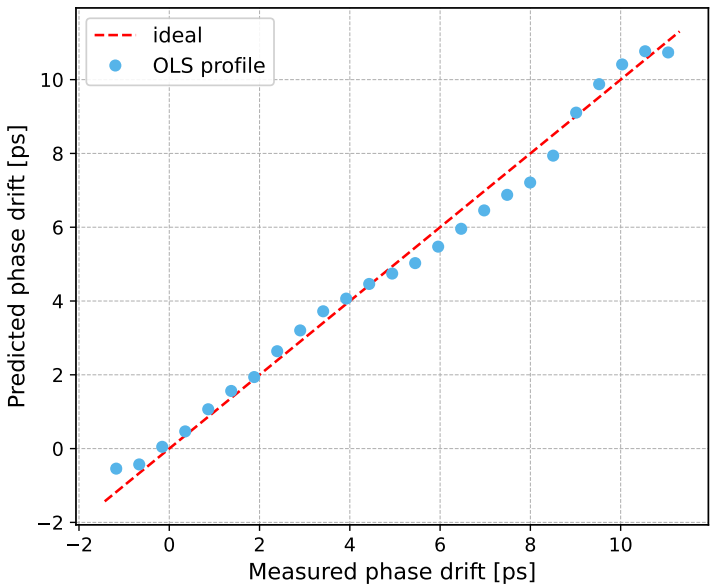

Figure 5.24: Comparison of measured and predicted phase drift using the temperature-based regression model.

The residuals provide further validation. Their distribution is approximately Gaussian with a mean of 0.02 ps, consistent with zero, and a standard deviation of about 0.8 ps (Figure 5.25a). This value matches the known resolution limit of the phase measurement

system and therefore represents the expected noise floor. When grouped as a function of the predicted drift, the residuals fluctuate symmetrically around zero without systematic trends (Figure 5.25b). Together, Figures 5.24 and 5.25 demonstrate both the agreement of the regression model with the data and the statistical consistency of the residuals.

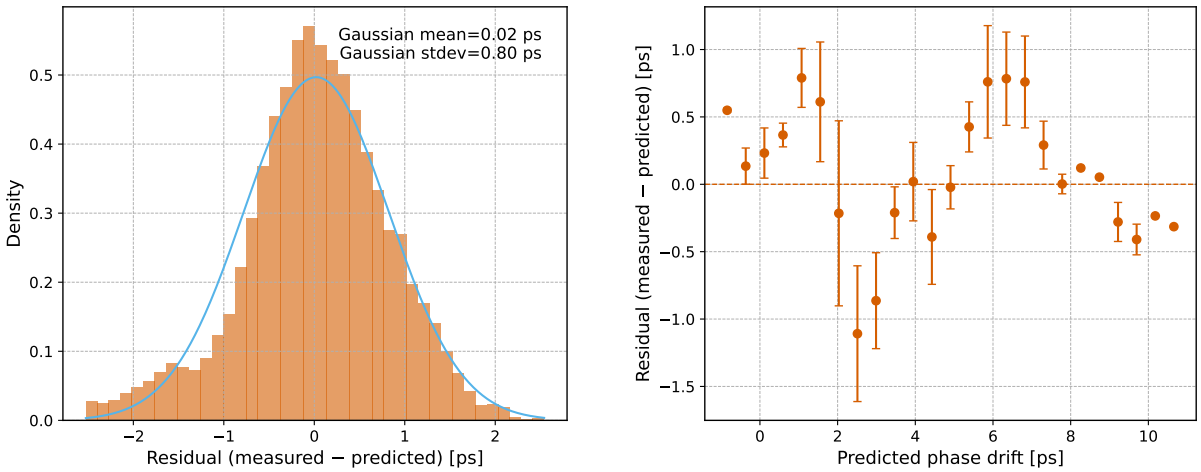

(a) Residuals distribution.

(b) Residuals vs predicted.

Figure 5.25: Residual analysis of the temperature-based regression model.

Before presenting the results of the temperature-based phase compensation, the algorithm's core steps can be summarized as follows:

1. Measure the reference FPGA temperatures on both LTI master and slave.
2. Update in a loop the FPGA temperature variations and the phase drift prediction.
3. If the phase drift is above a threshold, i.e. a single TxPI step size, an corrective action is triggered. The sign of the predicted phase drift determines whether the TxPI is incremented or decremented.

The Figures 5.26b and 5.26a show the compensation algorithm in action on a 8 °C positive and negative ramp. The y-axis is the recovered clock phase measured with both oscilloscope and DDMTD, after applying the correction. The x-axis is the FPGA temperature of the relevant board in the chamber. Figure 5.26b illustrates a scenario where DDMTD-based compensation could have failed. Between the start and 20 minutes into the test[9], the DDMTD lost its positive correlation with the phase, and would have triggered a TxPI correction with the wrong sign.

Tests with both LTIs in independent climatic chambers were not possible due to the ATCA shelf size and logistics (see Section 5.1.4). Although it was not possible to control both LTIs independently, the compensation model accounts for the individual con-

[9]Corresponding to FPGA temperatures from 41 °C to 42 °C.

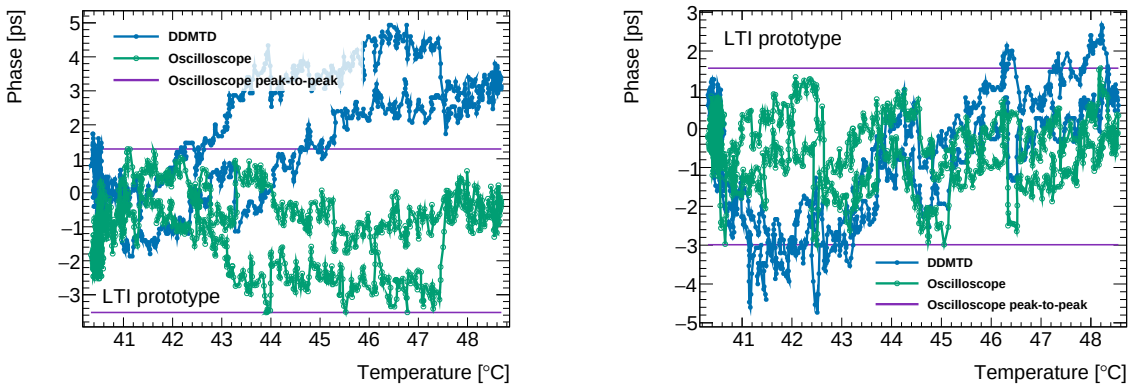

(a) LTI slave in the climatic chamber.

(b) LTI master in the climatic chamber.

Figure 5.26: Phase compensation algorithm in action, with either LTI master or slave in the climatic chamber.

tributions of master and slave temperatures, rather than their difference. As such, no fundamental limitation is expected, although this configuration was not tested.

Fibers were not included in the model because the typical temperature shift in the counting cavern over a run (0.5 °C) would not have a large impact on timing. Quantitatively, using an optical fiber of 40 m would result in a phase drift of 3.42 ps. A phase drift of this magnitude can be compensated with two TxPI step correction, triggered by a periodic reading of environmental monitoring data.

Table 5.4 summarizes the phase drift amplitudes observed with the oscilloscope on the recovered clock phase under different compensation strategies. The results are organized by subsystem. To suppress short-term measurement noise while preserving the underlying thermal drift, the phase data were smoothed using a moving average with a window size of three samples (corresponding to 15 s at the acquisition rate). This value was chosen as a compromise between noise reduction and temporal resolution: it is large enough to suppress random fluctuations due to instrument jitter, yet small enough to preserve transient features in the thermal response. Empirically, increasing the window further did not significantly improve stability metrics, while reducing it led to noisier estimates of the drift amplitude.

Table 5.4: Summary of compensation outcomes

	Oscilloscope phase drift amplitude [ps]		
In KB260	No corrections	DDMTD-based	Temperature-based
Slave	7.3 ± 2.1	5.0 ± 2.0	4.8 ± 1.8
Master	8.2 ± 2.2	6.3 ± 2.3	4.5 ± 2.0
Fibers	27.4 ± 1.9	5.5 ± 1.8	n/a

These results confirm that temperature-based compensation improves phase stability compared to DDMTD-only methods. However, residual drifts indicate that further refinement is required. When using only the DDMTD, a key limitation remains the difficulty of disentangling the origin of DDMTD variations, particularly when all boards and optical fibers are subject to correlated temperature changes. Notably, tests conducted across different firmware versions revealed variations in the regression coefficients of the temperature-based model, suggesting that internal signal routing and logic implementation can influence the thermal sensitivity of the measured phase. Future improvements are likely to come from combining multiple compensation strategies—leveraging both temperature and DDMTD data—to exploit the complementary strengths of each approach. Rather than relying on a single source of information, a hybrid algorithm that integrates all available observables may offer more robust performance across varying conditions. Additionally, these tests should be repeated on different hardware platforms to ensure that the observed behavior is not specific to a particular board revision or layout.

5.7 Integration tests

The first two integration tests of the LTI prototype with ATLAS Phase-2 sub-detectors took place in late June 2025. The first test was conducted with the LAr calorimeter during the LAr Integration Week (June 23–25) at the Electronics Maintenance Facility at CERN, followed by a second test with the HGTD during the FELIX expert week (June 25–27). These sessions focused on validating the physical and protocol-level interface compatibility, testing serial communication links, and performing signal integrity measurements.

5.7.1 Liquid Argon calorimeter

The LTI board was tested within the LAr timing distribution chain [78]. The setup consisted of **LTI → FELIX → lpGBT board → LAr Front-End Board prototype (FEB2)**. As a reminder, the lpGBT is a high-speed transceiver ASIC serving as interface between FELIX and the front-end electronics. The main objectives were to validate LTI signal transmission, synchronization, and compatibility with the front-end electronics.

The following aspects were tested:

- Installation and continuous operation of the LTI with loopback functionality to monitor outgoing TTC data.
- Programmable L0A and Pulse command generation from the LTI, with synchronization confirmed down to FELIX.

- Decoding of TTC signals (Bunch Counter Identification, L0A, Trigger Type, Pulse command by the LAr firmware on the lpGBT interface board and the FEB2.

- A preliminary estimate of the end-to-end timing jitter was obtained from oscilloscope measurements of the clock-edge differences between the LTI and the FEB2. The observed distribution had a full width at half maximum (FWHM) of about 50 ps, corresponding—under the assumption of Gaussian statistics—to a standard deviation of roughly 21 ps via FWHM $= 2.355\,\sigma$. This estimate should be regarded as approximate, as the measurement precision was limited by suboptimal probe conditions [78].

5.7.2 High-Granularity Timing Detector

During the HGTD FELIX expert week, the Local Trigger Interface (LTI) was used as the initial source in the clock distribution chain to the HGTD modules, following the path: LTI → FELIX → lpGBT → Module Flex (see Figure 5.27). The Module Flex is a flexible PCB interfacing the HGTD sensors to the optical TTC readout path (see Section 4.6). The main focus of the measurements was to study startup phase determinism, assessing how reliably the recovered clock phase was preserved through each component after system resets.

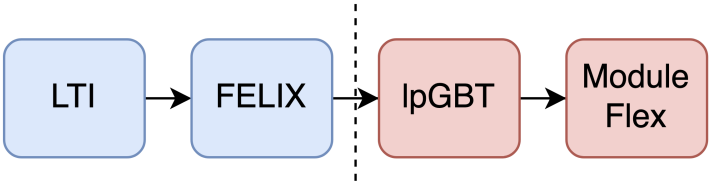

Figure 5.27: Clock distribution chain from LTI to HGTD Module Flex used for integration tests during the HGTD FELIX expert week. In blue the "off-detector" elements, while in red the "on-detector" electronics.

For this test, only the FELIX downlink RX was reset, while the TX was left untouched to preserve downstream link alignment. This procedure isolates the recovery behavior of the RX and lpGBT clock recovery, but does not test the determinism of a full FELIX reset sequence. After each RX reset, the phase difference between the recovered clock at the Module Flex and the reference clock from the LTI was measured directly with an oscilloscope. The resulting distribution across multiple resets, along with the corresponding two-dimensional histogram versus FPGA temperature, are shown in Figure 5.28. The two-dimensional histogram indicates that outlier phase values predominantly occurred when the FPGA temperature was outside the nominal range, corresponding to early conditions before the system reached thermal equilibrium [79].

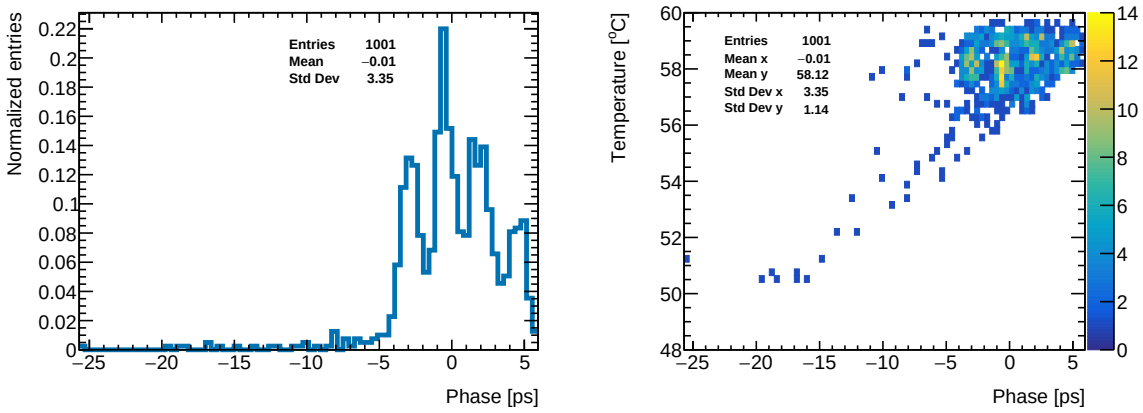

Figure 5.28: Left: Startup phase difference between LTI and Module Flex, measured with an oscilloscope after repeated resets of the FELIX RX transceiver. Right: Startup phase versus FPGA temperature.

In addition, the continuous RMS jitter of the FELIX→lpGBT→Module Flex segment was determined from phase-noise measurements integrated over the 0–1,MHz band, yielding 0.85 ps.

Based on the measured RMS jitter and the observed reset-induced phase spread, treated here as independent RMS-like contributions, a hypothetical full jitter and startup phase uncertainty budget—from the White Rabbit input at the CTP down to the HGTD Module Flex, and at stable temperature conditions—can be constructed. This estimate is summarized in Table 5.5.

The total combined contribution from the known terms, added in quadrature under the assumption of statistical independence, amounts to 6.34 ps. This includes the jitter of the White Rabbit → CTP reference (5.00 ps), the FELIX→lpGBT→Module Flex chain (0.85 ps), and conservative approximations of 1.00 ps each for the GTH transceivers in the CTP and LTI. However, this total does not yet account for contributions from the lpGBT and the Module Flex endpoints themselves, which are expected to be at the picosecond-level, but nonzero.

Assuming statistical independence among all effects, the residual margin σ_{margin} available for other instabilities (e.g. thermally induced phase drifts in electronics or optical fibers) can be estimated as:

$$\sigma_{\text{margin}} = \sqrt{(30\ \text{ps})^2 - (6.34\ \text{ps})^2} \simeq 29.3\ \text{ps}, \tag{5.8}$$

where the 30 ps corresponds to the per-track timing resolution target of the HGTD.

Including two additional independent terms of order 1 ps—one from the lpGBT and one from the Module Flex—raises the total estimated uncertainty from 6.34 ps to 6.50 ps. This results in the same residual quadrature margin at the first decimal digit:

$$\sigma_{\text{margin}} = \sqrt{(30\text{ ps})^2 - (6.50\text{ ps})^2} \simeq 29.3\text{ ps}. \tag{5.9}$$

Thus, under stable thermal conditions, the overall uncertainty budget remains robust. The upstream White Rabbit contribution (5 ps) alone accounts for nearly 80% of the total variance, making it the dominant term.

This estimate assumes that all listed contributions are statistically independent and that their RMS values can be combined in quadrature. Any correlation among sources of phase uncertainty—particularly in temperature-dependent effects—would invalidate this assumption and reduce the effective margin. In particular, the presence of thermal gradients or transient conditions may introduce coherent drifts across multiple components, significantly degrading the available timing budget if not properly compensated. The estimated margin should therefore be regarded as a best-case scenario, further motivating the need for a compensation algorithm that is as efficient and robust as possible.

Table 5.5: RMS jitter budget and startup phase spread to the HGTD Module Flex.

Segment	Continuous jitter [ps]	Startup spread [ps]	Source / Note
White Rabbit → CTP	5.00	–	Target jitter from White Rabbit system (Section 3.3.5).
CTP → LTI	0.50	1.00	Continuous jitter (Section 5.1.3); Conservative value phase uncertainty after reset from (Section 5.4).
LTI internal links	0.50	1.00	As above, values from Section 5.1.3 and Section 5.4.
FELIX → lpGBT → Module Flex	0.85	3.35	Continuous jitter from phase-noise measurement (0–1 MHz) [79]; startup spread from oscilloscope after repeated FELIX RX resets [79].
lpGBT / Module Flex reset (isolated)	–	0.84	Assumed from standalone lpGBT phase-stability tests [80]. No direct Module Flex-only measurement available.
Quadrature total	**5.12**	**3.73**	Sum in quadrature of contributions above.
Combined total	**6.34**		Quadrature of continuous jitter and startup spread totals.

Figures 5.29 and 5.30 show the experimental setup used during this HGTD-LTI integration test campaign, with the LTI integrated into the FELIX-based DAQ system.

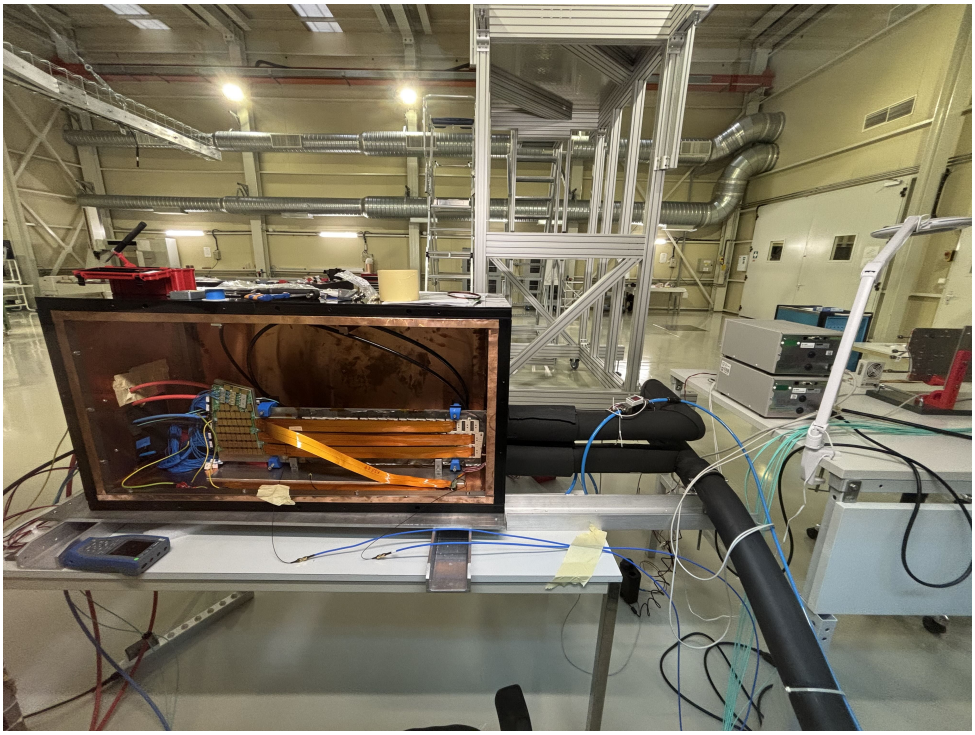

Figure 5.29: HGTD module setup inside a copper-shielded box. The orange flex PCBs (Module Flex) connected the lpGBT outputs to the timing detector modules.

Figure 5.30: FELIX-based data acquisition rack during integration tests. The LTI (bottom) distributed the clock to FELIX, while an oscilloscope monitored the startup phase of the recovered clock at the Module Flex against the LTI reference clock.

Chapter 6

Operational strategies and guidelines for the LTI software

This chapter outlines a proposed design for a software application that integrates the LTI phase monitoring and control in the ATLAS central remote control framework. The strategies and guidelines described in this chapter are not yet part of an official or finalized ATLAS TDAQ procedure. They reflect the author's development work on the LTI software, including contributions to low-level functions implementation and debugging, as well as insights from link bring-up tests. As such, the discussion should be regarded as a proposed integration concept, informed by initial experience rather than established operational documentation. In the following, it is assumed that this application implements a phase compensation algorithm based on both DDMTD and FPGA silicon die temperatures. In this discussion it is also assumed that an LTI is slave of the CTP, and master of a FELIX board. Figure 6.1 shows the interaction between the LTI control software, the ATLAS central remote control system, and other applications controlling the CTP and FELIX systems. The ATLAS operator in the control room of the experiment issues commands which coordinates the various applications via Finite States Machine (FSM) transitions and commands (e.g., `INITIALIZE`, `CONFIGURE`, `START`, `STOP`). Each application interfaces with its respective hardware using low-level function calls. In addition, each of these applications is interfaced with the configuration manager, the middleware layer for information sharing, and the error reporting system (see Section 2.2.11). The use of the information sharing layer enables a flexible implementation that remains highly maintainable and can readily evolve as operational experience is gained, allowing the system to grow and mature without requiring disruptive redesigns.

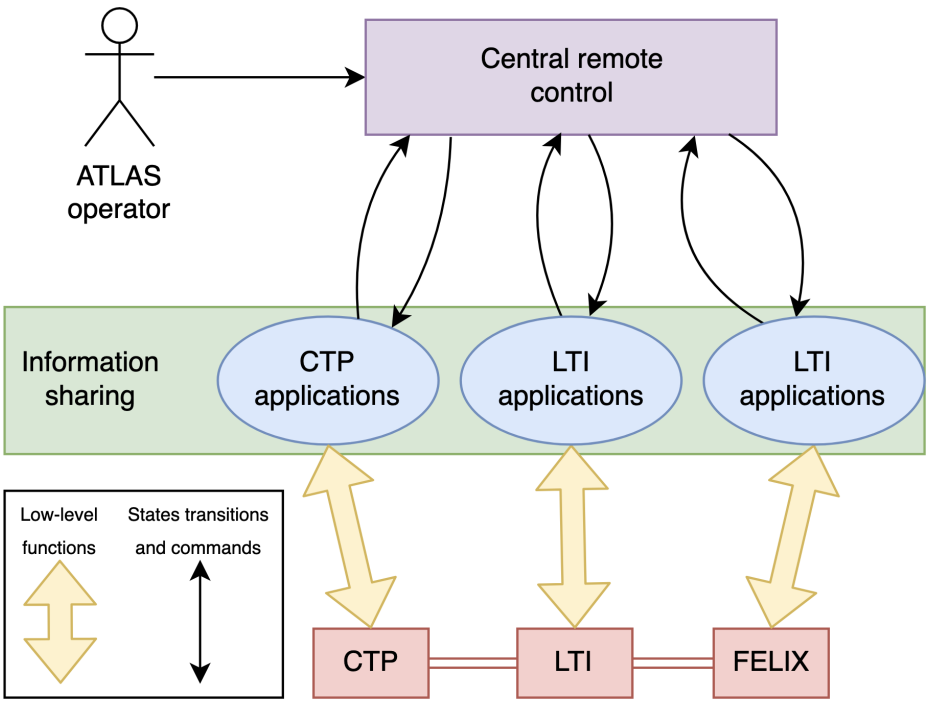

Figure 6.1: Context diagram of the LTI application within the ATLAS remote control system.

6.1 Link bring-up and board configuration

Prior to addressing the phase monitoring and control application, it is useful to first consider the integration of the LTI link bring-up procedure.

In the ATLAS TDAQ Finite-State model, it is common for hardware configuration tasks to be executed either when the application is in the `INITIAL` state or during the transition to `CONFIGURED`. The early execution of the LTI link bring-up is particularly important because many downstream hardware components rely on receiving a stable clock signal before their own configuration can complete. For this reason, the LTI should establish its links and begin clock distribution as soon as possible within the initialization sequence. The exact bring-up timing must be agreed with each sub-detector, and appropriate synchronization measures implemented, to ensure that all systems receive the clock early enough for proper operation. Once the configuration is complete, the partition enters in data-taking mode with the transition to `RUNNING`.

As described in Section 4.10.3, the link bring-up process requires coordination between two different boards. To ensure proper synchronization during this phase, the procedure should be integrated into the ATLAS middleware information sharing layer. This would allow both boards to publish their status information and subscribe to the counterpart's readiness signals, enabling a controlled and fault-tolerant initialization. In this way, any partial failures or delays can be detected in real time and reported to the central re-

mote control system, where the ATLAS operator can take appropriate action, while still allowing the system to proceed with the subset of operational links.

Configuration steps are also subject to timing constraints. In general, all configuration procedures should complete within a typical upper bound of about 30 s. This means that when ATLAS transitions to the `CONFIGURED` state, each LTI board is expected to establish communication across its optical links (60 LTI-out and 4 LTI-in) before moving on to the `RUNNING` state.

A single link initialization typically takes about 1 s. In the current implementation, the per-link bring-up routine contains blocking waits (e.g. `sleep` calls) for PLL/CDR lock and alignment. A naive sequential approach would therefore result in prohibitive total latency and introduce a single-point failure risk — if one link fails, the entire procedure stalls. Two mitigation strategies are proposed:

- **Multithreading:** Each group of links (e.g., those associated with the same FireFly module pair) is initialized by an independent thread. Since each thread operates on its own set of link and MGT registers, mutexes are not required, except when accessing shared resources such as the clock configuration, which must remain serialized. This solution allows blocking waits in different groups to overlap, reducing the total bring-up time to the duration of the slowest group.

- **Round-robin polling:** All links are started, and their status is polled sequentially until all are up. This method is only beneficial if bring-up is implemented in a non-blocking manner, where each link's status can be checked while others progress in the background. With the current blocking waits, round-robin reduces to sequential initialization. Removing the blocking behavior would make this method effective, as multiple links could then progress in parallel.

Only the link initialization process is parallelized; subsequent monitoring tasks can operate in a sequential round-robin or pipelined manner. If one or more links fail to establish communication within a programmable timeout, this condition is reported to the ATLAS remote control system. Importantly, the board is still marked as `CONFIGURED` and will operate with the functional subset of links, with the degraded status made visible to the operator through the monitoring interface.

6.2 Start of monitoring and compensation logic

After the board transitions to the `CONFIGURED` state, the monitoring and compensation application begins link monitoring, which can be triggered either by a direct command or automatically based on configuration parameters. The set of links to be monitored is

retrieved from the information sharing layer, ensuring that the monitoring logic remains flexible and easy to extend as operational needs evolve. At this point, the software:

- Publishes DDMTD phase measurements for all active LTI-out links;
- Publishes the local FPGA silicon die temperature;
- Subscribes to temperature data from downstream timing boards (e.g., FELIX).

The compensation algorithm processes both the DDMTD phase measurements and the temperature values from the slave board to estimate and, if necessary, correct for phase drift.

A visual summary of the monitoring control logic and FSM behavior is provided in Figures 6.3 and 6.2.

6.3 Periodic monitoring and compensation loop

In order to reduce complexity, phase monitoring and compensation are handled by a single periodic task scheduled every 3 s to 7 s. This probing interval replicates the intervals used in the experiments shown in Chapter 5. The task is executed in every state which is either `CONFIGURED`, `RUNNING`, performing in order: (i) acquisition of DDMTD measurements, (ii) acquisition of the local FPGA die temperature, and (iii) retrieval of remote temperature values via the information sharing layer. Based on these inputs, a simple decision rule determines whether a TxPI correction is required on each active link.

If one input is temporarily unavailable (e.g. missing subscribed information update), the loop proceeds with the available data and records a degraded mode flag in the published monitoring. Link-level timeouts and basic sanity checks prevent stale or inconsistent measurements from triggering TxPI actions. The loop publishes its inputs and decisions to the information sharing layer after each iteration. Statistics about application's activity are also published at a lower rate (e.g. every 5 minutes).

6.4 End of run and transition to idle

When ATLAS transitions out of the `RUNNING` state (e.g., into `STOPPED` or `UNCONFIGURED`), the monitoring application must gracefully disable all active measurement and compensation tasks. This includes:

- Halting the publication of DDMTD and temperature data;
- Notifying the remote control system of the transition.

This termination phase should be performed with care to avoid misalignments in timing that could affect subsequent reconfigurations. A clean shutdown ensures that the next `CONFIGURED` transition starts from a known, deterministic state.

6.5 Error handling and degraded operation

Given the distributed and asynchronous nature of the LTI software operations, various error scenarios must be considered. These include partial link failures, missing data from remote boards, and internal timing inconsistencies. This section summarizes how such conditions are detected and managed.

Link bring-up failures. As described in Section 4.10.3, each LTI link is initialized at the start of a run. If one or more links fail to come up within the predefined timeout, the application reports the affected links to the central remote control system, allowing the operator to decide whether to proceed to the `CONFIGURED` state without them or to follow up on the issue before continuing. This explicit reporting ensures that potentially significant data losses — such as a large fraction of a sub-detector being unavailable — are visible to the operator, who can then make an informed choice rather than unknowingly continuing in a degraded configuration.

DDMTD data loss or corruption. If DDMTD measurements are missing or corrupted, the software flags the affected link and temporarily excludes it from phase compensation, falling back to a reduced model. If the issue persists beyond a configurable threshold, a warning is raised to the operator interface. In the meantime, temperature-only compensation can continue as an alternative mode.

Remote temperature source unavailability. Temperature readings from connected boards (e.g., FELIX, CTP) are accessed via a publish-subscribe model. If these values are not received within a configurable window, the compensation algorithm falls back to a reduced model using only local measurements. This degradation is reported in the control interface, but it does not interrupt ongoing monitoring.

Thread-level exceptions. In the multithreaded link bring-up phase, unhandled exceptions or thread stalls are isolated using watchdog timers. The main control loop monitors thread status and logs any anomalies. Threads that fail to complete are force-terminated and marked as failed, but do not block global operation.

Run termination under faulty conditions. If the application loses all available inputs required for phase compensation (e.g., DDMTD measurements and temperature data from both local and remote boards) or encounters an unrecoverable software fault, it reports an error state to the remote control system. In this state, TxPI corrections are suspended, and the board continues operating with the last valid configuration to avoid introducing uncontrolled timing shifts. The condition is logged and displayed to the operator.

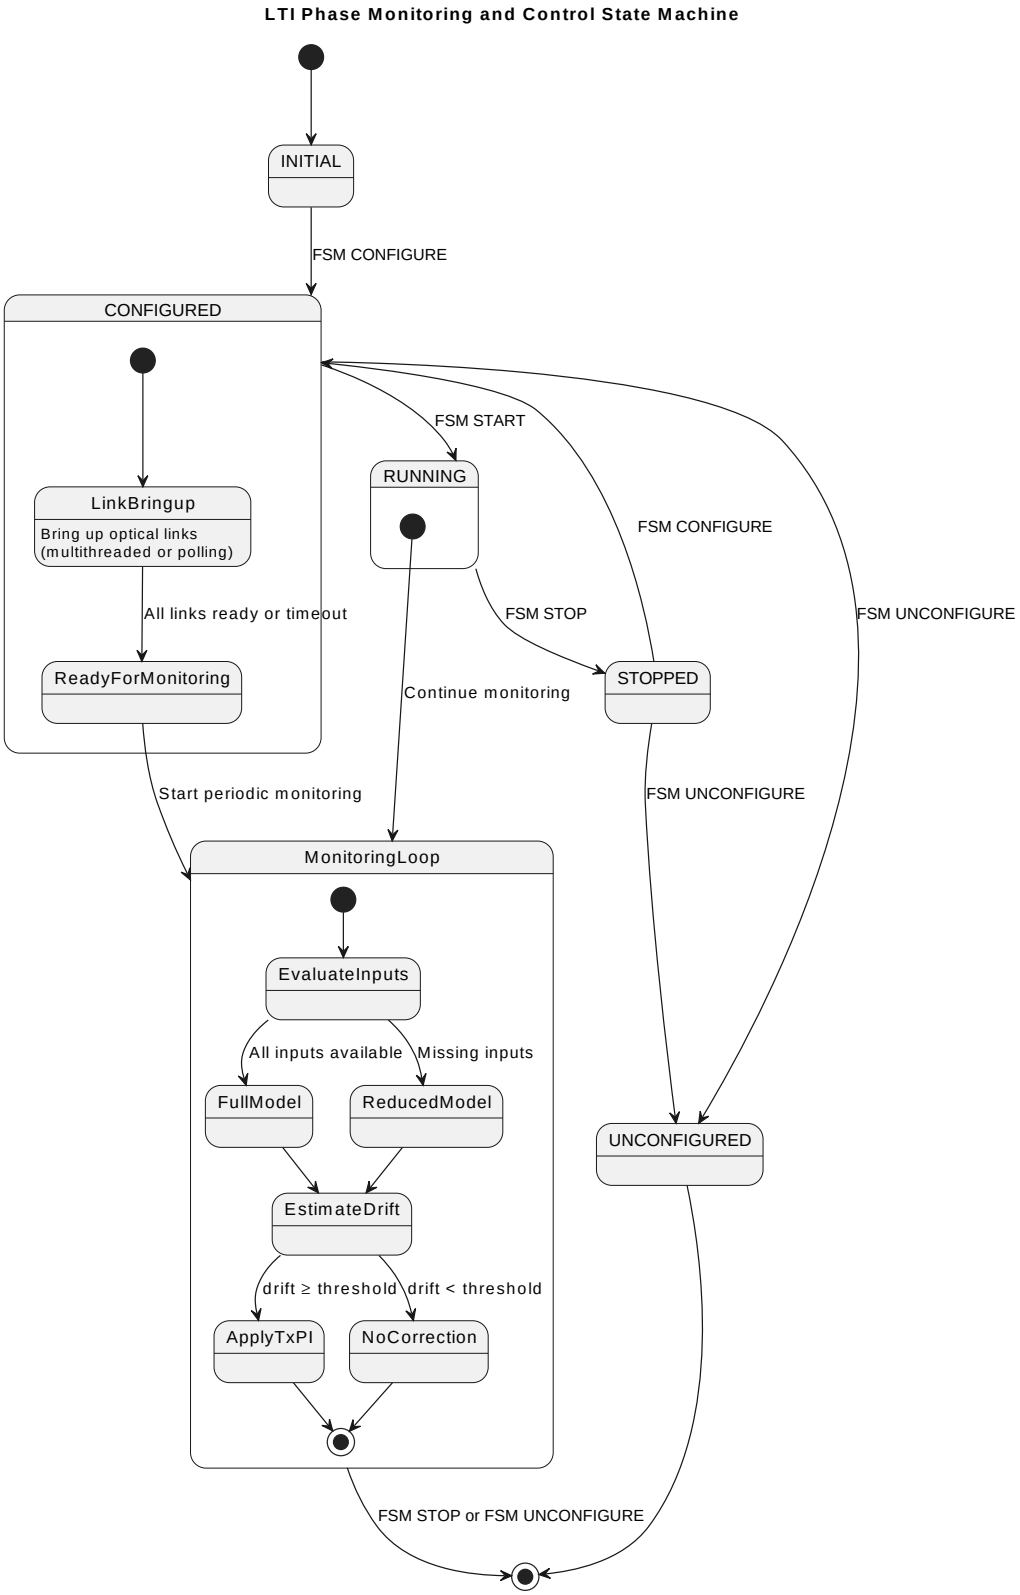

Figure 6.2: State machine diagram of the LTI software aligned with the Finite State Machine (FSM) transitions of the ATLAS remote control system.

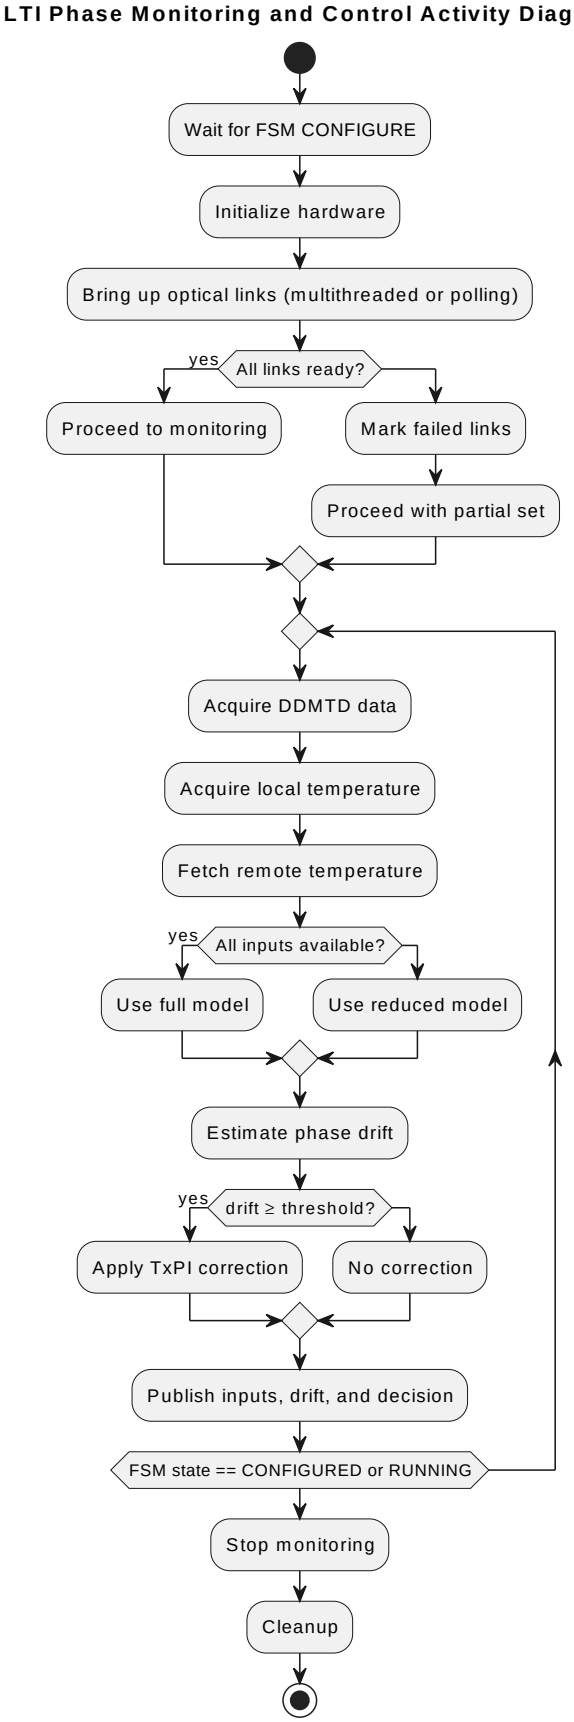

Figure 6.3: Activity diagram of the LTI phase monitoring and compensation loop.

Chapter 7

Conclusions & outlook

This thesis focused on the qualification and performance evaluation of the Local Trigger Interface (LTI) as the upgraded timing distribution element for the ATLAS Trigger and Timing Control system in preparation for Phase-2 operations at the High-Luminosity LHC. The increased luminosity and pile-up foreseen for the HL-LHC era impose strict requirements on detector synchronization, particularly for high-precision timing sub-detectors such as the High-Granularity Timing Detector, which requires a time resolution better than 30 ps.

Phase-stability studies identified two main sources of timing uncertainty: nondeterministic phase after transceiver reset, and phase drifts due to the FPGA silicon die temperature. The LTI's in-situ phase measurement and adjustment capabilities – the Digital Dual Mixer Time Difference and the Transmitter Phase Interpolator, respectively – were characterised in detail. Three mitigation strategies were investigated and validated, enabling deterministic link bring-up within measurement precision and active compensation of thermal drift. Laboratory tests and integration studies with ATLAS sub-detectors demonstrated that, with phase-compensation active, the LTI maintained the recovered clock phase within the resolution of the correction mechanism, requiring fewer than two correction actions over the measurement period. No measurable additional jitter was observed under integration test conditions.

Integration tests with the Liquid Argon calorimeter and the HGTD confirmed compatibility with existing sub-detector TTC receiver electronics and validated clock stability under realistic operating conditions. The results establish that the LTI meets the Phase-2 timing distribution requirements and that the developed monitoring and correction procedures are effective in preserving the phase stability essential for efficient event building and optimal sub-detector performance.

The performance evaluation and validation work reported here substantially reduces the technical risk associated with Phase-2 timing distribution in ATLAS. Nonetheless, further work remains before HL-LHC operations:

- **Decoupling functions from the FPGA:** The DDMTD hardware implementation could be moved on a separate chip, avoiding the non-linear dependance from the FPGA temperature. In the same manner, the CDR firmware block could be moved to the SoM.
- **Full-scale tests:** The tests illustrated in chapter 5 will need to be excercised on all the LTI-in and LTI-out links. Furthermore, with the incoming Long Shutdown 3, all the other relevant ATLAS Phase-2 systems will need integration tests with the LTI.
- **Control software:** Development of production-ready tools for phase monitoring, correction, and recovery, as outlined in the guidelines of chapter 6.
- **Error detection and resilience:** Automatic detection and recovery from TTC link failures or anomalous phase shifts without operator intervention.

These results demonstrate that a deterministic, low-jitter timing distribution with sub-30 ps stability is achievable in ATLAS Phase-2. The characterisation and validation methods presented here may also find application in other large-scale high-energy physics experiments requiring high-precision timing distribution.

Acknowledgements

I would like to express my deepest gratitude to my supervisors, Prof. Dr. Hans-Christian Schultz-Coulon and Dr. Thilo Pauly. Without their support, it would not have been possible to pursue this path. They always took the time to listen to me and to offer invaluable advice, for which I am truly grateful. I also wish to extend my sincere gratitude to Dr. Antoine Marzin, who was always there whenever I needed help and whose work I deeply admire.

I am very pleased to thank Dr. Tigran Mkrtchyan for his kindness and friendship, which I greatly appreciate. I am also grateful to Dr. Rosa Simoniello, whose constant care and readiness to help were a source of comfort, particularly in the most difficult moments.

I would especially like to thank Dr. Aimilianos Koulouris, who was always willing to take the time for discussions and to provide support whenever needed. I am also thankful to Dr. Kristina Mihule for her help and for the many stimulating discussions we shared.

I would like to express my appreciation to Dr. Stefan Haas and Dr. Ralf Spiwoks, whose assistance with hardware and readiness to explain and suggest solutions were of great help.

I am also very grateful to Dr. Yoav Afik, whose friendship has been a great source of encouragement.

I would like to warmly thank two fantastic friends I made during my time at CERN, Jan Brajer and Pratik Jawahar, whose companionship made this experience all the more rewarding.

I would also like to thank all the friendly people at the Kirchhoff Institute of Physics in Heidelberg, in particular Thomas Junkermann and Lisa Marie Baltes. I am also grateful to Ms. Petra Pfeifer, whose unfailing support with administrative matters was essential.

I extend my thanks to Prof. Dr. Fischer for acting as my referee and for taking the time to read my thesis.

I owe my deepest thanks to my mother, whose encouragement carried me through to this point. Finally, I am especially grateful to Lana for her constant support and trust in me. I look forward to continuing this journey together.

Bibliography

[1] R. L. Workman et al. "Review of Particle Physics". In: *Progress of Theoretical and Experimental Physics* 2024.8 (2024), p. 083C01. DOI: 10.1093/ptep/ptae003. URL: https://pdg.lbl.gov.

[2] CERN. *LHC Long-Term Schedule*. https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm. Accessed: 2025-07-04.

[3] CERN. *CERN Accelerator Complex [SVG file]*. Wikimedia Commons. Accessed: 3 July 2025. 2025. URL: https://commons.wikimedia.org/wiki/File:Cern-accelerator-complex.svg.

[4] Wikipedia contributors. *VMEbus*. https://en.wikipedia.org/wiki/VMEbus.

[5] Bruce Taylor. "TTC machine interface (TTCmi)". In: *CERN EP* (Feb. 2000).

[6] S Baron and A Monera-Martinez. "Status of the TTC upgrade". In: (2007). DOI: 10.5170/CERN-2007-001.236. URL: https://cds.cern.ch/record/1027471.

[7] ATLAS Collaboration. *ATLAS Liquid Argon Calorimeter: Technical Design Report*. CERN, 1996. URL: https://cds.cern.ch/record/331062.

[8] "Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV". In: *Eur. Phys. J. C* 80 (2020), p. 119. DOI: 10.1140/epjc/s10052-019-7599-0. URL: https://arxiv.org/abs/1906.12226.

[9] *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*. Tech. rep. Geneva: CERN, 2017. DOI: 10.17181/CERN.FOZZ.ZP3Q. URL: https://cds.cern.ch/record/2285585.

[10] *Technical Design Report: A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade*. Geneva: CERN, 2020. URL: https://cds.cern.ch/record/2719855.

[11] Yoav Afik et al. *Integration and commissioning of the ATLAS Muon-to-Central-Trigger-Processor Interface for Run-3*. Tech. rep. TWEPP2021 conference; deadline 31/10/2021. Geneva: CERN, 2021. URL: https://cds.cern.ch/record/2788530.

[12] Aimilianos Koulouris et al. *Phase-II Upgrade of the ATLAS L1 Central Trigger*. Tech. rep. Geneva: CERN, 2024. URL: https://cds.cern.ch/record/2905971.

[13] F. Bonini and L. Sanfilippo. *High-Precision Timing Distribution in the ATLAS Phase-2 Upgrade*. Presented at the poster session of the ATLAS week. On behalf of the ATLAS Level-1 Central Trigger Team. Oct. 2024. URL: https://indico.cern.ch/event/1447223/contributions/.

[14] G. Aad et al. "The ATLAS Experiment at the CERN Large Hadron Collider". In: *JINST* 3 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003. URL: https://doi.org/10.1088/1748-0221/3/08/S08003.

[15] G. Aad et al. "The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3". In: *Journal of Instrumentation* 19.05 (May 2024), P05063. DOI: 10.1088/1748-0221/19/05/p05063. URL: http://dx.doi.org/10.1088/1748-0221/19/05/P05063.

[16] *ATLAS magnet system*. Geneva: CERN, 1997. DOI: 10.17181/CERN.9O5C.VDTM. URL: https://cds.cern.ch/record/338080.

[17] Claus Grupen, Armin Böhrer, and Ludek Smolik. *Particle detectors*. Cambridge monographs on particle physics, nuclear physics, and cosmology. Cambridge: Cambridge University Press, 1996. URL: https://cds.cern.ch/record/306826.

[18] ATLAS Collaboration. *Insertable B-Layer Technical Design Report*. CERN, 2012. URL: https://cds.cern.ch/record/1291633.

[19] *ATLAS inner detector*. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331063.

[20] ATLAS Collaboration. "Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: tracker properties". In: *Journal of Instrumentation* 12.05 (May 2017), P05002. DOI: 10.1088/1748-0221/12/05/P05002. URL: https://dx.doi.org/10.1088/1748-0221/12/05/P05002.

[21] *ATLAS muon spectrometer*. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331068.

[22] "Commissioning of the ATLAS Muon Spectrometer with Cosmic Rays". In: *Eur. Phys. J. C* 70 (2010), pp. 875–916. DOI: 10.1140/epjc/s10052-010-1415-2. arXiv: 1006.4384. URL: https://cds.cern.ch/record/1275998.

[23] T. Kawamoto et al. *New Small Wheel Technical Design Report*. ATLAS New Small Wheel Technical Design Report. 2013. URL: https://cds.cern.ch/record/1552862.

[24] Savannah Clawson. "Overview of ATLAS forward proton detectors: status, performance, and new physics results". In: *European Physical Society Conference on High Energy Physics (EPS-HEP2023)*. CERN for the benefit of the ATLAS Collaboration. Hamburg, Germany, 2023. eprint: ATL-FWD-PROC-2023-003. URL: https://cds.cern.ch/record/2854602.

[25] C. Ohm and T. Pauly. "The ATLAS beam pick-up based timing system". In: *Nucl. Instrum. Methods Phys. Res., A* 623 (2010), pp. 558–560. DOI: 10.1016/j.nima.2010.03.069. URL: https://cds.cern.ch/record/1172170.

[26] Eva Barbara Holzer. "The LHC Beam Loss Monitoring System: Overview and Experience". In: *7th DITANET Topical Workshop on Beam Loss Monitoring*. DESY, Hamburg, 2011. URL: https://cds.cern.ch/record/1388844.

[27] Antoine Marzin. "L1 latency and rates for Run 3". Presentation, slides available as latency_run3.pdf. ATLAS Weekly meeting, CERN. Oct. 2020. URL: https://indico.cern.ch/event/963953/.

[28] *Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System*. Geneva: CERN, 2017. DOI: 10.17181/CERN.2LBB.4IAL. URL: https://cds.cern.ch/record/2285584.

[29] M Aleksa et al. *ATLAS Liquid Argon Calorimeter Phase-I Upgrade*. Final version presented to December 2013 LHCC. 2013. URL: https://cds.cern.ch/record/1602230.

[30] Shunichi Akatsuka et al. *Final Design Report of the ATLAS Level-1 Endcap Muon Trigger in the Phase-I Upgrade*. Version Draft 0.2. Last modified: March 7, 2017, 01:24. CERN, Mar. 2017. URL: https://indico.cern.ch/event/615176/attachments/1421502/2181543/L1MUPhase1FDR_0.2.pdf.

[31] R. Spiwoks. *The Busy Handling in the ATLAS Level-1 Central Trigger Processor*. ATLAS Internal Note ATL-DA-ER-0025. Version 1.2. ATLAS Collaboration, Oct. 2007. URL: https://edms.cern.ch/document/925303.

[32] ATLAS Collaboration. *Dead-time generation algorithms in the ATLAS Central Trigger Processor*. ATLAS internal note. CERN, 2024.

[33] G. Schuler and R. Spiwoks. *The Core Module of the Central Trigger Processor (CTP_CORE) — User Manual*. ATLAS Internal Note ATL-DA-ER-0023. Version 1.2. ATLAS Collaboration, Sept. 2006.

[34] ATLAS Level-1 Central Trigger Processor Group. *CTP_CORE Pre-scaled Clock and Random Triggers*. ATLAS Internal Note. Describes the configuration and operation of the CTP_CORE pre-scaled clock and random trigger generators. ATLAS Collaboration, 2007.

[35] Pedro Borrego Amaral. *CTPOUT User's Manual*. Tech. rep. EDA-00623. ATLAS Level-1 Central Trigger Processor. CERN, Aug. 2004. URL: https://edms.cern.ch/item/EDA-00623/0.

[36] S. Haas, M. Perantoni, and R. Spiwoks. *The Calibration Module of the Central Trigger Processor (CTP_CAL) - User Manual*. Tech. rep. ATL-DA-ER-0024. ATLAS Level-1 Central Trigger Processor. CERN, June 2007.

[37] Thilo Pauly et al. *ATLAS Level-1 Trigger Timing-In Strategies*. 2005. DOI: 10.5170/CERN-2005-011.274. URL: https://cds.cern.ch/record/898242.

[38] CERN, PH-ESE Group. *TTCrx Reference Manual, Version 3.10*. CERN Document Server: TTCrx Manual. CERN. Mar. 2005. URL: https://cds.cern.ch/record/837747.

[39] R. Jones et al. "The OKS Persistent In-memory Object Manager". In: *IEEE Transactions on Nuclear Science* 45.4 (1998), pp. 1958–1964. DOI: 10.1109/23.710894. URL: https://doi.org/10.1109/23.710894.

[40] Serguei Kolos. *Information Service User's Guide*. Version 1.4. User's Guide for the Information Service (IS) component of the ATLAS TDAQ Online Software. CERN. May 2011.

[41] G Avolio, M Caprini, and G Lehmann Miotto. "The integrated graphical user interface of the trigger and data acquisition system of the ATLAS experiment at the LHC". In: *Journal of Physics: Conference Series* 331.2 (Dec. 2011), p. 022032. DOI: 10.1088/1742-6596/331/2/022032. URL: https://dx.doi.org/10.1088/1742-6596/331/2/022032.

[42] Ralf Spiwoks. "Software Framework for the System-on-Chip of the ATLAS MUCTPI". In: *2nd System-on-Chip Workshop - CERN*. On behalf of the ATLAS Level-1 Central Trigger Team. Zoom, hosted by CERN, June 2021. URL: https://indico.cern.ch/event/1019078.

[43] Emily Anne Thompson. *ATLAS ITk Pixel Detector Overview*. Geneva: CERN, 2025. URL: https://cds.cern.ch/record/2928802.

[44] Laura Gonella and the ATLAS ITk Collaboration. "The ATLAS ITk detector system for the Phase-II LHC upgrade". In: *Nuclear Inst. and Methods in Physics Research, Section A* 1045 (2023), p. 167597. DOI: 10.1016/j.nima.2022.167597. URL: https://doi.org/10.1016/j.nima.2022.167597.

[45] Joaquim Pinol Bel and ATLAS Collaboration. "The High-Granularity Timing Detector for ATLAS at HL-LHC". In: *Journal Not Specified* 1 (2024), p. 0. DOI: 10.3390/1010000. URL: https://doi.org/10.3390/1010000.

[46] Teresa Barillari and ATLAS LAr Calorimeter Collaboration. "Development of the ATLAS Liquid Argon Calorimeter Readout Electronics for the HL-LHC". In: *42nd International Conference on High Energy Physics (ICHEP2024)* (2024). eprint: ATL-LARG-PROC-2024-002. URL: https://cds.cern.ch/record/2880042.

[47] Pavel Starovoitov and Tile Calorimeter System. "Upgrade of ATLAS Hadronic Tile Calorimeter for the High Luminosity LHC". In: *Preprints* 2022 (2022), pp. 1, 0. DOI: 10.3390/1010000. eprint: ATL-TILECAL-PROC-2022-010. URL: https://doi.org/10.3390/1010000.

[48] *Technical Design Report for the Phase-II Upgrade of the ATLAS Muon Spectrometer*. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2285580.

[49] F. Pastore and ATLAS Collaboration. *ATLAS TDAQ upgrades for Phase-2*. ATL-DAQ-PROC-2024-008. ATLAS Collaboration, CERN, 2024. URL: https://cds.cern.ch/record/2907275/files/ATL-DAQ-PROC-2024-008.pdf.

[50] N. Ellis et al. *Requirements and Specifications for the Phase-II Central Trigger Processor (CTP)*. Version 1.0. CERN, 2023. URL: https://edms.cern.ch/document/2787779.

[51] Thilo Pauly. Private communication, August 2025. 2025.

[52] Grégoire Hagmann. *LHC RF Signal Distribution*. Presentation slides, HPTD meeting #15. 29 January 2024. Jan. 2024.

[53] Philippa Hazell. *Phase Stability Testing of the White Rabbit based LHC RF Dependant Timing Distribution Backbone Upgrade*. Presentation slides, TWEPP 2024. Topical Workshop on Electronics for Particle Physics. Oct. 2024.

[54] Tomasz Włostowski. *White Rabbit-based LLRF upgrade for CERN's SPS*. Presentation slides. Brugg-Windisch, 11 October 2022. Oct. 2022.

[55] CERN White Rabbit Project. *White Rabbit – Sub-nanosecond timing over Ethernet*. Accessed: 2025-07-22. URL: https://white-rabbit.web.cern.ch/.

[56] N. Ellis et al. *Phase-II Local Trigger Interface (LTI) Specification*. Version 1.1. CERN, July 1, 2021. URL: https://edms.cern.ch/document/2379978.

[57] N. Ellis et al. *Phase-II Local Trigger Interface (LTI) User Requirements*. Version 1.1. CERN, July 1, 2021. URL: https://edms.cern.ch/document/2379976.

[58] ATLAS TDAQ Phase-II Upgrade Project. *Phase-II Local Trigger Interface (LTI) Hardware Implementation*. Version 1.0. Available at: Internal ATLAS documentation. CERN, Oct. 9, 2023.

[59] Samtec. *Firefly optical transceivers*. Accessed: 2025-07-22. URL: https://www.samtec.com/optics/systems/firefly/.

[60] CERN. *lpGBT Manual: lpGBT and lpGBT-FPGA Firmware Manual*. Version 1.5. CERN. Dec. 2023. URL: https://cds.cern.ch/record/2809058/files/lpGBT_manual.pdf.

[61] Eduardo Mendes, Sophie Baron, and Mark Taylor. "TCLink: A Timing Compensated High-Speed Optical Link for the HL-LHC experiments". In: *Proceedings of TWEPP 2020*. Mar. 2020, p. 057. DOI: 10.22323/1.370.0057.

[62] D. W. Allan and H. Daams. "Picosecond time difference measurement system". In: *29th Annual Symposium on Frequency Control*. Jan. 1975, pp. 404–411. URL: https://ui.adsabs.harvard.edu/abs/1975frco.symp..404A.

[63] Pedro Moreira et al. "Sub-nanosecond digital phase shifter for clock synchronization applications". In: *2012 IEEE International Frequency Control Symposium Proceedings*. 2012, pp. 1–6. DOI: 10.1109/FCS.2012.6243715.

[64] M. Rizzi et al. "White Rabbit Clock Characteristics". In: *Proceedings of the 2016 International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS)*. Stockholm, Sweden: IEEE, 2016, pp. 1–6.

[65] Xilinx / AMD. *UltraScale Architecture GTH Transceivers User Guide (UG576)*. Revision 1.7.1. Xilinx (now AMD). Aug. 2021. URL: https://docs.amd.com/v/u/en-US/ug576-ultrascale-gth-transceivers.

[66] Chris M. Lonvick and Tatu Ylonen. *The Secure Shell (SSH) Protocol Architecture*. RFC 4251. Jan. 2006. DOI: 10.17487/RFC4251. URL: https://www.rfc-editor.org/info/rfc4251.

[67] Keysight Technologies, Inc. *Keysight Infiniium MXR-Series Data Sheet (July 2022)*. Data Sheet MXR-Series Data Sheet ver. July 2022. Keysight Technologies, July 2022.

[68] Climats SAS. *Spirale3 User Manual EN – ind01*. France, Dec. 2019.

[69] CTS GmbH. *Operating Instructions: Temperature Test Cabinet Type T-40/500*. Lotzenäcker 21, D-72379 Hechingen, Germany, Nov. 2003.

[70] BINDER GmbH. *BINDER Interface Technical Specifications*. Post office box 102, 78502 Tuttlingen, Germany, Oct. 2024. URL: http://www.binder-world.com.

[71] *Resistance thermometer*. https://en.wikipedia.org/wiki/Resistance_thermometer. Accessed information on RTD types, construction, and performance. 2025. (Visited on 08/13/2025).

[72] Yoctopuce. *Yocto-PT100: Tiny USB interface for a Pt100 temperature probe*. https://www.yoctopuce.com/EN/products/usb-environmental-sensors/yocto-pt100. USB temperature sensor with high precision (supports 2/3/4-wire Pt100 probes). 2013–2025.

[73] Inc. Xilinx. *UltraScale Architecture System Monitor User Guide (UG580)*. v1.10.1. Document ID: UG580. Xilinx, Inc. San Jose, CA, USA, Sept. 2021. URL: https://www.xilinx.com/support/documentation/user_guides/ug580-ultrascale-sysmon.pdf.

[74] Wikipedia contributors. *Autocorrelation*. https://en.wikipedia.org/wiki/Autocorrelation. Accessed: 2025-08-10.

[75] Irmen de Jong. *Pyro5 Documentation*. Accessed: 2025-07-12. 2024. URL: https://pyro5.readthedocs.io/en/latest/.

[76] Wikipedia contributors. *Ordinary least squares — Wikipedia, The Free Encyclopedia*. 2025. URL: https://en.wikipedia.org/wiki/Ordinary_least_squares.

[77] Wikipedia contributors. *Huber loss — Wikipedia, The Free Encyclopedia*. 2025. URL: https://en.wikipedia.org/wiki/Huber_loss.

[78] Clément Camincher. *Highlights from Integration-week-09 @ CERN*. Presentation at LAr Weekly Meeting. Presented on behalf of the integration team. June 2025. URL: https://indico.cern.ch/event/1514079/#18-phase-2-news.

[79] Melvin Leguijt. *HGTD FELIX Expert Week – Clock Distribution Measurements*. Internal presentation. June 2025. URL: https://indico.cern.ch/event/1563087/.

[80] *Phase stability tests of the lpGBT*. VLDB+ Public Documentation. Available at: https://vldbplus.web.cern.ch/template/static/docs/Phase_stability_tests_of_the_lpGBT.pdf. 2024.