



The ATLAS Experiment at the CERN Large Hadron Collider: A Description of the Detector Configuration for Run 3

The ATLAS Collaboration

The ATLAS detector is installed in its experimental cavern at Point 1 of the CERN Large Hadron Collider. During Run 2 of the LHC, a luminosity of $\mathcal{L} = 2 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$ was routinely achieved at the start of fills, twice the design luminosity. For Run 3, accelerator improvements, notably luminosity levelling, allow sustained running at an instantaneous luminosity of $\mathcal{L} = 2 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$, with an average of up to 60 interactions per bunch crossing. The ATLAS detector has been upgraded to recover Run 1 single-lepton trigger thresholds while operating comfortably under Run 3 sustained pileup conditions. A fourth pixel layer 3.3 cm from the beam axis was added before Run 2 to improve vertex reconstruction and b -tagging performance. New Liquid Argon Calorimeter digital trigger electronics, with corresponding upgrades to the Trigger and Data Acquisition system, take advantage of a factor of 10 finer granularity to improve triggering on electrons, photons, taus, and hadronic signatures through increased pileup rejection. The inner muon endcap wheels were replaced by New Small Wheels with Micromegas and small-strip Thin Gap Chamber detectors, providing both precision tracking and Level-1 Muon trigger functionality. Tile Calorimeter scintillation counters were added to improve electron energy resolution and background rejection. Upgrades to Minimum Bias Trigger Scintillators and Forward Detectors improve luminosity monitoring and enable total proton-proton cross section, diffractive physics, and heavy ion measurements. These upgrades are all compatible with operation in the much harsher environment anticipated after the High-Luminosity upgrade of the LHC and are the first steps towards preparing ATLAS for the High-Luminosity upgrade of the LHC. This paper describes the Run 3 configuration of the ATLAS detector.

Contents

1 Overview	7
1.1 Brief history of the ATLAS detector performance and LHC roadmap	7
1.2 LHC Performance and Status	8
1.3 Brief overview of the ATLAS detector configuration for Run 3	11
1.3.1 Radiation and shielding	13
1.3.2 Tracking	13
1.3.3 Calorimetry	14
1.3.4 Muon system	15
1.3.5 Forward detectors	19
1.3.6 Trigger and Data Acquisition System	20
1.4 Physics and performance goals for Run 3	21
1.5 Outline of this Paper	23
2 Background Radiation and Shielding	25
2.1 LHC luminosity evolution and expectation	25
2.2 Radiation shielding	25
2.3 Characterisation of the radiation environment	27
2.4 Radiation monitoring	28
2.5 Radiation simulation codes and methods	29
2.6 Comparison of observed radiation levels with simulations	29
2.6.1 Inner detector	30
2.6.2 Calorimeter and muon regions	30
2.6.3 Simulation safety factors	31
2.7 Predicted Run 3 radiation levels	32
2.7.1 Radiation levels in the Inner detector	32
2.7.2 Radiation levels around the calorimeters	33
2.7.3 Radiation levels in the muon system and cavern	33
3 Inner Detector	38
3.1 Insertable <i>B</i> -layer	39
3.1.1 New beampipe	41
3.1.2 Sensors	41
3.1.3 Modules and staves	43
3.1.4 Integration and installation of the IBL in the ID	44
3.1.5 Readout of the IBL	45
3.1.6 Off-detector services and detector power supplies	46
3.1.7 Detector control and interlock systems	49
3.1.8 IBL CO ₂ cooling system	50
3.2 The ATLAS Pixel upgrade	52
3.2.1 Installation of new Service Quarter Panels	52
3.2.2 Readout upgrades during the winter shutdowns	52
3.3 Semiconductor Tracker readout upgrade	53
3.4 Transition Radiation Tracker upgrades	55
3.4.1 TRT DAQ upgrades in Run 2	55
3.4.2 TRT Gas Configuration Update	57

3.5	New ID thermosiphon system	57
3.6	Updated material description of the ID	58
3.7	Performance of the ID at the end of Run 2 and projections for Run 3	62
3.7.1	Pixel and IBL detector performance	62
3.7.2	SCT tracking performance	65
3.7.3	TRT tracking performance	66
4	Calorimeters	70
4.1	Liquid Argon Calorimeters	70
4.1.1	Modifications to the Original LAr System	70
4.1.2	New digital trigger path	73
4.1.2.1	New Baseplanes	73
4.1.2.2	Modifications of the FEB with new Layer Sum Boards	75
4.1.2.3	LAr Trigger Digitizer Board	75
4.1.2.4	New off-detector electronics	75
4.1.2.5	LAr Digital Processing Blades	76
4.2	Tile Calorimeter	77
4.2.1	Radiation environment	78
4.2.2	Cryostat counters	78
4.2.3	Minimum Bias Trigger Scintillator counters	81
4.2.4	TileCal Demonstrator for Run 4	82
5	Muon Spectrometer	84
5.1	Overview of Muon Spectrometer Upgrades	84
5.2	Endcap Upgrades: The New Small Wheels	88
5.2.1	Performance requirements for the NSW	89
5.2.2	Layout of the NSW	89
5.2.3	Mechanics and Shielding	91
5.2.4	Micromegas Technology	92
5.2.4.1	Micromegas Quadruplet and Wedge Structure	94
5.2.5	Small-Strip Thin Gap Chamber Technology	97
5.2.5.1	sTGC Quadruplet and Wedge Structure	98
5.2.6	Extending the Endcap Alignment System	100
5.2.7	NSW Electronics	101
5.2.7.1	CERN ASICs.	101
5.2.7.2	NSW Custom ASICs.	103
5.2.7.3	NSW Front-End Boards.	104
5.2.7.4	Trigger, Readout and Low voltage On-Detector Boards.	105
5.2.8	Trigger	107
5.2.8.1	sTGC Pad Trigger	107
5.2.8.2	sTGC Strip Trigger Inputs.	108
5.2.8.3	Micromegas Trigger Inputs.	108
5.2.8.4	NSW Trigger Processor.	109
5.2.9	Data Acquisition and Readout	109
5.2.10	Temperature and Magnetic Field Sensors	109
5.2.11	Services	109
5.2.11.1	Gas Distribution	110

5.2.11.2	Grounding	111
5.2.11.3	High Voltage Distribution	111
5.2.11.4	Low Voltage Distribution	111
5.2.11.5	Cooling Water	111
5.3	Upgrades to the Muon Spectrometer Barrel Coverage	112
5.3.1	Small-Diameter MDT Chambers	112
5.3.2	New Thin-Gap Resistive Plate Chambers	113
5.3.3	Improved Coverage in the Vicinity of Barrel Feet and Elevator Shafts	113
5.3.4	New BIS78 Muon Chambers	113
5.3.5	BIS78 Data Acquisition and Trigger System	117
6	Forward Detectors	118
6.1	LUCID Detector	118
6.1.1	LUCID 2 Detector Description	119
6.1.2	Prototypes for HL-LHC	121
6.1.3	JF detector	121
6.1.4	JN prototypes	122
6.1.5	Fibre detector	122
6.1.6	LUCID 2 Readout Electronics	123
6.1.7	LUCID absolute-luminosity calibration and luminosity-measurement methodology	124
6.2	ALFA Roman Pot Detector	127
6.3	AFP Detector	127
6.3.1	Silicon 3D tracker	129
6.3.2	AFP Time-of-Flight detector	129
6.3.3	AFP Trigger	129
6.3.4	AFP Data Acquisition	130
6.4	The Zero Degree Calorimeters	131
6.4.1	Overview of the ZDC upgrades	131
6.4.2	ZDC Readout and Trigger Electronics	133
6.4.3	ZDC Detector control	133
7	Trigger and Data Acquisition System	135
7.1	System Overview	135
7.1.1	Motivation for upgrades	137
7.1.2	Improvements for Run 2	138
7.2	Level-1 Calorimeter Trigger	139
7.2.1	Electron Feature Extractor (eFEX)	140
7.2.1.1	eFEX module design	140
7.2.1.2	eFEX algorithms	141
7.2.2	Jet Feature Extractor (jFEX)	143
7.2.2.1	jFEX module design	143
7.2.2.2	jFEX algorithms	144
7.2.3	Global Feature Extractor (gFEX)	146
7.2.3.1	gFEX module design	146
7.2.3.2	gFEX algorithms	146
7.2.4	L1Calo infrastructure	148
7.2.4.1	TREX	148

7.2.4.2	FOX	149
7.2.4.3	TopoFOX	149
7.2.4.4	Hub & ROD	150
7.3	L1 Muon Trigger	151
7.3.1	TGC EI-FI Coincidence	151
7.3.2	Tile-Muon Coincidence	151
7.3.3	NSW Trigger	153
7.3.3.1	NSW Trigger Processor requirements	153
7.3.3.2	NSW-TP mezzanine cards and input interfaces	154
7.3.3.3	NSW sTGC Strip Trigger Algorithm	155
7.3.3.4	Micromegas Trigger Algorithm	155
7.3.3.5	NSW-TP output interfaces	156
7.3.4	New Endcap Sector Logic	157
7.3.5	RPC Feet and Elevator Chambers	159
7.3.6	RPC BIS78	161
7.4	L1 Topological Trigger	161
7.5	Central Trigger & TTC	164
7.5.1	Muon to central trigger processor interface	164
7.5.1.1	Functional overview	164
7.5.1.2	System implementation	165
7.5.2	Central Trigger Processor	165
7.5.2.1	Functional overview	165
7.5.2.2	System implementation	167
7.6	DAQ/HLT	170
7.6.1	FELIX/SW ROD	170
7.6.1.1	FELIX & SW ROD Server Hardware	172
7.6.1.2	FELIX I/O Card	172
7.6.1.3	FELIX Firmware	172
7.6.1.4	FELIX Software	174
7.6.1.5	SW ROD Design	175
7.6.2	Network	175
7.6.3	HLT/AthenaMT	176
7.6.4	Online Software & Monitoring	180
7.7	Detector Control System	181
7.7.1	Standard Middleware OPC UA	182
7.7.2	Controls Software for GBT-SCA-based Front-End Electronics	183
7.7.2.1	GBT-SCA OPC UA Ecosystem	183
7.7.2.2	GBT-SCA Software Performance	185
7.7.3	Controls for ATCA Back-End Electronics	186
8	Outlook	188
	Acknowledgements	190
	Glossary	190

1 Overview

The Large Hadron Collider (LHC) at CERN and its detectors have ushered in a new era in particle physics, colliding protons at a centre-of-mass energy up to $\sqrt{s} = 13$ TeV and a peak instantaneous luminosity of $\mathcal{L} = 2.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The ATLAS detector, which was installed at the LHC Interaction Point (IP) 1 during the period 2000 – 2008, has performed extremely well during Run 1 and Run 2 of the LHC, recording 5, 21 and 147 fb^{-1} of proton-proton collision data at $\sqrt{s} = 7, 8$ and 13 TeV, respectively. The highlight of this period was the discovery of the Higgs boson at a mass of 125 GeV, which was announced by the ATLAS and CMS collaborations (each comprising nearly 3000 scientific authors) on July 4th, 2012 [1, 2]. The excellent performance of the ATLAS detector and of the LHC has enabled over 1000 publications exploring the nature of the Higgs boson, searching for new phenomena, and continuing to probe the Standard Model.

The ATLAS detector was designed for a peak instantaneous luminosity of $\mathcal{L} = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ at a proton-proton (pp) centre-of mass energy of $\sqrt{s} = 14$ TeV with about 25 interactions per bunch crossing, and of $\mathcal{L} = 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ of heavy ion (lead-lead: Pb+Pb, and proton-lead: p +Pb) collisions at 5.5 TeV per nucleon pair. The original configuration of the detector, as it was built for the start of the LHC, is described in Ref [3]. However, the LHC operating conditions differed from expectations; in particular, the instantaneous luminosity surpassed design, reaching a maximum value of $\mathcal{L} = 2.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ during Run 2. There were on average 33.7 interactions per bunch crossing during Run 2, with a peak value of over 60 interactions per bunch crossing recorded in 2017 and 2018. The ATLAS detector performed well despite these harsher conditions; performance and analysis techniques were adapted and improved to maintain the experiment’s excellent physics reach during Run 2. A substantial upgrade to the ATLAS detector, the “Phase-I Upgrade”, has consisted of improvements to the detector subsystems and their electronics in order to withstand the expected Run 3 conditions of $\langle \mu \rangle \approx 50$ (and at maximum 60) interactions per bunch crossing, and maintain the lowest achievable trigger thresholds, enabling the continued broad physics program planned for Run 3. The ATLAS detector configuration during Run 3 of the LHC is described in this paper.

1.1 Brief history of the ATLAS detector performance and LHC roadmap

The ATLAS detector has been in operation since autumn 2009. During this period, the LHC and ATLAS have alternated between distinct running periods and long shutdown periods:

Run 1, 2009 – 2013: ATLAS recorded 5 fb^{-1} of pp collision data at $\sqrt{s} = 7$ TeV and 21 fb^{-1} at $\sqrt{s} = 8$ TeV; $167 \text{ } \mu\text{b}^{-1}$ of Pb+Pb collisions were recorded in 2010 and 2011; 29.8 nb^{-1} of p +Pb collisions were recorded in 2013.

Long Shutdown 1 (LS1), 2013 – 2015: This shutdown was used to consolidate the LHC machine elements (repairing the magnet splices and upgrading the collimation scheme) to increase the centre-of-mass energy and reach the design luminosity. Upgrades to ATLAS installed and commissioned during LS1 are referred to as “Phase-0” Upgrades, and are described in this document; the most significant Phase-0 upgrade was the addition of the innermost layer of silicon pixel detectors described in Section 3.1.

Run 2, 2015 – 2018: ATLAS recorded 147 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV; 179.8 nb^{-1} of p +Pb collisions were recorded in 2016 and 1.76 nb^{-1} of Pb+Pb collisions were recorded in 2018.

The LHC reached a peak instantaneous luminosity of $\mathcal{L} = 2.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and, on average, delivered 33.7 interactions per bunch crossing.

Long Shutdown 2 (LS2), 2019 – 2022: During this shutdown, a new linear accelerator (LINAC-4) was connected into the injector complex [4], and the injection beam energy of the Proton Synchrotron (PS) Booster was upgraded in order to reduce the beam emittance. New cryogenics plants were installed to separate the cooling circuits of the superconducting radio frequency (RF) cavities from those of the superconducting magnets. The ATLAS Phase-I Upgrade was installed and commissioned during LS2 and its description constitutes the bulk of this document.

Run 3, 2022 – 2025: The LHC is expected to deliver a peak instantaneous luminosity of approximately $\mathcal{L} = 2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and an integrated luminosity of 250 fb^{-1} of proton-proton collision data at a centre-of-mass energy of $\sqrt{s} = 13.6 \text{ TeV}$.

Long Shutdown 3 (LS3), 2026 – 2028: The LHC will undergo a major upgrade of its components (e.g. low- β quadrupole triplets, crab cavities at the interaction regions). The ATLAS Phase-II Upgrade will be installed and commissioned.

High Luminosity Large Hadron Collider (HL-LHC), 2029 and beyond: The LHC complex is expected to deliver a levelled instantaneous luminosity of $\mathcal{L} = 5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and an annual integrated luminosity of approximately 250 fb^{-1} to reach a total dataset of 3000 fb^{-1} .

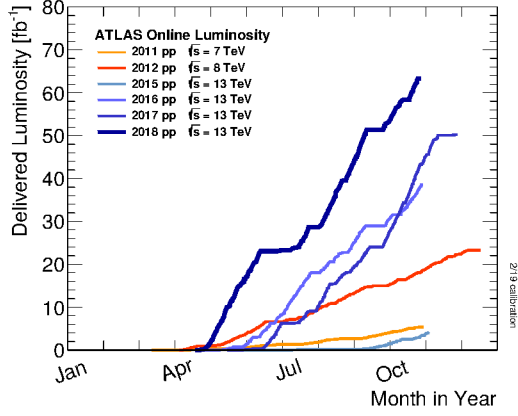
Figure 1.1 presents the distributions of the delivered luminosity for each production year between 2011 and 2018 and the distribution of the pile-up (μ , the number of pp collisions per LHC bunch crossing) for the four data-taking years of the LHC Run 2.

Over the ten years since the start of LHC, the ATLAS detector has operated reliably under the challenging conditions provided by the excellent performance of the LHC. Three indicators illustrate the detector performance: the fraction of operational channels, the data taking and the data quality efficiencies. The fraction of operational channels at the start of LS2 was above 99.5% for calorimeters and about 95% for tracking detectors. By the start of Run 3, some faulty channels had been repaired, while some detectors that had undergone extensive modifications were still being commissioned. The Resistive Plate Chambers (RPCs) required flushing with argon to remove debris that had accumulated during LS2, after which the operational fraction improved. The full details are shown in Table 1.1. ATLAS data-taking efficiency (the fraction of the time ATLAS collects data while LHC delivers collisions in stable conditions) improved from 90 to 95% in the course of the eight years of data taking. Over the same period, the fraction of collected data declared good for physics analysis has increased from 88.8 to 97.5% [5].

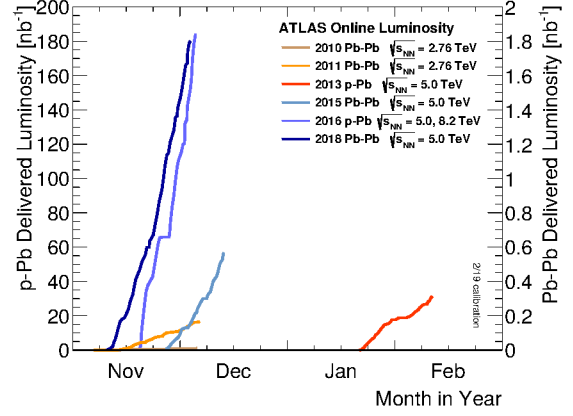
1.2 LHC Performance and Status

The LHC and its injector chain have undergone multiple major repairs, consolidations and upgrades since Run 1 as detailed in Ref. [4]. These changes facilitated the increase in collision energy and luminosity during Run 2 and are expected to provide further increases for Run 3 and beyond.


During LS1 the more than 10 000 high current splices between the LHC superconducting magnets were repaired and consolidated, 18 dipole magnets replaced and new safety systems added in order to safely increase the beam energy for Run 2. The LHC collimation system, injection kicker magnets and injection protection system were upgraded to handle long trains of bunches with 25 ns bunch spacing and higher



(a)



(b)



(c)

Figure 1.1: (a) Cumulative luminosity versus day delivered to ATLAS during stable beams for high energy pp collisions. (b) Cumulative luminosity versus day delivered to ATLAS during stable beams for $p+Pb$ and $Pb+Pb$ collisions (indicated on the left and right vertical axes, respectively). (c) Integrated luminosity as a function of the mean number of interactions per crossing, $\langle\mu\rangle$, for the 2015–2018 pp collision data at $\sqrt{s} = 13$ TeV. All data recorded by ATLAS during stable beams are shown, including the integrated luminosity and the mean μ value for each year. The mean number of interactions per bunch crossing, $\langle\mu\rangle$, corresponds to the calculated mean of the Poisson distribution of the number of inelastic interactions per bunch crossing.

Table 1.1: The number of channels vs. the approximate operational fraction for ATLAS subdetectors with more than 100 channels as of the end of Run 2 (March 2019) compared with the start of Run 3 (May 2022). The RPC operational fraction is expected to improve as commissioning continues, and compares favourably with the operational fraction at the start of Run 2. None of the 8704 RPC channels of the BIS78 were operational at the start of Run 3.

Subdetector	Run 2		Run 3	
	Number of Channels	Operational Fraction	Number of Channels	Operational Fraction
Pixel IBL	12M	99.3%	12M	96.7%
Pixel outer 3 layers	80M	94.8%	80M	96.7%
SCT	6.3M	98.6%	6.3M	98.3%
TRT	350k	97.2%	350k	96.6%
LAr Calorimeter (EM)	170k	100%	170k	100%
Tile Calorimeter	5200	99.5%	5200	99.2%
LAr Calorimeter (HEC)	5600	99.7%	5600	99.9%
LAr Calorimeter (FCal)	3500	99.8%	3500	99.8%
L1Calo Legacy Trigger	7160	99.9%	7160	99.9%
L1Calo Super Cell Trigger	not present	–	34k	100%
L1Muon RPC Trigger	383k	100%	392k	97.5%
L1Muon TGC Trigger	320k	99.9%	312k	100%
MDT	367k	99.7%	344k	99.7%
CSC	31k	93.0%	not present	–
RPC	383k	93.3%	392k	85.8%
TGC	320k	98.9%	312k	99.4%
sTGC	not present	–	358k	99.2%
Micromegas	not present	–	2.1M	98.0%
ALFA	10k	98.9%	10k	100%
AFP	430k	97.0%	430k	100%

bunch brightness. Additional improvements to the injection and beam dump systems were done in winter shutdowns during Run 2 to be able to further increase the beam intensity.

High-luminosity proton-proton collisions in Run 2 were delivered at $\sqrt{s} = 13$ TeV, while Pb-Pb collisions were delivered at $\sqrt{s} = 5.02$ TeV per nucleon-pair. Multiple improvements were deployed during the running period to increase the instantaneous luminosity a factor two beyond the increases brought by the higher beam energy and the larger number of bunches from shorter bunch spacing. As the beam optics and the understanding of collimation tolerances were improved, the interaction optics (β^*), which determines the transverse beam size at the interaction point, was gradually squeezed more and more. In 2016 brighter bunches were introduced through a change to the bunch-production scheme in the PS called Batch Compression Merging and Splitting which gave smaller transverse beam size [6]. In 2017 crossing angle anti-levelling was introduced, where the crossing angle in the IP was gradually reduced during each fill as bunch intensity decreased. In 2018 this was augmented with a β^* anti-levelling where a 15% additional squeeze at the IP was done towards the end of each fill. Both anti-levelling schemes increased the integrated luminosity without increasing the peak instantaneous luminosity, by running near the peak instantaneous luminosity for a much larger fraction of the total running time.

In LS2 a major consolidation campaign was carried out for the bypass diodes of the LHC superconducting

dipole magnets in order to more safely condition the magnets for highest beam energies. A large set of upgrades was carried out in the injector chain, including the connection of the new LINAC-4 to the injector complex and an increase of the PS Booster beam energy. This will allow very low emittance, high intensity bunches to be produced for the HL-LHC. The full intensity beams of the injector will only be usable after LS3, but upgrades and modifications during LS2 to the LHC collimators, injection kicker magnets and the beam dump will allow the LHC to use up to 60% more intense beams in Run 3. In the Super Proton Synchrotron (SPS), upgrades to the RF system will be used to produce ion beams with 50 ns bunch spacing using a technique called slip-stacking [7], thus enabling a larger beam intensity to be injected in the LHC for Pb-Pb collisions.

In Run 3, the proton-proton collision energy will increase to 13.6 TeV, but the peak instantaneous luminosity will remain limited to about $\mathcal{L} = 2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in ATLAS. The main limitation is luminosity-induced heating of the inner triplet magnets that provide the final focus before the beams reach the IP. The higher beam brightness from the LS2 injector upgrade will instead be used to provide extended periods of luminosity leveled at $\mathcal{L} = 2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ for as long as 10 hours per LHC fill. The corresponding peak pileup level is expected to be between 52 and 57 collisions per crossing during the luminosity levelling period. This is expected to give a yearly integrated luminosity of more than 80 fb^{-1} in the latter part of Run 3 and potentially more than 250 fb^{-1} for the full run. For ions, more than 6 nb^{-1} of Pb-Pb collisions is expected in Run 3.

1.3 Brief overview of the ATLAS detector configuration for Run 3

The ATLAS detector [3] at the LHC covers nearly the entire solid angle around the collision point¹.

It consists of an inner tracking detector surrounded by a thin superconducting solenoid, electromagnetic and hadronic calorimeters, and a muon spectrometer incorporating three large superconducting air-core toroidal magnets. The dimensions of the detector are 25 m in height and 44 m in length; the overall weight of the detector is approximately 7000 t.

The Inner Detector (ID) system is immersed in a 2 T axial magnetic field and provides charged-particle tracking in the range $|\eta| < 2.5$. The high-granularity silicon Pixel detector covers the vertex region and typically provides four measurements per track, the first hit normally being in the Insertable *B*-Layer (IBL) installed before Run 2 [8, 9]. It is followed by the Semiconductor Tracker (SCT), a silicon microstrip tracker that usually provides eight measurements per track. These silicon detectors are complemented by the Transition Radiation Tracker (TRT), which enables radially extended track reconstruction up to $|\eta| = 2.0$. The TRT also provides electron identification information based on the fraction of hits above a higher energy-deposit threshold corresponding to transition radiation.

The calorimeter system covers the pseudorapidity range $|\eta| < 4.9$. Within the region $|\eta| < 3.2$, electromagnetic calorimetry is provided by barrel and endcap high-granularity lead/Liquid Argon (LAr) calorimeters, with an additional thin LAr presampler covering $|\eta| < 1.8$ to correct for energy loss in material upstream of the calorimeters. Hadron calorimetry is provided by the steel/scintillator-tile calorimeter, segmented into three barrel structures within $|\eta| < 1.7$, and two copper/LAr hadron endcap calorimeters.

¹ ATLAS uses a right-handed coordinate system with its origin at the nominal IP in the centre of the detector and the *z*-axis along the beampipe. The *x*-axis points from the IP to the centre of the LHC ring, and the *y*-axis points upward. Cylindrical coordinates (*r*, ϕ) are used in the transverse plane, ϕ being the azimuthal angle around the *z*-axis. The pseudorapidity is defined in terms of the polar angle θ as $\eta = -\ln \tan(\theta/2)$. The two ends of the detector are labelled A for +*z* and C for -*z*, with B used for elements at $\eta = 0$.




Figure 1.2: Cut-away view of the Run 3 configuration of the ATLAS detector indicating the locations of the larger detector sub-systems.

The solid angle coverage is completed with forward copper/LAr and tungsten/LAr calorimeter modules optimised for electromagnetic and hadronic energy measurements respectively.

The Muon Spectrometer (MS) comprises separate trigger and high-precision tracking chambers measuring the deflection of muons in a magnetic field generated by the superconducting air-core toroidal magnets. The field integral of the toroids ranges between 2.0 and 6.0 Tm across most of the detector. Three stations of precision chambers, each consisting of layers of Monitored Drift Tubes (MDTs), cover the region $|\eta| < 2.7$, except in the innermost station of the endcaps, in the range $|\eta| > 1.3$, where the New Small Wheel (NSW) detectors, described below, have replaced the detectors used in Runs 1 and 2. The muon trigger system covers the range $|\eta| < 2.4$ with RPCs in the barrel $|\eta| < 1.0$, Thin Gap Chambers (TGCs) in the endcap $|\eta| > 1.0$ regions outside the toroids, and the NSWs between the cryostats of the endcap calorimeters and the endcap toroids.

Interesting events are selected by the first-level trigger system implemented in custom hardware, followed by selections made by algorithms implemented in software in the high-level trigger [10]. The first-level trigger accepts events from the 40 MHz bunch crossings at a rate below 100 kHz, which the high-level trigger further reduces in order to record events to disk at about 3 kHz.

An extensive software suite [11] is used in the reconstruction and analysis of real and simulated data, in detector operations, and in the trigger and data acquisition systems of the experiment.

Figure 1.2 illustrates the Run 3 configuration of the ATLAS detector; the main modifications implemented to the detector, its electronics, and the trigger and data acquisition system as well as their expected impact on the radiation levels are summarised here and will be described in detail in the following sections of this document.

1.3.1 Radiation and shielding

At the LHC, the primary source of radiation at full luminosity comes from collisions at the IP. In the ID, charged hadron secondaries from inelastic proton-proton interactions dominate the radiation backgrounds at small radii while at larger radii other sources, such as neutrons, become more important for the calorimeters and the MS. The expected background radiation levels for Run 3 are presented in Section 2, based on measurements from the LHC Run 1 and Run 2 and the modified and improved detector geometry.

In ATLAS, most of the energy from collisions at the IP is dumped into two regions: the collimators protecting LHC quadrupoles and the LAr Forward Calorimeter (FCal)s. The beam vacuum system spans the length of the detector, and in the forward region it is a major source of radiation backgrounds. Primary particles from the IP strike the beampipe at very shallow angles, such that the projected material depth is large. Studies have shown that the beam-line material contributes more than half of the radiation backgrounds in the muon system.

Details of the predicted Run 3 radiation levels and methods for monitoring and simulating the radiation in the ATLAS cavern are described in Section 2.

1.3.2 Tracking

The layout of the ID is illustrated in Figure 1.3 and detailed in Chapter 3. Its basic parameters are summarised in Table 1.2. The ID is immersed in a 2 T magnetic field generated by the central solenoid, which extends over a length of 5.3 m with a diameter of 2.5 m.

The precision tracking detectors (Pixel and SCT) cover the region $|\eta| < 2.5$. In the barrel region, they are arranged on concentric cylinders around the beam axis while in the endcap regions they are located on discs perpendicular to the beam axis. The highest granularity is achieved around the vertex region using silicon pixel detectors. The ATLAS Pixel detector consists of three barrel layers and three discs on each side, and has approximately 80 million readout channels. The pixel layers are segmented in $r\phi$ and z with typically three pixel layers crossed by each track. For these three outer layers, all pixel sensors are identical, with a pixel size of $50 \times 400 \mu\text{m}$. The intrinsic accuracies in the barrel are $10 \mu\text{m}$ ($r\phi$) and $115 \mu\text{m}$ (z) and in the discs are $10 \mu\text{m}$ ($r\phi$) and $115 \mu\text{m}$ (r). A fourth inner layer, the Insertable b -layer or IBL, was installed during LS1 and started to be operational at the start of Run 2 data taking. The IBL sensors have $50 \times 250 \mu\text{m}$ pixels and are at an average radius of 33.4 mm, adding an additional 12 million readout channels to the system (for a total of 92 million pixel channels). For the SCT, eight strip layers (four space points) are crossed by each track. In the barrel region, this detector uses small-angle (40 mrad) stereo strips to measure both $r\phi$ and z , with one set of strips in each layer parallel to the beam direction, measuring $r\phi$. They consist of two 6.4 cm long daisy-chained sensors with a strip pitch of $80 \mu\text{m}$. In the endcap region, the detectors have a set of strips running radially and a set of stereo strips at an angle of 40 mrad. The mean pitch of the strips is also approximately $80 \mu\text{m}$. The intrinsic accuracies per module in the barrel are $17 \mu\text{m}$ ($r\phi$) and $580 \mu\text{m}$ (z) and in the discs are $17 \mu\text{m}$ ($r\phi$) and $580 \mu\text{m}$ (r). There are approximately 6.3 million readout channels in the SCT.

The TRT is the outermost of the three tracking subsystems of the ID, and comprises several layers of gas-filled straw tubes interleaved with transition radiation material. The 300 000 thin-walled proportional-mode drift tubes provide on average 30 (r, ϕ) points with $130 \mu\text{m}$ resolution for charged particle tracks with $|\eta| < 2$ and $p_T > 0.5 \text{ GeV}$, contributing to the combined tracking system p_T resolution. Along with




Figure 1.3: Cut-away view of the ATLAS ID, which is designed to provide a high-precision reconstruction of charged-particle trajectories. The ID covers the pseudorapidity range $|\eta| < 2.5$ and has full coverage in ϕ . It consists of a silicon Pixel detector at the innermost radii surrounded by a silicon microstrip detector (SCT) and a straw-tube detector, the TRT, that combines continuous tracking capabilities with particle identification based on transition radiation.

continuous tracking, the TRT provides electron identification capability through the detection of transition radiation X-ray photons.

The updated ID material distribution as a function of pseudorapidity is presented in Figure 1.4. Compared with the equivalent distribution for the Run 1 detector, there is considerably more material at $|\eta| > 3.5$, due to the additional services for the IBL, but almost no change in the region covered by tracking.

1.3.3 Calorimetry

The ATLAS calorimeters measure the energies and positions of charged and neutral electromagnetically or strongly interacting particles. They are designed to absorb most of the particles coming from a collision, forcing them to deposit all of their energy and stop within the detector. The ATLAS calorimeters are sampling calorimeters, in that they consist of “absorbing” high-density materials that stop incoming particles, interleaved with layers of “active” media that measure the particle energies.

ATLAS uses two sampling calorimeter technologies: LAr [12] for the electromagnetic calorimeters and all of the endcap and forward calorimeters, and scintillating Tiles [13] for hadron calorimetry in the central region. The calorimeters are highlighted in Figure 1.5.

The LAr Calorimeter system consists of several subsystems, namely the LAr Electromagnetic Barrel Calorimeter (EMB), the LAr Electromagnetic Endcap Calorimeter (EMEC), the LAr Hadronic Endcap

Table 1.2: Main parameters of the Inner Detector system.

Item		Radial extension (mm)	Length (mm)
Overall ID envelope		$0 < r < 1150$	$0 < z < 3512$
Beampipe		$23.5 < r < 30$	
Pixel (IBL included)	Overall envelope	$31 < r < 242$	$0 < z < 3092$
4 cylindrical layers	Sensitive barrel	$33.5 < r < 122.5$	$0 < z < 400.5$
2 × 3 discs	Sensitive endcap	$88.8 < r < 149.6$	$495 < z < 650$
SCT	Overall envelope	$255 < r < 549$ (barrel)	$0 < z < 805$
		$251 < r < 610$ (endcap)	$810 < z < 2797$
4 cylindrical layers	Sensitive barrel	$299 < r < 514$	$0 < z < 749$
2 × 9 discs	Sensitive endcap	$275 < r < 560$	$839 < z < 2735$
TRT	Overall envelope	$554 < r < 1082$ (barrel)	$0 < z < 780$
		$617 < r < 1106$ (endcap)	$827 < z < 2744$
73 straw planes	Sensitive barrel	$563 < r < 1066$	$0 < z < 712$
160 straw planes	Sensitive endcap	$644 < r < 1004$	$848 < z < 2710$

Calorimeter (HEC), and the FCal. The ATLAS Tile Calorimeter covers the pseudorapidity region $|\eta| < 1.7$ using pseudo-projective calorimeter towers composed of scintillating tiles in a steel matrix, read out by wavelength-shifting fibers. The pseudorapidity coverage, granularity, and longitudinal segmentation of the calorimeters are summarised in Table 1.3. These calorimeters cover the range $|\eta| < 4.9$, using different geometries and absorber materials suited to the widely varying requirements of the physics processes of interest and of the radiation environment over this large η -range. Over the η region for which the inner detector provides tracking, the fine granularity of the electromagnetic (EM) calorimeter is ideally suited for precision measurements of electrons and photons. The coarser granularity of the rest of the calorimeter is sufficient to satisfy the physics requirements for jet reconstruction and missing transverse momentum (E_T^{miss}) measurements.

The ATLAS calorimeter detectors [14] require very few changes to run at luminosities substantially higher than the original design, and are expected to last the entire lifetime of the LHC and HL-LHC. The main calorimeter upgrades were performed on the LAr calorimeter electronics. A new digital trigger path provides finer granularity inputs to the upgraded trigger system and aims to better control the trigger rates by improving the selectivity of electron, photon, and tau lepton objects, the resolution of jets and E_T^{miss} trigger signatures, and the discrimination power against background emerging from both out-of-time and in-time pileup. These upgrades as well as improvements to Tile calorimeter cryostat scintillation counters and Minimum Bias Trigger Scintillators are described in Chapter 4.

1.3.4 Muon system

The muon spectrometer forms the large outer part of the ATLAS detector and detects charged particles exiting the barrel and endcap calorimeters, measuring their momentum in the pseudorapidity range

Table 1.3: Main parameters of the calorimeter system.

	Barrel		Endcap	
EM calorimeter				
Number of layers and $ \eta $ coverage				
Presampler	1	$ \eta < 1.52$	1	$1.5 < \eta < 1.8$
Calorimeter	3	$ \eta < 1.35$	2	$1.375 < \eta < 1.5$
	2	$1.35 < \eta < 1.475$	3	$1.5 < \eta < 2.5$
			2	$2.5 < \eta < 3.2$
Granularity $\Delta\eta \times \Delta\phi$ versus $ \eta $				
Presampler	0.025×0.1	$ \eta < 1.52$	0.025×0.1	$1.5 < \eta < 1.8$
Calorimeter 1st layer	$0.025/8 \times 0.1$	$ \eta < 1.40$	0.050×0.1	$1.375 < \eta < 1.425$
	0.025×0.025	$1.40 < \eta < 1.475$	0.025×0.1	$1.425 < \eta < 1.5$
			$0.025/8 \times 0.1$	$1.5 < \eta < 1.8$
			$0.025/6 \times 0.1$	$1.8 < \eta < 2.0$
			$0.025/4 \times 0.1$	$2.0 < \eta < 2.4$
			0.025×0.1	$2.4 < \eta < 2.5$
			0.1×0.1	$2.5 < \eta < 3.2$
Calorimeter 2nd layer	0.025×0.025	$ \eta < 1.40$	0.050×0.025	$1.375 < \eta < 1.425$
	0.075×0.025	$1.40 < \eta < 1.475$	0.025×0.025	$1.425 < \eta < 2.5$
			0.1×0.1	$2.5 < \eta < 3.2$
Calorimeter 3rd layer	0.050×0.025	$ \eta < 1.35$	0.050×0.025	$1.5 < \eta < 2.5$
Number of readout channels				
Presampler	7808		1536 (both sides)	
Calorimeter	101 760		62 208 (both sides)	
LAr hadronic endcap				
$ \eta $ coverage			$1.5 < \eta < 3.2$	
Number of layers			4	
Granularity $\Delta\eta \times \Delta\phi$			0.1×0.1	$1.5 < \eta < 2.5$
			0.2×0.2	$2.5 < \eta < 3.2$
Readout channels			5632 (both sides)	
LAr forward calorimeter				
$ \eta $ coverage			$3.1 < \eta < 4.9$	
Number of layers			3	
Granularity $\Delta x \times \Delta y$ (cm)			FCal1: 3.0×2.6	$3.15 < \eta < 4.30$
			FCal1: ~ four times finer	$3.10 < \eta < 3.15,$ $4.30 < \eta < 4.83$
			FCal2: 3.3×4.2	$3.24 < \eta < 4.50$
			FCal2: ~ four times finer	$3.20 < \eta < 3.24,$ $4.50 < \eta < 4.81$
			FCal3: 5.4×4.7	$3.32 < \eta < 4.60$
			FCal3: ~ four times finer	$3.29 < \eta < 3.32,$ $4.60 < \eta < 4.75$
Readout channels			3524 (both sides)	
Scintillator tile calorimeter				
	Barrel		Extended barrel	
$ \eta $ coverage	$ \eta < 1.0$		$0.8 < \eta < 1.7$	
Number of layers	3		3	
Granularity $\Delta\eta \times \Delta\phi$	0.1×0.1		0.1×0.1	
	Last layer 0.2×0.1		0.2×0.1	
Readout channels	5760		4092 (both sides)	

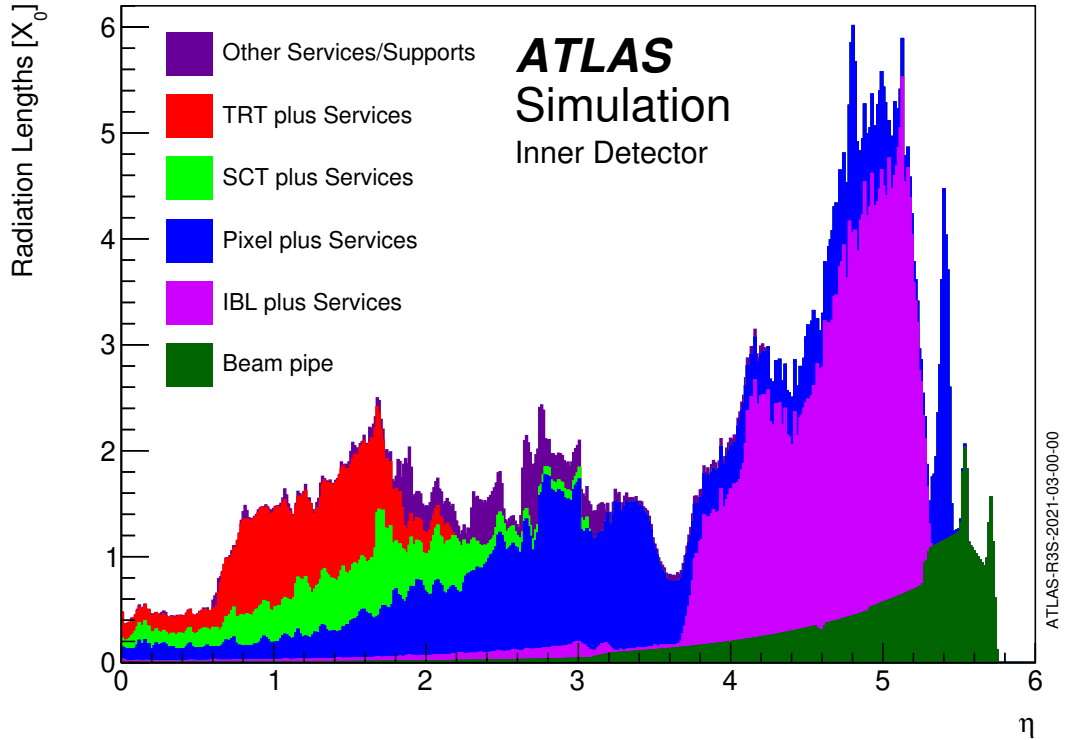


Figure 1.4: Radiation length (X_0) as a function of $|\eta|$ (averaged over ϕ) for the beampipe and different ID components as implemented in the ATLAS geometry model describing the Run 2 and Run 3 configuration. “Services” refers to supports, cooling infrastructure and cabling.

$|\eta| < 2.7$; the layout of the muon spectrometer is shown in Figure 1.6. The muon system is based on the magnetic deflection of muon tracks in the large superconducting air-core toroid magnets, instrumented with separate trigger and high-precision tracking chambers. Over the range $|\eta| < 1.4$, magnetic bending is provided by the large barrel toroid. For $1.6 < |\eta| < 2.7$, muon tracks are bent by two smaller endcap toroid magnets inserted into both ends of the barrel toroid. In between these two regions, $1.4 < |\eta| < 1.6$, magnetic deflection is provided by a combination of barrel and endcap fields. This magnet configuration provides a field which is mostly orthogonal to the muon trajectories, while minimising the degradation of resolution due to multiple scattering. The muon detector is included in the trigger system in the region $|\eta| < 2.4$.

The Muon Spectrometer (MS) comprises a “barrel”, consisting of three concentric, roughly cylindrical, stations (the inner, middle and outer barrels), and two endcaps, each consisting of three discs, referred to as the inner, middle, and outer endcap “wheels”, supplemented by an “extended” endcap ring of detectors positioned outside the radius of each endcap toroid cryostat. The anticipated high level of particle flux has had a major impact on the choice and design of the spectrometer instrumentation and its upgrade, affecting performance parameters such as rate capability, granularity, ageing properties, and radiation hardness.

The majority of the barrel detectors are unchanged from the original Run 1 configuration described in Ref. [15]: all three stations use multilayered MDT chambers for the precision measurements in the bending coordinate, and the outer and middle stations are also equipped with RPCs for triggering and to measure




Figure 1.5: Cut-away view of the ATLAS calorimeter system that measures the energies and positions of charged and neutral particles through interleaved absorber and active layers out to $|\eta| < 4.9$. LAr is used as the active medium for the electromagnetic calorimeters and all of the endcap and forward calorimeters; scintillating Tiles are used for hadron calorimetry in the central region.

the azimuthal coordinate of the tracks. The middle and outer wheels are unchanged from Run 1. The middle wheels are located on the far side of the endcap toroid cryostats and the outer wheels are mounted on scaffolding attached to the end-walls of the ATLAS cavern. The “extended” endcap rings provide a third measurement station between the inner and middle wheels for tracks with $1.05 < |\eta| < 1.3$, which are outside the acceptance of the outer wheel. The outer wheels contain only MDTs, while the middle wheels have both MDTs for precision tracking in the bending coordinate, and TGCs for triggering and for measuring the azimuthal coordinate. The inner wheels (often nicknamed “small wheels”), which are the main focus of the Phase-I upgrade, sit between the calorimeters and the endcap toroid cryostats, inside the barrel toroids. These inner wheels have been completely replaced by NSWs occupying the same position and providing tracking over the same polar angle range: $1.3 < |\eta| < 2.7$. The NSWs use two chamber technologies: small-strip TGCs (sTGCs) and micro-mesh gaseous structure (Micromegas) detectors, both fast enough for Level-1 trigger functionality, and both designed for precision tracking in the bending direction, as well as improved resolution in the azimuthal coordinate. These NSWs allow for improved p_T resolution in the trigger and increased background rejection, allowing for a low muon p_T threshold and manageable Level-1 trigger rate, thus maintaining the acceptance for many interesting physics processes.

The main parameters of the muon chambers are listed in Table 1.4; the upgrade of the muon system for Run 3 is described in detail in Chapter 5. The asymmetry of the RPC coverage is due to the installation of BIS78 detectors on only the $\eta > 0$ side of ATLAS for Run 3.




Figure 1.6: Cut-away view of the ATLAS muon system that detects charged particles exiting the barrel and endcap calorimeters and measures their momentum in the pseudorapidity range $|\eta| < 2.7$. The muon system is built around an air-core toroidal magnet system. The main muon detector components upgraded for Run 3 (e.g., the NSWs) are visible.

1.3.5 Forward detectors

Four smaller detector systems cover the ATLAS forward region (see Chapter 6). At 17 m from the IP lies LUCID (LUMinosity Cherenkov Integrating Detector), which detects inelastic proton-proton scattering in the forward direction, and is the main online and offline luminosity monitor for ATLAS. The second detector is Absolute Luminosity for ATLAS (ALFA) Roman Pot detector. Located at 240 m, it consists of scintillating fibre trackers housed inside Roman pots which are designed to approach as close as 1 mm to the beam. ALFA is used in dedicated low luminosity and high β^* running of the LHC and can provide a measurement of the total cross section for proton-proton interactions. The third system is the Zero Degree Calorimeters (ZDC), which plays a key role in determining the centrality of heavy-ion collisions. It is located at 140 m from the IP, just beyond the point where the common straight-section vacuum-pipe divides back into two independent beampipes. The ZDC modules consist of layers of alternating quartz rods and tungsten plates which measure neutral particles at pseudorapidities $|\eta| \geq 8.2$. The fourth system, which was installed in 2015 and 2016 is the ATLAS Forward Proton (AFP) detector, which consists of two arms, each with two stations, that are 210 m from the IP. The AFP is designed to study soft, hard, and central (exclusive) diffractive events at low luminosities using a silicon-based tracker for momentum measurements and a time-of-flight system to match protons from the two arms to a single inner detector vertex, thereby reducing the background from multiple proton-proton collisions.

Table 1.4: Main parameters of the muon spectrometer.

MDT	
η coverage	$ \eta < 2.7$ (innermost layer: $ \eta < 1.3$)
Number of modules	1098
Number of channels	355k
Function	Precision tracking
sTGC	
η coverage	$1.3 < \eta < 2.7$ (2.4 for trigger)
Number of quadruplets	192
Number of gas volumes	768
Number of channels	357k
Function	Trigger, precision tracking, 2nd coordinate
Micromegas	
η coverage	$1.3 < \eta < 2.7$ (2.4 for trigger)
Number of quadruplets	128
Number of gas volumes	512
Number of channels	2.05M
Function	Precision tracking, trigger, 2nd coordinate
RPC	
η coverage	$-1.05 < \eta < 1.3$
Number of modules	652
Number of channels	389k
Function	Trigger, 2nd coordinate
TGC	
η coverage	$1.05 < \eta < 2.7$ (2.4 for trigger)
Number of modules	1530
Number of gas volumes	3492
Number of channels	312k
Function	Trigger, 2nd coordinate

1.3.6 Trigger and Data Acquisition System

The ATLAS Trigger and Data Acquisition (TDAQ) system selects events with distinguishing characteristics (such as the presence of energetic leptons, photons, hadronic jets, τ leptons, or large missing energy) that make them interesting for physics analyses, and reads them out for further offline processing. It is based on a two-level event selection system: the Level-1 Trigger (L1) system, which consists of custom-built electronics, and the High-Level Trigger (HLT), which is a software-based system implemented on commercial computers. Interwoven with these levels is the Data Acquisition (DAQ) system, which transports data from custom subdetector electronics through to offline processing, according to the decisions made by the trigger. A diagram of the complete TDAQ system in Run 3 is shown in Figure 1.7.

The L1 trigger uses reduced-granularity information from the calorimeters and muon system to search for signatures of these events. The maximum L1 accept rate supported by the detector readout systems is 100 kHz, and a share of this rate budget is allocated to each underlying physics object according to the

physics goals of ATLAS. All processing for an event must be completed within the time window (latency) permitted by the detector electronics. This latency is 2.5 μs per event.

The HLT software is designed to reproduce the offline selection as closely as possible, a philosophy that will be taken one step further in Run 3 through the use of AthenaMT (see additional details in Section 7.6.3). On average, the event processing time at the HLT in 2018 was approximately 400 ms for runs with a peak luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The mean HLT event processing time in Run 3 is expected to be larger due to widespread use of full-detector HLT track reconstruction for hadronic signatures. Upgrades to the HLT farm, and algorithmic improvements to tracking will keep the HLT resource usage within the limits of the deployed system. The HLT reduces the event rate from 100 kHz after the L1 selection to approximately 3 kHz (averaged over the course of an LHC fill), after which the events are stored for offline analysis.

The DAQ system has also undergone an upgrade in order to read out the full detector during Run 3. The ReadOut System (ROS) must support the L1 accept rate of 100 kHz plus 20% contingency. The throughput of the ROS will increase in Run 3 due to the increased request rate at which the HLT will request data and a 30% increase in the average event size (the event size at $\langle\mu\rangle \approx 60$ is 2.1 MB). The system must accommodate an average rate of 3 kHz of physics events to mass storage with the flexibility to handle variations in such rate. The maximum throughput is 8 GB/s, which offers sufficient margin for these variations in rate. For example, at $\langle\mu\rangle \approx 60$, assuming a maximum physics rate to disk of 3.3 kHz and the above event size, the system throughput is 6.9 GB/s. This represents nearly a factor of two increase in the required performance compared to Run 2. In order to store the raw data volume for continuous operation at 3 kHz average output, guaranteeing 24 hours of storage in case of downtime of the CERN mass storage services, the contribution of file transfer and deletion latencies, and the margin needed to guarantee the required file system throughput characteristics, at least 1.4 PB of effective storage volume is available at Point 1.

1.4 Physics and performance goals for Run 3

The ATLAS physics goals for Run 3 will build on the successful discovery of the Higgs boson and take advantage of the 250 fb^{-1} proton-proton collisions planned by the LHC. The exploration of the mechanism of electroweak symmetry breaking includes an exploration of the phenomenology of the Higgs boson, precisely measuring its mass and width, the couplings to both fermions and bosons, and the observation of rare decay modes. Furthermore, ATLAS will continue to exploit the unique access to the energy frontier offered by the LHC through the study of rare Standard Model processes, flavour physics, and searches for new phenomena such as supersymmetry (SUSY) and exotic BSM scenarios that may reveal the nature of dark matter.

The Phase-I upgrades of the ATLAS detector have focused on the following improvements to the detector and trigger system, which are required to remain effective not only for Run 3, but throughout the lifetime of ATLAS. The first overall objective is the preservation, and in some cases improvement, of the low transverse-momentum electron and muon trigger thresholds that enabled the successful Run 1 and Run 2 physics program through the collection of a rich dataset of electroweak boson (H , W , and Z) decays. The second objective is to maintain sensitivity to electroweak-scale particles that produce hadronically-decaying tau leptons, jets and missing transverse momentum, the magnitude of which is referred to as $E_{\text{T}}^{\text{miss}}$. These objectives are accomplished through the detector upgrades described in this paper: calorimeter electronics upgrades that provide finer granularity and higher energy resolution to the trigger system, a new endcap muon detector that can tolerate the high background radiation environment that will be present in the

HL-LHC era, and TDAQ system upgrades that take advantage of the calorimeter and muon upgrades while rejecting background in a high pileup environment.

In addition to the Phase-I upgrades, the physics performance of the ATLAS detector in Run 3 critically depends on its ability to measure charged particle tracks in order to reconstruct primary and secondary vertices; this performance was significantly enhanced by the installation of the IBL described above.

The techniques for the measurement of the instantaneous and integrated luminosities developed in Run 2 will be refined further in Run 3 to ensure a continuously high-performant luminosity measurement. Run 3 provides, moreover, a testing ground for any new luminometer envisaged for the HL-LHC era to gain experience at moderate number of simultaneous particle interactions, before that number increases further for the HL-LHC. The interplay of different independent luminometers, carefully calibrated and together covering the full range in the number of simultaneous interactions over up to four orders of magnitude, is of key importance to achieve the targeted precision in the integrated luminosity in Run 3 of 1% or below.

The final Run 3 ATLAS detector configuration will enable the broad physics goals of the experiment described above through the enhanced selection of the signatures relevant to the ATLAS physics programme, including electrons, photons, muons, τ -leptons, jets, b -jets, B mesons, and E_T^{miss} . The trigger menu developed for Run 3 translates the physics priorities of the experiment into allocations of the total L1 and

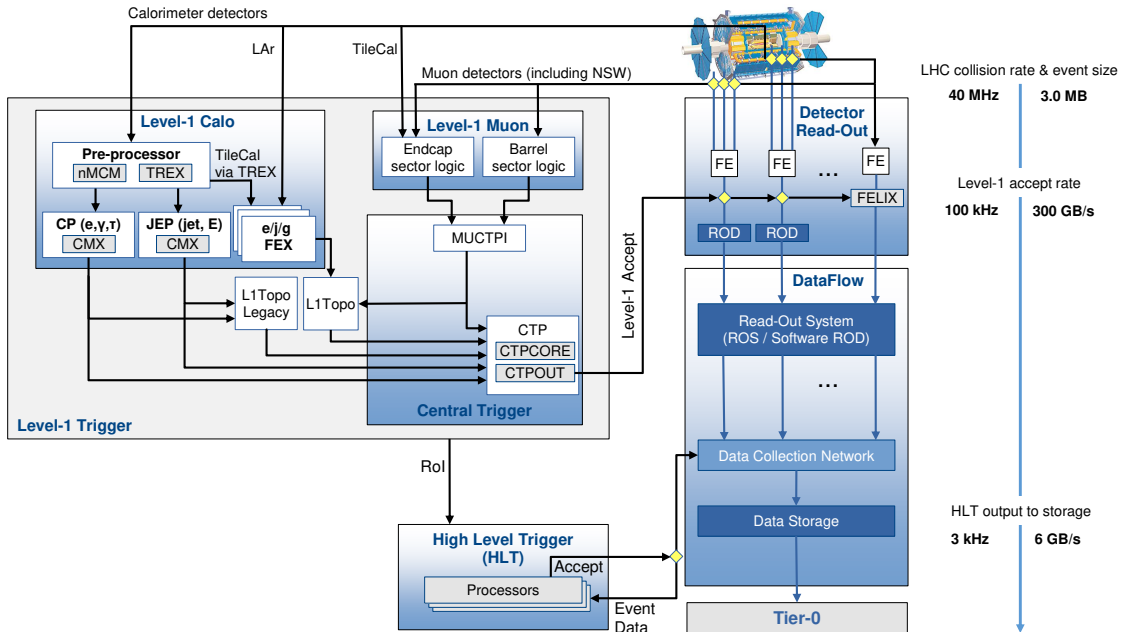


Figure 1.7: Schematic overview of the Trigger and DAQ system in Run 3.

HLT rates. At $\mathcal{L} = 2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ these rates are about 95 kHz and 3 kHz, respectively². The trigger menu comprises a list of trigger chains used for data-taking, where a chain consists of a L1 trigger item and a series of HLT algorithms that reconstruct physics objects and apply kinematic selections to them. Roughly equal shares of the overall rate are given to electron and muon trigger chains, with a large share of bandwidth reserved for jets, E_T^{miss} , taus, and multi-object triggers; rates for single electrons and muons are each limited to approximately 25 kHz or less. Improvements in the trigger performance allow for lower single-lepton p_T thresholds, yielding increased physics acceptance for a given trigger rate.


The expected single-lepton trigger performance is demonstrated in Figure 1.8. Figure 1.8(a) demonstrates the expected performance of the Run 3 single-electron trigger using $Z \rightarrow ee$ Monte Carlo simulation: a lower rate and improved efficiency is achieved compared to the Run 2 electron trigger. A threshold of 22 GeV is used for the Run 2 (black) and uncalibrated Run 3 (red) electron and photon trigger. The Run 3 isolation thresholds were tuned to give the lowest rate while introducing only a 2% inefficiency for electrons passing the L1 energy threshold; this isolation requirement is not applied for clusters with $E_T > 50(60)$ GeV in the Run 2 (Run 3) trigger. A layer- and η -dependent calibration is introduced (blue) to compensate for the varying detector response. The threshold on calibrated cluster energy is chosen to produce the same rate as the uncalibrated trigger, resulting in an improved efficiency. For muon triggers, the NSWs reduce the trigger rate through the use of trigger chambers that provide up to 16 space points for muons, allowing a constraint on both the position and direction of the candidate muon at the location of the NSW through a pointing segment matched to the primary trigger from the TGCs of the middle station of the MS (see also Section 5); this rate reduction is illustrated in Figure 1.8(b). In this figure, the expected Run 3 muon trigger performance is emulated by using the offline muon segments reconstructed by the MDTs and Cathode Strip Chambers (CSCs) in the old Small Wheels on so-called ‘‘enhanced bias’’ data³. The NSWs will provide even more important rate reductions in the HL-LHC era.

1.5 Outline of this Paper

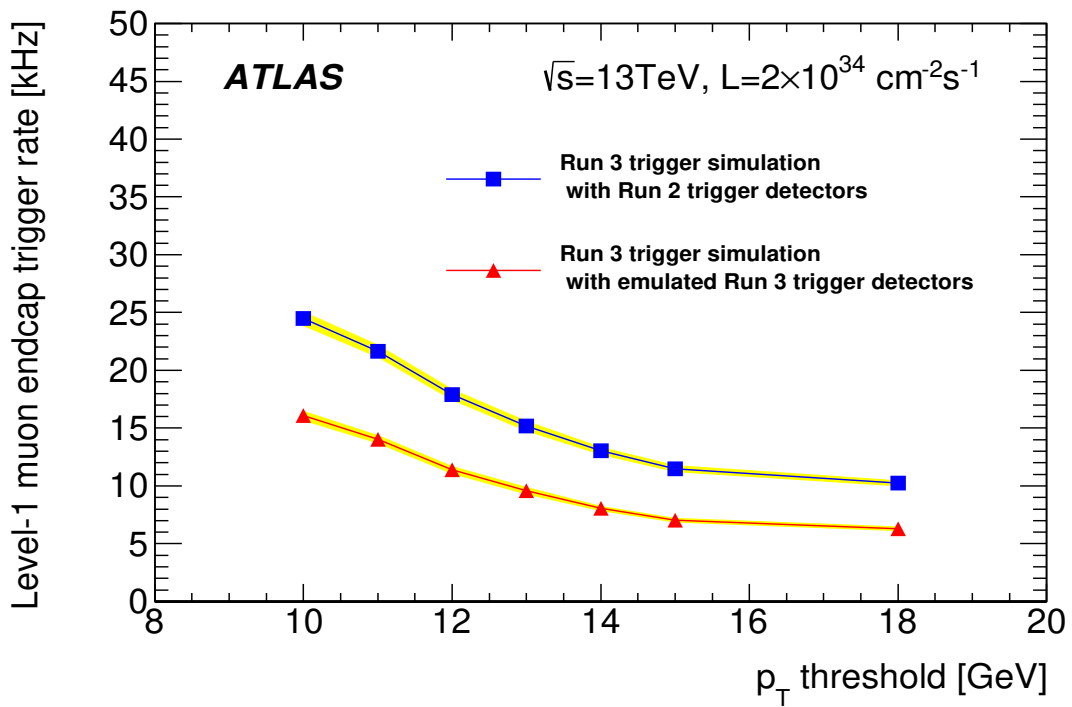
This paper describes the Run 3 configuration of the ATLAS detector. The background radiation environment and shielding are described in Section 2. Then the Inner Detector (ID), LAr and Tile calorimeters, and Muon systems are described in Sections 3, 4, and 5, respectively. The upgraded Forward Detectors are presented in Section 6, followed by a description of the Trigger and Data Acquisition (TDAQ) System in Section 7. Finally, an outlook to the Phase-II upgrades planned to prepare ATLAS for the HL-LHC are provided in Section 8.

² These include 10% (13%) of the rate at L1 (HLT) for support triggers, as was done in Run 2.

³ To assure statistical sensitivity in the most relevant kinematic regions, a mix of events is selected by the L1 trigger system that emphasises higher energies and object multiplicities. This sample, which is taken from Run 360026, is constructed in such a way that the selection bias is removable with event weights.



(a)



(b)

Figure 1.8: Expected (a) electron trigger efficiencies and (b) muon trigger rate reduction after the Phase-I upgrades.

2 Background Radiation and Shielding

The radiation field within and around the ATLAS detector is almost entirely created by the pp collisions, occurring at a rate of more than 10^9 s^{-1} at the IP. The secondary particles produced in these collisions distribute the total collision energy into different parts of the detector and the LHC ring.

About 50% of the particles produced are emitted into the acceptance of the central calorimeters ($|\eta| < 3$) but, on average, they carry only about 1% of the total energy. Around 5% of the collision energy is deposited in the forward calorimeters ($3 < |\eta| < 5$). Roughly a third is dissipated in the LHC machine elements within the experimental cavern, carried by 20% of the particles produced. The remainder escapes into the LHC ring and is of no consequence for the detector.

Without any shielding the secondary particles interacting with the beam-line elements and other material in the experimental area would result in prohibitively high radiation levels in the ATLAS cavern and detectors. Even with dedicated and carefully optimised shielding the radiation damage to sensors, electronics and power supplies is a serious concern and necessitates a careful evaluation of radiation tolerance and in many cases recourse to dedicated radiation-hard components. The margin between radiation exposure accumulated over the lifetime of the experiment and the maximum tolerated by the components is often narrow. Accurate predictions of radiation levels were therefore imperative for the original design of ATLAS and are ever more important for future upgrades.

2.1 LHC luminosity evolution and expectation

Figure 1.1 shows the luminosity which the LHC has delivered to ATLAS during proton-proton operation in the years 2011 to 2018. The 28 fb^{-1} (at $\sqrt{s} = 8 \text{ TeV}$) from Run 1 and 157 fb^{-1} (at $\sqrt{s} = 13 \text{ TeV}$) from Run 2 add up to 185 fb^{-1} . This was sufficient to cause non-negligible damage in detectors, especially the silicon sensors of the ID.

In Section 2.6 the observed radiation damage, accumulated between 2010 and 2018, is compared with simulation results in order to estimate the uncertainty of the simulation. For this the annealing (recovery) of the damage over the 7-year period, including the end-of-year technical stops and the long shutdown between Run 1 and Run 2, has to be corrected for.

The luminosity to be delivered during Run 3 might be as high as 250 fb^{-1} , but since this value is uncertain, the radiation values in this section will be kept generic by normalising them to 1 fb^{-1} . These values can be converted to instantaneous rates at, e.g. $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, by multiplying them by 10^{-5} fb/s .

2.2 Radiation shielding

The high energy particles emitted in the forward direction cross the LHC beampipe at very shallow angles and a significant fraction of them interact in the thin beampipe wall. The particle showers resulting from these interactions dissipate some of the energy locally while most continue into the dedicated absorbers (Target Absorber Secondaries (TAS)) at $|z| = 19 \text{ m}$ from the IP. These absorbers protect the superconducting final focus quadrupoles, e.g. Q1, of the LHC against the heat load due to the collision secondaries emerging from the IP. These TAS absorbers, two 1.8 m long copper cylinders with $r = 17 \text{ mm}$ apertures for the beam, are the most intense sources of background radiation in the ATLAS cavern. They are therefore surrounded by shielding made of cast ductile iron of up to 1.8 m thickness, lined with a layer




Figure 2.1: ATLAS forward radiation shielding, as described in the ϕ -symmetric FLUKA geometry. All dimensions are shown in cm. The z dimensions (horizontal axis) are measured from the nominal IP, which is to the left of the left-hand side of the drawing. The r dimensions (vertical axis) are measured from the centre of the beampipe (which is shown in the horizontal centre of the drawing). The role of the TAS absorber is to protect the focusing magnets, e.g. Q1, against particle debris from collisions at the IP. The shielding around the TAS, here labelled “TX1S”, is also referred to as “JN” in some chapters of this document. The new JD shielding that supports the NSW of the MS is located to the left of the toroid (outside the region shown in the drawing); the tip of its central hub can be seen projecting into the central opening in the toroid at the far left.

of boron-loaded polyethylene, as shown in Figure 2.1. In order to suppress the 470 keV photons from thermal neutron capture on boron, but also for fire protection, the polyethylene is enclosed in a steel cover. Because boron-10 (which makes up 20% of natural boron) has a thermal neutron capture cross-section of 3800 barn, a boron concentration of as little as a few percent is sufficient to remove thermal neutrons and to prevent them from being captured on other elements that would emit more energetic capture photons.

The beampipe within ATLAS represents a distributed source of radiation. Since calorimeters and muon stations are extended to high pseudorapidities, space left for dedicated shielding around the beampipe is very limited. In addition, any shielding has to be placed such that it does not intercept particles directed towards the TAS absorbers. Since the latter are embedded deep inside the shielding they are the preferred element to deposit waste energy. This is achieved by a conical inner bore of the shielding, as shown in Figure 2.1.

The purpose of the calorimeters is to measure the energy of particles, so the radiation load on them cannot be reduced by shielding and sufficiently radiation hard technologies must be applied. The liquid argon used in the ATLAS electromagnetic calorimeters, the hadronic endcaps and the forward calorimeters is

intrinsically radiation hard. The scintillators in the Tile calorimeter are subject to radiation damage, but their dose is reduced because they are shielded by the electromagnetic calorimeter. The electronics and power supplies are the most vulnerable parts and have to be protected by shielding, and by optimising their location.

The albedo from the calorimeters, thermal neutrons emitted back towards the IP, is a significant contribution to the radiation load of the inner tracking detector, adding to the irreducible particle flux emerging from collisions at the IP. The absorber material used in the ATLAS electromagnetic calorimeters is lead, which results in copious neutron emission upon hadronic interactions. The neutron albedo is, however, reduced significantly by polyethylene moderators lining the sides of the calorimeters which face the IP.

2.3 Characterisation of the radiation environment

The radiation environment, being composed of several particle types, all with a wide energy spectrum, cannot be described by a single quantity. On the other hand it is not generally feasible – and usually not necessary – to consider detailed particle spectra at all locations of interest. In most cases the characterisation of the radiation environment can be based on a few quantities that have proven to be related to the effects caused in various types of detectors or microelectronic circuits. The three generic damage mechanisms and the associated radiation quantities are:

- Bulk damage in silicon is usually assumed to be proportional to the 1 MeV neutron equivalent fluence ($\Phi_{\text{neq}}^{\text{Si}}$) [16], where fluence is defined as total track length per unit volume. For silicon this equivalence is defined through energy-dependent hardness factors [17], which depend on particle type and are not applicable to any other material⁴. These factors have been experimentally determined only over limited energy ranges for neutrons, protons and pions. For other particles they are extrapolated by theoretical calculations [18, 19]. In addition, different electrical properties, e.g. increase of leakage current and changes of the effective doping concentration, exhibit differences in particle type and material dependence [17]. For typical particle spectra encountered in ATLAS a systematic uncertainty of about 30% has to be assigned to the characterisation of the leakage current increase in silicon through $\Phi_{\text{neq}}^{\text{Si}}$. For estimates of the change of effective doping concentration, where the particle type and silicon properties (e.g. doping level and impurities) play a non-negligible role, it is better to consider $\Phi_{\text{neq}}^{\text{Si}}$ divided into two components: neutrons and other particles.
- Total Ionizing Dose (TID), measured as the amount of energy deposited via ionising processes per unit mass of material, leads to damage in electronics through charge trapped in oxide layers. For instance in transistors, the accumulation of this trapped charge can induce a shift of the threshold voltage needed to switch the state or increase lateral leakage current in the oxide layer. Scintillating materials and optical fibres also suffer from damage that to good approximation is proportional to the TID [20]. The damage manifests itself mainly as a reduction of the light transmission, but sometimes also as induced phosphorescence or changes of the scintillating properties.
- Single Event Effects (SEE) in electronics circuits are caused by large energy depositions close to sensitive regions of the chips. The released charge can be sufficient to flip the logic state of a transistor or, in the worst case, permanently damage the component. The amount of ionisation needed to cause a SEE can only be deposited by slow heavy ions. At the LHC such slow ions are nuclear fragments from hadronic interactions within the chip itself. The generally adopted characterisation

⁴ Hardness factor compilations also exist for other semiconductor materials, such as GaAs and diamond.

of the radiation field for SEE rate estimates is to use the total flux of hadrons with energies above 20 MeV [21] (Φ_{20}^{had}). Contrary to the two first mechanisms, which depend on the cumulative radiation exposure, the probability of a SEE occurring in any given time interval can, to a good approximation, be assumed to be independent of the irradiation history.

In addition to the characterisation by the damage potential of the radiation a further quantity is needed to describe the potential impact on detector performance: the instantaneous rate of charged particles at peak luminosity. These charged particles create hits in detectors and lead to an increase of occupancy and possibly spurious triggers which can compromise the available bandwidth and increase the dead-time. The muon system, which is designed for much more modest hit rates than the inner detector, is particularly sensitive to the detailed kinematics of the background which gives rise to fake muon trigger signals (see Section 5.1). Since low-energy charged particles have short ranges in matter, an accurate estimation of the hit rate in the sensitive volumes of the muon detectors requires a very detailed description of the detector and its shielding. General fluence simulations, as presented here, give only an indication of the hit rates to be expected.

2.4 Radiation monitoring

In view of the significant role which radiation damage was expected to play for ATLAS components, a monitoring of the radiation environment was foreseen early on. The RadMon system [22] consists of sensors at 14 locations inside and around the ATLAS detector. Each RadMon station comprises several Radiation-sensitive Field Effect Transistors (RadFet) with varying oxide thicknesses such that they together cover a range from few mGy up to 100 kGy of TID. In order to account for the fact that the radiation field around ATLAS is composed of many particle types and energies, the RadFets are calibrated at several facilities using different particle types: protons of different energies, X-rays, γ -rays and neutrons, allowing for an assessment of the differences in response for TID from different particle types. The uncertainty of the dose measurement with the RadFets was estimated at 20% [22].

The RadMon system also includes two different kinds of PIN-diodes which provide two ranges of sensitivity to measure $\Phi_{\text{neq}}^{\text{Si}}$: through an increase of the leakage current under reverse bias or through the resistivity change by monitoring the voltage under forward current. The diodes have different sensitivity ranges, such that the high-sensitivity ones provide a linear response from about 10^9 cm^{-2} up to few times 10^{12} cm^{-2} , where the low-sensitivity ones pick up and can measure up to around 10^{15} cm^{-2} . The diodes are also calibrated at various facilities, using protons, pions and neutrons, and no significant particle type dependence was uncovered. They are thus well suited for the mixed radiation field present in ATLAS.

In addition to the dedicated RadMon system, the radiation damage in the detectors themselves can be used to determine their total exposure. In particular the silicon sensors of the ID are ideal for such a study. Their irradiation and temperature (annealing) histories are recorded accurately and their leakage currents are constantly monitored. Since sensors are distributed over the entire Pixel and SCT volumes, they provide an excellent means to determine the spatial variation of $\Phi_{\text{neq}}^{\text{Si}}$ in the ID [23]; however, these sensors are not suited to measure the TID.

The scintillators of the Tile calorimeter have also suffered from radiation damage that has been compared with the simulated TID [24]. However, the radiation-induced light loss at the end of Run 2 was still very small, with large uncertainties, so the Tile scintillator data are not yet sufficient to constrain the simulation.

2.5 Radiation simulation codes and methods

The radiation environment in ATLAS has formerly been studied with the FLUKA [25–27] and G_{CALOR} [28] Monte Carlo simulation programs. Recent simulations, however, have been done with FLUKA and GEANT4 v10.4 [29]. While the first has a long history as a dedicated radiation simulation program, GEANT4 has only recently been employed for such tasks. The advantage of the GEANT4 simulation is that it benefits from a very detailed three-dimensional model of the ATLAS detector, built for physics performance simulations. Having two different simulation packages, utilising independent geometry models, provides a means to assess the systematic uncertainties due to geometry description and accuracy of the physics modelling.

The FLUKA description of the ATLAS detector is much simpler than that of GEANT4 and, with a few exceptions, ϕ -symmetric. Since, however, the FLUKA simulations are specifically targeted at estimating the radiation levels, great attention has been paid to the modelling details and material composition of the shielding. This is especially true for the ID, while for the calorimeter and cavern regions, where the FLUKA geometry lacks many details, the GEANT4 results are considered more reliable.

During the simulations several generic maps are generated, which illustrate the spatial variation of the intensity of the most relevant radiation quantities. All these maps are averaged over ϕ even if the underlying geometry is not perfectly ϕ -symmetric⁵ and mirror-symmetry with respect to $z = 0$ is assumed. The entire ATLAS cavern is covered by maps with 10x10 cm cell size in r and $|z|$. For the calorimeter regions a finer binning with 4x4 cm cell size is used. In the ID the radial gradient is steep and detector elements are thin. Here the radial bin size is 2 mm up to $r = 20$ cm and 2 cm beyond that. In $|z|$ the bin size in the ID region is 2 cm.

It is expected that the LHC will operate at a collision energy of 13.6 TeV throughout Run 3. Since an increase to 14 TeV is anticipated after LS3, the radiation simulations have been performed assuming the ultimate design energy of 14 TeV. The effect of this difference in the radiation levels is about 2%. The primary pp collisions were simulated with PYTHIA8 [30] and results were normalised assuming an inelastic cross section of 79.3 mb.

2.6 Comparison of observed radiation levels with simulations

Prior to LHC operation very few comparisons of radiation damage and simulations were available, so the initial radiation predictions [3] relied heavily on simulations and only rather vague estimations of the uncertainties could be given.

During the ten years of LHC operation the radiation field in and around ATLAS has been measured and monitored by the methods described in section 2.4. A comparison of these measurements allows for better estimates of the uncertainties associated with the simulations. The simulation models themselves have been further refined over the past decade and, most importantly, the available computing power has increased dramatically. These developments have made it possible to reduce the statistical uncertainties of the simulations to an extent that makes them negligible in almost all cases. The only exceptions are regions with very low radiation levels where the statistical fluctuations from individual energy depositions may be very large in a fully analogue simulation.

⁵ In a few situations, where the ϕ -structure has a significant impact on the radiation levels, the GEANT4 results have been averaged over restricted ϕ -ranges.

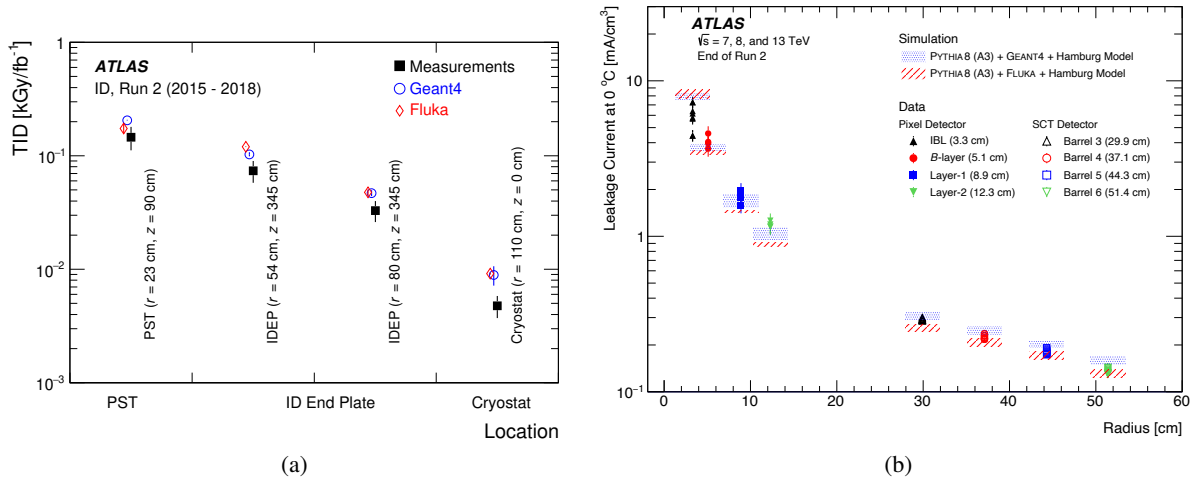


Figure 2.2: (a) Comparison of measured and simulated TID at several locations inside the ID volume. One set of monitors is close to the IP, fixed on the Pixel Support Tube (PST). Two sets are at different radii on the ID End Plate (IDEP) next to the endcap calorimeter, and a fourth set on the wall of the cryostat of the solenoid. The error bars reflect the statistical uncertainty of the simulations and the variation related to a position uncertainty of the RadMon sensors. (b) Comparison of measured and simulated leakage current in ID sensors as a function of radius [32]. The error bars on the symbols include various uncertainties affecting the measurement while the widths of the simulated bands reflect the statistical uncertainty and the variation as a function of sensor radius. Uncertainties of the silicon damage factors, discussed in Section 2.3, are not included.

2.6.1 Inner detector

Figure 2.2(a) shows the comparison of the RadMon measurements with FLUKA and GEANT4 simulations of the TID inside the ID volume. The contribution of Run 1 has been subtracted from the measurements in order to compare only the Run 2 dose, received at $\sqrt{s} = 13$ TeV. The simulations have a tendency to overestimate the TID. The difference with respect to measurements is 30 to 100% with the largest deviation on the cryostat wall where the dose is lowest.

Figure 2.2(b) shows a comparison of the leakage current measured in SCT sensors with FLUKA and GEANT4 predictions. The simulated $\Phi_{\text{neq}}^{\text{Si}}$ values are converted to a prediction of the leakage current using the Hamburg annealing model [31], which takes into account the detailed irradiation and temperature histories of the silicon modules. The agreement is better than about 20%, i.e. within the uncertainty assumed for the silicon hardness factors.

2.6.2 Calorimeter and muon regions

Figures 2.3(a) and 2.3(b) show the comparison of the RadMon measurements with TID and $\Phi_{\text{neq}}^{\text{Si}}$ predictions of FLUKA and GEANT4 at the locations of calorimeter front-end electronics and power supplies. The agreement between the simulated and measured TID is worse and less consistent than for the ID regions. At some locations the simulations are in satisfactory agreement with the measurements, but there are regions where the predictions are five times higher than the RadMon devices indicate. All locations shown in Figure 2.3 are behind a significant amount of material with respect to the IP. Some are close to, or

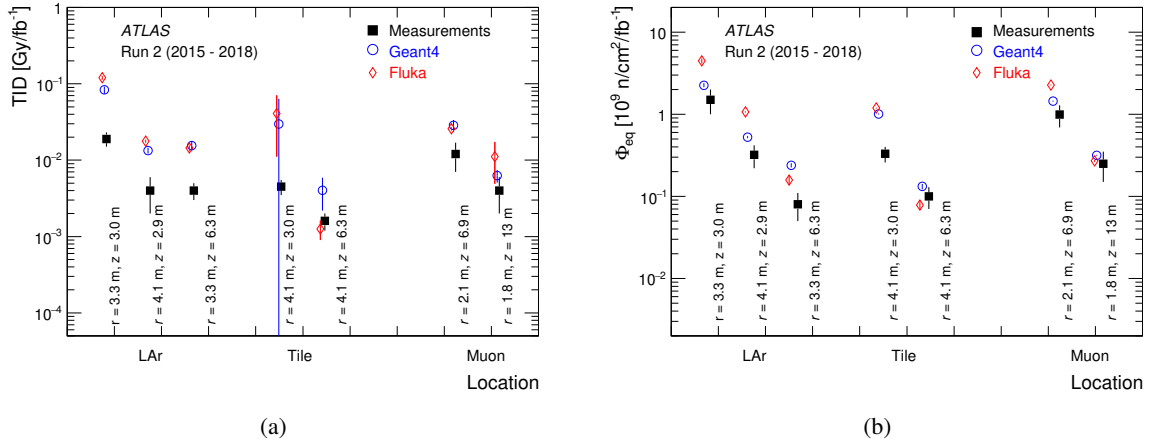


Figure 2.3: Comparison of measured (a) TID and (b) $\Phi_{\text{eq}}^{\text{Si}}$ in different locations around the calorimeters and the ATLAS muon system. The error bars reflect the statistical uncertainty of the simulations and the variation related to a position uncertainty of the RadMon sensors. Some of the very large error bars for TID are due to the fact that the quantities are recorded in a small well-shielded volume. Here a single deposition by an energetic electron can dominate and cause a large fluctuation.

even within, service channels where the exact amount of material is difficult to describe accurately in the simulation models. It is therefore very likely that the large deviations, which are all overestimations by the simulations, are caused by an underestimate of cables and other services. However, in many cases, TID predictions of FLUKA and GEANT4 are more consistent mutually than with data. It is possible that this indicates an issue with the calibration transfer of the RadFet devices, i.e. a systematic effect between the spectra at the calibration facilities and the wide particle spectrum at the location where the sensors are installed in ATLAS. This possibility needs further investigation and cross checks with additional monitors in Run 3.

For $\Phi_{\text{eq}}^{\text{Si}}$ the agreement is better with maximum discrepancies not exceeding a factor of two. This is quite plausible since the TID is more sensitive to the accurate description of material in the immediate vicinity of the RadMon detectors. Even small inaccuracies in the description of these immediate surroundings can result in a large deviation of the TID estimate while the other radiation quantities are less affected.

2.6.3 Simulation safety factors

Figures 2.2 and, especially, 2.3 indicate that the simulations can deviate from measurements by up to a factor of five. The larger deviations, however, are systematically overestimates of the simulations. All of these appear in regions behind a significant amount of material, and most likely indicate that some material is missing in the simulation models. In some regions of the inner detector the simulations underestimate the measurements, but by not more than 20%.

These comparisons are crucial input for defining the safety factors to be applied on top of the simulations, when designing the detectors and their electronics. Since the largest deviations systematically correspond to an overestimation of the simulations, ATLAS has adopted a unique simulation safety factor of 1.5 for all radiation quantities and detector regions.

2.7 Predicted Run 3 radiation levels

The radiation levels are mapped separately for each of the three main detector subsystems, as the very different sizes of the ID, the Calorimeters and the MS make different scales appropriate.

2.7.1 Radiation levels in the Inner detector

In the inner detector region the background originates from two main sources, of those discussed in Section 2.2:

- particles produced in the proton-proton collisions at the IP and secondaries from their interactions in the beampipe walls or the material of the ID.
- albedo from the electromagnetic calorimeters.

The first comprises a mix of all particle types, but for silicon bulk damage, charged pions are the most significant component. While the fluence of charged hadrons decreases rapidly with radius, the neutron fluence, which is dominated by calorimeter albedo, is more uniform throughout the Pixel and SCT volumes.

Figure 2.4 shows that, in terms of $\Phi_{\text{neq}}^{\text{Si}}$, the contributions of both generalised particle categories (neutrons, and all other particles) are equal at radii of 20 cm and 15 cm at $z = 0$ (in the central transverse plane of the detector) and at $|z| = 272$ cm (at the extremities of the SCT envelope), close to the endcap calorimeters, respectively. The bulk damage in the pixel detector is dominated by charged pions, but in the SCT by neutrons. Figure 2.4 shows that the neutron component is more significant closer to the endcap calorimeters. This is due to the intense neutron albedo from the lead absorbers, which cannot be entirely suppressed by the neutron moderators on the calorimeter face. While for older technologies, including all those installed for Run 3, $\Phi_{\text{neq}}^{\text{Si}}$ provides a sufficiently accurate parametrisation of the bulk damage in silicon, an additional particle-type dependence has been observed in some more recent devices.

In the absence of scattering and a magnetic field, it can be shown that particles emerging from the IP with a flat $|\eta|$ -distribution would result in a fluence independent of $|z|$ and dropping as r^{-2} . Figure 2.4 shows that this radial dependence models the data well up to $r \sim 20$ cm but in the SCT region the fluence drops somewhat more slowly due to particle production in inelastic interactions and possibly the curling up of tracks in the 2 T axial field. Comparing Figures 2.4(a) and 2.4(b) also reveals a small $|z|$ -dependence of the 1 MeV neutron-equivalent fluence. Besides the fact that the $|\eta|$ -distribution is not perfectly flat, this increase of $\Phi_{\text{neq}}^{\text{Si}}$ with $|z|$ may also be due to secondary production and the magnetic field. The increase of the neutron contribution closer to the endcap calorimeter, clearly seen when comparing the two plots, shows the impact of the calorimeter albedo, which is the main source of neutrons.

This r - and $|z|$ -dependence of $\Phi_{\text{neq}}^{\text{Si}}$ can also be appreciated qualitatively from Figure 2.5. Table 2.1 lists the three main radiation quantities averaged over the regions indicated by rectangles and labelled with letters in Figure 2.5. These areas are selected to cover representative regions of the Pixel, SCT and TRT detectors. The r and $|z|$ positions given in Table 2.1 correspond to the centre of the rectangle. The statistical uncertainties of the simulations are typically less than 1% but if the last digit given in Table 2.1 is affected, this is indicated. Systematic uncertainties are not included and, as discussed before, can amount to several tens of percent.




Figure 2.4: Radial dependence of $\Phi_{n_{eq}}^{Si}$ at (a) $|z| = 0$ and (b) $|z| = 272$ cm in the ID, as obtained from FLUKA simulations. The total $\Phi_{n_{eq}}^{Si}$ is divided into the contribution by neutrons and by all other particles. The r^{-2} dependence is normalised at $r = 3.1$ cm to the ‘By others’ value. The lower panels show the ratio of FLUKA and GEANT4 results for the total.

2.7.2 Radiation levels around the calorimeters

Figure 2.6 and the associated Table 2.2 show the $\Phi_{n_{eq}}^{Si}$ distribution in the calorimeter regions and the volume-averaged main radiation quantities, respectively. The regions for which the averages have been calculated correspond to the location of front-end and power supply electronics of the calorimeters, i.e. to equipment susceptible to suffer from radiation damage. In Figure 2.6 the streaming of radiation through the gaps between the calorimeter parts, especially at $|z| \approx 320$ cm, is clearly visible. In reality these gaps are filled with electrical cables, optical fibres and cooling pipes following complicated routing. These can be described only in an approximate and average way in the simulation codes, which is the possible cause for the overestimation of the radiation levels, seen in Figure 2.3. Although no direct measurements are available to support the predictions of Φ_{20}^{had} , high-energy hadrons are less sensitive to material than the other two radiation components. Therefore it can be assumed that the uncertainties are bounded by the uncertainty on $\Phi_{n_{eq}}^{Si}$, seen in Figure 2.3(b).

2.7.3 Radiation levels in the muon system and cavern

Since the cavern is efficiently shielded with respect to the beam-line, the available simulation samples have too much statistical fluctuation between bins to produce a contour plot of TID or charged particle fluences in the entire ATLAS experimental cavern. Instead Figure 2.7 shows the distribution of the photon fluence.




Figure 2.5: Radiation levels in the inner detector region. The contours show $\Phi_{\text{neq}}^{\text{Si}}$ obtained from FLUKA simulations. The radiation levels, averaged over the labelled volumes are given in Table 2.1 The Pixel and SCT layers are indicated by grey lines; the TRT and adjacent calorimeter and associated neutron moderator volumes are indicated by boxes.

Table 2.1: TID and characteristic particle fluences in the inner detector volume, obtained from FLUKA simulations at the locations indicated in Figure 2.5. The statistical uncertainty is in the last digit given, or smaller.

Region	$ z $ [cm]	r [cm]	TID [Gy/fb^{-1}]	$\Phi_{\text{neq}}^{\text{Si}}$ [$\text{cm}^{-2}/\text{fb}^{-1}$]	Φ_{20}^{had} [$\text{cm}^{-2}/\text{fb}^{-1}$]
IBL e	29	3	3.42×10^3	5.37×10^{12}	9.08×10^{12}
IBL c	4	3	3.12×10^3	6.19×10^{12}	8.49×10^{12}
PIX D3	65	9	821	1.13×10^{12}	1.45×10^{12}
PIX B3	4	12	377	7.85×10^{11}	7.22×10^{11}
SCT D6	178	30	171	3.58×10^{11}	2.49×10^{11}
SCT B1	71	30	118	2.93×10^{11}	2.00×10^{11}
SCT D9	273	46	99.1	3.36×10^{11}	1.58×10^{11}
TRT EC low	266	68	65.0	2.34×10^{11}	8.50×10^{10}
SCT B4	71	52	46.0	1.68×10^{11}	8.12×10^{10}
TRT B low	69	61	36.7	1.34×10^{11}	6.09×10^{10}
TRT EC up	266	94	35.2	1.86×10^{11}	5.10×10^{10}
TRT B up	69	103	12.9	8.27×10^{10}	2.37×10^{10}

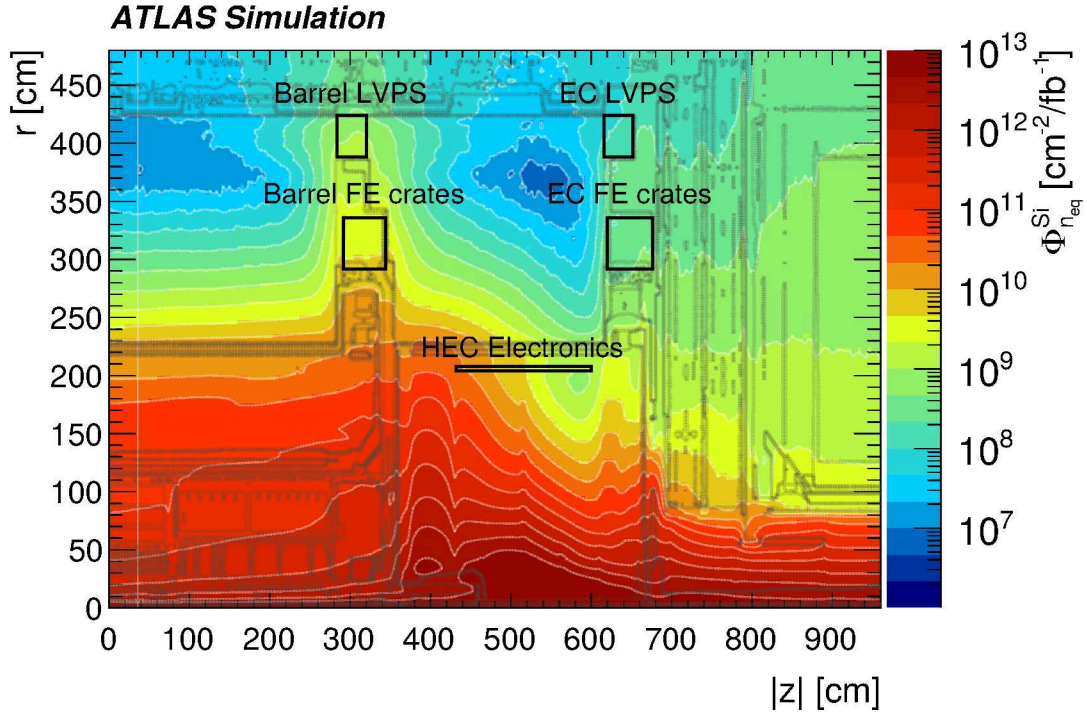


Figure 2.6: Radiation levels at the locations of calorimeter electronics. The contours show $\Phi_{\text{neq}}^{\text{Si}}$ obtained from GEANT4 simulations. The radiation levels in the indicated volumes are given in Table 2.2. The faint grey lines are density contours from the GEANT4 geometry model and only serve to give an indication of the detector geometry.

Table 2.2: TID and characteristic particle fluences obtained from GEANT4 simulations at the locations of calorimeter electronics and power supplies, i.e the regions indicated in Figure 2.6.

Region	z [cm]	r [cm]	TID [Gy/fb ⁻¹]	$\Phi_{\text{neq}}^{\text{Si}}$ [cm ⁻² /fb ⁻¹]	Φ_{20}^{had} [cm ⁻² /fb ⁻¹]
Barrel FE crates	318	314	1.42×10^{-1}	3.80×10^9	6.50×10^8
HEC Electronics	516	206	8.00×10^{-2}	7.15×10^9	4.44×10^8
Barrel LVPS	302	406	1.97×10^{-2}	1.00×10^9	1.45×10^8
EC FE crates	648	314	1.88×10^{-2}	3.40×10^8	4.14×10^7
EC LVPS	634	406	5.9×10^{-3}	1.31×10^8	1.33×10^7

This quantity is chosen for the illustration, since the photons are generating the e^\pm , which dominate the charge particle fluence and TID in the muon spectrometer region.

All major radiation quantities are collected in Table 2.3 as averages over representative volumes. The dose rates in the muon system vary from 0.59 mGy/fb^{-1} in the central barrel to 44 mGy/fb^{-1} at the innermost radius of the NSW region. At an instantaneous luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ these correspond to 12 nGy/s and $0.9 \text{ }\mu\text{Gy/s}$, respectively. Assuming all particles to be minimum ionising, these doses can be converted to rough charged particle flux estimates of about $50 \text{ cm}^{-2} \text{ s}^{-1}$ and $5000 \text{ cm}^{-2} \text{ s}^{-1}$, respectively. In reality the TID has contributions from particles with higher ionisation potential, which means that those estimates serve only as rough upper limits.

As shown in Figure 2.7, the LUCID detector is situated in the immediate proximity of the beampipe where it is also subject to intense albedo from the TAS regions. A comparison of the radiation levels in that region with those of the cavern gives a good indication of the performance of the forward shielding of ATLAS: the particle fluences are reduced by about four orders of magnitude and the TID by almost 6 orders of magnitude (e.g. comparing regions ‘LUCID’ and ‘EM-1’). This larger shielding efficiency for TID is due to the fact that charged particle showers, except for muons, developing around the beam-line all get suppressed and only neutrons with an associated photon component penetrate the shielding. Thus the radiation field in the cavern is almost entirely due to neutrons, photons from neutron capture and e^\pm from interactions of the photons. One notable exception in the radiation exposure of the muon system is the high- η edge of the NSW (NSWH in Table 2.3); here the space available for shielding is limited by the required acceptance and the clearance needed with respect to the beampipe. Even though dedicated shielding materials are employed, the radiation levels at the smallest radii of the NSW are almost two orders of magnitude higher than anywhere else in the muon system.

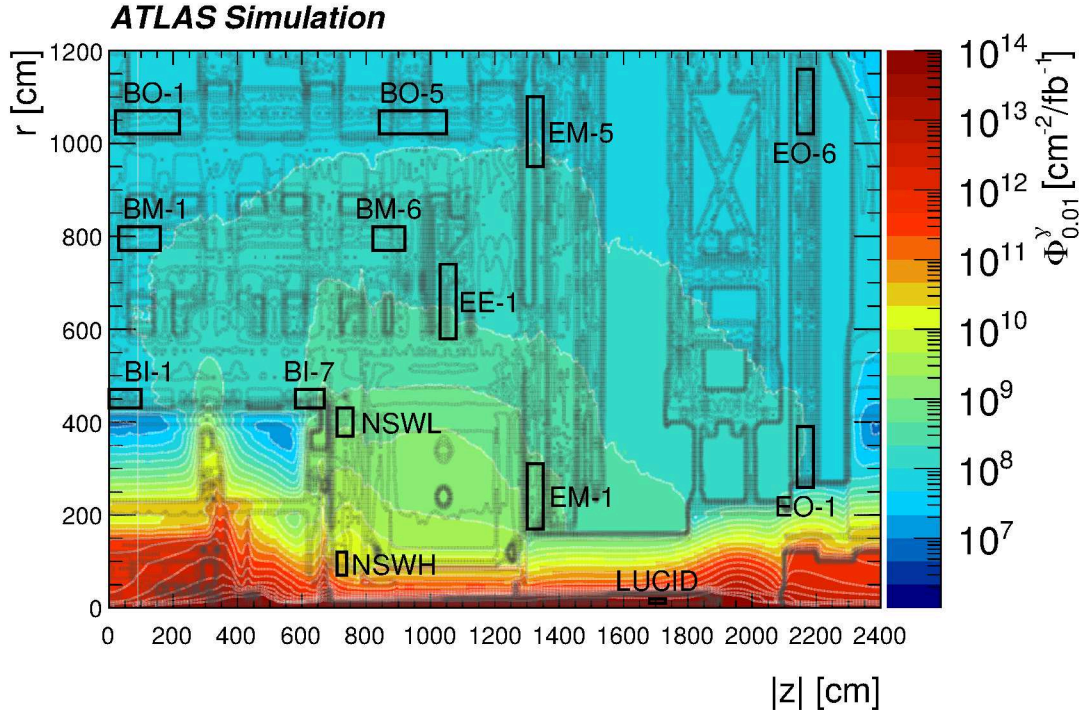


Figure 2.7: Contours of constant fluence for photons with $E > 10 \text{ keV}$ ($\Phi_{0.01}^y$) in the ATLAS cavern and muon spectrometer regions. The black boxes indicate regions over which the main radiation quantities given in Table 2.3 are averaged. More detailed explanations of the detectors in these regions can be found in Section 5. The faint grey lines are density contours from the GEANT4 geometry model and only serve to give an indication of the detector geometry.

Table 2.3: TID and $\Phi_{\text{neq}}^{\text{Si}}$ obtained from GEANT4 simulations in the ATLAS cavern and muon spectrometer as averages over the volumes indicated in Figure 2.7. The rightmost column gives the fluence of photons with $E > 10 \text{ keV}$ ($\Phi_{0.01}^y$). r and $|z|$ refer to the centre of the given region. The statistical uncertainty is in the last digit given, or smaller.

Region	$ z $ [cm]	r [cm]	TID [Gy/fb^{-1}]	$\Phi_{\text{neq}}^{\text{Si}}$ [$\text{cm}^{-2}/\text{fb}^{-1}$]	Φ_{20}^{had} [$\text{cm}^{-2}/\text{fb}^{-1}$]	$\Phi_{0.01}^y$ [$\text{cm}^{-2}/\text{fb}^{-1}$]
LUCID	1705	15	2.91×10^3	2.57×10^{12}	4.83×10^{11}	5.78×10^{13}
NSWH	725	95	4.40×10^{-1}	1.79×10^{10}	1.86×10^9	3.09×10^{10}
EM-1	1325	240	7.6×10^{-3}	3.03×10^8	6.45×10^7	6.17×10^8
NSWL	735	400	6.9×10^{-3}	2.02×10^8	4.81×10^7	9.03×10^8
BI-7	625	450	4.4×10^{-3}	9.05×10^7	1.13×10^7	3.42×10^8
EE-1	1055	660	2.43×10^{-3}	1.10×10^8	2.30×10^7	2.05×10^8
BM-6	870	795	1.63×10^{-3}	4.63×10^7	1.13×10^7	1.56×10^8
EO-1	2165	325	1.3×10^{-3}	3.19×10^7	3.99×10^6	9.58×10^7
EM-5	1325	1025	1.15×10^{-3}	3.21×10^7	8.34×10^6	9.32×10^7
BO-5	945	1045	1.07×10^{-3}	2.36×10^7	6.94×10^6	7.98×10^7
BM-1	95	795	8.1×10^{-4}	1.52×10^7	2.43×10^6	8.87×10^7
BI-1	50	450	7.3×10^{-4}	2.70×10^7	7.0×10^5	8.55×10^7
BO-1	120	1045	6.4×10^{-4}	1.03×10^7	2.59×10^6	5.69×10^7
EO-6	2165	1090	5.9×10^{-4}	1.36×10^7	3.23×10^6	5.82×10^7

3 Inner Detector

The Inner Detector (ID) is the primary tracking device for measuring the paths of all charged particles in ATLAS. It has been designed to provide hermetic and robust pattern recognition, excellent momentum resolution and both primary and secondary vertex measurements for charged-particle tracks within the pseudorapidity range $|\eta| < 2.5$. It is contained in a cylindrical envelope 7024 mm long and 1150 mm in radius, immersed in the 2 T field of a solenoidal magnet.

As described in Section 1.3.2, the ID consists of three complementary sub-detectors arranged coaxially around the beam line (see Figure 3.1): a high-resolution silicon Pixel detector [33] ($r < 122.5$ mm), the Semiconductor Tracker (SCT) [34] relying on stereo micro-strips ($299 < r < 514$ mm) and the Transition Radiation Tracker (TRT) [35] comprising several layers of gaseous straw tubes interleaved with transition radiation material ($563 < r < 1066$ mm).

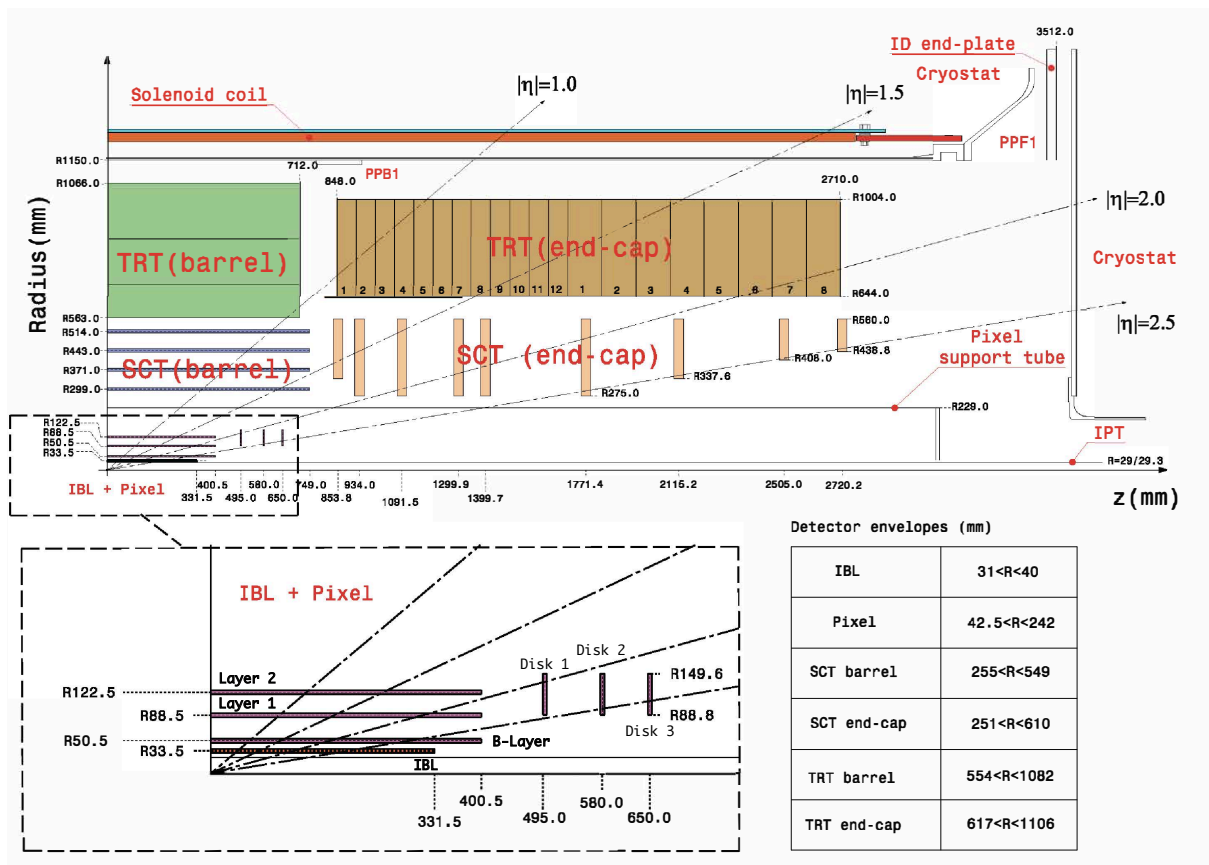


Figure 3.1: The layout of the ATLAS ID, including the IBL detector [9].

For many physics channels, particularly those involving relatively long-lived particles such as B -hadrons that decay and produce a secondary vertex inside the beampipe, the performance of the ATLAS experiment depends critically on the innermost layer of the Pixel detector. For this reason, during LS1, the detector underwent a major upgrade. During this period a fourth pixel layer, the Insertable B -Layer (IBL) [8, 9, 36], was added to the Pixel detector between a new, narrower beryllium beampipe and the previously existing innermost pixel layer (Pixel B -Layer). Figure 3.2 shows the transversal layout of the upgraded ATLAS ID

during Run 2 and Run 3. A description of the main concepts and characteristics of the IBL is given in Section 3.1.




Figure 3.2: Transverse view of the ATLAS ID including the IBL detector.

At the same time, part of the Pixel services was replaced by the new Service Quarter Panels (nSQPs), allowing for the optical-to-electrical converters to be moved into an area accessible for service by extending the electrical readout cables. A brief overview of the new Service Quarter Panels (nSQPs) installation together with a description of the Pixel off-detector readout upgrades in Run 2 is given in Section 3.2.

The readout systems of the SCT and TRT were also upgraded during Run 2, to cope with the increasing requirements coming from the LHC performance. A brief description of these upgrades is presented in Sections 3.3 and 3.4.

3.1 Insertable *B*-layer

The driving motivation for the Insertable *B*-Layer (IBL) detector was the consolidation and enhancement of the ID tracking performance in high luminosity scenarios. The reduced distance of the IBL from the beam axis (3.3 cm, compared to 5.0 cm for the Pixel *B*-Layer) increases the resolution of track impact parameters and thus enhances the vertex reconstruction and flavour-tagging performance of the tracking system. An additional layer of highly segmented pixel sensors helps to mitigate inefficiencies of the Pixel detector caused by overall radiation damage and irreparable module failures, and reduces the occurrence of fake tracks due to the combinatorics of high pileup backgrounds.

The reduced distance from the beam axis, combined with the increased luminosity of the machine, set stricter radiation hardness requirements. Simulations [8] predicted that IBL sensors and electronics would

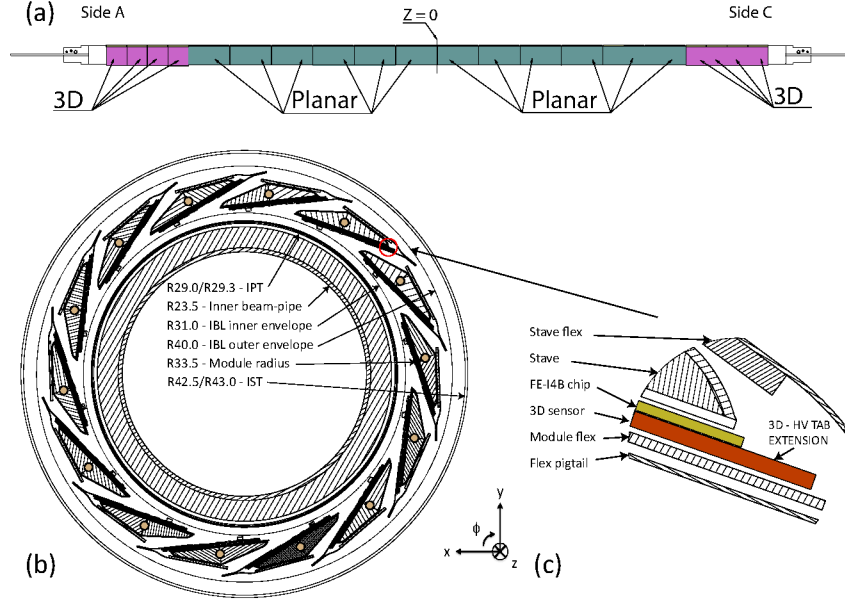


Figure 3.3: IBL layout [9]: (a) Longitudinal layout of planar and 3D modules on a stave. (b) An $r - \phi$ section showing the beampipe, the IPT, the staves of the IBL detector and the IST, as viewed from the C-side. (c) An expanded $r - \phi$ view of the corner of a 3D module fixed to the stave.

be subject to an overall fluence $\Phi_{\text{neq}}^{\text{Si}}$ of $3.3 \times 10^{15} \text{ cm}^{-2}$ for an integrated luminosity of 550 fb^{-1} , which was the original design requirement for Run 3. A planar technology adapted according to the Run 1 Pixel experience was used for the sensors in the central part while a new 3D technology was developed for the forward region. Section 3.1.2 describes these two technologies. For the electronics, a new front-end chip was developed on the basis of the Run 1 Pixel experience to cope with the more stringent requirements foreseen in future data taking scenarios.

The IBL layout is shown in Figure 3.3. The IBL consists of a single layer of pixel sensors assembled with their readout chips (modules) and arranged on 14 longitudinal supports (staves) fastened onto a high precision carbon fibre tubular structure, the IPT (outer radius of 29.3 mm), via the stave support ring. The staves surround a 23.5 mm-radius beryllium beampipe. The outer envelope is defined by a second carbon fibre cylinder, the IST (inner radius of 42.5 mm), fastened to the Pixel detector support structure.

The staves are tilted by 14° in a classic turbine design in order to ensure complete azimuthal coverage for high p_T tracks and to partially compensate for the Lorentz angle affecting the trajectory of the charges drifting in the planar sensors. With this arrangement, the average distance of the innermost pixel sensors from the beam-axis was reduced from 5.05 cm of the Pixel B -Layer down to 3.35 cm. Each stave hosts 20 modules aligned along the beam axis (z -direction): 12 two-chip modules covering the central part and eight single-chip modules dedicated to the high- $|\eta|$ regions extending to $|\eta| = 3.0$. The full detector coverage is effectively $|\eta| = 2.58$ if the 122 mm two-standard-deviation spread of the primary vertex distribution is taken into account. The main layout parameters of the IBL are presented in Table 3.1.

Limited IBL radial clearance prevents sensor shingling along a stave; for this reason, in order to maximise the coverage, thin-edge sensors were developed for the detector. An effective inactive edge width of $215 \mu\text{m}$ ($175 \mu\text{m}$) was measured for planar (3D) sensors, considerably reduced with respect to the $1100 \mu\text{m}$ of the

Table 3.1: Main layout parameters of the IBL. The nominal η coverage of the detector (assuming no spread on the longitudinal position of the IP) is $|\eta| < 3.0$ but a spread of the vertex position on the z -coordinate reduces the detector coverage down to $|\eta| < 2.58$.

Item	Value
Number of staves	14
Number of modules per stave	12 planar + 8 3D
Number of front-end chips per stave	32
Coverage in η – no vertex spread	$ \eta < 3.0$
Coverage in $\eta - 2\sigma$ (122 mm) vertex spread	$ \eta < 2.58$
Active $ z $ stave length	330.15 mm
Stave tilt in ϕ	14°
Overlap in ϕ	1.82°
Center of the sensor radius	33.5 mm

Pixel detector in Run 1. An air gap of $205\ \mu\text{m}$ is maintained between contiguous modules to provide adequate electrical insulation.

Cooling is achieved by means of a CO_2 two-phase system where the coolant is circulated within titanium pipes embedded in the stave structure; in order to optimise the thermal contact between the active components and the pipes, staves are filled with carbon foam, which contributes to the global stiffness of the mechanical supports. A detailed description of the cooling system is given in Section 3.1.8.

3.1.1 New beampipe

The ATLAS beampipe had to be replaced to allow sufficient radial clearance for the insertion of the IBL detector and its mechanical structure. The new beampipe, characterised by an inner radius of 23.5 mm (reduced from the previous 29 mm), consists of a 7100 mm-long beryllium section with an average wall thickness of approximately $870\ \mu\text{m}$; its extremities are welded to aluminium flanges, the size of which allow its insertion through the Inner Positioning Tube (IPT).

The inner surface of the beampipe is treated with a Non-Evaporative Getter (NEG) coating to improve the vacuum quality by bonding to gas molecules remaining within the partial vacuum; this getter coating had to be activated with a bake-out procedure using heaters wrapped around the beampipe surface. To mitigate the effect of extreme heat on the silicon sensors, a layer of aluminium was interleaved between the beampipe and the IBL to reduce the infra-red emissivity of the beampipe. The heater temperature was carefully monitored during the process.

3.1.2 Sensors

Two sensor technologies are used in the IBL: an improved version of the ATLAS Run 1 Pixel planar sensors [33], able to cope with unprecedented operational conditions, and an innovative design of 3D devices [37] a first-ever application in high energy physics.

The two sensor implementations share a common footprint in order to be interfaced to the same readout chip, with two-chip planar sensor tiles having an equal longitudinal dimension of two single-chip 3D sensors in

Table 3.2: Comparison of the main characteristics of the IBL pixels with the original Pixel detector layers.

Technical Characteristic	Pixel	IBL
Active surface (m ²)	1.73	0.15
Number of channels (×10 ⁶)	80.26	12.04
Pixel size (μm ²) - nominal/long	50x400/600	50x250/500
Pixel array (columns×rows)	160x18	336x80
Front-end chip size (mm ²)	7.6x10.8	20.2x19.2
Active surface fraction (%)	74	89
Analogue current (μA)	26	10
Digital current (μA)	17	10
Analogue voltage (V)	1.6	1.4
Digital voltage (V)	2.0	1.2
Data out transmission (Mbps)	40 – 160	160
Sensor type	planar	planar/3D
Sensor thickness (μm)	250	200/230
Layer thickness (X ₀)	2.8	1.9
Cooling fluid	C ₃ F ₈	CO ₂

order to fit a common mechanical layout. Both technologies are characterised by the same granularity, with 250x50 μm pixels organised in 80x336 arrays (to be compared with the 18x160 matrix configuration of 400x50 μm pixels of the three outermost Pixel layers). Pixels of sizes different from the nominal are also present in order to fill the gaps between front-ends. A comparison of the main characteristics of the Pixel and IBL detectors is shown in Table 3.2.

IBL planar sensors are fabricated with a n-in-n design on a 200 μm-thick substrate, thinner than the 250 μm devices used in the rest of the Pixel detector. Slim inactive edges (200 μm wide) are achieved by shifting the guard rings on the p-side underneath the pixel implantations. Planar sensors are produced in two-chip tiles of an overall dimension of 18.59x41.32 mm; the operational voltage is expected to evolve during the device lifetime from 80 V up to 1000 V.

IBL 3D sensors rely on a 230 μm-thick p-type substrate and are fabricated with a process in which 10 μm-diameter columnar electrodes are implanted by double-sided Deep Reactive Ion Etching (DRIE) – i.e. n-type and p-type columns penetrate the substrate from opposite sides. Columnar electrodes are either passing-through or 210 μm deep according to the different designs implemented by the vendors (FBK⁶ and CNM⁷) at the time of production; see Figure 3.4 for a visual comparison of the two designs. The pixel layout consists of two n-type readout electrodes connected at the wafer surface and surrounded by six p-type ohmic electrodes which are shared with the neighbouring cells. Edge isolation is achieved by a fence of ohmic electrodes corresponding to an inactive area of 200 μm deep. Due to the characteristic design, which decouples the sensor thickness from the drift distance, 3D sensors are expected to be particularly radiation tolerant and for this reason suitable to equip detectors exposed to high irradiation levels. The operational bias of these sensors started with 20 V at the beginning of Run 2 and is expected not to exceed 200 V during the device lifetime.

⁶ Fondazione Bruno Kessler, Trento, Italy

⁷ Centro Nacional de Microelectronica, Barcelona, Spain




Figure 3.4: Design of the columns of (a) FBK and (b) CNM 3D sensors, including the location of the bump used to bond the sensor to the front-end chip [9]. This sketch is for illustration only and is not to scale.

3.1.3 Modules and staves

The connection between sensor and front-end chip is achieved by means of fine-pitch bump-bonding and flip-chip technology. Due to more stringent radiation and occupancy requirements, the front-end electronics of the IBL has been completely redesigned with respect to the integrated circuit used in the rest of the Pixel detector in Run 1 (the FE-I3). The new front-end chip, the FE-I4B [38], was manufactured in 130 nm Complementary metal-oxide-semiconductor (CMOS) technology, featuring a large footprint (20.3x19.2 mm) and 26 880 readout channels replicating the pixel matrix. Each cell comprises an independent amplifier with adjustable shaping and a discriminator with an individually adjustable threshold. The analogue threshold (typically 1500-2500 electrons) sets the minimum collected charge to be processed (and converted to Time over Threshold (ToT)) by the front-end chip; the digital threshold (in units of ToT) sets the minimum digital collected charge amplitude to be transmitted by the module and is typically 2Bunch Crossing (BC). The collected charge amplitude is measured as ToT in units of the 25 ns LHC bunch crossing period with a four-bit resolution. The FE-I4B chips (see Section 3.1.5 for details about the readout architecture) are slimmed to 150 μm in order to reduce the material budget.

A double-sided, flexible printed circuit (known as the module flex hybrid) is glued to the back side of the sensor (see Figure 3.5) and connects the module to external electrical services. Read-out and power lines connections between the module flex hybrid, the FE-I4B chips and the sensors are wire-bonded. The module temperature monitoring and interlock is achieved with a Negative Temperature Coefficient (NTC) thermistor mounted on the module flex hybrid.

IBL modules are fastened by two glue dots to their supporting staves. Staves are carbon foam structures glued to V-shaped carbon fiber laminates; the carbon foam provides efficient thermal coupling to the titanium cooling pipes that are embedded within the staves. Cooling is obtained by a CO₂ evaporative system, capable of maintaining the silicon sensors at an operating temperature of $-15\text{ }^{\circ}\text{C}$; the minimum value achievable for future operation is $-35\text{ }^{\circ}\text{C}$. Power, data acquisition and module configuration are routed through a multi-layer bus (the stave flex hybrid) laminated on the back of the staves, from which they are distributed to modules by means of the module flex hybrids.

During the first year of the IBL operation in 2015, a significant increase of the low voltage current of the front-end chip and a detuning of its parameters (threshold and ToT) were observed in relation to the received TID. The increase of the low voltage current of the FE-I4B chip and the drifting of its tuning parameters were traced back to the generation of a leakage current in n-type metal-oxide-semiconductor (NMOS) transistors induced by radiation. The radiation induces positive charges that are quickly trapped into the shallow-trench-isolation oxide at the edge of the transistor. Their accumulation builds up an electric




Figure 3.5: Photograph of an IBL planar module [9]. The full size of the silicon sensor is 18.8x41.3 mm.

field sufficient to open a source-drain channel where the leakage current flows. If the accumulation of positive charges is relatively fast, the formation of interface states is a slower process. The negative charges trapped into interface states start to compete with the oxide-trapped charges with a delay. This is what gives origin to the TID effect at low dose [39].

In order to study the dependence of the low voltage current increase on temperature and dose rate, several irradiation tests were performed by setting one of those variables and changing the other. Given the results of these tests, it was decided to increase the IBL operational set point from $-10\text{ }^{\circ}\text{C}$ to $+15\text{ }^{\circ}\text{C}$ for a few months in early 2016 to limit the TID effects during the first radiation exposure, before reducing it to $-20\text{ }^{\circ}\text{C}$ for the remainder of Run 2. In addition, the digital supply voltage (V_D) was lowered from 1.2 V to 1 V to decrease the low voltage current.

3.1.4 Integration and installation of the IBL in the ID

A total of 20 staves were assembled during the production phase. An issue concerning wire-bond corrosion [9] was identified on most staves mid-way through the production. This resulted from a combination of extreme susceptibility to corrosion due to chemical contamination on the flex and accidental exposure to humidity during the temperature cycling tests after stave loading; all but two staves were fully repaired. Because the wire bonds of the IBL were left unprotected (potting the bond foot or the use of spray coatings such as polyurethane were initially considered as options to protect them), it was crucial that the temperature be maintained well above the dew point at all stages of the integration and installation, and it remains crucial for operation.

Of the 18 staves available for the integration, 14 were selected on the basis of quality criteria and installed in the detector. These criteria rely on a weighted pixel failure fraction aimed at minimising the geometrical inefficiencies and clusterisation of defects. The average bad pixel fraction for the integrated IBL staves at the time of installation amounted to 0.09% (0.07% for $|\eta| < 2.5$).

The IBL was integrated as a single package that was lowered into the ATLAS experimental cavern and finally inserted inside the Pixel detector. Once all staves were loaded, an annular support was installed and clipped to each stave center in order to provide additional radial stiffness. This constraint did not eliminate other degrees of freedom, most notably a rotation around the beam axis arising from the mismatch of

thermal expansion coefficients of the stave and stave flex and their asymmetric assembly; this resulted in a distortion of the staves of the order of a few $\mu\text{m} / ^\circ\text{C}$ that required an enhanced temperature stability at the level of 0.2°C and the development of in-run alignment correction procedures [40].

3.1.5 Readout of the IBL

The IBL readout system was designed to be fully efficient at a few times the nominal LHC peak luminosity (up to $2\text{--}3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$) with a L1 trigger rate of 100 kHz. The readout can be naturally divided into two parts, the on-detector and the off-detector systems, which communicate via optical fibres in both directions.

The on-detector system consists of the FE-I4B front-end readout chips and the service electronics, which include the electro-optical converter boards (optoboards). The optoboards are connected to the front-end chips by the FE-I4B input and output lines.

The building block of the readout system, the IBL DAQ module, is formed by two neighbouring front-end chips. Each IBL DAQ module shares the same clock and command lines (including the L1 trigger signals) that are distributed to the pair of front-end chips.

The FE-I4B input lines run at 40 Mbps. When charge deposition is detected by the discriminator in each pixel, the hit and its timestamp are briefly buffered in the pixel cells. The FE-I4B pixel array is organised in double-columns like the FE-I3 chip, but the readout architecture is very different. Instead of moving all hits from the pixel array to a global shared memory structure for later trigger processing, the FE-I4B double-columns are further divided in 2×2 pixel regions where hits are stored locally. This results in an enhanced capability to cope with high hit rates and reduced inefficiency. Each region contains four identical analogue pixels and one shared memory and logic block called the Pixel Digital Region. This memory can store up to five events.

When the Level-1 Trigger Accept (L1A) arrives, any event for which the timestamp matches is read out via a serial Low-voltage Differential Signaling (LVDS) output that operates at 160 Mbps. There is one LVDS output line per front-end.

The optoboards are responsible for translating the optical signals received from the off-detector electronics into electrical signals for the FE-I4B inputs; vice versa they translate the electrical signals received from the FE-I4B outputs to optical signals before transmitting them to the off-detector electronics. PIN diodes are used to convert optical into electrical signals. Vertical-Cavity Surface Emitting Laser (VCSEL) arrays convert electrical into optical signals.

The off-detector electronics consists of commercial optical Rx plugins and Tx plugins, used to interface the readout hardware with the optical fibres. The plugins are mounted in the Back of Crate (BOC) cards, used to transmit the clock, triggers and commands to the modules after Bi-Phase Mark (BPM) encoding, and to receive and decode 8b/10b-encoded data from the front-ends. The BOC communicates through the VME backplane to the ROD cards. The ROD has several tasks: it distributes the L1 trigger upstream, along with signals received from the ATLAS Trigger, Timing and Control system (TTC) Interface Module (TIM), reformats the data received from the front-ends (via the BOC), and finally transmits the generated event fragments to the ATLAS ROS [41] via CERN (S-LINKs) [42]. This transmission is done by forwarding the reformatted data back through the backplane to dedicated optical transceivers (Quad-SFP [43]) on the BOC. Each S-LINK has a bandwidth of 1.28 Gbps. As shown in Figure 3.6, a single IBL ROD/BOC pair hosts four S-LINKs for a total output bandwidth towards the ATLAS ROS units of 5.12 Gbps. The ROD




Figure 3.6: The IBL ROD/BOC readout cards [9] (used also for the Pixel readout upgrades during three consecutive winter shutdowns) are located at the off-detector side of the optical link. (a) The BOC card and (b) the ROD card are paired in a VME crate via its back-plane.

is also used to calibrate and tune the IBL detector. A schematic of the entire readout chain is shown in Figure 3.7.

Each IBL ROD/BOC pair reads out an entire IBL stave via 32 FE-I4B chips. Two Tx plugins are used to communicate with 16 double-chip modules via 16 independent serial lines, eight lines for each Tx plugin. Four Rx plugins read back the data from each of the 32 front-ends individually; each Rx plugin receives data from eight front-ends. A total of 14 IBL ROD/BOC pairs have been equipped in a 9U VME crate located in the ATLAS service cavern (USA15).

The same type of readout board was used to upgrade the outer Pixel layers during Run 2. A detailed description of these upgrades will be given in Section 3.2.2.

3.1.6 Off-detector services and detector power supplies

The basic element for the segmentation of the IBL electrical services is the half stave. The end of each half-stave is connected to a set of corrugated flex Printed Circuit Boards (PCBs) which on the other side connect to a PCB called the cable board as is schematically shown in Figure 3.8. The cable board is part of a cable harness, called a type-1 cable, which makes the connection out to the ID endplate region. Each harness is divided into two individual sub-cables, one for data transmission and one for power and the Detector Control System (DCS).

The type-1 cable transmission bundle consists of 24 polyimide-clad copper twisted pairs. For the data from the front-ends 28 American Wire Gauge (AWG) wire is used due to the higher rate of 160 Mbps, while for the 40 MHz clock and the 40 Mbps command transmission 36 AWG wire is used. Each command and clock line is shared by two frontend chips. The standard for all signals is LVDS. Four 36 AWG ground wires in each bundle ensure a common ground level between transmitters and receivers. The transmission cable is approximately 5 m long and terminates at the optobox in the outer part of the ID endplate region.

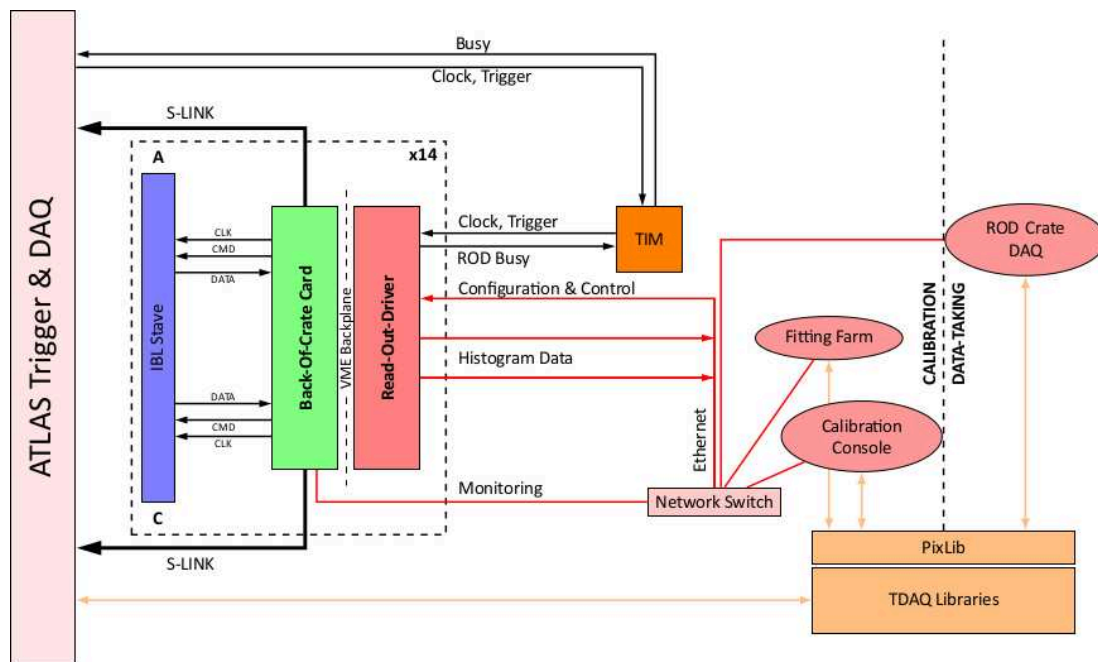


Figure 3.7: Schematic of the IBL read-out system [9]. The IBL Stave is the only on-detector component in this diagram.

The second cable in the type-1 cable harness contains 61 polyimide-clad copper wires: 32 low voltage 24 AWG wires, eight wires (24/28 AWG) for low-voltage sensing, 16 32 AWG DCS wires, four 32 AWG wires for high voltage and one 26 AWG drain wire. The cable is 3.5 m long and terminates in a custom 67-pin AXON33 connector of 21 mm diameter.

In the ID endplate area the type-1 power sub-cables connect to 9 m-long type-2 cables which end in the Patch Panel 2 (PP2) situated inside the MS volume. For the entire IBL there are four PP2 crates, two per detector side. The PP2 contains regulators for the low voltage and is passive for High Voltage (HV) and DCS. Type-3 cables, 75 m long, connect the PP2 with the HV and Low Voltage (LV) power supplies and the DCS electronics inside the counting room.

Commercial power supplies provide high and low voltage for the detector. The maximum high voltage for planar sensors is 1000 V and for 3D sensors 500 V, at a maximum current of 8 mA (planar) and 10 mA (3D). The low voltage from the power supplies is regulated down to around 2 V by the PP2 regulator boards (see Table 3.3). Four front-end chips share one HV channel, as well as one LV channel on a regulator board. Each PP2 regulator board, supplying one half-stave, is fed by a single primary voltage supply channel.

The optoboard LV supply chain replicates the frontend LV chain design from power supply to the detector-side end of the type-2 cable at which point a direct connection to the optobox is made. In addition to LV, the optoboards require a bias voltage for the PIN diodes of 5 V and a control voltage for the VCSELs of about 0.8 V (see Table 3.3). All voltages, as well as a reset signal, are provided by custom Supply and Control for the Optical Link (SC-OL) modules in the counting room. Each optoboard which reads out half a stave is powered separately.

The DCS lines coming from each half stave are used to monitor the temperature on every fourth frontend

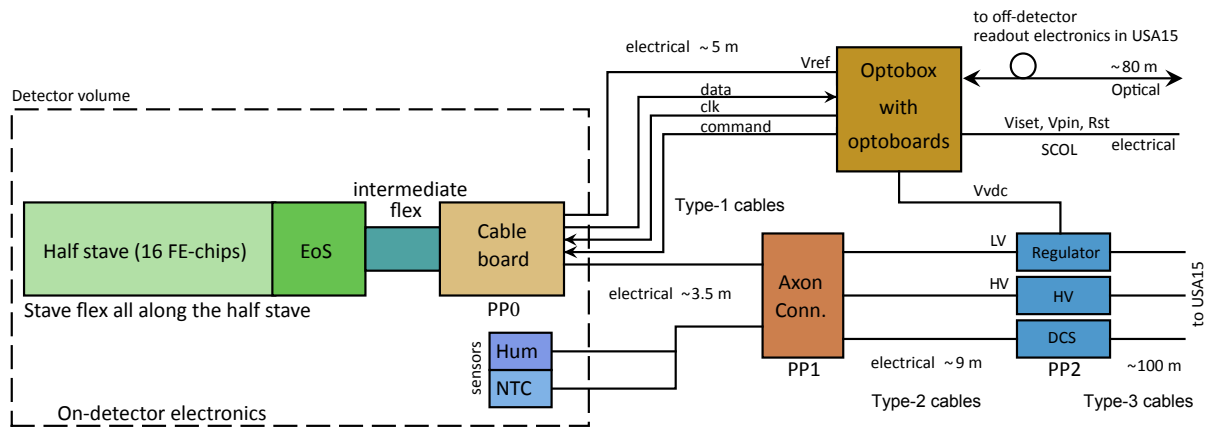


Figure 3.8: Block diagram of on-detector and off-detector electrical services for one half-stave of the IBL detector [9].

Table 3.3: Summary of the silicon sensor bias voltage and front-end electronics requirements and granularity. The different front-end electronics supply voltages are described in the text. The power consumption expected during initial operation is also shown. Following irradiation, the power consumption increases significantly in the sensors, front-end electronics and cables.

		Pixel	IBL Planar (3D)
Bias voltage supplies	Voltage (maximum)	600 V	1000 V (500 V)
	Voltage (start of Run 2)	250 V	80 V (20 V)
	Voltage (end of Run 2)	400 V	400 V (40 V)
	Current (maximum)	4 mA	8 mA (10 mA)
	Segmentation	one per module	one per 4 FE-I4B
Front-end electronics low-voltage supplies	Voltages	1.7 to 2.1 V analogue 2.1 to 2.5 V digital 0.8, 2.5 and 5 V opto-device	1.2 to 2.2 V feeding both analogue and digital 0.8, 2.5 and 5 V opto-device
	Current	3.7 kA	209 A
	Segmentation of bulk supply	One per 6 to 7 modules	One per 4 FE-I4B
	Segmentation of regulated supplies	One regulator per module	One regulator per 4 FE-I4B
Power	Front-end electronics power	~ 6 kW	~ 0.3 kW
	Cables plus regulators	~ 18 kW	~ 0.8 kW
	Total power (start of Run 2)	~ 24 kW	~ 1.1 kW
	Total power (end of Run 2)	~ 32 kW	~ 1.6 kW

chip, on the cable board, on the cooling pipe, and for a fraction of the staves on the type-1 bundle while for the rest of the staves this temperature sensor is replaced by a humidity sensor on the cable board. The temperature monitoring uses NTC sensors. The humidity sensors were not radiation hard; for this reason they were only used during the installation phase. In the counting room the DCS lines are connected to a custom crate that processes the signals.

Table 3.3 summarises the power requirements for Pixel and IBL. The power consumption of the IBL was about 1.6 kW for the low voltage at the end of Run 2, while for the high voltage it was about 15 W. The power consumption is expected to increase further towards the end of life due to a rise in leakage current (HV) and LV current caused by radiation damage.

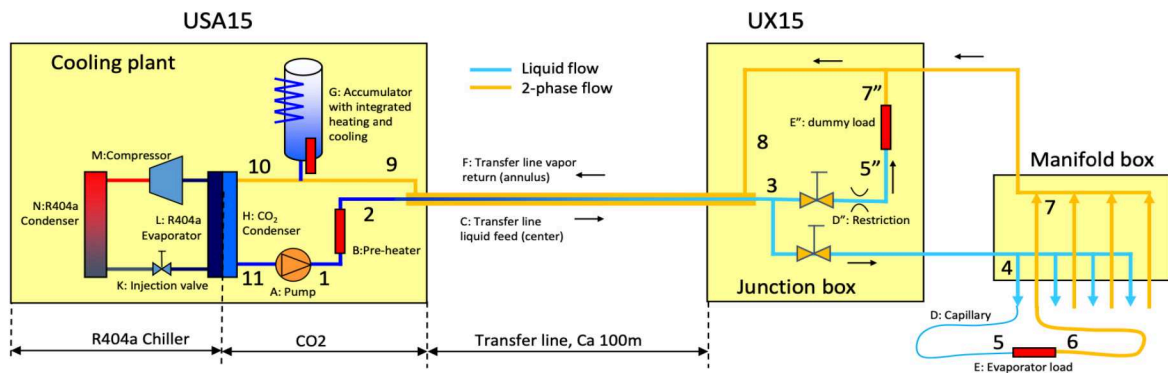


Figure 3.10: Schematic representation of the IBL CO₂ cooling system [9]. The numbers indicate the loop direction.

The IBL-specific DCS is organised in three layers: the lowest level implements the control at a hardware based layer and provides the necessary setup and configuration. The middle level maps the hardware channels to the detector units. The smallest units that can be controlled independently are groups of four front-ends, or a DCS module, and the optoboards each serving four DCS modules. At the top level, detector oriented control and monitoring is provided by a Finite State Machine (FSM) hierarchy implemented in a framework provided by the CERN Joint Controls Project (JCOP) and ATLAS groups. Thanks to this tree structure, shifters can easily monitor the detector state by observing the changes of the upper level of the tree; the detector can be controlled via FSM commands by non-DCS-experts. The lowest level of the FSM hierarchy is based on modules and optoboards and follows the detector geometry of readout units and staves building the innermost Pixel layer.

For detector safety, a dedicated hardware interlock system is installed. It is fully independent from the DCS software, but its input and output signals are monitored by the DCS. The interlock system protects the detector and certain electronics against overheating by generating interlock signals from analogue temperature monitoring with a high granularity. In addition, it receives signals from DSS and distributes the corresponding interlocks, and protects the IBL from dangerous beam conditions by receiving the Stable Beams signal from the LHC and using it to provide a software-generated beam injection permit signal.

3.1.8 IBL CO₂ cooling system

The IBL is cooled by a two-phase CO₂ system which guarantees an operational temperature in the range between -35 to $+15$ °C, thanks to a total design cooling power of about 3 kW. The cooling prevents thermal runaway of the irradiated sensors. Figure 3.10 shows a simplified schematic view of the system.

The CO₂ is cooled by commercial chillers based on a hydrofluorocarbon mixture (R404A) and pumped to the experimental cavern through a concentric transfer line of a length of about 100 m. The liquid supply line is inside the two-phase return line to pre-heat the sub-cooled liquid from the plant to approximately the same temperature as the returning two-phase fluid. This also protects the supply liquid from possible ambient heating. The transfer line is vacuum insulated with a stationary vacuum shield of 63.5 mm. At the end of the transfer line the fluid is distributed to the detector staves through a manifold box (see the left side of Figure 3.11) where 14 concentric vacuum flex lines (seven for each side) bring the fluid to the detector through the same routing as the cables on the ID end plate (see the red arrows in Figure 3.11 on

the right). The flex line vacuum is obtained thanks to an active vacuum pump system installed in UX15 and monitored by a dedicated Programmable Logic Controller (PLC). Next to the manifold box there is a junction box where temperature and pressure sensors are installed to monitor the CO₂ conditions closest to the detector. Also a 3 kW heater is present in order to simulate the detector load running in bypass mode for commissioning or test of the system.

The evaporation temperature in the detector is regulated by the pressure in the accumulator (CO₂ tank providing coolant reservoir). The return pressure (saturation temperature) is controlled by heating or cooling of the two-phase mixture inside the accumulator. By controlling the accumulator it is possible to control in a very precise way the evaporation temperature in the detector and to provide the required cooling conditions and stability. The diameter of the return lines is chosen to be slightly larger than that of the IBL cooling pipes, to ensure that the pressure drop along these lines remains sufficiently low that the accumulator pressure directly determines the evaporator pressure.




Figure 3.11: (a) IBL cooling lines and (b) the ID end-plate region, with the path of the cooling lines highlighted.

In USA15, two redundant cooling units are installed and running, one providing cooling to the detector with the other running cold in stand-by mode through a by-pass. The purpose of the stand-by mode is to keep the pump of the inactive system cold to allow for a fast start in case it is required. One unit is running with normal electrical power while the other is also supplied by Uninterruptible Power Supply (UPS) batteries: in case of a failure of the normal power, there is an automatic swap between the two units which is transparent for the detector, and only a very small and short temperature variation is visible. The logic of installing two redundant units allows uninterrupted operation in case of failure or servicing of one line. The two cooling units are controlled by two independent PLCs.

Since the start of operation, the IBL cooling has operated at three different temperatures. In May 2015 the IBL was operated at -10°C , until the start of the study of the high current drawn by the readout chips discussed in Section 3.1.3, when the cooling system was run at the much higher temperature of $+15^{\circ}\text{C}$ until June 2016. From then on, the operational set point of the plant was reduced to -20°C .

The IBL cooling system has been shown to operate stably over the full range between room temperature and its lowest verified limit, which is -35°C . Its reliability was demonstrated during external events like power cuts, glitches and primary cooling failures. No downtime occurred during data taking from the

beginning of Run 2 and, on the occasion of the only system failure, the stand-by unit kept the detector in the required conditions.

The temperature stability, needed to limit the detector bowing effect described in Section 3.1.4, was better than 0.2 °C (with less than 0.05 °C RMS) and made it possible to correct for bowing through alignment corrections.

3.2 The ATLAS Pixel upgrade

3.2.1 Installation of new Service Quarter Panels

The installation of the nSQPs during LS1 enabled the relocation of the optoboards outside the ID volume to an accessible area, making future repairs possible without extraction of the Pixel detector. All defects originating from broken data transmission lines and faulty optoboards were therefore repaired during LS1. Additionally, all faulty connections outside the active Pixel detector volume were repaired during the process of reconnection after the nSQPs installation. Faulty connections within the active volume were not accessible and thus could not be repaired.

The full detector package was removed, taken to the surface and tested there, before being re-installed in ATLAS in December 2013. The refurbished three-layer Pixel detector was then reconnected and tested. The number of disabled modules was decreased to 33, 1.9% of the total. The biggest improvement was achieved in the *B*-Layer, where the disabled fraction was reduced from 6.3% to 1.4%, and Layer 2, where the 7.0% faulty fraction was reduced to 1.9%.

The nSQPs included additional data fibres dedicated to the readout of Layer 1 (and part of the Discs); this eventually resulted in doubling the bandwidth of the transmitted data once new IBL ROD/BOC card pairs were deployed at a later stage of the project (see Section 3.2.2). For Layer 1 and part of the Discs, the readout upgrade increased the bandwidth from 80 to 160 Mbps, and for Layer 2 from 40 to 80 Mbps. This increase of the bandwidth by a factor two would allow the detector to run without bandwidth limitation up to an instantaneous luminosity of $\mathcal{L} = 3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. Further details of the readout upgrade are given in the Section 3.2.2.

Because of the easier maintenance afforded by the installation of the nSQPs, another optoboard replacement campaign was conducted during LS2. The detector status at the start of Run 3 was as follows: 10 modules (3.5%) disabled in the *B*-layer, 14 modules (2.8%) disabled in Layer 1, 32 modules (4.7%) disabled in Layer 2, and 10 modules (3.5%) disabled in the Discs. In addition, three IBL front ends were disabled.

3.2.2 Readout upgrades during the winter shutdowns

New readout components developed for the IBL (see Section 3.1.5 for details) were used to upgrade the off-detector readout hardware of the outer Pixel layers, making use of the new fibre installation provided with the nSQPs. The nSQP project was an opportunity to increase the bandwidth for some of the Pixel layers and remove some limitations that would have arisen with the expected increase of luminosity. The strategy was to gradually upgrade the various layers during three consecutive winter shutdowns between 2015 and 2018. The timeline of these upgrades was as follows:

2015/16 Layer 2 + 12.5% of Layer 1, a total of **22 ROD/BOC** card pairs

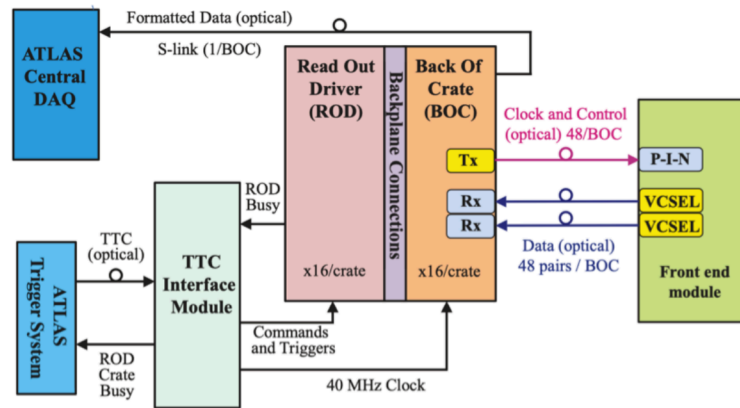


Figure 3.12: Schematic of the SCT Data Acquisition System.

2016/17 Remaining part of Layer 1 (87.5%), a total of **42 ROD/BOC** card pairs

2017/18 *B*-Layer and Disc, a total of **38 ROD/BOC** card pairs

By the end of Run 2, in 2018, both Pixel and IBL finally shared the same type of readout system, making it easier to operate the detector and offering more powerful debugging capabilities. No further changes to the off-detector readout system are foreseen during Run 3 operation.

3.3 Semiconductor Tracker readout upgrade

The SCT data acquisition system in Run 1 comprised 90 ROD and BOC cards. The BOC provides the optical interface between a ROD and up to 48 SCT modules; for each module the BOC transmits the clock and trigger via a single command stream and receives data back via two optical links from the six ASICs on each side of the module. The ROD processes incoming data from the 96 data links at the L1 trigger rate, and combines those data into a single event fragment which is broadcast on a single optical S-LINK to the ATLAS DAQ. Figure 3.12 shows a schematic of the SCT DAQ.

The SCT was designed to operate with 0.2 to 0.5% occupancy from its 6.3 million sampled strips with an LHC pileup of up to about $\mu = 23$ *pp* collisions per BC. The fundamental bottlenecks arising from higher pileup that restrict increasing occupancy are the bandwidth limitations of the data links from the modules (which transmit bit streams at 40 Mbps) and the S-LINKs (which transmit 32-bit words at 40 MHz, a throughput of 1.28 Mbps). Extrapolations of occupancy with increasing pileup during Run 1 suggested that the data flow in the data links and S-LINKs would exceed bandwidth limitations at pileups of around $\mu = 87$ and 33 *pp* collisions per BC respectively, at a L1 trigger rate of 100 kHz. These projections used the optimum on-chip data compression mode (where only hits exceeding a threshold from the in-time BC are read out, while simultaneously vetoing hits from the preceding BC), and the standard data packing schema of the ROD which included readout of the hits from each strip over three consecutive BC.

A number of mitigation steps were deployed to prepare the SCT for the anticipated higher pileup during Run 2. Furthermore, it was discovered early in Run 2 that SCT hit occupancy had increased significantly compared to Run 1 data, which is now attributed to secondary interactions with the newly-installed IBL services. The increased occupancy led to a revision of the pileup limits for SCT operation, leading to




Figure 3.13: Number of S-LINKs that would be operating above the bandwidth occupancy threshold as a function of that threshold. The 2012 configuration was the original 90-ROD system, the 2015 configuration used 128 RODs. Neither of these configurations was suitable for Run 3 conditions; however, the 2016 configuration incorporated supercondensed mode and the fibre reshuffling, and this allows all the S-LINKs to operate below the bandwidth occupancy threshold. A further minor cable remapping was performed in 2019, improving matters even more.

further mitigation steps during Run 2 in order to achieve a comfortable margin for operation with $\mu \sim 60$. These mitigation steps were:

Expanded DAQ To address the pileup limit imposed by the S-LINKs, the numbers of RODs and BOCs were increased in 2014, during LS1, thereby reducing the number of modules serviced by each ROD from 48 to 36. During LS2 the DAQ system was expanded from 90 to 128 cards (the full complement of cards possible in each ROD crate), increasing the total number of S-LINKs from 90 to 128.

Cable remapping In 2015, the cabling of the data links to the RODs was rearranged to harmonize the data load on each ROD. This resulted in a flatter occupancy distribution across the S-LINKs, removing spikes from the highest occupancy S-LINKs and thereby improving the pileup limit imposed by the S-LINKs.

Data compression A new highly efficient data packing scheme with compression on the ROD (“supercondensed mode”) was developed during LS1 and was deployed routinely from 2016. The new scheme resulted in a roughly 25% data size reduction, at the expense of losing the 3-bit timing information.

Chip Masking From 2017, the noisiest chips were masked to reduce data throughput on the data links. This worked by using a reference table of the noisiest chips as a function of pileup, and masking the 128 channels of those chips. The mechanism worked “on the fly”, gradually reducing to zero the number of masked chips as instantaneous luminosity decreased during a run. Masking was flagged in the data to avoid an arbitrary loss in hit efficiency. Typically, even at the highest pileup of $\langle \mu \rangle \sim 60$, well under 1% of chips were masked.

Figure 3.13 shows the impact of the different mitigation steps related to S-LINKs, showing the number of S-LINKs which operate above a defined fraction of the available bandwidth (occupancy threshold) as a function of that threshold. The plot illustrates, for example, that no S-LINKs used more than about 90% of the bandwidth from 2016.

Following these mitigation steps, pileup limits at 100 kHz trigger rate are around $\langle\mu\rangle = 70$ for both the data links and the S-LINKs, which is sufficient for Run 2 and Run 3 operation. Hot chip masking is not deployed routinely, but remains an option for future operation.

3.4 Transition Radiation Tracker upgrades

3.4.1 TRT DAQ upgrades in Run 2

The TRT [35, 47, 48] DAQ consists of custom analogue/digital front-end electronics and the back-end RODs [49]. An analogue component of the front-end, the Amplifier, Shaper, Discriminator and BaseLine Restorer (ASDBLR), detects signals from two drift straws using two different thresholds: a tracking threshold of about 200 eV, the Low Threshold (LT), and the transition radiation threshold of about 6 keV, the High Threshold (HT). During each 25 ns bunch-crossing, the digital component of the electronics, the Drift Time Measurement Read-Out Chip (DTMROC), samples the output of the ASDBLRs eight times, creating one bit for each 3.125 ns indicating whether the straw signal exceeds the low threshold, and one additional bit to record whether the signal crosses the high threshold at any time during the 25 ns. Upon receiving a L1A signal, the data from three bunch-crossings are sent to the RODs on the detector back-end. These 27 bits of information constitute a “straw word”, which is the fundamental block of data from a readout and compression perspective. Each ROD validates and formats the data from up to 120 DTMROCs, building a “data fragment” by serialising the straw words. The fragment is then compressed with an entropy-based Huffman algorithm [50] implemented in the ROD firmware and is sent to the ATLAS central ROS system via an S-LINK.

With this setup, the TRT DAQ system was capable of operating at L1 input rates up to 80 kHz at the average pileup of $\langle\mu\rangle = 35$ in Run 1; however, as the LHC exceeded its initial design luminosity, the goal for the final years of Run 2 (2017-18) and for Run 3 became running at an L1 input rate of 100 kHz with pileup up to $\langle\mu\rangle = 60 - 70$ pp collisions per BC. This pileup corresponds to an average per-event straw hit occupancy of 75 – 85% in the high- z slices of the TRT endcaps, which have the highest occupancies due to their high- $|\eta|$ positions. Under these more strenuous conditions, bottlenecks were observed when reading data both into and out of the RODs. Together, these bandwidth saturations limited the system to a L1 rate of 90 kHz and a maximum straw hit occupancy of 50%, for which the following DAQ upgrades were implemented between the end of Run 1 and early Run 2 (2015–16).

To alleviate the ROD input bandwidth saturation, the four trailing bits in the 27-bit straw word were truncated off in the DTMROCs. This corresponds to discarding hits with the longest drift time where the ionizing particles barely cross the edge of the straws, which are rather rare and are less important to tracking. This truncation reduces the straw word size sent to the RODs by 15% with no impact on tracking performance, enabling the TRT to run without saturating the ROD input bandwidth at up to 104 kHz with evenly spaced triggers, and up to 102 kHz in physics data taking where the rate can instantaneously exceed 104 kHz.

To cope with the S-LINK bandwidth saturation when reading data out of the RODs at high rates, a series of ROD firmware updates and hardware upgrades were performed. The firmware updates include the introduction of a validity gate to remove words leaving no LT hits within the 18.75 to 56.25 ns time window. Such hits are most likely to originate from a neighbouring bunch-crossing, and only about 3% of the hits from the bunch of interest are lost to this validity gate. Similarly, the HT bits from the first and third bunch-crossings were masked to zero to allow more efficient compression, which keeps 94% of HT hits

from the bunch of interest. Additionally, non-vital information was suppressed in the error blocks attached to the ROD data fragments. An auxiliary word-caching hash table was also implemented to relieve the heavy use of the Network Search Engine ASICs on the RODs [49] during the Huffman compression. Finally, the Huffman compression table was updated to maintain optimal compression performance in the harsher run conditions of late Run 2 and Run 3. Together, these firmware updates resulted in a 40% reduction of the event fragment size sent off the RODs, without significant loss of performance and operational stability. On the hardware side, the High-speed Optical Link for ATLAS (HOLA) S-LINK interface cards [51] used to transfer data from the TRT RODs to the ROSs were upgraded with a faster clock crystal (60 MHz vs 40 MHz) and a faster FPGA (2.5 Gbps vs 2.0 Gbps). This created 50% additional bandwidth for the S-LINKs.

Figure 3.14 summarises the effect of the updated compression table and HOLA cards. The projected maximum L1 rate at which the TRT can read out its data fragments without saturating the S-LINK is shown as a function of the detector occupancy. These limits are calculated from the average Huffman codeword length for the hit straws, the number of empty straws, and the S-LINK bandwidth. Only the high- z endcap RODs are used in this calculation, as they have the highest occupancies and are closest to the bandwidth limit. The bands represent the uncertainty in the fragment size due to the variable length of each fragment's error block, which can range from 0 to 60 32-bit words; the upper (lower) edge of the shaded area corresponds to the minimum (maximum) error block lengths. Under the final Run 2 conditions, the error block sizes were typically around 15 words or less. After all the upgrades, the system can handle occupancies up to 76-86% at the target L1 rate of 100 kHz, depending on the error block size.




Figure 3.14: Projected maximum achievable L1 input rate as a function of TRT straw occupancy (dotted orange) before and (blue) after increasing the S-LINK bandwidth and reoptimising the compression table. Only the high- z endcap RODs are considered in the calculation, as they experience the highest occupancies. The bands represent the uncertainty due to the variable size of the error blocks in the event fragments; the upper (lower) edge of the area corresponds to the minimum (maximum) error block. Two data points are shown for the ROD closest to the S-LINK bandwidth limit. These data points are taken from the beginning of a 2018 run with $\langle\mu\rangle = 57$, resulting in an occupancy of 76% in this ROD. The red cross marks the L1 rate at which this ROD's data was collected, while the black cross shows the maximum rate at which this data could have been collected before saturating the S-LINK, extrapolated based on the average word length after the compression.

3.4.2 TRT Gas Configuration Update

By the end of Run 1, several leaks developed in flexible active gas exhaust pipes made of PolyEther Ether Ketone (PEEK). The leaks originated from cracks occurring in places of local stress on the PEEK pipes due to a reaction with ozone produced in the active gas during TRT operation. An attempt to repair the leaking pipes during LS1 in 2013 was partially successful for the TRT endcaps. However, the leaks in the TRT barrel and some of those in the TRT endcaps are located in inaccessible areas and therefore repairs are not possible. Since the number of leaks is expected to grow with increasing LHC luminosity, it became unaffordable to operate the entire detector with the baseline xenon-based gas mixture. In order to understand the TRT performance with a lower-cost argon-based gas mixture, dedicated studies were performed during the proton-lead collision running at the end of Run 1. In that study, leaking modules in the barrel and endcaps were operated with the argon-based gas mixture. The study found that the transition radiation based Particle Identification (PID) performance of the detector was significantly reduced but that the tracking properties were preserved [52].

During Run 2 many leaking modules were changed to the argon-based gas mixture in stages. In 2015 all modules of the barrel inner layer and one wheel from each endcap were operated with the argon-based mixture. Starting in 2016 additional leaks led to a configuration where the two inner barrel module layers were supplied with the argon-based gas mixture and only the outer module layer was operated with the xenon-based gas mixture. This configuration still left half of the straws in the TRT barrel volume with full transition radiation information. In the endcaps, two wheels on side A (4 and 6) and three wheels on side C (3, 6, 9) were supplied with the argon-based gas mixture, as can be seen in Figure 3.15(a). In this configuration, the PID properties of the TRT endcaps were not significantly affected.

For Run 3 it is planned to use the argon-based mixture for the entire barrel and for a few more endcap wheels on the C side, as illustrated in Figure 3.15(b). This will significantly reduce the loss of xenon-based gas mixture while maintaining a stable gas configuration hopefully for the entire Run 3. In the Run 3 gas configuration the PID performance of the endcaps is largely preserved. In the barrel the PID function is significantly reduced because of a poor absorption of transition radiation photons by the argon gas but in a combination with dE/dx measurements it still contributes to the ATLAS electron identification particularly at particle energies below 10 GeV [52].

3.5 New ID thermosiphon system

From 2008 to 2017, the ATLAS ID was cooled using an oil-free compressor plant. This system requires maintenance every year and has a suction pressure limitation of 1 bar. It was observed that below this limit the failure rate of one of the compressors increased drastically. In parallel, between 2012 and 2016, a thermosiphon system was developed using the height difference between the surface and the experimental cavern (see Figure 3.16 for the general layout), with the advantage of having all active elements on the surface and therefore easily accessible for maintenance. This system is composed of four major units (see Figure 3.17 for the full circuit schematic):

Water circuit providing cold water to the condenser of the chiller first stage, sourced from the ATLAS cooling towers at approximately 25 °C.

Chiller circuit using a two-stage vapour compression cycle to cool down a brine (C_6F_{14}) to -70 °C; the chiller operates in a cascade of two refrigerant fluids, using R404A at the first stage and R23 for the second one.

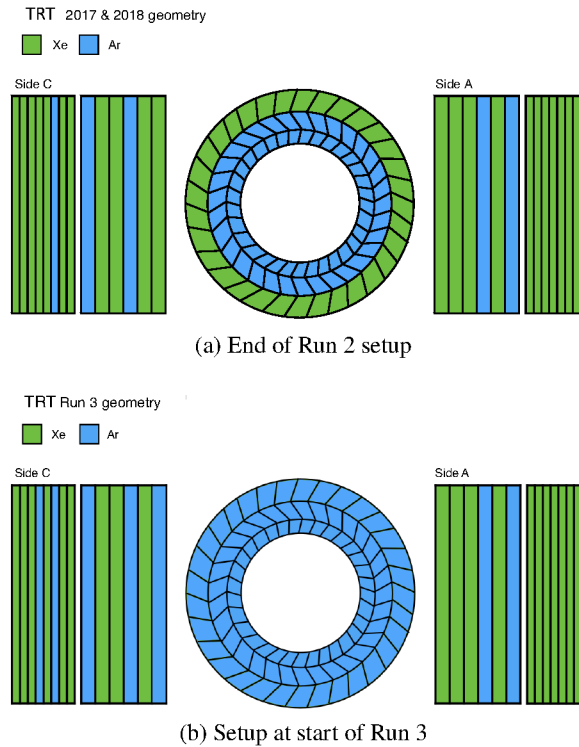


Figure 3.15: TRT gas geometries (a) at the end of Run 2 and (b) at the start of Run 3. Blue represents TRT sections supplied with the Ar-based gas mixture and green represents the sections supplied with the Xe-based gas mixture.

Brine circuit using a C_6F_{14} closed loop to condense the C_3F_8 of the primary circuit through heat exchange across the tubes of the thermosiphon condenser. This loop is able to run at about 40 kg s^{-1} to provide enough flux to condense the 1.2 kg s^{-1} of C_3F_8 corresponding to the 60 kW of heat dissipated by the ATLAS inner detector.

Thermosiphon primary circuit condensing the C_3F_8 at the surface of ATLAS to produce a natural hydrostatic liquid column over the 92 m of height difference from the surface to the cavern. The liquid evaporates in the (unchanged) on-detector cooling channels and returns to the surface as vapour by differential pressure. The thermosiphon circuit was developed to provide the same performance as the existing compressor plant.

The compressor plant and the thermosiphon are interconnected; since 2018 the thermosiphon is used as the main cooling system and the compressor plant as a back-up in case of failure. The two systems run with the same pressure parameters, allowing for a fully transparent swap from one to the other.

3.6 Updated material description of the ID

Obtaining an accurate description of the material is essential to understand the performance of the detector. The IBL insertion during LS1 and the modifications to the layout of the cables and support structures of the existing pixel detector (nSQP upgrade) required a new analysis of the ID material in Run 2.

Three complementary techniques were applied to measure the material in the ID:




Figure 3.16: General layout of the ID thermosiphon system distributed between the ATLAS Surface area, the ATLAS service cavern (USA15) and the ATLAS detector cavern (UX15). Colour-coded labels identify the main elements of the system, which are highlighted in the drawing. The 92 m column of coolant is contained in a pipe extending the full depth of the PX15 service shaft.

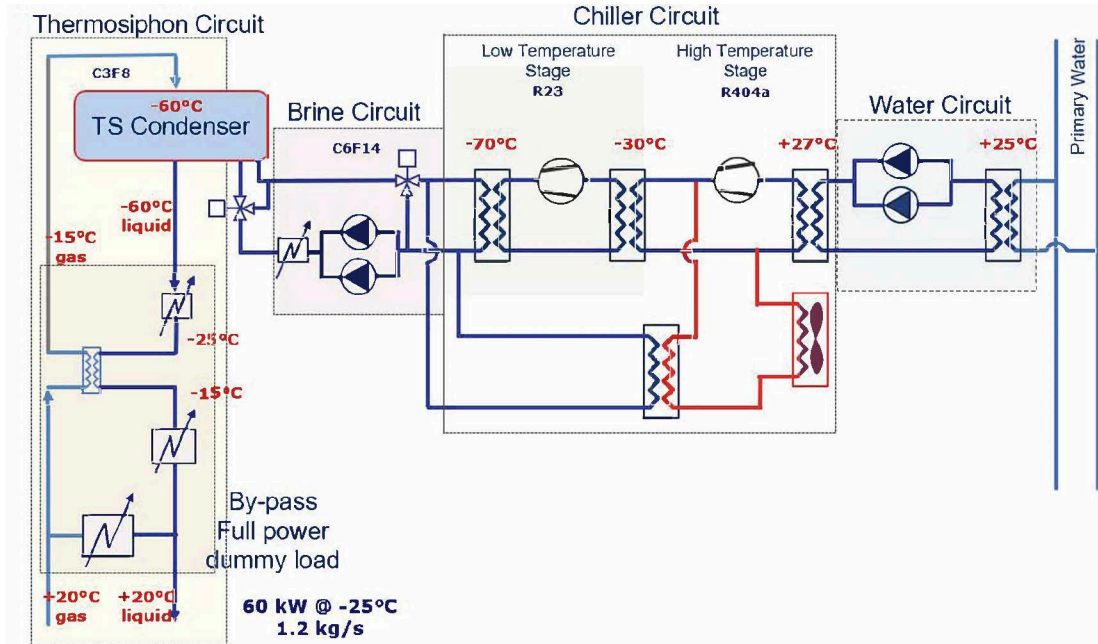


Figure 3.17: Full schematic of the thermosiphon circuit showing, from left to right, the primary thermosiphon circuit using the full depth of the shaft to cool through the natural pressure differential, the brine circuit used to condense the coolant in the primary circuit, the chiller circuit that cools the brine, and the water circuit that provides cool water from the ATLAS cooling towers to the condenser of the chiller first stage.

- **photon conversion vertex reconstruction**, taking advantage of the precise theoretical understanding of the electromagnetic interaction processes.
- **hadronic interaction vertex reconstruction**, sensitive to the material through nuclear interactions and offering much better resolution in the radial position of the vertex compared to the photon conversion. However, the description of hadronic interactions is quite complex and only phenomenologically modelled in the simulation.
- **track-extension efficiency method**, complementary approach applicable to the full tracking acceptance to measure the nuclear interaction rate of charged hadrons through hadronic interactions by matching track segments reconstructed in the Pixel detector with tracks that are also reconstructed in the SCT and TRT detectors, where the unmatched Pixel segments are assumed to be associated with charged hadron that interact hadronically while traversing the region between the Pixel and SCT detectors.

All were studied in a dedicated set of analyses using a low-luminosity $\sqrt{s} = 13$ TeV pp collision sample corresponding to around 2.0 nb^{-1} collected in 2015 [53].

While the first two methods probe the barrel region of the inner detector, in particular the new detector components installed before Run 2 (the beampipe, the IBL and the supporting tubes of IPT and Inner Support Tube (IST)), the track-extension efficiency method is also sensitive in the endcap regions of $1.0 < |\eta| < 2.5$ where most of the refurbished pixel services reside.

The precision of each measurement varies depending on the detector region. All of these approaches are used together to measure a large part of the inner detector's volume and cross-check individual measurements. The description of the geometry model was examined in detail both in radial and longitudinal distributions of the rate of reconstructed hadronic interaction and photon conversion vertices.

In the central barrel region, a significant amount of missing material in the IBL front-end electronics for the flex bus, surface mounted devices on the front-end chips and the IPT and IST was identified in the original geometry model that was used for ATLAS Monte-Carlo (MC) simulation in 2015. Figure 3.18 compares the updated geometry model in simulation with the 2015 data, showing good agreement.

The results of these studies have been taken into account in an improved description of the material in the ATLAS inner detector simulation, resulting in a reduction in the uncertainties associated with the charged-particle reconstruction efficiency determined from simulation. The updated geometry model, which was created to resolve the above discrepancies, provides a much better description of the material in the ATLAS ID simulation and is used in analyses [40].

The beampipe is found to be very accurately described except the central region ($|z| < 40$ mm). The simulated material in the IBL within the updated geometry model is found to be consistent with that observed in data, within less than 10%, mainly due to uncertainties of the hadronic interaction and conversion measurements (see Figure 3.19).

The Pixel barrel layers are found to be described well, and the results from the analyses using the hadronic interactions and photon conversions agree within the systematic uncertainties. They confirm the results of the previous hadronic interaction analysis obtained with the Run 1 data set.

The updated geometry model provides reasonable agreement with the data in the ratio of the rate measurements of hadronic interactions and photon conversions within the uncertainties of the measurements. The measured rates of photon conversions and hadronic interactions reconstructed in data are found to agree to within 7% – 18% with those predicted by simulation, based on the updated geometry model, out

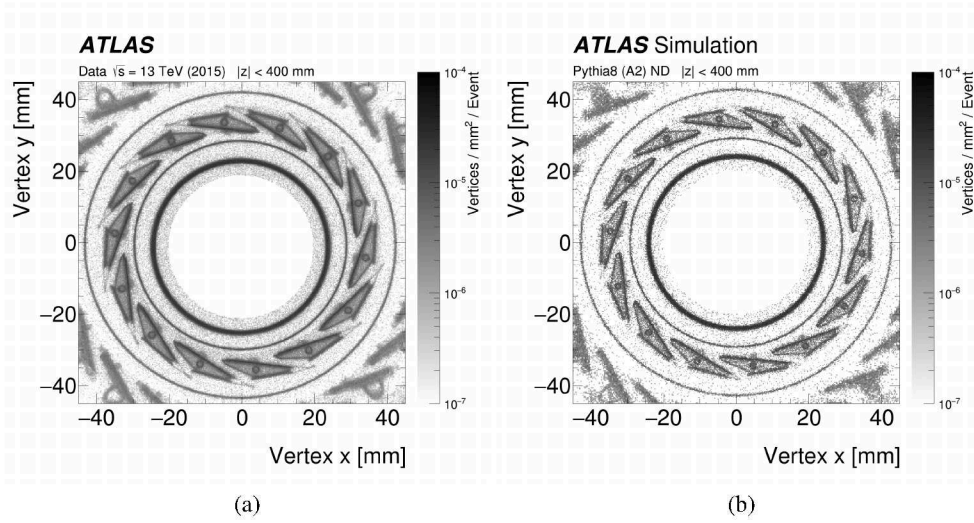


Figure 3.18: Distribution of hadronic-interaction vertex candidates with $|\eta| < 2.4$ and $|z| < 400$ mm (a) for data and (b) for the PYTHIA8 MC simulation with the updated geometry model [53].

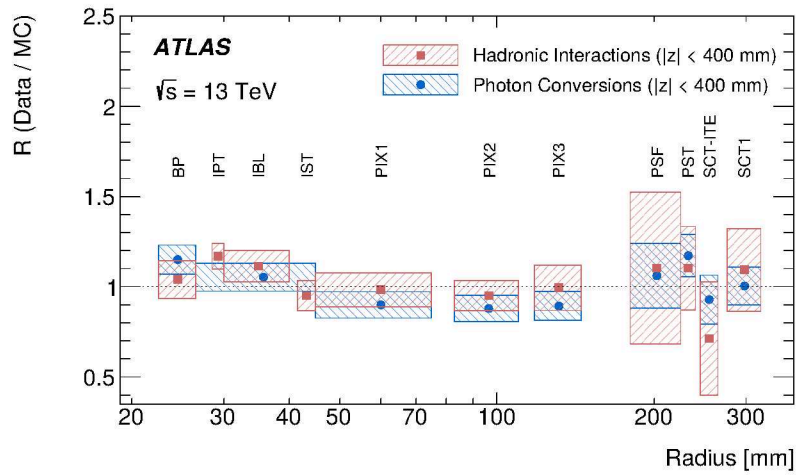


Figure 3.19: Comparison of the rate ratio, denoted R , between data and MC simulation, for hadronic interactions and photon conversions as a function of radius [53]. The horizontal range of each marker represents the radial range of vertices used in each measurement, while the vertical range represents the total uncertainty. The radial regions used for comparing data to MC simulation are the following: BP (beampipe), IPT, IBL (IBL staves), IST, PBX1-2-3 (Pixel Barrel layers 1-2-3), PSF (Pixel Support Frame), PST, SCT-ITE (SCT Inner Thermal Enclosure), SCT1 (first SCT barrel layer).

to the outer envelope of the Pixel detector. This is also supported by a study of the transverse impact parameter resolution below $p_T = 1$ GeV, where multiple scattering is dominant (see also Ref. [53]).

In the forward region, the material in the pixel service region is found to be underestimated in the geometry model by up to $\delta N_{\lambda j} = (3.7 \pm 0.9)\%$ at some values of η . This corresponds to roughly 10% of the material in the pixel services in the corresponding regions. Furthermore, in the very forward region ($3.1 < |\eta| < 5$), outside of the tracking acceptance and corresponding to the FCal acceptance, an extra contribution of material was identified. The mismatch concerns the IBL Type-1 low voltage cable; the original idea to use copper-clad aluminium wires was dropped in favour of a more robust option of all-copper wires, offering better connectability and routing flexibility, and very similar resistance, but higher material density.

3.7 Performance of the ID at the end of Run 2 and projections for Run 3

The performance of the Pixel and IBL detectors during Run 2 (discussed in Section 3.7.1) is indicative of the performance expected during Run 3, as these detectors were designed for high instantaneous luminosity and pileup conditions. The SCT and TRT required more substantial changes to cope with sustained periods of peak instantaneous luminosity; the expected impact of these changes is discussed in Sections 3.7.2 and 3.7.3.

3.7.1 Pixel and IBL detector performance

The inclusion of the IBL, adding a new point measurement with higher precision at a radius just outside the beampipe, has significantly improved the tracking performance over the whole acceptance, and this improvement is seen most clearly in the impact parameter resolution. Figure 3.20 shows this benefit in terms of impact parameter resolution by comparing two measurements performed in early Run 2 and at the end of Run 1.

The global enhancement, which is in line with expectations, can be explained in terms of proximity to the IP of the innermost layer (which drives the performance at low track p_T) and smaller pixel pitch in the z direction. Furthermore, these results demonstrate the robustness of the in-run alignment procedure that was put in place in order to mitigate the distortion of the IBL described in Section 3.1.4.

During Run 2, the pixel detector sensors accumulated a significant fluence that reached $\Phi_{\text{neq}}^{\text{Si}} \sim 10^{15} \text{ cm}^{-2}$ at the location of the IBL, as predicted by the simulations shown in Table 2.1 for the 147 fb^{-1} of pp collision data taken during Run 2. The resulting radiation damage required a careful tuning of the detector operating parameters in order to minimise its impact on the pixel performance in the reconstruction of charged particle tracks. The observed decrease of the charge collection has been countered with the increase of the bias voltage; this can be seen in Figure 3.21, where the cluster size and the dE/dx decreasing trends for increasing integrated luminosity were mitigated during the run.

The effect of this countermeasure of increasing the bias voltage can be observed in Figure 3.22, where the IBL resolution in the transverse coordinate ($r\phi$) shows a nearly constant behaviour as a function of the integrated luminosity.

In the B -Layer, the decrease of the efficiency has been minimised by suitable combinations of analogue (applied on the charge that is deposited in each pixel by its individual discriminator) and digital (or ToT) thresholds (see Figure 3.23). Different threshold schemes along the η coordinate have been applied to take

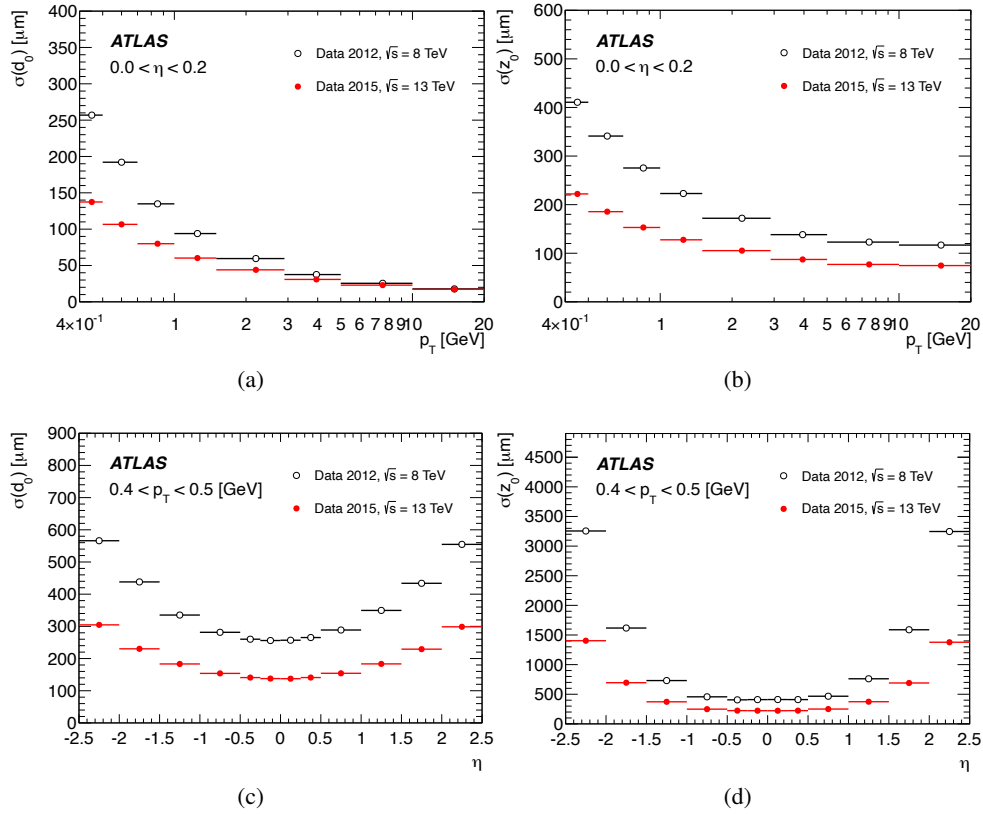


Figure 3.20: Comparison between 2012 Run 1 and 2015 Run 2 detector configuration for (left) the transverse impact parameter as functions of (a) p_T and (c) η , and (right) the longitudinal impact parameter as functions of (b) p_T and (d) η .

into account the uneven distribution of the radiation fluence and readout bandwidth occupation, leading to a hybrid threshold scenario. Despite a decrease of the collected charge of up to 35% on the central B -Layer modules, these mitigation actions resulted in no change of the pixel efficiency by the end of Run 2 (see Figure 3.23).

A similar strategy is foreseen for Run 3. The planned operating conditions in Run 3 follow from those adopted towards the end of Run 2. The bias voltage will be set at up to 500 V for the IBL planar sensors, at up to 100 V for the IBL 3D sensors and at up to 600 V for the other Pixel sensors. At these large values of the depletion voltage the Lorentz angle is small (from 250 mrad at 80 V at the beginning of Run 2 to 100 mrad at 400 V in 2018 on the IBL) but charge sharing between neighbouring pixels is still ensured by the non-zero angle of incidence of particle tracks. This ensures stable spatial resolution. The tuning to lower and, if needed, η -dependent thresholds, successfully adopted for the B -layer at the beginning of 2018, may also be extended to the IBL if the response of the detector efficiency along z does not otherwise remain uniform.

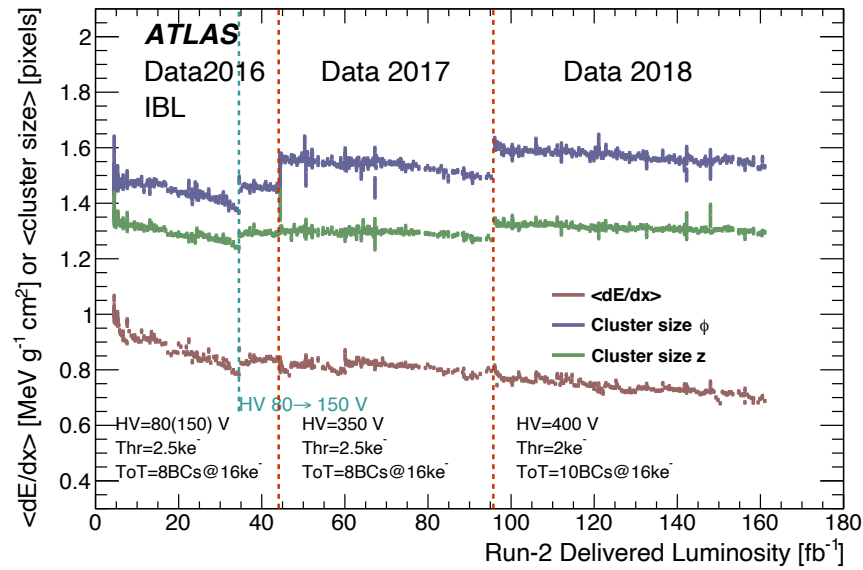


Figure 3.21: The dependence on the delivered luminosity of the average cluster size and the dE/dx measured with the IBL. Each point represents a single run, and only runs recorded in 2016, 2017 and 2018 are shown. Clusters are selected which match exactly one reconstructed charged track with $p_T > 10$ GeV and $|\eta| < 1.4$, associated to jets with $p_T > 200$ GeV by $0.1 < \Delta R(\text{track}, \text{jet}) < 0.4$. The lower cut is to reduce contamination from two particle clusters. The impact of changing the bias voltage in the IBL is clearly visible. The gradual decrease of the measured dE/dx is due to the reduced charge collection fraction due to radiation damage. Red dotted lines indicate the different data-taking years.

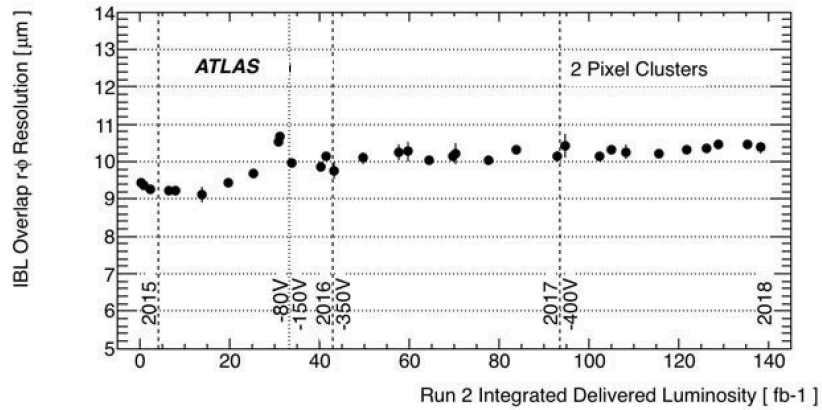


Figure 3.22: IBL spatial resolution in the transverse coordinate (small pixel pitch) as a function of the integrated luminosity in Run 2. The resolution is determined from the corrected residuals of pairs of reconstructed IBL clusters associated to tracks from charged particles traversing the detector in the region of the module overlaps. Only clusters with two pixels are considered here to remove the dependence of the resolution on the cluster shape. The dashed lines indicate the different data-taking years and the dotted lines the change of depletion voltage during LHC operation.




Figure 3.23: Efficiency for B -Layer clusters associated to a reconstructed particle track as a function of the track $|\eta|$ for different years in Run 2 and different threshold settings at the beginning of 2018. The use of a hybrid setting with lower thresholds for the central modules exposed to higher particle fluence and higher thresholds for forward modules affected by larger occupancy recovers the efficiency obtained in operations at the beginning of Run 2.

3.7.2 SCT tracking performance

In Run 2 the LHC delivered high instantaneous luminosity and pileup conditions that were far in excess of the original SCT design goals. The SCT DAQ system had to be re-optimised to mitigate bandwidth limitations as described in Section 3.3. As a result of these changes, the SCT can now provide efficient tracking with a pileup μ of up to 70 pp interactions per BC and a L1 trigger rate up to 100 kHz. During Run 2, the first significant operational impacts arising from radiation damage to the sensors and to the on-detector electronics in the SCT were observed. All SCT p^+ -on- n silicon sensors underwent type inversion, followed by a continuous increase in depletion voltage. Consequently, higher operating voltages are being applied progressively to ensure full depletion of the sensors and to maintain hit efficiency [54].

The maximum μ in Run 3 will be regulated by β levelling so as not to exceed the values seen at the end of Run 2 and the maximum trigger rate is expected to be 100 kHz [54]. Given the optimisation to address the Run 2 conditions, the SCT DAQ is expected to operate smoothly in Run 3. Radiation damage in SCT will continue to evolve in Run 3, primarily in the form of increases of the full depletion voltage (V_{FD}) and the leakage current in the sensor. Figure 3.24 shows V_{FD} and the leakage current observed in Run 2 for SCT Barrel Layer 3 (the innermost SCT layer, henceforth abbreviated to Barrel 3). The data points in the figure match well with the predictions of the Hamburg model [31]. The projection for V_{FD} and leakage current has been extended for Run 3 with the proposed luminosity delivery plan, showing the anticipated increase of V_{FD} and leakage current. Figure 3.25(a) shows measured hit efficiencies as a function of HV during Run 2 for modules close to $z = 0$ on Barrel 3 (designated $|\eta_{\text{index}}| = 1$). As this position corresponds to the maximum level of accumulated radiation, it required higher HV to maintain hit efficiency. Figure 3.25(b) shows the mean cluster width as a function of HV for tracks passing through Barrel 3 modules for $-5.0^\circ < \phi_{\text{inc}} < -4.5^\circ$, where ϕ_{inc} is the incident angle with respect to the sensor

surface. The required HV to obtain a good cluster width increased toward the end of Run 2. Figure 3.26 shows the time evolution of the incident angle with minimum cluster width, ϕ_{MCW} (a good estimator for the Lorentz angle) [54], in Barrel 3 modules: the change of the electrical field in the sensor due to the increase of V_{FD} affected the measurement. The big jumps in 2018 are due to increases of HV from 150 V to 200 V. Those measurements will be repeated periodically and if necessary HV will continue to be increased to maintain optimal tracking conditions.

3.7.3 TRT tracking performance

The expected LHC operation mode in Run 3 is to use luminosity levelling at $\mathcal{L} = 2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. Under these conditions the average hit occupancy of the TRT straws will be significantly higher than in previous runs. The track occupancy in the TRT, defined as the hit occupancy in straws in the path of a track of interest, is expected to be up to 0.75 in Run 3, while it was typically less than 0.5 during Run 2. In order to ensure optimal TRT performance, some modifications to the TRT reconstruction software were implemented. For the TRT track reconstruction these modifications have been implemented, and include a tighter requirement for hits to have a significant weight in the track fit. Figure 3.27 shows the resulting hit position measurement accuracy in the TRT straws and the relative transverse momentum resolution in the ATLAS ID for tracks with $p_T > 20 \text{ GeV}$ as a function of the TRT track occupancy for simulated $Z \rightarrow \mu\mu$ events. The dependencies in Figure 3.27 are shown for both the TRT hit weights used in Run 2 and for recalculated weights intended for use during Run 3. The new hit weights preserve excellent track reconstruction up to the highest occupancy expected in Run 3.

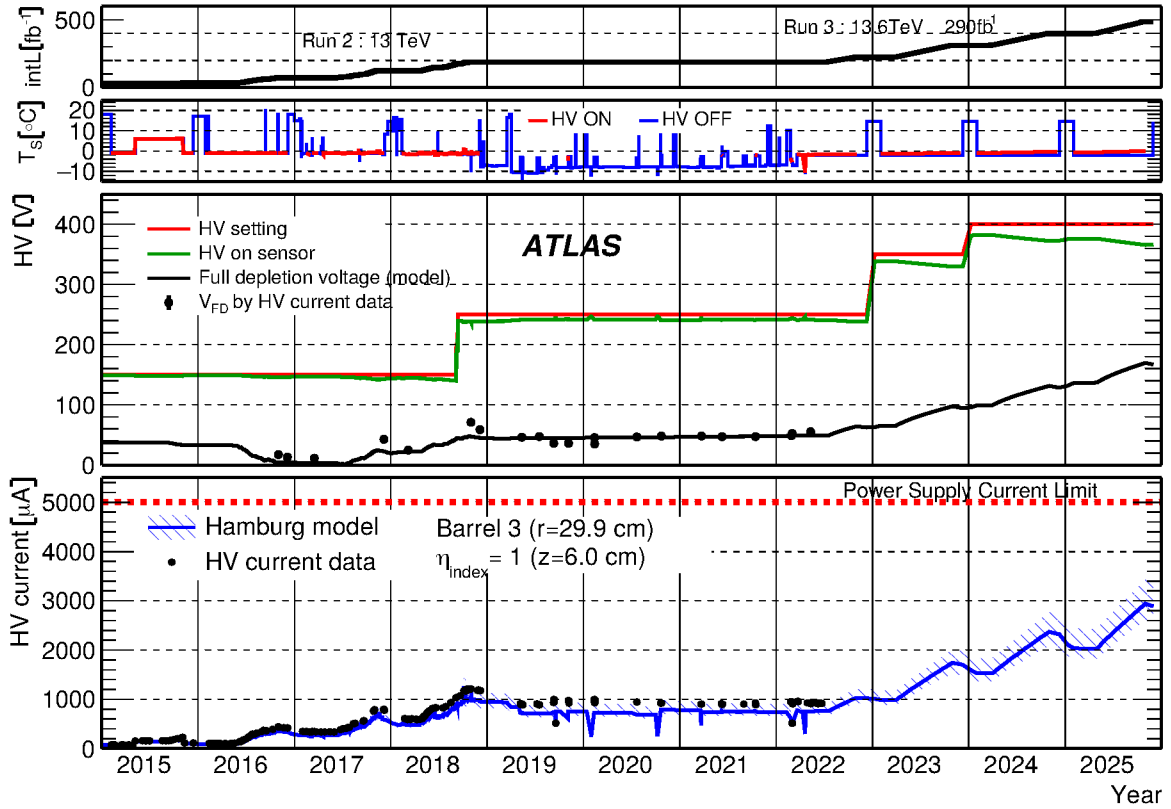


Figure 3.24: Projection of the module HV for the central Barrel 3 modules from 2015 to the end of 2025. Top plot: The integrated luminosity vs. time; in Run 3 the integrated luminosity is assumed to be 33 fb^{-1} in 2022 and 86 fb^{-1} per year in 2023-2025 at 13.6 TeV. The second plot shows the expected time profile of the sensor temperature with the HV on (red) and off (blue). The red line indicates the expected temperature rise due to bulk-heating by leakage current. The third plot shows the HV setting (red), the actual HV values on the sensor (green), and the full depletion voltage (black) estimated by the Hamburg model. The black points show measurements obtained from I-V scans, which show a kink that is an indicator for the full depletion voltage. The bottom plot shows the evolution of leakage current using the Hamburg model (blue line) compared to data (black points). Hashed areas indicated model uncertainties which do not include (unknown) errors of the luminosity to fluence conversion factors.

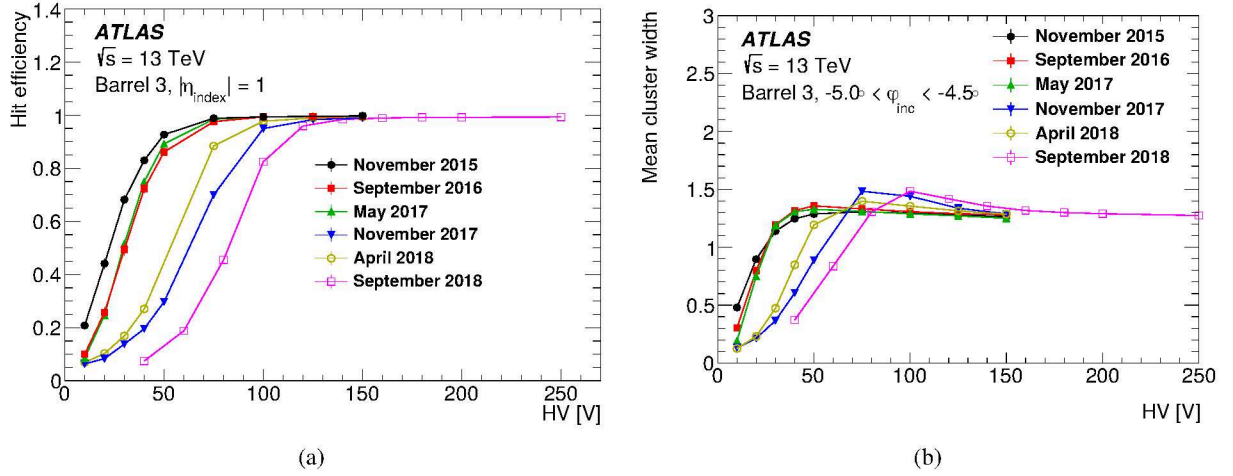


Figure 3.25: (a) Hit efficiency as a function of HV for central-most (designated $|\eta_{\text{index}}| = 1$) modules in SCT Barrel 3, measured from November 2015 to September 2018 [54]. (b) Evolution of mean cluster width as a function of HV for Barrel 3 during Run 2 [54].


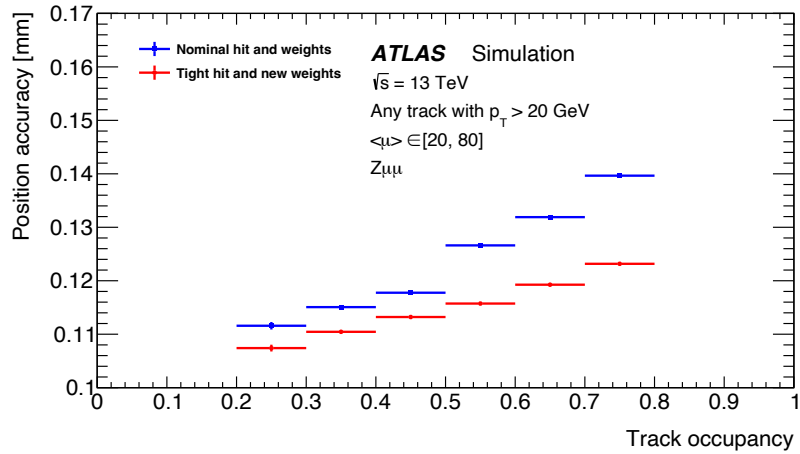



Figure 3.26: Time dependence of ϕ_{MCW} between 2015 and 2018 in Barrel 3. Only the sides with no stereo angle are shown. The error bars are statistical, while the shaded bands show the systematic uncertainties. The nominal operational HV was raised from 150 to 250 V at the time in 2018 indicated by the vertical dashed line [54]. $\langle 111 \rangle$ and $\langle 100 \rangle$ are Miller indices, indicating the crystal lattice orientation of the silicon wafers.



(a) Position accuracy



(b) p_T resolution

Figure 3.27: (a) The TRT hit position measurement accuracy and (b) the relative transverse momentum resolution of the ID for tracks with $p_T > 20 \text{ GeV}$ in simulated $Z \rightarrow \mu\mu$ events as a function of TRT track occupancy for both the Run 2 hit weights (“nominal”) and the expected Run 3 TRT hit weights (“new”).

4 Calorimeters

ATLAS uses two sampling calorimeter technologies: Liquid Argon [12] for the electromagnetic calorimeters and all the endcap and forward calorimeters, and scintillating Tiles [13] for hadron calorimetry in the central region. The ATLAS calorimeters [14] were designed to last for the entire lifetime of the LHC, and the detectors themselves require very few modifications to run at higher luminosity. This section contains a discussion of the changes in the Liquid Argon (Section 4.1) and Tile (Section 4.2) Calorimeters for Run 3.

4.1 Liquid Argon Calorimeters

The Liquid Argon (LAr) system consists of several subsystems, namely the LAr Electromagnetic Barrel Calorimeter (EMB), the LAr Electromagnetic Endcap Calorimeter (EMEC), the LAr Hadronic Endcap Calorimeter (HEC), and the LAr Forward Calorimeter (FCal). Using the full granularity of the calorimeters, read out with precision readout electronics, the LAr Calorimeter system measures the energy of electrons, photons, τ leptons, and jets as they are slowed by the dense calorimeter material, and contributes to the identification of these physics objects. The system also participates in the calculation of missing transverse energy (E_T^{miss}). In addition, it provides lower-granularity information to the L1 calorimeter trigger system in order to select events potentially containing electrons, photons, τ leptons, jets or E_T^{miss} .

During Runs 1 and 2, these signals to the trigger system for most of the calorimeter consisted of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ Calorimeter Trigger Towers (TTs), groups of elementary calorimeter cells for which the readout signals are analogue sums of the signals in the longitudinal layers of the calorimeters. For Run 3 and beyond, a new digital trigger readout path was implemented [55, 56], increasing the granularity of the summed signals by up to a factor of ten: Super Cells group together elementary calorimeter cells to provide information for the energy deposition in each individual electromagnetic calorimeter layer and with finer granularity for the front and middle layers for $|\eta| < 2.5$. For the HEC, the granularity of the signals for the digital trigger is the same as the Runs 1 and 2 trigger towers. For the FCal, where the division in η - ϕ trigger cells is less regular, the increase in granularity depends on the layer. With the increased granularity of the new trigger signals, the efficiency in selecting events with interesting signatures and the discrimination power against jets are expected to improve. The legacy (Runs 1 and 2) analogue trigger path [55] will be kept operating in parallel to the new digital trigger path at least until the new system has been fully commissioned with the first collision data in Run 3 and its performance is equal to or exceeds the performance of the old trigger system.

The modifications and improvements to the original LAr calorimeter system that are not directly related to the new digital trigger path are described in Section 4.1.1. The implementation of the new digital trigger path is then presented in Section 4.1.2.

4.1.1 Modifications to the Original LAr System

The readout of the LAr Calorimeter remains largely unchanged; an independent digital trigger path has been added which operates in parallel to the legacy electronics. An updated schematic block diagram of the upgraded LAr readout electronics architecture for Run 3 is shown in Figure 4.1. The legacy electronics include the precision main readout, the calibration system, and the analogue trigger path

providing lower-granularity energy sums to the Level-1 Calorimeter (L1Calo) system. As can be seen on the diagram, the new digital trigger can be configured and read out independently of the legacy system.

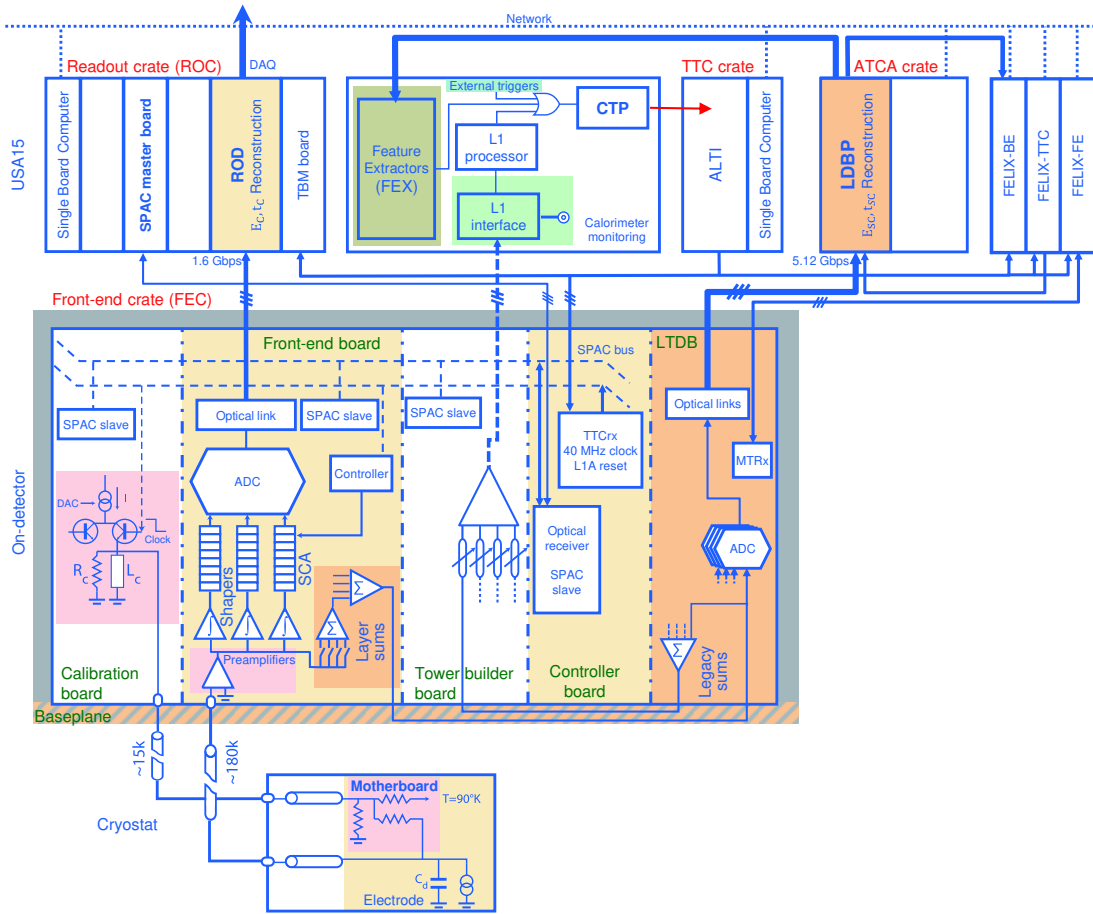


Figure 4.1: The schematic block diagram of the Run 3 LAr readout electronics architecture. In this representation, the LAr ionisation signal proceeds upwards, through the front-end crates mounted on the detector to the off-detector electronics. The elements added during LS2 are highlighted in orange. This diagram is valid for the electromagnetic calorimeters, where only the very minimal boards shown at the bottom of the diagram are located inside the cryostats; the HEC and FCal have additional electronics inside the endcap cryostats.

In the main readout, the Front End Boards (FEBs) installed in Front End Crates (FECs) on the detector shape, amplify, and sample the LAr ionisation signals in three overlapping gain scales at 40 MHz. Upon a L1A, the FEBs typically digitise four samples corresponding to the signal of the triggered collision and transmit them to the off-detector RODs which employ Digital Signal Processors (DSPs) to calculate the energy deposited in each cell and the peaking time of the ionisation signal using the Optimal Filtering (OF) technique [57]. For every L1A, the calculated energies for each calorimeter cell are sent to the ROS for eventual integration with the ATLAS event. Under certain configurable conditions (such as for cells above a given (programmable) energy threshold), the peaking time, a pulse quality factor, and/or the raw Analog-to-Digital Converter (ADC) samples are also transmitted.

The calibration system relies on the front-end calibration boards, to inject pulses of known amplitude and of known shape directly into the signal path, close to the calorimeter cells. The calibration pulse

then propagates to the electronics through the same path as the signals from the ionisation of the LAr, from which a prediction of the physics pulse shape, required for the calculation of the Optimal Filtering Coefficients (OFCs), and electronics calibration constants are calculated. Calibration campaigns are typically performed several times per week during periods between LHC physics fills and the constants are updated periodically to ensure calorimeter response stability better than 1%. As an example, Figure 4.2 shows the evolution of the pedestal measurements in calibration runs over the course of Run 2, for the electromagnetic calorimeter and for the highest gain scale.

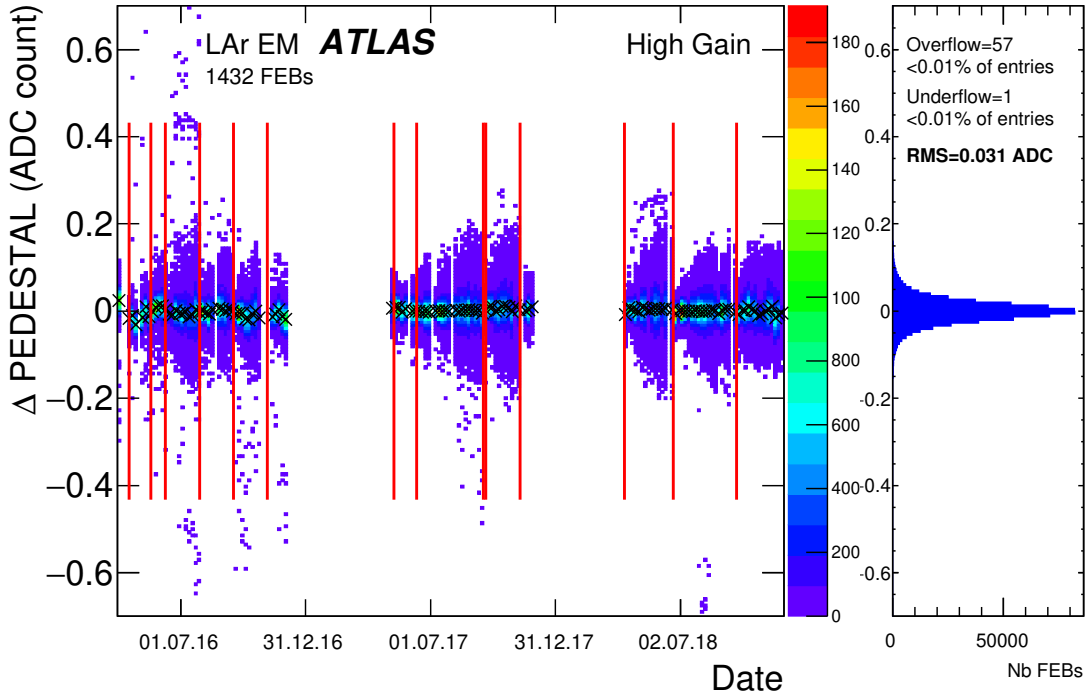


Figure 4.2: The absolute deviation of the pedestal per-FEB average from running reference values as a function of time for the LAr EM calorimeters in High Gain. The dates on which the running reference values were obtained are marked with red lines. The right panel is a projection plot of the same data; the RMS values shown include all entries. The overflow and underflow entries reported are due to a few problematic FEBs which were replaced.

Finally, the analogue trigger path begins with the Trigger Tower Builder Boards (TBBs), also installed in the FECs, which create fixed-size projective energy sums for each TT by adding together the signal sums from each of the up to four LAr calorimeter layers. The layer sum signals are aligned in time using delays applied to each TBB input. The TT sums are routed with twisted pair copper cables to receivers in the off-detector electronics cavern and are digitised and processed by dedicated L1Calo trigger electronics.

The LAr calorimeters operated very reliably during Runs 1 and 2 with a data-taking efficiency ranging from 99.2% in 2011 [58] to better than 99.7% in 2018 [5]. The improved performance can be partly attributed to optimised monitoring and recovery procedures. Examples include the more efficient identification and vetoing of events with large-scale coherent noise (noise bursts) [5, 58] and the online identification of noisy calorimeter cells contributing to high L1 trigger rates and their automatic removal from the TT energy sums. Improvements of some components in the LAr system have also contributed to the performance. One such improvement was the upgrade of the original high-voltage units with more robust high-voltage




Figure 4.3: Photographs of on-detector electronics components for the new LAr digital trigger path: a new baseplane (left) of the type installed in the LAr calorimeter front-end crates during LS2 to accommodate the installation of the LTDB (right).

modules during LS2. The new modules can operate in current-control (current-limiting) mode avoiding high-voltage trips and the needed recovery time.

An additional modification to the original LAr system was the replacement of the modules for TTC, discussed in more detail in Section 7.5.1. The legacy off-detector TTC distribution system consisted of several VME modules chained together with flat-ribbon cable connections to receive the TTC signals from the ATLAS Central Trigger Processor (CTP) and propagate them to the LAr front-end and off-detector electronics. The original TTC chain was replaced during LS2 with the new ATLAS Local Trigger Interface (ALTI), a single-board VME module developed by ATLAS. The ALTI modules are less susceptible to electromagnetic interference and eliminate multiple potential points of failure in the TTC chain.

Finally, during the course of the installation of the new digital trigger path, the FEC water cooling system was refurbished. In particular, it had been observed that the material of the cooling hoses circulating water to the FEBs and to the FEC Low Voltage Power Supplies (LVPSs) was ageing. The cooling hoses were therefore replaced, to mitigate the risk of leaks during Run 3.

4.1.2 New digital trigger path

Several new hardware components are required both on- and off-detector to implement the new digital trigger path: a new trigger front-end board, the LAr Trigger Digitizer Board (LTDB) was designed and constructed to transmit the higher-granularity trigger signals off the detector, where they are read out and processed by a completely new readout system. The detailed Super Cell granularity is presented in Table 4.1. Some modifications to the on-detector electronics in the FECs and off-detector electronics were also required. These modifications and new components, all of which were installed during LS2, are briefly described below.

4.1.2.1 New Baseplanes To support the new LTDBs, the baseplanes in all the FECs had to be replaced. The purpose of the baseplanes is to support the boards in the FEC and to deliver the signals arriving

Table 4.1: Size of the Super Cells in the LAr Calorimeter digital trigger path, in terms of both elementary cells and $\Delta\eta$ and $\Delta\phi$. The $|\eta|$ ranges correspond to the Super Cell granularity changes. The numbers of elementary cells in a Super Cell is given by n_η and n_ϕ . The notation (0.05)0.025 stands for a Super Cell composed of 1 cell of $\Delta\eta = 0.05$ and 3 cells of $\Delta\eta = 0.025$. In the HEC, the Super Cell granularity is the same as in the legacy system; the layers are summed. The FCal modules are built with a non-pointing $x - y$ geometry; therefore, the Super Cell geometry is somewhat irregular in shape and size and only approximate constant $\eta - \phi$ regions can be defined.

$ \eta $ -range	Layer	Elementary cell	Super Cell	
		$\Delta\eta \times \Delta\phi$	$n_\eta \times n_\phi$	$\Delta\eta \times \Delta\phi$
EM Barrel				
0–1.4	Presampler	0.025×0.1	4×1	0.1×0.1
	Front	0.003125×0.1	8×1	0.025×0.1
	Middle	0.025×0.025	1×4	0.025×0.1
	Back	0.05×0.025	2×4	0.1×0.1
1.4–1.5	Presampler	0.025×0.1	4×1	0.1×0.1
	Front	0.025×0.025	1×4	0.025×0.1
	Middle	0.075×0.025	1×4	0.075×0.1
EMEC: Outer Wheel				
1.375–1.5	Front	$(0.05)0.025 \times 0.1$	4×1	0.125×0.1
	Middle	$(0.05)0.025 \times 0.025$	1×4	$(0.05)0.025 \times 0.1$
1.5–1.8	Presampler	0.025×0.1	4×1	0.1×0.1
	Front	0.003125×0.1	8×1	0.025×0.1
	Middle	0.025×0.025	1×4	0.025×0.1
	Back	0.05×0.025	2×4	0.1×0.1
1.8–2.0	Front	0.004167×0.1	4×1	0.0167×0.1
	Middle	0.025×0.025	1×4	0.025×0.1
	Back	0.05×0.025	2×4	0.1×0.1
2.0–2.4	Front	0.00625×0.1	4×1	0.025×0.1
	Middle	0.025×0.025	1×4	0.025×0.1
	Back	0.05×0.025	2×4	0.1×0.1
2.4–2.5	Front	0.025×0.1	4×1	0.1×0.1
	Middle	0.025×0.025	1×4	0.025×0.1
	Back	0.05×0.025	2×4	0.1×0.1
EMEC: Inner Wheel				
2.5–3.1	Front	0.1×0.1	2×2	0.2×0.2
	Middle	0.1×0.1	2×2	0.2×0.2
3.1–3.2	Front	0.1×0.1	1×2	0.1×0.2
	Middle	0.1×0.1	1×2	0.1×0.2
HEC				
1.5–2.5	Summed	0.1×0.1	1×1	0.1×0.1
2.5–3.2	Summed	0.2×0.2	1×1	0.2×0.2
FCal				
3.1–3.5	0	x,y-various	various	$\approx 0.1 \times 0.4$
3.5–4.0	0	x,y-various	various	$\approx 0.1-0.15 \times 0.4$
4.0–4.9	0	x,y-various	various	$\approx 0.15-0.2 \times 0.4$
3.1–4.9	1	x,y-various	various	$0.1-0.3 \times 0.4$
3.1–4.9	2	x,y-various	various	$0.4-0.5 \times 0.4$

from the detector at the back of the baseplane to the readout FEBs. In addition, the baseplanes route the analogue sums of the readout signals provided by the FEBs and the LTDBs to the appropriate boards. The new baseplanes (Figure 4.3 left), have slots to receive the LTDBs in addition to the slots in which the legacy boards are seated. The intricate, multilayered design of the new baseplanes enables boards with the same footprint as the original baseplanes to provide the many additional traces required to route the finer granularity analogue sums to the LTDBs and subsequently, in most cases, after an additional sum by the LTDB (recreating the legacy analogue sums) to the legacy Trigger TBBs. Analogue sums that have the same granularity as the legacy sums (typically the back layer sums, the presampler sums and in some regions first layer sums) are delivered to the TBBs independently of the LTDBs, while the rest of the sums are not available if the LTDB is not installed. In total, 114 new baseplanes of six different types are needed to equip all the FECs of the LAr Calorimeter system.

4.1.2.2 Modifications of the FEB with new Layer Sum Boards The FEBs perform the first stage of analogue summing of the signals to reduce the granularity used in the trigger system. During Runs 1 and 2, this first sum was performed by a daughter board, the Layer Summing Board (LSB), which summed the signals from all cells in the same calorimeter layer belonging to the same trigger tower. A series of new LSBs was produced in order to provide sums of the signals from cells in the same Super Cell, thereby increasing the granularity of the sums by up to a factor of four. All 1524 FEBs installed on the detector were removed and refurbished with the new LSBs. In addition, their cooling plates and cooling hoses were replaced and tested thoroughly for leaks. After the baseplanes were replaced, the FEBs were reinstalled in the FECs together with the new LTDBs and the rest of the legacy boards.

4.1.2.3 LAr Trigger Digitizer Board The LTDB transmits LAr pulse samples for Super Cells in four layers at 40 MHz to the off-detector electronics. The LTDB is also responsible for the second stage of analogue summing in order to recreate the legacy trigger sums and provide them to the legacy TBBs through the baseplane. The board is designed so that the summing can be performed independently of the configuration state of its digital part. To cover the entire LAr calorimeter system, 124 boards of seven different types are needed. The assembled printed circuit board for one type of LTDB is shown in Figure 4.3. On the LTDB, the pulse for each Super Cell is sampled every 25 ns and the samples are digitised with a custom-designed 12-bit, 4-channel ADC. The ADC is implemented in a 130 nm CMOS technology and achieves an Effective Number Of Bits (ENOB) of 11 with a dynamic range of 11.7 bits while consuming less than 50 mW per channel. Each LTDB can handle up to 320 Super Cells, so 80 ADC chips are installed on each LTDB. The digitised samples are serialised and transmitted with custom-designed optoelectronics with ASICs implemented in a 250 nm Silicon-on-Sapphire (SoS) process. Up to 40 optical fibres operating at 5.12 Gbps transfer the samples to the off-detector electronics.

4.1.2.4 New off-detector electronics The configuration, monitoring, and readout of the LTDB are performed with new off-detector electronics installed in Underground Service Area USA15 in 2020. A diagram of the architecture of the new off-detector electronics, operated independently of the legacy electronics, is shown in Figure 4.4. The configuration and monitoring of the LTDBs is performed using the Front-End Link eXchange (FELIX) system over optical links interfaced to Gigabit Transceiver Slow Control Adapter (GBT-SCA) devices (see Section 7.7.2 for further details). For the readout, the LAr Digital Processing System (LDPS) comprises 30 LAr Digital Processing Blades (LDPBs) implemented in ATCA technology and installed in three shelves.

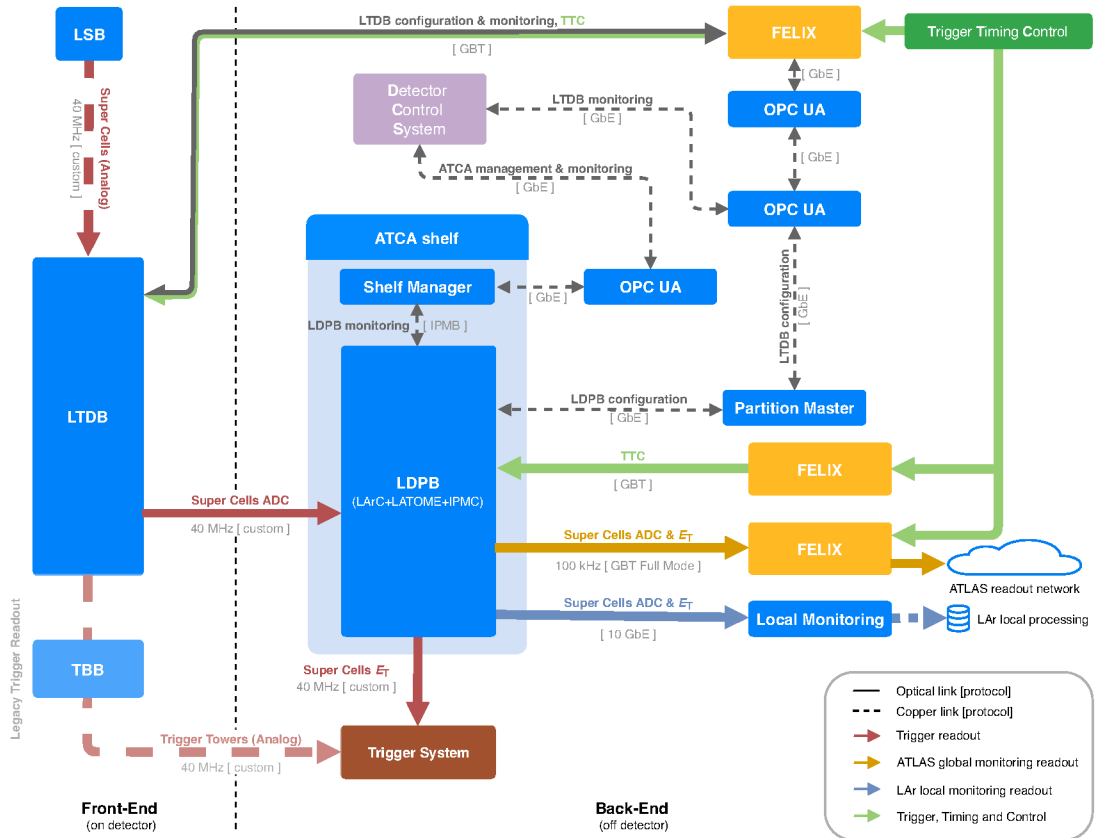


Figure 4.4: The LAr calorimeter digital trigger readout system installed during LS2.

4.1.2.5 LAr Digital Processing Blades Each LDPB consists of a LArC motherboard that houses up to four LAr Trigger Processing Mezzanine (LATOME) Advanced Mezzanine Card (AMC) daughterboards as shown in Figure 4.5. Up to 48 input fibres can be connected to each LATOME routing the digitised pulse samples from the LTDBs at 5.12 Gbps, receiving 12-bit ADC samples from up to 320 Super Cells at 40 MHz. The samples are processed on the LATOME in real time by custom-designed firmware running on an Arria10 FPGA. An OF procedure, similar to the one used for the precision main-readout, is used to calculate the energy deposited in each Super Cell at 40 MHz for each bunch crossing. The resulting calibrated transverse energy values for each Super Cell (E_T^{SC}) are routed to the L1Calo Feature EXtractor (FEX) processors (see Section 7.2) directly from the LATOMEs over 48 output fibres each operating at 11.2 Gbps. In some cases, several copies of these values are transmitted to the different FEXs. In total, the LDPS receives 25.2 Tbps of input from the front-end and streams up to 41.1 Tbps to the L1 system. The LArC facilitates the control, configuration, and monitoring of the blade and the hosted LATOMEs; it connects to the ATCA shelf, distributes the power to the various elements of the LDPB, and provides a connection to the ATCA shelf manager via an Intelligent Platform Management Bus (IPMB) for the low-level monitoring of the status of the blade, including voltage and temperature measurements. Additional direct Gigabit Ethernet (GbE) connections allow the reprogramming, loading of configuration constants, and further monitoring of the LDPB. The hardware control and monitoring of each LArC blade proceeds via an Intelligent Platform Management Controller (IPMC). The LDPB also implements two additional paths for reading out processed data from the LATOMEs: the “TDAQ path” transmits on L1A

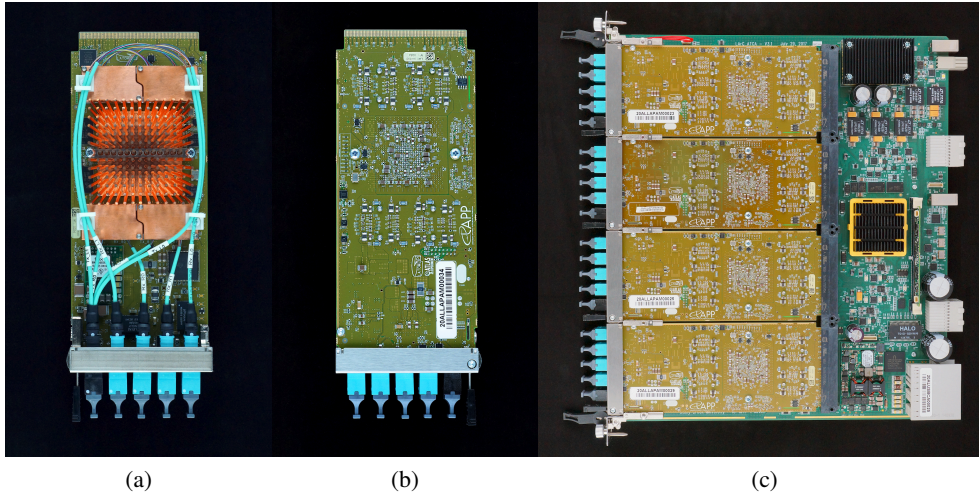


Figure 4.5: Photographs of off-detector electronics components for the new LAr digital trigger path. (a) and (b) show the front and back faces of a LAr Trigger Processing Mezzanine (LATOME), respectively, while (c) shows a LAr Carrier (LArC) fully equipped with four LATOMEs.

(therefore with a rate of 100 kHz) for each Super Cell the E_T^{SC} and the ADC values of the five time slices used in its calculation. The data are transmitted over FULL mode links (see Section 7.6.1) to FELIX to be integrated with the ATLAS event for recording. The additional data throughput to TDAQ is of the order of 400 Gbps and can be used for debugging the trigger decision, comparisons to the main readout, and calorimeter noise studies. The second additional data path implemented on the LArC is the “monitoring path” which provides an independent data flow of order of a few 10 Gbps for local monitoring of the data, transmitted to a local computer farm over 10 GbE links using the ATCA infrastructure. Finally, the LArC receives and distributes the ATLAS TTC signals to the LATOMEs.

Further technical details of the new LAr digital trigger system including its integration and commissioning, as well as the procedures used to validate its proper functioning and the performance achieved during the commissioning of the system, are available in Ref. [56].

4.2 Tile Calorimeter

The ATLAS Tile Calorimeter (TileCal) covers the region $|\eta| < 1.7$ using pseudo-projective calorimeter towers composed of scintillating tiles in a steel matrix, read out by wavelength-shifting fibres. The TileCal comprises a cylindrical Barrel section that surrounds the LAr EMB cryostat and two Extended Barrels that surround the LAr endcap cryostats. The cryostat scintillation counters sit between the central and endcap cryostats. They are mounted to the TileCal Extended Barrel, and are read out through the TileCal electronic system. They are used for luminosity monitoring and to correct for energy losses in the outer wall of the barrel cryostat and ID services in this region where there is a large amount of passive material. The counters present in Runs 1 and 2 covered the region $1.2 < |\eta| < 1.6$, while for Run 3 this coverage has been extended to $1.2 < |\eta| < 1.72$.

The Minimum Bias Trigger Scintillators (MBTS) cover $2.0 < |\eta| < 4.0$ and are used for triggering and for luminosity monitoring. As they are read out through the TileCal electronic system, they have been

historically considered as part of the TileCal.

The only upgrades in Run 3 for the TileCal are in these two systems: the cryostat and MBTS scintillation counters. The upgrades have been motivated primarily by the radiation damage that affected these systems in Run 2. Additionally, there has been development of a TileCal electronics demonstrator in preparation for Run 4, which is also described in this section.

4.2.1 Radiation environment


The $1.2 < |\eta| < 4.0$ region in the ATLAS detector is a high-radiation environment for plastic scintillators. At a distance of about $z = \pm 3500$ mm from the IP, the average ionisation dose in the pseudorapidity region from 3.0–4.0 is of the order of 830 Gy/fb^{-1} while the section from pseudorapidity 2.0–3.0 sees an average dose of the order of 95 Gy/fb^{-1} , with the dose in each rapidity interval dropping rapidly with increasing distance from the beam-line (for example from 2700 Gy/fb^{-1} at $|\eta| = 4$, to 213 Gy/fb^{-1} at $|\eta| = 3$). For a projected total luminosity in Run 3 of 250 fb^{-1} , the dose is expected to be of the order of 200 kGy for $3 < |\eta| < 4$, and 20 kGy for $2 < |\eta| < 3$. The dose at a pseudorapidity near 1.7 (E4) is of the order of 10 Gy/fb^{-1} , or a total of 2.5 kGy for 250 fb^{-1} , again dropping rapidly with increasing radius. Significant light loss can be expected in plastic scintillators for radiation doses greater than 5 kGy [59, 60]. To give an indication of the impact of the light loss in MBTS, at the start of Run 2 in 2015, the two-track efficiency was 98%; in 2018 when the light loss reached 95% in the outer counters, this two-track efficiency dropped to 75%.

4.2.2 Cryostat counters


In the region occupied by the cryostat counters, the fraction of passive material along the particle path is high (as shown in Figure 4.6), primarily due to the cryostat walls. In the region of maximum passive material, there is up to 10 radiation lengths of material in front of the cryostat counters. The electromagnetic showers are insufficiently sampled in this region, with the first active layer located close to the shower maximum. Using the cryostat counters to correct for the energy loss in the passive material by adding the weighted energy deposit in the scintillators leads to a partial recovery of the electron and photon energy resolution in the affected areas [61].

In Runs 1 and 2, each counter was divided into two segments, covering $1.2 < |\eta| < 1.4$ (E3) and $1.4 < |\eta| < 1.6$ (E4). However, a significant energy resolution degradation is also present in the η region 1.6–1.75. Before the start of Run 2, two of the 128 cryostat scintillators were re-built with an extension to cover the η range from 1.6–1.75. The energy resolution in the region covered by this extension was significantly improved with the corrections provided by the energy deposited in this scintillator. The replacement of the cryostat scintillators for Run 3 allows for an extension of the $|\eta|$ coverage to 1.72; an extension to 1.75 is not possible due to interference with the liquid argon cryostat heaters in some ϕ regions. The number of electronics channels for the cryostat scintillators is fixed, so the segmentation for E3 was modified to an η range of 1.2–1.6, and E4 to 1.6–1.72. Simulations indicated that the energy resolution in the η range from 1.2–1.6 would not be significantly degraded by joining the two segments, and that the new region 1.6–1.72 would survive the radiation expected in Run 3.

There are 64 cryostat counters corresponding to the 64 TileCal modules, each counter containing an E3 and an E4 channel. The cryostat counters are mounted at their outer radius to the extended barrel TileCal modules using specially designed brackets. Neighbouring counters are attached at the inner radius using



(a)



(b)

Figure 4.6: (a) Tile E1-E4 counter location (in yellow) and (b) amount of material (in radiation lengths) in front of the cryostat counters.

connectors designed for the purpose. During Runs 1 and 2, the counters were aligned (in ϕ) at the outer radius, where they are connected to the brackets. Due to imperfections in the circularity of the extended barrels, this led to the appearance of gaps between some of the counters at the inner radius, resulting in additional uninstrumented areas, distributed non-uniformly. The alignment scheme was redesigned for Run 3, with the alignment carried out at the inner radius (with each counter maintaining a separation of 1 mm from its neighbours), with the impact of any non-circularity taken into account in the brackets at the outer radius. The old E3/E4 counters were removed, and the new ones installed, in 2019–2020.

The cryostat counters use 6 mm-thick scintillators (EJ-208 from Eljen, a scintillator designed with an improved radiation hardness) sandwiched between two trapezoidal aluminium shells. Care was taken during installation to avoid any electrical contact between the aluminium shells and the endcap cryostat wall. The scintillating plates are wrapped in Tyvek sheets, to prevent damage to the surface of the scintillator and also to reflect back some of the light that escapes from the top/bottom of the scintillator. Opaque spacers separate the two pieces of scintillator comprising E3 and E4. For light-tightness, the counters are sealed with both black electrical tape and aluminium tape.

The EJ-208 scintillator emits light in the blue wavelength region. Y-11(200)MSJ Wavelength-Shifting (WLS) fibres from Kuraray coupled to one side of each scintillator capture the blue light and then re-emit light in the green portion of the spectrum, a fraction of which is captured and then transmitted to optical connectors glued into the tops of the counters. From there, optical cables using BCF098 clear fibres from St. Gobain transmit the light to photo-tubes in the extended barrel TileCal modules, where the signals are read out through the standard TileCal data acquisition chain. The area above the optical cables is covered with aluminium foil panels in order to improve the light-tightness. Before installation, each assembled counter was checked for light yield and uniformity using a strontium β source. Figure 4.7 shows the end surface of one of the ATLAS TileCal extended barrels and the LAr endcap cryostat, where the MBTS counters, the cryostat scintillators, and the aluminium foil panels can be seen.




Figure 4.7: The surface of one of the ATLAS TileCal extended barrels and the liquid argon endcap cryostat where the MBTS counters (protected by a white cover) and cryostat scintillation counters can be seen. Also shown are the aluminium foil panels used for light-tightness. The black tubing contains clear fibres which are used to transmit the light signals of the MBTS counters.

In summary, the E3/E4 cryostat scintillators are very useful in improving the lepton energy resolution and luminosity monitoring in the region $1.2 < |\eta| < 1.72$. New counters have been constructed for Run 3, extending the rapidity range and using the more radiation-hard scintillator described above. The Run 3 degradation is estimated to be in the range 28 – 43% (for E3–E4, respectively) after 250 fb^{-1} .

4.2.3 Minimum Bias Trigger Scintillator counters




Figure 4.8: The two scintillator segments from one of the MBTS counters present in ATLAS during Run 2. The effects of the radiation exposure in Run 2 can be seen as yellowing at the small end.

The MBTS counters cover the high-rapidity interval from 2.0–4.0 on both sides of the ATLAS detector. The MBTS counters are divided into eight ϕ segments per endcap, with each ϕ segment having two radial subdivisions in η , 2.0–3.0 and 3.0–4.0. The MBTS channels are read out through the standard TileCal data acquisition chain. In addition, a high-gain trigger signal is available from each channel, allowing the MBTS signals to be used in the trigger. Initially, the MBTS counters were intended to be used only in the early running of ATLAS in order to determine whether a particular beam crossing contained a proton-proton collision, needed in the extremely low luminosity collisions. They were not expected to be needed, nor functioning, in high luminosity conditions. However, they did maintain functionality, and proved to be useful for luminosity monitoring, van der Meer scans, and for triggering during the heavy ion running.

The high radiation environment in this rapidity region required the MBTS counters to be replaced at the end of Run 1, with the expectation that they would also need to be replaced at the end of Run 2. The rapidity interval covered by the inner MBTS counters has some of the highest radiation levels for which plastic scintillators have ever been used. At the end of Run 2, the inner section of the MBTS counters had lost 99% of the original light yield, and the outer section had lost 95% of its original light yield, so care was taken in the design for their replacement in Run 3, to improve the radiation hardness.

Light loss in polystyrene-based scintillators from radiation exposure results primarily from damage to the polystyrene matrix, resulting in the formation of radicals and other color centres, which absorb the light produced by the secondary fluors. A picture of one of the counters from Run 2 is shown in Figure 4.8, showing the coloration pattern due to the damage from radiation. The bonds can be re-formed over time (annealing). The annealing can be accelerated by the presence of oxygen. For Run 3, polystyrene-based scintillator is used for both the inner and outer segments, but different dopants were used in the two areas. In particular, the dopants paraterphenyl (PTP) (primary) and BBQ (secondary) were used in the inner region. With these dopants, the output scintillation light of the inner segment is in the green region of the spectrum, and is less sensitive to absorption by the color centres produced by radiation exposure. The outer segment used the standard PTP and POPOP ((1,4-bis-(2-(5-phenyloxazolyl))-benzene) dopants, producing blue scintillator light.

Wavelength-shifting fibres are placed in σ -shaped grooves cut into the scintillator, with the WLS fibres glued into optical connectors at the outer radius of the counter. The same type of WLS fibres are used in

the cryostat scintillators and at the outer radius of the MBTS counters. Green scintillation light needs to be shifted to orange in the WLS fibres used in the inner segment MBTS counters. This requires a different type of WLS fibre than, for example, in the cryostat scintillators or in the outer segment of the MBTS counters. O2 fibres from Kuraray, shifting the green light to orange are used. Optical cables transmit the light from the MBTS counters to the standard optical path of the TileCal readout, with the difference that red-sensitive photo-multipliers (Hamamatsu R7600-20 ERMA) are used.

For both the inner and outer MBTS regions, the scintillator is divided into four layers, each 5 mm thick (rather than the single layer of 2 cm used in Runs 1 and 2), individually wrapped with Tyvek sheets. This layered structure presents more surface area to allow for the diffusion of oxygen, accelerating the annealing. Each ϕ -segment is placed inside its own aluminium can, with the edges sealed with both black electrical and aluminium tape. The aluminium cans are bolted onto a boronated polyethylene moderator attached to the front of the endcap cryostat, as shown in Figure 4.9. The outside of the aluminium can is protected by a white cover, as seen in Figure 4.7 .




Figure 4.9: (a) An assembled Run 3 MBTS counter without the top cover. (b) Arrangement of wavelength-shifting fibres from the scintillator sandwich to the optical connector.

The new MBTS scintillator, along with the new WLS fibre readout used in the inner segment, has been shown to be less sensitive to radiation effects. Using four separate 5 mm thick sections of scintillator instead of a solid 2 cm thick piece of scintillator should allow for more oxygen diffusion into the scintillator, and thus a faster annealing of the color centres formed upon radiation exposure.

4.2.4 TileCal Demonstrator for Run 4

The readout electronics for the Tile Calorimeter for the HL-LHC running will have to be upgraded to deal with the increased radiation levels and the increased out-of-time pileup. Substantial progress towards this Phase II upgrade was already made by the start of Run 3. In order to gain experience with the new readout scheme prior to the high-luminosity running, a hybrid demonstrator module, combining the new TileCal

module read-out scheme, but still compatible with the present (to be used in Run 3) readout system, was constructed to be used for one barrel module in Run 3. The demonstrator electronics were extensively tested in test beams, and were inserted into one barrel module on the ATLAS detector in 2019. The plan is for this demonstrator to remain in place for at least part of Run 3. The goal is to carry out further studies under realistic ATLAS running conditions.

The demonstrator consists of a super-drawer that partitions a legacy TileCal drawer into four mini-drawers, each servicing up to 12 Photo-Multiplier Tube (PMT) channels. The super-drawer continuously digitises PMT signals using two gains, resulting in an effective 17-bit dynamic range, and sends the sampled data to off-detector systems at a rate of 40 MHz.

The current L1 trigger system (Run 3) is analogue, however. Compatibility with the TileCal legacy trigger system is achieved by the use of an adder-based board that groups the PMT analogue trigger signals in pseudo-projective towers, and sends the analogue sums to the legacy L1 trigger system. The in-situ tests (since 2019) have been successful and have resulted in useful information for the future installation of the full Phase II upgrade.

5 Muon Spectrometer

5.1 Overview of Muon Spectrometer Upgrades

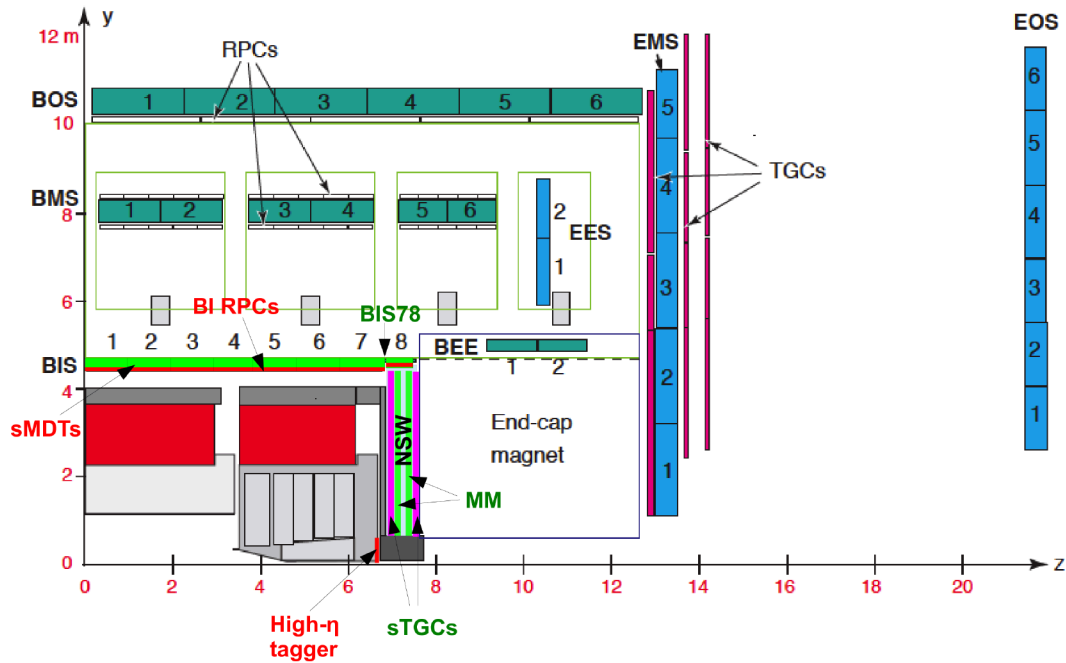
The Muon Spectrometer (MS) forms the large outer part of the ATLAS detector and detects charged particles exiting the barrel and endcap calorimeters, measuring their momentum in the pseudorapidity range $|\eta| < 2.7$. It also triggers on these particles in the region $|\eta| < 2.4$. The driving performance goal is a stand-alone transverse momentum resolution better than 15% for 1 TeV tracks. This translates into a requirement to measure an effective sagitta of about 500 μm with a resolution of about 75 μm . In the MS, this sagitta must be determined from tracks of particles passing through an inhomogeneous field, and measured at just three stations spanning a distance ranging from about 6 m in the central barrel region to more than 15 m in the endcaps.

Since the volume of the MS is very large, it is impossible to provide continuous tracking; instead, tracks are reconstructed from straight segments reconstructed at (usually) three points: an inner station close to where the tracks exit the calorimeters and upstream of the toroid magnetic field, a middle station inside or immediately downstream from the toroid field, and an outer station well outside the magnetic field of ATLAS. The detectors in each station are multilayered, and provide at least six points along the muon trajectory that can be reconstructed as a straight track segment with a well-measured orientation in space, especially in the bending direction of the toroid magnets. Tracks are then constructed by fitting curved paths to the three “pointing segments” and matching them to tracks from the ID.


The MS comprises a Barrel, consisting of three concentric, roughly cylindrical, stations (Barrel Outer: BO, Barrel Middle: BM and Barrel Inner: BI), and two Endcaps, each consisting of three discs, referred to as Wheels (Endcap Outer, Middle and Inner: EO, EM, and EI), as well as an Extended Endcap Ring (EE) positioned outside the radius of each endcap toroid cryostat. The EI wheels, also called the Small Wheels, sit between the calorimeters and the endcap toroid cryostats, inside the barrel toroids (see Figures 5.1 and 1.6). The EM wheels (or Big Wheels), are located on the far side of the endcap toroid cryostats. The EO wheels, similar in size to the EM wheels, are mounted on metal scaffolding structures attached to the end-walls of the ATLAS cavern. The EE rings provide a third measurement station in planes between the EI and EM wheels for tracks in the intermediate pseudorapidity region, $1.05 < |\eta| < 1.3$, not covered by the EO wheel. The outermost (EIL4) stations of the EI region covering the range $1.0 < |\eta| < 1.25$ are not mounted on the EI wheels; they are instead permanently fixed between the barrel toroid coils. The MS has an eight-fold symmetry: eight Small (S) sectors are aligned with the eight coils of the barrel toroid magnet, and eight Large (L) sectors cover the regions between the coils.⁸

The majority of the barrel detectors are unchanged from the original Run 1 configuration described in Ref. [15]: all three stations use multilayered Monitored Drift Tube (MDT) chambers for the precision measurements in the bending coordinate, and the Outer (BO) and Middle (BM) stations are also equipped with Resistive Plate Chambers (RPCs) for triggering and to measure the azimuthal coordinate of the tracks. Additional barrel chambers added since Run 1 are described in Section 5.3.3. In response to the increasing number of gas leaks due to cracks in the gas inlets of the RPC system that developed over time during

⁸ Sectors are numbered from 1 to 16, starting with the Large sector at the x -axis of the ATLAS coordinate system and proceeding azimuthally (thus, Sector 5 is at the top of the detector, and Sector 13 at the bottom). Modules within each sector are numbered sequentially according to their distance from the IP: thus, BIL1 chambers sit near the centre ($z = 0$) of the Barrel Inner layer in the Large sectors, while EOS6 chambers are at the outer radius ($r \sim 11$ m) of Small sectors in the Endcap Outer wheel, as shown in Figure 5.1. The only chambers at $z = 0$ are the BOL0 chambers in Sectors 12 and 14. Other sectors have a gap at $z = 0$ to allow services to exit from the calorimeters.



(a) Cut through Small sectors



(b) Cut through Large sectors

Figure 5.1: Schematic quadrant view showing locations of the main components of the MS for Run 3. (a) shows a cut through a Small sector and (b) a cut through a Large sector. The new detectors installed in LS2 are the NSW (the new EI wheel), visible on both cuts in lime green and magenta, and the BIS chambers in positions 7 and 8 on Side A, labelled in green as BIS78. The blue and dark green rectangles are MDTs. The large red rectangles represent the Tile Calorimeter. Components indicated with red labels will not be installed until Run 4 [62].

Run 2, significant work was undertaken during LS2 to reinforce the RPC gas inlets and recover a large number of channels that had become inactive. Inlets were repaired and no-return valves installed. The gas inlets on 100 service boxes (two out of the four on the corners of each of 50 chambers) were reinforced by foam injection in 2022. Preliminary results indicate that this prevents new leaks from developing, and it will be done for all the accessible service boxes during winter shutdowns over the course of Run 3.

During Runs 1 and 2, the endcap MS comprised three technologies: MDTs, as in the barrel, Cathode Strip Chambers (CSCs) in the innermost region of the EI wheel, and Thin Gap Chambers (TGCs), used for triggering and to provide the azimuthal coordinate of muons in the endcaps of the ATLAS MS [3]. The EO and EM wheels and the EE rings are unchanged from Run 1. The EO wheels contain only MDTs, and therefore measure only in the (precision) bending coordinate, while the EM wheels each comprise four distinct wheels: one of MDT chambers for precision tracking in the bending coordinate, sandwiched between one of triplet TGCs and two of doublet TGCs for triggering and for measuring the azimuthal coordinate. The 528 TGC doublet modules and 216 triplet modules in the three EM TGC Wheels of each endcap continue to play the same essential roles in the L1 muon trigger and in offline tracking as they did in Runs 1 and 2. During LS2, 24 faulty TGC modules in the EM wheels were replaced by spares, increasing the number of operational channels. The EE rings consist only of MDT detectors for tracking, and do not participate in the trigger. In the inner (EI) wheels of the endcap MS from Runs 1 and 2 [3], however, the original TGC doublets did not have sufficient resolution in the bending direction, nor enough detector layers, to form pointing segments for the L1 trigger. For the Run 3 trigger, the EI wheels are required to provide pointing segments of sufficient angular precision to test for matches with muon candidates in the EM wheels and for consistency with muon tracks originating from the IP. The original EI wheels from Runs 1 and 2 have therefore been completely replaced by the New Small Wheels (NSWs), and are the main subject of this chapter. The sixteen EIL4 assemblies, comprising eight MDT detectors and 21 TGC doublets on each side of ATLAS, fixed between the barrel toroid magnets, are retained for Run 3. EM-wheel trigger candidates without corresponding NSW segments satisfying these matching criteria can be dropped, and this is expected to reduce the false-trigger rate to an acceptable level.

The main objective of the Phase-I ATLAS upgrades [63] is to sharpen the L1 trigger threshold turn-ons and discriminate against background while maintaining the L1 rate at a manageable level. This must be achieved without pre-scaling or raising p_T thresholds for the single-muon L1 triggers, which would result in a significant loss of acceptance for many interesting physics processes.

When the design of the Phase-I upgrades began, the estimated total rate of the L1 trigger for single muons with $p_T > 20$ GeV was expected to rise above 50 kHz in Run 3 if no measures were taken, while ATLAS can allocate only 25 kHz for muon triggers out of a total L1 bandwidth of 100 kHz [64]. This scenario will be relevant for the HL-LHC but, because of luminosity-levelling, will probably not be experienced during Run 3.

The MS upgrades focus chiefly on the endcap regions. The Run 1 endcap muon trigger relied only on the TGCs of the EM wheels. There, a substantial background arises (see Figure 5.2), proportional to the instantaneous luminosity, due to relatively low-energy charged particles emerging from hadronic showers in the forward shielding, that enter the endcap toroid cryostats without passing through the EI muon wheel. The paths of these particles are then bent in the toroids, and a fraction of them enter the EM wheels on trajectories that closely mimic those of muons coming from the IP. This led to a large fake rate when only the EM wheels were used to trigger [15]. The number of Regions of Interest (RoIs) identified by the L1 single muon trigger thus increases sharply for $|\eta| > 1$, while the η distribution of reconstructed muons is almost flat. Most of the RoIs for $|\eta| > 1$ in Run 1 were due to this shielding-induced background.

The asymmetry seen in the trigger between the two ends of the detector is due to the opposite bending directions of the positively charged protons originating from the shielding interactions.




Figure 5.2: The expected muon trigger rate reduction from the NSW, BIS78 and Level-1 Muon (L1Muon) endcap upgrades; however, the BIS78 on the negative- η side, where the fake trigger rate is lower, will only be installed after Run 3.

Since the background rate in the endcap region (including the transition region between the barrel and endcaps) increases with luminosity, it will eventually exceed the allocated Level-1 Muon (L1Muon) trigger rate. Failure to suppress this background would result in the need to increase the p_T threshold of the lowest unrescaled muon trigger. The first steps to reduce the charged-particle background were taken in Run 2, when coincidences with the TGC doublets in the original EI wheels, or with the Tile Calorimeter, were required in the L1 trigger to ensure that particles causing L1Muon triggers in the endcap regions were on trajectories consistent with an origin at the IP. However, the background rates in the region of the EI stations are much higher than in the EM stations (as discussed in Section 2.6.2), so the relatively coarse granularity of the legacy EI TGCs implied that at high luminosity, nearly all EM triggers would have had corresponding background hits in the original EI wheels, rendering the coincidence requirement less effective. Moreover, the 30 mm-diameter MDTs used in the original EI wheels are limited by space charge build-up effects [65] at high background rates, and their precision tracking capability was expected to degrade at the HL-LHC. The goal of the Phase I upgrade was therefore to replace the EI wheels with detectors able to provide very fast and precise track-segment matching, while maintaining the offline tracking performance of the original EI wheels from Runs 1 and 2.

To accomplish this, the original EI wheels have been completely replaced by New Small Wheels (NSWs) occupying the same envelopes and providing tracking over the same polar angle range: $1.3 < |\eta| < 2.7$. These are described in Section 5.2. The large-radius region of the inner station $1.0 < |\eta| < 1.3$ remains covered by the EIL4 MDT and doublet TGC detectors of the original muon endcap system, which are mounted between the toroid coils and not attached to the EI wheels. The TGC doublets in EIL4 will continue to provide adequate confirmation under Run 3 conditions that a particle has traversed the endcap toroid zone, reducing the fake endcap triggers in this region; they are scheduled for replacement by triplet TGCs during LS3 to improve their efficiency in rejecting fake triggers.

The EIL4 chambers, however, cover only about 70% [62] of the full azimuthal angle in a ring surrounding

the EI wheels (now the NSWs), with gaps for the barrel toroid coils in the Small sectors. To fill these azimuthal gaps in trigger coverage, the original BIS7 and BIS8 MDT chambers (those farthest from the IP) on Side A of ATLAS were replaced with new, smaller, BIS78 chambers using small-diameter MDTs (sMDTs), leaving enough space to add new RPCs with thinner gas gaps and improved electronics (“new RPCs”) to complete the trigger coverage in the Small sectors of this transition region (described in Section 5.3.4).

The sMDT technology used in the BIS78 upgrade was first introduced in chambers used to fill several small holes in the original coverage of the MS during winter shutdowns in the course of Run 2. These upgrades are briefly described in Section 5.3.3, along with some additional (original technology) MDT chambers added during the first LHC long shutdown.

5.2 Endcap Upgrades: The New Small Wheels

The NSW detectors operate in a high-background radiation region (where detected hit rates could potentially increase to as much as 20 kHz/cm^2 in the small region closest to the beamline, at the upper-limit estimate of HL-LHC luminosity of $\mathcal{L} = 7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$), reconstructing muon tracks with high precision and furnishing information for the L1 trigger [15].

The NSWs use two chamber technologies: small-strip TGC (sTGC) detectors, and micro-mesh gaseous structure (Micromegas) detectors. Both can be used for triggering: the single BC identification capability (more than 95% of signals are collected within one BC) of the sTGCs makes them particularly suitable for the primary triggering role, and the small drift gap in the Micromegas chambers means that they are also very fast, with very little dead time. Both technologies have precision tracking capabilities, with ultimate resolutions in the bending direction of the order of $100 \mu\text{m}$. In the Micromegas detectors the resolution is due in part to their very fine strip pitch (less than 0.5 mm), while the sTGCs (with a strip pitch of 3.2 mm) rely more heavily on charge sharing. Such precision is crucial to maintain the current muon momentum resolution in the high background environment of the upgraded LHC when the full track reconstruction is run. Second-coordinate resolution (in the azimuthal direction) is greatly improved with respect to the original EI wheels, with good trigger-level second-coordinate resolution from sTGC pads and Micromegas stereo roads, and excellent resolution from the sTGC wire groups and full Micromegas stereo strips in the final offline reconstruction. The sTGC–Micromegas chamber technology combination thus forms a fully redundant detector system for tracking, both for the trigger and offline. This detector combination is designed to provide robust, fast, high-resolution performance for the HL-LHC, as well as meeting the immediate requirements of Run 3.

The demanding performance requirements are summarised in Section 5.2.1. A brief description of the layout of the NSW follows in Section 5.2.2. The Micromegas and sTGC technologies are described in Sections 5.2.4 and 5.2.5 respectively. The mechanics of the supporting disc, spokes and shielding hub are summarised in Section 5.2.3, and Section 5.2.6 covers the extensions to the endcap optical alignment system. The trigger hardware mounted on the detector (along with a broad outline of the logic of its firmware) is discussed briefly in Section 5.2.8, while the off-detector elements of the trigger are described later, in Section 7.3. Changes to the services for the EI wheel to accommodate the new detectors are described in Section 5.2.11.

5.2.1 Performance requirements for the NSW

The NSW has the same precision requirements as the original EI wheel: it must provide track segments that, in combination with measurements from the ID and the EM and EO wheels, allow the endcap MS to measure the transverse momentum (p_T) of passing muons with a precision of better than 15% for 1 TeV muons in the full pseudorapidity coverage $|\eta| < 2.7$. To satisfy this requirement, the NSW must reconstruct track segments with a position resolution of about $50\ \mu\text{m}$ in the bending plane, irrespective of background conditions, and even if some detector planes are not operational. With sixteen detector planes, of which twelve measure the precision coordinate (the radial, or bending, direction), this implies a resolution of the order of 150 to $175\ \mu\text{m}$ is required for each plane of detector strips.

The L1 trigger track segments must be reconstructed online with an angular resolution of approximately 1 mrad, matching the ultimate (Run 4) angular resolution of the EM trigger segments. The angular resolution requirement for Run 3 comes from the η -resolution of the EM-wheel TGCs, and is about 3 mrad, but new electronics for the MDT chambers of the EM Wheels in the Phase II upgrade [62] will allow the MDTs to participate in the L1 trigger at the HL-LHC, improving the resolution of the EM-wheel segments.

Segment finding efficiencies must be better than 97% for muons with p_T greater than 10 GeV to equal the performance of the original EI wheels. Efficiencies and resolutions are required not to degrade for high-momentum muons that emit δ -rays or showers.

The muon L1 trigger rate must be kept below 25 kHz and fit within the overall ATLAS L1 fixed latency budget for Run 3 of approximately $2\ \mu\text{s}$. The trigger design is also required to meet the ATLAS HL-LHC Level-0 Trigger (L0) ⁹ latency budget of $10\ \mu\text{s}$, for a single-muon L0 trigger rate expected to be around 45 kHz [66].

5.2.2 Layout of the NSW

The geometry of the NSW retains the eight-fold symmetry of most of the legacy MS. Each NSW consists of 16 sectors: eight Large, and eight Small (see Figure 5.3 and Table 5.1). The Small sectors are aligned with the barrel toroid coils, and form a plane close against the shielding disc (the new JD in Figure 5.5). The Large sectors form a second plane, slightly farther from the IP. The Large and Small sectors overlap mechanically, to avoid uninstrumented regions between the sectors; the instrumented areas of the Micromegas overlap between adjacent Large and Small sectors, but the instrumented sTGC areas of adjacent sectors do not overlap.

Both the Micromegas and the sTGC detectors are built in (mostly) trapezoidal modules with four layers of gas gaps corresponding to active detector planes. These are referred to as Quadruplets. Quadruplets of various sizes are assembled radially into Large and Small sTGC wedges and Micromegas double wedges, and these are assembled to build Large and Small sectors (as shown in Figure 5.4(a)). A sector thus comprises 16 active detector layers in total: eight from the sTGCs and eight from the Micromegas.

Each sector is built on a central spacer frame that is kinematically mounted to the spokes of the wheel (see Section 5.2.3). Each spacer frame has a wedge comprising two Micromegas quadruplets secured to each side using in-plane sliding attachments, and this structure is the Micromegas double wedge. A wedge built

⁹ During Run 3, there is no difference between the L0 and L1 trigger rates, but the readout is designed to function at the HL-LHC, where a two-level trigger with L0 running at 1 MHz with up to $10\ \mu\text{s}$ latency will be used.

Table 5.1: Modularity of the sTGC and Micromegas in the NSW. There are two Micromegas wedges and two sTGC wedges in each sector (Large or Small), and there are eight Large sectors and eight Small sectors in each wheel.

Module	Per Quadruplet	Per Wedge	Full ATLAS (2 wheels)
Large sectors (28° azimuthal)			16
Small sectors (17° azimuthal)			16
Total sectors			32
sTGC Wedges			64
sTGC Quadruplets		3	192
sTGC Gas Volumes	4	12	768
Micromegas Wedges			64
Micromegas Quadruplets		2	128
Micromegas Gas Volumes	4	8	512




Figure 5.3: New Small Wheels: (a) shows the structure: the wheel on the left is seen from the cavern end-wall, with the A-frame and shielding hub visible, along with the eight Large sectors and the eight spokes supporting them; the wheel on the right is seen from the IP, and since the new JD shielding disc is not shown, the eight Small sectors, aligned with the barrel toroid coils, are visible, along with the spokes supporting them. Some of the structural elements can be seen more clearly in Figure 5.5. (b) shows the first NSW being lowered into the ATLAS cavern; it is shown from the side with the Large sectors and the A-frame.

from three sTGC quadruplets, positioned precisely relative to each other with glued fibreglass frames, is kinematically mounted on each side of the spacer frame, sandwiching the Micromegas assembly. This arrangement maximises the distance between the sTGC wedges (see Figure 5.4(b)), which provide the primary trigger. The relative alignment of the Micromegas and sTGC quadruplet modules is monitored with an optical alignment monitoring system. This system takes readings approximately every two hours to determine the displacements, rotations, and certain deformation modes of each individual quadruplet. The extension of the optical endcap alignment monitoring system to the NSW is explained in more detail in Section 5.2.6.




Figure 5.4: (a) Structure of Small and Large NSW sectors. Each sector consists of two wedges of sTGC (green) hung kinematically on either side of a central spacer frame, to which Micromegas wedges (brown) are directly bolted on each side. Each sTGC wedge consists of three quadruplets, and each Micromegas wedge consists of two quadruplets. (b) rz -plane view of an NSW with z -scale exaggerated, showing from left to right (IP to cavern wall): the new JD in grey, the Small sectors, and finally the Large sectors. Each sector contains 16 active layers: eight of Micromegas and eight of sTGC. The sTGC wedges are shown in blue and the Micromegas wedges in brown.

5.2.3 Mechanics and Shielding

The increased mass and different geometry of the NSWs with respect to the original EI wheels of Runs 1 and 2 required an entirely new mounting structure; moreover, the increased background rates anticipated at the HL-LHC required a different shielding configuration. The shielding disc and central hub of the EI Wheels were therefore entirely replaced, and the NSWs are mounted on the new JD shielding discs. These are made of structural steel, up to 90 mm thick. The central structural axle supporting the detector wheel is the NSW shielding hub, made of an external shell of stainless steel containing several interlocking pieces of forged copper, with no interior voids. The NSW supporting structure (see Figure 5.5) provides the mechanical network from which the Large and Small Sectors of the NSW are suspended. It is built from many separate bolted and welded elements, made from either aluminium profiles and plates, or from austenitic steel. The main elements of each NSW structure are two “foot spoke” assemblies, which support the weight of the NSW, six “standard” Small-sector spoke assemblies, and eight Large-sector spoke assemblies, the latter comprising both “inner” and “outer” parts. Each spoke contains adjustable supports for an alignment bar (see Section 5.2.6) and adjustable systems for the kinematic mounts that hold the sectors. Since the radiation hub constitutes the majority of the mass of the NSW, large mechanical deformations are avoided by supporting the structure by a bridge-like assembly, of which the shielding disc forms one side, the other side consisting of a bolted A-frame connecting the two foot-spoke anchorages to the hub and to the aluminium structure.

Three kinematic mounts secure each sector. Two of them suspend the sector from the two neighbouring spokes, one allowing only rotation in the plane of the sector about a fixed axis, and the second allowing only translation in a single direction in the plane of the detector. The third kinematic mount allows only

small translations in the plane. The distribution of the kinematic mounts on the spokes is driven by the position of the maximum force applied on the fixed kinematic support.

Since the NSWs are required to move out of their running position inside the barrel toroids during maintenance periods to allow access to the calorimeters, services including power, gas and cooling water are routed through four flexible cable trays (the “flexible chains”) attached to each NSW that allow the wheel to move without being uncabled.




Figure 5.5: Support structures for the NSWs, including the new JD discs (which are mounted on the side closer to the IP).

5.2.4 Micromegas Technology

Micromegas technology was developed in the mid-1990s [67]. It permits the construction of thin planar gaseous particle detectors where the traditional planes of HV wires (which can be fragile) are replaced by a thin metallic micro-mesh. Micromegas detectors consist of a planar (drift) electrode, a gas gap of a few millimetres acting as both the conversion and the drift region, and the micro-mesh at a distance of 120 to 130 μm from the readout electrode, creating the amplification region. The Micromegas operating principle, as implemented in ATLAS, is illustrated in Figure 5.6. The HV potentials are chosen such that the electric field is a few hundred V/cm in the drift region, and 40 to 50 kV/cm in the amplification region. These are adjusted to set the drift velocity close to its local maximum and plateau, while minimising longitudinal diffusion [68] and ion back-diffusion. The field in the amplification region is adjusted, according to the gas mixture used, to get a gas gain close to 10 000. Charged particles traversing the drift space ionise the gas; the electrons liberated by the ionisation process drift towards the micro-mesh. With an electric field in the amplification region 50 to 100 times stronger than the drift field, the micro-mesh is transparent to more than 95% of the electrons. The electron avalanche takes place in the thin amplification region, immediately above the readout electrode. The drift of the electrons in the conversion gap to reach the micro-mesh is a comparatively slow process (though still very fast compared with most other detector techniques): it depends on the drift gas, drift distance, and electric field, and in the configurations used in ATLAS can take up to around 100 ns. The amplification process, however, happens in a fraction of a nanosecond, resulting

in a fast pulse of electrons on the readout strip. The ions produced in the avalanche move in the opposite direction to the electrons, back to the micro-mesh. Most of the ions are produced in the last avalanche step, close to the readout strip. It is the fast evacuation of the positive ions that particularly suits the Micromegas to operation at very high particle fluxes.




Figure 5.6: Layout and operating principle of the ATLAS NSW Micromegas detectors. The micro-mesh is integrated into the drift panel, and stretched over the grid of support pillars on the readout panel when the two panels are pressed together. The red arrows indicate trajectories of incoming muons, the large red dots show the positions of primary ionization clusters, the diffusion paths of the ionization electrons are shown in yellow, and the small red dots in the amplification gap indicate the gain stages of the avalanche. Parameters shown are for the inner quadruplets and baseline gas mixture; see Table 5.2 for additional operating parameters.

The main challenge addressed in the ATLAS Micromegas design is electrical discharging [68] in large-area detectors. Electrical discharges occur when the total number of electrons in the avalanche reaches a few tens of millions (the Raether limit [69]). These discharges may damage the detector and readout electronics, and the voltage breakdown can lead to large dead times. High detection efficiency for minimum ionising muons requires gas amplification factors of the order of 10 000, meaning that ionisation processes producing more than 1000 electrons over distances comparable to the typical lateral extent of an avalanche (a few hundred microns) carry the risk of discharging (see, for example, Ref. [70]). Such ionisation levels are easily reached by low-energy alpha particles or slowly moving charged debris from neutron (or other) interactions in the detector gas or detector materials.

The electrical discharge protection system developed for the NSW Micromegas detectors consists of adding a layer of resistive strips on top of the readout strips, separated by a thin insulator, as shown in Figure 5.6. The readout electrode is no longer directly exposed to the charge created in the amplification region, but capacitively coupled to the signals [15]. Some fraction of the signal height is lost, but the chamber can be operated at higher gas gain, with electrical discharge intensities reduced by about three orders of magnitude.

To create the amplification and drift potentials, positive HV is applied to the resistive strips, and the micro-mesh is connected to ground, while the drift cathode is held at a negative potential [68]. This scheme helps to reduce electrical discharges and allows for more stable operation of the detectors, as voltage

breakdowns remain local and charge can be evacuated very quickly to ground through the micro-mesh, and the micro-mesh potential does not change. This HV scheme allows for easy segmentation of the HV connection scheme, which simplifies detector fabrication.

The readout strips are etched on 0.5 mm thick PCBs covered by a 64 μm -thick layer of insulator, and are overlaid with the resistive strips. The micro-mesh support pillars are deposited on top of the resistive strips in a $7\times 7\times 7$ mm triangular grid. Unlike most Micromegas detectors now in operation, in the ATLAS implementation the micro-mesh is not integrated into the readout structure, but rather connected to the drift panel, and stretched across the support pillars of the readout panel by mechanical tension [71–73]. The precision direction of the Micromegas coincides with the bending direction of the toroidal magnetic field, while the second coordinate is determined using strips with a small stereo angle. The main detector and operating parameters of the Micromegas detectors for ATLAS are summarised in Table 5.2.

The HV stability of the Micromegas is found to be strongly correlated with the actual resistance of the resistive strips of the readout anode. Lower resistance values are often found very close to the panel border, giving rise to specific discharge points in those areas. To combat this, a process of edge passivation [72–74] was applied to the readout panels prior to their assembly. A thin film of epoxy was applied to the active region around the border of readout panels if the measured resistance was below a defined threshold. The breadths of the passivated border regions were up to a few centimetres. This procedure improved the HV stability of the chambers at the expense of a very small reduction in their active area.

Two gas mixtures have been extensively tested. The originally proposed and tested gas mixture, now considered the backup mixture, is 93% argon and 7% CO_2 ; this is the same mixture used for the ATLAS MDTs, except that the MDT gas requires active humidification up to the level of a few hundred ppm of water, while Micromegas need a very dry gas. Superior HV stability is, however, obtained by using an argon mixture in which 2% of CO_2 is replaced by 2% of isobutane, as the chambers reach the plateau of full efficiency for a lower HV than without the isobutane. The mixture remains non-flammable. This isobutane mixture, now considered the baseline, will be used for Run 3, and if ageing tests confirm that it does not give rise to deposits on the electrodes, it will remain the default gas for the lifetime of the experiment.

The Micromegas require a positive potential of about 500 to 600 V on the resistive strips, and a negative potential of about 250 V on the drift cathode, producing an amplification field of about 40 to 45 kV/cm and a drift field of around 500 V/cm, respectively, as illustrated in Figure 5.6. The precise choices for the amplification potential depend on the gas choice, and are given in Table 5.2. One of the great advantages of adding isobutane to the gas mixture is that it allows the chambers to run at full efficiency with lower amplification potentials, improving stability and reducing the number and intensity of electrical discharges. The inner quadruplets, which are subject to higher background rates, may be run at a lower potential than the outer ones. All the drift electrodes of each quadruplet can be supplied by a single HV channel, so each wedge requires just two channels for the negative drift voltage. Since the resistive strips are split in the middle, positive voltage is supplied separately to each side of each readout PCB. Each of these half-PCBs is isolated by its own HV capacitor. As each layer of a wedge has eight PCBs, this results in 16 readout HV sections per wedge layer; however, the half-PCBs of one layer of one quadruplet share a common HV line, so there are two independent HV inputs on each side of each wedge layer (one to each quadruplet). A splitter box distributes the input HV to the groups of readout HV sections.

5.2.4.1 Micromegas Quadruplet and Wedge Structure Each sector of the NSW, Large or Small, contains two Micromegas wedges attached to the faces of the central aluminium spacer frame as described

Table 5.2: Main Micromegas detector and operating parameters.

Item/Parameter	Characteristics	Value
Micro-mesh	Stainless steel; mesh separate from readout board	
Micro-mesh	Wire diameter	30 μm
Micro-mesh	Gap between wires	71 μm
	Micro-mesh is separate from readout board.	
Amplification gap		120 to 130 μm
Drift/conversion gap		5 mm
Resistive strips	Interconnected	R=10 to 20 M Ω /cm
Readout strip width		0.3 mm
Readout strip pitch	Inner / Outer Modules	0.425 / 0.45 mm
Stereo angle	one + and one – slope layer per quadruplet	$\pm 1.5^\circ$
Total number of strips		2.1M
Baseline Gas	Ar:CO ₂ :iC ₄ H ₁₀	93:5:2
Backup Gas	Ar:CO ₂	93:7
HV on resistive strips	Baseline gas	500 V
HV on resistive strips	Backup gas	570 V
HV on drift panel		–240 V
Amplification field		~ 40 kV/cm
Drift field		480 V/cm
Bending Coordinate Resolution	single-plane, η -strips centroid fit	100 to 200 μm
Second Coordinate Resolution	ϕ , single-plane, stereo strips centroid fit	2.7 mm
NSW-TP Bending Coord. Res.	from NSW-TP fitter	300 μm
NSW-TP Second Coord. Res.	$r\phi$, from NSW-TP fitter	11 to 12 mm

in Section 5.2.3. The spacer frame is hollow, with integrated water cooling channels, and most of the Micromegas services are routed through this central volume.

Each Micromegas wedge is assembled from two separate quadruplet modules, as shown in Figure 5.4. There are thus four sizes of Micromegas quadruplets: Inner and Outer (radially) for the Small wedges, and similarly for the Large wedges. Each quadruplet consists of five stiff panels, 1 cm thick. Three are drift panels (the outer, single-sided, and the central, double-sided). The other two are the double-sided readout panels. Between the panels are the four active layers, corresponding to independent gas volumes 5 mm thick, as shown in Figure 5.7. A particle traversing a sector of the NSW thus passes through eight Micromegas drift regions (four in each of the two wedges, as shown in Figure 5.4(b)), inducing charges on readout strips on the corresponding readout panels. Two of these layers on each wedge have readout strips perpendicular to the radial centre-line of the wedge, (“eta strips”) and two have strips at a small “stereo” angle of $\pm 1.5^\circ$ (one positive and one negative) with respect to the eta strips.

The double-sided drift panel forms the centre of each quadruplet. Double-sided readout panels mounted on either side of the central drift panel form the two central gas volumes. Single-sided drift panels mounted on the outside complete the two outer gas volumes. The four active layers of the quadruplet are grouped in two pairs, each pair sharing one double-sided readout panel in a back-to-back configuration. This ensures that background tracks not synchronous with the bunch crossing will not be collinear in the two neighbouring planes, as the drift in the two adjacent detectors is in opposite directions, and the offsets of the track segments due to assuming the wrong start time for the drift go in opposite directions.

Out-of-time background can thus be rejected. The back-to-back structure also ensures that systematic shifts of reconstructed particle positions due to the deviation of the electrons' drift path in a magnetic field (the Lorentz angle) cancel out.

The readout panels are assembled from separate PCBs spanning the full azimuthal width of the sector, but with radial dimensions in the range 435 to 475 mm to allow for production on standard PCB machines. The inner quadruplet in each wedge has five radial PCB segments in each readout panel, and the outer quadruplet has three. Readout strips must be precisely aligned from one face of each readout panel to the other, and from one readout panel to the next. Several complementary methods were used to locate and align the readout strips with precision [75], including a mechanical hole-and-slot pair, a pair of targets included in the copper strip pattern and aligned with the board axis and strips, and four RASNIK [76] masks etched on the edge of each readout PCB. RASNIK masks consist of a small “chessboard” pattern of dark and light squares, with absolute positions encoded by inverting the colours of a few squares, to be read out with a CCD. RASNIK masks etched on the PCBs on both faces of the readout panels can be read out simultaneously with a dedicated calibrated device called a RasFork [72]. The RASNIK mask positions are read out during various stages of chamber assembly, with an accuracy of 3 to 5 μm , such that the strip positions on the PCBs are known with an accuracy better than 50 μm at the end of the construction process. Subsequent deformations of the Micromegas quadruplets due to temperature and humidity can affect the positions of the strips, increasing this uncertainty to around 100 μm . Alignment source platforms, described in more detail in Section 5.2.6, are glued to the outer surface of each Micromegas wedge.




Figure 5.7: Arrangement of readout and drift panels in a Micromegas quadruplet. An interconnect is shown in cross-section at the front of the drawing; these interconnects limit bulging of the panels due to the slight (3 mbar) overpressure of the gas volumes.

Each Micromegas wedge is read out by 64 Micromegas Front-End Boards (MMFE8s) [77] (described in Section 5.2.7.3) connected directly to the ends of the readout strips with flexible pressure connectors. Also mounted on the wedge are eight Level-1 Data Driver Cards (L1DDCs) for data readout (described in

Section 5.2.7.4), each connected to eight of the MMFE8s, and eight ART Data Driver Cards (ADDCs) for the Micromegas trigger (described in Section 5.2.7.4), also each connected to eight MMFE8s.

The single-plane spatial resolution obtainable in the precision (“eta”) direction has been estimated using small Micromegas bulk prototype chambers (with different electronics) in a test beam to be about 90 to 100 μm per plane; by using the “ μTPC method” [68] to reconstruct the position from inclined tracks, it is expected to be possible to obtain a single-plane resolution within that range for all inclinations between 0° and 40° . A second-coordinate resolution of about 2.7 mm can be obtained using the stereo strips [78], and this has been demonstrated using the centroid method in a test beam [79] using prototype detectors (1x0.5 m) with a slightly smaller strip pitch (415 μm) than the ATLAS Micromegas, positioned perpendicular to the beam axis. Similar results for the spatial resolution of perpendicularly incident 120 to 150 GeV muons and pions were achieved in the precision direction and second coordinate in SPS experiments using three Module 0 chambers and one series production chamber which was read out using the VMM electronics. The results of the first test with a Module 0 chamber appear in Refs. [80] and [81]; the others have not yet been published. Run 3 ATLAS data will be required to determine the ultimate resolution obtainable with the full-sized detectors and non-perpendicular tracks.

5.2.5 Small-Strip Thin Gap Chamber Technology

The small-strip TGC (sTGC) is a new development of the TGC technology first developed in the early 1980s [82], to allow for very fast on-line tracking that can be used in the L1 (hardware) trigger, with sufficient precision for the offline muon tracking as well. Wherever possible, the sTGC retain the characteristics of the TGC already used in ATLAS: two FR4 cathodes coated with a resistive graphite-resin mixture span a gap 2.8 mm thick. The gap is filled with a mixture of 55% CO_2 and 45% n-pentane. The anode plane in the middle of the gap is strung with 50 μm gold-plated tungsten wires at a pitch of 1.8 mm. The wire direction is parallel to the central radial axis of the chamber (contrary to the existing TGC in ATLAS where the wires are strung azimuthally). The wires are ganged together, with groups of typically 20 sharing a high-voltage capacitor.

The resistivity of the cathode coating in the sTGC is significantly lower than for the other TGC in ATLAS: about $150 \text{ k}\Omega/\square$ ¹⁰ for the chambers closest to the IP, and $200 \text{ k}\Omega/\square$ for the others. This low resistivity allows for rapid clearing of charge on the cathode surface. Signals are capacitively induced on copper readout strips (running perpendicular to the wires) on one side of the gap, and on large copper pads on the other side. The signal layer is under a thin (150 or 200 μm) insulating pre-preg layer directly under the graphite cathode, increasing the capacitance compared with the other ATLAS TGC (where the signal strips are on the back of the FR4), to keep the same transparency for fast signals. This readout layer is supported by a 1.3 to 1.4 mm-thick PCB backed by a copper grounding skin to limit cross-talk. The pad board has an additional pre-preg layer between the core of the PCB and the ground skin to support the readout traces, which are connected by conductive vias to the pads.

The main detector and operating parameters of the sTGC detectors for ATLAS are summarised in Table 5.3 and illustrated in Figure 5.8.

The readout strips are 2.7 mm wide, with a pitch of 3.2 mm, much finer than the TGC strips. The strip width and pitch required for the trigger and tracking functions were both optimised by studying the position

¹⁰ Surface resistivity is the resistance between two opposite sides of a square of a thin-film resistor, and is independent of the size of the square; it is measured in Ω , but in order to distinguish it from resistance and emphasise its two-dimensional property, it is usually written as Ω/\square or Ω/square .

Table 5.3: Main sTGC detector and operating parameters. Total numbers of channels are for both wheels together. “Outer Quadruplets” refers to both Middle and Outer Quadruplets; the Inner Quadruplets are exposed to substantially higher background rates, and therefore have finer pad granularity and less resistive cathodes for faster charge evacuation. The wire groups in the innermost part of the Inner Quadruplets are not read out.

Item/Parameter	Characteristics	Value
Readout strip pitch		3.2 mm
Readout strip width		2.7 mm
Total number of strips		282×10^3
Typical pad azimuth	Inner / Outer Quadruplets	$5^\circ / 7.5^\circ$
Typical pad radial height		80 mm
Range of “full” pad areas		61 to 519 cm ²
Total number of pads		46 656
Anode-cathode gap		1.4 mm
Wire pitch		1.8 mm
Wire diameter	Gold-plated tungsten	50 μ m
Total number of wires	Ganged in groups of 20, group boundaries offset by 5 wires between layers	6 390 296
Number of wire groups	Total / Read out	31 776 / 28 704
HV on wires	Positive polarity	2.8 kV
Cathode resistivity	Inner / Outer Quadruplets	150 / 200 k Ω /□
Pre-preg thickness between readout and cathode	Inner / Outer Quadruplets	150 / 200 μ m
Gas	n-pentane:CO ₂	45:55
Bending Coordinate Resolution	single-plane η from strips	100 to 200 μ m
Azimuthal Resolution	single-plane $r\phi$ (from wire groups)	2.6 mm
NSW-TP Bending Coordinate Res.	L1 η from centroid fit to strips in band	< 1 mrad
Pad Trigger Azimuthal Res.	ϕ at L1 from pad towers (Inner/Outer)	7 mrad / 10 mrad
	$r\phi$ (at L1 from pad towers)	7 to 38 mm

resolution of small detectors built with a range of strip widths and separations, to reduce capacitive coupling between channels and non-linear corrections as much as possible for trigger purposes, while maintaining the smallest possible number of readout channels. In combination with a reduction of the HV and ground decoupling resistors, this allows the sTGC to remain efficient for minimum ionising particles even in large-surface detectors subject to detected rates of up to 20 kHz/cm² over the full surface [15]. The intrinsic position resolution achievable with the final strip configuration was measured using a quadruplet made with full-sized prototype boards, in conjunction with a silicon pixel telescope (used to determine independently the individual particle trajectories), in a 32 GeV pion testbeam at Fermilab [83], and was found to be around 45 μ m for particles incident perpendicular to the plane of the quadruplet. The resolution obtainable in practice is limited by the precision with which the as-built positions of the strips are known.

5.2.5.1 sTGC Quadruplet and Wedge Structure In order to provide a track segment of sufficiently good pointing resolution, as well as for trigger efficiency, the sTGC gas gaps are assembled in quadruplets, with four gas gaps separated from each other by 5 mm-thick fibreglass frames and paper honeycomb spacers, as shown in Figure 5.8(b). The dimensions of the quadruplets were dictated by the maximum size




Figure 5.8: The sTGC (a) detector operating principle and (b) arrangement of gas gaps (in green, each comprising a strip cathode and a pad cathode separated by two frames clamping the wire plane) to form a quadruplet, shown to scale at the corner of a quadruplet, with brass alignment features inset in the strip cathodes. The pre-preg thickness is 150 μm for the innermost quadruplets and 200 μm for the middle and outer ones.

of cathode boards that could be produced industrially: the largest quadruplets are about 2.2 m wide (the full width of the Large sector at its outer radius). Three trapezoidal¹¹ quadruplets of different sizes are precisely glued together with a pair of fibreglass frames (visible in Figure 5.3(a)) to form the wedges shown in Figure 5.4. The fibreglass frames incorporate three pairs of projecting fibreglass bars that are used to support the mount points from which the wedges are hung on the sectors. The positions of these projections were mechanically constrained during the gluing of the wedges, relative to alignment features¹² on the cathode strip boards of all four detector planes, using precisely machined cylindrical pins, placed with an accuracy of a few tens of μm on the surface plates used for wedge assembly. The same brass alignment features were used to position the alignment source platforms (described in Section 5.2.6) on the surface of the sTGC wedge that faces the Micromegas double-wedge. The accuracy with which the positions of the brass features are known with respect to the platform positions (combined with the generally lower accuracy with which the strip positions are known relative to the brass features) limits the precision with which the absolute strip positions can be determined. The kinematic mounts are designed to allow for adjustment between the two sTGC wedges to achieve parallelism of their strips, as defined by the alignment platforms, to better than 1 mrad.

There are six sizes of sTGC quadruplets: three for the Small-sector wedges and three for the Large. Each sector has two sTGC wedges, hung kinematically from the sector spacer frames, outside the Micromegas wedges. This maximises the lever arm between the track segments provided by the two sTGC wedges

¹¹ In the case of the outermost quadruplets in the Large sectors, the trapezoids are modified, with a rectangular section at the outer radius (see Figure 5.4(a)).

¹² Each strip board has two machined brass alignment features on one edge, allowing for alignment against a pair of precisely positioned cylindrical pins: a “V”-shaped feature allowing only rotation in the plane about the pin defines a fixed origin, and a flat feature parallel to the edge of the board defines a direction with respect to the first pin while avoiding over-constraining the board in case of small scaling variations due to temperature or humidity changes.

(refer to Figure 5.4(b)).

The HV wires of the innermost quadruplets in each wedge (those closest to the beampipe) are split in two, creating two HV regions within a common gas volume. This separates the high-background region ($2.4 < |\eta| < 2.7$) and allows it to operate at the same effective amplification as the rest of the detector, by means of the HV divider network. Only the higher-radius portion of the wires ($|\eta| \lesssim 2.4$) is read out for azimuthal coordinate measurements; since the azimuthal granularity of the pad cathode boards is comparable to that of the wire groups for $|\eta| > 2.4$, the pads provide the azimuthal coordinate in the region with no wire readout.

No active components are mounted on the cathode boards. All pad and strip readout channels are connected by soldered jumper wires to long multilayer adapter boards attached to both sides of each gas gap. Wire groups are read out through narrow adapter boards soldered to the longer parallel edge of each quadruplet, and their signals routed through the pad adapter boards by means of a ribbon cable. A pad or strip front-end board, described in Section 5.2.8, is connected to each pad or strip adapter board respectively, using a high-density low-profile connector.

As for all TGCs in ATLAS, each sTGC wedge is enclosed around its periphery by a gas-tight envelope that is continuously flushed with CO₂. This is done to maintain a dry atmosphere around the HV elements, as well as to dilute any possible leak of the operating gas. If traces of n-pentane are detected in the CO₂ stream at the output of a wedge, HV and LV as well as gas supplies are automatically switched off, and an alarm is activated. The gas channel also serves (in conjunction with the copper ground skins of the detectors themselves) to complete the Faraday cage surrounding the detectors, adapter boards and front-end electronics. The connectors for the front-end boards protrude through sealed slots in the Faraday cage; the front-end boards thus sit outside the gas channel. The Faraday cage is completed by electrical shielding mounted around the front-end boards, to which the cooling pipes (described in Section 5.2.11) are connected by a system of thermally conductive clips.

5.2.6 Extending the Endcap Alignment System

The ATLAS muon endcap alignment system[84] is based on the concept of calibrated alignment bars connected by optical lines. Each alignment bar is a hollow aluminium tube, very precisely measured, containing precisely calibrated instrumentation to monitor its own length and deformations.

The smallest alignable units of the NSW are the Micromegas and sTGC quadruplets. The alignment system monitors the rotations and translations of the quadruplets, as well as a limited number of deformation modes. Within the individual quadruplets, it is assumed that the internal geometry does not change (apart from the allowed deformations); however, it is not assumed to be ideal: data taken during the construction of the quadruplets allow the as-built position of each individual strip to be reconstructed with respect to the coordinate system of the quadruplet.

Each NSW has 16 alignment bars, one in each sector, unlike the EM and EO Wheels (and the original EI wheels), which have only eight. This doubling allows Brandeis CCD Angle Monitors (BCAMs) on the alignment bars to see both sTGC wedges and both Micromegas wedges in each sector. The NSW alignment bars are mounted in the Large-sector and Small-sector spokes (see Section 5.2.3), and connected to each other and to the adjacent chambers by a network of optical lines, where CCD cameras in BCAMs [84] mounted on the alignment bars monitor the positions of light sources mounted on the chambers and on BCAMs on the neighbouring alignment bars. The NSWs are aligned with respect to the other wheels of the

endcap MS using the existing *Polar* alignment corridors, which traverse tightly constrained unobstructed paths through holes and gaps in detectors and support structures, and through dedicated hollow tubes passing through the endcap toroid cryostats. Polar BCAMs are therefore positioned on the eight Large-sector alignment bars, in the same positions as the former EI polar BCAMs. *Azimuthal* alignment corridors connect all the alignment bars to their neighbours, as shown in Figure 5.9(a), with BCAMs on the alignment bars pointing at light sources on the nearest neighbour BCAMs in both the adjacent Large and Small sectors (Large-to-Large and Large-to-Small, or Small-to-Small and Small-To-Large). The optical connections among the 16 alignment bars and the 64 wedges in each endcap are illustrated in Figure 5.9.

Quadruplets on each sTGC or Micromegas wedge must be aligned with respect to the adjacent alignment bars. *Proximity* lines connect the BCAMs on the alignment bars to light sources on the sTGC and Micromegas wedge surfaces inside the 27 mm gaps between the sTGC and Micromegas wedges in each sector. Light injectors located at the periphery (“rim”) of the NSW illuminate optical fibres. The ends of these fibres are held in precisely determined positions on the alignment source platforms (mentioned in Sections 5.2.4 and 5.2.5 above) installed on the sTGC and Micromegas wedge surfaces. Each wedge has 18 platforms glued to its surface. In the sTGC wedges, six of these (in two lines of three) are glued to the outermost quadruplet in the wedge, nine (in three lines) to the middle quadruplet, and three (in one line) to the innermost quadruplet. In the Micromegas wedges nine platforms (in three lines) are glued to the outer quadruplet and nine (in three lines) to the inner quadruplet. Each platform holds two or four optical fibre ends. The Micromegas and sTGC platforms are sufficiently close to each other that the BCAMs on the alignment bars can each see ten light sources: two each from two Micromegas platforms and two from each of the corresponding sTGC platforms, and two from its partner proximity BCAM on the opposite side of the sector. The positions of the source platforms are precisely measured using photogrammetry with a resolution of about 50 μm . Their positions are measured with respect to the RASNIK masks that locate the strip positions on the Micromegas readout boards, and to the brass edge-features (see Figure 5.8(b)) that determine the strip-pattern alignment of the sTGC cathodes; thus, for both Micromegas and sTGC the source platform positions are directly related to the positions of the readout strips. The positions of the illuminated fibre ends on the source platforms are continuously monitored by BCAMs on the alignment bars. The alignment system can detect and measure chamber movements as small as 40 μm [84].

5.2.7 NSW Electronics

The trigger and readout electronics for the NSW are closely intertwined. Three interconnected pathways are illustrated in Figure 5.10: the Trigger and Timing path, described in Section 5.2.8, the Data Acquisition and Readout path, described in Section 5.2.9, and the Configuration and Monitoring Path (see Section 7.7). Each of these pathways has elements located directly on the sTGC and Micromegas wedges, in the “Rim Crates” mounted at the outer radius of the NSW, and in the service cavern USA15. This chapter describes only what is located on the NSWs: the front-end electronics, the sTGC Pad Trigger, and the additional elements that provide inputs to the NSW-TP for the sTGC and Micromegas strip triggers, and to the data acquisition system. The NSW-TP and FELIX aggregator that deliver the data to the ATLAS trigger and data acquisition systems are described in Sections 7.3.3 and 7.6.1. More detailed descriptions of the architecture, design considerations, and functionality of all the electronics described in this section are provided in Ref. [85].

5.2.7.1 CERN ASICs. The following ASICs developed at CERN and used extensively in other experiments as well as ATLAS are fundamental elements of the NSW electronics design:

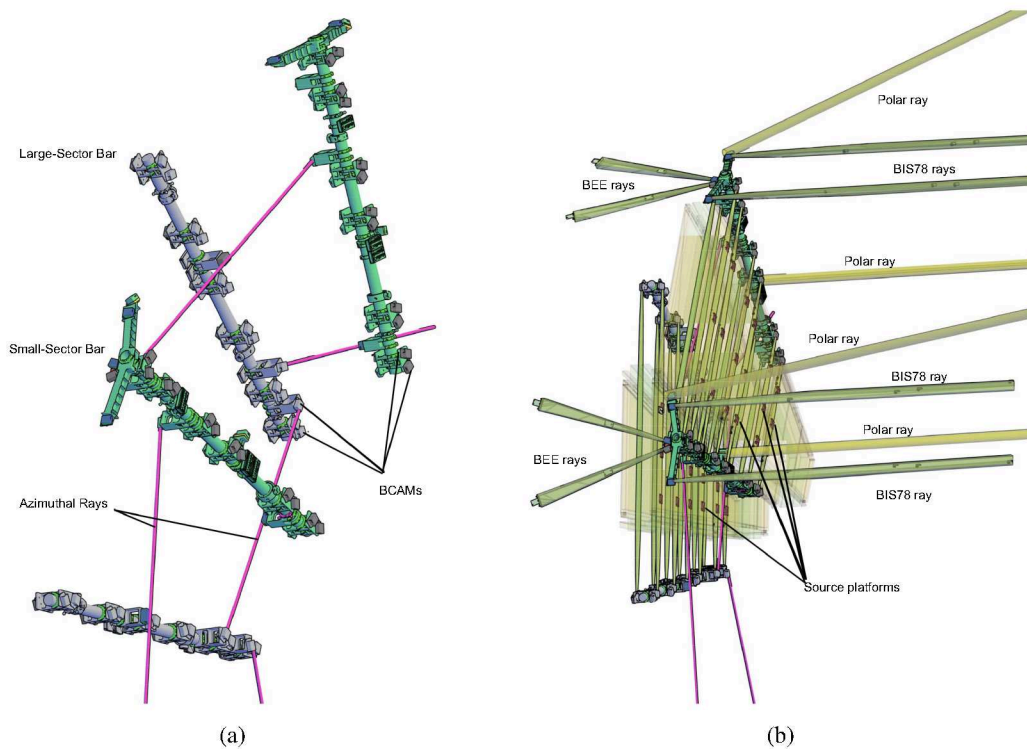


Figure 5.9: Alignment rays in the NSW. (a) Alignment bars (green) in the spokes supporting the Large sectors occupy the positions previously filled by the old EI alignment bars. They are called the Small-sector bars because the spokes supporting Large sectors sit in the Small sectors on either side of each Large sector. The Small-sector bars are connected by polar rays to the EM alignment bars exactly as the original alignment bars were. They are connected to each other and to the Large-sector alignment bars (grey, located in the spokes supporting the Small sectors) by azimuthal rays (magenta) connecting BCAMs (grey boxes) with integrated light sources to each other. (b) The Small-sector bars are connected to the Large-sector sTGC and Micromegas quadruplets (shown transparent), and the Large-sector bars are connected to the Small-sector sTGC and Micromegas quadruplets, by proximity rays (yellow) connecting BCAMs on the bars to light sources on platforms mounted on the quadruplet surfaces (small magenta rectangles). The Small-sector bars also have BCAMs that connect them to EM (Big Wheel) alignment bars via the polar alignment rays, shown in yellow and pointing toward the right, and to sources on the BIS78 chambers (grey-yellow rays pointing to the right), and to sources on the BEE chambers (grey-yellow rays pointing to the left).

- GBTx: the Gigabit Transceiver [86] aggregates many slow serial E-LINKs into a single serial link running at 4.8 Gb/s
- GBT-SCA: every NSW board includes a uniquely identified GBT-SCA [87, 88] (see Section 7.7) for configuration, calibration and monitoring using the GBT-SCA OPC UA Server. Each GBT-SCA includes a 31-channel 12-bit ADC for monitoring the temperature sensors, on-chip temperatures and voltage power levels of the front-end boards. It is configured via the FELIX (Section 7.6.1).
- FEAST: The FEAST [89] is an integrated DC-to-DC convertor providing up to 10 W of power at the voltages required by the ASICs on the NSW boards, in the range 1.2 to 3.3 V, from an input supply in the range of 5 to 12 V DC. All the NSW boards are powered by the FEAST.
- Versatile Links: the Versatile Link Transceiver (VTRx) and Unidirectional Twin-Transmitter version of the VTRx (VTTx) are radiation tolerant optical link interfaces, transmitting data between

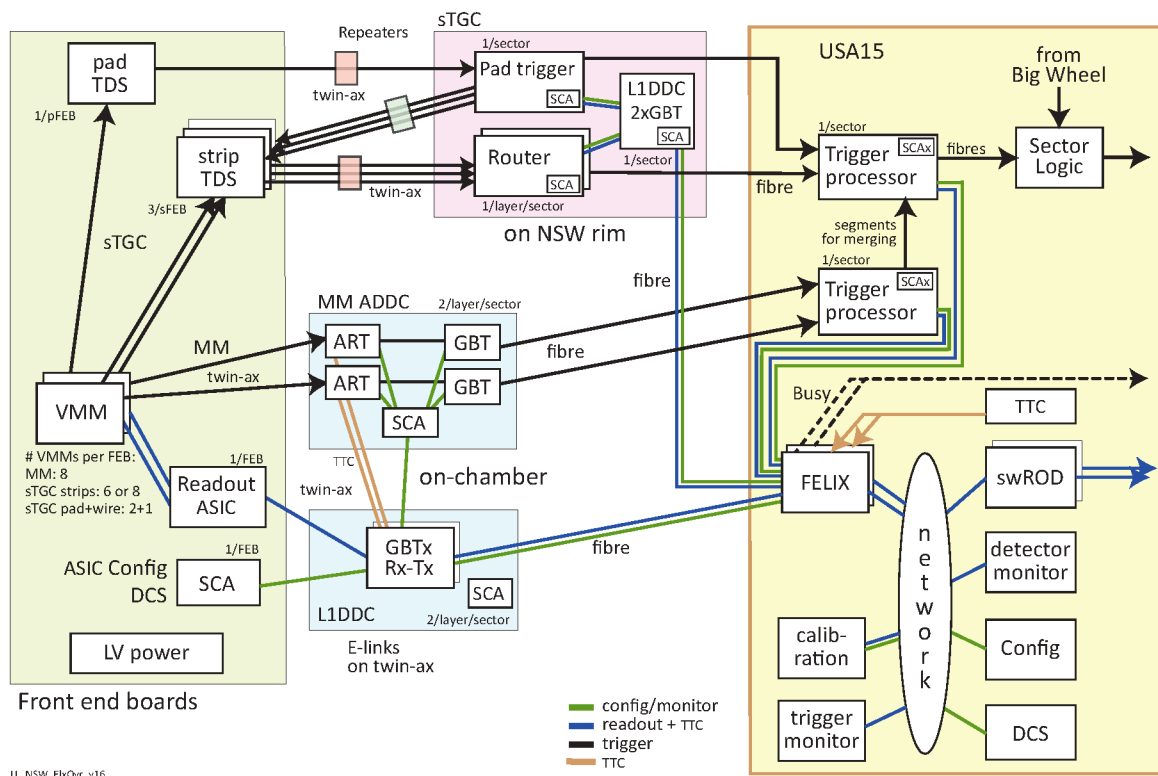


Figure 5.10: Overview of the NSW electronics on- and off-detector electronics and their connectivity. The front-end boards are connected directly to the detectors (left box). Also mounted directly on the wedges are the ADDC boards for the Micromegas trigger (upper blue box) and the L1DDC used for readout by both detectors (lower blue box). The boards for the sTGC Pad Trigger are located in the Rim Crates of the NSW (central pink box), which also contain the Routers for the sTGC trigger, the Rim-L1DDC and an LVDB. Every NSW electronics board installed in the ATLAS cavern has a GBT-SCA for configuring and monitoring its components. Off the detectors in USA15 (right box) are the NSW-TP and the FELIX aggregator that deliver the data to the ATLAS trigger and data acquisition systems, and are described in Section 7.

on-detector and off-detector electronics at up to 5 Gb/s [90].

5.2.7.2 NSW Custom ASICs. The NSW electronics rely on a number of custom ASICs, described in this section. The VMM and ROC are designed to work in both NSW subsystems, while the TDS and ART are specific to the sTGC and Micromegas, respectively, and supply the inputs required by their respective NSW-TP algorithms.

VMM. A new, highly configurable custom Amplifier-Shaper-Discriminator (ASD) front-end ASIC, the VMM3a [91], was developed for both the sTGC and Micromegas front-end boards. The VMM can read out, amplify and shape, and provide peak finding and digitisation for up to 64 channels from the negative anode strip signals of the Micromegas, or the negative wire-group signals, or positive cathode strip or cathode pad signals of the sTGCs. Thresholds can be set independently for each channel, and a configurable option can also read out the peak charge of the nearest neighbours of strips with signals above threshold, so that a cluster of at least three strips can always be used for centroid-finding. The VMM also provides trigger outputs from the sTGC pads or strips, or from the Micromegas.

The VMM has four independent output paths, of which three are used in the NSW:

- Precision (10-bit) amplitude and (effective) 20-bit timestamp readout at L1A ¹³ with 250 ns deadtime per channel and a 64-deep FIFO per channel guaranteeing no data loss for a trigger latency up to 16.0 μ s;
- Serial out Address in Real Time (ART) synchronised to a 160 MHz clock, used to provide a 6-bit address of the first strip above threshold for the Micromegas trigger;
- Parallel prompt outputs from all 64 channels in a variety of selectable formats for the sTGC trigger, including a 6-bit ADC. When a peak is found in the sTGC strips, its amplitude is digitised in 6-bits and immediately sent serially, one line per channel, to the strip-Trigger Data Serializer (TDS).

Readout Controller. Present on all the FEBs, the ReadOut Controller (ROC) [92] receives L0A and L1A bits from the trigger (see Section 7.3). On receipt of a L1A ¹⁴, the ROC reads out data from the VMMs. The VMM writes to the ROC with an output bandwidth of 640 Mb/s (512 Mb/s before 8b/10b encoding). Each ROC receives and processes data from up to eight VMMs, and transmits the output to the L1DDC.

sTGC Trigger Data Serialiser. The TDS [93] is a single ASIC designed with two operation modes to handle pad and strip signals from the VMMs on the sTGC FEBs, respectively denoted “pad-TDS” and “strip-TDS”. In either mode, the TDS consists of 64-channel interfaces to two VMMs, as well as a Preprocessor and a Serialiser. A pad-TDS reads out up to 104 pads, and a strip-TDS up to 128 strips.

ART. The ART chips [94] on the Micromegas ADDC boards receive the six-bit address of the strip that fired first in a BC from each of 32 VMMs, and aggregate the addresses (and only the addresses) of the fired strips that are the inputs to the Micromegas strip trigger algorithm in the NSW-TP.

5.2.7.3 NSW Front-End Boards. As far as possible, the Micromegas and sTGC share front-end electronics components: the highly configurable VMM was designed to meet the requirements of both technologies, and all the front-end boards use the same ROC to buffer and process signals from the VMMs for readout and transmit them to the L1DDCs, as described in Section 5.2.9. All use the GBT-SCA for slow controls. The FEBs are all described in detail in Ref. [85]. The left-hand box of Figure 5.10 illustrates the major similarities and differences between the three types of FEB. The FEBs all receive power from Low Voltage Distributor Boards (LVDBs), described in Section 5.2.7.4.

sTGC Front-End Boards. The sTGCs have two types of front-end boards: the strip Front-End Board (sFEB) to read out the strip signals, and the pad Front-End Board (pFEB) to read out the pad and wire signals. Each sFEB has six (outer and middle quadruplets) or eight (inner quadruplets) VMMs, depending on the number of strips – one VMM can read out up to 64 strip channels. Each pFEB has three VMMs: two for the pads and one for the wires. Both types of FEB use the ROC to prepare VMM signals for readout via an L1DDC (Section 5.2.7.4). Both use TDS ASICs to digitise signals for the NSW-TP, as described in Section 5.2.8. There is one TDS on the pFEB, connected to two VMMs (the third VMM

¹³ Level-0 Trigger Accept (L0A) after Phase-II upgrades.

¹⁴ L0A for HL-LHC.

on the pFEB reads out the wires, which are not used in the trigger, and therefore requires no TDS). The pad-TDS transmits serialised pad hits to the Pad Trigger. Every sFEB has a TDS connected to each pair of VMMs that tags all active strip charges with the Bunch Crossing IDentifier (BCID) and saves them in its ring buffer, awaiting possible selection by the Pad Trigger for transmission to Routers in the Rim Crates. Communications between the ROC and the L1DDC, and between the pad-TDS, the strip-TDS and the Pad Trigger and Router, all use miniSAS twin-axial ribbon cables.

Micromegas Front-End Board. The Micromegas have a single type of front-end board, the MMFE8, that acts as an interface between the Micromegas detectors and the trigger (ADDC) and data acquisition (L1DDC) electronics cards. Each MMFE8 has eight VMM chips to read out the strips. Each VMM sends signals for the trigger to an ART ASIC on an ADDC board. A ROC formats and buffers VMM strip signals for transmission to an L1DDC. All signals to the L1DDC and ADDC are sent over miniSAS twin-axial ribbon cables.

5.2.7.4 Trigger, Readout and Low voltage On-Detector Boards. The sTGC and Micromegas provide different types of data to the trigger and to the DAQ readout. The Level-1 Data Driver Cards (L1DDCs) aggregate and transmit the data acquisition readout for both technologies. The Micromegas have an additional set of cards mounted directly on the detectors, the ADDCs, which transmit the addresses of the strips with the first hits to the NSW-TP (see Section 5.2.8.3). This section also describes the LVDBs used to distribute power, the Router cards used by the sTGC strip trigger, the sTGC Pad Trigger Boards, and the Serial and LVDS Repeaters.

Level-1 Data Driver Card (L1DDC). The L1DDC [95] is an intermediate element in the DAQ system for both the Micromegas and sTGC detectors. It is transparent to the type of data being transmitted or received. It combines the three pathways listed above (Trigger and Timing, DAQ and Readout, Configuration and Monitoring) into one or more bidirectional optical links (VTRx). Each L1DDC collects detector signals and monitoring data through E-LINKs to a GBTx [90], which then transmits the data to the relevant back-end systems through bidirectional (VTRx) or unidirectional (VTTx) optical fibre links at 4.8 Gbps each. The links are powered by FEAST DC-DC converters. The L1DDC also distributes synchronous clocks, and trigger and configuration data from the back-end systems, to the FEBs.

The NSW uses three different types of L1DDC boards, all made with the same custom radiation-tolerant ASICs, including GBTx for high-speed serialisation and deserialisation, Gigabit VTRx optical transceiver and transmitter modules, and a GBT-SCA to monitor voltages and temperatures and to configure the ROC on the FEBs. The Micromegas and sTGC detectors have slightly different L1DDC boards mounted on the wedges to read out the strips, as well as the sTGC pads and wire groups. When a L1A initiates the transfer of data buffered on a ROC, the data are transferred to a GBTx that, in turn, drives the output to FELIX over the optical links. The Micromegas L1DDC contains three GBTx ASICs so it can read out eight MMFE8s within one detector layer. A common L1DDC board design, with room for two GBTx chips, is used for sTGC strips and for sTGC pads and wires, although the L1DDCs for pads and wires have only one of the two GBTx chips installed. Each sTGC L1DDC reads out the three FEBs on one side of one detector layer within a wedge. The Rim-L1DDC is mounted in the NSW Rim Crates (see Figure 5.10) and provides readout of the sTGC Pad Trigger input hits and its output decisions. The GBT-SCA on the Rim-Crate L1DDC configures the Pad Trigger and the eight Router cards in the sTGC trigger path.

Low-jitter synchronous clock signals for the Pad Trigger and Routers can be provided either from the GBTx on the Rim-Crate L1DDC or from a direct fibre from USA15.

Each Micromegas wedge layer requires two L1DDCs (eight in total per wedge) mounted at intervals along the outside edges of the wedge, on both edges. The sTGC on-wedge L1DDCs are also mounted on both edges. There are eight on each wedge in total, one each for the pFEBs and sFEBs of each layer.

ART Data Driver Card (ADDC). The Micromegas ADDC contains two ART ASICs, each receiving data as described above in Section 5.2.7.2 from four MMFE8s, and thus from 32 VMMS, all treated as independent data streams by the ART. Each ART transmits its data to a GBTx, which in turn transmits it to the NSW-TP through one of the two transmission channels of a VTTx.

sTGC Pad Trigger Board. In order to place the track finding and extrapolation logic in a less harsh and more accessible location, a coincidence, typically requiring at most a three-out-of-four coincidence of pads in each of the four-layer quadruplets, is used to choose the relevant bands of strips to be sent off-detector (as described in Section 5.2.8.1). This substantially reduces both the required bandwidth and the amount of centroid and track-finding logic. The board responsible for this decision-making is the sTGC Pad Trigger [85], located in the Rim Crate for each sector, and implemented in radiation- and magnetic field-tolerant electronics. Placing the boards at the rim of the wheel, where radiation is less intense, allows the flexibility of forming the pad tower coincidence in programmable logic on FPGAs. Each Pad Trigger board receives 24 input links (three per layer) at 4.8 Gb/s from the pad-TDS ASICs on the pFEBs of the sector. The Pad Trigger sends outputs of its logic to the strip-TDSs on the sFEBs via MiniSAS ribbon cables, to the NSW-TP via optical VTTx, as well as sending the pad hit readout data directly to FELIX via E-LINK.

sTGC Router. The sTGC Router [96] serves as a packet switch for routing strip charge information from the strip-TDS on the sFEBs to the sTGC Trigger Processor (TP) [85]. There is one per layer for each sector. They are implemented in radiation- and magnetic field-tolerant electronics and sit in the sector's Rim Crate with the Pad Trigger board.

sTGC Serial and LVDS Repeaters. Signals on the fast copper TwinAx connecting the FEBs to the Pad Trigger and the Router suffer excessive attenuation if the cable length exceeds about 5 m. The pad-TDS to Pad Trigger and strip-TDS to Router links operate at a speed of 4.8 Gb/s; the Pad Trigger to strip-TDS at 640 Mb/s, based on LVDS signals. Serial and LVDS Repeaters [97] (as appropriate) are therefore placed on all these lines to mitigate the attenuation.

Low Voltage Distributor Board. The input power to the FEASTs on all the NSW on-detector electronics is supplied by New Generation Power Supplies from CAEN (NGPSs) located in US15. The NGPSs provide 280 V to the Intermediate Conversion Stage (ICS) [85] modules (see Section 5.2.11.4), which convert it to 11 V and supply the electronics through the LVDBs. One LVDB can supply up to eight front-end boards (analogue section) and up to four digital boards (digital section). Each Micromegas wedge has eight LVDBs which act as splitters for the channels from the ICS. Every input channel is connected to eight MMFE8s, with an individual fuse for each MMFE8. There are two types of sTGC LVDB: one mounted on the Faraday cage of each wedge (64 in total), and one in each Rim-Crate (of which there are 32, one per

sector). The on-wedge LVDBs have fuses for each FEB, and each of the four LV power supply channels corresponds to all the FEBs on one detector layer (three pFEBs and three sFEBs); they also power the on-wedge L1DDCs. The Rim-Crate LVDBs aggregate two power supply channels each, and power all the Rim-Crate electronics. The Serial Repeaters typically draw power from a nearby L1DDC or, in a few cases, from a nearby LVDS Repeater.

5.2.8 Trigger

The overall requirements for the trigger electronics (and for the data acquisition — Section 5.2.9) are stringent: the trigger decision has to arrive at the Sector Logic (SL) within 1075 ns (43 BCIDs) of the collision, and nearly half this time is required for fibre and cable delays. The on-detector electronics must be radiation-tolerant to 3000 Gy, and operate in a highly inhomogeneous magnetic field, exceeding 0.5 T in places. Once the detectors are assembled, access to their electronics is extremely limited: most of the FEBs are completely inaccessible, and most of the Rim Crates are only accessible during shutdowns long enough for the NSW to be moved out of its running position. The trigger is therefore designed so that it still functions efficiently even if only a subset of the detector layers are operational, and some key components are redundant. Despite the fact that the NSW consists of two very different detector technologies, many components were designed to be used by both. A defining aspect of the architecture is the use of FELIX (see Section 7.6.1) for the Readout, Configuration and TTC distribution paths.

The L1 Muon trigger for Run 3 is described in Section 7.3. As in Runs 1 and 2, forward muon triggers require track segments in the TGCs of the EM-TGCs; however, for Run 3, in order to reduce the fake rate, the requirement of a matching segment in the NSW is added to the L1 Muon Sector Logic trigger decision.

This section focuses on relevant features of the detectors designed for NSW triggering, and on the trigger functionality of the front-end electronics located on the detector, and on the rim of the NSW. Three distinct trigger pathways identify potential signals from the NSW: the sTGC Pad Trigger, the sTGC Strip Trigger and the Micromegas Trigger. The Pad Trigger (described in Section 5.2.8.1) provides an initial RoI defined by an eight-layer coincidence “tower” of pads, defining an azimuthal range (identified by a “ ϕ -ID”) that is used directly by the TP, and a radial range (the “Band-ID”) that uniquely determines which group of sTGC strips should be read out by the sTGC Strip Trigger. The sTGC Strip Trigger performs fast cluster-finding using only strips within the range of the Band-ID indicated by the Pad Trigger, as described in Section 5.2.8.2. The Micromegas Trigger runs independently of the sTGC triggers, looking for particle “roads” based around the first Micromegas strip channel to fire for each VMM; this is explained in Section 5.2.8.3.

The Pad Trigger runs in the Rim Crates on the NSWs, while the sTGC Strip Trigger and the Micromegas Trigger algorithms run in the NSW-TP, located off the detector in the Service Cavern USA15. At each bunch crossing, the NSW-TP looks for local track segments that point to the IP, and sends them to the Sector Logic to corroborate muon triggers from the EM-TGC. The track segments are inputs to the L1-Muon trigger decision, described in Section 7.3.

5.2.8.1 sTGC Pad Trigger The VMM sends a discriminated pulse to the pad-TDS from its threshold-crossing circuit as described in Section 5.2.7.2.

The pad signals are captured every BC. The pad-TDS has per-channel programmable delays that compensate for the different time-of-flight and pad trace lengths to the VMM. At the end of the BC, the firing status of all the channels is multiplexed and serialised by the pad-TDS, and the vector of pad “hit” bits in a gas gap is sent within one BC to the Pad Trigger Board in the Rim Crate for the sector.

To reduce the number of strip channels to be read out, while keeping the number of physical pads small, pad patterns are staggered across the four layers of each detector wedge by half a pad in both directions to make logical towers corresponding to RoIs built from virtual pads one quarter the area of a physical pad (about 4 cm in radial extension), typically corresponding to about 13 strips). The geometry of the pad offsets between the two wedges in each sector ensures that these logical towers point toward the IP.

Because a pad-TDS reads out at most 104 pads per cathode board, this is the maximum that can be used in the trigger. Some pad cathode boards of the innermost quadruplets have up to 112 pads, and in these cases the pads at the smallest radii are not used in the trigger; however these pads are not in the region $|\eta| < 2.4$ covered by the EM-TGC trigger that the NSW trigger is required to confirm.

The Pad Trigger uses the inputs from the pad-TDS on each of the four layers of the three quadruplets in each of the two sTGC wedges of its sector to tag a bunch crossing and identify the band of strips passing through the triggered tower in each layer that must be selected for readout from the strip-TDS. Every pattern corresponding to a logical pad tower is checked to see if at least three out of four layers in both quadruplets have hits. There are approximately 4000 possible trigger “candidates” (logical pad towers corresponding to tracks from the IP) in a Large sector and about 1800 in a Small sector, which are stored in LUTs.

The ϕ -ID of towers satisfying these criteria is sent to the NSW-TP, as well as directly to the FELIX for readout. The Band-ID is sent to the strip-TDS on the sFEBs to request the corresponding band of strips in each layer for their associated BCID.

5.2.8.2 sTGC Strip Trigger Inputs. Each layer of each quadruplet has three (four for the inner quadruplets) strip-TDS chips on its sFEB. At each BC, the Pad Trigger can request up to four strip-TDS chips per wedge-layer, and at most one candidate per strip-TDS, to select and send data for a band of strips. The strip-TDS in each layer that holds a selected band transmits its digitised strip charges for that BCID to a Router in the Rim Crate. The data transmitted include the BCID, band-ID, ϕ -ID and the strip charges. Although data from 17 strips (the 13 in the RoI and two neighbouring strips on either side) are serialised, there is only enough time to send 120 bits of data from the strip-TDS in one BC. Since the size of a muon cluster is typically around four to five active strips, it is possible to reduce the number of strips transmitted from 17 to 14 and add one bit indicating whether the highest or lowest 14 are selected. The Router has inputs from nine strip-TDS chips; however, only four TDS inputs are active for any given BCID. For each BCID, the Router selects the four active strip-TDS inputs requested by the Pad Trigger and sends their data to the centroid-finding and track-extrapolation logic in the NSW-TP via four optical fibres per layer. The sTGC strip trigger algorithm in the NSW-TP is described in Section 7.3.3.3.

5.2.8.3 Micromegas Trigger Inputs. The address of the first strip channel in a BC to fire for each VMM is transmitted to the ADDC, which aggregates and forwards the address data (as described in Section 5.2.7.4) to the NSW-TP. At most eight hits per ART are forwarded to the NSW-TP per BCID (see Section 5.2.7.4). The Micromegas trigger algorithm in the NSW-TP is described in Section 7.3.3.4.

5.2.8.4 NSW Trigger Processor. The right-hand box of Figure 5.10 illustrates the connectivity of the NSW-TP, which is located off the detector, and described in detail in Section 7.3.3. The sTGC strip trigger and the Micromegas trigger algorithms run independently on separate FPGAs in the NSW-TP, which then merges the sTGC and Micromegas candidate segments and sends at most eight merged NSW candidate segments for each BCID to the Sector Logic to be compared with EM-TGC muon candidates. Inputs and outputs are also sent to the FELIX.

5.2.9 Data Acquisition and Readout

Micromegas strips are read out [85] through the MMFE8; sTGC Strips are read out through the sFEBs, and pads and wire groups through the pFEBs. All data from all Micromegas readout strips and sTGC pads, wire groups and strips with signals above threshold (along with the signals of the neighbouring channels in the case of the strips) are digitised using the 10-bit ADC of the VMM ASICs on the respective FEBs. The digitised signals are buffered on the VMMs until a L1A¹⁵ is received from the CTP. Zero data loss is guaranteed for a maximum latency of 16 μ s, but if the VMM readout is delayed for longer, new data are lost; however, this far exceeds the latency requirements for Run 3, unchanged from Runs 1 and 2. The ROC ASIC on all of these FEBs (described above in Section 5.2.7), receives L0A and L1A bits from the CTP via FELIX (see Section 7.3). On receipt of a L1A¹⁶, the ROC reads out data from the VMMs and transmits it to the FELIX via the GBTx on the L1DDCs mounted on the wedge.

5.2.10 Temperature and Magnetic Field Sensors

The magnetic field is complicated and highly non-uniform in the region occupied by the NSW. To measure it *in situ*, twelve 3D Hall probe cards per Large sector, similar to those that were mounted on the MDT chambers of the original EI wheels [3], are mounted on the inner (IP-facing) faces of each of the Large-sector sTGC wedges in the plane farthest from the IP (four per quadruplet).

All Micromegas wedges have nine temperature sensors (NTC 10K thermistors) mounted on the same surface as the alignment source platforms (the surface facing toward the corresponding sTGC wedge). All sTGC wedges also have nine of these temperature sensors, epoxied in locations spread over the full surface of the non-alignment (outward) face. All Micromegas and sTGC wedges have four additional NTC sensors, mounted on the cooling water inlet and outlet pipes on either edge of the wedge (Section 5.2.11).

The B-field sensors and temperature sensors are all read out through ELMB modules [44] in crates at the rim of the NSW. ELMB is a general-purpose plug-on board developed by the ATLAS collaboration for various detector control tasks, and used since the beginning of Run 1.

5.2.11 Services

The NSWs re-use, as far as possible, the services installed for the original EI wheels, including the gas racks that supplied the TGCs, which use the same n-pentane/CO₂ mixture as the TGCs, MDTs and CSCs. The MDT gas system was modified to re-purpose the racks supplying the original EI wheel. The former

¹⁵ From Run 4, this will be a L0A.

¹⁶ L0A at HL-LHC.

CSC gas system was modified to supply the Micromegas detectors with a new gas mixture. The TGC gas system distribution racks were modified to allow the supply of the new chambers in the NSW.

The LV requirements of the NSW are different from those of the original EI wheels, and required a new system, described in Section 5.2.11.4. Its greater power consumption required upgrades to the chilled water cooling systems of ATLAS, which now provide 175 kW to each side of ATLAS.

5.2.11.1 Gas Distribution

Micromegas Gas The Micromegas use a mixture of 93% argon, 5% CO₂ and 2% isobutane (iC₄H₁₀). The mixer is required to supply a nominal flow of 600 L/h, with 32 flow cells calibrated to receive the same mixture at a flow of 0 to 30 L/h per flow cell. The Micromegas gas is supplied to each wheel through 16 input channels, each serving two wedges in the same plane of azimuthally adjacent Large or Small sectors. A bundle of eight inlet pipes goes to the upper half of each wheel and another to the lower half. Four of the eight pipes serve one quadrant, each supplying two adjacent Large or two adjacent Small wedges. Each pipe then splits into four at a manifold, with an input line going to each edge of each of the two wedges served. Gas enters the Micromegas wedges at the outer rim. Impedances in the input lines and manifolds respectively ensure uniform gas flow through every wedge of the same type and through every layer of a wedge. The gas is provided to each detector layer in every sector via two intermediate buffer volumes. The average gas pressure in the chambers is 0.1 to 0.2 kPa above atmospheric pressure, depending on the height of each chamber. Gas is provided simultaneously to both inlets of the outer module and distributed by two small pipes to the inner module. The outlet of the inner module goes to the return line. Outlet lines serving the same groupings of detectors as the inlet lines are connected at the inner end of each wedge.

sTGC Gas The sTGCs use the same strongly quenching mixture of 45% n-pentane and 55% CO₂ as the existing TGCs, also at atmospheric pressure. The TGC gas system was modified to add 12 channels for the n-pentane:CO₂ mixture and two channels for the CO₂ envelope. The CO₂ channel surrounding the sTGC gas volumes, described in Section 5.2.5.1, is exhausted to a line equipped with flammable gas sniffers, so that in case flammable gas should leak from the detectors, it would be detected in the output of the CO₂ channel, and not escape into the ATLAS cavern.

The maximum flow rate for the entire TGC gas system for Run 3, including the sTGC, remains 5000 L/h. This does not constrain the sTGC flow rate, and will change after LS3, when the gas-mixing system will be modified. Gas is supplied to each wheel through two main gas input channels, each connected to a manifold distributing the gas to the eight sectors of the upper or lower half of the wheel. Calibrated flow restrictors ensure that all wedges, Large or Small, receive the same number of volume exchanges per day, irrespective of their position on the wheel. To ensure uniform flow and avoid damaging multiple layers in case of accidental contamination with dirt or debris, each layer of the wedge is separately supplied. Gas enters each layer at one side of the small end of the innermost quadruplet and follows a serpentine path around the ends of the wire supports inside the gas gap, until it reaches the corner of the large end of the same side of the quadruplet, passes through a tube into the second quadruplet, and then passes in the same way to the outermost quadruplet, finally exiting at the outer radius of the wedge. There are thus four gas input channels entering the wedge at the inner radius and four gas outlet channels leaving it at the outer radius, all on the same edge of the wedge.

5.2.11.2 Grounding The grounding of the MS was designed [85] to follow the overall ATLAS policy [98]: for the definition of the DC level, and for safety, all wedges are connected from one point to the ATLAS structure ground. This “star” grounding does not allow currents to circulate in ground loops. It is, however, difficult to adhere strictly to this ideal in practice, and additional grounding connections were added in many places during construction and commissioning of the sTGC and Micromegas detectors to suppress noise.

5.2.11.3 High Voltage Distribution The sTGC anode wires are biased at 2.8 kV during running for optimal efficiency. A single HV is supplied to all the wire groups in a single gas volume of each quadruplet, except for the innermost quadruplets, where the wires are divided in two, and the inner- and outer-radius ends of each gas volume receive separate HV. Each of the 32 wedges in each NSW thus requires sixteen HV channels in total, connected to the wedge with a single, multi-channel Lemo-Redel connector. The Safe High Voltage (SHV) cables from the distributors connect to patch-panels on the rim of the wheel. Like all the HV and LV systems of the ATLAS MDT, RPC and TGC, the NSW HV distribution uses the CAEN EASY (Embedded Assembly SYstem) Crate system, with EASY 3000 crates containing the power supplies and HV modules placed in racks in UX15, and mainframes and branch controllers in the USA15 cavern. The AC-DC power supplies in the US15 service cavern are the same that were used for the TGC of the original EI-wheel, and the original power cables are also reused. The HV distributors occupy the two racks in the main UX15 experimental cavern previously used for the original EI wheel TGCs. The HV modules that powered the original EI wheel TGCs are re-used, supplemented with new, additional radiation-hard modules. Their cables were all replaced with new ones for Run 3.

5.2.11.4 Low Voltage Distribution The 7500 on-detector electronics boards of the NSW draw about 110 kW of power in aggregate. The boards are all powered by FEAST ASICs [89], which can operate from an input supply of at most 12 V. As mechanical and thermal considerations limit the permissible cable bulk and require a supply of at least 48 V DC, a power conversion stage on the detector is therefore necessary: the ICS.

The ICS is required to withstand up to 96 Gy of ionising radiation, up to 5.8×10^{12} of 1 MeV-equivalent n/cm^2 , and single-event hadron fluences exceeding 10^{12} p/cm^2 ($E > 20$ MeV), and to function in a magnetic field of up to (or perhaps exceeding) 0.5 T. The ICSs are installed near the rim of the NSW, water-cooled, monitored and controlled remotely.

Each ICS module provides eight LV channels. The Micromegas system requires 80 ICSs (five per sector), and the sTGC system needs 48 ICSs (three per sector), for a total of 1024 LV channels in all.

5.2.11.5 Cooling Water Numerous chips on the front-end boards of both Micromegas and sTGC wedges require water cooling, as do the L1DDCs of both technologies and the ADDCs of the Micromegas, as well as the ICS for the LV power, the repeaters, and the electronics in the Rim Crates at the edge of the wheel. For each wheel, two cooling loops supply the eight upper and lower sectors of the wheel respectively. The maximum height difference between elements supplied by the same loop is therefore 3.6 m. Cooling pipes from the cooling station enter each wheel via the flexible chains. A manifold splits the input flow into twenty individual lines, eight for the Micromegas sectors, eight for the sTGC sectors, and four supplying the rim electronics. In each line, flow can be switched off by means of a manual cut-off valve, and adjusted by means of a needle valve. An identical manifold, with the cut-off valves but without needle valves, combines the return lines from the sectors and rim electronics.

The Micromegas wedges take advantage of the hollow spacer frame between them, which contains integrated stainless steel cooling channels. The sTGC have no access to the central spacer, so their cooling water circulates along the radii of the wedges: each sTGC wedge has two cooling loops, one on each long edge, outside the main Faraday cage, with both the inlet and outlet at the outer radius. The cooling pipes are made of copper, bent to route them through the hot zones of all the FEBs, and thermally coupled via conductive foam and metal connectors to the heat-emitting chips. They then loop through the L1DDC at the outer radius before exiting the wedge. To minimise the risk of leaks, the cooling loop on each side is a single continuous loop, with electrically insulating connectors placed as far as possible from the wedge electronics.

The FEBs have built-in temperature monitoring, and temperature sensors are also mounted on the input and outputs of the cooling pipes and on the outer surfaces of the sTGC and Micromegas wedges, as described in Section 5.2.10, ensuring that a cooling failure would be rapidly detected.

5.3 Upgrades to the Muon Spectrometer Barrel Coverage

This section describes all upgrades to the Barrel MS since Run 1. MDT chambers serve as the primary precision tracker for the entire barrel of the ATLAS MS, as well as for the EM and EO stations of the endcaps; however, because of the relatively large 30 mm diameter of the individual drift tubes, they have long drift times, are subject to high occupancies in regions where background rates are high, and there are some regions of the detector where space is very tight. The RPCs that provide the barrel muon trigger are subject to similar constraints. In response to these challenges, new, smaller-volume versions of both technologies have been developed to fill gaps in coverage in the original MS, and improve trigger coverage in the complicated transition region, where tracks pass through both barrel and endcap chambers. The sMDT and new RPC technologies are described in Sections 5.3.1 and 5.3.2 respectively. Chambers added during LS1 and Run 2 are described in Section 5.3.3, and those added in LS2 in Section 5.3.4.

5.3.1 Small-Diameter MDT Chambers

For regions where high backgrounds or tight space constraints make the use of traditional MDTs impractical, new small-diameter MDT (sMDT) [99, 100] chambers have been developed with a 15 mm tube diameter: half the diameter of the drift tubes of the MDT chambers. Operated with the same mixture of 93% argon and 7% CO₂ gas at 3 bar, with the same gas gain of 20 000, and read out by the same eight-channel ASD chips, the spatial resolution of the sMDT as a function of the drift radius is similar to that of the MDT [101] (see Section 5.3.4). The average sMDT drift tube resolution is expected to remain below or around 110 μm , even for the maximum background rates expected at the HL-LHC in the inner barrel region of ATLAS.

The sMDT spatial resolution and efficiency remain comparable to those obtained with MDTs in the absence of any background rate, even when rates are increased by roughly an order of magnitude, because the smaller tube diameter reduces the tube cross section exposed to background radiation and the maximum drift time from about 720 ns to 175 ns and leads to a strong suppression of the effects of space charge caused by the photon and neutron background radiation in the cavern [99, 100]. sMDT chambers have shown no ageing up to a charge accumulation on the wire of 9 C/cm [99, 100], equivalent to several HL-LHC lifetimes in all regions of the detector.

5.3.2 New Thin-Gap Resistive Plate Chambers

Like the MDTs, the RPCs used in the trigger can also be made more compact and robust, and faster, by shrinking the gas gap. Each singlet of the ATLAS new RPC with thinner gas gaps and improved electronics, (“new RPC”), consists of a gas gap of 1 mm between two 1.2 mm-thick graphite-coated HPL (Bakelite) electrodes, with the precise gap distance maintained by a grid of insulating polycarbonate spacers placed at 7 cm intervals.

The new chambers are operated with the same gas mixture¹⁷ as the original ATLAS RPCs, a mixture of $C_2H_2F_4$ (94.7%)- C_4H_{10} (5%)- SF_6 (0.3%), and reach full efficiency at a relatively modest HV of 5.6 kV with 7 to 10 times smaller average avalanche charge than the legacy RPCs. The rate capability and lifetime are correspondingly enhanced by a factor 10 with respect to the legacy RPCs [103], such that the chambers can sustain a 1 kHz/cm^2 counting rate over the lifetime of the HL-LHC, at least twice the requirement for the barrel inner layer [104]. The time resolution of the new RPCs is better than 350 ps due to the smaller avalanches, resulting in a triplet time resolution of about 200 ps.

5.3.3 Improved Coverage in the Vicinity of Barrel Feet and Elevator Shafts

In the original Run 1 configuration, there were acceptance holes in the bottom sector (Sector 13) of the muon spectrometer barrel middle and outer layers due to the gaps required to allow elevators to pass. In 2013, during LS1, these acceptance gaps in the outer barrel layer were projectively covered by the “BOE” chambers, mounted below the outer barrel layer. The BOE are standard MDT and RPC detectors of the original types used throughout the ATLAS MS.

In addition, two “BME” chambers were installed in Sector 13 in 2014, also during LS1, to cover the elevator holes in the barrel middle layer. The BME stations use sMDT (operated with spare electronics from the original MDT construction) and new RPC, with new analogue front-end electronics, similar to the BIS78 chambers described in Section 5.3.4. They are mounted on rails, which allow them to move into a parking position when the elevators are in use. The BME sMDT chambers were operated in ATLAS throughout Run 2, improving the muon momentum resolution in these sectors [105, 106] by providing a third track point measurement.

Twelve additional sMDT chambers (BMG) were then installed in the barrel middle layer during the 2016/17 winter shutdown to fill acceptance gaps inside the detector feet in Sectors 12 and 14 that support the calorimeters, the NSW and the endcap toroid cryostats. These were the first sMDT chambers operated using the newer front-end electronics also used for the BIS78 upgrade, described in Section 5.3.4. The locations of all the chambers described in this section are shown in Figure 5.11.

5.3.4 New BIS78 Muon Chambers

While the requirement of a matching segment in the NSW will reduce the L1 single-muon trigger rate in the $|\eta| > 1.3$ region, a similar background exists in the region covered partly by barrel and partly by endcap chambers ($1.0 < |\eta| < 1.3$). In the original Run 1 configuration the L1 trigger rate for muons with $p_T > 20 \text{ GeV}$ originating from this small region was about 20% of the total L1 trigger rate. As in the region covered by the NSW, this fake-muon background can also be reduced by requiring a coincidence between

¹⁷ Studies are in progress to find a suitable replacement to conform to environmental requirements at the HL-LHC [102].

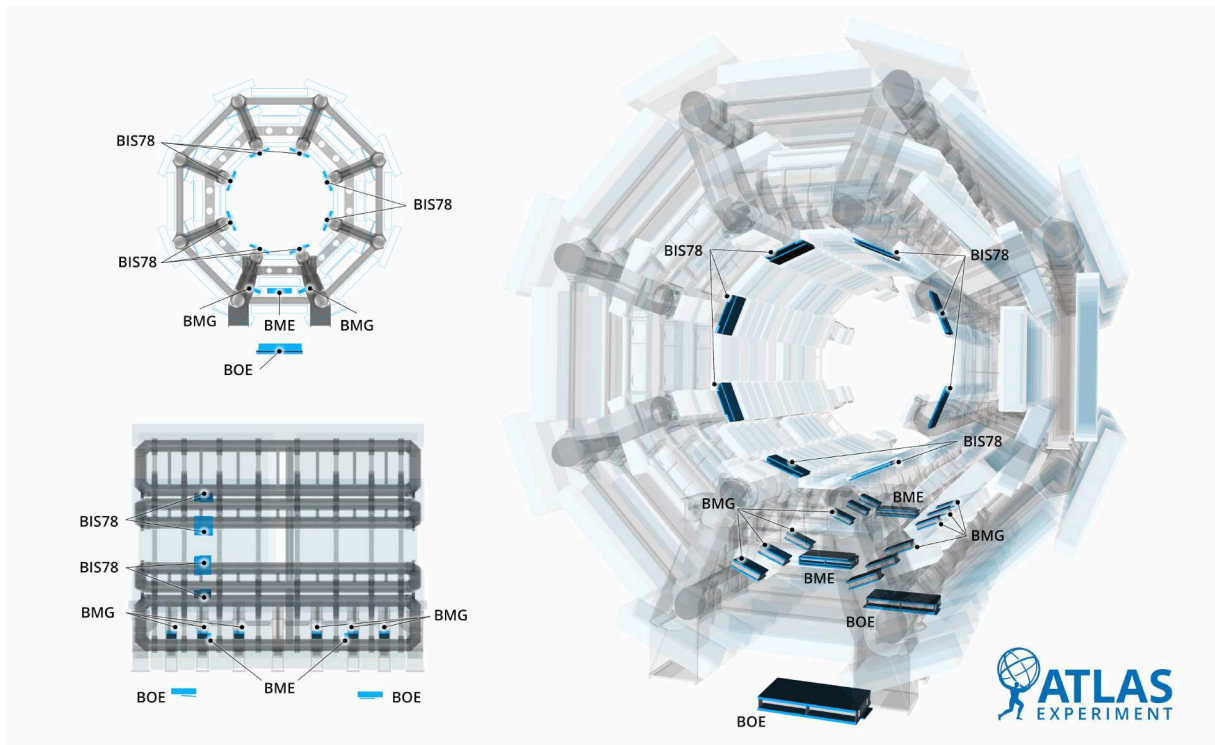


Figure 5.11: Upgraded barrel muon chambers are shown in solid blue, and include the two BOE MDT chambers and the two BME sMDT chambers, covering the elevator shafts, the twelve BMG sMDT chambers in the calorimeter feet, and the eight sMDT and new RPC integrated BIS78 modules on Side A of the transition region between the barrel and endcap.

the trigger chambers of the EM wheels and an additional inner-layer chamber (see Figure 5.2). In the Large sectors of the MS transition region, the additional trigger inner layer is provided by the EIL4 TGC chambers (mounted between the barrel toroid coils as shown in Figure 5.1(b)), but in the Small sectors, the barrel toroid coils preclude the existence of “EIS” chambers.

To fill these gaps in the Small sectors of the inner layer of the endcap, the barrel inner layer extends into the transition region, with BIS chambers installed on the surface of the barrel toroid magnet coils. During Runs 1 and 2, these BIS7 and BIS8 chambers served only for tracking, and were not part of the trigger: like the other barrel-inner chambers, they consisted only of MDTs. Because of the limited space between the original EI wheel and the barrel toroid magnet coils (see Figure 5.12(a)) the original BIS8 chambers comprised only a single multilayer of MDTs.

In order to add triggering capabilities in this region and thus reduce the trigger rate in the Small sectors, the MDTs had to be replaced by a more compact technology with space to add RPCs. Both the BIS7 and BIS8 MDT chambers on side A of ATLAS were removed during LS2, and replaced by eight new integrated modules combining integrated sMDT BIS7/BIS8 chambers and separate BIS7 and BIS8 new RPCs in a single mechanical structure [62] (see Figure 5.12(b)). Because the new integrated modules cover the areas formerly covered by the BIS7 and BIS8 chambers (see Figure 5.11), they are called “BIS78”. The corresponding BIS7 and BIS8 chambers on side C will be replaced by BIS78 during LS3 for operation at HL-LHC.

The coverage of the range $1 < \eta < 1.3$ with the new RPCs in BIS78 is illustrated in Figure 5.13, while the




Figure 5.12: (a) Placement of the BIS78 sMDT and new RPC chambers in the Small sectors of the inner endcap layer at the interface with the NSW and (b) 3D model of a BIS78 module.

resulting anticipated rate reduction is shown in Figure 5.2. The new, smaller-volume technologies are also better able to withstand the high background rates anticipated in this region at the HL-LHC. The BIS78 chambers constitute a pilot project for the complete replacement of the BIS MDT layer by the same type of integrated sMDT and new RPC chambers and the addition of new RPC triplets to the BIL layer after Run 3 for operation at HL-LHC [62]. The impact of the improved time resolution that could be obtained at the HL-LHC by using new RPCs more extensively in the BI layer will also be studied with the BIS78 chambers in Run 3.

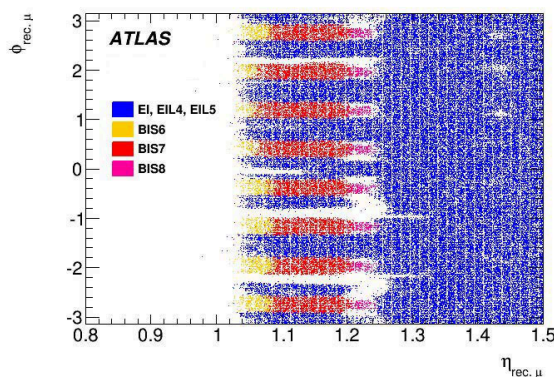


Figure 5.13: The figure shows the $\eta - \phi$ distribution of reconstructed muons with $p_T > 20$ GeV and associated to an endcap trigger (from the TGCs of the EM wheels) from track segments in the EI chambers (blue) and in the BIS chambers (all other colours).

Like the BIS7 MDT chambers they replace, the BIS78 sMDT chambers consist of two four-layer multilayers of small-diameter aluminium drift tubes. The two multilayers are separated by a 45.5 mm-high aluminium spacer frame, for an overall sMDT module height of 240 mm (see Figure 5.14). The radially outer multilayer, ML2, covers the full area formerly covered by the BIS7 and BIS8 chambers, while the inner multilayer, ML1, covers only the BIS7 area, and is thus always shorter than ML2 by 30 tubes per layer, to make space for the BIS8 new RPC as shown in the upper part of Figure 5.14. The BIS78 outer multilayers, ML2,




Figure 5.14: Cross-sectional view of BIS78 module, with inset 3D view (bottom right) and zoomed view of the new RPC triplet (bottom left).

mounted directly against the barrel toroid coils, carry the rail supports, the in-plane alignment monitoring system, the external optical alignment sensors connecting chambers to each other (which are the same as for the Runs 1 and 2 alignment configuration) and the magnetic field sensors. The new light sources of optical monitoring systems connecting the barrel to the NSWs are mounted on the inner multilayer, ML1, in the gap between the sMDT and the new RPC chambers. The new chambers have twice as many electronics channels as the previous MDTs, in total almost 11 000.

Most tubes in the BIS78 sMDT chambers are 1660 mm long, the full azimuthal width of the chambers. The final 12 tubes per layer at the BIS8 end are in all cases shortened to 1000 mm, giving the chambers the "T"-shape seen in the lower right inset of Figure 5.14 and in Figure 5.12(b) that allows them to interleave with the NSW. In order to accommodate NSW support structures, however, chambers in Sectors 2, 12, 14 and 16 have cutouts where 30 tubes per layer in the BIS7 region have a reduced length of 1530 mm. Because the NSW support structures are not radially symmetric, the lengths of the BIS78 chambers are also different from one sector to another, with 108 tubes per ML1 layer and 78 per ML2 layer for Sectors 2, 4 and 6 and, for all the other sectors, 96 per ML1 layer and 66 per ML2 layer. The chambers in the uppermost Sectors, 4 and 6 are identical, as are those in Sectors 8 and 10; each of the four bottom chambers is unique. There are thus six different sMDT chamber types in the eight Small azimuthal sectors. The chambers that will eventually be installed on Side C of the detector (after LS3) are mirror images of the corresponding Side A chambers with respect to the mid-plane at $z = 0$.

The BIS7 and BIS8 new RPC each consist of three new RPC detector singlets, as described in Section 5.3.2. With the small gas gap and electrode thickness, these new RPC singlets are just 11.8 mm thick. Their readout panels consist of two 0.3 mm-thick PCBs carrying the ground plane and the strip pattern, respectively, which are glued to either side of a 3 mm-thick FOREX[®] polyvinyl chloride (PVC) foam spacer. The ground planes of the two panels constitute the singlet Faraday cage into which the frontend electronics are integrated, in order to fully exploit its high sensitivity. The RPC singlets are electrically completely decoupled from each other, and each new RPC singlet is individually read out on each side of the detector plane with η and ϕ strip panels, respectively. Each of the eight BIS78 new RPC stations has 544 η and 544 ϕ strips, so there are, in total, 8704 new RPC readout channels.

The new RPC in the BIS78 chambers are equipped with new highly sensitive readout electronics using silicon Bipolar Junction Transistor (BJT) charge preamplifiers with highly increased sensitivity of better than 4 mV/fC at 1700 electrons RMS noise amplitude, less than 600 ps rise time and up to 100 MHz

bandwidth [107]. The frontend boards also contain new fast 4-channel discriminator ASICs in SiGe Bipolar CMOS (BiCMOS) technology with 3 mV minimum threshold and 500 MHz bandwidth. The discriminated LVDS pulses from each BIS78 module are transmitted through a 5 m-long cable to a set of 18 external Time-to-Digital Converter (TDC) boards employing High Performance Time to Digital Converter (HPTDC) ASICs [108]. The design of the front-end electronics has been made particularly robust, with low power consumption of less than 20 mW/channel, radiation hardness, and Electrostatic Discharge (ESD) protection, as they cannot be replaced after the chamber has been constructed.

The new RPC triplets fit into aluminium support frames 55 mm high, (see bottom left image of Figure 5.12(a)) and are pressed together by 2 mm-thick pre-bent aluminium plates at the top and the bottom, which distribute the force. sMDT chambers and new RPC modules are mounted independently, and electrically insulated from each other on the BIS MDT rail system using support brackets passing between the chambers in the z direction. They fit within the allowed height envelope of 290 mm. The plastic gas inlets of the new RPC are designed as part of the gas-gap frame, increasing their mechanical stability and reducing external stress in order to prevent cracking. The BIS7 new RPCs all have an azimuthal dimension of 1840 mm. Those in Sectors 2, 4 and 6 have a z dimension of 1180 mm, while for Sectors 8, 10, 12, 14 and 16 the z length is reduced to 990 mm. The BIS8 new RPCs all measure 440 mm in z . Those in Sectors 4, 6, 8 and 10 occupy the full 1840 mm azimuthally, while those in Sectors 2, 12, 14 and 16 are reduced in azimuth to 1720 mm.

5.3.5 BIS78 Data Acquisition and Trigger System

The BIS78 sMDT are read out in a very similar way to the legacy MDT, except that new HPTDC ASICs [108] replaced the original TDC (as for the BMG chambers described in Section 5.3.3). The front-end boards of the sMDT chambers each contain three eight-channel ASD chips [109] to read out 24 tubes (as for the legacy MDTs) and one HPTDC ASIC. A radiation-tolerant Actel ProASIC3 FPGA is used on each board to configure the ASD and HPTDC chips. The HPTDCs send the hit information via twisted pair cables to an MDT Chamber Service Module (CSM), of which there are two per chamber, mounted in accessible places on the barrel toroid magnet coils. The CSMs transmit the serialised data via optical fibres to the legacy Muon ReadOut Drivers (MRODs) and also send the ATLAS clock and trigger signals to the HPTDCs.

The readout for the RPCs is based on the new FELIX [110] system described in Section 7.6.1, with new electronics of the type that will be used in the Phase II upgrades of the RPC system. The digitised hit times from the strips of a BIS78 new RPC chamber are serialised and transmitted via a Gigabit Optical Link (GOL) [111] to the trigger Pad board, which is mounted on the chamber. The Pad board uses a Xilinx Kintex-7 FPGA [112] to determine the local trigger coincidences from the RPC triplet hits. Radiation-tolerant firmware design employing Soft Error Mitigation (SEM) core and triple redundancy logic is used to ensure robustness against Single Event Upsets (SEUs) and TID radiation effects. The RPC data transmission to the FELIX data acquisition system is handled by a GBTx chip on the Pad board, with a GBT-SCA for monitoring and configuration, while the trigger candidate information is transmitted by the FPGA via an optical fibre to the off-detector endcap Sector Logic board. The endcap Sector Logic boards combine the trigger information of the BIS78 RPCs, the TGCs and the NSWs to generate the L1 endcap muon trigger, as described in Chapter 7.

6 Forward Detectors

The ATLAS Forward detectors are a set of four detectors installed along the LHC beampipe at different distances from the IP. These four detectors are used for different purposes, all of them using the information carried by particles whose detection is missed by the rapidity coverage of the ATLAS central detector. The physics topics that can be studied with these detectors range from total cross section measurement (ALFA, described in Section 6.2), to diffractive physics (AFP — Section 6.3 — and ALFA), to heavy ion physics (ZDC, Section 6.4), to absolute luminosity determination (LUCID, Section 6.1). The layout of the forward detectors is shown in Figure 6.1.

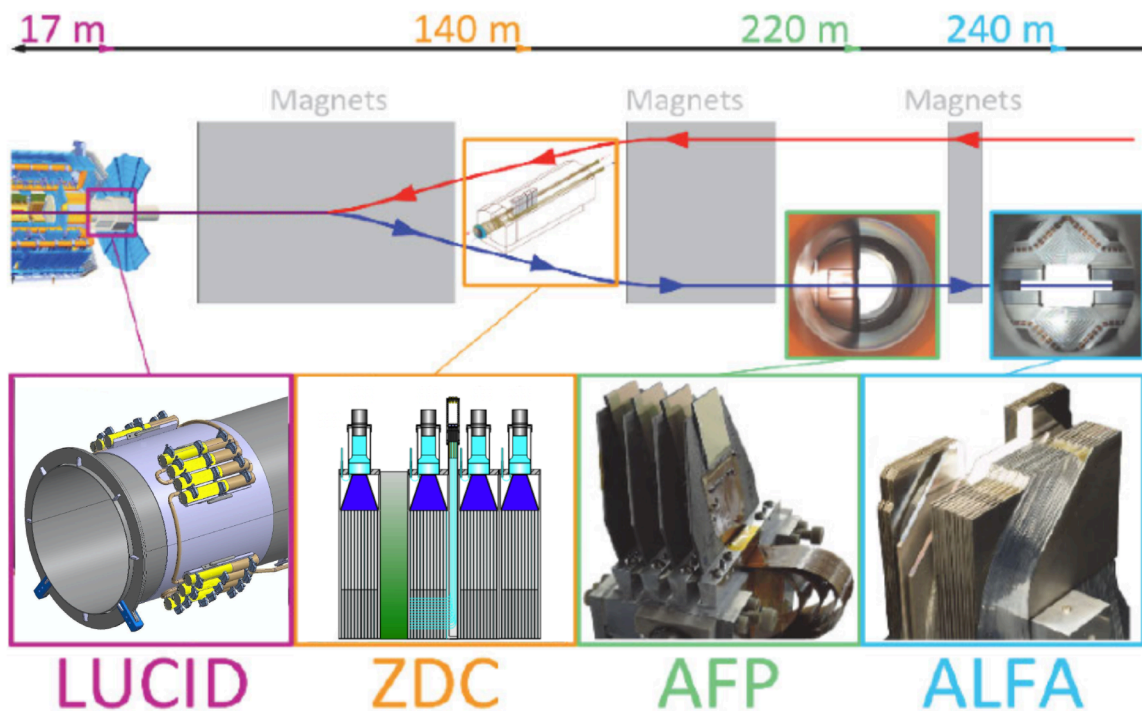


Figure 6.1: Layout of the forward detectors. The LUCID detector inside the main ATLAS volume provides a high-precision luminosity measurement; the ZDC, located just beyond the point where the two beams diverge, measure the energy of “spectator” neutrons in heavy-ion collisions; the AFP detectors measure time-of-flight and can trigger on bunch crossings producing forward protons in special low-pileup runs; and the ALFA Roman Pot detectors determine the total proton-proton cross-section by measuring elastic proton scattering at very small angles.

6.1 LUCID Detector

LUCID (LUminosity Cherenkov Integrating Detector) is designed for high-precision measurement of the luminosity delivered to the ATLAS experiment by the LHC, both in real time and, with final calibrations and a more refined analysis, for offline reconstruction and physics analyses.

LUCID is the only ATLAS luminometer capable of delivering per-bunch absolute-luminosity measurements under all the various LHC beam conditions: high-luminosity pp collisions at top energy, heavy-ion collisions,

special-purpose pp collisions at lower energies or very low luminosity conditions, and is therefore the reference luminosity detector for ATLAS. LUCID is complemented by several luminometers with different capabilities and covering various luminosity regimes: the Beam Conditions Monitor (BCM) [113], and the MBTS detector with per-bunch absolute-luminosity measurement under specific low-luminosity conditions; the per-bunch relative-luminosity measurements based on counting charged-particle tracks reconstructed in the ID at low and high luminosities, or based on counting pixel-clusters in the Pixel detector (under development); the bunch-integrated luminometers based on currents in the LAr gaps of a subset of EMEC and Forward Calorimeter (FCal) cells, and on integrated currents of the scintillator-PMTs in a selection of TileCal extended-barrel modules, mostly for high luminosity conditions. For more details, see Section 6.1.7. The luminosity reported by these luminometers is assumed to be proportional to the flux of charged particles striking the corresponding detector, integrated over a time interval ranging from one second to around one minute.

Online LUCID instantaneous luminosity measurements (bunch-integrated) are provided as fast feedback on a one- to two-second time scale to the ATLAS online-monitoring software and to the LHC control system, as are luminosity measurements from the other experiments (ALICE, CMS and LHCb). This short latency and high sampling rate are of primary importance for the LHC operators to monitor beam conditions and optimise the accelerator performance. The absolute accuracy of the reported luminosity is limited to about 5% (typically with larger uncertainties early in the running period, and improvements over the course of the data-taking year).

The *offline* determination of the absolute luminosity scale is based on the van der Meer (vdM) method [114], and carried out in dedicated runs at low luminosity ($\mathcal{L} \sim 10^{31} \text{cm}^{-2} \text{s}^{-1}$), and under accelerator conditions optimised to reduce the systematic uncertainties in the luminosity calibration [115, 116]. This calibration is extrapolated to high-luminosity physics conditions and monitored throughout the data-taking year using a methodology that involves most of the ATLAS luminometers, and that is outlined in Section 6.1.7. The overall uncertainty in the luminosity for proton collisions at $\sqrt{s} = 13 \text{ TeV}$ was determined to be 0.83% [117]. Specialised pp datasets and heavy-ion periods require their own calibrations, with uncertainties as low as 1.5% in Run 2.

The design of the LUCID detector evolved over the course of LHC Runs 1 and 2. A first version of LUCID, described in Ref. [3], was used during Run 1; the Run 2 version [118, 119], called LUCID 2, will also be used for Run 3. The changes in the detector design were dictated by the rapidly changing LHC beam conditions, and by the increase in the annual radiation dose impinging on the detector. LUCID 2 has been shown to perform very well at the highest instantaneous luminosity delivered by the LHC. It was able to provide a highly accurate luminosity determination over the entire Run 2 period, thanks in part to its built-in redundancy – each LUCID PMT can provide a luminosity measurement either alone or in combination with other PMTs – and also to the extensive use of the numerous independent, complementary and well-understood luminosity measurements provided by ID-, LAr- and TileCal-based systems, that were used to cross-check or correct the LUCID response.

6.1.1 LUCID 2 Detector Description

The LUCID 2 detector is composed of two arms placed symmetrically at about $\pm 17 \text{ m}$ on either side of the ATLAS IP. The mechanical design of one arm is shown in Figure 6.2.

The PMT detector is formed by four groups of four PMTs each, arranged symmetrically around the beampipe. The PMTs are of the R760 model manufactured by Hamamatsu. The quartz entrance window




Figure 6.2: Mechanical design of the LUCID 2 detector. The left of the diagram is toward the ATLAS cavern wall, and the cone (referred to in the text as the “VJ cone”) points toward the IP, holding LUCID inside the JF shielding, approximately midway between the EM and EO wheels of the MS. The insets show pictures of a group of PMTs and details of one PMT and its magnetic shield case.

of each PMT is used as a Cherenkov radiator. Charged particles with momentum above the Cherenkov threshold in the quartz window produce a number of photons which, on average, correspond to about 35 photoelectrons (p.e.). This signal is well above the noise threshold (1 p.e.) of the readout chain.

The PMTs are mounted around an aluminium cooling cylinder, which is in thermal contact with a water cooling pipe, keeping the environmental temperature at an acceptable level for the PMTs during the bake-out of the LHC beampipe. The PMT detectors are surrounded by a carbon fibre support cylinder that can be opened at four access points to allow easy replacement of the PMTs and their high voltage divider bases for each of the four PMT groups. The main conceptual difference in the LUCID 2 detector for Run 3 with respect to the Run 2 configuration is that the bases are now, like the PMTs, completely replaceable during the winter shutdowns of the LHC schedule. The reason for this was that about half of the PMT channels developed contact problems between the bases and the PMT pins after three years of running. It is possible that this connector problem was due to radiation, due to the location of the connectors; however, radiation tests with a gamma source have not confirmed this hypothesis. The possibility of completely replacing a PMT and its base during a winter shutdown will help to maintain a high percentage of working readout channels during each one-year running period.

Each PMT base is of the boosted type, with four HV feeding points to ensure good linearity of the PMT response up to about $100 \mu\text{A}$ of current. The whole LUCID HV system consists of 160 independent channels powered by a CAEN SY1527 module [120] housing four HV cards of type A1535N and three of type A1535SN. The LUCID 2 detector is fully integrated into the ATLAS DCS (see Section 7.7.1), which continuously monitors the HV and current of every channel as well as the temperature in various locations of the detector.

6.1.2 Prototypes for HL-LHC




Figure 6.3: The three LUCID 3 prototypes for Run 4 installed. From top left, clockwise: the JN PMT, the side C fibre detector, the side A/C JF PMTs. JN in this chapter is the same as TX1S in Figure 2.1.

In addition to the LUCID 2 PMT detector, three types of prototypes are present in the LUCID Run 3 configuration, aiming to both measure luminosity in Run 3 and to validate the LUCID upgrade project for HL-LHC (called LUCID 3). Refer to Figure 2.1 (where JN is labelled as “TX1S”) for the locations of the shielding components. The three prototypes are:

- the JF prototype is a PMT-based detector located in the inner channel of the JF shielding at a larger distance from the beampipe than LUCID 2 .
- the JN prototype is a low-rate PMT-based detector located in a shielded region behind the JN shielding surrounding the TAS;
- the Fibre prototype is a fibre-based detector located around the VJ-cone surrounding LUCID 2

Figure 6.3 contains sketch of the three prototypes. A short description of the prototypes follows, while an extensive description of the LUCID 3 upgrade project can be found in Ref. [121].

6.1.3 JF detector

Two main limitations will prevent LUCID 2 from measuring the luminosity at the HL-LHC: the increased pileup, which will lead to the saturation of the event/hit-counting algorithms, and the total radiation to which the detector will be exposed, and the resulting activation of material, which will make it unsafe to carry out the necessary annual maintenance of the PMTs. In addition, the new LHC vacuum pumping system to be installed during LS3 (VAX) would strongly interfere with LUCID if it remained in its present location. Both aspects can be mitigated by moving the detector further from the beampipe. A possible location was identified in the inner channel of the JF shielding hosting the beampipe, as shown in Figure 6.3 (bottom images), where the PMTs will be at a radius of about 30 cm, compared to about 12 cm for LUCID 2.

In the new location, both the radiation dose and the charged particle fluxes will be reduced by about 30% with respect to the LUCID 2 location. Moreover, no interference with the VAX will be present. Finally, the JF shielding is removed and brought to the surface at each end-of-year shutdown, allowing for maintenance in a safer location than that of LUCID 2 (which cannot be removed from the beampipe). The prototype detectors consist of four PMTs per side. Two PMTs of each group will be read out by a LUCROD (LUCID ReadOut Driver) card while the other two can be considered as spares, and read out in case faults develop in the initial two. On Side A, four standard R760 PMTs are installed, while on Side C three R760 and one R1635 PMT are installed. The R1635 is a smaller acceptance PMT (8 mm diameter instead of 10 mm for the R760) which has been custom-modified by Hamamatsu for LUCID to be equipped with a quartz window, needed both for the Cherenkov light production and for its radiation hardness. The smaller acceptance will allow a further limitation of the hit-counting algorithm saturation, assuming the new PMT proves suitable for LUCID use.

6.1.4 JN prototypes

A further reduction of the particle flux and radiation exposure can be obtained by placing PMTs in the shadow of the forward shielding: here the levels are reduced to about 10% with respect to the LUCID 2 location, potentially avoiding any hit-counting saturation and possibly reducing or eliminating the non-linearity with the luminosity which has been observed in LUCID 2. The proposed location is shown in Figure 6.3 (top left). The two R760 PMTs are at different radii, therefore with different levels of shadowing from the JN, and are expected to have quite different acceptances despite the small distance between them, due to the predicted steep radial dependence of the particle flux with the position behind the shielding. In this way it will be possible to have a direct measurement of the relative particle fluxes and the final PMT position in LUCID 3 can be optimised.

6.1.5 Fibre detector

During Run 2, four fibre bundles per side constituted the LUCID Fibre detector. The Cherenkov light produced by charged particles along the fibres is routed to PMTs located in a low-radiation area in the vertical channel behind the JN shielding. The fibre detector did not behave as expected, in particular due to poor long term stability. The reason for this poor behaviour is not entirely clear, but may be related both to the monitoring system of the PMT gain (based not on the Bismuth source, but on light-emitting diode (LED)s) and to the opacification of the fibres, which was not monitored. On the other hand, the charge algorithms in general, and in particular that of the fibre detector, showed remarkable linearity with the luminosity, which was not the case for the hit-counting algorithms. For this reason, a new attempt to exploit this technology is being made in Run 3 but with major modifications aimed at solving the main problems of the former fibre detector:

- Quartz fibres of type UVNSS 600/624/660 are now used, motivated by the improved radiation hardness of their Fluorine-doped silica cladding compared to that of the silicon cladding of the original fibres. There are two bundles of fibres and each bundle is read out at the end by a PMT;
- A more effective calibration and monitoring system obtained by the combination of a Bismuth source deposited onto the window of the readout PMT and an LED system to monitor ageing of the fibres.

The fibre detector is shown in Figure 6.3 (top right) in which one can see the routing of the two fibre bundles and the placement of the two readout PMTs, on top of the support structure for the VJ cone (in black in the figure) in a region of strongly reduced radiation dose. Each fibre bundle starts from the PMT window, reaches the LUCID 2 detector area and then returns back, with an open end used to inject LED light to monitor fibre ageing. A set of six LEDs with wavelength variable from the green to the UV injects light both directly to the PMT front window (prompt signal) and to the opposite end of the fibre bundle (delayed signal). The prompt and delayed signal have enough time separation to be clearly distinguished, allowing their relative amplitudes to be monitored. Changes in the relative amplitude of the two signals are expected to provide an estimate of the decrease in the light transmission of the fibre bundles due to radiation damage: this information will be used to make offline corrections to the luminosity measured during the run to compensate for this effect. The readout PMT is a Hamamatsu R7459, which has a large area window, allowing it to accommodate both the fibre bundle end and the Bismuth radioactive source used for calibration purposes.

6.1.6 LUCID 2 Readout Electronics

LUCID 2 is designed to measure the luminosity delivered in each of the 3564 bunch crossings during one LHC orbit. In order to achieve this, a fast analogue readout and signal processing chain has been designed. The signals delivered by each PMT are fed into a 15 m low-loss coaxial cable and then received by the powerful LUCROD processing card pictured in Figure 6.4(a). The relatively short cable length does not significantly degrade the fast PMT signal, which can then be amplified by a 350 MHz bandwidth amplifier and fed to a fast 12-bit ADC, sampling at a rate of 320 MHz (eight samples per LHC machine BC). Each digital waveform signal is processed in one FPGA providing two types of information:

- Signal amplitudes above a preset threshold, called *hits* in the following;
- Signal amplitudes integrated over the eight waveform samples, called *charge* in the following.

The hit signals of each readout channel are then stored in the other FPGA where they are delivered to a second processing card, the LUMAT (LUMinosity And Trigger), pictured in Figure 6.4(b), which counts the total number of hits per detector arm, runs several hit-based algorithms, and provides two L1 Trigger signals (the AND and OR of the two detector arms) to the ATLAS CTP.

The LUCROD and LUMAT boards are both implemented in the VME 9U standard. The LUCROD card accepts 16 analogue input channels and the connection to the LUMAT card is over a fast optical fibre (S-LINK [42]). Each detector arm is served by two LUCROD cards. All LUCROD cards continuously send hits to two LUMAT cards which combine the information from the two detector arms. The LUCROD cards sit in one VME crate on each detector side, placed on a service platform close to the detector, in order to limit the length of the signal cables. The two LUMAT cards sit in a VME crate placed in the ATLAS USA15 service area.

Since LUCID has to provide online luminosity measurements not only during ATLAS data-taking but also during LHC machine tuning, it can run continuously in standalone mode, independent of the ATLAS TDAQ system. A fraction of the fully reconstructed waveforms is stored locally for possible cross-checks or to study detector performance.




Figure 6.4: Photographs of the LUCID 2 readout cards: (a) LUCROD and (b) LUMAT.

6.1.7 LUCID absolute-luminosity calibration and luminosity-measurement methodology

A fundamental ingredient for a stable luminosity estimate based on LUCID is the control, at the 1% level, of the stability of the LUCID PMT gains. This is accomplished by a very effective calibration system based on the radiation emitted by the (approximately) 40 kBq ^{207}Bi source deposited on each PMT window. This source emits, among other decay products, monochromatic electrons of about 1 MeV kinetic energy, which is enough for them to cross the 1.2 mm quartz window, and to produce roughly the same amount of light as a high-energy charged particle traversing that same window. At low luminosity, for example, during heavy-ion running or in the tails of van der Meer (vdM)-calibration scans, the $\mathcal{O}(10\text{ kHz})$ rate from this monochromatic electron source represents a non-negligible background that must be subtracted from the raw LUCID signal.




Figure 6.5: Gain variation with respect to the first day of calibration in 2017. The PMT HV was corrected between collision runs aiming to keep ^{207}Bi signals average amplitude constant (red line).

The PMT gains are monitored on an approximately daily basis by recording the amplitude spectrum of the ^{207}Bi signal from each PMT during periods without beam. The typical amplitude and charge inferred from these spectra, normalised to unity at the beginning of the running period, are shown in Figure 6.5 as functions of time over the year 2017. The red points refer to the average amplitude, and the blue ones to the

average charge. The average amplitudes of the calibration spectra are kept constant within a few percent throughout the entire period, while the average charge diverges systematically away from unity, decreasing by 8% over the six-month time span shown in the figure. Since the relationship between signal amplitude and charge changes when the HV setting is increased to keep the gain constant [119], only one of the two can be stabilised throughout the data-taking year.

Keeping the amplitude constant is essential for the hit algorithms, since their luminosity measurement is based on the fraction of beam crossings in which this amplitude exceeds a given threshold (defining a “hit”); charge algorithms, in contrast, rely on the stability of the charge measurement. In early Run 2, the charge algorithms were originally perceived as the most promising. In 2015 and 2016, therefore, the choice was made to keep the average charge constant. Experience, however, demonstrated the superior performance of the hit algorithms, so that in 2017 and 2018 the average amplitude was stabilised instead. Offline corrections to the hit-counting efficiency that had been introduced in 2015 and 2016 became obsolete.

In hit-counting algorithms, the average fraction f of BCs containing a hit (typically a signal above a given threshold) is assumed to obey binomial statistics (while the pileup parameter μ is assumed to be Poisson-distributed). The luminosity per bunch crossing, \mathcal{L}_{BC} , is given by

$$\mathcal{L}_{BC} = \frac{-\ln(1-f)}{\sigma_{vis}} \cdot f_{rev}, \quad (1)$$

where f_{rev} is the LHC revolution frequency, and σ_{vis} is a calibration constant called the visible cross section that is measured by the vdM method. A hit-counting algorithm based on one of the two independent sets of eight PMTs denoted as Bi1 or Bi2, is typically chosen as the “preferred” algorithm both for online luminosity monitoring and for a first rough estimate of the offline luminosity. For the final offline luminosity, the preferred algorithm is chosen in each data-taking period (typically one year) based on a detailed analysis of its linearity and relative stability across that period. In Run 2, the preferred offline algorithm was based either on a subset of eight PMTs (in 2016 and 2017), or on a single PMT (in 2015 and 2018).

For charge algorithms, the average charge collected in a given BCID, is intrinsically proportional to the bunch luminosity \mathcal{L}_{BC} associated with that BCID:

$$\mathcal{L}_{BC} = \frac{C}{K_{cal}}, \quad (2)$$

where C is the charge averaged over a time interval, denoted as a Luminosity Block (LB)¹⁸, collected in the BCID considered¹⁹, and K_{cal} is a calibration constant to be determined by the vdM method. Because of their intrinsic linearity, charge algorithms are in principle better behaved than hit counting algorithms in high pileup conditions. The downside is that the instantaneous luminosity inferred from charge measurements is directly proportional to the PMT gains, which have been observed to vary by up to a few percent over the course of a single LHC fill. Solutions to compensate for this disadvantage are still under study.

The absolute luminosity scale of each LUCID algorithm, or equivalently the corresponding visible cross-section σ_{vis} , is measured by the vdM method under experimental conditions optimised to minimise systematic uncertainties: low pileup ($\mu \sim 0.5$) to eliminate LUCID non-linearities, at most 150 isolated bunch pairs colliding in the LHC to avoid long-range beam-beam crossings or out-of-time pileup in the luminometer electronics, and therefore low instantaneous luminosity ($\mathcal{L} \sim 10^{31} \text{ cm}^{-2} \text{ s}^{-1}$). The

¹⁸ LBs have a typical duration of about one minute, within which the instantaneous luminosity and data-taking configuration are considered to be stable, and serve as an approximate “time-unit” in ATLAS luminosity measurement.

¹⁹ Care must be taken to exclude those bunch crossings in which the LUCID analogue readout chain may be saturating.

bunch intensity is lowered to around 0.8×10^{11} protons/bunch, and the injected emittance increased to 2.5–3.5 $\mu\text{m rad}$, so as to minimise longitudinal charge leakage out of the nominally filled positions along the LHC-ring, as well as beam-beam–induced distortions of the vdM-scan curves. In addition, the crossing angle is set to zero to minimise orbit-drift and beam-beam correction uncertainties, and the β function²⁰ at the IP is increased to $\beta^* = 19\text{ m}$ to widen the luminous region in order to facilitate non-factorisation corrections [115, 116].

These luminosity calibrations, that are obtained in the so-called “vdM regime”, cannot be directly applied to physics data-taking conditions which exhibit non-zero crossing angle, with β^* values in the range 0.3 to 0.6 m, 40% lower emittance, 50% higher bunch intensity, two orders of magnitude higher pileup, up to around 2500 bunches grouped in trains, as well as three orders of magnitude larger instantaneous luminosity. This is primarily because LUCID suffers from significant pileup-dependent non-linearities that result in an overestimate of the luminosity of up to 10% at $\mu \sim 50$. These non-linearities are corrected, separately for each algorithm, using a calibration-transfer procedure based on counting custom-reconstructed charged-particle tracks emerging from inelastic collisions in randomly selected colliding-bunch crossings – an observable proportional to the number of pp interactions per BC and thus to the instantaneous luminosity.

The custom track reconstruction uses only hits in the IBL, the Pixel and the SCT detectors; the track-selection criteria are optimised for luminosity monitoring over the full range of pileup levels encountered during pp running ($0.01 < \mu < 60$). During both Run 1 and Run 2, track-counting (TC) demonstrated very good linearity with respect to μ , as well as excellent long-term stability under evolving ID conditions. Because the per-bunch TC-based luminosity is statistically limited under the low- μ conditions of vdM scans, this algorithm is not independently calibrated by the vdM method, but cross-calibrated to LUCID during an extended head-on collision period in the same LHC fill in which the vdM scans are performed. TC is used to transfer the absolute-luminosity calibration of LUCID to physics data-taking conditions, in a reference, high-luminosity LHC fill in which the measured TC/LUCID luminosity ratio is used to parameterise the μ dependence of the uncorrected LUCID response.

The linearity of the TC algorithm is cross-validated against relative-luminosity measurements based on the TileCal and LAr calorimeters. In the case of TileCal, the dynamic range of the cryostat scintillators E3 and E4 (see Figure 4.6(a)) provides sufficiently sensitive bunch-integrated luminosity measurements during head-on collisions all the way from the vdM to the physics regime. For EMEC and FCal, specialised fills with isolated bunches that cover the full range of μ values are used to compare the TC- and TileCal-based measurements with those inferred from LAr energy-flow measurements.

During Run 2, the calibration-transfer procedure was applied from one to three times per year, and proved essential to ensure the year-long stability of the LUCID response. This stability is quantified by comparing, over the course of each data-taking year, the TC-corrected, bunch-integrated LUCID luminosity with that inferred from the PMT currents in the D6 cells of the TileCal, and from the LAr-gap currents in the EMEC and the FCal. Since these luminometers are not sensitive enough to be calibrated in the low-luminosity LHC fill in which vdM scans are recorded, they are cross-calibrated to TC in a small subset of high-luminosity fills close in time to the vdM-calibration session. Additional luminometers, such as TimePix sensors [122], provide further cross-checks. Finally, monitoring the rate of reconstructed $Z \rightarrow \ell\ell$ decays, known as Z-counting, allows a fully independent check of the relative long-term stability of the ATLAS luminosity

²⁰ The β function describes the single-particle betatron motion around the central orbit, and in particular the variation of the transverse beam envelope along the beam trajectory. β^* is the value of the β function at the IP, and indicates how squeezed the beams are at the IP, by giving the distance along the beam direction after which the β function (the transverse beam size) has been multiplied by a factor of 2 ($\sqrt{2}$) compared to its value at the IP. A large β^* indicates wide and almost parallel beams, a small β^* indicates narrow, but divergent beams.

measurements, both within each running year and across multiple data-taking years [123]. This technique, which does not rely on theoretical predictions of the absolute Z cross-section, is also used to compare the integrated luminosity delivered to the ATLAS and CMS experiments.

The luminosity measurement methodology outlined above provided among the most accurate absolute-luminosity measurements at any hadron collider to date [117], and serves as a starting point for high-precision luminosity determination at the ATLAS IP in Run 3.

6.2 ALFA Roman Pot Detector

The Absolute Luminosity for ATLAS (ALFA) Roman Pot detector consists of four stations, each equipped with two scintillator detectors in an upper and lower Roman Pot. Two stations are placed on each side of the ATLAS detector, about 240 m from the IP, in the long straight sections of the LHC tunnel. ALFA is described in Ref. [124] as it was operated during special runs in Run 2 for elastic cross-section and diffraction studies at β^* values from 11 m to 2.5 km. The following paragraphs describe only the changes to allow operation of ALFA in Run 3.

The readout electronics of ALFA have aged and some readout parts were already replaced during Run 2 because of a water leak in the LHC tunnel. The scintillating fibre detectors and readout electronics are expected to age further during Run 3. Moreover, the LHC TCL6 collimator settings during insertions of the AFP Roman pots in the last years of Run 2 increased the total radiation dose of the ALFA detectors significantly, by a factor of 5 to 10 compared to the situation with no AFP insertion (and TCL6 closure).

Irradiation studies have shown that the readout motherboards and trigger boards start malfunctioning after about 500 Gy and, in particular, that the embedded ELMBs start failing after 50 Gy [125]. The Kuraray SCSF-78-SJ scintillating fibres only start to degrade above 10 kGy, with an efficiency loss of 20% and higher. The dose accumulated in Run 2 is about 20 Gy for the motherboards, whereas the dose to the fibres, strongly dependent on the distance to the beam, was measured to be in the range of 1 to 3 kGy in the last year of Run 2. To minimise radiation damage, it is therefore important that an ALFA high- β^* run occur early in Run 3. To reduce the exposure of ALFA in Run 3, a 40 cm thick iron shielding wall was erected at 221 m after TCL6 and before Q6; and two pairs of similar walls between Q6 and each of the two ALFA stations at 235 m and 243 m, as shown in Figure 6.6. The shielding reduces the radiation load to about half, which should allow ALFA to operate until the end of Run 3.

Five new motherboards were produced to replace the most irradiated boards in the outer stations and to have spares in Run 3.

During LS2 the Roman Pot movement system [124] was refurbished and all the stopper and limit switches were readjusted. In Run 2, the ALFA trigger, which operates outside the standard ATLAS latency, was connected directly to the ATLAS CTP (via a specialised ALFA-CTPIN module) to minimise the trigger path length. The firmware of this CTPIN module has been modified to provide a direct path to the CTP for the Level 1 Topological Processor (L1Topo) trigger, and to be switched to the ALFA trigger only for ALFA special runs. The Readout system is unchanged.

6.3 AFP Detector

The ATLAS Forward Proton (AFP) detector Phase I upgrade project [126] consists of four detector stations, two on either side of the ATLAS IP at ± 205 m (near stations) and ± 217 m (far stations). Together, the near




Figure 6.6: The pair of shielding walls between the ALFA far station (left, before Q7) and near (right) station in the LHC tunnel in LSS1. Another pair is situated before the near station (not shown).

and far stations on one side constitute an Arm of the detector. Each station contains a single horizontal (x) Roman pot housing the AFP detectors. The pots present a $300\ \mu\text{m}$ -thin window to the LHC beam. Each near pot contains a four-plane Silicon Tracker (SiT), which is a 3D pixel detector similar to the IBL except for a $150\ \mu\text{m}$ thin edge parallel to the long pixel direction (y) on the side of the LHC beam [127, 128]. Each far pot contains, in addition, a Time-of-Flight (ToF) detector [129] consisting of novel quartz Cherenkov hodoscopes with radiation-hard electronics behind the tracker to measure the vertex position of double-proton final states with 3 to 5 mm precision.

During the 2015-2016 winter shutdown, the C-side arm of AFP was installed, but without the ToF. The AFP detector was completed with the ToF detectors during the 2016-2017 winter shutdown. Over the course of 2017, ATLAS collected data where AFP was read out corresponding to an integrated luminosity of $20\ \text{fb}^{-1}$. From July 2017 onwards, the Roman pots were operated at a distance of $11.5\sigma_{\text{beam}} + 0.3\ \text{mm}$ from the LHC beam, which is around 1.7 mm and 2.8 mm for the far and near stations respectively for $\beta^* = 40\ \text{cm}$.

In preparation for Run 3, the AFP detector was refurbished during LS2: the movement mechanics were re-tuned and all switches were readjusted. The PCI eXtensions for Instrumentation (PXI) movement controller hardware and software were updated. Because of radiation damage during Run 2, all silicon detectors were replaced by newly produced 3D silicon pixel tracker modules of the same type, as described in Section 6.3.1. The ToF detector underwent a design change to prevent corona discharge problems in vacuum, described in Section 6.3.2. The trigger and the readout electronics and software were updated, as discussed in Section 6.3.3.

6.3.1 Silicon 3D tracker

The AFP slim-edge 3D silicon detector planes are located in the near pots. During LS2, all of them were replaced with new 3D modules of the same type. The irradiation pattern on the SiT planes is a characteristic narrow band that extends away from the LHC beam and either diagonally up or down depending on the sign of the vertical crossing angle of the colliding beams. Thus, switching SiT modules between the arms (or changing the sign of the crossing angle), exposes a fresh area of the sensor module to diffractive protons. It is expected that the new modules will operate efficiently for the first two full years of Run 3. After one year of running or an accumulated dose corresponding to about 50 fb^{-1} , a swap of detectors between the two detector arms is foreseen. New SiT detector modules may be required for the last years of Run 3 and different solutions (based on RD53 [130] or TimePix4 [131] ASICs) are under consideration.

The cooling for the SiT is based on a two-stage vortex tube system, using compressed air [132]. The 3D tracker planes are mounted in good thermal contact on top of a heat exchanger box, and cold air at -30°C circulates through the box. For Run 3, the heat exchanger efficiency was improved by about 30% by filling it with open-cell metal foam.

6.3.2 AFP Time-of-Flight detector

The AFP ToF detectors are located in the far Roman Pots. The latest version consists of solid L-shaped fused silica bars and customized Micro-Channel Plate PMTs (MCP-PMTs) with an extended lifetime which operate at low gains (order of 1000) [133]. The improvements were aimed to increase the efficiency, the lifetime as well as the radiation hardness of the detector which has been designed to operate in high radiation areas (above 400 kGy/year). They have a time resolution in the expected range of 20 to 26 ps per proton [134]. In 2017, the ToF efficiency was very low (2 to 9%), and this was understood to be because the integrated charge received by the PMTs during that first year exceeded the limits for which the PMTs were intended. In 2018, it was not possible to install the ToF refurbished with long-life PMTs because of unresolved high-voltage breakdown when the MCP-PMTs were operated in the secondary vacuum of the pot. In a change of approach, the new ToF is now read through a quartz window separating the Cherenkov light guides (inside the pot) from the MCP-PMT (outside the pot). This requires a “Cherenkov light-feedthrough” on the pot’s flange (which serves as the detector platform) which is adjustable in depth to match the specific depth of the pot. The design is shown in Figure 6.7. In this view, forward protons enter the detectors from right-top and the Cherenkov radiators are oriented at 48° with respect to the direction of the forward protons. Vacuum bellows seal the tube holding the square 4×4 multi-anode MCP-PMT [135] to the flange, allowing for height adjustment of the ToF structure.

6.3.3 AFP Trigger

While at high luminosity most bunch crossings in ATLAS produce at least one forward proton in one of the AFP detectors (and it is therefore usually unnecessary to use this information in the trigger), for special low-pileup runs with $\langle\mu\rangle \sim 1$, it is useful to trigger on events with forward protons. For Run 2 this was accomplished with a SiT-plane hit-based majority trigger which had, however, a rather long dead time of about 250 ns. A similar, but ToF-based, option was provided in the Run 2 TDC module, using the on-board FPGA to select events where a majority of ToF Cherenkov bars in a train of four successive bars fired. This option is also available in the new picoTDC module.




Figure 6.7: (a) The AFP far detector platform flange holding the 4-layer Silicon 3D Tracker (SiT, top right) and the Time-of-Flight detector (ToF, bottom left). The SiT is mounted on a cold-air heat exchanger. The normal to the SiT planes makes a 14° angle with the incoming protons to eliminate dead area between pixels. The ToF multi-anode MCP-PMT is mounted in air inside a tubular feedthrough housing. (b) shows the ends of the 4×4 L-shaped Suprasil (“quartz”) radiators (LQbars). Cherenkov light from the LQbars reaches the MCP-PMT through a quartz window that separates the PMT from the secondary safety vacuum of the detector environment.

In addition, a digital trigger module was added for use in Run 3, inserted in the digital data path between the constant-fraction discriminator modules and the TDC [136], improving the TDC resolution from 18 ps to 4 ps. This new trigger module provides programmable majority logic for each ToF train, and outputs the observed trigger pattern as a sequence of 1 to 2 ns-wide NIM pulses (a start/trigger pulse followed by the trigger status of each of the four trains) on the fast air-core trigger cable to the CTP.

6.3.4 AFP Data Acquisition

New DAQ boards [137] were designed for the AFP. They are housed in an ATCA crate. The DCS interface to the ATCA system (see Section 7.7.3) was implemented for control and monitoring, and the AFP DCS [138] was updated to control and monitor the voltages, currents, and temperatures of the new front-ends, digital Trigger modules, picoTDCs, and Versatile Link demonstrator boards (VLDBs). The OptoBoards in the LHC tunnel that were used for digital-to-optical conversion in Run 2 were replaced by faster VLDBs developed for Run 3 and beyond.

Front-end status information is collected by the DAQ, for instance the number of parity errors or module busy signals. This information is monitored and in case the DAQ decides that some SiT or ToF device is corrupted, e.g. by a single event upset, a request is sent to the DCS for power-cycling the affected device.

A new package for AFP data quality monitoring was written for Run 3 and operates within the general ATLAS multi-threaded data quality monitoring system and conforms to the ATLAS guidelines.

6.4 The Zero Degree Calorimeters

The ZDCs play a central role in the ATLAS heavy ion physics program. Their primary function is to measure the energy of “spectator” neutrons which do not participate in hadronic processes as the nuclei collide, to determine impact parameters when the nuclei do not overlap completely. The neutrons propagate in the original beam direction, with minimal deflection, carrying on average the per-nucleon beam energy. Because of its ability to observe neutrons from nuclear breakup, the ZDC plays several critical roles during heavy ion operations and data analysis:

- It is an integral component of the heavy ion triggering scheme and of some BSM searches, as it allows the deliberate enhancement and suppression of electromagnetic processes, including photonuclear and elementary photon-photon processes such as light-by-light scattering.
- It is the only part of the detector that records an unbiased selection of peripheral collisions at low multiplicities, essential for studies of angular correlations in heavy ion collisions.
- Correlation of the ZDC energy with transverse energy in the central detector demonstrates the clear geometric nature of particle production in heavy ion collisions. It also provides a well-tested means, using two-dimensional cuts, to reject both in-time and out-of-time pileup, which is key for precise physics measurements.

6.4.1 Overview of the ZDC upgrades

In Runs 1 and 2, the ZDC detector was designed to nearly fill the space available in the neutral beam absorber (TAN) region, which sits 140 m from the nominal IP, just after the beams diverge, and protects one of the LHC beamline dipoles. As shown in Figure 6.8, there are four modules on each side (forward and backward), each of which consists of 1.1 nuclear interaction lengths of tungsten plates (11 plates, each of 1 cm thickness) interleaved with a row of 1.5 mm-diameter quartz rods. Light produced by Cherenkov radiation from shower products is transported up to the top of the detector, guided through a trapezoidal prism light guide and into a Hamamatsu H6559 PMT [139] assembly. The signals from each PMT are carried by cables of approximately 200 m to USA15, where they are digitised and included in the ZDC trigger. Run 2 began with a detector very similar to that used in Run 1, after periodic replacements of the commodity-grade quartz, which was inexpensive but susceptible to radiation damage, particularly during proton-proton operation. However, the increased rates relative to Run 1 led to large anode currents, and consequently to large changes in the observed detector gain. To mitigate these effects, the PMTs were replaced with a modified version (H6559 MOD) with high voltage boosters in the last three dynode stages.

Several modifications to the detector were implemented for Run 3:

- To minimise the impact of radiation damage on the signal properties, the quartz rods were replaced with fused silica rods of the same dimensions. A combination of irradiation studies and comparisons of fused silica samples placed in the TAN (in a prototype of the Beam Rate of Neutrals relative luminosity monitors (BRAN) detector utilised by the LHC for luminosity measurements [140]), were used to determine that high concentrations of OH and H₂ doping provide additional radiation hardness. Studies of the expected dose in Run 3 have led to a final choice of fused silica rods doped with H₂, which provides good stability in the ultraviolet range up to several million Gy.

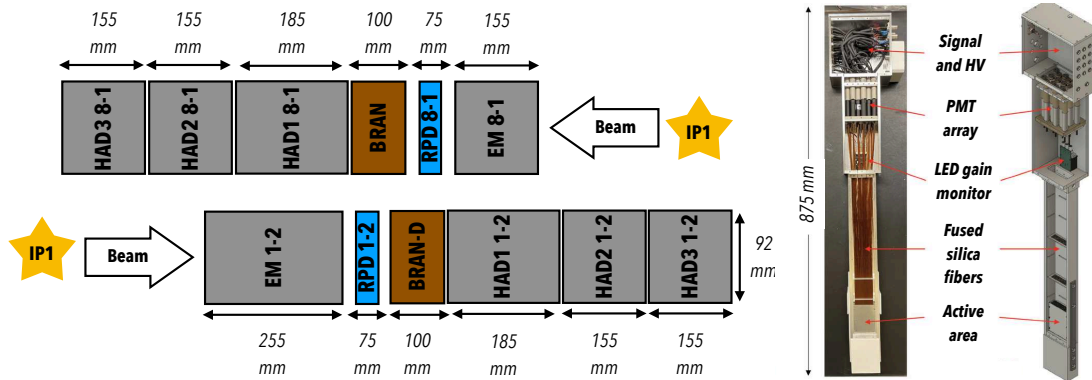


Figure 6.8: The left-hand diagram shows the arrangement of the ATLAS ZDC arms on Side A (top) and Side C (bottom), including the new RPD and indicating the different BRAN detectors being used in the first year of Run 3. The new ATLAS RPD detector is shown on the right, with a built detector shown head-on, next to a CAD rendering showing the various cable connections on the side.

- Digitisation is handled by the L1Calo Pre-processor Modules (PPMs), which run at a sampling rate of either 40 or 80 MHz. This is well matched to the signal shape induced by the long coaxial cables, which transform a fast signal into one with a rise time of 4 ns and an exponential decay time of about 25 ns, for a total signal length of about 70 ns. This was already just barely acceptable for the 75 ns bunch spacing reached in 2018, but is too long for the 50 ns expected for heavy ion operations in Run 3. To provide a faster signal with less attenuation, the coaxial cables used in Run 2 have been replaced by air-core coaxial cables similar to those used in Run 2 for the AFP detector. These reduce the attenuation by a factor of nearly 10 and the signal shaping is substantially reduced, such that nearly 97% of the charge is contained within a single bunch crossing.
- The necessary updates to the digitisation scheme will be discussed in the next section. The sampling rate is increased by a factor of 4 to 6, making the after-pulsing of the PMTs, observed in test beam data with much shorter cables, visible in typical events. While it would be desirable to find a photomultiplier tube specifically designed to reduce the impact of after-pulsing, it turns out that the reduction in the effective aperture of the photocathode is strong enough to make this option unnecessary in the next LHC run.
- A new Reaction Plane Detector (RPD) has been built, and tested in an SPS test beam, to determine the direction of the event-by-event deflection of the cloud of spectator neutrons on each side, and to measure the correlation of the directions between the two sides. The two sides are expected to show a clear anti-correlation of the deflection direction, reflecting the angle of the internuclear impact parameter, referred to as the Reaction Plane. The reaction plane angle can be utilised to study the directed flow of emitted particles, and also provides sensitivity to the presence of strong magnetic fields in the initial state of the nuclear collision. The RPD consists of a 4x4 cm array of rad-hard optical fibres of different lengths read out by Hamamatsu R1635 PMTs at the top of the TAN. These provide an effective 4x4 array of 1x1 cm cells in two dimensions. Since the longer fibres overlap the shorter ones, machine learning approaches are needed to disentangle the correlated signals between the fibres. In simulations, the resolution of the reaction plane angle is found to be excellent and comparable to similar detectors, for example in the STAR experiment at the RHIC accelerator.

6.4.2 ZDC Readout and Trigger Electronics

The shorter signals from the air-core cables dramatically mitigate the impact of both in-time and out-of-time pileup on the measured ZDC waveforms, nearly eliminating the overlapping signals from adjacent filled bunch crossings and even collisions closer in time due to debunched beam. However, this requires a much faster sampling rate than the Run 2 readout system (based on electronics originally developed for the L1 calorimeter system) can provide. The LUCROD card designed for LUCID, as described in Section 6.1, was adapted to rework the entire trigger and readout approach for the ZDC detectors. The new card has modified firmware and is referred to as the LUCROD/ZDC.

The Run 2 ZDC trigger and readout scheme split the incoming signals into two paths, one for digitising the signal and the other for providing L1 triggers for ATLAS. In the trigger path, the incoming analogue signals were summed in hardware and a threshold was applied just below the single neutron peak for each side, using NIM discriminators. The readout path involved a fourfold splitting of the incoming signals into both high and low gain, to obtain an effective 12-bit dynamic range, and applying a 12.5 ns delay to achieve an 80 MHz effective sampling rate even when using the L1Calo PPMs, which sampled at 40 MHz.

For Run 3, the LUCROD card, with special firmware for the ZDC, combines digitisation, triggering and event logging into the same card, with waveform sampling performed at 320 MHz. Each LUCROD/ZDC card provides 8 channels, sampled with 12-bit flash ADCs, with an FPGA for each channel (FPGACH), one FPGA (FPGAV) which receives data from all 8 channels simultaneously to perform trigger calculations, and one main FPGA (FPGAM) which serialises and transports the channel data to ATLAS. One LUCROD is used to digitise the signals without amplification, while providing an analogue copy to a second card which amplifies the signals by a factor of 16. Combined, the two cards provide an effective 16-bit dynamic range. The high-gain LUCROD/ZDC also runs trigger algorithms that provide the functionality previously provided by the Nuclear Instrumentation Module (NIM) modules. Peak finding is performed for each channel in the FPGACH using a threshold on the second derivative of the waveform, while FPGAV performs digital sums and uses two sets of lookup tables to form a L1 decision for the ATLAS CTP. A lookup table for each side is used to generate multiple energy thresholds – specifically, distinguishing one neutron from no activity, which tags nuclear breakup, and distinguishing five neutrons or more, which distinguishes hadronic and electromagnetic processes. A combined lookup table then incorporates the threshold level from the two sides, and provides a three-bit output word to the ATLAS CTP which reflects the correlations between the energies of the two sides, distinguishing (for example) between events with energy on only one side and those with substantial energy on both sides. Upon receipt of a L1A from ATLAS, the FPGAM serialises the data from the channel FPGACHs, prepares packets for transmission over S-LINK and manages S-LINK transmission to the ATLAS ROS. With this implementation, the LUCRODs replaces essentially all of the ZDC electronics utilised for Run 2.

6.4.3 ZDC Detector control

The ZDC PMTs are controlled by a CAEN mainframe located in USA15, with long HV cables running to distribution boxes located under the TAN, and finally to the detectors, with standard SHV cables to drive the primary PMT HV, and with multicore HV cables to drive the three boosters per PMT. The ATLAS DCS (as described in Section 7.7.1) is used to set and monitor the high voltage and current for each channel. It is crucial to track the booster currents at higher rates, to prevent damage to the PMTs. As the signals were calibrated in short intervals of several minutes throughout Run 2, based on a calibration peak from events with one neutron – at full beam energy – in either ZDC, a stable response was not required

over longer timescales, and the Run 2 detector was simply continually monitored for degradation of the observed energy of the one-neutron peak. The HV in one or more PMTs on each side was then increased when the fluctuations of that peak began to approach the noise peak seen in empty events. In Run 3, the radiation-hard fused silica rods replacing the quartz should provide good stability throughout each heavy ion run period, and minimise or eliminate any need for HV adjustments.

7 Trigger and Data Acquisition System

7.1 System Overview

As described in Section 1.3.6 and shown in Figure 7.1, the Trigger and Data Acquisition (TDAQ) system is based on a two-level trigger (event selection) system served by the Data Acquisition (DAQ) system that transports triggered data from custom subdetector electronics through to offline processing. The Level-1 Trigger (L1) trigger system is based on custom-built electronics and the High-Level Trigger (HLT) is based on software implemented on commodity computers.

The L1 Trigger, shown in Figure 7.2, uses reduced-granularity information from the calorimeters and muon system to search for signatures from high- p_T muons, electrons, photons, jets, and τ -leptons decaying into hadrons, as well as events with large missing transverse energy or large total transverse energy. It consists of the Level-1 Calorimeter (L1Calo) and Level-1 Muon (L1Muon) trigger systems, the MUCTPI, the L1Topo and the CTP. The L1Calo trigger system (Section 7.2) identifies high- E_T objects such as electrons, photons, jets, and τ -leptons. Isolation requirements may also be applied to these objects. It also selects events with large E_T^{miss} or large total transverse energy. The L1Muon trigger system (Section 7.3) selects events containing high- p_T muons, based on inputs from the RPCs in the barrel region and from the TGCs and NSWs in the endcaps, and then transmits data to the CTP via the MUCTPI. The L1Topo trigger

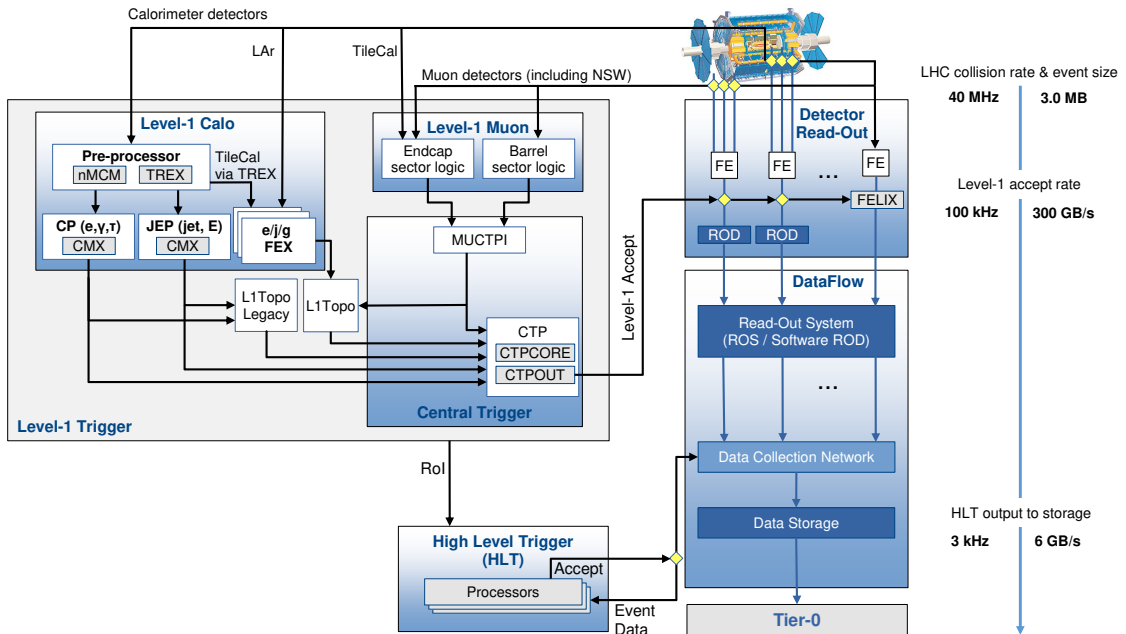


Figure 7.1: Schematic overview of the Trigger and DAQ system at the beginning of Run 3.

system (Section 7.4) takes input Trigger Objects (TOBs) containing kinematic information (e.g. E_T and $\eta - \phi$ coordinates) from the L1Calo and L1Muon systems and applies topological selections. As shown in the figure, both the legacy L1Topo and upgraded L1Topo systems are in use during the commissioning phase.

The final L1 trigger decision is made by the CTP (Section 7.5.2). The CTP receives hit multiplicities from the L1Calo, L1Muon, and L1Topo systems and accepts events satisfying requirements based on object type and threshold multiplicity. Up to 512 distinct L1 trigger items may be configured in the CTP. The L1 trigger decision, as well as the 40.08 MHz LHC bunch-crossing clock, is distributed to the detector front-end and readout systems via the TTC system [141].

The parameters of the L1 trigger system depend on the pipeline memories, as specified and built for the original construction of the detector, in custom electronics located on or near the detector to store information while the trigger decision is in progress. This necessitates that the Run 3 L1 latency, the time from a given bunch crossing to the trigger decision, must adhere to the original specification of the detector. The design of the trigger and detector front-end systems requires that the L1 latency be less than $2.5 \mu\text{s}$. The maximum L1A rate supported by the detector readout systems is 100 kHz, and so the menu of the trigger selections is tuned to allocate a share of the rate budget to each underlying physics object according to ATLAS analysis goals.

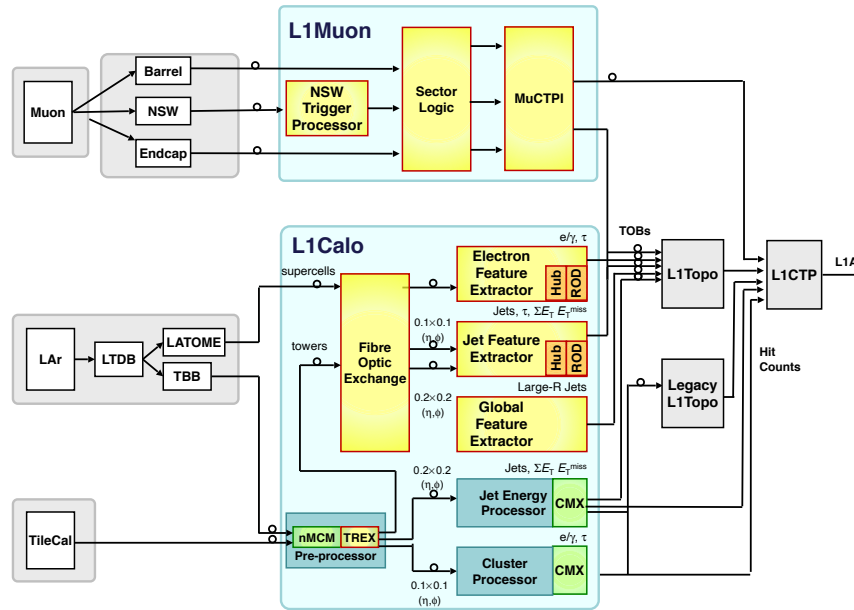


Figure 7.2: Schematic overview of the L1 trigger system in Run 3. The new and refurbished elements of the L1Muon Phase-I trigger system are shown in yellow. During the commissioning period, both the legacy (shown in green) and Phase-I L1Calo (shown in yellow) modules will run in parallel. Once the Phase-I system has been commissioned, the legacy L1Calo Jet/Energy Processor (JEP), Cluster Processor (CP) and L1Topo modules will be removed; the Phase-I L1Calo system and upgraded L1Topo module will provide triggers for physics in Run 3.

If accepted by the L1 trigger, events are then sent to the HLT. At this level, software processes reconstruct the event at higher levels of detail than L1, seeded by RoIs, which are regions of the detector in which possible trigger objects have been identified by the L1 trigger. HLT algorithms use full-granularity information from the calorimeters and the muon and tracking systems to provide better energy and momentum resolution for

threshold selections, as well as precision tracking for particle identification. The overall selection criteria for the HLT are once again provided by the trigger menu, tuned so as to maximise the data available for key analyses for the given resources.

The HLT software is designed to reproduce the offline selection as closely as possible. On average, the event processing time at the HLT in 2018 was approximately 400 ms for runs with a peak luminosity of $\mathcal{L} = 2.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. In Run 3 the HLT reduces the event rate from 100 kHz after the L1 selection to approximately 3 kHz (averaged over the course of an LHC fill); the data are then stored for offline analysis. The HLT software has undergone a significant redesign in order to take advantage of multithreaded CPU architectures; this framework, which is also used in the offline software, is called AthenaMT and is described in Section 7.6.3.

Underpinning the operation of the two trigger levels, the data acquisition system begins with detector-specific on- and off-detector electronics which perform a variety of data processing and monitoring features before passing events either to the L1 trigger or to the downstream system. In Runs 1 and 2 the off-detector stage was performed in (typically VME-based) detector-specific custom hardware modules called RODs. On receipt of a L1A, these RODs read out their data to the first common stage of the DAQ system, the ROS. The ROS buffers event data and serves it to HLT nodes on request to facilitate a trigger decision. As is described in Section 7.6.1, Run 3 will see the introduction of the new FELIX and SW ROD systems into the readout path, with the goal to reduce the amount of custom hardware in the system and perform more processing in the software domain. In the new system the ROD and ROS are replaced by FELIX, which receives readout data from L1 trigger processors and detector front-end electronics via point-to-point links and then routes them to software applications running on commodity servers (the SW RODs) that perform the same tasks previously handled in the hardware RODs, plus the buffering function of the ROS. By Run 4, all of ATLAS readout will be via this new mechanism, but in Run 3 it has been rolled out for those systems with new or upgraded front-end electronics. The interface to the HLT is identical between ROS and SW ROD, with events routed on demand to HLT processing nodes in both cases.

The final stage of the DAQ system, once HLT processing has been completed, is for accepted events to be sent to a dedicated cluster of servers, known for historical reasons as Sub-Farm Outputs (SFOs), for packing, compression, and finally transfer to offline storage.

7.1.1 Motivation for upgrades

Since Run 1, the instantaneous luminosity and number of interactions per bunch crossing (pileup) have increased substantially, as illustrated in Figure 1.1. In Run 3, the improvements in the LHC (see Section 1.2) and luminosity levelling will allow a much larger fraction of each run to be near the peak instantaneous luminosity of $\mathcal{L} = 2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, causing the average pileup to increase to $\langle \mu \rangle \approx 50$ or beyond. Increased pileup degrades the calorimeter resolution and isolation of single particles, which leads to decreased trigger efficiency and necessitates the use of higher trigger thresholds to mitigate the resulting increase in rate. One of the main goals of the ATLAS trigger upgrade is to reduce the impact of pileup in order to maintain low thresholds, especially at L1, thereby maximising the dataset recorded for precision measurements (e.g. of the Higgs boson) as well as for searches for physics beyond the Standard Model. Since the events passing the L1 selection will contain a larger number of pileup interactions, further upgrades at the HLT, including improved precision tracking and the expansion of the computing farm, are intended to maintain a reasonable output rate within bandwidth and storage limitations.

7.1.2 Improvements for Run 2

Substantial upgrades to the ATLAS TDAQ system were made prior to the beginning of Run 2 in 2015. These changes are described in detail in [10], but are briefly summarised here for completeness.

During LS1, the L1Muon barrel system was equipped with additional trigger electronics to improve the acceptance in the feet and elevator regions (where new detectors were also installed, as discussed in Section 5.3.3). In the endcap region, new trigger coincidence logic was deployed in order to improve the rejection of background originating outside the IP. This logic also made use of new signals derived from the outermost Tile calorimeter cells.

Several significant upgrades were made to the L1Calo system as shown in Figure 7.2. The PPMs [142] were upgraded to provide improved bunch-crossing identification and pedestal subtraction capabilities, resulting in better pileup suppression and lower L1 jet and E_T^{miss} trigger rates. The CP and JEP systems were upgraded to allow an increased data transmission rate. Additional improvements were made to the CP to allow for energy-dependent isolation requirements to be applied on EM object candidates. New extended Common Merger Extended Modules (CMXs) [10] were added to replace the existing Cluster Merger Modules (CMMs), which transmitted the threshold multiplicities to the CTP. In addition to this functionality, the CMX permits the transmission of TOBs, containing such information as the position and energy of physics objects passing the L1 trigger, to the new L1Topo system.

The L1 Topological trigger system (L1Topo) [143], described in Section 7.4, was installed and commissioned during Run 2. The system uses TOBs from L1Calo and L1Muon and applies topological selections in order to reduce background rates while preserving signal efficiency.

For the DAQ system, Run 2 saw the introduction of a newer version of the Input/Output (I/O) card housed in the ROS servers which provides the interface to front-end optical links and buffers data throughout the HLT selection process. The move to this new board, the RobinNP, resulted in a system up to 4 times denser than its predecessor, saving significant rack space, while also providing 10 times higher buffering capacity per data-link than before. The performance of the upgraded ROS was also nearly 4 times that of its predecessor in terms of supported maximum data request rate from the trigger.

Alongside the new ROS, the HLT itself was significantly redesigned for Run 2. What was previously two separate farms, one focused on RoI-driven selection and one studying full events, was merged into one with a more dynamic event-building approach. The combined farm eliminated the previously complex load balancing required to manage the two-level version, while also minimising any duplication of data requests to the ROS. Thanks to this redesign, it was also possible to greatly simplify the data collection network connecting the HLT to the ROS and SFO. What was previously two separate networks became one high-performance implementation, with 20 times higher bandwidth than before (40 GbE vs 2 GbE).

Finally, the SFO systems were also upgraded for Run 2. While the functionality remained the same, the new servers provided significantly improved staging and streaming capability, able to buffer more than twice as much data as before, making it possible to stably process and write much higher data volumes to offline storage.

All of these upgrades combined allowed the ATLAS trigger system in Run 2 to execute more complex algorithms with higher overall processing time and to writing out larger and richer datasets, making it possible to significantly improve the sophistication of the trigger menu and greatly enhance the ATLAS physics programme.

7.2 Level-1 Calorimeter Trigger

The Level-1 Calorimeter (L1Calo) trigger, pictured in Figure 7.2, identifies events containing electrons, photons, jets, and τ -leptons decaying into hadrons, as well as events with large missing transverse energy (E_T^{miss}) or large total transverse energy ($\sum E_T$) using custom hardware processors based on FPGA technology. To manage increasing rates in this regime, while preserving an effective and efficient calorimeter trigger, significant upgrades are necessary. These upgrades include more refined processing of electromagnetic calorimeter information at higher granularity, which provides improved identification of isolated electrons and photons and improved rejection of jets, thereby maintaining low trigger thresholds, which are especially important for electroweak physics. Improved processing at higher granularity and new algorithms will also enable better identification of jets and discrimination against pileup, which will benefit searches for physics involving boosted objects and E_T^{miss} .

The Run 3 system includes new electromagnetic and jet feature extractors: the electron Feature EXtractor (eFEX) (Section 7.2.1) identifies electron, photon, and τ objects, while the jet Feature EXtractor (jFEX) (Section 7.2.2) identifies jets, E_T^{miss} , and τ s. A global feature extractor (global Feature EXtractor (gFEX), Section 7.2.3) identifies large-area jets, E_T^{miss} , and $\sum E_T$. The new FEX modules are shown in Figure 7.3.




Figure 7.3: (a) A production eFEX module. (b) A production jFEX module. (c) A production gFEX module.

Compared to Runs 1 and 2, L1Calo receives finer-granularity input data from the LAr calorimeter in Run 3. In Runs 1 and 2, these inputs consisted of “trigger towers” spanning 0.1×0.1 in η and ϕ . In the Run 3 system, the LAr Digital Processing System (LDPS; Section 4.1.2) provides electromagnetic calorimeter information in the form of Super Cells containing sums of four or eight calorimeter cells. The detailed mapping is provided in Table 4.1 and summarised here. Each trigger tower contains 10 Super Cells, as shown in Figure 7.4: one in the presampler, four in each of the first and second layers, where the majority of shower energy is deposited, and one in the third layer. The full granularity, corresponding to $\Delta\eta \times \Delta\phi = 0.025 \times 0.1$ in the first and second layers, is available to the e/γ and τ triggers, while the jet, E_T , and E_T^{miss} triggers use coarser granularity, comparable to the size of the 0.1×0.1 trigger towers. For comparison, in Run 2, L1 jet triggers were based on 0.2×0.2 jet elements.

In Runs 1 and 2, the calorimeter data arrived at L1Calo in analogue format, and was then digitised and calibrated by the pre-processors. In Run 3, the digitisation and calibration of the LAr calorimeter data




Figure 7.4: The finer granularity of the Run 3 L1Calo trigger towers after the upgrade of the liquid argon calorimeter electronics. Each tower is divided into ten Super Cells, each providing an E_T value.

are now performed in the LAr calorimeter electronics (4.1.2). The Tile calorimeter data are still received in analogue format and are digitised by the pre-processors, and then the new Tile REciever eXchange (TRES) (Tile Rear EXtension; Section 7.2.4) modules transmit them to the FEXs and the legacy system. It is planned to upgrade the Tile calorimeter system to perform the signal digitisation for running at the HL-LHC [144].

The digitised and calibrated calorimeter inputs are sent to L1Calo via optical fibres. Mapping and routing of the fibres to the FEXs is performed by the Fibre-Optic eXchange (FOX) optical plant (Section 7.2.4).

The output from L1Calo consists of TOBs which include the position, E_T , object type, and energy sum. These are sent to L1Topo, which computes the multiplicity of TOBs passing a given threshold and sends the information to the CTP. Another optical plant (TopoFOX; Section 7.2.4) is used to map the output fibres from L1Calo to L1Topo. In addition, L1Calo sends RoI information to the HLT and readout data to the FELIX and SW ROD (Section 7.6.1).

During the start-up and commissioning phase of Run 3, triggers for physics will be provided by the legacy (Runs 1 and 2) L1Calo system. Once the Phase-1 system has been commissioned and validated, it will be enabled to provide physics triggers and the legacy system will be removed.

7.2.1 Electron Feature Extractor (eFEX)

7.2.1.1 eFEX module design The eFEX system consists of 24 modules located within two ATCA shelves, with each module containing four algorithm-processing FPGAs and one control FPGA. A block diagram of an eFEX module is shown in Figure 7.5, and a photograph in Figure 7.3(a).

The system covers a region of $|\eta| \leq 2.5$ and the full ϕ range. Each of the 24 modules receives data on up to 136 fibre links (11.2 Gb/s), covering a calorimeter area of up to $\Delta\eta \times \Delta\phi = 1.7 \times 1.0$. Each FPGA processes data for 32 algorithm cores of area $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$, for a total area of $\Delta\eta \times \Delta\phi = 0.4 \times 0.8$ per FPGA. On the borders, the total number of algorithm cores processed per FPGA increases to 40. The input to each algorithm core includes the surrounding environment, which measures $\Delta\eta \times \Delta\phi = 0.3 \times 0.3$. Each FPGA contains an algorithm module which produces e/γ and τ TOBs. The internal clock for processing is 200 MHz and the output rate is 280 MHz. TOBs are first sorted locally on each FPGA before being sent to dedicated e/γ and τ sorter modules on two of the four FPGAs; this sorting is also performed at 280 MHz.

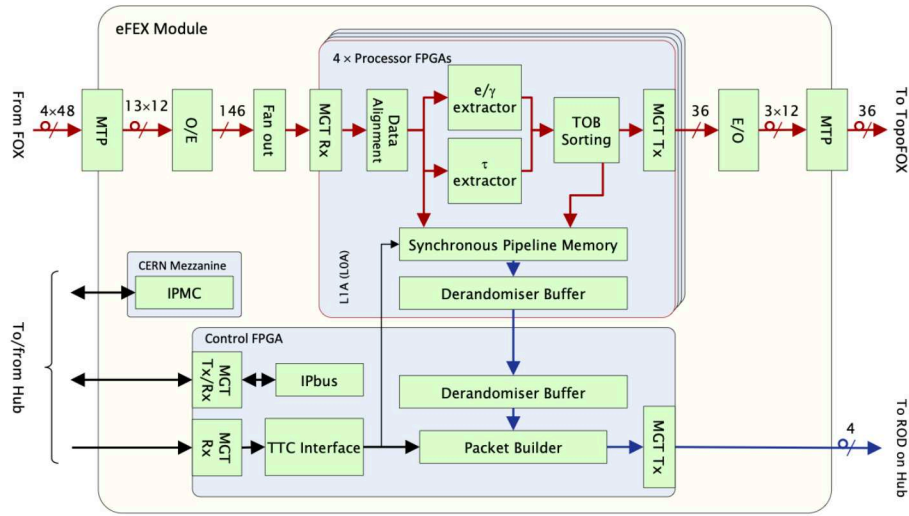


Figure 7.5: Block diagram of an eFEX module, illustrating the real-time and readout paths. Control and monitoring signals, except for the L1A, are not shown.

The TOBs computed by the eFEX algorithms are merged and sorted across the entire system. A maximum of six TOBs per algorithm (e/γ and τ) per module may be sent to L1Topo; in case more than six TOBs are identified, the six with the highest transverse energy TOBs are sent.

Timing and readout TTC signals are received on the backplane from the Hub (see Section 7.2.4.4). Upon receiving a L1A, TOBs and input data (prescaled to limit the bandwidth usage) are sent to the ROD (described in Section 7.2.4.4). Each eFEX module sends its readout data over 8 electrical links (6.4 Gb/s) on the ATCA backplane using the multi-lane Aurora protocol [145]. From the ROD, the readout data are then sent to the FELIX and SW ROD (Section 7.6.1).

Configuration, control, and monitoring Configuration is performed by the control FPGA. Monitoring and control are handled by the IPbus [146] firmware implemented in the control FPGA. A CERN-standard IPMC module [147] performs low-level control functions.

7.2.1.2 eFEX algorithms

e/γ algorithm The eFEX e/γ algorithm searches for a “seed”, i.e. a local energy maximum, by comparing the four Super Cells in a given algorithm core in the second layer of the electromagnetic calorimeter. This process is illustrated in Figure 7.6(a). The ϕ -direction of the cluster is established by finding the highest-energy Super Cell adjacent to the seed. A cluster of 3×2 , Super Cells (see Figure 7.6(b)), is formed around the seed, and the cluster energy is then computed by adding the energy of the corresponding Super Cells in the presampler, first, and third layers of the calorimeter.

To improve background rejection, isolation requirements may be applied. These are generally defined as a ratio between the energy sum in the cluster and that in the surrounding area. Such requirements are applied in the form of selections based on several electromagnetic shower observables; the following observables can be computed in the eFEX:

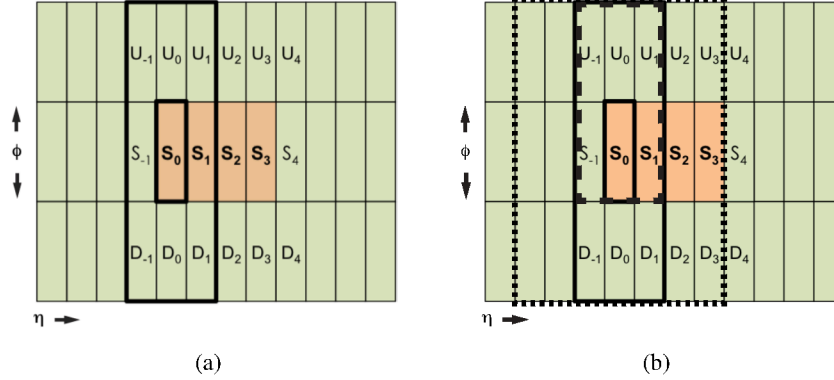


Figure 7.6: (a) The L1Calo eFEX seed finder algorithm. The four potential seeds (S) in the algorithm core are compared in order to determine the local energy maximum. The directionality of the cluster, upward (U) or downward (D) is determined by finding the largest energy deposit among the Super Cells surrounding the seed. (b) The dashed rectangle shows of 3×2 Super Cells used to compute the cluster energy, assuming that the local energy maximum was at cell S_0 and that the most energetic directly neighbouring Super Cells to S_0 was U_0 . The dotted rectangle shows the 7×3 Super Cells area used for the calculation of the shower width in the second calorimeter layer.

- Shower width in the second calorimeter layer

$$R_\eta = 1 - \frac{E_{T,3 \times 2}}{E_{T,7 \times 3}}, \quad (3)$$

where $E_{T,3 \times 2}$ is the transverse energy sum in an area of 3×2 Super Cells in the second layer of the calorimeter centred on the local maximum and its highest- E_T neighbour in ϕ , and $E_{T,7 \times 3}$ is the transverse energy sum in an area of 7×3 Super Cells in the second layer of the calorimeter centred on the same locus.

- Hadronic fraction

$$R_{\text{had}} = \frac{E_{T,\text{had}}}{E_{T,\text{EM}} + E_{T,\text{had}}} \quad (4)$$

where $E_{T,\text{had}}$ is the transverse energy sum in a window of 0.3×0.3 in the hadronic calorimeter and $E_{T,\text{EM}}$ is the transverse energy sum in the four layers of the electromagnetic calorimeter, computed in a window of 3×2 (1×3) Super Cells in the presampler and third layers, centred as in the case of R_η .

- Shower width in the first calorimeter layer

$$w_{s,\text{tot}} = \sqrt{\frac{\sum E_{T,i} \times (i - i_{\text{max}})^2}{\sum E_{T,i}}}, \quad (5)$$

where i runs over five Super Cells in η , centred on the most energetic Super Cell in the first layer of the calorimeter.

τ algorithm The eFEX τ algorithm has been designed to identify prompt hadronic decays of τ -leptons. First, the cluster seed tower ($\Delta\eta \times \Delta\phi = 0.1 \times 0.1$) is chosen by identifying a local maximum in an environment of 3×3 towers. From this seed tower, the seed SuperCell is then chosen by identifying

its constituent SuperCell with the highest energy; this search is performed only in the second layer of the EM calorimeter. A cluster is then formed in the direction (up or down in ϕ) of the highest-energy adjacent SuperCell. In the fine-granularity layers of the EM calorimeter (the first and second layers), the cluster spans a region of 5×2 Super Cells ($\Delta\eta \times \Delta\phi = 0.125 \times 0.2$). In the coarse-granularity layers (the presampler and the third layer of the EM calorimeter, plus the hadronic calorimeter), the cluster comprises 3×2 Super Cells ($\Delta\eta \times \Delta\phi = 0.3 \times 0.2$). The energies in each of these regions are summed to compute the total cluster energy. The τ clustering algorithm is illustrated in Figure 7.7.




Figure 7.7: (a) The L1Calo eFEX tau cluster, as defined in the first and second EM calorimeter layers. The seed Super Cell is shown in red, while the other Super Cells making up the cluster are shown in orange. The surrounding environment, which is not included in the cluster energy computation, is shown in blue. (b) The L1Calo eFEX tau cluster, as defined in the presampler and third EM calorimeter layers, as well as in the hadronic layer.

The eFEX τ algorithm can also apply an isolation criterion based on the shower shape in the EM calorimeter. The discriminating variable F_{core} is used; this variable is defined as the ratio of the transverse energy deposited in a region of 3×2 Super Cells surrounding the seed ($E_{T,3 \times 2}$) to that deposited in a region of 9×2 Super Cells surrounding the seed ($E_{T,9 \times 2}$):

$$F_{\text{core}} = \frac{E_{T,3 \times 2}}{E_{T,9 \times 2}}. \quad (6)$$

7.2.2 Jet Feature Extractor (jFEX)

7.2.2.1 jFEX module design The jFEX system identifies jets, hadronically-decaying τ -leptons, $E_{\text{T}}^{\text{miss}}$, and $\sum E_{\text{T}}$ in the range $|\eta| \leq 4.9$. The system consists of one ATCA shelf, equipped with six jFEX modules. Digitised electromagnetic and hadronic calorimeter data arrive at the jFEX via the FOX (Section 7.2.4) and are received by optoelectronic devices (MiniPODs). Each module contains 24 MiniPODs (20 receivers and four transmitters) and four Xilinx® Ultrascale+™ processor FPGAs. Each processor FPGA contains 120 Multi-Gigabit Transceivers (MGTs) and can handle an input bandwidth of up to 3.6 Tb/s. TOBs identified by the jFEX are sent to L1Topo. A block diagram of a jFEX module is shown in Figure 7.8 and a photograph in Figure 7.3(b).

Timing and readout Timing and TTC signals are received on the backplane from the Hub (Section 7.2.4.4). For each L1A, a single readout data packet is constructed per event from the processor FPGA data and is sent to the Hub via high-speed links on the backplane. Upon receiving a L1A, TOBs and input data (prescaled to limit the bandwidth usage) are sent to the ROD (described in Section 7.2.4.4), which then sends the data to FELIX and the SW ROD (Section 7.6.1).

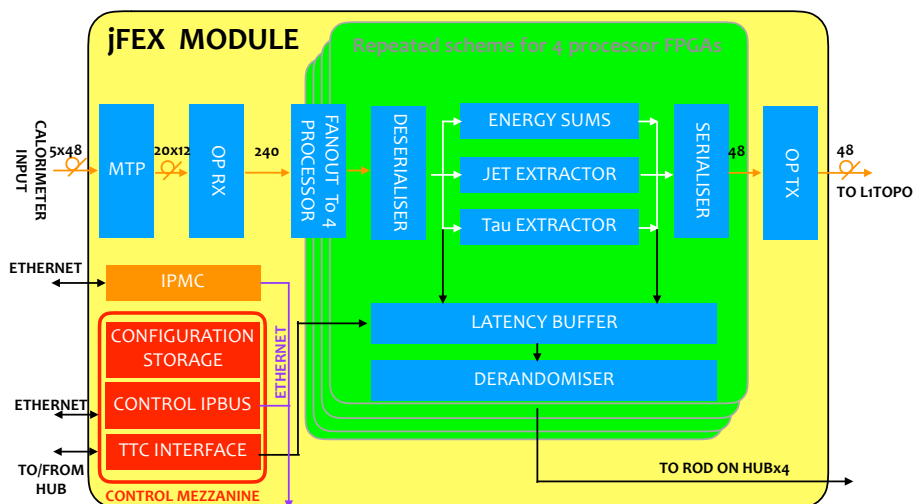


Figure 7.8: Block diagram of a jFEX module, illustrating the real-time and readout paths. Control and monitoring signals are only shown for the L1A.

Configuration, control, and monitoring The processor FPGAs are configured by the control FPGA. Control and monitoring of the FPGAs is done via IPbus [146], which executes such functions as setting parameters for the algorithms, controlling high-speed links, access of playback and spy memories, and environmental monitoring. Low-level control, compliant with ATCA standards, is performed by an IPMC module [147].

7.2.2.2 jFEX algorithms

Each jFEX module provides full $\phi = 2\pi$ coverage. Four jFEX modules cover the barrel region, spanning $|\eta| \leq 1.6$, and two modules cover the endcap and forward regions corresponding to $1.6 < |\eta| < 4.9$. Each of the four processor FPGAs in a given module covers a core of $\Delta\eta \times \Delta\phi = 0.8 \times 1.6$ in the barrel region or 3.3×1.6 in the endcap and forward regions. In order to properly identify TOBs on the edges of a single FPGA, each FPGA is supplied with a copy of the data from the surrounding environment, corresponding to a so-called “overlap” area of 2.4×3.2 .

Calorimeter data inputs to the jFEX consist of 0.1×0.1 trigger towers in the region $|\eta| < 2.5$. The granularity is slightly coarser in the endcap and forward regions: 0.2×0.2 in the region $2.5 < |\eta| < 3.1$ and 0.1×0.2 in the region $3.1 < |\eta| < 3.2$. In the forward region, $3.1 < |\eta| < 4.9$, the inputs consist of Super Cells; in this region, the granularity is irregular, corresponding to the positions of the calorimeter electrodes; this is illustrated in Figure 7.9.

Pileup subtraction and noise cuts. The pileup density ρ is computed and subtracted from each trigger tower E_T . Noise cuts may also be applied to individual trigger towers. Separate noise cuts can be applied for the computation of E_T^{miss} versus all other objects.

jFEX small-radius jet algorithm. The jFEX small-radius jet algorithm is illustrated in Figure 7.10. Small-radius jets are identified using a sliding-window algorithm in a search window of 5×5 towers, spanning $\eta \times \phi = 0.5 \times 0.5$ (equivalent to $\Delta R < 0.3$). “Seeds” of 3×3 towers (equivalent to $\Delta R < 0.2$) are constructed around each tower in the search window. To identify local maxima without double counting and without missing jet objects, the energy sums of the seeds in the search window are compared using




Figure 7.9: Granularity of trigger tower ($|\eta| < 3.1$) and Super Cell ($3.1 < |\eta| < 4.9$) inputs to the jFEX. The different colours serve to highlight the different granularities, as described in the text. The dashed lines illustrate the boundaries between the core region of a particular FPGA ($1.6 < |\eta| < 4.9$, $0 < \phi < \pi/2$) and the surrounding environment.

either \geq or $>$ operators. The central tower of the seed with the largest energy sum in a given search window is then chosen as the centre of the jet. The energy sum of the seed is added to the sum of the tower energies in a ring of radius $0.2 \leq R < 0.4$ to form a “round” jet object consisting of 45 towers.




Figure 7.10: (a) The seeding process by which local maxima are identified in the jFEX. (b) Comparative operators used to identify local maxima in a given search window. (c) A small-radius jet as defined by the jFEX.

jFEX large-radius jet algorithm. The large-radius jet algorithm also uses a sliding-window algorithm. In order to reduce resource usage, it uses some elements of the small-radius jet algorithm. The energy sum of a small-radius jet is added to the energy sum of a ring of radius $0.4 \leq R < 0.8$ to form a large-radius jet of $R < 0.8$ consisting of 148 trigger towers.

jFEX E_T^{miss} and $\sum E_T$ algorithms The total transverse energy sum ($\sum E_T$) is computed by adding the EM and hadronic trigger tower E_T values in slices of constant η . Each $\sum E_T$ TOB contains two sums corresponding to different regions in $|\eta|$. The final energy summation over the full detector area is done by L1Topo.

E_T^{miss} is computed in a manner analogous to $\sum E_T$: transverse energy values are summed in slices of

constant ϕ and then weighted by $\cos(\phi)$ and $\sin(\phi)$ in order to determine the x - and y -components, respectively.

jFEX τ algorithm jFEX identifies hadronically-decaying τ -leptons in the range $|\eta| \leq 2.5$. Seeds are identified as for jets, using a search window of size $R < 0.2$. In order to minimise resource usage, the algorithm reuses jet seeds ($R < 0.2$) as τ cluster energies. The isolation is computed using the energy ring of $0.2 \leq R < 0.4$ used by the small-radius jet algorithm.

jFEX forward electron algorithm jFEX identifies electrons in the region $2.3 \leq |\eta| < 4.9$. This region extends beyond the acceptance of eFEX whose coverage is limited to that of the tracking acceptance. The seeding procedure, which uses only EM trigger towers as inputs, is analogous to that used for jets and τ -leptons. The search window varies in size depending on the calorimeter granularity: $R < 0.2$ in the region $2.3 \leq |\eta| < 2.5$ and $R < 0.3$ in the region $2.5 \leq |\eta| < 4.9$. The cluster energy is defined as the energy sum of the electron seed (a single EM trigger tower) and the most energetic neighbouring tower. The EM isolation is defined as the transverse energy sum of all EM towers within $R < 0.4$, excluding the electron cluster itself.

7.2.3 Global Feature Extractor (gFEX)

7.2.3.1 gFEX module design

The gFEX has been designed such that the data from the entire calorimeter can be processed on a single module, thus permitting the use of full-scan algorithms. This functionality is intended to facilitate identification of boosted objects and global observables such as E_T^{miss} , which are of particular interest in a number of searches for new physics. Jet substructure algorithms may be employed to discriminate between signals from boosted boson or top quark decay products and the Quantum Chromodynamics (QCD) multi-jet background. The gFEX also incorporates pileup subtraction capabilities, which provide robustness against pileup for E_T^{miss} triggers. E_T^{miss} trigger rates are particularly sensitive to pileup and are of great importance for searches for SUSY, dark matter, and Higgs bosons decaying to invisible final states.

The gFEX module, illustrated as a block diagram in Figure 7.11 and with a photograph in Figure 7.3(c), includes three processor FPGAs and a Zynq+ System-on-Chip (SoC) which performs control and monitoring functions. Each processor FPGA covers the entire ϕ ring; two cover the barrel region $|\eta| < 2.5$ and the remaining FPGA covers the endcap and forward regions. Data arrive at the gFEX from the LDPS (LAr – Section 4.1.2) and TREX (Tile) via the FOX (Section 7.2.4) at a transmission speed of 11.2 Gb/s.

Timing and readout TTC information is transmitted to and from the Zynq+ FPGA via FELIX (described in Section 7.6.1). Upon receipt of a L1A, readout data and TOBs are sent to FELIX. From the FELIX, readout data are sent to the SW ROD, which builds event fragments to be sent to the HLT via the DAQ system.

Configuration, control, and monitoring The IPbus protocol [146] is used to control access to playback and spy memories on the board for diagnostic and commissioning purposes. Monitoring and control of the module are performed via IPMC [147], as well as by DCS (see Section 7.7).

7.2.3.2 gFEX algorithms

Inputs to the gFEX are combined into “gTowers” by summing transverse energy from the electromagnetic and hadronic calorimeters and subsequently applying a calibration with a lookup table. The granularity in the barrel region is 0.2×0.2 , increasing in the endcaps and forward region as illustrated in Figure 7.12.

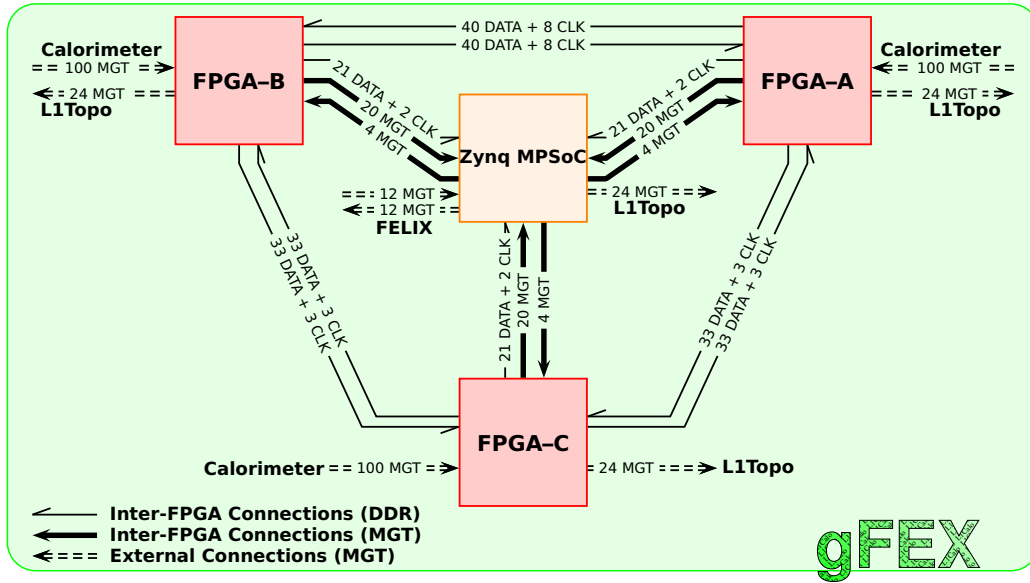


Figure 7.11: Block diagram of the L1Calo global Feature EXtractor (gFEX).

Contiguous groups of 3×3 gTowers (2×3 gTowers on the boundaries between FPGAs and FPGA divisions) are called “gBlocks”. gBlocks have configurable thresholds and can be used as seeds for large-radius jet identification as described below.

Two types of TOB are identified by the gFEX. Jet TOBs are produced on a single processor FPGA and include large-radius jets, gBlocks, the local pileup energy density (ρ) and optionally substructure information. Global TOBs are produced on the Zynq FPGA using information from all of the processor FPGAs and can include such global observables as E_T^{miss} , $\sum E_T$, and H_T . All TOBs computed by the gFEX are sent to L1Topo.

gFEX pileup suppression algorithm. Pileup suppression is performed on an event-by-event basis in the firmware. The energy density ρ is computed per processor FPGA on an event-by-event basis, and then multiplied by a factor of 69 to take into account the 69 gTowers that make up a large-radius jet as defined below. This value is then subtracted from the E_T of the large-radius jet.

gFEX large-radius jet algorithm. Large-radius jets (“gJets”) are formed using a cone algorithm by summing the gTower E_T in a “circular” 1.8×1.8 area ($R < 0.9$) around each gTower. A total of 69 gTowers are used for each large-radius jet as shown in Figure 7.12. Since gJets may span the area covered by two different processor FPGAs, information is shared between processor FPGAs in the form of partial energy sums. gJets are permitted to overlap. Large-radius jets and gBlocks TOBs are separately sorted on both the right and left η divisions of the A and B processor FPGAs and sent to L1Topo on separate fibres. Initially, only the most energetic gJet in each division will be transmitted. If latency allows, a sub-leading non-overlapping gJet will be transmitted as well. A non-zero seed threshold can be applied to the central gBlock of the large-radius gJets before sorting. For the gBlocks, the sub-leading gBlock is required to have its centre outside of the area of the leading gBlock.

gFEX “jets-without-jets” E_T^{miss} algorithm. The “jets without jets” (JwoJ) algorithm [148] is based on the concept that jet observables may be transformed into global event-shape observables. The JwoJ algorithm is motivated by the fact that diffuse transverse energy is likely to be associated with pileup interactions




Figure 7.12: Schematic of the gFEX processor FPGA coverage, illustrating the granularity of the gTower inputs. 3×3 (2×3 on FPGA region boundaries) groups of gTowers make up gBlocks. Note that both gBlocks and large-radius jets are permitted to overlap. All possible large-radius ($R < 0.9$) jets are constructed and those passing an optional seed requirement are sorted by energy and transmitted separately to L1Topo in two η regions per FPGA.

whereas clustered transverse energy is more likely to be associated with the interaction of interest. The gFEX computes E_T^{miss} by separating the E_T sums into “hard” (MHT) and “soft” (MST) terms, where the hard term consists of the E_T sum of towers with the associated gBlock satisfying $E_T > 25$ GeV and the soft term consists of the E_T sum of the remaining towers. The E_T^{miss} is then computed as a linear combination of the hard and soft terms:

$$E_{T,x,y}^{\text{miss}} = a_{x,y} \text{MHT}_{x,y} + b_{x,y} \text{MST}_{x,y}, \quad (7)$$

where the a and b coefficients can be optimised for resolution and overall performance.

7.2.4 L1Calo infrastructure

7.2.4.1 TREX In Run 3, the output signals from the Tile calorimeter will still be in analogue format (Section 4.2) but inputs to the FEXs must be in digital format. The digitisation is performed in the PPMs of the legacy L1Calo system [142], and new TREX modules, located on the PPMs, process and duplicate the digitised Tile signals and send them to the Phase-1 L1Calo system via optical links at 11.2 Gb/s. For backwards compatibility with the legacy system, the Tile signals are sent electrically to the legacy CP and JEP. The readout data are sent to both the FELIX and the legacy ROD via optical links (9.6 Gb/s and 960 Mb/s, respectively). TTC info is received from the PPMs via a VME-P2 connector and from the FELIX via an optical link at 4.8 Gb/s.

The TREX module, pictured in Figure 7.13, consists of a pre-processor data collector (PREDATOR) Xilinx® Ultrascale™ FPGA, and four LVDS data in-out (DINO) FPGAs (Xilinx® Artix® -7). It includes

six Samtec FireFly™ optical transceivers: four 12-channel transmitters to the FEXs and two four-channel Duplex transceivers to handle TTC and readout signals. Monitoring of environmental conditions is performed using a Zynq® Ultrascale+™ MPSoC (multi-processor system-on-chip) device, and the data are sent to DCS via Gigabit Ethernet interface. Configuration, control, and monitoring of the module are performed via the PPMs using interfaces to VME.

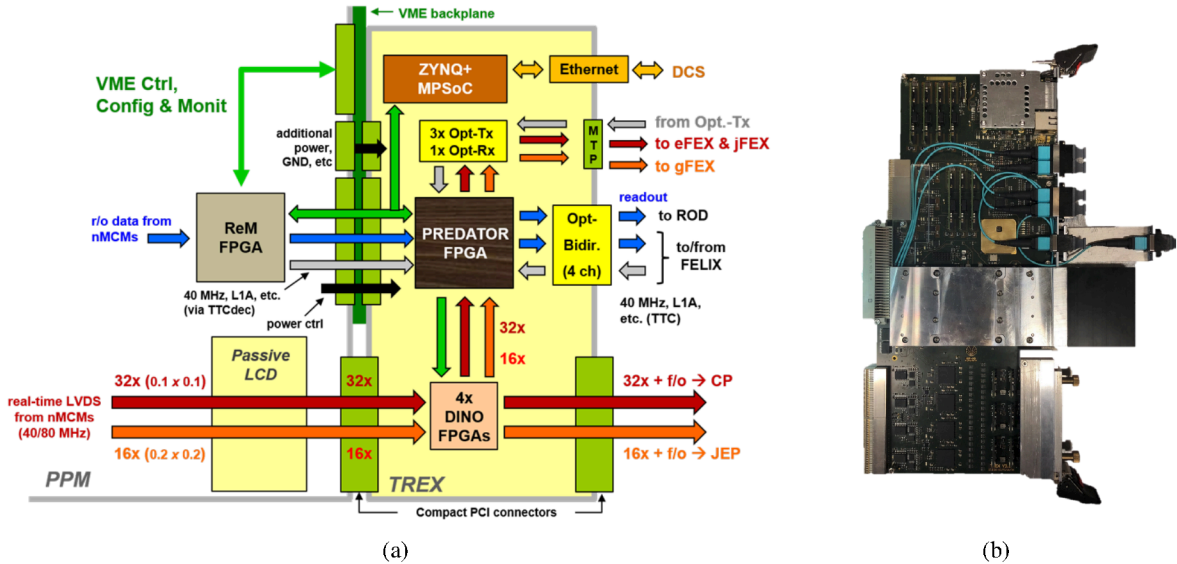


Figure 7.13: (a) Schematic diagram of the TREX. (b) A photo of a TREX module.

7.2.4.2 FOX The digitised inputs to L1Calo from the LAr calorimeter and the TREX arrive via the FOX optical plant. The FOX separates the fibre bundles (approximately 7500 fibres in total) and re-routes them according to the necessary mapping required by the L1Calo FEXes. It also routes the input signals as required for data sharing between modules covering contiguous areas. Duplicate output connectors are provided in order to allow for spares for the gFEX board.

The FOX is internally divided into LArFOX and TileFOX components, which route the fibres from the LAr calorimeter and TREX, respectively. The LArFOX consists of four “FOX boxes”, which route the EM and hadronic inputs to the eFEX, while the TileFOX consists of two FOX boxes, which route the EM and hadronic inputs to the jFEX and gFEX. Fibres may also be routed from one FOX box to another before finally being sent to a given FEX. A FOX box consists of a 19-inch rack-mount “2U” chassis containing internal shuffle modules specified by their input and output Multi-fiber Termination Push-on (MTP) connectors and the input-to-output fibre mapping.

7.2.4.3 TopoFOX Inputs to L1Topo from the L1Calo FEXes and MUCTPI are handled by the TopoFOX, which performs fibre mapping and routing to the three L1Topo processor modules in a manner analogous to the FOX. Each of the four processor FPGAs in a given FEX module sends a 12-fibre ribbon to the 48-fibre output connector of that module, while each of the side-A and side-C FPGAs of the MUCTPI send two 12-fibre ribbons to the 48-fibre muon output bundle. The TopoFOX receives and sorts these fibre outputs, combining them into twelve 72-fibre bundles, two of which are sent to each of the six FPGAs making up the three L1Topo modules.

7.2.4.4 Hub & ROD Common communications functionalities for the eFEX, jFEX, and L1Topo ATCA shelves are provided by the Hub modules. The Hub supports the system readout via a ROD daughter card contained on a mezzanine on the module, provides switching functionality for module control and DCS, and distributes timing and control signals to the modules. Each ATCA shelf contains two Hub modules. In total, the L1Calo system contains seven Hub modules: four for the eFEX system, two for the jFEX system, and one for L1Topo.

In a given shelf, the Hub module located in logical slot 1 provides switching capability for module control signals; it also receives the LHC clock and TTC information from the FELIX via a 12-channel MiniPOD optical receiver. This information is then fanned out to the ROD daughter card contained on the Hub, the FEX modules and the second Hub in the shelf. This second Hub module, located in logical slot 2, provides switching for the DCS network. Readout control data from its ROD is sent to the Hub in slot 1 to be included in the combined readout data stream. At the HL-LHC, FEX data will be multiplexed between the two RODs in the shelf in order to cope with the higher data rate.

High-level control of the Hub is performed via an IPbus interface [146]. The Hub also connects to the IPMB [149] via an IPMC card [147] located on the module. An I²C bus is used to manage communications on the Hub itself.

Readout data from the eFEX, jFEX, and L1Topo modules are sent to the new FELIX and SW ROD readout system (described in Section 7.6.1) by the ROD daughter card mounted on the Hub module. The gFEX sends its data directly to FELIX. The ROD receives the readout data from all FEX modules in a given shelf. The data are then decoded, the checksum evaluated, and finally merged into a single packet and buffered before being transmitted to FELIX.

The ROD includes a Xilinx® Virtex® -7 FPGA and four MiniPOD optical transceivers, which transmit the readout data to FELIX. A photo of a Hub module, with the ROD mounted, is shown in 7.14.




Figure 7.14: A Hub module, with the ROD mezzanine mounted (at the top of the module, outlined in pink).

7.3 L1 Muon Trigger

In Run 1, the Level-1 Muon (L1Muon) trigger decision in the endcap region ($1.05 < |\eta| < 2.4$) was based on the coincidence of hits in the TGC stations of the endcap middle layer (EM-TGC), called the Big Wheels. The EM-TGC have three stations (TGC-M1, TGC-M2, and TGC-M3) per side, with the M1 station consisting of three layers and the outer two stations (M2 and M3) each consisting of two layers, for a total of seven layers.

To improve the rejection of fake muons in the full η range of the EM-TGC, several upgrades have been performed, with the aim of reducing the L1 trigger rate while keeping the efficiency high [150]. During Run 2, an additional requirement of a coincidence between the EM-TGC and the D-layer cells of the Tile calorimeter (see Figure 7.15) in the range $1.05 < |\eta| < 1.3$, was introduced.

Further rate reduction (see Figure 1.8(b) for the improvement in L1 Muon endcap trigger rate) in the range $1.3 < |\eta| < 2.4$ is achieved in Run 3 by replacing the existing muon endcap inner stations (the Small Wheels) by the NSW described in Section 5.2, comprising sTGC and Micromegas detectors with high-rate tolerance and improved resolution. The upgrades to the L1Muon endcap trigger are illustrated in Figure 7.15. At $\mathcal{L} = 2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, the total L1 trigger rate for single muons with $p_T > 20 \text{ GeV}$ was expected to be 18 kHz before these improvements; with these improvements, the expected rate is 13 kHz. The additional fake muon rejection provided by the upgrade will be even more important during Run 4 and beyond. The expected performance of the combined upgrades to the L1Muon system is also shown in Figure 5.2.

In the L1Muon trigger, muon candidates are identified by measuring the degree to which their paths deviate (in both R and ϕ) from the pattern of hits expected from a muon with infinite momentum; this deviation is inversely proportional to the p_T of the muon. Using the deviations expected for different muon p_T , so-called “muon roads” can be defined for different p_T thresholds. The Run 2 trigger electronics permitted the use of six programmable p_T thresholds; in Run 3, the upgraded sector logic (see Section 7.3.4) will allow 15 programmable thresholds in the endcap.

7.3.1 TGC EI-FI Coincidence

The main background source in the L1Muon endcap trigger is low-momentum charged particles produced in the endcap toroid magnets and beam shielding. In order to reject backgrounds due to these particles, a coincidence requirement between the Big Wheel and the TGC-FI chambers of the old EI wheels was introduced in 2015, in part as a proof of principle for the NSW, which ultimately replaced the EI wheels. An additional coincidence requirement between the Big Wheel and the EIL4 TGC chambers was introduced in 2016, and remains in effect for Run 3 and beyond. Run 2 papers refer to the combination of these two vetoes as the “TGC EI-FI coincidence”; however, for Run 3 only the “TGC EIL4 coincidence” remains.

7.3.2 Tile-Muon Coincidence

In the region $1.05 < |\eta| < 1.3$, where the inner layer of the muon system provides incomplete coverage due to the presence of the toroid magnets, a coincidence between the EM-TGC chambers and the Tile calorimeter D-cell layers is required in order to improve the rejection of fake muons, which consist primarily of low- p_T protons. The distribution of this background can be seen in Figure 7.16(a).

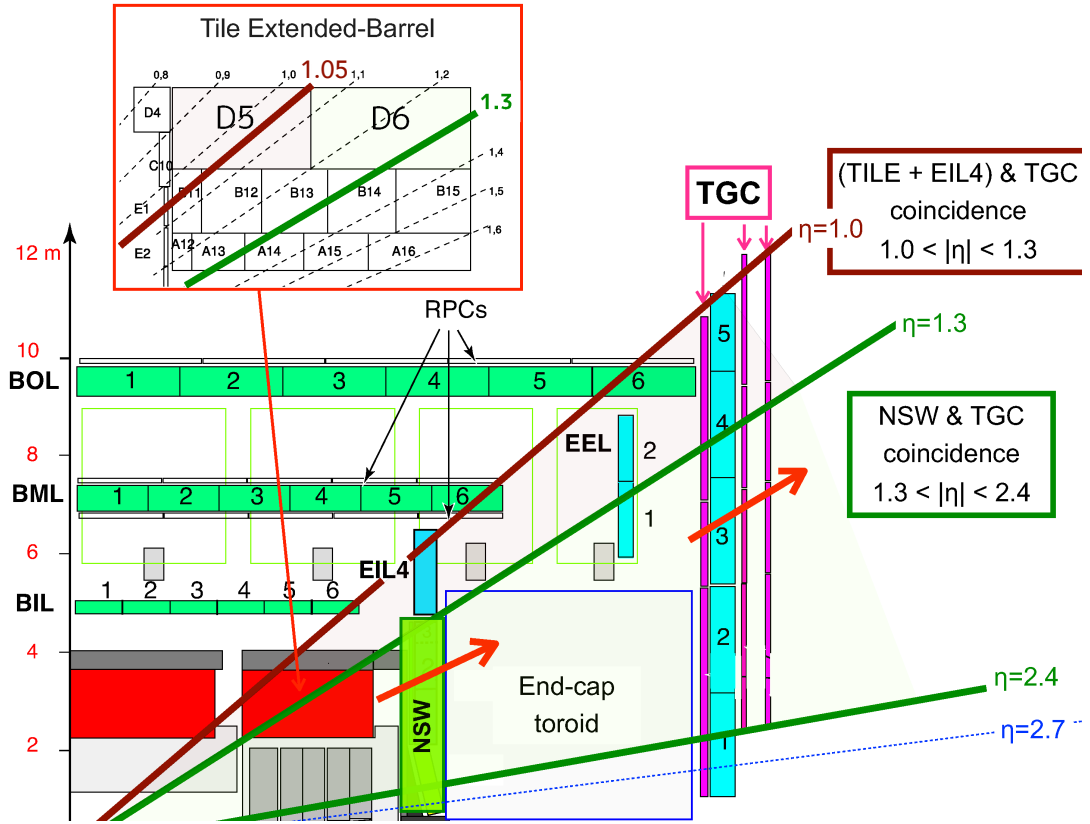


Figure 7.15: Schematic overview of the upgraded L1Muon endcap system, illustrating the NSW and Tile-Muon coincidence.

As shown in Figure 7.15, a high- p_T muon originating from the IP and passing through the endcap in the region $1.0 < |\eta| < 1.3$ can be expected to traverse the D5 or D6 cells of the Tile calorimeter extended barrel. These modules have a granularity of $\eta \times \phi = 0.2 \times 0.1$, thus providing finely segmented energy measurements. Muon candidates with $p_T > 20$ GeV are required to coincide with an energy deposit satisfying a pre-determined threshold in at least one of the corresponding Tile modules mapped in ϕ to the associated muon trigger sector.

New TMDBs were installed in Run 2 to perform matching in ϕ between the Tile extended barrel modules and the L1Muon endcap sector logic. Each extended barrel consists of 64 modules in ϕ (128 in total), while the L1Muon endcap region consists of 48 trigger sectors. Each TMDB receives the D5 and D6 inputs from eight Tile modules and three L1Muon endcap sector logic boards. Hence, 16 TMDBs in total are required. The TMDBs are 9U VME modules situated in two crates in USA15. They perform the following tasks:

- Receive and digitise the analogue signals from the Tile D5 and D6 cells.
- Provide calibration for the Tile signals.
- Provide signal detection for each Tile cell.
- Provide BCID information using timing information from the TTC receiver [141].

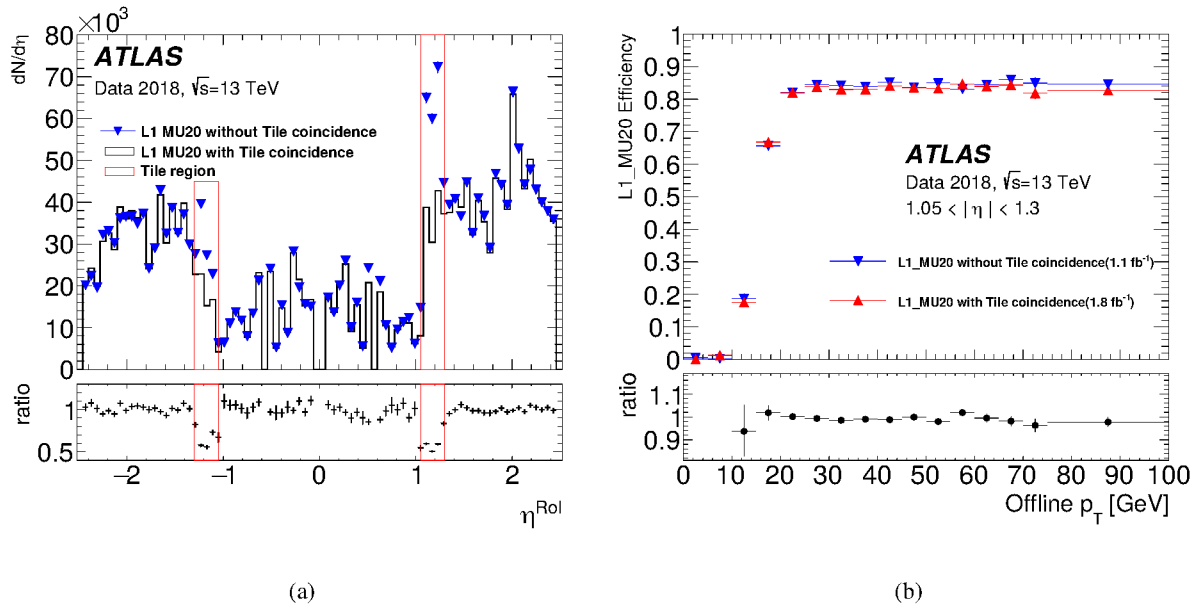


Figure 7.16: (a) The pseudorapidity distribution of Run 2 L1Muon RoIs with $p_T > 20$ GeV before and after the deployment of the Tile calorimeter coincidence requirement. The reduction in the range $1.05 < |\eta_{\text{RoI}}| < 1.3$, where the coincidence requirement is applied, is highlighted. (b) The efficiency of the Run 2 L1 trigger selection for muon RoIs with $p_T > 20$ GeV in this pseudorapidity range.

- Provide η , ϕ , and BCID information from cells in which a signal has been detected to the corresponding three muon sector logic boards.
- Share information with neighbouring receiver boards to accommodate for non-perfect matching in ϕ between the eight Tile modules and three muon sector logic boards.
- Provide readout data fragments to the DAQ system.

For single muons with $p_T > 20$ GeV, the trigger rate was reduced by 6% after this coincidence requirement was introduced, with only a 2.5% efficiency loss in the range $1.05 < |\eta| < 1.3$, as shown in Figure 7.16.

7.3.3 NSW Trigger

Trigger inputs from both the sTGC and the Micromegas chambers in a sector are processed by algorithms running on the NSW-TP cards (Figure 7.17).

7.3.3.1 NSW Trigger Processor requirements The NSW-TP provides track segments from the NSW detectors (as described in Section 5.2.8) to be matched in the Sector Logic with coincidences found in the EM-TGC. Each track segment is characterised by its radial position η , its angle $\Delta\theta$ with respect to an ideal infinite-momentum track (a line from the IP to the segment’s radial position in the NSW), measured by the precision strips of the sTGC or Micromegas detectors, and by its azimuthal position, ϕ , measured by the sTGC pad towers (see also Section 5.2.8.1) or by the output of Micromegas “diamond” fitter (explained in Section 7.3.3.4). The required η -resolution of both the sTGC and Micromegas strip triggers is 0.005. The

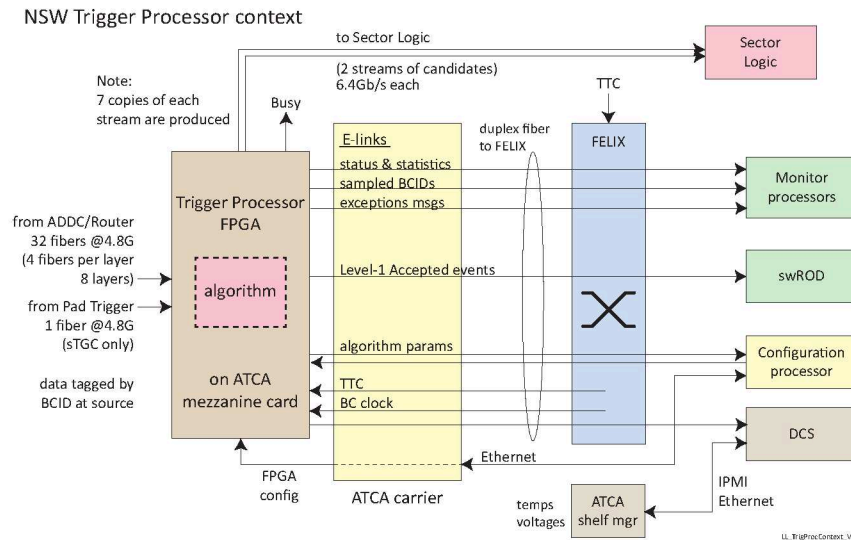


Figure 7.17: Functional diagram of the NSW-TP.

sTGC pad towers have an $r\phi$ -resolution of 7 to 38 mm, increasing with radius (based on the size of the logical pads), and the Micromegas trigger diamonds provide an $r\phi$ -resolution of better than 11 mm at all radii. The angular resolution of the NSW segments required for Phase-II is 1 mrad in order to match the expected angular resolution of the EM-TGC. For Run 3, $\Delta\theta$ is passed to the Sector Logic but not used in the Phase-I trigger decision.

7.3.3.2 NSW-TP mezzanine cards and input interfaces The NSW-TP is housed in two ATCA [151] crates in USA15. The NSW-TP processes information from the Micromegas and sTGC layers in the sector in separate algorithms (one for the eight Micromegas layers, and one for the eight sTGC layers) in separate Virtex-7 XC7V690 FPGAs on a mezzanine card. Each blade contains two mezzanine cards, to serve one NSW octant (one Large sector and one adjacent Small sector); a crate of eight blades serves a full wheel.

One FPGA on each mezzanine card takes 32 Micromegas fibre inputs (from the sixteen ADDCs per sector). Each ADDC GigaBit transceiver packet can contain ART data from up to eight triggered VMMs. When a packet is received, the ART data are decoded, for each strip, into a strip number and the slope associated with a line from the hit strips to the IP. The decoded data are provided to the Micromegas TP algorithm (Section 7.3.3.4), which runs on the FPGA.

The other FPGA on the mezzanine takes 32 sTGC fibre inputs from the Routers (four for each of the eight layers), plus two redundant fibres from the sTGC Pad Trigger (see Section 5.2.8.1) for the sector. Only data from strips passing through the tower selected by the sTGC Pad Trigger (see Section 5.2.8.1) are transmitted to the NSW-TP. A maximum of four such track segments per sector may be transmitted per BCID. The data from the Routers consist of a Band-ID identified by the Pad Trigger, the 6-bit ADC values from 14 strips, a flag indicating whether the 14 high or low strips of the 17 strips in the band are transmitted, along with the six low bits of the BCID, and the ϕ -ID (see Section 5.2.8.2). The decoded data are used by the sTGC TP algorithm (Section 7.3.3.3), which runs on the FPGA.

7.3.3.3 NSW sTGC Strip Trigger Algorithm The sTGC strip trigger uses the digitised peak data describing the quantity of charge deposited on each active strip selected by the Pad Trigger to compute a centroid for each of the eight sTGC layers and, from the centroids in the eight layers, track segments. A configurable charge threshold is used to define the “active” strips used in the centroid calculation. A configurable look-up table defines which patterns of active strips are accepted: wide clusters due to neutrons or δ -rays can be rejected; certain patterns with isolated hits near the cluster may be configured to be accepted. An average centroid is calculated from the valid layer centroids in each of the two wedges. An R -index, defined as the radial position of the trajectory projected onto a virtual plane passing through the nominal z -coordinate of the most downstream wire plane of the NSW, and $\Delta\theta$ are calculated directly from the values of the two wedge-centroids, using LUTs.

The sTGC trigger produces at most four candidates per sector; the Micromegas trigger can transmit up to eight candidates. The FPGAs are connected by a high-speed, low-latency 64-bit LVDS link. Since the Micromegas TP algorithm results are available sooner, Micromegas candidates are sent to the sTGC TP for merging, with no impact on the latency due to the data transfer. The merging algorithm limits the total number of merged candidates to at most eight, removing duplicates. Since ϕ from the Micromegas is more precise and R from the sTGC is more precise (see Tables 5.2 and 5.3), a duplicate segment can take its ϕ from the Micromegas and its R -index and $\Delta\theta$ from sTGC. Further details of the merging algorithm are given in Ref. [85].

7.3.3.4 Micromegas Trigger Algorithm The Micromegas trigger algorithm first converts incoming hits to slope values, based on the slope of the line between the hit strip and the ATLAS IP. Hits are stored in a buffer with a tuneable number of bunch crossings (currently five). A set of overlapping roads, eight strips wide, is formed for the η strips; the number of strip hits within a road required for further processing can be set between two and four (currently three).

The Micromegas trigger identifies “roads”: sets of strips in the eight Micromegas layers of a sector that are consistent with the path of a particle (with momentum above a given threshold) coming from the IP. In practice, Micromegas roads are twelve strips wide (5.1 mm or 5.4 mm), with the four outer strips at each side overlapping the neighbouring roads. Each wedge contains two layers of “eta” strips and two layers of “stereo” strips (one with a stereo angle of 1.5° and the other -1.5° with respect to the eta strips). When a coincidence of three η strips is found (of the four η layers in the sector), there is a range of possible stereo hits in both stereo projections. The intersection of two opposite-angle roads of stereo strips (also 12 strips wide) forms a diamond (see Figure 7.18(a)). The number of these diamonds required to span the full azimuthal extent of the chamber (see Figure 7.18(b)) depends on the radius. Seven diamonds are sufficient to span the full azimuth at the innermost radius of the Large sectors, while seventeen are needed at the outer radius. For the η road, all diamonds are sequentially parsed for possible stereo hits. The number of stereo hits required is also tuneable from two of opposite direction to a maximum of four. Currently the default is three, but there must be at least two with opposite inclination angles. Once the coincidences in the η roads and diamonds satisfy the coincidence requirements, the hit strips are then used to calculate slopes. See also Figures (a) and (b). Candidate diamonds are required to have hits in a configurable number of the eta layers in the sector and a configurable number of the corresponding stereo layers. A high-efficiency configuration would be at least two out of four hits in the eta layers and at least two out of four in the stereo layers, whereas a high-rejection configuration would require three out of four in both. Coincidences are considered within a configurable sliding window of time, in the range of four to eight BCs. These configurable parameters are chosen based on detector and electronics performance, and on LHC conditions.




Figure 7.18: The Micromegas trigger is built from “roads” 12 strips wide. (a) shows an η -road (grey), and two opposite-angled stereo roads (red and blue) which intersect to form a diamond. (b) shows one η -road (grey), and five positive and five negative-angled overlapping stereo roads, forming five diamonds spanning the full length of the η -road. Both axes on each plot use units of strip pitch. The η -road overlaps by four strip pitches with the η -roads immediately above and below it. The stereo roads also overlap by four strip pitches with their immediate neighbours, as shown. The trigger diamonds are the intersections of twelve stereo strips in each direction. The diamonds overlap with their neighbours, as shown by the black outlines. The Micromegas TP FPGA performs a fit to the strips in the roads within the selected diamonds. This fit determines the azimuthal resolution of the Micromegas trigger.

A local slope is calculated from the hits on the horizontal η strips and used to calculate $\Delta\theta$. The stereo hits are used to find the ϕ of the track segment. If a fit is found to be consistent with a projection back to the IP using $\Delta\theta$, its R -index is determined using a lookup table and sent along with $\Delta\theta$ for merging with the sTGC segments.

7.3.3.5 NSW-TP output interfaces The outputs of both TP algorithms are segment candidates. The NSW-TP requires a valid track segment to be within a configurable angle (up to ± 15 mrad) of the corresponding infinite momentum track through the IP. Each 24-bit segment includes a radial R -index, an azimuthal ϕ -index, a track angle $\Delta\theta$, and resolution flags that indicate whether the segment is from the sTGC (sTGC segments have one bit lower ϕ resolution than Micromegas), or whether one quadruplet had only a 3-out-of-4 coincidence, which implies worse $\Delta\theta$ resolution. The sTGC TP algorithm produces up to four segment candidates, and then accepts up to eight segment candidates from the output of the Micromegas TP algorithm running on the same mezzanine card. The TP algorithm running on the sTGC FPGA then merges the two lists of segments. The merging procedure removes duplicates and drops segments beyond the eight allowed. Priority is currently given to sTGC segments. There are options to ignore one or the other of Micromegas or sTGC segments and in case of duplicates, to take the ϕ -id from Micromegas and the other variables from the sTGC segment. [85]

The surviving segments are then sent via fast serial links from the NSW-TP to the Sector Logic described in Section 7.3.4, where endcap Muon candidates are formed by successfully matching NSW track segments with coincidences found in the EM-TGC (Big Wheels). Each NSW sector sends the track segment data to

the Sector Logic via optical fibres: up to eight candidates per sector per BCID on two fibres, each running at 6.4 Gb/s, four candidates per fibre, including the sector ID and BCID. Seven copies (14 links) are sent out from the NSW-TP sTGC FPGA to up to seven different endcap Sector Logic modules. (A single Sector Logic board receives data from at most six NSW trigger sectors, but each NSW trigger processor may need to deliver data to up to seven Sector Logic boards. This is needed to cover the overlap of NSW sectors with EM-TGC sectors, with multiple scattering, misalignments and magnetic field bending also taken into account.) For an overview of the NSW trigger-path electronics see Figure 5.10. The whole process repeats for every BCID with a fixed latency of 1075 ns from the time of collision to the time the signal reaches the Sector Logic.

Via the ATCA card's two FPGAs and Rear Transition Module, each mezzanine has two fibre connections to FELIX, one for sTGC and one for Micromegas, carrying, on different E-LINKs, the L1A event readout, exception messages, statistics, sampled events for monitoring, algorithm parameters, and TTC signals. The BC clock recovered from these links is distributed to the various FPGA and transceiver clocks via jitter cleaners.

A Zynq SoC FPGA on the ATCA card handles configuration of all the FPGAs and communicates various temperatures and voltages from elements of the NSW-TP via Ethernet to monitoring applications. An IPMC card on the ATCA card communicates critical temperatures and voltages to DCS via the ATCA Shelf Manager. If any temperature exceeds a configurable threshold, the shelf manager can shutdown a complete ATCA board or specific mezzanines.

7.3.4 New Endcap Sector Logic

Muon track candidates are reconstructed based on coincidence logic and assigned to transverse momentum intervals by the sector logic boards. For Run 3, new sector logic boards have been designed in order to receive the information from new detector components, and apply additional coincidence requirements. Inputs to the new sector logic boards include information from the legacy EM-TGC, EIL4 TGC, and Tile calorimeters, as well as the new RPC-BIS78 and NSW chambers. The new sector logic boards utilise a Xilinx® Kintex® -7 FPGA, featuring Multi-Gigabit Transceiver technology (GTX); each board has 12 GTX interfaces and 14 G-Link receivers. A sector logic board is shown in Figure 7.19.

The “endcap” region ($|\eta| < 1.9$) is subdivided into 48 independent trigger sectors per side in ϕ , as shown in Figure 7.20, and 24 sectors per side in the “forward” ($|\eta| > 1.9$) region. A sector logic board processes information from two adjacent sectors. In total, 72 new sector logic boards are required to cover the entire L1Muon endcap trigger system.

A block diagram of the L1Muon trigger sector logic is shown in Figure 7.21. Hit positions in the third station of the EM-TGC (η_{M3}, ϕ_{M3}) are extrapolated back to the IP, as shown in Figure 7.22, thus determining the “road” corresponding to the straight track of a muon with infinite momentum. Deviations from the centre of these straight roads ($\Delta R, \Delta\phi$) depend on the track momentum and are computed at the first TGC station (TGC-M1). Inputs from the NSW to the new sector logic include the position of the NSW track segment (η_{NSW}, ϕ_{NSW}) and the deviation of the reconstructed track segment with respect to the IP ($\Delta\theta$).

Look-up tables are used to determine the $R - \phi$ coincidence between signals from the EM-TGC and the NSW, EIL4 TGC chambers, Tile calorimeter D-layer cells, or RPC-BIS78 chambers. Tracks accepted by the sector logic are divided into intervals corresponding to the programmable p_T thresholds. Up to four of




Figure 7.19: A L1Muon sector logic board.




Figure 7.20: Schematic of the L1Muon endcap trigger segmentation for one octant of the detector, showing independent trigger sectors outlined in green and RoIs outlined in red. An example of the grouping of the detector elements into a trigger sector is shown on the left side of the figure.

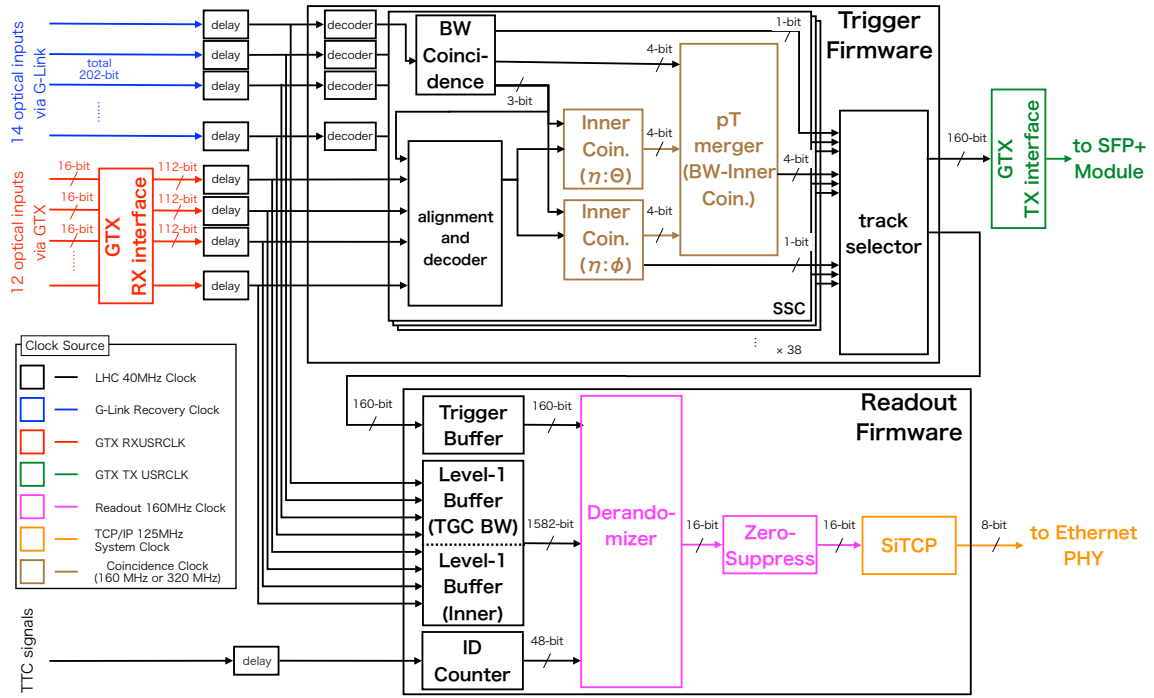


Figure 7.21: Block diagram of the L1Muon sector logic. The Big Wheel (G-Link) inputs are shown in blue and the NSW (GTX) inputs are shown in red.

the highest- p_T tracks per sector are sent to the MUCTPI. RoIs, spanning an area of $\eta \times \phi = 0.025 \times 0.033$, are sent to the HLT.

7.3.5 RPC Feet and Elevator Chambers

Muon triggers in the barrel region ($|\eta| < 1.05$) are provided by three concentric layers of RPC doublets. The L1 trigger decision in the barrel region relies on coincidence logic. For the low- p_T thresholds, a coincidence of three out of four layers in the middle station is required. The high- p_T trigger thresholds require the low- p_T trigger logic to be satisfied, as well as an additional hit in one of two layers in the outer barrel station.

During Run 1, the sectors 12, 13 and 14 of the barrel spectrometer had a trigger coverage approximately 20% lower than the other sectors due to the presence of the toroid feet support structure (sectors 12 and 14, corresponding to $-2.16 < \phi < -1.77$ and $-1.37 < \phi < -0.98$) and the elevator shaft ($|\eta| \simeq 0.7$, $\phi \simeq -1.57$). In LS1, additional chambers were installed in the toroid feet and elevator regions in order to increase the L1Muon barrel trigger coverage by about 3%, as described in Section 5.3.3.

In the feet region, the new chambers consist of two layers of RPC doublets, instead of the usual three layers. A two-layer coincidence is required for the high- p_T thresholds, but since there are fewer doublet layers in this region, the fraction of fake muon triggers in this region is higher than in the rest of the barrel. The impact of the new chambers in the feet region can be seen in Figure 7.23.

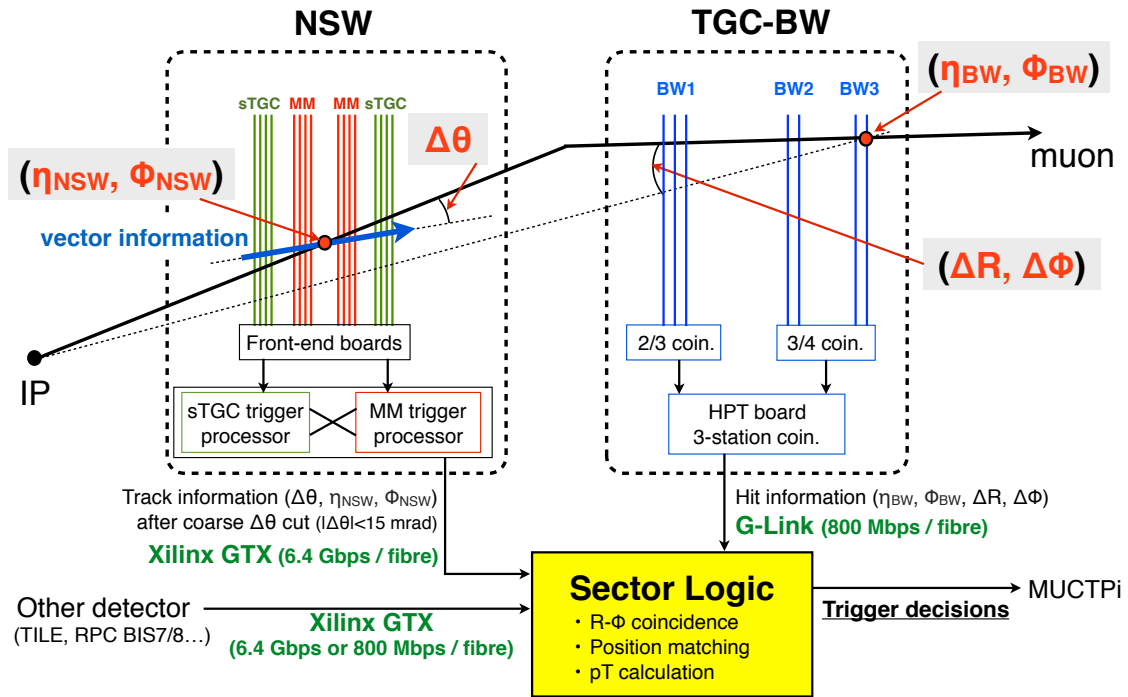


Figure 7.22: Overview of the L1Muon endcap trigger system, showing the tracking information used to compute the trigger decision. The solid line shows the trajectory of a muon with finite momentum, the dotted line shows the trajectory of a hypothetical infinite momentum muon used to define the centre of the trigger road.

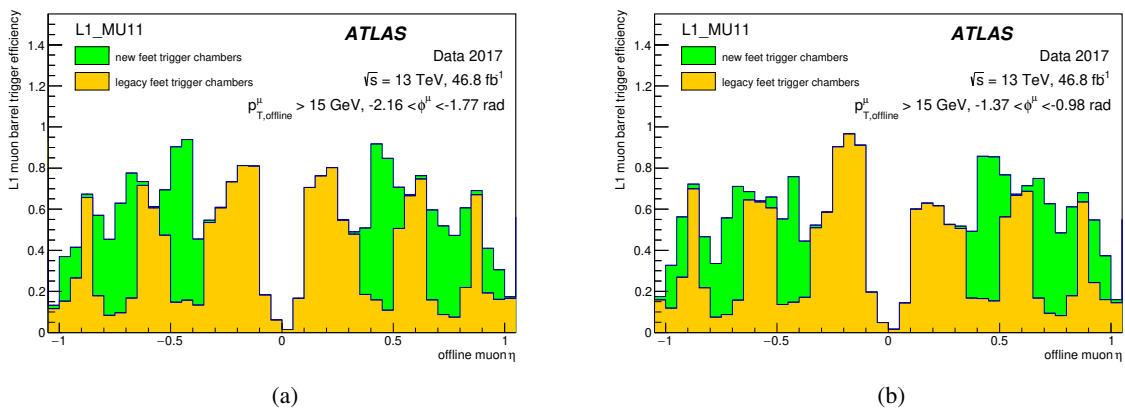


Figure 7.23: Trigger efficiency for muons with $p_T > 15$ GeV and satisfying a three-station coincidence requirement in the barrel region. The increased efficiency due to the additional chambers in the feet region can be seen for the two ϕ -sectors where the chambers were added: (a) Sector 12 and (b) Sector 14.

7.3.6 RPC BIS78

In the regions of the inner barrel (BI) that receive the highest background rates, additional RPC modules have been deployed, as discussed in Section 5.3.4. These new BIS78 (Barrel Inner Small [62]) chambers, covering η stations 7 and 8, are located on the A-side of ATLAS, in the transition region between the barrel and endcap ($1.0 < \eta < 1.3$). The high background rate in this region is due to secondary charged tracks originating from beam halo protons and a lack of detector shielding and instrumentation. The BIS78 chambers provide stand-alone detection and localisation of charged tracks and have an angular resolution of 3 mrad, which provides improved trigger rejection of low- p_T muons. These new chambers are a pilot project for the Phase-2 BI upgrade, which is planned to compensate for the reduced efficiency of the RPCs in the HL-LHC environment, and increase the geometrical acceptance of the muon trigger in the barrel with BI-BO combinations.

The trigger logic for the BIS78 chambers is performed on Pad boards [152], which are installed on each station. The Pad is an FPGA-based board which collects the BIS78 RPC hit data from the front-end electronics over 18 serial links at 1.6 Gb/s. It selects muon trigger candidates by requiring a local 2/3 coincidence of the RPC hits, applies a zero-suppression algorithm, and then sends the trigger information to the endcap sector logic board (described in Section 7.3.4), located off the detector. Readout data are sent to FELIX (Section 7.6.1) through optical links via a gigabit transceiver (GBTx) chip [153].

The Pad board utilises a Xilinx® Kintex® -7 FPGA, which supports optical transmission with fixed latency and provides robustness against radiation, including both single-event upsets and total ionising dose effects.

The expected rate reduction expected from the BIS78 RPCs is illustrated in Figure 5.2.

7.4 L1 Topological Trigger

The ATLAS physics programme relies significantly on electroweak-scale processes involving hadronically-decaying tau leptons, jets, and E_T^{miss} , such as $H \rightarrow \tau\tau$ and $ZH \rightarrow \nu\bar{\nu}b\bar{b}$. Also of importance are processes with unique topologies, such as B - or J/ψ -meson decays to low- p_T leptons or vector-boson fusion production of a Higgs boson, which then decays invisibly (Vector Boson Fusion (VBF) $H \rightarrow$ invisible). These signatures have large multi-jet backgrounds, and so the ability to reject the background and improve signal purity while still maintaining low trigger thresholds is of critical importance.

To accomplish this, a new L1Topo system was introduced in Run 2 and commissioned in 2016. It consisted of a single ATCA shelf with two FPGA-based processor modules which performed selections based on geometric or kinematic observables of TOBs received from the L1Calo and / or L1Muon systems, e.g $\Delta\eta$, $\Delta\phi$, ΔR , and invariant mass (M_{inv}). The details of the hardware design for the Run 2 system may be found in [143].

The L1Topo system has been redesigned for Run 3. It consists of three ATCA modules: one module computes L1Calo trigger multiplicities (Topo1) and the other two (Topo2 and Topo3) apply topological selections as in Run 2. (In the legacy system, the multiplicity task was performed by the L1Calo CMX [10].) Each module includes two Xilinx® Ultrascale+ processor FPGAs for algorithm computation, which provide increased processing power compared to the Run 2 version. Each FPGA has 118 input and 24 output fibres. Each module also contains two mezzanine cards: one for power and control, and another for communication with the CTP. An L1Topo module is shown in Figure 7.24.




Figure 7.24: A production L1Topo module.

L1Topo receives TOBs for jets, E_T^{miss} , e/γ clusters, and muons from the L1Calo and L1Muon systems via the TopoFOX optical plant (described in Section 7.2.4). The data received include the object type, (η, ϕ) coordinates, and transverse energy of the object, and its isolation (in the case of e/γ objects). Table 7.1 summarises the number of TOBs per object type sent to L1Topo.

Table 7.1: A summary of the maximum number of TOBs sent to L1Topo for each object type.

Object	Number of TOBs
e/γ (eFEX)	144
Tau (eFEX)	144
Forward e/γ (jFEX)	5
Large Tau (jFEX)	6
Small-radius jet (jFEX)	168
Large-radius jet (jFEX)	24
E_T^{miss} (jFEX)	7
E_T (jFEX)	7
Small-radius jet (gFEX)	6
Large-radius jet (gFEX)	3
E_T^{miss} (gFEX)	3
$\sum E_T$ (gFEX)	1
Muon	32

The L1Topo algorithms are implemented in a pipelined, deadtime-free manner using a fixed latency. Algorithms are distributed across the FPGAs as evenly as possible such that optimal resource usage is achieved. The configuration of algorithms, as well as their configurable parameters, are stored in the trigger menu.

The real-time output data sent by L1Topo to the CTP consist of individual bits indicating the algorithm decisions as well as an overflow bit for the topological algorithms. The output of the multiplicity algorithms

includes multiple bits indicating the number of TOBs fulfilling the requirements of a given algorithm. The outputs are sent to the CTP via optical or electrical cables.

Upon receipt of an L1A, the L1Topo real-time output data are captured and sent to the DAQ system. Readout to FELIX and the SW ROD is done via the ROD as described in Section 7.2.4.4. The clock and TTC signals are received from the Hub module via the backplane.

Control and configuration are performed via an IPbus interface on every FPGA, as well as on a control FPGA which controls the entire module. Voltage and temperature are monitored by the IPMC via an Inter-integrated Circuit (I²C) bus and are provided to DCS by the shelf manager.

Three types of L1Topo algorithms exist:

- **Sort / select / no-sort algorithms** These algorithms take all input TOBs and convert the various TOB formats to a single global data format. They also sort the TOBs by E_T or p_T , or select all TOBs with E_T or p_T exceeding a configurable threshold value, or satisfying object-specific criteria, such as isolation or substructure requirements. The output consists of reduced lists of sorted TOBs, which may then be used as inputs to the decision algorithms, described below. This reduction is necessary to handle the otherwise excessively large combinatorics. The final lists contain six TOBs in the case of “sort” algorithms and 10 TOBs for “select” algorithms. “No sort” algorithms only perform the conversion to a global data format and are used as inputs to the multiplicity algorithms. The latency of these algorithms is two bunch crossings (50 ns).
- **Decision algorithms** These algorithms determine whether a given trigger condition has been satisfied; examples include selections on $\Delta\eta$, $\Delta\phi$, and M_{inv} . Their outputs include decision and overflow bits which are sent to the CTP. The latency of the decision algorithms is limited to one bunch crossing (25 ns).
- **Multiplicity algorithms** These algorithms count the number of objects passing a given threshold, e.g. E_T , or located within a given region in η . The output multiplicity bits are transmitted to the CTP. The latency of the multiplicity algorithms is limited to three bunch crossings (75 ns).

Configurable parameters (e.g. threshold values) of the algorithms are specified in the trigger menu. The three algorithm types are illustrated in block format in Figure 7.25.

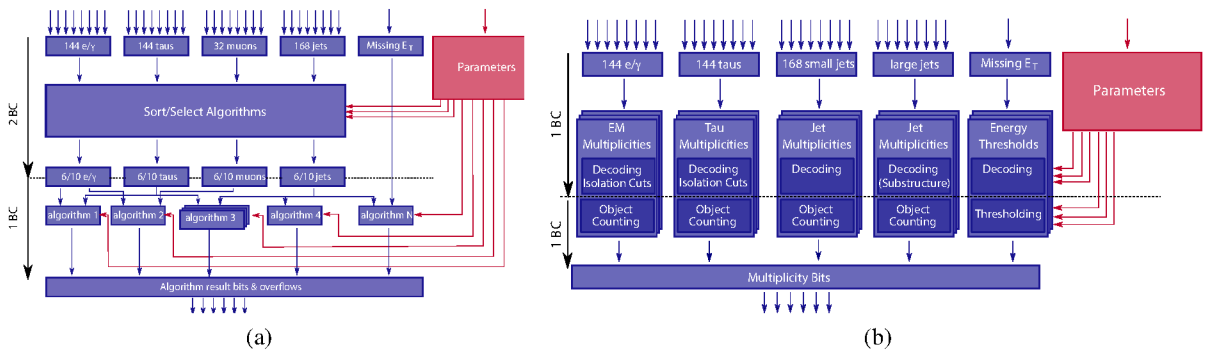


Figure 7.25: The L1Topo algorithm structure is illustrated in (a). TOB inputs are first passed to sort / select / no sort algorithms, which produce shortened lists of TOBs, which are in turn passed to the decision algorithms. Subfigure (b) shows the multiplicity algorithm structure which has as input all TOBs produced by eFEX, jFEX and gFEX. These TOB inputs are first passed to “no-sort” algorithms (labelled as “Decoding” in the figure). The resulting lists of TOBs in a global format are then passed to the multiplicity algorithms; the output multiplicity bits are sent to the CTP.

7.5 Central Trigger & TTC

The L1 CTP receives trigger information from the other L1 trigger systems and forward detectors and executes the final L1 trigger decision. Muon trigger inputs are received via the MUCTPI, which is described in Section 7.5.1. The CTP itself is described in Section 7.5.2.

7.5.1 Muon to central trigger processor interface

7.5.1.1 Functional overview The muon barrel and endcap trigger processors send their results as input to the MUCTPI, which has been redesigned and replaced during the Phase-I upgrade [154]. The replacement was necessary in order to provide full-granularity muon information at the BC rate to the L1 topological processor and to be able to interface to the sector logic modules using high-speed optical links.

The MUCTPI receives information for up to four muon track candidates per muon trigger sector from the endcap trigger sector logic modules, and up to two candidates from the barrel trigger sector logic modules. The information includes the position and p_T threshold passed by the track candidates (15 p_T thresholds for the endcap, and six p_T thresholds for the barrel), along with additional information, such as track quality flags, geometrical flags, and in the case of the endcap, a flag indicating the electric charge of the muon candidate.

The MUCTPI processes the information in three parallel paths:

- It combines the information from all trigger sectors to calculate the total multiplicity of muon candidates per muon threshold and sends the multiplicities to the CTP for each bunch crossing. Up to 64 bits of multiplicity information can be sent to the CTP for each bunch crossing. The multiplicity can either be a 2- or 3-bit value, and the maximum multiplicity value for a given p_T threshold includes the cases with even more muon candidates.
- It sends muon position and transverse momentum information of selected muon candidates to the L1 topological trigger system, to be used in subsequent topological trigger algorithms in combination with calorimeter information. The muon candidates sent are ordered according to decreasing p_T .
- It can apply muon-only topological trigger algorithms to be sent to the CTP.

Care is taken to avoid double-counting of muons which traverse more than one detector region due to geometrical overlap of the chambers and the deflection of the muons in the magnetic field. Many cases of overlaps are resolved within the barrel and endcap muon trigger processors. The remaining overlaps to be treated by the MUCTPI are those between neighbouring trigger sectors.

The MUCTPI also provides data to the HLT and to the data acquisition system for events selected at L1. A subset of the muon candidate information, ordered by p_T is sent to the HLT to be used as RoIs for further processing. The DAQ system records a more complete set of information, including the computed multiplicity values, which is used to monitor the functions of the MUCTPI.

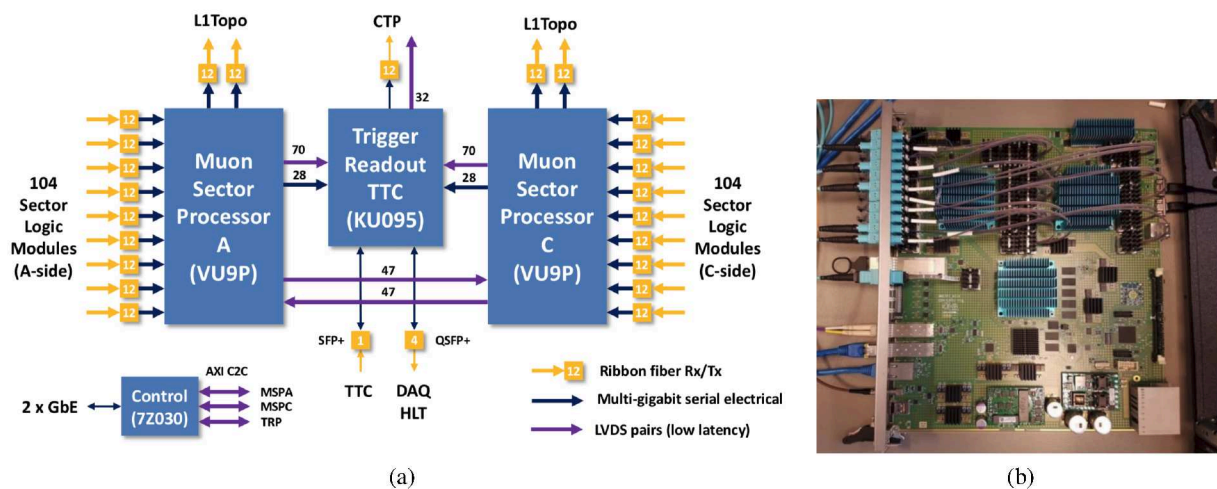


Figure 7.26: The MUCTPI module shown (a) as a block diagram and (b) in a photograph.

7.5.1.2 System implementation The Phase-I MUCTPI, shown in Figure 7.26, is integrated on a single ATCA blade, replacing the 18 9U VME cards of the Run 2 system. The architecture is based on a highly integrated generation of FPGAs, featuring a large number of on-chip MGs as well as 12-channel ribbon fibre optics receiver and transmitter modules (MiniPODs) for the data transfer.

The two Muon Sector Processor FPGAs – each taking care of one detector hemisphere – receive and process muon trigger data from 208 inputs from sector logic modules connected through high-speed serial optical links using MiniPOD receiver modules. The Muon Sector Processor FPGAs also copy information on selected muon trigger objects to several L1Topo modules using MiniPOD transmitter modules.

The Trigger and Readout Processor FPGA merges the muon multiplicity information received from the two Muon Sector Processor FPGAs and sends the results to the CTP. In addition, it can also implement muon topological trigger algorithms. This is possible because all the trigger information is available in a single module with low latency. The same FPGA also receives, decodes and distributes the TTC information.

A SoC is used for configuration, control and monitoring of the module. The device integrates a programmable logic part with a dual-core ARM processor subsystem. The processor subsystem runs the required software to interface the MUCTPI to the ATLAS run control system through a GbE interface. It will also be used for environmental monitoring of components of the board such as the power supply, optical modules, and FPGAs. The values read include voltages, currents, temperatures, optical input power, clock status, etc.

7.5.2 Central Trigger Processor

7.5.2.1 Functional overview The CTP is the last stage of processing of the Level-1 trigger system. It receives digital trigger information from the L1Topo, legacy L1Calo, MUCTPI and ALFA systems, and from various forward detectors. The CTP system used during Run 1 has been described in [3] and has since been upgraded.

Table 7.2 shows the various inputs of the CTP. Three different input paths are available:

- The traditional electrical trigger path via the CTPIN modules is used for trigger signals coming from various forward detectors, for calibration signals from some sub-detectors, and for special triggers such as a filled-bunch trigger based on beam-pickup monitors, and a minimum-bias trigger based on scintillation counters. It is also used for the legacy L1 calorimeter trigger system during the start-up phase of Run 3 and the commissioning of the L1 topological trigger system.
- The electrical low-latency input via the CTPCORE+ module, which is used for latency-critical trigger signals, such as those from the L1 topological processor, muon-only topological triggers from the MUCTPI system, and signals from the ALFA detector. For the L1 topological processor, these are signals from the Topo2 and Topo3 modules: Topo2 provides topological algorithms that combine jets with one or more other objects, comprehensive combinations of jet and τ , and jet and EM clusters, along with muons and missing transverse energy where appropriate, while Topo3 is largely dedicated to topological triggers typically involving electrons, ranging from the standard model performance triggers, to exotics triggers involving many lepton flavours.
- The optical input via the CTPCORE+ module is used for the muon threshold multiplicities from the MUCTPI and for non-latency critical L1 topological algorithms, mainly coming from the Topo1 module. These include simple multiplicity triggers on electromagnetic, hadronic τ , and jet objects, energy sums, missing transverse energy, and other fast topological algorithms.

Table 7.2: Overview of the trigger inputs to the CTP in Run 3.

CTP input	Cable origin	Number of bits	Trigger information
CTPIN	Various sources	36	Forward detectors, calibration, special triggers
CTPCORE electrical	MUCTPI/ALFA	77–149	Muon-only topological algorithms and ALFA
	Topo2, Topo3		Topological algorithms
CTPCORE optical	MUCTPI	64	Muon threshold multiplicities
	Topo1	263–335	Simple multiplicity triggers, energy sums, missing energy

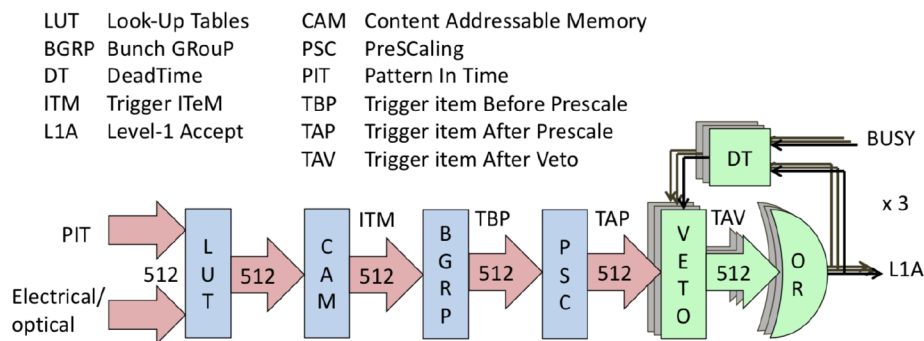


Figure 7.27: Block diagram of the trigger path in the CTP.

The trigger path of the CTP is shown in Figure 7.27. Up to 512 trigger inputs can be used in programmable

look-up tables to form trigger conditions from the input signals, as described in more detail below.

Each trigger item can be put in coincidence with any combination of the 16 bunch groups, which are programmable collections of bunch-crossings.

In the next step, the rate of each trigger item is controlled by pre-scaling, where an algorithm rejects pseudo-randomly on average a certain fraction of the trigger item signals. The fraction is called the pre-scale factor and can be individually set for each trigger item. Each trigger item can subsequently be enabled or disabled to be used in the trigger decision. The L1A signal generated by the CTP is the logical OR of all enabled trigger items.

For each L1A signal, the CTP provides an 8-bit trigger-type word signal, whose bits indicate certain categories of types of triggers and can be used to select options in the event data processing and monitoring in the sub-detector readout chain. For each L1A signal, the CTP also sends information about the trigger decision for all trigger items to the HLT and to the data acquisition system. The CTP can support up to three data-taking partitions for calibration and test runs. The trigger items after pre-scaling are partitioned and assigned to the three partitions, each of which individually enables the trigger items of interest, gates them with the veto signal, forms the L1A of the partition, and calculates the deadtime taking into account the L1A and the BUSY signals of the corresponding partition.

In addition to its function as trigger processor, the CTP is also the timing reference of the detector. It receives the beam-synchronous clock and timing signals from the LHC and distributes them together with the L1A signal to the ATLAS sub-detectors, preserving their timing information. It also interfaces with the LHC Global Positioning System (GPS) timing system and attaches a nano-second precision absolute time-stamp to the read-out data of each L1-accepted event. In addition, it is in charge of generating luminosity blocks, which divide a data acquisition run into small time intervals of typically one minute. These luminosity blocks are the shortest time interval for which the integrated luminosity, corrected for deadtime and pre-scale effects, can be determined. They are also used as the time interval for online and offline data quality assessments.

The CTP features many monitoring facilities that allow the monitoring of the rate of each incoming trigger signal, and of each trigger item at the different stages of trigger processing in the CTP. These rates are determined, published online, and preserved for offline usage, every few seconds and on luminosity block boundaries. Some of the trigger rates, such as those before and after veto, are used to correct the integrated luminosity for deadtime losses. The deadtime and its components are monitored directly by corresponding deadtime counters. In addition, the CTP has counter facilities that allow the monitoring of trigger rates and deadtime fractions per BCID and thus the study of bunch-by-bunch effects of the trigger rates and deadtime.

7.5.2.2 System implementation The CTP for Run 1 has been described in detail in [3] and has since been modified during LS1 and LS2. The upgraded CTP system consists of six types of modules which are housed in a single 9U VME crate, as shown in Figure 7.28:

- The CTPMI module as the machine interface module.
- Three CTPIN modules to receive trigger input signals.
- The CTPCORE+ module to receive additional direct trigger input signals and implement the trigger logic.

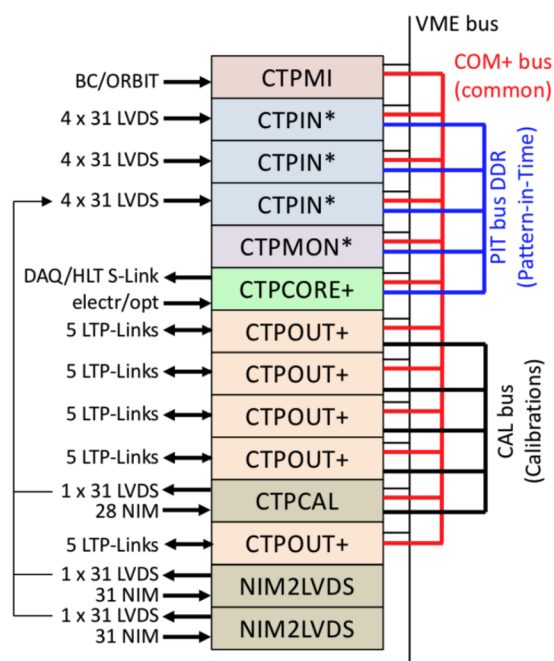


Figure 7.28: Block diagram of the various components of the CTP. Components marked with + were newly built during LS1, while components marked with * use the Run 1 hardware, but with upgraded firmware. The firmware of the CTPCORE+ was further upgraded during LS2.

- The CTPMON module for per-bunch-monitoring of trigger input signals from the CTPIN.
- Five CTPOUT+ modules as interfaces to the sub-detector TTC distribution.
- One CTPCAL module to receive calibration request signals from sub-detectors and receive electrical signals.
- Two NIM to LVDS conversion (NIM2LVDS) boards to receive further electrical signals.

The controller CPU interacts with the modules via the VME backplane. Internal communication between the modules proceeds by custom bus systems implemented as custom-built backplanes installed in the CTP crate: The PITbus is used for routing synchronised and aligned trigger inputs from the CTPIN boards to the CTPCORE+ module, the COMbus+ is for the interchange of common timing, trigger and control signals between all the modules, and the CALbus is used for sub-detector calibration requests.

For Run 2, the following components of the CTP were upgraded:

- The firmware of the 3 CTPIN modules was upgraded to drive trigger input signals to the PITbus in double-data rate at 80 MHz, hence doubling the number of input signals available on the PITbus from 160 to 320.
- The CTPCORE board was replaced by the newly designed CTPCORE+ board, which allows trigger logic to be applied to up to 512 trigger inputs (to be compared with 256 inputs in Run 1). The inputs are a programmable selection of signals from the PITbus and additional 192 direct electrical inputs from the L1Topo trigger processor and the ALFA sub-detector. The CTPCORE+ can form up to 512

trigger items (256 in Run 1) that can be put into coincidence with up to 16 bunch groups (eight in Run 1). It also features the possibility to support up to three data-taking partitions for calibration and test runs.

- The four CTPOUT boards have been replaced by five newly designed CTPOUT+ boards, each supporting running with up to three concurrent data-taking partitions, and featuring additional monitoring and testing facilities.
- The COMbus backplane was replaced by the newly designed COMbus+ backplane, which allows the distribution of timing signals for up to three concurrent data-taking partitions, and includes the communication with a fifth CTPOUT+ board.
- The CTPMON firmware was updated to select up to 160 trigger inputs from the PITbus for per-bunch monitoring.

For Run 3, the CTPCORE+ firmware has been updated to receive optical trigger signals from L1Topo via up to 12 optical fibres running at a link speed of 6.4 GBaud, corresponding to 96 trigger bits per optical fibre. A programmable switch matrix maps the incoming trigger signals to up to 512 usable trigger inputs.

The timing signals from the LHC are received by the CTPMI, which can also generate these signals internally for stand-alone running. This board also controls and monitors the internal and external busy signals; for example, the busy signal transmitted from a sub-detector in case of overload on its data acquisition system. The module sends the timing signals to the COMbus+, thereby making them available to all of the other modules in the CTP.

The CTPIN modules receive trigger inputs from forward detectors, the legacy L1Calo system, and from various other sources, such as calibration systems, the beam-pickup monitors, and a minimum-bias trigger based on scintillation counters. The input boards select and route the trigger inputs to the PITbus, after synchronising them to the clock signal and aligning them with respect to the bunch-crossing. Three boards with four connectors of 31 trigger input signals each allow for a total of 372 input signals to be connected, of which up to 320 can be made available on the PITbus at any given time, using double-data rate at 80 MHz.

The trigger decision module (CTPCORE+) receives the trigger inputs from the PITbus and from two external sources: via three electrical connectors of 64 trigger input signals each, and via up to 12 optical fibres containing up to 96 trigger bits each. A selection of these trigger input signals is made using a programmable switch matrix. The CTPCORE+ module combines the trigger input signals from the PITbus, the electrical and the optical CTPCORE+ inputs, and internally generated signals, using several programmable look-up tables to form up to 512 trigger conditions. In a further step the trigger conditions are combined using content-addressable memories to form up to 512 trigger items. Any of the up to 512 trigger conditions may participate in any of the up to 512 trigger items. The trigger items are put into coincidence with individual combinations of up to 16 bunch groups, subsequently undergo pseudo-random prescaling and deadtime gating, and can be individually enabled or disabled to take part in the final L1A trigger decision. The trigger results are transmitted to the COMbus+. The CTPCORE+ module also acts as the readout driver of the system, sending information to the HLT trigger and the data acquisition for each accepted event.

The output module (CTPOUT+) receives the timing and trigger signals from the COMbus+ and fans them out to the sub-detectors. The module receives back from the sub-systems the busy signals, which are sent to the COMbus+, and 3-bit calibration trigger requests, which are routed to the CALbus. The calibration module time-multiplexes the calibration requests on the CALbus and sends them via a front-panel cable to

one of the input modules. The calibration module also has front-panel inputs for beam pick-up monitors, minimum-bias scintillators, other forward detectors, and test triggers. In addition, two NIM2LVDS modules provide further front-panel inputs for such trigger signals.

The TTC signals of the CTP are distributed electrically to the sub-detectors via the CTPOUT+ links, through which in return the sub-detector busy signals and some calibration request signals are received. A local trigger system acts as an interface between the CTP and the TTC distribution of each sub-detector, and imitates the function of the CTP during stand-alone data taking of the sub-detector.

During Runs 1 and 2, the local trigger system consisted for all sub-detectors of a series of custom 6U VME electronics boards [141]: an LTP module as the TTC switch board and generator of local signals, an LTPI module for interconnections between sub-detector partitions, a TTCvi for the serialisation of the TTC information, and a TTCex for the optical transmission to the sub-detector front- and back-end electronics, where TTC receiver chips (TTCrx) decode the transmitted information and make it available as electrical signals for further use.

During LS2, the ALTI has been introduced as the local trigger system for new Phase-I sub-detector partitions and to replace the legacy set of TTC modules (LTPI, LTP, TTCvi and TTCex) for all sub-detectors, except for the RPC, TGC, and LUCID detectors.

The ALTI module provides the same electrical and optical interfaces as the legacy set of TTC modules in a single 6U VME electronics board. Using a single modern FPGA, all the functions of the legacy set of TTC modules are replicated. New, useful functions were added, such as an optical input and memory to analyse optical TTC signals, memories to store incoming electrical signals, additional monitoring features including per-bunch monitoring of trigger signals, and a range of specific functions serving to generate L1A sequences similar to the ones from the CTP. These latter functions include a programmable look-up table for defining the trigger logic of the trigger inputs and generated trigger signals of the ALTI board, pseudo-random trigger generators, pseudo-random pre-scaling, bunch group masking, and simple and complex deadtime generation. Where possible, the same algorithms as in the CTP were used.

From the ALTI module or the legacy set of TTC modules, the timing signals are distributed to the detector electronics using the TTC system. The implementation and use of the TTC system is sub-system specific. As an example, the muon trigger systems use TTC standard components to transmit the timing signals all the way to the electronics mounted on the chambers, while in case of the inner tracking detector a custom-built distribution system is used to transmit the signals from the counting rooms to the cavern.

7.6 DAQ/HLT

7.6.1 FELIX/SW ROD

FELIX and the SW ROD are new detector readout components introduced into the ATLAS DAQ system for Run 3. FELIX is designed to act as a configurable data router, receiving packets from detector front-end electronics and transferring them to peers on a commodity high-bandwidth ethernet network. Whereas previous detector readout implementations relied on diverse custom hardware platforms, known as RODs, as the interface between the detector electronics and the common DAQ chain, the motivation for FELIX is to unify all readout across one well supported and flexible platform. As well as its readout function, FELIX will also serve as a relay for trigger accept and clock information from the TTC system to front-end

electronics. It will also be possible to use FELIX to send general purpose control data to front-end electronics to manage modules throughout data taking and calibration.

Another key aspect of the FELIX concept is to take advantage of advances in technology, for example: larger and faster FPGAs, the advent of multicore CPUs and high-performance networking, to move tasks which were previously performed in customised hardware (the aforementioned RODs) into the more flexible firmware and software domains. In this new architecture, with FELIX acting as a simple router, detector data processing, monitoring and control functions are instead implemented in software hosted by commodity server systems subscribed to FELIX data. The primary peer on this network will be the SW ROD, which will perform functions such as event fragment building, data formatting and other detector-specific processing to prepare data and facilitate online selection. The SW ROD will also buffer event fragments and supply them on request to HLT nodes, via an identical interface to the legacy ROS.

FELIX and SW ROD-based readout paths will exist alongside the legacy ROS system for the duration of Run 3. An overview of the architecture is presented in Figure 7.29. FELIX and SW ROD installations were deployed during LS2 for systems undergoing significant detector or readout upgrades in preparation for Run 3 operation. These are the NSW, LAr digital readout and new L1Calo systems described earlier in this chapter. Smaller scale demonstrators for upgraded Barrel RPC-BIS78 and the Tile Calorimeter were also installed during this shutdown. In total the Run 3 installation comprises about 60 FELIX servers (hosting 100 I/O cards between them), with 30 software ROD servers. This is similar in size to the legacy ROS system in Run 3, which consists of approximately 100 servers. The remaining ATLAS systems will then all be migrated to FELIX during the next long shutdown before the HL-LHC, resulting in the final decommissioning of the legacy ROS.

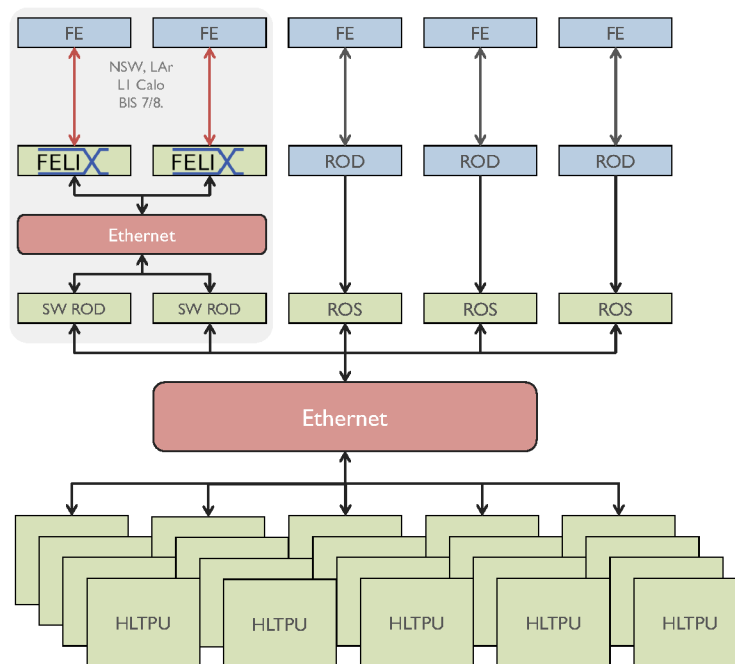


Figure 7.29: Comparison of Readout Architecture in Run 2 (right-hand side, with hardware RODs) and Run 3 (pink box on left-hand side) showing new FELIX readout paths. High-Level Trigger Processing Unit (HLTPU)s are able to sample event data from both paths via an identical interface.

FELIX systems are able to interface with front-end electronics over one of two optical link protocols: GigaBit Transceiver (GBT [153]), a radiation-hard standard developed at CERN, where multiple lower speed links (E-links) from separate pieces of electronics can be aggregated into a single 4.8 Gb/s link; and FULL mode, an in-house design with no link substructure for higher bandwidth (9.6 Gb/s) communication between FPGAs. Data streams for either protocol can be configured to use different encoding, although 8b10b is typically used for normal dataflow.

7.6.1.1 FELIX & SW ROD Server Hardware Each FELIX server is 2U high and hosts custom I/O cards (FLX-712, described in Section 7.6.1.2), with firmware to interface with either of the two link protocols. For the GBT case each server hosts two cards, for the higher bandwidth FULL mode case each server hosts one card (driven primarily by the number of available Peripheral Component Interconnect Express (PCIe) lanes). Each server also hosts high-bandwidth network interface cards (dual-port 25 GbE for GBT, dual-port 100 GbE for FULL mode). Each FELIX server has an Intel® Xeon® E5-1660 V4 CPU (8 cores 3.2 GHz) and 32 GB of Double Data Rate (DDR)4 Error Correction Code (ECC) Random Access Memory (RAM). Each SW ROD server is 1U high, featuring dual Intel® Xeon® Gold 5218 CPUs (16 cores 2.3 GHz) and 96 GB of DDR4 ECC RAM. Each server also hosts a dual-port 100 GbE network interface.

7.6.1.2 FELIX I/O Card The FLX-712, shown in Figure 7.30, is a PCIe card supporting a 16-lane Gen 3 interface, able to reach a throughput of up to 100 Gb/s. An MTP 24 or 48 coupler provides the interface to external data fibres, after which the light is internally routed to one of eight MiniPOD transceivers (four for reception and four for transmission) handling 12 links each. A maximum of 48 bi-directional optical links can therefore be connected to each board. A Xilinx® Kintex® Ultrascale™ (XCKU115) FPGA provides the platform for all on-board firmware features. An on-board PEX8732 PCIe switch makes it possible to map two separate PCIe 8-lane endpoints into one 16-lane bus. A JTAG connector is provided to facilitate FPGA configuration, though this may also be stored in an on-board FLASH chip. FPGA programming and card health monitoring and control are also possible over PCIe.

Finally, an interchangeable mezzanine card provides an interface for a number of timing and control systems. The ATLAS TTC system in Run 3 connects to FELIX via an optical fibre for the distribution trigger and clock information and a LEMO connector for the receipt of BUSY signals. An SI5345 jitter cleaner on-board the FLX-712 itself ensures a sufficiently good quality clock for all FELIX use cases.

During LS2, approximately 60 FELIX servers hosting a total of 100 FLX-712 cards were deployed, routing data to 30 SW ROD systems. A significantly larger number (of order 6 times more) will be deployed in LS3 to service all remaining ATLAS systems.

7.6.1.3 FELIX Firmware The FELIX firmware, a diagram of which can be found in Figure 7.31, is designed to be modular and flexible. Separate components manage different key functions, such as the link wrapper (GBT or FULL mode) and the PCIe and Direct Memory Access (DMA) [155] engines. Between these lies the Central Router module, which performs the most data intensive workload. Here data arriving over different links are decomposed according to protocol and converted into regular 1 kB elements for optimal DMA transfer to the host server's memory. In order to optimise FPGA resource utilisation and timing, the FELIX firmware deployed in the FLX-712 consists of two duplicate paths with identical modules, each servicing half the input links and reading out to an 8-lane PCIe interface. As such, each FLX-712 card appears to the host server as two 8-lane devices.




Figure 7.30: Run 3 FELIX hardware platform (FLX-712), with key components labelled.

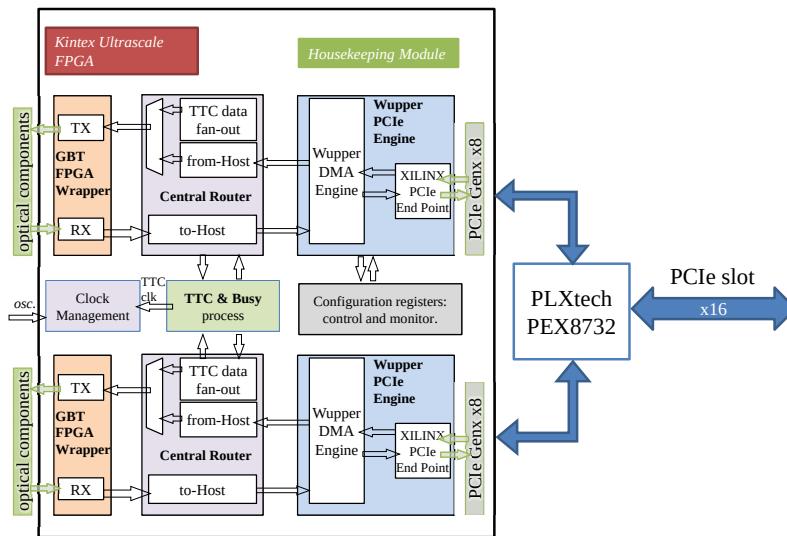


Figure 7.31: Diagram of the firmware deployed on the FPGA of the FLX-712 in GBT mode. In FULL mode the incoming (RX) GBT link wrapper is replaced with a dedicated FULL mode module. The Central Router is also much simpler in this case, as the data processing requirements are less severe.

Alongside the primary dataflow path, a separate common module interacts with the TTC system, injecting trigger and clock signals into the data path and relaying BUSY signals back to the central trigger should operating conditions require a pause in dataflow. The TTC module is also responsible for generating an information packet for each trigger accept received on a special stream for downstream subscribers to use to facilitate event fragment building and synchronisation. Other common modules also manage configuration registers, clock control/distribution and general housekeeping.

By re-using the basic blocks above it is possible to flexibly produce firmware designs for different use cases. For ATLAS, separate designs are maintained for both GBT and FULL mode, where the primary differences are the link wrapper and the complexity of the Central Router (which is significantly lower for FULL mode). Due to FPGA resource utilisation constraints the maximum number of GBT links which can be supported for primary dataflow is 24. For FULL mode the limitation comes from the PCIe bandwidth of the FLX-712, which can accommodate a maximum of 12 links. However, the standard approach is to build firmware to support up to 24 links, this giving the option to operate more than 12 at lower occupancy. Thus a higher link density can be provided within the same bus constraints.

7.6.1.4 FELIX Software The FELIX software suite comprises high- and low-level components. Alongside a dedicated device driver, low-level tools make it possible to test all firmware features in a laboratory setting and debug any issues which may arise. At a higher level, a high-performance daemon operates in an “always on” fashion in order to receive data from the FLX-712 over DMA and provide onward routing. DMA transfers are received via a separate ring buffer for each PCIe device visible to the host server (hence two per FLX-712). The software daemon, based on the NetIO architecture [156] is designed to be event driven, able to react to hardware interrupts from the card, indicating incoming data, or signals from the network interface or operating system. The design is such that copies of the data in memory are kept to a bare minimum to maximise throughput. Finally, Remote Direct Memory Access (RDMA) technology is used to transfer data to connected network peers without intermediate processing. An overview of the complete software stack is presented in Figure 7.32.

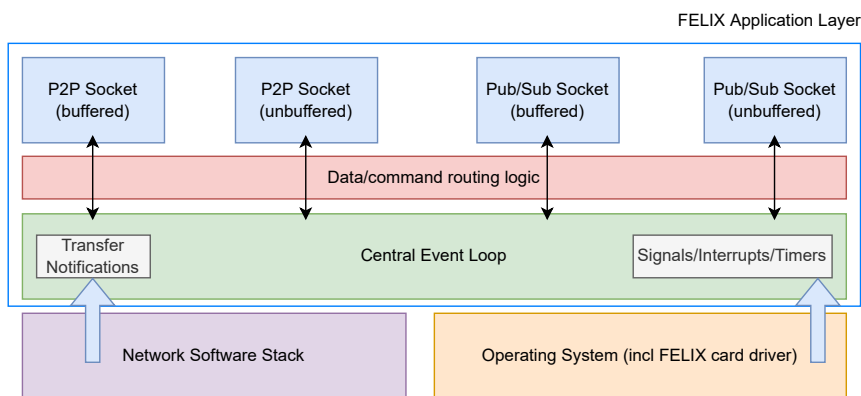


Figure 7.32: Diagram showing the different layers of the FELIX software architecture, with the central event loop (shown in green) able to interact with low level network software (supporting Remote Direct Memory Access (RDMA)) and operating system signals (including interrupts from the FELIX I/O card). The higher-level layer implements application-specific logic and handles subscriptions from external clients via dedicated sockets, available both in peer-to-peer or publisher/subscriber mode. Sockets can also either operate in “buffered” mode, whereby data are packed into an output buffer before transfer across the network (useful to optimise throughput for small packets) or “unbuffered” mode, where data are transferred immediately (useful for larger packets).

7.6.1.5 SW ROD Design The SW ROD is implemented as an ATLAS Run Control-aware application running on dedicated servers. The application is designed to be able to support flexible workloads, with configurable event fragment building and processing. The overall architecture of the SW ROD application is shown in Figure 7.33. At the input stage, dedicated “Reader” threads subscribe to FELIX data streams and write incoming packets into “slice” buffers, with each buffer made up of “vectors” of data packets corresponding to a specific Level-1 Trigger Identifier (L1ID) for a specific input link. From here packets are placed into an associative “ReadOut Buffer (ROB)Fragment assembly” map structure, with all data in a previously assembled slice associated with their corresponding L1ID. Thus the aggregation from many packets per L1ID to one ROBFragment object happens in two steps. At this point, if required, fragments are placed into a “Ready” queue for handling by a subdetector-specific plugin, operating in a dedicated user processing thread. Within this thread, a variety of fragment analysis, error checking, monitoring and statistics gathering can take place, making it possible to replicate the actions of the old generation of hardware RODs which the SW ROD replaces. Should such processing not be required fragments can be sent directly to the output stage. At this stage, fragments are held in a “ROBFragment” map while the HLT performs selection operations on the events of which they form part. A pool of request processing threads transfers any requested fragment data from the map to the HLT during this period. This final stage is designed to be functionally identical to the legacy ROS buffers, making use of a common interface to the HLT, such that any HLT node requesting data makes no distinction as to whether it comes from a ROS or a SW ROD.

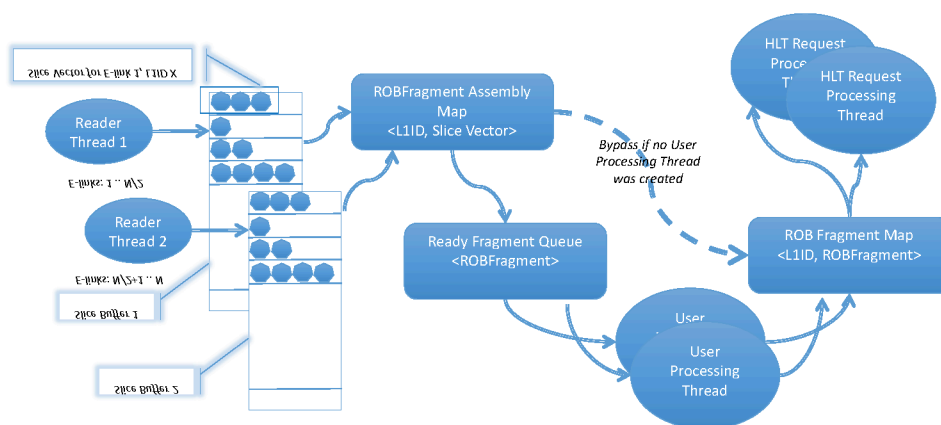


Figure 7.33: SW ROD Architecture with subdetector plugins. Data are aggregated by common L1ID in the input stage before being passed (if needed) to dedicated subdetector code, before being buffered and transferred onward as needed during HLT processing. In this example only two Reader Threads are used for clarity, but in principle more can be configured as needed.

7.6.2 Network

As part of the preparations for Run 3, the TDAQ network is undergoing a major programme of updates. The overall design is based on two physical levels, with machines operating in the electronics cavern (FELIX, SW ROD, ROS and various DCS and detector infrastructure nodes) communicating with systems

on the surface (HLT, core infrastructure, control calibration and monitoring nodes) via a 40 GbE backbone. Different communication workloads are supported, from bulk dataflow to diverse control and monitoring traffic.

While the overall core throughput between the electronics cavern and the surface has not changed since Run 2 (remaining at 40 GbE), the overall number of nodes connected to the network has increased significantly, with more extensive deployment of virtual networks to serve different functions, eliminating the need for further physical interconnects. The network's router infrastructure has also been replaced, both at the core router level (now two Juniper QFX10016s) and all client switches at various levels (now made up of 50 Juniper QFX5100), with improved performance and the introduction of active-active redundancy (based on MC-LAG technology). All routers now make use of the next generation of ASICs (Broadcom Trident II in the client switches and Juniper Q5 in the core routers). A mixture of copper and optical fibre connectivity is used, so as to optimise the cost for each use case. The overall architecture of the network for Run 3 is shown in Figure 7.34, along with a more detailed description of node connectivity.

The majority of new clients on the TDAQ network belong to the new FELIX and SW ROD readout paths. Each FELIX server hosts a dual 25 or 100 GbE network interface card, with each SW ROD hosting a dual 100 GbE interface. The servers are connected via dedicated high-performance switches (a total of six Juniper QFX 5200), which feature lossless operation over RDMA and active-active redundancy. Each SW ROD then features a separate dual 40 GbE interface to the common dataflow network, for control and bulk dataflow through to the HLT and beyond. Further new nodes may also be added to the network in future as the HLT farm continues to expand on the surface. As such, additional contingency has been factored into the specification of the core network to allow for such evolution to take place.

As mentioned before, the physical network is subdivided into separate Virtual Local Area Networks (VLANs), each with a dedicated quality of service policy, to enable efficient management of separate workloads without extra cabling. There are five separate VLANs in operation: Management, for network control and monitoring; Control, for server control (DHCP, NFS, DNS, etc.) and DAQ control; Data, for DAQ traffic; Monitoring, for Low-priority monitoring and backup (operated on a best-effort basis) and finally SimP1: for simulation jobs running in the HLT farm. With this implementation, virtual router instances also provide traffic isolation and enhanced security (e.g. the ATLAS technical network is protected against traffic flooding).

Beyond the upgrades described above, significant effort has gone into reworking the cabling structure in the surface cavern and improving uninterruptible power supply coverage for the DAQ network.

7.6.3 HLT/AthenaMT

Event processing frameworks for most LHC experiments have traditionally been designed to process single events serially on a single CPU core, with events distributed between independent processing nodes. Unfortunately, for offline processing, such a model no longer matches trends in computer architecture. Crucially, while the number of CPU cores available in a standard compute node has increased, the amount of memory per core is increasing at a lower rate. In order to effectively make use of the increasing core count, the average memory utilisation per job must therefore be reduced. The most effective way to achieve this is to share memory between cores. This can be achieved by sharing event processing between multiple threads, i.e. multi-threading, thus reducing the overall memory footprint per core. Memory has never been a limiting factor for the ATLAS HLT, due to the different operational model, but the system can still benefit

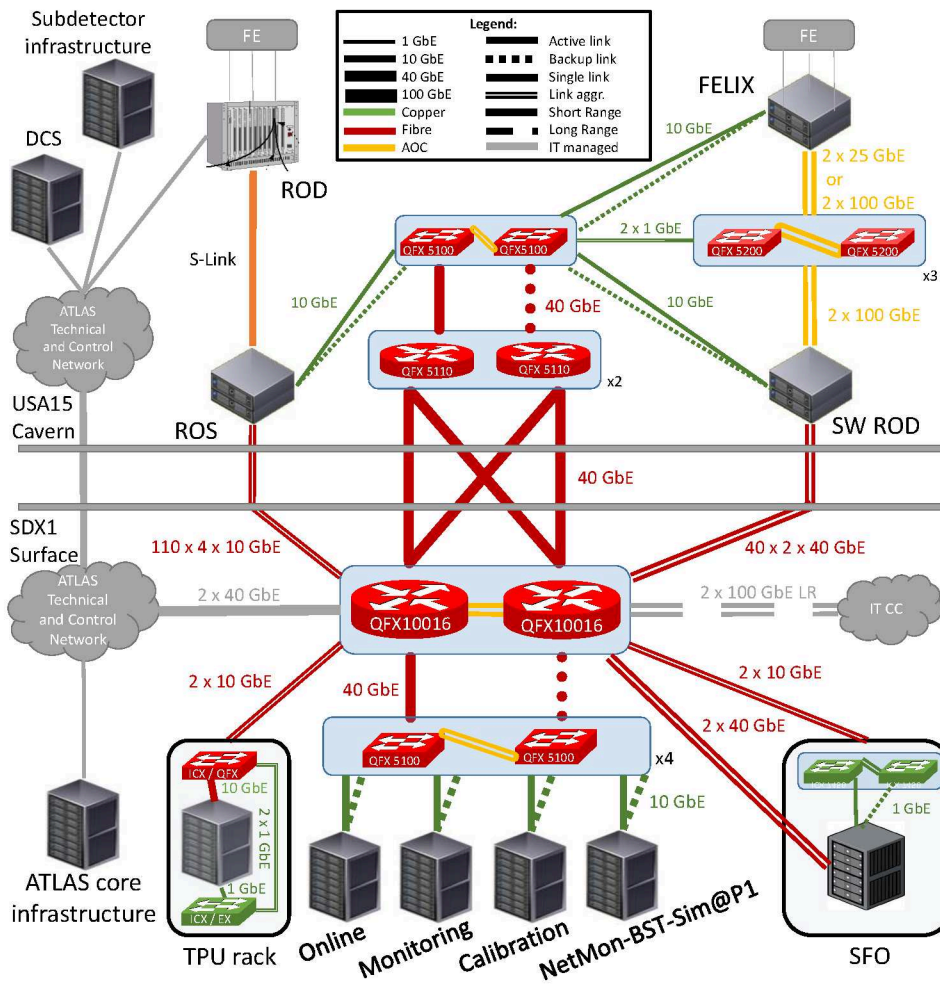


Figure 7.34: TDAQ Network for Run 3, showing upgraded core routers (QFX10016 in the centre) and new FELIX/SW ROD network (top right). Client routers (QFX5100) act as the interface between the core and the processing farms performing online control, monitoring and calibration functions. Nodes participating in primary dataflow (such as ROS, SW ROD, TPUs and SFO) are connected directly to the core. In the case of ROS, FELIX, and SW ROD, separate connections for the purposes of node control are also present. The ATLAS Control Network (ATCN) operates in parallel, connecting DCS and ATLAS core infrastructure, and interfaces to the TDAQ network via the core and via connections to hardware ROD components.

from any optimisation to gain additional margin for future evolution, and to aid with other workflows with differing resource limitations such as HLT simulation as part of MC event generation.

To tackle the memory issue, a significant redesign of the ATLAS software framework (Athena) was undertaken to allow it to process events across multiple threads. The new AthenaMT framework to be used for Run 3 positions ATLAS software to optimally exploit future server technology evolution. An additional feature of AthenaMT is its built-in hooks to facilitate potential future use of co-processors, such as GPUs and FPGAs. Asynchronously offloading compute-intensive tasks to these devices can free up CPU cores for work better suited to the CPU. While the use of co-processors is not planned for the start of Run 3, the structure of AthenaMT means they can be integrated in future without major architectural changes.

The Athena framework is mainly written in C++, with a configuration layer in Python. The underlying Gaudi framework [157] is shared with the LHCb experiment. The Gaudi layer defines the basic classes used for event processing, and also provides a component known as the scheduler, which is responsible for optimal algorithm execution. In Run 2, the Gaudi scheduler was used by offline reconstruction, but the HLT used a custom implementation: “Trigger Steering”. The steering layer was needed to implement additional HLT functionality, to facilitate features such as RoI-based processing and the sequencing of trigger selection criteria to arrive at a decision for each event. The Trigger Steering layer fulfilled these requirements successfully during Runs 1 and 2, but with a significant development and maintenance overhead.

Taking into consideration the above, the key drivers of the design of the new AthenaMT framework were to reduce maintenance overhead and to make effective use of hardware. As already discussed, effective use of modern multi-core machines requires multi-threading. From the beginning, the new framework was designed to meet both offline and trigger requirements, eliminating the need for a custom trigger-specific layer. Data and control flow, as well as regional reconstruction, were designed to be part of the scheduler. AthenaMT makes it possible to implement three different kinds of parallelism: inter-event, where multiple events are processed in parallel; intra-event, where multiple algorithms can run in parallel for an event and in-algorithm, where algorithms can utilise multi-threading and vectorisation. A pictorial representation of the different options is presented in Figure 7.35.

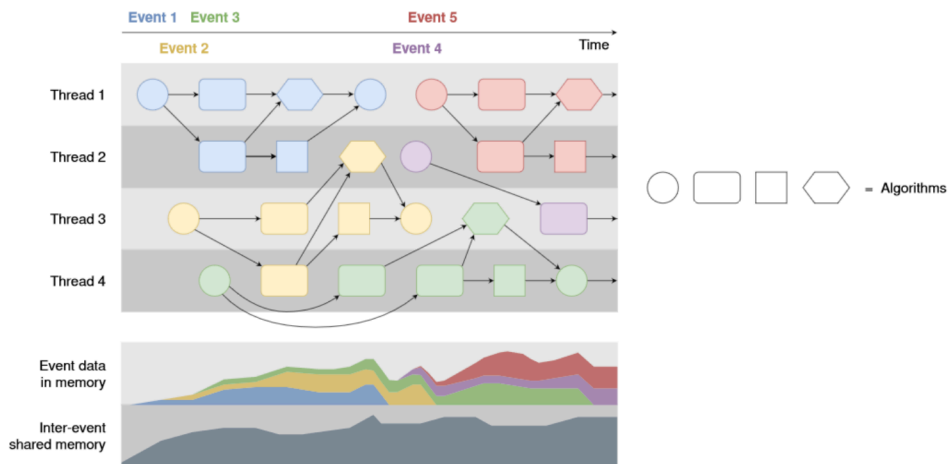


Figure 7.35: Diagram demonstrating the three different kinds of parallelism available with AthenaMT: inter-event, where multiple events are processed in parallel; intra-event, where multiple algorithms can run in parallel for an event and in-algorithm, where algorithms can utilise multi-threading and vectorisation.

As described in Section 1.4, the trigger menu translates the physics priorities of the experiment into allocations of the total L1 and HLT rates. Each unique selection combination is called a chain, with the set of chains making up the menu. During Run 2, more than 2000 unique combinations of trigger selection criteria were employed to distinguish interesting events. The configuration for each selection was generated from a large Python “TriggerMenu” package. As it is important to know exactly which menu is used both for data-taking and production of simulated samples, the configuration is stored in a database. As part of the software upgrade for Run 3, both the actual configuration and its database storage mechanism have been redesigned. The key requirement was to reduce the complexity of the configuration and to ease maintenance and development overhead while also improving the performance of the configuration serialisation/deserialisation. The reduced complexity is the result of a redesign of the configuration (both in the trigger and across Athena as a whole), while the improved serialisation/deserialisation performance comes from storing the configuration in a JSON [158] representation rather than in a relational format.

The AthenaMT software framework must interact with the online trigger and data-acquisition infrastructure, as well as new hardware triggers installed for Run 3. During Run 2, each node in the HLT forked multiple sub-processes, thus sharing memory between them via the copy-on-write mechanism. In Run 3, each forked sub-process will additionally run with multiple threads, and potentially process multiple events across these threads. This design implies that tuning of the number of forks, threads and concurrent events will be necessary to ensure maximum performance. It also implies that if a sub-process crashes or takes too long, all concurrently processed events must be saved for offline reprocessing. A comparison of the Run 2 and Run 3 architectures is presented in Figure 7.36.

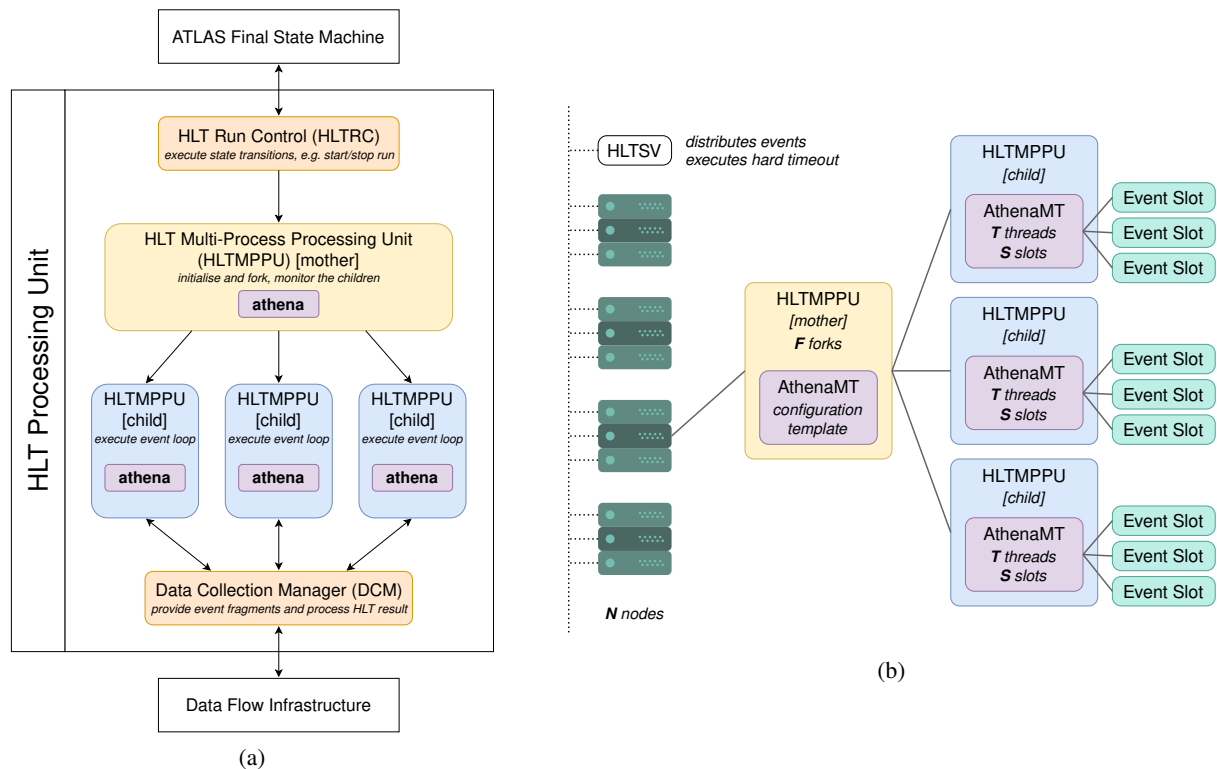


Figure 7.36: (a) Athena process operating within the wider HLT processing framework in Run 2, where multiple forked instances performed all processing and (b) Run 3 equivalent, potentially with fewer forks, with each fork capable of spawning multiple processing threads²¹.

7.6.4 Online Software & Monitoring

The ATLAS Control & Configuration software suite underwent a significant overhaul for the start of Run 2. The features introduced have since been refined throughout data taking; thus a modern and flexible system was in place at the the start of Run 3. The goals of the upgrade were three-fold: first, to properly accommodate additional requirements that could not be seamlessly included during data taking; to re-factor software that had been repeatedly modified to include new features, and make it more maintainable; and to take the opportunity to modernise the software base, making the most of the rapid evolution in information technology during the lifetime of the LHC. This upgrade was carried out retaining the important constraint to minimally impact the operational mode of the system or any public Application Programming Interfaces (APIs), thus making it as easy as possible for the large user community to integrate all the changes. An example of the changes is the introduction of a completely redesigned online monitoring archiving engine (P-BEAST) able to aggregate and serve data to a number of clients (principally a Grafana web interface for the production of time series displays). P-BEAST interfaces with the existing statistics distribution system (IS) without API changes in the latter, but provides enormous improvements in user experience and monitoring capability. An example of a typical online monitoring time series display with P-BEAST and Grafana is shown in Figure 7.37.



Figure 7.37: Example dashboard produced with the Grafana web interface.

Alongside the tools already mentioned, ATLAS also deploys specific Data Quality monitoring tools to aid shifters in the task of monitoring the performance of a data-taking session. The Data Quality Monitoring Framework automates data quality assessment by applying data quality algorithms to the histograms produced by event selection applications running in the HLT. The results produced by these algorithms are displayed to the shift crew in a hierarchical manner, so that the shift operator can quickly spot problems and easily identify their origin. This framework was further enhanced throughout Run 2 with the addition the Data Quality Monitoring Archiver, a new feature to store results in ROOT format for offline analysis. An example of a web display produced by the tool's integrated archive viewer is shown in Figure 7.38.

²¹ Image includes an icon made by Smashicons from flaticon.com.

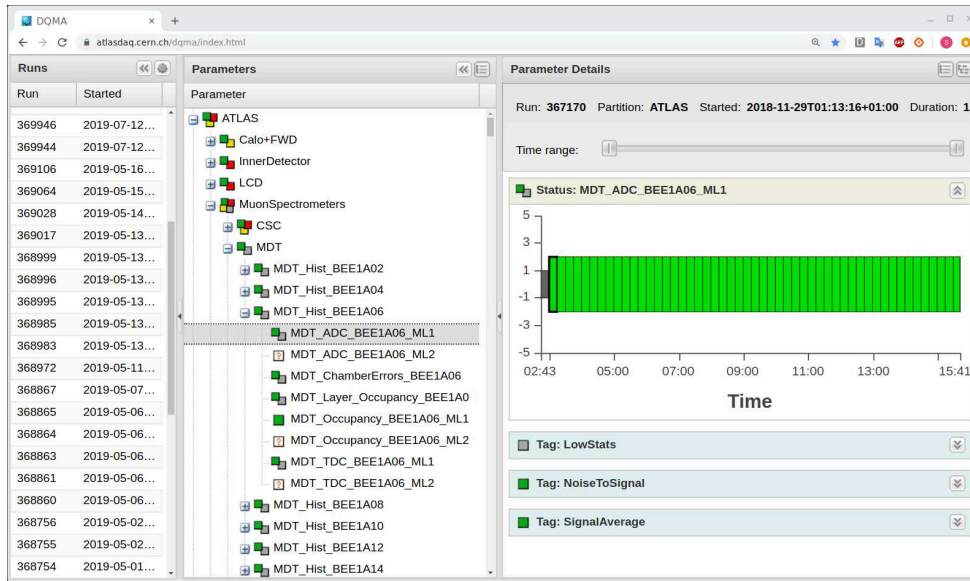


Figure 7.38: Example screenshot of the Data Quality Monitoring Application.

7.7 Detector Control System

The purpose of the DCS [159] is to ensure coherent and safe operation of ATLAS and to serve as a homogeneous interface to all subdetectors and to the technical infrastructure of the experiment. The DCS must be able to bring the detector into any desired operational state, to continuously monitor and archive the operational parameters, to signal any abnormal behaviour to the operator, and to allow manual or automatic actions to be taken. It must also provide bi-directional communication between DCS and run control to synchronise the state of the detector with the operation of the data acquisition system. The DCS also handles interactions between ATLAS sub-detectors and other systems that are controlled externally, such as the LHC accelerator, CERN technical services, the ATLAS magnets and the DSS.

Apart from consolidation and modernisation efforts for the DCS hardware and software components, several conceptual upgrades have been implemented since Run 1. First, the standard middleware employed by all systems was migrated from OPC Data Access to the Open Platform Communications Unified Architecture (OPC UA) standard [160]. Middleware applications interface front-end components with the DCS back-end software applications based on the WinCC OA SCADA package [45] (formerly known as PVSS). Second, for the new detector systems introduced for Run 3, DCS functions of most of the new or upgraded front-end electronics components are covered by the GBT-SCA ASIC[87, 88] which communicates with the back-end electronics using the same optical links as the readout or trigger data processed by the FELIX system. These monitoring and control data are relayed to and from the DCS back-end by the GBT-SCA software suite, a dedicated set of middleware applications. Finally, a significant part of the new back-end electronics of the upgrade systems is based on crates and blades following the ATCA standard for which a dedicated integration solution was conceived. The ATCA equipment significantly extends the DCS functionality compared to the previous implementation based on VME crates. These three aspects are described in more detail in the following three sub-sections.

7.7.1 Standard Middleware OPC UA

The middleware standard for DCS systems is Open Platform Communications Unified Architecture (OPC UA), an industry standard for machine-to-machine communication in the controls domain, allowing independence from the operating system and development in various programming environments. Features like robust data modelling, custom hardware embedding and secure communications are among the advantages of this standard.

In order to reduce development and maintenance efforts, a framework for OPC UA server creation is available – the Quick OPC UA Server Generation Framework (quasar) [161, 162]. Development starts with creation of a design file, in EXTensible Markup Language (XML) format, describing an object-oriented information model of the target system or device. Using this model, the framework generates an executable OPC UA server application, which exposes the per-design OPC UA address space, without the developer writing a single line of code. Furthermore, the framework generates skeleton code into which the developer adds the required target device/system integration logic. This approach allows both developers unfamiliar with the OPC UA standard, and advanced OPC UA developers, to create servers for the systems they are experts in while greatly reducing design and development effort as compared to developments based purely on COTS OPC UA toolkits. Higher level software may further benefit from the explicit device model by using the XML design description as the basis for generating client connectivity configuration and server data representation. Moreover, having the XML design description at hand facilitates automatic generation of validation tools.

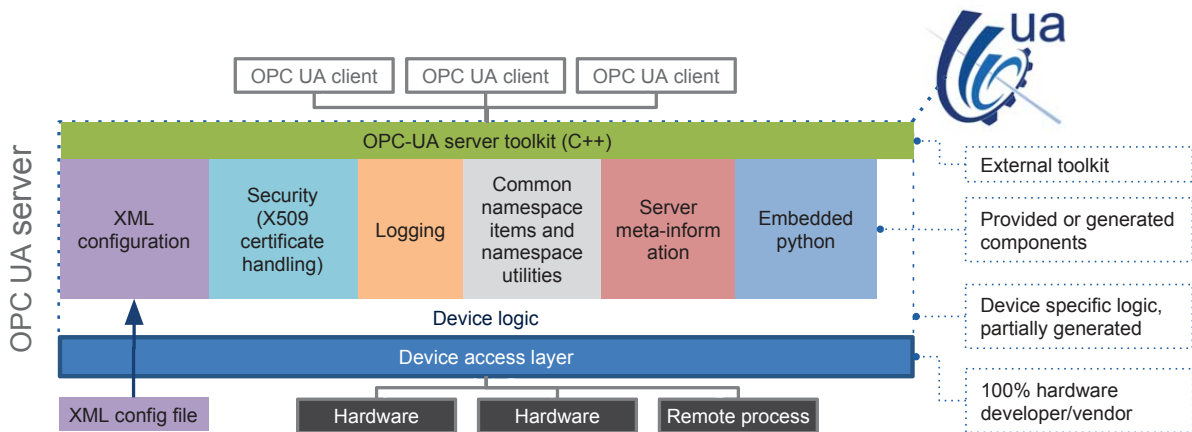


Figure 7.39: OPC UA server architecture within the quasar framework.

Figure 7.39 gives an overview of the different layers of quasar put into context. Controllable devices or systems are accessed using their specific access layer – often provided together with the specific device. The device logic layer functions as an interface with the high level layers of quasar, which comes in several modules covering different functional aspects. The address space module sits in the OPC UA end of the server, exposing data towards OPC UA clients, and is implemented using an OPC UA back-end layer with exchangeable back-end implementations. A configuration module facilitates address space and device instantiation and the definition of their relations. XML is used as the configuration format backed by XML schema definitions. A XML schema to C++ mapping generator is used to build actual instances from configuration files. An additional subsystem called Calculated Items, operating entirely in the address

space, enables creation of new variables which are derived from existing ones using mathematical functions. quasar further comes with optional modules such as component based logging, certificate handling, server meta-data, embedded Python processing, WinCC OA integration tools and Structured Query Language (SQL)/Not Only SQL (SQL) archiving with historic data access. A ready-to-use build system based on CMake along with pre-configured tool-chains for several platforms such as x86_64 or ARM-based Linux and Microsoft Windows are provided. Finally, an OPC UA client generation facility called UaoForQuasar is available for building C++ clients for quasar-based servers.

By the start of Run 3, all middleware applications for Java Card OpenPlatform-supported devices such as power supplies and the ELMB [44], as well as their WinCC OA integration components were migrated to the OPC UA standard. Furthermore, at the time of writing, more than 20 different OPC UA server implementations for numerous custom device types used in ATLAS have been developed and integrated, including servers running on embedded platforms such as system-on-chip devices.

7.7.2 Controls Software for GBT-SCA-based Front-End Electronics

The GBT-SCA is a radiation-tolerant ASIC and part of a chip-set of the GBT project, providing simultaneous transfer of readout data, timing and trigger signals as well as slow control and monitoring data, by multiplexing multiple logical electrical data links, onto a single optical link using the rad-tolerant GBTx ASIC on the front-end side. The GBT-SCA serves as an interface to the control and monitoring signals of front-end electronics on the detectors, using two redundant E-LINKs to connect to a GBTx.

The GBT-SCA software suite [163] provides a high level of abstraction and an interface to all communication channels of a GBT-SCA, profiting from the hardware parallelism between independent channels. To ensure reliability, the software does the necessary bookkeeping for the synchronous communication and transaction tracking.

Moreover, the software achieves high performance and low latency, including features such as grouping requests for lengthy operations requiring transfers of large amounts of data over JTAG, such as FPGA programming. Since thousands of GBT-SCAs are used in the detector systems, scalability is an important design aspect. At the same time, monitoring and control tasks require availability close to 100%, implying the need for a high level of robustness. For the final production system, the GBT-SCA software is interfaced with the optical link receiver system, FELIX, via a dedicated communication link called netIO. Figure 7.40 shows an overview of the GBT-SCA integration chain, illustrating the interplay of the components of the GBT-SCA software suite.

7.7.2.1 GBT-SCA OPC UA Ecosystem In the GBT-SCA software package core there is a library that is structured in modules that implement the required functionality in various layers. The library was designed to be flexible and easily adaptable to the diverse systems intended to use it by its polymorphic high-level data link control back-end. Moreover, the GBT-SCA software package contains the Demonstrators, which are tools that directly use the library and are used for testing and for low-level diagnostics. Finally, as part of the package, a GBT-SCA Simulator was developed that is able to generate GBT-SCA traffic, simulating realistic GBT-SCA behaviour, in order to allow for development and testing without real hardware.

The GBT-SCA OPC UA server was implemented using the quasar framework, taking advantage of quasar's built-in features such as calculated variables, threading, different types of variables and methods. The server architecture divides the GBT-SCA channels into device classes, according to their respective

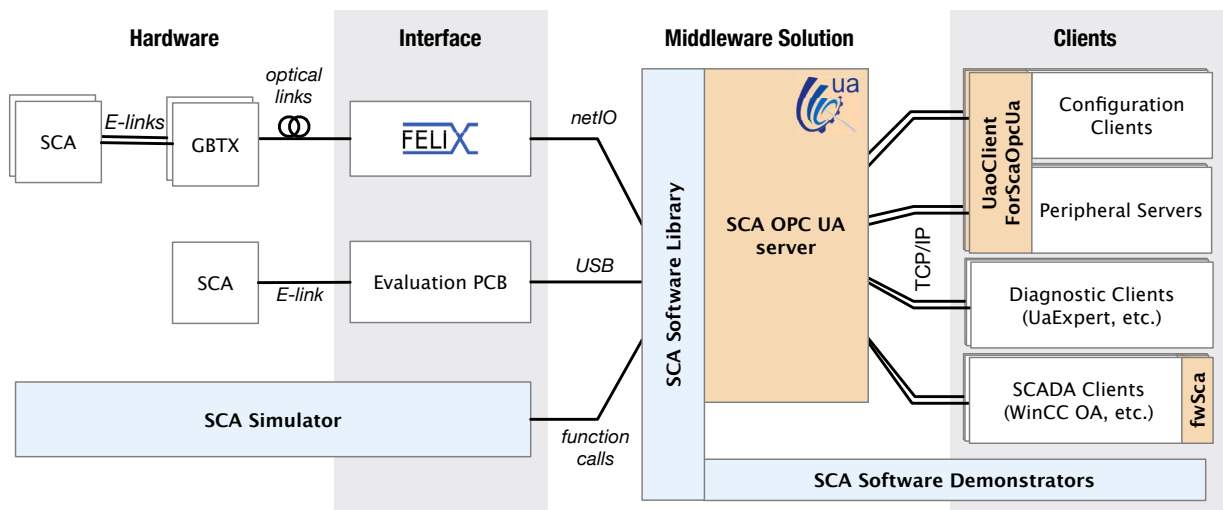


Figure 7.40: Global picture of the software suite. The GBT-SCA Software package, in light blue, comprises the GBT-SCA Software API to communicate with the GBT-SCA via different back-ends, the GBT-SCA Simulator to emulate GBT-SCA traffic for testing and development, and the Demonstrator tools which are used for standalone operations. The GBT-SCA OPC UA server and its ecosystem, in orange, is the middleware of choice to exchange data with the front-ends. UaoClientForScaOpcUa is a library that clients use to communicate with the GBT-SCA server. Finally, the fwSca module automates the integration of the server data into SCADA systems.

hardware functions. In addition, a Global Statistician module was developed to collect and measure general statistics across the setup and to expose the collected metrics to the clients. Finally, a GBT-SCA Supervisor software module oversees the state of the system and provides supervisory functionality such as automatic recovery from communication loss with GBT-SCAs, GBT-SCA ID validation and other administrative tasks.

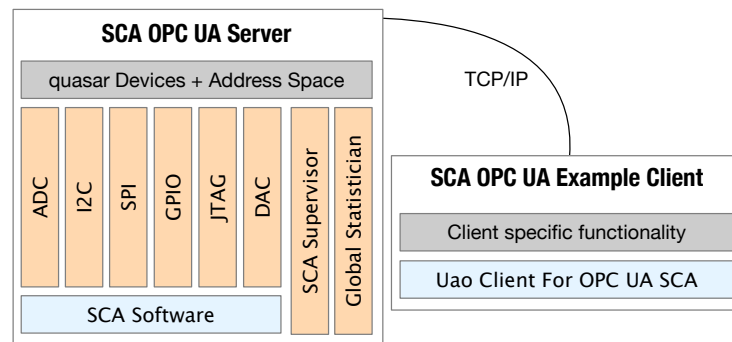


Figure 7.41: GBT-SCA OPC UA stack.

For any sub-system application using the GBT-SCA functionality – the GBT-SCA OPC UA Clients – the GBT-SCA OPC UA server is used as the hub to transmit and synchronise the data, and a choice was made to decentralise all specialised applications (as shown in Figure 7.40). This choice facilitates maintenance, has the advantage of dividing responsibilities among different communities, allows for staging by creating higher-level wrapper applications, and allows for interoperability between diverse clients. To support the concept, a quasar-generated C++ library, namely UaoClientForOpcUaSca is provided for building ad-hoc OPC UA clients. This library supplies the interface to the GBT-SCA OPC UA server and is created based on information sourced from the design of the server. Applications in ATLAS, that use the aforementioned

library, are TDAQ OPC UA clients used for configuration, or peripheral servers which perform sub-detector specific higher-level operations. The simplified architecture of the OPC UA GBT-SCA server and an example of a GBT-SCA OPC UA client is depicted in Figure 7.41. The server allows for the usage of any general-purpose OPC UA clients for diagnostic purposes such as the publicly available *UaExpert* tool [164]. Finally, the most common way that an OPC UA server is used by the DCS is through SCADA WinCC OA-based systems. These systems employ OPC UA clients which connect to the servers to retrieve the data and visualise the information in a user interface, usually deployed in the control room where the system is monitored by a shifter. In the case of WinCC OA, the OPC UA connectivity is realised via a software module, *fwSca*, which is supplied by the GBT-SCA software suite. This module allows for fast integration as it creates all the necessary configuration in the WinCC OA application based on a priori information of the GBT-SCA OPC UA information schema.

7.7.2.2 GBT-SCA Software Performance The server has been designed to serve setups of various sizes and types. The biggest challenge is the NSW detector upgrade (see Section 5.2). In this system, 6976 GBT-SCAs are employed, distributed over different types of front-end electronics boards. The traffic of the GBT-SCAs is handled by 30 FELIX hosts with 18 optical fibre connections each, and a corresponding number of GBT-SCA OPC UA servers.

In an early integration setup, the GBT-SCA OPC UA server was tested against a full-sector slice of the NSW Micromegas sub-system with eight detector layers fully equipped with their front-end electronics. The slice serves 160 GBT-SCAs, handled by a single server which was used in various realistic scenarios. The GBT-SCAs are separated in three categories/types of electronics with different functionality and interfaces as described in Table 7.3.

The setup used one FELIX host equipped with an Intel(R) Xeon(R) CPU E5-1650 v4 @3.60 GHz. The GBT-SCA OPC UA server runs in the FELIX host machine along with FELIX software. In a first constant-throughput scenario, a WinCC OA SCADA application monitored the analogue inputs from a separate host while three OPC UA clients were used for diagnostics. In the second burst traffic scenario, the server was used by 128 additional configuration clients.

Table 7.3: GBT-SCA channel usage in the ATLAS NSW Micromegas full-sector slice. The setup was used to evaluate the performance of the server.

Board Name	MMFE8	ADDC	L1DDC
Functionality	readout	trigger aggregator	data aggregator
GBT-SCA Numbers	128	16	16
ADC Inputs	15	10	9
Calculated variables	15	10	9
I ² C Master	2	6	2
I ² C Slave	44+60	6	2
Serial Peripheral Interface (SPI) Slave	8	-	-
General Purpose I/O (I/O)	19	18	-

Constant-throughput Monitoring Traffic: Even when no configuration activities are performed, the server is used constantly to provide monitoring data from the detector electronics. These data, mostly from

analogue inputs, correspond to the minimum possible activity of the server. In the Micromegas full-sector slice setup the global request rate was measured to be around 7800 requests/s for 2192 ADC inputs (each analogue input consists of two GBT-SCA requests). That resulted in an actual refresh rate of about 2 Hz per analogue input. The CPU usage of the server reached about 25% on average and the share of available physical memory used was 340 MB, a metric that is stable and not dependent on the usage.

On-demand GBT-SCA Traffic - Front-end Configuration: The most challenging aspect, in terms of open sessions and process complexity, is the configuration of the front-end boards. Emulating the cold start of the NSW Micromegas detector, a full-sector configuration was attempted in addition to the constant-throughput monitoring traffic as described above. During the process, up to 58 concurrent sessions were established from various OPC UA clients. The configuration clients were programming the front-end electronics using a combination of interleaving operations between GPIO, I²C, and SPI totalling around 2700 requests for each GBT-SCA. The global request rate reached about 35 000 requests/s. The instantaneous CPU usage peaked at 218%. The total time to initialise all the front-ends was measured to be 10 s. In ATLAS, each sector is independent, so these times and rates are also applicable to the full detector.

7.7.3 Controls for ATCA Back-End Electronics

The ATCA standard is employed as a back-end platform by many of the system upgrades for Run 3 (the NSW Trigger Processor, the LAr LATOME, and the TDAQ L1 gFEX, eFEX, MUCTPI, jFEX, and L1Topo systems), replacing the VME standard as the preferred back-end technology for new electronics systems. To allow the DCS to work with ATCA, an integration solution based on the common OPC UA toolset was developed, managing ATCA shelves via their shelf manager Simple Network Management Protocol (SNMP) interface and providing control and monitoring of shelf and blade functions. The solution covers the “CERN-standard” Pigeon Point Shelf Managers (ShMM 500, ShMM 700R) and, while designed for ATCA, is compatible with the broader device family. An overview of this scheme is illustrated in Figure 7.42.

The ATCA OPC UA server, a modularised software application (see also 7.7.1), models selected parts of the ATCA standard functionality for compliant devices in an object-oriented design and interfaces with the shelf manager SNMP agent via the experiment controls network. Code generation techniques are used to implement the selected device functions based on the SNMP Management Information Base provided by the shelf manufacturer. The SNMP back-end within the server implementation is a C++ wrapper of the Net-SNMP open-source library and provides a generic interface to any SNMP device. In addition, the ATCA OPC UA server provides features such as automatic hardware discovery that queries a given set of shelf managers and automatically creates a map of all blades, fan trays, power supplies and other field-replaceable units, and makes their functionality available to OPC UA clients. The supervision of custom blade functions such as on-board sensors controlled by the blade IPMC is also supported.

Finally, a set of associated tools, allowing for easy client integration with the WinCC OADCS applications, are part of the ATCA software solution. These tools take care of datapoint creation, alarm handling, archiving, and Finite State Machine integration, and facilitate the creation of operator interfaces with pre-built generic user interface panels, making shelf integration into the ATLAS DCS an easy and efficient task.

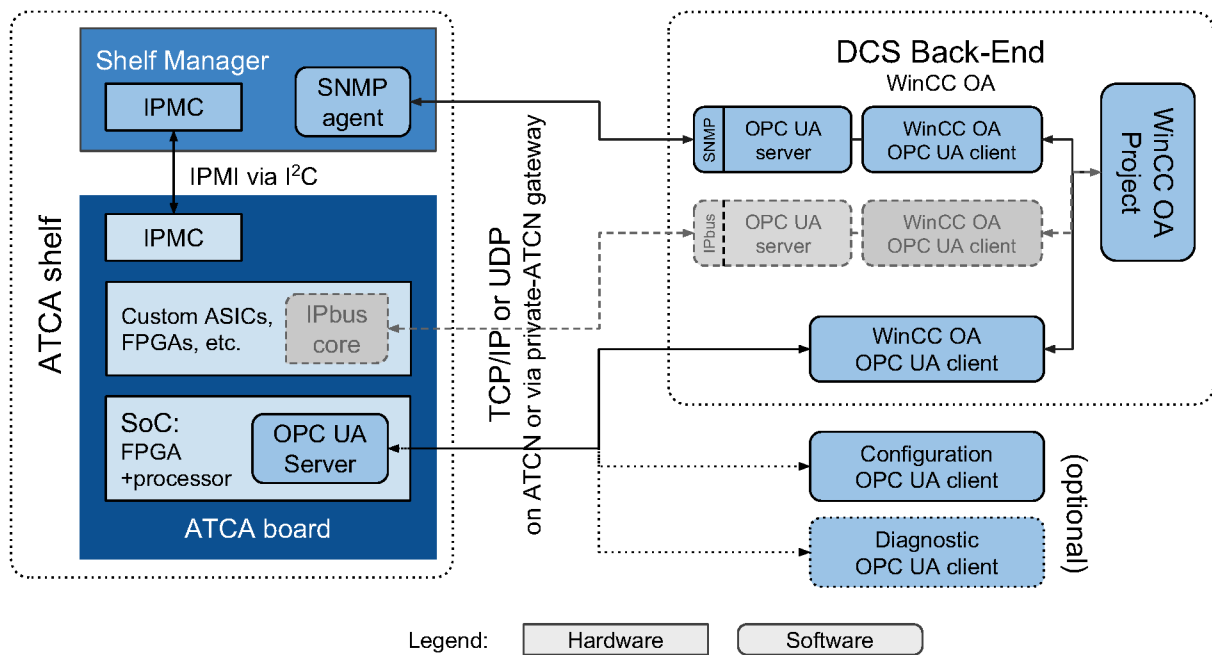


Figure 7.42: Overview of the scheme based on OPC UA to manage ATCA shelves via their shelf manager interface. During Run 3, some systems still use the IPbus-based pathway shown in grey, but this is deprecated and will be phased out by the end of LS3.

The LAr Calorimeter back-end electronics (see also Section 4.1.2) and the TDAQ Level-1 Calorimeter Trigger electronics (see also Section 7.2) use this method to monitor onboard component parameters such as temperature and voltage.

8 Outlook

Extensive upgrades were performed to ready the ATLAS detector for Run 3 based on the experience gained from Runs 1 and 2. While the running parameters foreseen for Run 3 are not expected to subject the detector to more extreme conditions than it has experienced already, it is expected that with luminosity levelling, ATLAS will perpetually experience conditions close to the extremes of Runs 1 and 2 during Run 3. The emphasis of the Phase-I upgrades was therefore on making the detector and its trigger as robust as possible, so that ATLAS can run comfortably during Run 3 with un-prescaled single-lepton trigger thresholds comparable to those of Run 1. The Phase-I upgrades were designed to last for the remaining lifetime of ATLAS, and constitute the first step in preparing ATLAS for the rigours of data-taking at the HL-LHC, when it is expected that the instantaneous luminosity could rise as high as $\mathcal{L} = 7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, with an average pileup of $\langle \mu \rangle \approx 200$.

The remaining upgrades required for ATLAS to run at the HL-LHC are being prepared now, and will be installed during LS3, at the end of Run 3. The largest of these is the complete replacement of the ATLAS ID by the new, all-silicon, Inner Tracker (ITk) [165, 166] that is designed for a similar or even improved tracking performance compared to the current ATLAS ID, but in the challenging pileup conditions of the HL-LHC. The ITk pixel system, which builds on the technology used in the IBL, will be composed of five barrel layers plus endcap rings covering up to $|\eta| = 4$. It will be surrounded by a large-area strip detector composed of four barrel layers and six endcap discs, extending up to a maximum radius of 1 m and covering up to $|\eta| < 2.7$. There will be substantial upgrades to the electronics of numerous subsystems for Phase-II, with a focus on radiation-hard readout electronics for the LAr [167] and Tile [144] Calorimeters that must operate at the trigger rates and latencies required for Phase-II luminosities. The already excellent timing resolution of the calorimeters will be augmented with the addition of a dedicated High-Granularity Timing Detector (HGTD) [168]. The Muon Spectrometer [62] will see further improvements to trigger coverage and redundancy, with the replacement of on- and off-detector readout and trigger electronics, the replacement of the TGC EIL4 doublet modules by more efficient triplets, and with the addition of RPCs to the entire inner barrel layer, building on the BIS78 pilot project described in Section 5.3.4. Information from the MDT detectors will be incorporated in the first level of the Muon trigger, increasing its granularity and improving the sharpness of the trigger transverse momentum thresholds. The replacement of the front-end electronics of the existing ATLAS detector systems during the Phase-I and Phase-II upgrades allows for a substantially higher hardware-level trigger rate and a longer latency. The TDAQ system for ATLAS in the HL-LHC [66, 169] will consist of a single-level hardware trigger that analyses calorimeter and muon detector information at 40 MHz within 10 μs latency. After the hardware-based trigger decision, the resulting full detector and trigger data will be read out at a rate of 1 MHz. An upgraded high-level trigger system will be implemented using commodity hardware, refining the trigger objects in order to achieve a maximum output rate of 10 kHz.

The Phase-I upgrade project is an essential part of a broad upgrade program through the lifetime of the Large Hadron Collider and is fully compatible with the future Phase-II upgrade program of the ATLAS experiment. The collection of $\sqrt{s} = 13.6 \text{ TeV}$ collisions in July 2022 is the realisation of the Run 3 configuration of the ATLAS detector (see Figure 8.1) and has marked the start of the vibrant physics programme planned for Run 3 and beyond.




Figure 8.1: Event display (Run 427514, Event 68319093) of a collision event recorded in ATLAS on 7 July 2022, when stable beams of 6.8 TeV protons were delivered to ATLAS by the LHC. The red line shows a muon candidate with a transverse momentum of 15 GeV reconstructed using information from the inner tracking detectors and the ATLAS Muon Spectrometer (MS) endcap. The muon candidate was among the first reconstructed using hits in the Micromegas chambers of the NSW on side C. The NSW is outlined in white and the Micromegas hits are shown as orange lines. Additional muon chambers associated with the track are shown as green (MDT endcap) and purple (TGC endcap) boxes.

Acknowledgements

We thank CERN for the very successful operation of the LHC, as well as the support staff from our institutions without whom ATLAS could not be operated efficiently.

We acknowledge the support of ANPCyT, Argentina; YerPhI, Armenia; ARC, Australia; BMFWF and FWF, Austria; ANAS, Azerbaijan; CNPq and FAPESP, Brazil; NSERC, NRC and CFI, Canada; CERN; ANID, Chile; CAS, MOST and NSFC, China; Minciencias, Colombia; MEYS CR, Czech Republic; DNRF and DNSRC, Denmark; IN2P3-CNRS and CEA-DRF/IRFU, France; SRNSFG, Georgia; BMBF, HGF and MPG, Germany; GSRI, Greece; RGC and Hong Kong SAR, China; ISF and Benoziyo Center, Israel; INFN, Italy; MEXT and JSPS, Japan; CNRST, Morocco; NWO, Netherlands; RCN, Norway; MEiN, Poland; FCT, Portugal; MNE/IFA, Romania; MESTD, Serbia; MSSR, Slovakia; ARRS and MIZŠ, Slovenia; DSI/NRF, South Africa; MICINN, Spain; SRC and Wallenberg Foundation, Sweden; SERI, SNSF and Cantons of Bern and Geneva, Switzerland; MOST, Taiwan; TENMAK, Türkiye; STFC, United Kingdom; DOE and NSF, United States of America. In addition, individual groups and members have received support from BCKDF, CANARIE, Compute Canada and CRC, Canada; PRIMUS 21/SCI/017 and UNCE SCI/013, Czech Republic; COST, ERC, ERDF, Horizon 2020 and Marie Skłodowska-Curie Actions, European Union; Investissements d’Avenir Labex, Investissements d’Avenir Idex and ANR, France; DFG and AvH Foundation, Germany; Herakleitos, Thales and Aristeia programmes co-financed by EU-ESF and the Greek NSRF, Greece; BSF-NSF and MINERVA, Israel; Norwegian Financial Mechanism 2014-2021, Norway; NCN and NAWA, Poland; La Caixa Banking Foundation, CERCA Programme Generalitat de Catalunya and PROMETEO and GenT Programmes Generalitat Valenciana, Spain; Göran Gustafssons Stiftelse, Sweden; The Royal Society and Leverhulme Trust, United Kingdom.

The crucial computing support from all WLCG partners is acknowledged gratefully, in particular from CERN, the ATLAS Tier-1 facilities at TRIUMF (Canada), NDGF (Denmark, Norway, Sweden), CC-IN2P3 (France), KIT/GridKA (Germany), INFN-CNAF (Italy), NL-T1 (Netherlands), PIC (Spain), ASGC (Taiwan), RAL (UK) and BNL (USA), the Tier-2 facilities worldwide and large non-WLCG resource providers. Major contributors of computing resources are listed in Ref. [170].

Glossary

ADC Analog-to-Digital Converter	71, 75, 76, 77, 102, 104, 109, 133, 154, 185, 186
ADDC ART Data Driver Card	97, 103, 104, 105, 106, 108, 111, 154, 185
AFP ATLAS Forward Proton	4, 10, 19, 118, 127, 128, 129, 130, 132, 199
ALFA Absolute Luminosity for ATLAS	4, 10, 19, 118, 127, 128, 165, 166, 168
ALICE A Large Ion Collider Experiment at LHC	119
ALTI ATLAS Local Trigger Interface	73, 170
AMC Advanced Mezzanine Card	76
API Application Programming Interface	180, 184
ARM Advanced RISC Machines	165

ART Address in Real Time	103, 104, 105, 106, 108, 154, 190
ASD Amplifier-Shaper-Discriminator	103, 117
ASDBLR Amplifier, Shaper, Discriminator and BaseLine Restorer	55
ASIC Application-Specific Integrated Circuits .53, 56, 75, 101, 102, 103, 104, 105, 106, 109, 111, 117, 129, 176, 181, 183, 200	
ATCA Advanced Telecommunications Architecture 5, 75, 76, 77, 130, 140, 141, 143, 144, 150, 154, 157, 161, 165, 181, 186, 187	
ATCN ATLAS Control Network	177
AthenaMT multithreaded ATLAS software framework	137, 178, 179
AWG American Wire Gauge	46, 47
BC Bunch Crossing	43, 53, 55, 65, 88, 104, 108, 123, 125, 126, 155, 157, 164
BCAM Brandeis CCD Angle Monitor	100, 101, 102
BCID Bunch Crossing IDentifier	105, 107, 108, 109, 125, 152, 153, 154, 157, 167
BCM Beam Conditions Monitor	119
BiCMOS Bipolar CMOS	117
BIS78 Integrated Muon BIS7 and BIS8 sMDT chambers supplemented with new RPCs; see also Section 5.3.4	10, 18, 85, 87, 88, 102, 113, 114, 115, 116, 117, 161, 188
BJT Bipolar Junction Transistor	116
BOC Back of Crate	45, 46, 52, 53, 54
BPM Bi-Phase Mark	45
BRAN Beam Rate of Neutrals relative luminosity monitors	131, 132
BSM Beyond the Standard Model	21, 131
CALbus custom bus used for sub-detector calibration requests in the CTP	169
CAN Controller Area Network	49
CMM Cluster Merger Module	138
CMOS Complementary metal-oxide-semiconductor	43
CMS Compact Muon Solenoid experiment at LHC	7, 119, 127
CMX Common Merger Extended Module	138, 161
COMbus+ custom TTC bus used in the CTP	168, 169
COTS Commercial Off-The-Shelf	182
CP Cluster Processor	136, 138, 148

CPU Central Processing Unit	137, 168, 171, 176, 178
CSC Cathode Strip Chamber	10, 23, 86, 109, 110
CSM Chamber Service Module	117
CTP Central Trigger Processor 73, 109, 123, 127, 130, 133, 135, 136, 138, 140, 161, 162, 163, 164, 165, 166, 167, 168, 169, 170, 191, 192, 196	
CTPCAL CTP module that receives calibration requests from subdetectors	168
CTPCORE+ Upgraded CTP CORE Processor	166, 167, 168, 169, 197
CTPIN CTP Input Interface	127, 166, 167, 168, 169, 192, 197
CTPMI CTP Machine Interface	167, 169
CTPMON CTP module that monitors per-bunch input signals from the CTPIN	168, 169
CTPOUT+ CTP output modules that interface to sub-detector TTC	168, 169, 170
DDR Double Data Rate	172
D-layer Outermost longitudinal layer of the TileCal calorimeter, with segmentation of $\Delta\eta \times \Delta\phi = 0.2 \times 0.1$	151
DAQ Data Acquisition System .2, 20, 21, 22, 45, 53, 54, 55, 65, 105, 130, 135, 137, 138, 146, 153, 163, 164, 170, 176	
DCS Detector Control System 46, 47, 48, 49, 50, 120, 130, 133, 146, 149, 150, 157, 163, 175, 177, 181, 182, 185, 186	
DMA Direct Memory Access	172, 174
DRIE Deep Reactive Ion Etching	42
DSP Digital Signal Processor	71
DSS Detector Safety System	49, 50, 181
DTMROC Drift Time Measurement Read-Out Chip	55
EASY Embedded Assembly SYstem from CAEN	111
ECC Error Correction Code	172
EI Muon detector inner wheels	151, 192
E-LINK Electrical chip-to-chip interconnect	102, 105, 106, 157, 183
eFEX electron Feature EXtractor	4, 139, 140, 141, 142, 143, 146, 149, 150, 163, 186
EIL4 4th TGC chamber of the Muon detector EI Large Sectors, counting radially; see also Chapter 5 84, 86, 87, 114, 151, 157, 188	
ELMB Embedded Local Monitor Board	49, 109, 127, 183
EM electromagnetic	10, 15, 16, 74, 138, 143, 145, 146, 149, 166

EM-TGC	TGC in the endcap middle layer, also called the Big Wheels	107, 108, 109, 151, 153, 154, 156, 157
EMB	LAr Electromagnetic Barrel Calorimeter	14, 70, 77
EMEC	LAr Electromagnetic Endcap Calorimeter	14, 70, 74, 119, 126
ENOB	Effective Number Of Bits	75
ESD	Electrostatic Discharge	117
FCal	LAr Forward Calorimeter	10, 13, 15, 16, 62, 70, 71, 74
FCal	Forward Calorimeter	119, 126
FEAST	DC-DC converter	102, 105, 106, 111
FEB	Front End Board	71, 73, 75, 104, 105, 106, 107, 109, 112
FEC	Front End Crate	71, 72, 73, 75
FELIX	Front-End LInk eXchange	5, 75, 77, 101, 102, 103, 105, 106, 107, 108, 109, 117, 137, 140, 141, 143, 146, 148, 150, 157, 161, 163, 170, 171, 172, 173, 174, 175, 176, 177, 181, 183, 185
FEX	L1Calo Feature EXtractor	76, 139, 140, 148, 149, 150
FOX	Fibre-Optic eXchange	140, 143, 146, 149
FPGA	Field Programmable Gate Array	56, 76, 106, 109, 117, 123, 129, 133, 139, 140, 141, 143, 144, 145, 146, 147, 148, 149, 150, 154, 156, 157, 161, 162, 163, 165, 170, 171, 172, 173, 174, 178, 183, 194, 195
FPGACH	FPGA for each channel of the ZDC on the LUCROD/ZDC	133
FPGAM	main FPGA for the LUCROD/ZDC	133
FPGAV	FPGA performing simultaneous trigger calculations on all 8 channels of the LUCROD/ZDC	133
FR4	glass-reinforced epoxy laminate	97
FSM	Finite State Machine	50
GbE	Gigabit Ethernet	76, 77, 138, 165, 172, 176
GBT	Gigabit Bidirectional Trigger and Data Link	172, 174, 183, 193
GBT-SCA	Gigabit Transceiver Slow Control Adapter	5, 75, 102, 103, 104, 105, 117, 181, 183, 184, 185, 186
GBTx	GBT Transmitter	102, 105, 106, 109, 117, 183
gFEX	global Feature EXtractor	4, 139, 146, 147, 148, 149, 150, 163, 186
GOL	Gigabit Optical Link	117
GPIO	General Purpose I/O	185, 186
GPS	Global Positioning System	167

GPU Graphics Processing Unit. Usable for general purpose calculations (also known as GPGPU) .	178
GTX MGT in the Xilinx FPGA architecture. GTX transceivers support up to 12 Gb/s.	157
HEC LAr Hadronic Endcap Calorimeter	10, 14, 70, 71, 74
HGTD High-Granularity Timing Detector	188
HL-LHC High Luminosity Large Hadron Collider	8, 11, 15, 22, 23, 82, 86, 87, 88, 89, 91, 104, 109, 112, 113, 114, 115, 121, 140, 150, 161, 171, 188
HLT High-Level Trigger	20, 21, 23, 135, 136, 137, 138, 140, 146, 159, 164, 167, 169, 171, 175, 176, 178, 179, 180
HLTPU High-Level Trigger Processing Unit	171
HOLA High-speed Optical Link for ATLAS	56
HPL high-pressure phenolic laminate, also known as Bakelite	113
HPTDC High Performance Time to Digital Converter	117
HT High Threshold	55
Hub Common readout infrastructure for L1Calo. The ROD is a daughter card.	141, 143, 150, 163, 198
HV High Voltage .	47, 48, 49, 65, 66, 67, 68, 92, 93, 94, 95, 98, 100, 111, 113, 120, 124, 125, 133, 134
I²C Inter-integrated Circuit	163, 185, 186
I/O Input/Output	138, 172, 174, 185, 193
IBL Insertable <i>B</i> -layer	2, 3, 10, 11, 13, 14, 15, 22, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 58, 60, 61, 62, 63, 64, 126, 128, 188
ICS Intermediate Conversion Stage for NSW LV	106, 111
ID Inner Detector	2, 3, 11, 13, 14, 15, 17, 23, 25, 28, 29, 30, 32, 33, 38, 39, 44, 46, 47, 50, 51, 52, 57, 58, 59, 60, 62, 66, 69, 77, 84, 89, 119, 126, 188, 194
IDEP ID End Plate	30
IP Interaction Point	7, 10, 11, 13, 19, 25, 26, 27, 30, 32, 41, 62, 78, 84, 86, 87, 88, 89, 90, 91, 92, 97, 107, 108, 109, 118, 119, 120, 126, 127, 131, 138, 152, 153, 154, 155, 156, 157
IPMB Intelligent Platform Management Bus	76, 150
IPMC Intelligent Platform Management Controller	76, 141, 144, 146, 150, 157, 163, 186
IPT Inner Positioning Tube	40, 41, 60, 61
IST Inner Support Tube	40, 60, 61
ITk Inner Tracker	188
JCOP Joint COntrols Project	50

JEP Jet/Energy Processor	136, 138, 148
jFEX jet Feature EXtractor	4, 139, 143, 144, 145, 146, 149, 150, 163, 186
JTAG Joint Test Action Group industry standard, typically used for programming and debugging FPGAs. first	172, 183
L0 Level-0 Trigger	89
L0A Level-0 Trigger Accept	104, 109
L1 Level-1 Trigger 20, 21, 22, 23, 45, 53, 55, 56, 65, 70, 72, 76, 83, 86, 87, 88, 89, 97, 98, 107, 113, 117, 123, 133, 135, 136, 137, 138, 139, 151, 153, 159, 164, 166, 167, 179, 186	
L1A Level-1 Trigger Accept 45, 55, 71, 76, 104, 105, 109, 133, 136, 137, 141, 143, 144, 146, 157, 163, 167, 169, 170	
L1Calo Level-1 Calorimeter Trigger . 10, 71, 72, 133, 135, 136, 138, 139, 140, 142, 143, 147, 148, 149, 150, 161, 162, 165, 169, 171, 194	
L1DDC Level-1 Data Driver Card for Micromegas and sTGC 96, 103, 104, 105, 106, 107, 109, 111, 112, 185	
L1ID Level-1 Trigger IDentifier	175
L1Muon Level-1 Muon Trigger 10, 87, 135, 136, 138, 151, 152, 153, 157, 158, 159, 160, 161, 162	
L1Topo Level 1 Topological Processor 127, 135, 136, 138, 140, 141, 143, 145, 147, 148, 149, 150, 161, 162, 163, 165, 168, 169, 186	
LAr Liquid Argon Calorimeter 10, 11, 12, 14, 15, 16, 18, 23, 70, 71, 72, 73, 74, 75, 76, 77, 80, 119, 126, 139, 140, 149, 171, 186, 187, 188	
LArC LAr Carrier	76, 77
LATOME LAr Trigger prOcessing MEzzanine. Mezzanine card for LAr Carrier Board. Together these form the LDPB.	76, 77, 186
LB Luminosity Block	125
LDPB LAr Digital Processing Blade	75, 76
LDPS LAr Digital Processing System	75, 76, 139, 146
LED light-emitting diode	122, 123
LHC Large Hadron Collider 7, 8, 10, 11, 13, 15, 19, 21, 25, 27, 29, 39, 43, 45, 49, 50, 53, 55, 57, 64, 65, 66, 70, 72, 88, 118, 119, 120, 121, 123, 125, 126, 127, 128, 129, 130, 131, 132, 136, 137, 150, 155, 167, 169, 176, 180, 181, 189	
LHCb Large Hadron Collider beauty experiment	119, 178
LINAC-4 Linear Accelerator at the start of the injector chain for the LHC	8, 11
LS1 Long Shutdown 1	7, 8, 13, 38, 52, 54, 57, 58, 112, 113, 138, 159, 167, 168

LS2 Long Shutdown 2	8, 10, 11, 52, 54, 71, 73, 76, 85, 86, 112, 114, 127, 128, 129, 167, 168, 170, 171, 172
LS3 Long Shutdown 3	8, 11, 29, 87, 110, 114, 116, 121, 172, 187, 188
LSB Layer Summing Board	75
LT Low Threshold	55
LTDB LAr Trigger Digitizer Board	73, 75, 76
LTP Local Trigger Processor	170
LTPI Local Trigger Processor Interface	170
LUCID LUMinosity Cherenkov Integrating Detector	4, 19, 36, 118, 119, 120, 121, 122, 123, 124, 125, 126, 133, 170
LUCROD LUCID ReadOut Driver	122, 123, 124, 133
LUMAT LUMinosity And Trigger	123, 124
LUT Lookup Table. An array that replaces runtime computation with a simpler array indexing operation.	108, 155
LV Low Voltage	47, 48, 49, 100, 107, 110, 111, 194
LVDB Low Voltage Distributor Board for NSW	103, 104, 105, 106, 107
LVDS Low-voltage Differential Signaling	45, 46, 105, 106, 107, 117, 148, 155, 168, 197
LVPS Low Voltage Power Supply	73
MBTS Minimum Bias Trigger Scintillators	77, 78, 80, 81, 82, 119
MC Monte-Carlo	60, 61
MCP-PMT Micro-Channel Plate PMTs	129, 130
MDT Monitored Drift Tube	10, 12, 17, 18, 20, 23, 84, 85, 86, 87, 88, 89, 94, 109, 111, 112, 113, 114, 115, 116, 117, 188, 189, 199
MGT Multi-Gigabit Transceiver	143, 165, 194
Micromegas MICRO MESH GAseous Structure	3, 10, 18, 20, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97, 99, 100, 101, 102, 103, 104, 105, 106, 107, 108, 109, 110, 111, 112, 151, 153, 154, 155, 156, 157, 185, 186, 189, 195, 196
MMFE8 Front-end Board for Micromegas	96, 97, 105, 106, 109, 185
MROD Muon ReadOut Driver	117
MS Muon Spectrometer	12, 13, 17, 23, 26, 32, 47, 84, 85, 86, 88, 89, 101, 111, 112, 113, 114, 120, 189, 197
MTP Multi-fiber Termination Push-on	149, 172
MUCTPI Muon-CTP Interface	135, 149, 159, 164, 165, 166, 186

NEG Non-Evaporative Getter	41
new RPC new version of RPC with a thinner gas gap and improved electronics .88, 112, 113, 114, 115, 116, 117, 191, 198	
new JD shielding disc supporting each NSW	26, 89, 90, 91, 92
NGPS New Generation Power Supply from CAEN	106
NIM Nuclear Instrumentation Module	133, 168, 197
NIM2LVDS NIM standard to LVDS conversion	168, 170
NMOS n-type metal-oxide-semiconductor	43
NoSQL Not Only SQL	183
nSQP new Service Quarter Panel	39, 52, 58
NSW New Small Wheel 12, 18, 19, 23, 26, 36, 85, 86, 87, 88, 89, 90, 91, 92, 93, 94, 95, 100, 101, 102, 103, 104, 105, 106, 107, 108, 109, 110, 111, 113, 115, 116, 117, 135, 151, 152, 153, 154, 155, 156, 157, 159, 171, 185, 186, 189, 194, 196, 197, 200	
NSW-TP NSW Trigger Processor ...95, 98, 101, 103, 104, 105, 106, 107, 108, 109, 153, 154, 156, 157	
NTC Negative Temperature Coefficient	43, 48, 109
OF Optimal Filtering	71, 76
OFC Optimal Filtering Coefficient	72
OPC Open Platform Communications	49, 181
OPC UA Open Platform Communications Unified Architecture 5, 102, 181, 182, 183, 184, 185, 186, 187	
original EI wheel original MS EI wheel from Runs 1 and 2 .86, 87, 88, 89, 91, 100, 109, 110, 111, 114	
PCB Printed Circuit Board	46, 94, 96, 97, 116
PCIe Peripheral Component Interconnect Express	172, 174
PEEK PolyEther Ether Ketone	57
pFEB Front-end Board for sTGC pad and wire signals	104, 105, 106, 107, 109
PID Particle Identification	57
PIN a diode with an undoped intrinsic semiconductor region between p-type semiconductor n-type semiconductor regions	28, 45, 47
PITbus custom Pattern In Time bus used to route synchronised and aligned trigger inputs from the CTPIN boards to the CTPCORE+ module	168, 169
Pixel Silicon Pixel Detector	11, 13, 14, 15, 126
PLC Programmable Logic Controller	51
PMT Photo-Multiplier Tube 83, 119, 120, 121, 122, 123, 124, 125, 126, 129, 130, 131, 132, 133, 134, 196	

PP2 Patch Panel 2	47, 49
PPM L1Calo Pre-processor Module	132, 133, 138, 148, 149
PS Proton Synchrotron (serving in the injector chain for the LHC)	8, 10, 11
PST Pixel Support Tube	30, 61
PVC polyvinyl chloride	116
PXI PCI eXtensions for Instrumentation	128
QCD Quantum Chromodynamics	146
Quad-SFP Quad Small Form-factor Pluggable Transceiver Specification	45
quasar Quick OPC UA Server Generation Framework	182, 183, 184
RadFet Radiation-sensitive Field Effect Transistors	28
RAM Random Access Memory	172
RASNIK A 3-point optical displacement monitor with sub-nanometre precision consisting of a camera, a lens and a coded mask	96, 101
RDMA Remote Direct Memory Access	174, 176
RF radio frequency	8
ROB ReadOut Buffer	175
ROC ReadOut Controller	103, 104, 105, 109
ROD Readout Driver. Prior to Run 3, custom hardware Readout Drivers were used by detector subsystems. In Run 3, the SW ROD has taken on much of this functionality, however a newly designed daughter card for the Hub is also referred to as a ROD. 45, 46, 52, 53, 54, 55, 56, 71, 137, 141, 143, 148, 150, 163, 170, 171, 175, 177, 194, 198, 199	
RoI Region of Interest	86, 107, 108, 136, 138, 140, 153, 158, 159, 164
ROS ReadOut System	21, 45, 55, 56, 71, 133, 137, 138, 171, 175, 177, 199
RPC Resistive Plate Chamber . 8, 10, 12, 17, 18, 20, 84, 86, 88, 111, 112, 113, 114, 116, 117, 135, 159, 161, 170, 188, 197	
RPC-BIS78 new RPCs in the integrated Muon BIS7 and BIS8 chambers; see also Section 5.3.4	157, 171
RPD Reaction Plane Detector in the ATLAS Forward system	132
Rx Receiver	45, 46
SCADA Supervisory Control And Data Acquisition	49, 181, 184, 185
SC-OL Supply and Control for the Optical Link	47, 49
Sector Logic Muon Trigger Sector Logic	107, 109, 153, 154, 156, 157

SEM Soft Error Mitigation	117
SEU Single Event Upset	117
sFEB Front-end Board for sTGC strip signals	104, 105, 106, 107, 108, 109
SFO Sub-Farm Output	137, 138, 177
SHV Safe High Voltage	111, 133
SiT Silicon Tracker for AFP based on 3-D pixels	128, 129, 130
S-LINK fast optical chip-to-chip interconnect	45, 53, 54, 55, 56, 123, 133
SCT Semiconductor Tracker	3, 10, 11, 13, 14, 15, 28, 30, 32, 34, 38, 39, 53, 60, 61, 62, 65, 68, 126
SEE Single Event Effects	27, 28
sMDT Small-diameter MDT	88, 112, 113, 114, 115, 116, 117, 191
SNMP Simple Network Management Protocol	186
SoC System-on-Chip	146, 157, 165
SPI Serial Peripheral Interface	185, 186
SPS Super Proton Synchrotron (serving as the injector for the LHC)	11, 97, 132
SQL Structured Query Language	183, 197
sTGC small-strip TGC	10, 18, 20, 88, 89, 90, 91, 97, 98, 99, 100, 101, 102, 103, 104, 105, 106, 107, 108, 109, 110, 111, 112, 151, 153, 154, 155, 156, 157, 195
Super Cell Intermediate sums of four or eight calorimeter cells for use in the trigger	10, 70, 73, 74, 75, 76, 77, 139, 140, 141, 142, 143, 144, 145
SUSY Supersymmetry	21, 146
SW ROD Software ROD. The software RODs perform the same tasks previously handled in the hardware RODs prior to Run 3, plus the buffering function of the ROS.	5, 137, 140, 141, 143, 146, 150, 163, 170, 171, 172, 175, 176, 177, 198
TAN Neutral Beam Absorber	131, 132, 133
TAS Target Absorber Secondaries	25, 26, 36, 121
TBB Tower Builder Board	72, 75
TC track-counting	126
TDAQ Trigger and Data Acquisition System	20, 22, 23, 76, 77, 123, 135, 138, 175, 176, 177, 185, 186, 187, 188
TDC Time-to-Digital Converter	117, 129, 130
TDS Trigger Data Serializer	103, 104, 105, 106, 107, 108
TGC Thin Gap Chamber	10, 12, 18, 20, 23, 86, 87, 89, 97, 100, 107, 109, 110, 111, 114, 115, 117, 135, 151, 157, 170, 188, 189, 192, 193, 199, 200

TGC-FI TGC Forward Inner chambers	151
TID Total Ionizing Dose	27, 28, 30, 31, 33, 34, 35, 36, 37, 43, 44, 117
TileCal Tile Calorimeter	77, 78, 80, 81, 82, 83, 119, 126, 192
TIM TTC Interface Module	45
TMDB Tile Muon Digitiser Board	152
TOB Trigger Object	136, 138, 140, 141, 143, 144, 145, 146, 147, 161, 162, 163
ToF Time-of-Flight	128, 129, 130
ToT Time over Threshold	43, 62
TP Trigger Processor	106, 107, 154, 155, 156
TREX Tile REciever eXchange	140, 146, 148, 149
TRT Transition Radiation Tracker .	2, 3, 10, 11, 13, 14, 15, 32, 34, 38, 39, 55, 56, 57, 58, 60, 62, 66, 69
TT Calorimeter Trigger Tower	70, 72
TTC Trigger, Timing, and Control system	45, 73, 77, 107, 136, 141, 143, 146, 148, 149, 150, 152, 157, 163, 165, 168, 170, 172, 174, 191, 192, 200
Tx Transmitter	45, 46
UPS Uninterruptible Power Supply	51
US15 Underground service area	106, 111
USA15 Main underground electronics cavern	46, 51, 59, 75, 101, 103, 106, 107, 111, 123, 131, 133, 152, 154
UX15 ATLAS underground experimental cavern	51, 59, 111
VAX new LHC vacuum pumping system to be installed in LS3	121, 122
VBF Vector Boson Fusion	161
VCSEL Vertical-Cavity Surface Emitting Laser	45
vdM van der Meer	119, 124, 125, 126
VLAN Virtual Local Area Network	176
VLDB Versatile Link demonstrator board	130
VME Versa Module Europa; a computer bus standard	45, 46, 73, 123, 137, 148, 149, 152, 165, 167, 168, 170, 181, 186
VMM Custom front-end ASIC for the Muon NSW	97, 103, 104, 105, 106, 107, 108, 109, 154
VTRx Versatile Link Transceiver (bi-directional version)	102, 105
VTTx Versatile Link Transceiver, unidirectional twin transmitter version	102, 105, 106

WinCC OA Simatic WinCC Open Architecture	49, 181, 183, 185, 186
WLS Wavelength-Shifting	80, 81, 82
XML EXtensible Markup Language	182
ZDC Zero Degree Calorimeters	4, 19, 118, 131, 132, 133

References

- [1] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1, arXiv: [1207.7214](https://arxiv.org/abs/1207.7214) [[hep-ex](#)].
- [2] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30, arXiv: [1207.7235](https://arxiv.org/abs/1207.7235) [[hep-ex](#)].
- [3] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003.
- [4] G. Arduini et al., *LHC Upgrades in preparation of Run 3*, *JINST* (2023), To be published in the same volume as this paper.
- [5] ATLAS Collaboration, *ATLAS data quality operations and performance for 2015–2018 data-taking*, *JINST* **15** (2020) P04003, arXiv: [1911.04632](https://arxiv.org/abs/1911.04632) [[physics.ins-det](#)].
- [6] R. Steerenberg, *Batch Compression Merging and Splitting (BCMS)*, (2017), CERN General Photo, URL: <https://cds.cern.ch/record/2259071>.
- [7] C. Ankenbrandt, *Slip Stacking: A New Method of Momentum Stacking*, FERMILAB-FN-0352, 1981, URL: <https://inspirehep.net/literature/169350>.
- [8] ATLAS Collaboration, *ATLAS Insertable B-Layer: Technical Design Report*, ATLAS-TDR-19; CERN-LHCC-2010-013, 2010, URL: <https://cds.cern.ch/record/1291633>, Addendum: ATLAS-TDR-19-ADD-1; CERN-LHCC-2012-009, 2012, URL: <https://cds.cern.ch/record/1451888>.
- [9] B. Abbott et al., *Production and integration of the ATLAS Insertable B-Layer*, *JINST* **13** (2018) T05008, arXiv: [1803.00844](https://arxiv.org/abs/1803.00844) [[physics.ins-det](#)].
- [10] ATLAS Collaboration, *Performance of the ATLAS trigger system in 2015*, *Eur. Phys. J. C* **77** (2017) 317, arXiv: [1611.09661](https://arxiv.org/abs/1611.09661) [[hep-ex](#)].
- [11] ATLAS Collaboration, *The ATLAS Collaboration Software and Firmware*, ATL-SOFT-PUB-2021-001, 2021, URL: <https://cds.cern.ch/record/2767187>.
- [12] ATLAS Collaboration, *ATLAS Liquid Argon Calorimeter: Technical Design Report*, ATLAS-TDR-2; CERN-LHCC-96-041, 1996, URL: <https://cds.cern.ch/record/331061>.
- [13] ATLAS Collaboration, *ATLAS Tile Calorimeter: Technical Design Report*, ATLAS-TDR-3; CERN-LHCC-96-042, 1996, URL: <https://cds.cern.ch/record/331062>.

- [14] ATLAS Collaboration, *Readiness of the ATLAS liquid argon calorimeter for LHC collisions*, *Eur. Phys. J. C* **70** (2010) 723, arXiv: 0912.2642 [hep-ex].
- [15] ATLAS Collaboration, *ATLAS New Small Wheel: Technical Design Report*, ATLAS-TDR-020; CERN-LHCC-2013-006, 2013, URL: <https://cds.cern.ch/record/1552862>.
- [16] G. P. Summers et al., *Correlation of Particle-Induced Displacement Damage in Silicon*, *IEEE Transactions on Nuclear Science* **34** (1987) 1133, URL: <https://ieeexplore-ieee-org.ezproxy.cern.ch/document/4337442>.
- [17] M. Moll, *Displacement Damage in Silicon Detectors for High Energy Physics*, *IEEE Trans. Nucl. Sci.* **65** (2018) 1561.
- [18] M. Huhtinen and P. Aarnio, *Pion induced displacement damage in silicon devices*, *Nucl. Instrum. Meth. A* **335** (1993) 580, URL: <http://www.sciencedirect.com/science/article/pii/016890029391246J>.
- [19] M. Huhtinen, *Simulation of non-ionising energy loss and defect formation in silicon*, *Nucl. Instrum. Meth. A* **491** (2002) 194, URL: <http://www.sciencedirect.com/science/article/pii/S0168900202012275>.
- [20] R.-y. Zhu, *Radiation damage in scintillating crystals*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **413** (1998) 297, URL: <https://www.sciencedirect.com/science/article/pii/S0168900298004987>.
- [21] M. Huhtinen and F. Faccio, *Computational method to estimate Single Event Upset rates in an accelerator environment*, *Nucl. Instrum. Meth. A* **450** (2000) 155, URL: <http://www.sciencedirect.com/science/article/pii/S0168900200001558>.
- [22] J. Hartert et al., ‘The ATLAS radiation dose measurement system and its extension to SLHC experiments’, *Proceedings of Topical Workshop on Electronics for Particle Physics, Naxos, Greece, 2008*, URL: <http://cds.cern.ch/record/1158640?ln=en>.
- [23] ATLAS Collaboration, *Modelling radiation damage to pixel sensors in the ATLAS detector*, *JINST* **14** (2019) P06012, arXiv: 1905.03739 [hep-ex].
- [24] R. Pedro on behalf of the ATLAS Collaboration, *Optics robustness of the ATLAS Tile Calorimeter*, *J. Phys.: Conf. Ser.* **1162** (2019) 012004, URL: <https://iopscience.iop.org/article/10.1088/1742-6596/1162/1/012004>.
- [25] A. Ferrari, P. R. Sala, A. Fasso and J. Ranft, *FLUKA: A multi-particle transport code*, CERN-2005-010, 2005, URL: <https://cds.cern.ch/record/898301>.
- [26] T. T. Böhlen et al., *The FLUKA Code: Developments and Challenges for High Energy and Medical Applications*, *Nucl. Data Sheets* **120** (2014) 211.
- [27] A. Ferrari, K. M. Potter, S. Rollet and P. R. Sala, *Radiation calculations for the ATLAS detector and experimental hall*, CERN-EST-96-001, 1996, URL: <https://cds.cern.ch/record/300336>.

- [28] C. Zeitnitz and T. A. Gabriel,
The GEANT-CALOR interface and benchmark calculations of ZEUS test calorimeters,
Nucl. Instrum. Meth. A **349** (1994) 106,
URL: <http://www.sciencedirect.com/science/article/pii/S0168900294906130>.
- [29] GEANT4 Collaboration, S. Agostinelli et al., *GEANT4 – a simulation toolkit*,
Nucl. Instrum. Meth. A **506** (2003) 250.
- [30] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159,
arXiv: [1410.3012](https://arxiv.org/abs/1410.3012) [[hep-ph](#)].
- [31] M. Moll,
Radiation damage in silicon particle detectors: Microscopic defects and macroscopic properties,
Ph.D. Thesis: Universität Hamburg, 1999,
URL: <https://bib-pubdb1.desy.de/record/300958>.
- [32] ATLAS Collaboration,
Measurements of sensor radiation damage in the ATLAS inner detector using leakage currents,
JINST **16** (2021) P08025, arXiv: [2106.09287](https://arxiv.org/abs/2106.09287) [[hep-ex](#)].
- [33] ATLAS Collaboration, *ATLAS pixel detector electronics and sensors*, *JINST* **3** (2008) P07007.
- [34] ATLAS Collaboration, *Operation and performance of the ATLAS semiconductor tracker*,
JINST **9** (2014) P08009, arXiv: [1404.7473](https://arxiv.org/abs/1404.7473) [[hep-ex](#)].
- [35] The ATLAS TRT collaboration,
The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: design and performance,
JINST **3** (2008) P02013.
- [36] ATLAS Collaboration, ATLAS-TDR-19-ADD-1; CERN-LHCC-2012-009, 2012,
URL: <https://cds.cern.ch/record/1451888>.
- [37] S. Parker, C. Kenney and J. Segal,
3D — A proposed new architecture for solid-state radiation detectors,
Nucl. Instrum. Meth. A **395** (1997) 328, Proceedings of the Third International Workshop on
Semiconductor Pixel Detectors for Particles and X-rays, ISSN: 0168-9002,
URL: <https://www.sciencedirect.com/science/article/pii/S0168900297006943>.
- [38] M. Garcia-Sciveres et al., *The FE-I4 pixel readout integrated circuit*,
Nucl. Instrum. Meth. A **636** (2011) S155, 7th International "Hiroshima" Symposium on the
Development and Application of Semiconductor Tracking Detectors, ISSN: 0168-9002,
URL: <https://www.sciencedirect.com/science/article/pii/S0168900210009551>.
- [39] Ian Dawson (ed.),
Radiation effects in the LHC experiments: Impact on detector performance and operation,
CERN Yellow Reports: Monographs (2021),
URL: <https://doi.org/10.23731/CYRM-2021-001>.
- [40] ATLAS Collaboration, *Alignment of the ATLAS Inner Detector in Run-2*,
Eur. Phys. J. C **80** (2020) 1194, arXiv: [2007.07624](https://arxiv.org/abs/2007.07624) [[hep-ex](#)].
- [41] G. Crone et al.,
The ATLAS ReadOut System Performance with first data and perspective for the future,
Nucl. Instrum. Meth. A **623** (2010) 534, 1st International Conference on Technology and
Instrumentation in Particle Physics, ISSN: 0168-9002,
URL: <http://www.sciencedirect.com/science/article/pii/S016890021000625X>.

- [42] E. van der Bij and S. Haas, *CERN S-LINK homepage*, 2003,
URL: <http://hsi.web.cern.ch/HSI/s-link/>.
- [43] SFF Committee, *Quad Small Form-factor Pluggable (QSFP) Transceiver Specification*, 2006,
URL: <https://www.gigalight.com/downloads/standards/QSFP-MSA.pdf>.
- [44] B. Hallgren, H. Boterenbrood, H.J. Burckhart and H. Kvedalen,
'The embedded local monitor board (ELMB) in the LHC front-end I/O control system',
Proceedings of the 7th Workshop on Electronics for LHC Experiments: Stockholm, Sweden 10 - 14 Sep 2001. Ed. by C. Isabella, CERN, Geneva: CERN, 2001 325,
URL: <https://cds.cern.ch/record/530675>.
- [45] ETM professional control GmbH, *SIMATIC WinCC Open Architecture SCADA System*,
URL: <http://www.winccoa.com/>.
- [46] O. Beltramello et al., *The Detector Safety System of the ATLAS experiment*,
JINST **4** (2009) P09012, URL: <https://dx.doi.org/10.1088/1748-0221/4/09/P09012>.
- [47] The ATLAS TRT collaboration, *The ATLAS TRT Barrel Detector*, *JINST* **3** (2008) P02014.
- [48] The ATLAS TRT collaboration, *The ATLAS TRT end-cap detectors*, *JINST* **3** (2008) P10003.
- [49] The ATLAS TRT collaboration, *The ATLAS TRT electronics*, *JINST* **3** (2008) P06007.
- [50] D. A. Huffman, *A Method for the Construction of Minimum-Redundancy Codes*,
Proceedings of the IRE **40** (1952) 1098.
- [51] A. Ruiz, E. van der Bij and S. Haas, *HOLA High-speed Optical Link for Atlas*, 2003,
URL: <http://hsi.web.cern.ch/hsi/s-link/devices/hola/>.
- [52] ATLAS Collaboration,
Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: tracker properties,
JINST **12** (2017) P05002, arXiv: [1702.06473](https://arxiv.org/abs/1702.06473) [hep-ex].
- [53] ATLAS Collaboration, *Study of the material of the ATLAS inner detector for Run 2 of the LHC*,
JINST **12** (2017) P12009, arXiv: [1707.02826](https://arxiv.org/abs/1707.02826) [hep-ex].
- [54] ATLAS Collaboration,
Operation and performance of the ATLAS semiconductor tracker in LHC Run 2,
JINST **17** (2021) P01013, arXiv: [2109.02591](https://arxiv.org/abs/2109.02591) [physics.ins-det].
- [55] ATLAS Collaboration,
ATLAS Liquid Argon Calorimeter Phase-I Upgrade: Technical Design Report,
ATLAS-TDR-022; CERN-LHCC-2013-017, 2013,
URL: <https://cds.cern.ch/record/1602230>.
- [56] G. Aad et al.,
The Phase-I trigger readout electronics upgrade of the ATLAS Liquid Argon calorimeters,
JINST **17** (2022) P05024, arXiv: [2202.07384](https://arxiv.org/abs/2202.07384) [physics.ins-det].
- [57] W. Cleland and E. Stern,
Signal processing considerations for liquid ionization calorimeters in a high rate environment,
Nucl. Instrum. Meth. A **338** (1994) 467,
URL: <https://www.sciencedirect.com/science/article/pii/0168900294913323>.
- [58] ATLAS Collaboration,
Monitoring and data quality assessment of the ATLAS liquid argon calorimeter,
JINST **9** (2014) P07024, arXiv: [1405.3768](https://arxiv.org/abs/1405.3768) [hep-ex].

- [59] H. Jivan et al., *Radiation hardness of plastic scintillators for the Tile Calorimeter of the ATLAS detector*, *Journal of Physics: Conference Series* **645** (2015) 012019, URL: <https://doi.org/10.1088/1742-6596/645/1/012019>.
- [60] S. Liao et al., *A comparative study of the radiation hardness of plastic scintillators for the upgrade of the Tile Calorimeter of the ATLAS detector*, *Journal of Physics: Conference Series* **645** (2015) 012021, URL: <https://doi.org/10.1088/1742-6596/645/1/012021>.
- [61] ATLAS Collaboration, *Electron and photon energy calibration with the ATLAS detector using data collected in 2015 at $\sqrt{s} = 13$ TeV*, ATL-PHYS-PUB-2016-015, 2016, URL: <https://cds.cern.ch/record/2203514>.
- [62] ATLAS Collaboration, *ATLAS Muon Spectrometer Phase-II Upgrade: Technical Design Report*, ATLAS-TDR-026; CERN-LHCC-2017-017, 2017, URL: <https://cds.cern.ch/record/2285580>.
- [63] ATLAS Collaboration, *Letter of Intent for the Phase-I Upgrade of the ATLAS Experiment*, CERN-LHCC-2011-012. LHCC-I-020, 2011.
- [64] ATLAS Collaboration, *ATLAS TDAQ System Phase-I Upgrade: Technical Design Report*, ATLAS-TDR-023; CERN-LHCC-2013-018, 2013, URL: <https://cds.cern.ch/record/1602235>.
- [65] J. Dubbert et al., ‘Development of precision drift tube detectors for very high background rates at the super-LHC’, *2007 IEEE Nuclear Science Symposium Conference Record*, vol. 3, 2007 1822.
- [66] ATLAS Collaboration, *ATLAS TDAQ Phase-II Upgrade: Technical Design Report*, ATLAS-TDR-029; CERN-LHCC-2017-020, 2017, URL: <https://cds.cern.ch/record/2285584>.
- [67] Y. Giomataris, P. Rebourgeard, J. Robert and G. Charpak, *MICROMEGAS: a high-granularity position-sensitive gaseous detector for high particle-flux environments*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **376** (1996) 29, URL: <https://www.sciencedirect.com/science/article/pii/S0168900296001751>.
- [68] T. Alexopoulos et al., *Performance studies of resistive-strip bulk micromegas detectors in view of the ATLAS New Small Wheel upgrade*, *Nucl. Instrum. Meth. A* **937** (2019) 125, URL: <http://www.sciencedirect.com/science/article/pii/S0168900219305194>.
- [69] H. Raether, *Die Entwicklung der Elektronenlawine in den Funkenkanal*, *Z. Phys.* **112** (1939) 464.
- [70] V. Peskov, M. Cortesi, R. Chechik and A. Breskin, *Further evaluation of a THGEM UV-photon detector for RICH - comparison with MWPC*, *JINST* **5** (2010) P11004, arXiv: [1008.0151](https://arxiv.org/abs/1008.0151) [[physics.ins-det](https://arxiv.org/abs/1008.0151)].
- [71] E. Farina and on behalf of the ATLAS Muon Collaboration, *ATLAS NSW Micromegas readout boards: Industrialisation and Quality Control and Quality Assurance*, *J. Phys.:Conf. Ser.* **1498** (2020) 012052, URL: <https://dx.doi.org/10.1088/1742-6596/1498/1/012052>.

- [72] J. Allard et al., *The large inner Micromegas modules for the ATLAS Muon Spectrometer upgrade: Construction, quality control and characterization*, *Nucl. Instrum. Meth. A* **1026** (2022) 166143, URL: <https://www.sciencedirect.com/science/article/pii/S0168900221010330>.
- [73] J. Agarwala et al., *Construction and test of the SM1 type Micromegas chambers for the upgrade of the ATLAS forward muon spectrometer*, *Nucl. Instrum. Meth. A* **1040** (2022) 167285, URL: <https://www.sciencedirect.com/science/article/pii/S0168900222006118>.
- [74] I. Gnesi, *Micromegas chambers for the ATLAS New Small Wheel upgrade*, *JINST* **15** (2020) C09019, URL: <https://dx.doi.org/10.1088/1748-0221/15/09/C09019>.
- [75] F. Kuger, *Production and quality control of Micromegas anode PCBs for the ATLAS NSW upgrade*, *JINST* **11** (2016) C11010, URL: <https://doi.org/10.1088/1748-0221/11/11/C11010>.
- [76] M. Beker et al., *The Rasnik 3-point optical alignment system*, *JINST* **14** (2019) P08010, URL: <https://doi.org/10.1088/1748-0221/14/12/P08010>.
- [77] G. Iakovidis, 'VMM ASIC', *J. Phys.:Conf. Ser.* Vol. 1498, MicroPattern Gaseous Detectors Conference 2019, La Rochelle, France, 2020 012051, URL: <https://doi.org/10.1088/1742-6596/1498/1/012051>.
- [78] ATLAS Collaboration, *Stereo Information in Micromegas Detectors*, ATL-MUON-PUB-2015-001, 2015, URL: <https://cds.cern.ch/record/2052206>.
- [79] K. Ntekas, *Performance characterization of the Micromegas detector for the New Small Wheel upgrade and Development and improvement of the Muon Spectrometer Detector Control System in the ATLAS experiment*, (pp. 180–185), PhD thesis: National Technical University, Athens, 2016, URL: <https://cds.cern.ch/record/2143887?ln=en>.
- [80] M. Iodice, *Resistive Micromegas for the Muon Spectrometer Upgrade of the ATLAS Experiment*, *PoS ICHEP2016* (2017) 275.
- [81] T. Alexopoulos et al., *Construction techniques and performances of a full-size prototype Micromegas chamber for the ATLAS muon spectrometer upgrade*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **955** (2020) 162086, URL: <https://www.sciencedirect.com/science/article/pii/S0168900219304966>.
- [82] S. Majewski, G. Charpak, A. Breskin and G. Mikenberg, *A thin multiwire chamber operating in the high multiplication mode*, *Nucl. Instrum. Meth.* **217** (1983) 265, URL: <https://www.sciencedirect.com/science/article/pii/0167508783901461>.
- [83] A. Abusleme et al., *Performance of a full-size small-strip thin gap chamber prototype for the ATLAS new small wheel muon upgrade*, *Nucl. Instrum. Meth. A* **817** (2016) 85, URL: <http://www.sciencedirect.com/science/article/pii/S0168900216001285>.
- [84] S. Aefsky et al., *The Optical Alignment System of the ATLAS Muon Spectrometer Endcaps*, *JINST* **3** (2008) P11005, URL: <https://doi.org/10.1088/1748-0221/3/11/P11005>.
- [85] G. Iakovidis, L. Levinson et al., *The New Small Wheel electronics*, *JINST* **18** (2023) P05012, arXiv: 2303.12571 [hep-ex], URL: <https://dx.doi.org/10.1088/1748-0221/18/05/P05012>.

- [86] P. Moreira, *GBTX MANUAL V0.15*, 2018,
URL: <https://espace.cern.ch/GBT-Project/GBTX/Manuals/gbtxManual.pdf>.
- [87] S. Bonacini, A. Caratelli et al., *GBT-SCA User Manual V. 2019-002*, (online), 2019, URL: https://espace.cern.ch/GBT-Project/GBT-SCA/Manuals/GBT-SCA_Manual_2019.002.pdf.
- [88] A. Caratelli et al., *The GBT-SCA, a radiation tolerant ASIC for detector control and monitoring applications in HEP experiments*, *JINST* **10** (2015) C03034,
URL: <https://doi.org/10.1088/1748-0221/10/03/c03034>.
- [89] C. Fuentes et al., *Optimization of DC-DC Converters for Improved Electromagnetic Compatibility With High Energy Physics Front-End Electronics*, *IEEE Transactions on Nuclear Science* **58** (2011) 2024, For more details about the FEAST, see <https://espace.cern.ch/project-DCDC-new/>.
- [90] P. Moreira, A. Marchioro and Kloukinas, *The GBT: A proposed architecture for multi-Gb/s data transmission in high energy physics*, 2007,
URL: <http://cds.cern.ch/record/1091474>.
- [91] G. de Geronimo, G. Iakovidis, S. Martoiu and V. Polychronakos, *The VMM3a ASIC*, *IEEE Transactions on Nuclear Science* **69** (2022) 976.
- [92] R. Coliban et al., *The Read Out Controller for the ATLAS New Small Wheel*, *JINST* **11** (2016) C02069.
- [93] J. Wang, L. Guan, J. W. Chapman, B. Zhou and J. Zhu, *Design of a Trigger Data Serializer ASIC for the Upgrade of the ATLAS Forward Muon Spectrometer*, *IEEE Transactions on Nuclear Science* **64** (2017) 2958.
- [94] S. Tang et al., ‘The Development and Production of the ADDC for the Micromegas Detector of the ATLAS New Small Wheel Upgrade’, *2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 2019 1.
- [95] P. Gkoutoumis, *LEVEL-1 DATA DRIVER CARD - A high bandwidth radiation tolerant aggregator board for detectors*, *PoS MPGD2017* (2019) 036.
- [96] J. Wang et al., *FPGA Implementation of a Fixed Latency Scheme in a Signal Packet Router for the Upgrade of ATLAS Forward Muon Trigger Electronics*, *IEEE Transactions on Nuclear Science* **62** (2015) 2194.
- [97] ATLAS Collaboration,
The Serial and LVDS repeaters for the ATLAS New Small Wheel sTGC trigger,
ATL-MUON-PUB-2022-003, 2022, URL: <https://cds.cern.ch/record/2841848>.
- [98] G. Blanchot, *Grounding of the ATLAS experiment*,
ATL-ELEC-PUB-2007-002, ATL-COM-ELEC-2007-003, 2007,
URL: <https://cds.cern.ch/record/1073170>.
- [99] H. Kroha et al., *Construction and test of a full prototype drift-tube chamber for the upgrade of the ATLAS muon spectrometer at high LHC luminosities*, *Nucl. Instrum. Meth. A* **718** (2013) 427,
ed. by F. Cervelli, G. Chiarelli, F. Forti, M. Grassi and A. Scribano.
- [100] H. Kroha, R. Fakhroutdinov and A. Kozhin,
New high-precision drift-tube detectors for the ATLAS muon spectrometer, **12** (2017) C06007,
URL: <https://dx.doi.org/10.1088/1748-0221/12/06/C06007>.

- [101] B. Bittner et al., *Performance of drift-tube detectors at high counting rates for high-luminosity LHC upgrades*, *Nucl. Instrum. Meth. A* **732** (2013) 250, ed. by T. Bergauer et al., arXiv: [1603.09508](https://arxiv.org/abs/1603.09508) [[physics.ins-det](#)].
- [102] B. Mandelli, R. Guida and G. Rigoletti, *Performance studies of RPC detectors operated with new environmentally-friendly gas mixtures in presence of LHC-like radiation background*, *PoS ICHEP2020* (2021) 857.
- [103] G. Aielli et al., *New results on ATLAS RPC's aging at CERN's GIF*, *IEEE Trans. Nucl. Sci.* **53** (2006) 567.
- [104] Q. Li et al., *Performance study of HL-LHC ATLAS RPC prototype*, *JINST* **14** (2019) C09022.
- [105] C. Ferretti and H. Kroha, *Upgrades of the ATLAS Muon Spectrometer With sMDT Chambers*, *Nucl. Instr. and Meth. A* **824** (2016) 538.
- [106] O. Kortner et al., *Upgrades of the ATLAS Muon Spectrometer with New Small-Diameter Drift Tube Chambers*, *Nucl. Instr. and Meth. A* **936** (2019) 509.
- [107] L. Pizzimento et al., *Development of a new Front End electronics in Silicon and Silicon-Germanium technology for the Resistive Plate Chamber detector for high rate experiments*, *JINST* **14** (2019) C10010.
- [108] J. Christiansen, *HPTDC – High Performance Time to Digital Converter*, 2004, URL: <https://cds.cern.ch/record/1067476/files/cer-002723234.pdf>.
- [109] Y. Arai et al., *ATLAS Muon Drift Tube Electronics*, *JINST* **3** (2008) P09001.
- [110] J. Anderson et al., 'FELIX: The new approach for interfacing to front-end electronics for the ATLAS experiment', *IEEE NSS Real Time Conference*, <https://ieeexplore.ieee.org/document/7543142/>, 2016.
- [111] P. Moreira et al., *A radiation tolerant gigabit serializer for LHC data transmission*, (2001), URL: <https://cds.cern.ch/record/588665>.
- [112] *Xilinx Kintex-7 FPGA product information*, URL: <https://www.xilinx.com/products/silicon-devices/fpga/kintex-7.html>.
- [113] V. Cindro et al., *The ATLAS beam conditions monitor*, *JINST* **3** (2008) P02004.
- [114] S. van der Meer, *Calibration of the effective beam height in the ISR*, CERN-ISR-PO-68-31, 1968, URL: <https://cds.cern.ch/record/296752>.
- [115] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC*, *Eur. Phys. J. C* **76** (2016) 653, arXiv: [1608.03953](https://arxiv.org/abs/1608.03953) [[hep-ex](#)].
- [116] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, ATLAS-CONF-2019-021, 2019, URL: <https://cds.cern.ch/record/2677054>.
- [117] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, (2022), arXiv: [2212.09379](https://arxiv.org/abs/2212.09379) [[hep-ex](#)].

- [118] G. Alberghi et al., *Choice and characterization of photomultipliers for the new ATLAS LUCID detector*, *JINST* **11** (2016) P05014.
- [119] G. Avoni et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, *JINST* **13** (2018) P07017.
- [120] *Caen SY 1527 Technical Information Manual*, datasheet, URL: https://www.caen.it/documents/Events/4/opcrelease3x_rev10.pdf.
- [121] ATLAS Collaboration, *The LUCID 3 detector for the ATLAS Phase-II Upgrade*, (2021), URL: <https://cds.cern.ch/record/2780604>.
- [122] A. Sopczak et al., *Precision Luminosity of LHC Proton-Proton Collisions at 13 TeV Using Hit Counting With TPX Pixel Devices*, *IEEE Trans. Nucl. Sci.* **64** (2017) 915, arXiv: [1702.00711](https://arxiv.org/abs/1702.00711) [[physics.ins-det](https://arxiv.org/abs/1702.00711)].
- [123] ATLAS Collaboration, *Luminosity monitoring using $Z \rightarrow \ell^+\ell^-$ events at $\sqrt{s} = 13$ TeV with the ATLAS detector*, ATL-DAPR-PUB-2021-001, 2021, URL: <https://cds.cern.ch/record/2752951>.
- [124] S. Abdel Khalek, B. Allongue, F. Anghinolfi et al., *The ALFA Roman Pot Detectors of ATLAS*, *JINST* **11** (2016) 11013, arXiv: [1609.00249](https://arxiv.org/abs/1609.00249) [[physics.ins-det](https://arxiv.org/abs/1609.00249)].
- [125] A. Almeghari, *Radiation resistance of the ATLAS-ALFA electronics and trigger efficiency analysis at $\sqrt{s} = 13$ TeV and $\beta^*=2.5$ km*, MA thesis: Islamic University of Gaza, 2018, URL: <https://library.iugaza.edu.ps/thesis/124502.pdf>.
- [126] ATLAS Collaboration, *ATLAS Forward Proton Phase-I Upgrade: Technical Design Report*, ATLAS-TDR-024; CERN-LHCC-2015-009, 2015, URL: <https://cds.cern.ch/record/2017378>.
- [127] J. Lange, E. Cavallaro, S. Grinstein and I. L. Paz, *3D silicon pixel detectors for the ATLAS Forward Physics experiment*, *JINST* **10** (2015) C03031, URL: [https://doi.org/10.1088/1748-0221/10_03_c03031](https://doi.org/10.1088/1748-0221/10/03/c03031).
- [128] J. Lange et al., *Beam tests of an integrated prototype of the ATLAS Forward Proton detector*, *JINST* **11** (2016) P09005.
- [129] C. Ng et al., *Engineering Change Request - Installation of the ATLAS/AFP stations, Phase-2*, tech. rep., CERN, 2016.
- [130] M. Garcia-Sciveres, *The RD53A Integrated Circuit*, tech. rep., CERN, 2017, URL: <https://cds.cern.ch/record/2287593>.
- [131] X. Llopart et al., *Timepix4, a large area pixel detector readout chip which can be tiled on 4 sides providing sub-200 ps timestamp binning*, *JINST* **17** (2022) C01044, URL: <https://dx.doi.org/10.1088/1748-0221/17/01/C01044>.
- [132] D. Houška, V. Vacek, M. Doubek and M. Haubner, *Commissioning of cooling system AIRCOOLER SPLIT before installation for TOTEM detectors*, tech. rep. internal CTU note, CTU, Prague, 2013, URL: <https://pdfs.semanticscholar.org/59d0/a841d8f11ba217421c35296193de45491813.pdf>.
- [133] L. Nozka et al., *Upgraded Cherenkov time-of-flight detector for the AFP project*, *Opt. Express* **31** (2023) 3998, URL: <https://opg.optica.org/oe/abstract.cfm?URI=oe-31-3-3998>.

- [134] K. Cerny et al., *Performance study of the ATLAS Forward Proton Time-of-Flight Detector System*, Proceedings of Science (2020), URL: <https://cds.cern.ch/record/2703903>.
- [135] *Mini PLANACON XPM85112-S-R2D2: a long-life, high-rate, 4x4-channel multi-anode multi-channel plate photomultiplier*, PHOTONIS Defense, Inc., 1000 New Holland Ave., Lancaster, PA 17601-5688, USA, URL: www.photonis.com.
- [136] J. Zich, V. Georgiev, M. Holik, V. Pavlicek and O. Vavroch, ‘Multichannel Majority Coincidence Circuit for ToF AFP Detector’, *27th Telecommunications Forum TELFOR 2019*, Belgrade, Serbia, 2019.
- [137] R. Herbst et al., *Design of the SLAC RCE Platform: A general purpose ATCA based data acquisition system*, tech. rep. SLAC-PUB-16182, SLAC, 2014, URL: <https://www.slac.stanford.edu/pubs/slacpubs/16000/slac-pub-16182.pdf>.
- [138] E. Banaś et al., *Detector Control System for the AFP detector in ATLAS experiment at CERN*, *J. Phys. Conf. Ser.* **898** (2017) 032002, URL: <http://iopscience.iop.org/article/10.1088/1742-6596/898/3/032022>.
- [139] *H6559 Photomultiplier tube assembly*, Hamamatsu, 325-6, Sunayama-cho, Naka-ku, Hamamatsu City, Shizuoka Pref., 430-8587, Japan, URL: www.hamamatsu.com/eu/en/product/type/H6559/index.html.
- [140] H. Matis et al., *The BRAN luminosity detectors for the LHC*, *Nucl. Instrum. Meth. A* **848** (2017) 114, URL: <https://www.sciencedirect.com/science/article/pii/S0168900216312797>.
- [141] S. Ask et al., *The ATLAS central level-1 trigger logic and TTC system*, *JINST* **3** (2008) P08002.
- [142] ATLAS Collaboration, *Performance of the upgraded PreProcessor of the ATLAS Level-1 Calorimeter Trigger*, *JINST* **15** (2020) P11016, arXiv: 2005.04179 [hep-ex].
- [143] ATLAS Collaboration, *Performance of the ATLAS Level-1 topological trigger in Run 2*, *Eur. Phys. J. C* **82** (2021) 7, arXiv: 2105.01416 [hep-ex].
- [144] ATLAS Collaboration, *ATLAS Tile Calorimeter Phase-II Upgrade: Technical Design Report*, ATLAS-TDR-028; CERN-LHCC-2017-019, 2017, URL: <https://cds.cern.ch/record/2285583>.
- [145] Xilinx®, *Aurora 8B/10B Protocol Specification*, https://www.xilinx.com/content/dam/xilinx/support/documentation/ip_documentation/aurora_8b10b_protocol_spec_sp002.pdf, Accessed: 2021-11-04.
- [146] C. Ghabrous Larrea, K. Harder, D. Newbold, D. Sankey, A. Rose, A. Thea and T. Williams, *IPbus: a flexible Ethernet-based control system for xTCA hardware*, *JINST* **10** (2015) C02019.
- [147] CERN, *CERN-IPMC*, <https://cern-ipmc.web.cern.ch/doc>, Accessed: 2021-11-04.
- [148] D. Bertolini, T. Chan, and J. Thaler, *Jet Observables Without Jet Algorithms*, *JHEP* **1404** (2014) 013, arXiv: 1310.7584 [hep-ph].

- [149] Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Inc., *Intelligent Platform Management Interface Specification, v. 2.0*, <https://www.intel.com/content/dam/www/public/us/en/documents/specification-updates/ipmi-intelligent-platform-mgt-interface-spec-2nd-gen-v2-0-spec-update.pdf>, Accessed: 2021-11-04, 2015.
- [150] ATLAS Collaboration, *Performance of the ATLAS muon triggers in Run 2*, *JINST* **15** (2020) P09015, arXiv: 2004.13447 [hep-ex].
- [151] *Advanced Telecom Computing Architecture*, URL: <http://www.picmg.org/openstandards/advancedtca/>.
- [152] Salvatore Loffredo, ATLAS Muon Collaboration, *The BIS78 Pad Trigger Board for the Phase-I Upgrade of the Level-1 Muon Trigger of the ATLAS Experiment at the LHC*, *2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC) (2019) 1*, URL: <https://ieeexplore.ieee.org/document/9059637>.
- [153] K. Wyllie et al., *A Gigabit Transceiver for Data Transmission in Future High Energy Physics Experiments*, *Physics Procedia* **37** (2012) 1561, Proceedings of the 2nd International Conference on Technology and Instrumentation in Particle Physics (TIPP 2011), ISSN: 1875-3892, URL: <http://www.sciencedirect.com/science/article/pii/S1875389212018706>.
- [154] A. Armbruster et al., *The ATLAS Muon to Central Trigger Processor Interface Upgrade for the Run 3 of the LHC*, *2017 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC) (2017)*.
- [155] J. Jeffers and J. Reinders, 'Chapter 8 - Coprocessor Architecture', *Intel Xeon Phi Coprocessor High Performance Programming*, ed. by J. Jeffers and J. Reinders, Boston: Morgan Kaufmann, 2013 243, ISBN: 978-0-12-410414-3, URL: <https://www.sciencedirect.com/science/article/pii/B9780124104143000086>.
- [156] J. Schumacher, 'Utilizing HPC Network Technologies in High Energy Physics Experiments', *2017 IEEE 25th Annual Symposium on High-Performance Interconnects (HOTI)*, 2017 57.
- [157] G. Barrand et al., *GAUDI — A software architecture and framework for building HEP data processing applications*, *Computer Physics Communications* **140** (2001) 45, CHEP2000, URL: <https://www.sciencedirect.com/science/article/pii/S0010465501002545>.
- [158] Bray, T., ed., *The JavaScript Object Notation (JSON) Data Interchange Format, STD 90, RFC 8259*, (online), 2017, URL: <https://www.rfc-editor.org/info/std90>.
- [159] A. Barriuso Poy et al., *The detector control system of the ATLAS experiment*, *JINST* **3** (2008) P0500.
- [160] *OPC UA Online Reference*, (online), URL: <https://reference.opcfoundation.org/>.
- [161] S. Schlenker et al., 'quasar - A Generic Framework for Rapid Development of OPC UA Servers', *Proceedings, 15th International Conference on Accelerator and Large Experimental Physics Control Systems (ICALPCS 2015): Melbourne, Australia, October 17-23, 2015*, 2015 WEB3002, URL: <http://jacow.org/icalpcs2015/papers/web3o02.pdf>.

- [162] P. P. Nikiel, B. Farnham, V. Filimonov and S. Schlenker, *Generic OPC UA Server Framework*, *Journal of Physics: Conference Series* **664** (2015) 082039, URL: <https://doi.org/10.1088%2F1742-6596%2F664%2F8%2F082039>.
- [163] P. Moschovakos, H. Boterenbrood, A. Koulouris, P. Nikiel and S. Schlenker, *A Software Suite for the Radiation Tolerant Giga-bit Transceiver - Slow Control Adapter*, *International Conference on Accelerator and Large Experimental Physics Control Systems* (2020) 1333, <https://doi.org/10.18429/JACoW-ICALEPCS2019-WEPHA102>, URL: <https://jacow.org/icalepcs2019/papers/wepha102.pdf>.
- [164] *UaExpert—A Full-Featured OPC UA Client*, (online), URL: <https://www.unified-automation.com/products/development-tools/uaexpert.html>.
- [165] ATLAS Collaboration, *ATLAS Inner Tracker Strip Detector: Technical Design Report*, ATLAS-TDR-025; CERN-LHCC-2017-005, 2017, URL: <https://cds.cern.ch/record/2257755>.
- [166] ATLAS Collaboration, *ATLAS Inner Tracker Pixel Detector: Technical Design Report*, ATLAS-TDR-030; CERN-LHCC-2017-021, 2017, URL: <https://cds.cern.ch/record/2285585>.
- [167] ATLAS Collaboration, *ATLAS LAr Calorimeter Phase-II Upgrade: Technical Design Report*, ATLAS-TDR-027; CERN-LHCC-2017-018, 2017, URL: <https://cds.cern.ch/record/2285582>.
- [168] ATLAS Collaboration, *A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade: Technical Design Report*, ATLAS-TDR-031; CERN-LHCC-2020-007, 2020, URL: <https://cds.cern.ch/record/2719855>.
- [169] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS Trigger and Data Acquisition System — EF Tracking Amendment*, ATLAS-TDR-029; CERN-LHCC-2022-004, 2022, URL: <http://cds.cern.ch/record/2802799>.
- [170] ATLAS Collaboration, *ATLAS Computing Acknowledgements*, ATL-SOFT-PUB-2021-003, 2021, URL: <https://cds.cern.ch/record/2776662>.

The ATLAS Collaboration

G. Aad ¹⁰², B. Abbott ¹²⁰, D.C. Abbott ¹⁰³, J. Abdallah⁸, K. Abeling ⁵⁵, S.H. Abidi ²⁹,
A. Abouhorma ^{35e}, S. Abovyan¹¹⁰, H. Abramowicz ¹⁵¹, H. Abreu ¹⁵⁰, Y. Abulaiti ¹¹⁷,
A.C. Abusleme Hoffman ^{137a}, B.S. Acharya ^{69a,69b,p}, C. Adam Bourdarios ⁴, L. Adamczyk ^{85a},
L. Adamek ¹⁵⁵, S.V. Addepalli ²⁶, J. Adelman ¹¹⁵, M. Adersberger¹⁰⁹, A. Adiguzel ^{21c},
S. Adorni ⁵⁶, T. Adye ¹³⁴, A.A. Affolder ¹³⁶, Y. Afik ³⁶, M.N. Agaras ¹³, J. Agarwala ^{73a,73b},
A. Aggarwal ¹⁰⁰, C. Agheorghiesei ^{27c}, J.A. Aguilar-Saavedra ^{130f}, A. Ahmad ³⁶,
F. Ahmadov ^{38,z}, W.S. Ahmed ¹⁰⁴, S. Ahuja ⁹⁵, X. Ai ⁴⁸, G. Aielli ^{76a,76b}, M. Ait Tamlihat ^{35e},
B. Aitbenkikh ^{35a}, I. Aizenberg ¹⁷¹, M. Akbiyik ¹⁰⁰, T.P.A. Åkesson ⁹⁸, G. Akhperjanyan¹⁷⁵,
A.V. Akimov ³⁷, K. Al Khoury ⁴¹, G.L. Alberghi ^{23b}, J. Albert ¹⁶⁷, P. Albicocco ⁵³,
S. Alderweireldt ⁵², M. Aleksa ³⁶, I.N. Aleksandrov ³⁸, C. Alexa ^{27b}, T. Alexopoulos ¹⁰,
A. Alfonsi ¹¹⁴, F. Alfonsi ^{23b}, M. Alhroob ¹²⁰, B. Ali ¹³², S. Ali ¹⁴⁸, M. Aliev ³⁷,
G. Alimonti ^{71a}, W. Alkakh ⁵⁵, C. Allaire ⁶⁶, J. Allard¹³⁵, B.M.M. Allbrooke ¹⁴⁶,
C.A. Allendes Flores ^{137f}, P.P. Allport ²⁰, A. Aloisio ^{72a,72b}, F. Alonso ⁹⁰, C. Alpigiani ¹³⁸,
M. Alvarez Estevez ⁹⁹, B. Alvarez Gonzalez ³⁶, M.G. Alvigi ^{72a,72b}, M. Aly ¹⁰¹,
Y. Amaral Coutinho ^{82b}, A. Ambler ¹⁰⁴, C. Amelung³⁶, M. Amerl ¹, C.G. Ames ¹⁰⁹,
D. Amidei ¹⁰⁶, S.P. Amor Dos Santos ^{130a}, K.R. Amos ¹⁶⁵, V. Ananiev ¹²⁵, C. Anastopoulos ¹³⁹,
N. Andari ¹³⁵, T. Andeen ¹¹, J.K. Anders ³⁶, S.Y. Andrean ^{47a,47b}, A. Andreazza ^{71a,71b},
C.R. Anelli¹⁶⁷, S. Angelidakis ⁹, A. Angerami ^{41,ab}, A.V. Anisenkov ³⁷, A. Annovi ^{74a},
C. Antel ⁵⁶, M.T. Anthony ¹³⁹, E. Antipov ¹²¹, M. Antonelli ⁵³, M. Antonescu^{27e},
D.J.A. Antrim ^{17a}, F. Anulli ^{75a}, M. Aoki ⁸³, T. Aoki ¹⁵³, J.A. Aparisi Pozo ¹⁶⁵, M.A. Aparo ¹⁴⁶,
L. Aperio Bella ⁴⁸, C. Appelt ¹⁸, N. Aranzabal ³⁶, V. Araujo Ferraz ^{82a}, C. Arcangeletti ⁵³,
A.T.H. Arce ⁵¹, E. Arena ⁹², J-F. Arguin ¹⁰⁸, A. Argyris³⁶, S. Argyropoulos ⁵⁴, J.-H. Arling ⁴⁸,
A.J. Armbruster ³⁶, C.E. Armijo⁷, O. Arnaez ¹⁵⁵, H. Arnold ¹¹⁴, Z.P. Arrubarrena Tame¹⁰⁹,
G. Artoni ^{75a,75b}, H. Asada ¹¹¹, K. Asai ¹¹⁸, S. Asai ¹⁵³, N.A. Asbah ⁶¹, J. Assahsah ^{35d},
K. Assamagan ²⁹, R. Astalos ^{28a}, R.J. Atkin ^{33a}, M. Atkinson¹⁶³, N.B. Atlay ¹⁸, H. Atmani^{62b},
P.A. Atlasiddha ¹⁰⁶, E. Auberon¹³⁵, K. Augsten ¹³², S. Aune¹³⁵, S. Auricchio ^{72a,72b},
A.D. Auriol ²⁰, F. Aust¹⁰⁹, V.A. Austrup ¹⁷³, G. Avner ¹⁵⁰, G. Avolio ³⁶, G. Avoni ^{23a},
D. Axen¹⁶⁶, K. Axiotis ⁵⁶, P. Aydinier¹⁰⁹, M.K. Ayoub ^{14c}, T. Azaryan ¹⁵⁸, G. Azuelos ^{108,ae},
D. Babal ^{28a}, H. Bachacou ¹³⁵, K. Bachas ^{152,r}, A. Bachi ³⁴, F. Backman ^{47a,47b}, A. Badea ⁶¹,
P. Bagnaia ^{75a,75b}, M. Bahmani ¹⁸, A.J. Bailey ¹⁶⁵, V.R. Bailey ¹⁶³, J.T. Baines ¹³⁴,
C. Bakalis ¹⁰, O.K. Baker ¹⁷⁴, P.J. Bakker ¹¹⁴, E. Bakos ¹⁵, D. Bakshi Gupta ⁸, S. Balaji ¹⁴⁷,
R. Balasubramanian ¹¹⁴, G. Balbi ^{23a}, E.M. Baldin ³⁷, P. Balek ¹³³, R. Ball¹⁰⁶,
E. Ballabene ^{71a,71b}, J. Ballansat⁴, F. Balli ¹³⁵, L.M. Baltos ^{63a}, W.K. Balunas ³², J. Balz ¹⁰⁰,
J. Ban ⁴¹, E. Banas ⁸⁶, M. Bandieramonte ¹²⁹, A. Bandyopadhyay ²⁴, S. Bansal ²⁴, L. Barak ¹⁵¹,
E.L. Barberio ¹⁰⁵, D. Barberis ^{57b,57a}, M. Barbero ¹⁰², G. Barbier⁵⁶, G. Barbour⁹⁶, L. Bardo³⁶,
K.N. Barends ^{33a}, A. Barfusser ¹⁰⁹, T. Barillari ¹¹⁰, M-S. Barisits ³⁶, T. Barklow ¹⁴³,
R.M. Barnett ^{17a}, P. Baron ¹²², D.A. Baron Moreno ¹⁰¹, A. Baroncelli ^{62a}, G. Barone ²⁹,
A.J. Barr ¹²⁶, L. Barranco Navarro ^{47a,47b}, F. Barreiro ⁹⁹, J. Barreiro Guimarães da Costa ^{14a},
U. Barron ¹⁵¹, M.G. Barros Teixeira ^{130a}, S. Barsov ³⁷, F. Bartels ^{63a}, R. Bartoldus ¹⁴³,
A.E. Barton ⁹¹, P. Bartos ^{28a}, A. Basalae ⁴⁸, A. Basan ¹⁰⁰, M. Baselga ⁴⁹, I. Bashta ^{77a,77b},
A. Bassalat ⁶⁶, M.J. Basso ¹⁵⁵, C.R. Basson ¹⁰¹, R.L. Bates ⁵⁹, S. Batlamous^{35e}, J.R. Batley ³²,
B. Batool ¹⁴¹, M. Battaglia ¹³⁶, D. Battulga ¹⁸, M. Bause ^{75a,75b}, P. Bauer ²⁴, A. Bayirli ^{21a},
J.B. Beacham ⁵¹, T. Beau ¹²⁷, B.C. Beauchamp¹²¹, P.H. Beauchemin ¹⁵⁸, R. Beccherle ^{74b},
F. Becherer ⁵⁴, P. Bechtle ²⁴, H.P. Beck ^{19,q}, K. Becker ¹⁶⁹, A.J. Beddall ^{21d},

V.A. Bednyakov ³⁸, C.P. Bee ¹⁴⁵, L.J. Beemster ¹⁵, T.A. Beermann ³⁶, M. Begalli ^{82d,82d}, M. Begel ²⁹, A. Behera ¹⁴⁵, J.K. Behr ⁴⁸, C. Beirao Da Cruz E Silva ³⁶, J.F. Beirer ^{55,36}, F. Beisiegel ²⁴, C. Belanger-Champagne ¹⁰⁴, M. Belfkir ¹⁵⁹, F. Belhesan ¹⁷¹, G. Bella ¹⁵¹, F. Bellachia ⁴, L. Bellagamba ^{23b}, A. Bellerive ³⁴, P. Bellos ²⁰, K. Beloborodov ³⁷, K. Belotskiy ³⁷, J. Beltramelli ¹³⁵, N.L. Belyaev ³⁷, M. Ben Moshe ²⁶, D. Bencheqroun ^{35a}, F. Bendebba ^{35a}, J. Bendotti ³⁶, Y. Benhammou ¹⁵¹, D.P. Benjamin ²⁹, M. Benoit ²⁹, T.A. Benoit ¹³⁵, J.R. Bensingher ²⁶, S. Bentvelsen ¹¹⁴, L. Beresford ³⁶, M. Beretta ⁵³, E. Bergeas Kuutmann ¹⁶², N. Berger ⁴, B. Bergmann ¹³², J. Beringer ^{17a}, S. Berlendis ⁷, G. Bernardi ⁵, C. Bernius ¹⁴³, F.U. Bernlochner ²⁴, F. Bernon ³⁶, T. Berry ⁹⁵, P. Berta ¹³³, A. Berthold ⁵⁰, I.A. Bertram ⁹¹, H. Bervas ¹³⁵, D. Besin ¹³⁵, I. Bessudo ¹⁵¹, S. Bethke ¹¹⁰, A. Betti ^{75a,75b}, A.J. Bevan ⁹⁴, Th. Bey ¹³⁵, M. Bhamjee ^{33c}, S. Bhatta ¹⁴⁵, D.S. Bhattacharya ¹⁶⁸, P. Bhattarai ²⁶, V.S. Bhopatkar ¹²¹, R. Bi ^{29,ah}, R.M. Bianchi ¹²⁹, Y. Bianga ⁵⁰, M. Biaut ¹⁰², O. Biebel ¹⁰⁹, R. Bielski ¹²³, M. Biglietti ^{77a}, T.R.V. Billoud ¹³², M. Bindi ⁵⁵, A. Bingul ^{21b}, C. Bini ^{75a,75b}, A. Biondini ⁹², C. Bira ^{27e}, C.J. Birch-sykes ¹⁰¹, G.A. Bird ^{20,134}, M. Birman ¹⁷¹, P. Birney ^{156a}, M. Biro ¹³³, T. Bisanz ³⁶, E. Bisceglie ^{43b,43a}, D. Biswas ^{172,1}, D. Bitá ¹⁶¹, A. Bitadze ¹⁰¹, K. Bjørke ¹²⁵, T.P. Blaszczyk ³⁶, I. Bloch ⁴⁸, C. Blocker ²⁶, A. Blue ⁵⁹, U. Blumenschein ⁹⁴, J. Blumenthal ¹⁰⁰, G.J. Bobbink ¹¹⁴, V.S. Bobrovnikov ³⁷, M. Boehler ⁵⁴, B. Boehm ¹⁶⁸, D. Bogavac ³⁶, A.G. Bogdanchikov ³⁷, C. Bohm ^{47a}, V. Boisvert ⁹⁵, P. Bokan ⁴⁸, T. Bold ^{85a}, D.D. Boline ¹⁴⁵, M. Bomben ⁵, M. Bona ⁹⁴, F. Bonini ²⁹, M. Boonekamp ¹³⁵, C.D. Booth ⁹⁵, A.G. Borbély ⁵⁹, H.M. Borecka-Bielska ¹⁰⁸, L.S. Borgna ⁹⁶, G. Borissov ⁹¹, J. Bortfeldt ³⁶, D. Bortoletto ¹²⁶, C. Bortolin ³⁶, D. Boscherini ^{23b}, M. Bosman ¹³, J.D. Bossio Sola ³⁶, J.M. Botte ³⁴, K. Bouaouda ^{35a}, S. Bouaziz ¹³⁵, N. Bouchhar ¹⁶⁵, J. Boudreau ¹²⁹, T. Bouedo ⁴, E.V. Bouhova-Thacker ⁹¹, D. Boumediene ⁴⁰, R. Bouquet ⁵, A. Boveia ¹¹⁹, J. Boyd ³⁶, D. Boye ²⁹, I.R. Boyko ³⁸, N. Braam ¹⁶⁷, J. Bracnik ²⁰, P.H. Braga Lisboa ^{82b}, N. Brahimi ^{62d}, G. Brandt ¹⁷³, O. Brandt ³², F. Braren ⁴⁸, B. Brau ¹⁰³, J.E. Brau ¹²³, I.P. Brawn ¹³⁴, K. Brendlinger ⁴⁸, R. Brenner ¹⁷¹, L. Brenner ¹¹⁴, R. Brenner ¹⁶², S. Bressler ¹⁷¹, P. Breugnon ¹⁰², D. Britton ⁵⁹, D. Britzger ¹¹⁰, I. Brock ²⁴, G. Brooijmans ⁴¹, W.K. Brooks ^{137f}, E. Brost ²⁹, L.M. Brown ¹⁶⁷, L.E. Bruce ⁶¹, T.L. Bruckler ¹²⁶, P.A. Bruckman de Renstrom ⁸⁶, B. Brüers ⁴⁸, D. Bruncko ^{28b,*}, A. Bruni ^{23b}, G. Bruni ^{23b}, K.M. Brunner ¹⁰⁴, M. Bruschi ^{23b}, N. Brusino ^{75a,75b}, T. Buanes ¹⁶, Q. Buat ¹³⁸, P. Buchholz ¹⁴¹, A.G. Buckley ⁵⁹, S.I. Buda ⁵⁴, I.A. Budagov ^{38,*}, M.K. Bugge ¹²⁵, O. Bulekov ³⁷, B.A. Bullard ¹⁴³, S. Burdin ⁹², C.D. Burgard ⁴⁹, A.M. Burger ⁴⁰, B. Burghgrave ⁸, J.T.P. Burr ³², C.D. Burton ¹¹, J.C. Burzynski ¹⁴², E.L. Busch ⁴¹, V. Büscher ¹⁰⁰, P.J. Bussey ⁵⁹, J.M. Butler ²⁵, C.M. Buttar ⁵⁹, J.M. Butterworth ⁹⁶, W. Buttinger ¹³⁴, C.J. Buxo Vazquez ¹⁰⁷, A.R. Buzykaev ³⁷, G. Cabras ^{23b}, S. Cabrera Urbán ¹⁶⁵, F. Cadoux ⁵⁶, D. Caforio ⁵⁸, H. Cai ¹²⁹, Y. Cai ^{14a,14d}, V.M.M. Cairo ³⁶, O. Cakir ^{3a}, D. Calabro ^{73b}, N. Calace ³⁶, P. Calafiura ^{17a}, G. Calderini ¹²⁷, P. Calfayan ⁶⁸, G. Callea ⁵⁹, L.P. Caloba ^{82b}, D. Calvet ⁴⁰, S. Calvet ⁴⁰, T.P. Calvet ¹⁰², M. Calvetti ^{74a,74b}, R. Camacho Toro ¹²⁷, S. Camarda ³⁶, D. Camarero Munoz ²⁶, P. Camarri ^{76a,76b}, M.T. Camerlingo ^{72a,72b}, D. Cameron ¹²⁵, C. Camincher ¹⁶⁷, M. Campanelli ⁹⁶, A. Camplani ⁴², V. Canale ^{72a,72b}, A. Canesse ¹⁰⁴, M. Cano Bret ⁸⁰, J. Cantero ¹⁶⁵, N.Y.Y. Cao ⁶¹, Y. Cao ¹⁶³, S. Cap ⁴, E. Capitolo ⁵³, F. Capocasa ²⁶, G. Capradossi ^{75a,75b}, M. Capua ^{43b,43a}, G.B. Cara ¹³⁵, A. Carbone ^{71a,71b}, R. Cardarelli ^{76a}, J.C.J. Cardenas ⁸, F. Cardillo ¹⁶⁵, C.A. Cardot ³⁶, T. Carli ³⁶, G. Carlino ^{72a}, M.D. Carlotta ³⁶, J.I. Carlotta ¹³, B.T. Carlson ^{129,s}, E.M. Carlson ^{167,156a}, K.J. Carlson ^{156a}, L. Carminati ^{71a,71b}, M. Carnesale ^{75a,75b}, S. Caron ¹¹³, E. Carquin ^{137f}, S. Carrá ^{71a,71b}, G. Carratta ^{23b,23a}, F. Carri Argos ^{33g}, J.W.S. Carter ¹⁵⁵, T.M. Carter ⁵², M.P. Casado ^{13,i}, A. Caserio ^{73a,73b}, A.F. Casha ¹⁵⁵, C. Cassese ^{72b}, E.G. Castiglia ¹⁷⁴, F.L. Castillo ^{63a}, L. Castillo Garcia ¹³, V. Castillo Gimenez ¹⁶⁵, N.F. Castro ^{130a,130e},

A. Catinaccio ³⁶, J.R. Catmore ¹²⁵, V. Cavaliere ²⁹, N. Cavalli ^{23b,23a}, V. Cavasinni ^{74a,74b},
 E. Celebi ^{21a}, F. Celli ¹²⁶, M.S. Centonze ^{70a,70b}, F. Ceradini ^{77b}, K. Cerny ¹²²,
 A.S. Cerqueira ^{82a}, A. Cerri ¹⁴⁶, L. Cerrito ^{76a,76b}, F. Cerutti ^{17a}, A. Cervelli ^{23b}, S.A. Cetin ^{21d},
 Z. Chadi ^{35a}, D. Chakraborty ¹¹⁵, M. Chala ^{130f}, Th. Chaleil ¹³⁵, J. Chan ¹⁷², S.K. Chan ⁶¹,
 W.Y. Chan ¹⁵³, J.D. Chapman ³², J.W. Chapman ¹⁰⁶, B. Chargeishvili ^{149b}, D.G. Charlton ²⁰,
 T.P. Charman ⁹⁴, M. Chatterjee ¹⁹, C.C. Chau ³⁴, S. Chekanov ⁶, S.V. Chekulaev ^{156a},
 G.A. Chelkov ^{38,a}, A. Chen ¹⁰⁶, B. Chen ¹⁵¹, B. Chen ¹⁶⁷, H. Chen ^{14c}, H. Chen ²⁹,
 J. Chen ^{62c}, J. Chen ¹⁴², K. Chen ²⁹, O. Chen ³⁴, S. Chen ¹⁵³, S.J. Chen ^{14c}, X. Chen ^{62c},
 X. Chen ^{14b,ad}, Y. Chen ^{62a}, C.L. Cheng ¹⁷², H.C. Cheng ^{64a}, S. Cheong ¹⁴³, A. Cheplakov ³⁸,
 E. Cheremushkina ⁴⁸, E. Cherepanova ¹¹⁴, R. Cherkaoui El Moursli ^{35e}, E. Cheu ⁷, K. Cheung ⁶⁵,
 L. Chevalier ¹³⁵, N. Chevillot ⁴, V. Chiarella ⁵³, G. Chiarelli ^{74a}, N. Chiedde ¹⁰², G. Chiodini ^{70a},
 A.S. Chisholm ²⁰, A. Chitan ^{27b}, M. Chitishvili ¹⁶⁵, Y.H. Chiu ¹⁶⁷, M.V. Chizhov ³⁸, K. Choi ¹¹,
 A.R. Chomont ^{75a,75b}, Y. Chou ¹⁰³, E.Y.S. Chow ¹¹⁴, T. Chowdhury ^{33g}, L.D. Christopher ^{33g},
 A. Chrul ³⁶, K.L. Chu ^{64a}, M.C. Chu ^{64a}, X. Chu ^{14a,14d}, J. Chudoba ¹³¹, J.J. Chwastowski ⁸⁶,
 G. Ciapetti ^{75b}, M. Ciapetti ³⁶, R.P. Ciecko ³⁶, D. Cieri ¹¹⁰, K.M. Ciesla ^{85a}, V. Cindro ⁹³,
 A. Ciocio ^{17a}, F. Cirotto ^{72a,72b}, Z.H. Citron ^{171,m}, M. Citterio ^{71a}, D.A. Ciubotaru ^{27b},
 B.M. Ciungu ¹⁵⁵, A. Clark ⁵⁶, B.L. Clark ⁶¹, P.J. Clark ⁵², J.M. Clavijo Columbie ⁴⁸,
 S.E. Clawson ¹⁰¹, W. Cleland ¹²⁹, J.C. Clemens ¹⁰², C. Clement ^{47a,47b}, J. Clercx ⁴⁸, L. Clissa ^{23b,23a},
 Y. Coadou ¹⁰², M. Cobal ^{69a,69c}, A. Coccaro ^{57b}, R.F. Coelho Barrue ^{130a},
 R. Coelho Lopes De Sa ¹⁰³, S. Coelli ^{71a}, G. Cohen ¹⁷¹, H. Cohen ¹⁵¹, A.E.C. Coimbra ^{71a,71b},
 B. Cole ⁴¹, R.M. Coliban ^{27a}, J. Collot ⁶⁰, P. Conde Muiño ^{130a,130g}, M.P. Connell ^{33c},
 S.H. Connell ^{33c}, I.A. Connelly ⁵⁹, E.I. Conroy ¹²⁶, M. Constable ^{156a,156b}, F. Conventi ^{72a,af},
 H.G. Cooke ²⁰, A.M. Cooper-Sarkar ¹²⁶, F. Corbaz ³⁶, F. Cormier ¹⁶⁶, L.D. Corpe ³⁶,
 M. Corradi ^{75a,75b}, E.E. Corrigan ⁹⁸, F. Corriveau ^{104,x}, S. Corsetti ¹⁰⁶, A. Cortes-Gonzalez ¹⁸,
 M.J. Costa ¹⁶⁵, T.C.P. Costa De Paiva ¹⁰³, F. Costanza ⁴, D. Costanzo ¹³⁹, B.M. Cote ¹¹⁹,
 G. Cowan ⁹⁵, J.W. Cowley ³², K. Cranmer ¹¹⁷, S. Crépe-Renaudin ⁶⁰, F. Crescioli ¹²⁷,
 O. Crespo-Lopez ³⁶, M. Cristinziani ¹⁴¹, M. Cristoforetti ^{78a,78b,c}, V. Croft ¹⁵⁸, G. Crosetti ^{43b,43a},
 A. Cueto ³⁶, T. Cuhadar Donszelmann ¹⁶⁰, H. Cui ^{14a,14d}, Z. Cui ⁷, W.R. Cunningham ⁵⁹,
 F. Curcio ^{43b,43a}, P. Czodrowski ³⁶, M.M. Czurylo ^{63b}, M.J. Da Cunha Sargedas De Sousa ^{62a},
 J.V. Da Fonseca Pinto ^{82b}, C. Da Via ¹⁰¹, W. Dabrowski ^{85a}, T. Dado ⁴⁹, J. Daguin ³⁶,
 S. Dahbi ^{33g}, T. Dai ¹⁰⁶, C. Dallapiccola ¹⁰³, M. Dam ⁴², G. D'amen ²⁹, V. D'Amico ¹⁰⁹,
 J. Damp ¹⁰⁰, J.R. Dandoy ¹²⁸, M.F. Daneri ³⁰, H.O. Danielsson ³⁶, V. Danielyan ¹¹⁰,
 E. Danilevich ³⁷, M. Danninger ¹⁴², V. Dao ³⁶, G. Darbo ^{57b}, S. Darmora ⁶, S.J. Das ²⁹,
 S. D'Auria ^{71a,71b}, C. David ^{156b}, P. David ⁴, T. Davidek ¹³³, D.R. Davis ⁵¹, P.M. Davis ²,
 B. Davis-Purcell ³⁴, L. Davoine ³⁶, I. Dawson ⁹⁴, K. De ⁸, R. De Asmundis ^{72a}, M. De Beurs ¹¹⁴,
 N. De Biase ⁴⁸, S. De Castro ^{23b,23a}, S. De Cecco ^{75b}, B. De Fazio ^{72b}, G. De Geronimo ¹⁰⁶,
 N. De Groot ¹¹³, P. de Jong ¹¹⁴, S.R. de Jong ¹⁶⁷, H. De la Torre ¹⁰⁷, A. De Maria ^{14c},
 A. De Salvo ^{75a}, U. De Sanctis ^{76a,76b}, A. De Santo ¹⁴⁶, J.B. De Vivie De Regie ⁶⁰,
 R. De Oliveira ³⁶, G. Decock ¹³⁵, D.V. Dedovich ³⁸, J. Degens ¹¹⁴, J. Degrange ³⁶, A.M. Deiana ⁴⁴,
 F. Del Corso ^{23b,23a}, J. Del Peso ⁹⁹, F. Del Rio ^{63a}, P. Delebecque ⁴, F. Deliot ¹³⁵,
 C.M. Delitzsch ⁴⁹, M. Della Pietra ^{72a,72b}, D. Della Volpe ⁵⁶, A. Dell'Acqua ³⁶,
 L. Dell'Asta ^{71a,71b}, M. Delmastro ⁴, P.A. Delsart ⁶⁰, S. Demers ¹⁷⁴, M. Demichev ³⁸, B. Deng ⁴⁴,
 S.P. Denisov ³⁷, L. D'Eramo ¹¹⁵, D. Derendarz ⁸⁶, F. Derue ¹²⁷, P. Dervan ⁹², K. Desch ²⁴,
 H. Deschamps ¹³⁵, D. Desforge ¹³⁵, K. Dette ¹⁵⁵, C. Deutsch ²⁴, F.A. Di Bello ^{57b,57a},
 A. Di Ciaccio ^{76a,76b}, L. Di Ciaccio ⁴, A. Di Domenico ^{75a,75b}, C. Di Donato ^{72a,72b},
 A. Di Girolamo ³⁶, G. Di Gregorio ⁵, A. Di Luca ^{78a,78b}, B. Di Micco ^{77a,77b}, R. Di Nardo ^{77a,77b},
 K.F. Di Petrillo ⁶¹, L. Di Stante ^{76b}, C. Diaconu ¹⁰², F.A. Dias ¹¹⁴, T. Dias Do Vale ¹⁴²,

M.A. Diaz [id](#)^{137a,137b}, F.G. Diaz Capriles [id](#)²⁴, M. Didenko [id](#)¹⁶⁵, E.B. Diehl [id](#)¹⁰⁶, L. Diehl [id](#)⁵⁴,
W. Dietsche²⁴, S. Díez Cornell [id](#)⁴⁸, C. Diez Pardos [id](#)¹⁴¹, A. Dik³⁷, N. Dikic³⁶, K. Dima¹⁵²,
C. Dimitriadi [id](#)^{24,162}, A. Dimitrievska [id](#)^{17a}, W. Ding²⁹, J. Dingfelder [id](#)²⁴, B. Dinkespiler¹⁰²,
I-M. Dinu [id](#)^{27b}, G. Dissert¹³⁵, S.J. Dittmeier [id](#)^{63b}, F. Dittus [id](#)³⁶, F. Djama [id](#)¹⁰², T. Djobava [id](#)^{149b},
J.I. Djuvslund [id](#)¹⁶, D. Dobos³⁶, C. Doglioni [id](#)^{101,98}, J. Dolejsi [id](#)¹³³, Z. Dolezal [id](#)¹³³, M. Donadelli [id](#)^{82c},
P. Dondero [id](#)^{73b}, B. Dong [id](#)¹⁰⁷, J. Donini [id](#)⁴⁰, A. D’Onofrio [id](#)^{77a,77b}, M. D’Onofrio [id](#)⁹², J. Dopke [id](#)¹³⁴,
O. Dorholt¹²⁵, A. Doria [id](#)^{72a}, M. Doubek [id](#)³⁶, M.T. Dova [id](#)⁹⁰, A.T. Doyle [id](#)⁵⁹, M.A. Draguet [id](#)¹²⁶,
E. Drechsler [id](#)¹⁴², E. Dreyer [id](#)¹⁷¹, I. Drivas-koulouris [id](#)¹⁰, A.S. Drobac [id](#)¹⁵⁸, M. Drozdova [id](#)⁵⁶,
D. Du [id](#)^{62a}, Y. Du [id](#)^{62b}, T.A. du Pree [id](#)¹¹⁴, Y. Duan [id](#)^{62d}, F. Dubinin [id](#)³⁷, M. Dubovsky [id](#)^{28a},
E. Duchovni [id](#)¹⁷¹, G. Duckeck [id](#)¹⁰⁹, O.A. Ducu [id](#)^{27b}, D. Duda [id](#)¹¹⁰, A. Dudarev [id](#)³⁶, A.C. Dudder [id](#)¹⁰⁰,
D. D’Uffizi^{75b}, M. D’uffizi [id](#)¹⁰¹, L. Duflo [id](#)⁶⁶, M. Dührssen [id](#)³⁶, R.M. Duim [id](#)^{156a}, C. Dülsen [id](#)¹⁷³,
A.E. Dumitriu [id](#)^{27b}, N. Dumont Dayot⁴, M. Dunford [id](#)^{63a}, M.G. Dunford [id](#)³⁴, S. Dungs [id](#)⁴⁹,
K. Dunne [id](#)^{47a,47b}, A. Duperrin [id](#)¹⁰², H. Duran Yildiz [id](#)^{3a}, M. Düren [id](#)⁵⁸, A. Durglishvili [id](#)^{149b},
B. Dury¹⁶⁶, A. Dushkin¹⁵⁸, B.L. Dwyer [id](#)¹¹⁵, G.I. Dyckes [id](#)^{17a}, M. Dyndal [id](#)^{85a}, S. Dysch [id](#)¹⁰¹,
B.S. Dziedzic [id](#)⁸⁶, P. Dziurdzia³⁶, Z.O. Earnshaw [id](#)¹⁴⁶, B. Eckerova [id](#)^{28a}, S. Eggebrecht [id](#)⁵⁵,
M.G. Eggleston⁵¹, E. Egidio Purcino De Souza [id](#)¹²⁷, L.F. Ehrke [id](#)⁵⁶, G. Eigen [id](#)¹⁶, K. Einsweiler [id](#)^{17a},
T. Ekelof [id](#)¹⁶², P.A. Ekman [id](#)⁹⁸, Y. El Ghazali [id](#)^{35b}, H. El Jarrari [id](#)^{35e,148}, A. El Moussaouy [id](#)^{35a},
V. Ellajosyula [id](#)¹⁶², M. Ellert [id](#)¹⁶², S. Elles⁴, F. Ellinghaus [id](#)¹⁷³, A.A. Elliot [id](#)⁹⁴, N. Ellis [id](#)³⁶,
J. Elmsheuser [id](#)²⁹, M. Elsing [id](#)³⁶, D. Emelianov [id](#)¹³⁴, A. Emerman [id](#)⁴¹, Y. Enari [id](#)¹⁵³, I. Ene [id](#)^{17a},
S. Epari [id](#)¹³, J. Erdmann [id](#)^{49,ac}, A. Ereditato [id](#)¹⁹, P.A. Erland [id](#)⁸⁶, M. Errenst [id](#)¹⁷³, M. Escalier [id](#)⁶⁶,
C. Escobar [id](#)¹⁶⁵, E. Etzion [id](#)¹⁵¹, G. Evans [id](#)^{130a}, H. Evans [id](#)⁶⁸, M.O. Evans [id](#)¹⁴⁶, A. Eyring²⁴,
A. Ezhilov [id](#)³⁷, S. Ezzarqtouni [id](#)^{35a}, F. Fabbri [id](#)⁵⁹, L. Fabbri [id](#)^{23b,23a}, G. Facini [id](#)⁹⁶, V. Fadeyev [id](#)¹³⁶,
R.M. Fakhruddinov [id](#)³⁷, D. Falchieri [id](#)^{23a}, S. Falciano [id](#)^{75a}, L.F. Falda Ulhoa Coelho [id](#)³⁶, P.J. Falke [id](#)²⁴,
S. Falke [id](#)³⁶, A.C. Falou⁶⁶, G. Falsetti^{43a}, J. Faltova [id](#)¹³³, Y. Fan [id](#)^{14a}, Y. Fang [id](#)^{14a,14d},
G. Fanourakis [id](#)⁴⁶, M. Fanti [id](#)^{71a,71b}, M. Faraj [id](#)^{69a,69b}, Z. Farazpay⁹⁷, A. Farbin [id](#)⁸, A. Farilla [id](#)^{77a},
E.M. Farina [id](#)^{73b}, T. Farooque [id](#)¹⁰⁷, J. Farrell²⁹, S.M. Farrington [id](#)⁵², F. Fassi [id](#)^{35e}, D. Fassouliotis [id](#)⁹,
W. Faszler^{156a}, M. Faucci Giannelli [id](#)^{76a,76b}, C. Fausten¹⁷³, A. Favareto [id](#)^{57a}, W.J. Fawcett [id](#)³²,
L. Fayard [id](#)⁶⁶, D.R. Febvre³⁶, P. Federicova [id](#)¹³¹, O.L. Fedin [id](#)^{37,a}, G. Fedotov [id](#)³⁷, M. Feickert [id](#)¹⁷²,
L. Feligioni [id](#)¹⁰², A. Fell [id](#)¹³⁹, D.E. Fellers [id](#)¹²³, N.F. Felt⁶¹, C. Feng [id](#)^{62b}, M. Feng [id](#)^{14b}, Z. Feng [id](#)¹¹⁴,
M.J. Fenton [id](#)¹⁶⁰, A.B. Fenyuk³⁷, L. Ferencz [id](#)⁴⁸, R.A.M. Ferguson [id](#)⁹¹, S.I. Fernandez Luengo [id](#)^{137f},
J. Ferrando [id](#)⁴⁸, A. Ferrari [id](#)¹⁶², P. Ferrari [id](#)^{114,113}, R. Ferrari [id](#)^{73a}, U.R. Ferreira Dias [id](#)^{82a},
E. Ferrer Ribas [id](#)¹³⁵, D. Ferrere [id](#)⁵⁶, C. Ferretti [id](#)¹⁰⁶, F. Fiedler [id](#)¹⁰⁰, W. Fielitz²⁹, Y. Filippov³⁸,
A. Filipčič [id](#)⁹³, E.K. Filmer [id](#)¹, F. Filthaut [id](#)¹¹³, M.C.N. Fiolhais [id](#)^{130a,130c,b}, G. Fiore [id](#)^{70b},
L. Fiorini [id](#)¹⁶⁵, A. Fischer¹¹⁰, F. Fischer [id](#)¹⁴¹, W.C. Fisher [id](#)¹⁰⁷, T. Fitschen [id](#)¹⁰¹, I. Fleck [id](#)¹⁴¹,
P. Fleischmann [id](#)¹⁰⁶, T. Flick [id](#)¹⁷³, B.M. Flierl [id](#)¹⁰⁹, C. Flores [id](#)^{137a}, L. Flores [id](#)¹²⁸, M. Flores [id](#)^{33d},
L.R. Flores Castillo [id](#)^{64a}, F.M. Follega [id](#)^{78a,78b}, I. Fomichev³⁸, N. Fomin [id](#)¹⁶, M. Fontaine¹³⁵,
J.H. Foo [id](#)¹⁵⁵, B.C. Forland⁶⁸, L.F. Formenti [id](#)¹⁰⁴, A. Formica [id](#)¹³⁵, A.C. Forti [id](#)¹⁰¹, E. Fortin [id](#)¹⁰²,
A.W. Fortman [id](#)⁶¹, M.G. Foti [id](#)^{17a}, D. Fougeron¹⁰², L. Fountas [id](#)⁹, D. Fournier [id](#)⁶⁶, H. Fox [id](#)⁹¹,
J.C. Fragnaud⁴, P. Francavilla [id](#)^{74a,74b}, S. Francescato [id](#)⁶¹, S. Franchellucci [id](#)⁵⁶, M. Franchini [id](#)^{23b,23a},
S. Franchino [id](#)^{63a}, D. Francis³⁶, L. Franco [id](#)¹¹³, L. Franconi [id](#)¹⁹, N. Frank³⁶, M. Franklin [id](#)⁶¹,
M. Fras [id](#)¹¹⁰, G. Frattari [id](#)²⁶, A. Freddi^{73b}, A.C. Freegard [id](#)⁹⁴, P.M. Freeman²⁰, W.S. Freund [id](#)^{82b},
N. Fritzsche [id](#)⁵⁰, A. Froch [id](#)⁵⁴, D. Froidevaux [id](#)³⁶, J.A. Frost [id](#)¹²⁶, Y. Fu [id](#)^{62a}, M. Fujimoto [id](#)¹¹⁸,
E. Fullana Torregrosa [id](#)^{165,*}, M. Furukawa [id](#)¹⁵³, J. Fuster [id](#)¹⁶⁵, A. Gabrielli [id](#)^{23b,23a}, A. Gabrielli [id](#)¹⁵⁵,
P. Gadow [id](#)⁴⁸, G. Gagliardi [id](#)^{57b,57a}, L.G. Gagnon [id](#)^{17a}, G.E. Gallardo [id](#)¹²⁶, E.J. Gallas [id](#)¹²⁶,
R.B. Galleguillos Silva [id](#)^{137b,137d}, B.J. Gallop [id](#)¹³⁴, S.J. Galuszka³⁶, R. Gamboa Goni [id](#)⁹⁴, K.K. Gan [id](#)¹¹⁹,
S. Ganguly [id](#)¹⁵³, L.M. Gantel⁴, J. Gao [id](#)^{62a}, W. Gao [id](#)^{62b}, Y. Gao [id](#)⁵², F.M. Garay Walls [id](#)^{137a,137b},
B. Garcia^{29,ah}, C. García [id](#)¹⁶⁵, J.E. García Navarro [id](#)¹⁶⁵, M. Garcia-Sciveres [id](#)^{17a}, R.W. Gardner [id](#)³⁹,

M. Gareau ^{156a}, N. Garelli ¹⁵⁸, D. Garg ⁸⁰, R.B. Garg ¹⁴³, G.G. Gariano ^{57a}, C.A. Garner ¹⁵⁵,
V. Garonne ²⁹, S.J. Gasiorowski ¹³⁸, P. Gaspar ^{82b}, A. Gaudiello ^{57a}, G. Gaudio ^{73a}, V. Gautam ¹³,
P. Gauzzi ^{75a,75b}, I.L. Gavrilenko ³⁷, A. Gavrilyuk ³⁷, C. Gay ¹⁶⁶, G. Gaycken ⁴⁸, E.N. Gazis ¹⁰,
A.A. Geanta ^{27b,27e}, M. Gebyehu ¹¹⁴, C.M. Gee ¹³⁶, C.N.P. Gee ¹³⁴, J. Geisen ⁹⁸, C. Gemme ^{57b},
M.H. Genest ⁶⁰, S. Gentile ^{75a,75b}, M. George ⁵⁵, S. George ⁹⁵, W.F. George ²⁰, V. Georgiev ¹⁶⁴,
T. Geralis ⁴⁶, L.O. Gerlach ⁵⁵, P. Gessinger-Befurt ³⁶, M. Ghasemi Bostanabad ¹⁶⁷,
M. Ghneimat ¹⁴¹, K. Ghorbanian ⁹⁴, A. Ghosal ¹⁴¹, A. Ghosh ¹⁶⁰, A. Ghosh ⁷, B. Giacobbe ^{23b},
S. Giagu ^{75a,75b}, P. Giannetti ^{74a}, A. Giannini ^{62a}, S.M. Gibson ⁹⁵, A. Giganon ¹³⁵, K. Gigliotti ⁷,
M. Gignac ¹³⁶, D.T. Gil ^{85b}, A.K. Gilbert ^{85a}, B.J. Gilbert ⁴¹, D. Gillberg ³⁴, G. Gilles ¹¹⁴,
N.E.K. Gillwald ⁴⁸, L. Ginabat ¹²⁷, D.M. Gingrich ^{2,ae}, A. Giokaris ¹⁰, M.P. Giordani ^{69a,69c},
J. Giraud ¹³⁵, P.F. Giraud ¹³⁵, P. Giromini ⁶¹, G. Giugliarelli ^{69a,69c}, D. Giugni ^{71a}, F. Giuli ³⁶,
H. Gjersdal ¹²⁵, I. Gkialas ^{9,j}, E.L. Gkoukousis ¹³, P. Gkoutoumis ¹⁰, L.K. Gladilin ³⁷,
C. Glasman ⁹⁹, G.R. Gledhill ¹²³, M. Glisic ¹²³, K.W. Glitza ¹⁷³, G. Glonti ¹³⁵, I. Gnesi ^{43b,f},
Y. Go ^{29,ah}, C. Goblin ¹³⁵, M. Goblirsch-Kolb ²⁶, B. Gocke ⁴⁹, D. Godin ¹⁰⁸, S. Godiot ¹⁰²,
B. Gokturk ^{21a}, S. Goldfarb ¹⁰⁵, T. Golling ⁵⁶, M.G.D. Gololo ^{33g}, D. Golubkov ³⁷, D. Golyzniak ³⁶,
J.P. Gombas ¹⁰⁷, A. Gomes ^{130a,130b}, G. Gomes Da Silva ¹⁴¹, A.J. Gomez Delegido ¹⁶⁵,
R. Goncalves Gama ⁵⁵, R. Gonçalo ^{130a,130c}, G. Gonella ¹²³, L. Gonella ²⁰, D. Gong ⁴⁴,
A. Gongadze ³⁸, F. Gonnella ²⁰, J.L. Gonski ⁴¹, R.Y. González Andana ⁵²,
S. González de la Hoz ¹⁶⁵, S. Gonzalez Fernandez ¹³, R. Gonzalez Lopez ⁹²,
C. Gonzalez Renteria ^{17a}, R. Gonzalez Suarez ¹⁶², S. Gonzalez-Sevilla ⁵⁶,
G.R. Gonzalvo Rodriguez ¹⁶⁵, L. Goossens ³⁶, N.A. Gorasia ²⁰, P.A. Gorbounov ³⁷, B. Gorini ³⁶,
E. Gorini ^{70a,70b}, A. Gorišek ⁹³, A.T. Goshaw ⁵¹, C. Gössling ⁴⁹, M.I. Gostkin ³⁸, S. Goswami ¹²¹,
C.A. Gottardo ³⁶, S.A. Gotz ¹⁰⁹, M. Goughri ^{35b}, V. Goumarre ⁴⁸, A.G. Goussiou ¹³⁸,
N. Govender ^{33c}, C. Goy ⁴, A.M. Grabas ¹³⁵, I. Grabowska-Bold ^{85a}, K. Graham ³⁴,
E. Gramstad ¹²⁵, S. Grancagnolo ¹⁸, M. Grandi ¹⁴⁶, V. Gratchev ^{37,*}, P.A. Gravelle ³⁴,
P.M. Gravila ^{27f}, F.G. Gravila ^{70a,70b}, H.M. Gray ^{17a}, I. Grayzman ¹⁷¹, M. Greco ^{70a,70b},
C. Grefe ²⁴, I.M. Gregor ⁴⁸, P. Grenier ¹⁴³, C. Grieco ¹³, A.A. Grillo ¹³⁶, K. Grimm ^{31,n},
S. Grinstein ^{13,u}, J.-F. Grivaz ⁶⁶, J.P. Grohs ⁵⁰, E. Gross ¹⁷¹, J. Grosse-Knetter ⁵⁵, C. Grud ¹⁰⁶,
A. Grummer ¹¹², J.C. Grundy ¹²⁶, L. Guan ¹⁰⁶, W. Guan ¹⁷², C. Gubbels ¹⁶⁶,
S. Guelfo Gigli ^{73a,73b}, J.G.R. Guerrero Rojas ¹⁶⁵, G. Guerrieri ^{69a,69b}, F. Guescini ¹¹⁰,
N.E.H. Guettouche ¹⁰², R. Gugel ¹⁰⁰, J.A.M. Guhit ¹⁰⁶, A. Guida ⁴⁸, E. Guido ^{57a}, J. Guillard ¹³⁵,
T. Guillemin ⁴, E. Guilloton ^{169,134}, S. Guindon ³⁶, D. Guo ⁴⁴, F. Guo ^{14a,14d}, J. Guo ^{62c},
L. Guo ⁶⁶, Y. Guo ¹⁰⁶, R. Gupta ⁴⁸, S. Gurbuz ²⁴, S.S. Gurdasani ⁵⁴, G. Gustavino ³⁶,
M. Guth ⁵⁶, P. Gutierrez ¹²⁰, L.F. Gutierrez Zagazeta ¹²⁸, C. Gutschow ⁹⁶, C. Guyot ¹³⁵,
C. Gwenlan ¹²⁶, C.B. Gwilliam ⁹², E.S. Haaland ¹²⁵, A. Haas ¹¹⁷, S. Haas ³⁶, M. Habedank ⁴⁸,
C. Haber ^{17a}, J. Habring ¹¹⁰, H.K. Hadavand ⁸, A. Hadeef ¹⁰⁰, S. Hadzic ¹¹⁰, E.H. Haines ⁹⁶,
M. Haleem ¹⁶⁸, J. Haley ¹²¹, J.J. Hall ¹³⁹, G.D. Hallowell ¹⁰², L. Halser ¹⁹, K. Hamano ¹⁶⁷,
H. Hamdaoui ^{35e}, M. Hamer ²⁴, G.N. Hamity ⁵², J. Han ^{62b}, K. Han ^{62a}, L. Han ^{14c},
L. Han ^{62a}, S. Han ^{17a}, Y.F. Han ¹⁵⁵, K. Hanagaki ⁸³, M. Hance ¹³⁶, D.A. Hangal ^{41,ab},
H. Hanif ¹⁴², M.D. Hank ³⁹, R. Hankache ¹⁰¹, J.B. Hansen ⁴², J.D. Hansen ⁴², P.H. Hansen ⁴²,
K. Hara ¹⁵⁷, D. Harada ⁵⁶, T. Harenberg ¹⁷³, S. Harkusha ³⁷, Y.T. Harris ¹²⁶, N.M. Harrison ¹¹⁹,
P.F. Harrison ¹⁶⁹, N.M. Hartman ¹⁴³, N.M. Hartmann ¹⁰⁹, P. Hartung ¹⁰⁹, Y. Hasegawa ¹⁴⁰,
K. Hashemi ²⁶, A. Hasib ⁵², L.A. Hasley ⁴⁴, T.G. Haubold ¹¹⁰, S. Haug ¹⁹, R. Hauser ¹⁰⁷,
M. Havranek ¹³², C.M. Hawkes ²⁰, R.J. Hawkins ³⁶, Y. Hayashi ¹⁵³, S. Hayashida ¹¹¹,
D. Hayden ¹⁰⁷, C. Hayes ¹⁰⁶, R.L. Hayes ¹⁶⁶, C.P. Hays ¹²⁶, J.M. Hays ⁹⁴, H.S. Hayward ⁹²,
F. He ^{62a}, Y. He ¹⁵⁴, Y. He ¹²⁷, M.P. Heath ⁵², V. Hedberg ⁹⁸, A.L. Heggelund ¹²⁵,
N.D. Hehir ⁹⁴, C. Heidegger ⁵⁴, K.K. Heidegger ⁵⁴, W.D. Heidorn ⁸¹, J. Heilman ³⁴,

S. Heim ⁴⁸, T. Heim ^{17a}, J.G. Heinlein ¹²⁸, J.J. Heinrich ¹²³, L. Heinrich ¹¹⁰, J. Hejbal ¹³¹,
 L. Helary ⁴⁸, A. Held ¹⁷², S. Hellesund ¹²⁵, C.M. Helling ¹⁶⁶, S. Hellman ^{47a,47b}, C. Helsens ³⁶,
 T. Hemperek ²⁴, R.C.W. Henderson ⁹¹, L. Henkelmann ³², A.M. Henriques Correia ³⁶, R.G. Hentges ⁵⁰,
 H. Herde ⁹⁸, Y. Hernández Jiménez ¹⁴⁵, L.M. Herrmann ²⁴, M.G. Herrmann ¹⁰⁹, T. Herrmann ⁵⁰,
 G. Herten ⁵⁴, R. Hertenberger ¹⁰⁹, L. Hervas ³⁶, N.P. Hessey ^{156a}, H. Hibi ⁸⁴,
 E. Higón-Rodríguez ¹⁶⁵, S.J. Hillier ²⁰, M. Hils ⁵⁰, I. Hinchliffe ^{17a}, F. Hinterkeuser ²⁴,
 M. Hirose ¹²⁴, S. Hirose ¹⁵⁷, D. Hirschbuehl ¹⁷³, T.G. Hitchings ¹⁰¹, B. Hiti ⁹³, J. Hobbs ¹⁴⁵,
 R. Hobincu ^{27e}, N. Hod ¹⁷¹, M.C. Hodgkinson ¹³⁹, B.H. Hodgkinson ³², A. Hoecker ³⁶,
 M.R. Hoferkamp ¹¹², J. Hofer ⁴⁸, A.E. Hoffmann ²⁹, D. Hoffmann ¹⁰², D. Hohn ⁵⁴, D. Hohov ⁶⁶,
 T. Holm ²⁴, M. Holzbock ¹¹⁰, L.B.A.H. Hommels ³², B.P. Honan ¹⁰¹, J. Hong ^{62c},
 T.M. Hong ¹²⁹, J.C. Honig ⁵⁴, A. Hönle ¹¹⁰, B.H. Hooberman ¹⁶³, W.H. Hopkins ⁶, Y. Horii ¹¹¹,
 P. Horn ⁵⁰, S. Hou ¹⁴⁸, A.S. Howard ⁹³, J. Howarth ⁵⁹, J. Hoya ⁶, M. Hrabovsky ¹²²,
 A. Hrynevich ⁴⁸, T. Hryn'ova ⁴, P.J. Hsu ⁶⁵, S.-C. Hsu ¹³⁸, K. Hu ^{62b}, Q. Hu ⁴¹, X. Hu ¹⁰⁶,
 Y.F. Hu ^{14a,14d,ag}, D.P. Huang ⁹⁶, F. Huang ¹⁶⁶, S. Huang ^{64b}, X. Huang ^{14c}, Y. Huang ^{62a},
 Y. Huang ^{14a}, Z. Huang ¹⁰¹, Z. Hubacek ¹³², M. Huebner ²⁴, F. Huegging ²⁴, T.B. Huffman ¹²⁶,
 M. Huhtinen ³⁶, S.K. Huiberts ¹⁶, W.K. Hulek ³⁶, R. Hulsken ¹⁰⁴, N. Huseynov ^{12a}, J. Huston ¹⁰⁷,
 J. Huth ⁶¹, R. Hyneman ¹⁴³, S. Hyrych ^{28a}, G. Iacobucci ⁵⁶, G. Iakovidis ²⁹, K. Iakovidis ³⁶,
 V. Iankovskaia ¹⁶⁷, B. Iankovski ¹⁷¹, I. Ibragimov ¹⁴¹, L. Iconomidou-Fayard ⁶⁶, P. Iengo ^{72a,72b},
 R. Iguchi ¹⁵³, T. Iizawa ⁵⁶, Y. Ikegami ⁸³, M. Ikeno ⁸³, A. Ilg ¹⁹, N. Ilic ¹⁵⁵, I. Ilyashenko ³⁷,
 H. Imam ^{35a}, T. Ingebretsen Carlson ^{47a,47b}, A. Innocente ^{70b}, G. Introzzi ^{73a,73b}, M. Iodice ^{77a},
 V. Ippolito ^{75a,75b}, M. Ishino ¹⁵³, W. Islam ¹⁷², C. Issever ^{18,48}, S. Issinski ^{156a}, S. Istin ^{21a},
 H. Ito ¹⁷⁰, J.M. Iturbe Ponce ^{64a}, R. Iuppa ^{78a,78b}, O. Iurikovskii ³⁸, M. Ivanovici ^{27a}, A. Ivina ¹⁷¹,
 J.M. Izen ⁴⁵, V. Izzo ^{72a}, P. Jacka ^{131,132}, P. Jackson ¹, R.M. Jacobs ⁴⁸, B.P. Jaeger ¹⁴²,
 C.S. Jagfeld ¹⁰⁹, P. Jain ⁵⁴, G. Jäkel ¹⁷³, K. Jakobs ⁵⁴, T. Jakoubek ¹⁷¹, J. Jamieson ⁵⁹,
 K.W. Janas ^{85a}, J. Janssen ²⁴, G. Jarlskog ⁹⁸, A.E. Jaspan ⁹², S. Javello ¹³⁵, M. Javurkova ¹⁰³,
 F. Jeanneau ¹³⁵, L. Jeanty ¹²³, J. Jejelava ^{149a,aa}, P. Jenni ^{54,g}, J. Jentzsch ⁴⁹, C.E. Jessiman ³⁴,
 S. Jézéquel ⁴, C. Jia ^{62b}, J. Jia ¹⁴⁵, X. Jia ⁶¹, X. Jia ^{14a,14d}, Z. Jia ^{14c}, Y. Jiang ^{62a}, S. Jiggins ⁵²,
 J. Jimenez Pena ¹¹⁰, G. Jin ^{62a}, S. Jin ^{14c}, A. Jinaru ^{27b}, O. Jinnouchi ¹⁵⁴, P. Johansson ¹³⁹,
 K.A. Johns ⁷, J.W. Johnson ¹³⁶, D.M. Jones ³², E. Jones ¹⁶⁹, P. Jones ³², R.W.L. Jones ⁹¹,
 T.J. Jones ⁹², M. Joos ³⁶, J. Joseph ^{17a}, R. Joshi ¹¹⁹, D. Jourde ¹³⁵, J. Jovicevic ¹⁵, X. Ju ^{17a},
 J.J. Junggeburth ³⁶, T. Junkermann ^{63a}, A. Juste Rozas ^{13,u}, S. Kabana ^{137e}, A. Kaczmarzka ⁸⁶,
 M. Kado ^{75a,75b}, H. Kagan ¹¹⁹, M. Kagan ¹⁴³, A. Kahn ⁴¹, A. Kahn ¹²⁸, C. Kahra ¹⁰⁰, T. Kaji ¹⁷⁰,
 E. Kajomovitz ¹⁵⁰, N. Kakati ¹⁷¹, C.W. Kalderon ²⁹, A. Kallitsopoulou ¹⁵², A. Kamenshchikov ¹⁵⁵,
 S. Kanayama ¹⁵⁴, N. Kanellos ¹⁰, N.J. Kang ¹³⁶, K.K. Kapusciaik ^{85a}, D. Kar ^{33g}, K. Karava ¹²⁶,
 M.J. Kareem ^{156b}, E. Karentzos ⁵⁴, I. Karkanas ^{152,e}, S.N. Karpov ³⁸, Z.M. Karpova ³⁸,
 V. Kartvelishvili ⁹¹, A.N. Karyukhin ³⁷, E. Kasimi ^{152,e}, C. Kato ^{62d}, S. Katunin ³⁷, J. Katzy ⁴⁸,
 S. Kaur ³⁴, K. Kawade ¹⁴⁰, K. Kawagoe ⁸⁹, T. Kawamoto ¹³⁵, G. Kawamura ⁵⁵, E.F. Kay ¹⁶⁷,
 F.I. Kaya ¹⁵⁸, S. Kazakos ¹³, V.F. Kazanin ³⁷, A. Kazarov ^{33c}, Y. Ke ¹⁴⁵, J.M. Keaveney ^{33a},
 M. Keberri ¹³⁵, R. Keeler ¹⁶⁷, G.V. Kehris ⁶¹, J.S. Keller ³⁴, A.S. Kelly ⁹⁶, D. Kelsey ¹⁴⁶,
 J.J. Kempster ²⁰, K.E. Kennedy ⁴¹, P.D. Kennedy ¹⁰⁰, O. Kepka ¹³¹, B.P. Kerridge ¹⁶⁹,
 S. Kersten ¹⁷³, B.P. Kerševan ⁹³, S. Keshri ⁶⁶, L. Keszeghova ^{28a}, S. Ketabchi Haghight ¹⁵⁵,
 M. Khandoga ¹²⁷, A. Khanov ¹²¹, A.G. Kharlamov ³⁷, T. Kharlamova ³⁷, E.E. Khoda ¹³⁸,
 T.J. Khoo ¹⁸, G. Khoraiuli ¹⁶⁸, J. Khubua ^{149b}, Y.A.R. Khwaira ⁶⁶, M. Kiehn ³⁶,
 A. Kilgallon ¹²³, D.W. Kim ^{47a,47b}, E. Kim ¹⁵⁴, Y.K. Kim ³⁹, N. Kimura ⁹⁶, P. Kind ¹⁷³,
 P. Kinget ^{aj,41}, A. Kirchhoff ⁵⁵, D. Kirchmeier ⁵⁰, C. Kirfel ²⁴, J. Kirk ¹³⁴, A.E. Kiryunin ¹¹⁰,
 T. Kishimoto ¹⁵³, I. Kiskiras ⁴⁶, D.P. Kisliuk ¹⁵⁵, C. Kitsaki ¹⁰, O. Kivernyk ²⁴,
 T. Klapdor-Kleingrothaus ⁵⁴, M. Klassen ^{63a}, C. Klein ³⁴, L. Klein ¹⁶⁸, M.H. Klein ¹⁰⁶,

M. Klein [ID⁹²](#), S.B. Klein [ID⁵⁶](#), U. Klein [ID⁹²](#), P.R. Klemm [ID¹⁰⁹](#), P. Klimek [ID³⁶](#), A. Klimentov [ID²⁹](#),
F. Klimpel [ID¹¹⁰](#), T. Klioutchnikova [ID³⁶](#), F.F. Klitzner [ID¹⁰⁹](#), P. Kluit [ID¹¹⁴](#), S. Kluth [ID¹¹⁰](#),
E. Kneringer [ID⁷⁹](#), T.M. Knight [ID¹⁵⁵](#), E.B.F.G. Knoops [ID¹⁰²](#), A. Knue [ID⁵⁴](#), D. Kobayashi [ID⁸⁹](#),
R. Kobayashi [ID⁸⁷](#), M. Kocian [ID¹⁴³](#), P. Kodyš [ID¹³³](#), D.M. Koeck [ID¹⁴⁶](#), P.T. Koenig [ID²⁴](#), T. Koffas [ID³⁴](#),
M. Kolb [ID¹³⁵](#), A. Kolbasin [ID³⁷](#), I. Koletsou [ID⁴](#), F. Kolitsi [ID¹⁶¹](#), T. Komarek [ID¹²²](#), S. Kompogiannis [ID¹⁵²](#),
K. Köneke [ID⁵⁴](#), A.X.Y. Kong [ID¹](#), M. Kongsore [ID¹⁰⁶](#), T. Kono [ID¹¹⁸](#), N. Konstantinidis [ID⁹⁶](#), B. Konya [ID⁹⁸](#),
R. Kopeliansky [ID⁶⁸](#), S. Koperny [ID^{85a}](#), K. Korcyl [ID⁸⁶](#), K. Kordas [ID^{152,e}](#), G. Koren [ID¹⁵¹](#), A. Korn [ID⁹⁶](#),
S. Korn [ID⁵⁵](#), I. Korolkov [ID¹³](#), N. Korotkova [ID³⁷](#), B. Kortman [ID¹¹⁴](#), O. Kortner [ID¹¹⁰](#), S. Kortner [ID¹¹⁰](#),
W.H. Kostecka [ID¹¹⁵](#), V.V. Kostyukhin [ID¹⁴¹](#), A. Kotsokechagia [ID¹³⁵](#), A. Kotwal [ID⁵¹](#), A. Koulouris [ID³⁶](#),
A. Kourkoumeli-Charalampidi [ID^{73a,73b}](#), C. Kourkoumelis [ID⁹](#), E. Kourlitis [ID⁶](#), G. Koutelieris [ID¹⁰](#),
T. Koutsosimos [ID¹⁵²](#), D.F. Kouyoumdjian [ID^{137f}](#), S. Kovalenko [ID¹²¹](#), O. Kovanda [ID¹⁴⁶](#), I. Koveshnikov [ID^{156a}](#),
R. Kowalewski [ID¹⁶⁷](#), W. Kozanecki [ID¹³⁵](#), A.S. Kozhin [ID³⁷](#), V.A. Kramarenko [ID³⁷](#), G. Kramberger [ID⁹³](#),
P. Kramer [ID¹⁰⁰](#), M.W. Krasny [ID¹²⁷](#), A. Krasznahorkay [ID³⁶](#), J.A. Kremer [ID¹⁰⁰](#), T. Kresse [ID⁵⁰](#),
J. Kretzschmar [ID⁹²](#), K. Kreul [ID¹⁸](#), N. Krieger [ID⁵⁵](#), P. Krieger [ID¹⁵⁵](#), S. Krishnamurthy [ID¹⁰³](#), M. Krivos [ID¹³³](#),
K. Krizka [ID^{17a}](#), K. Kroeninger [ID⁴⁹](#), H. Kroha [ID¹¹⁰](#), J. Kroll [ID¹³¹](#), J. Kroll [ID¹²⁸](#), K.S. Krowpman [ID¹⁰⁷](#),
U. Kruchonak [ID³⁸](#), H. Krüger [ID²⁴](#), N. Krumnack [ID⁸¹](#), M.C. Kruse [ID⁵¹](#), J.A. Krzysiak [ID⁸⁶](#), Z. Kubik [ID¹⁶⁴](#),
O. Kuchinskaia [ID³⁷](#), S. Kuday [ID^{3a}](#), D. Kuechler [ID⁴⁸](#), J.T. Kuechler [ID⁴⁸](#), S. Kuehn [ID³⁶](#), R. Kuesters [ID⁵⁴](#),
F. Kuger [ID¹⁶⁸](#), T. Kuhl [ID⁴⁸](#), V. Kukhtin [ID³⁸](#), Y. Kulchitsky [ID^{37,a}](#), S. Kuleshov [ID^{137d,137b}](#),
A.M. Kulinska [ID³⁶](#), P. Kulka [ID^{28b}](#), M. Kumar [ID^{33g}](#), N. Kumari [ID¹⁰²](#), B.M. Kunkler [ID⁶⁸](#), A. Kupco [ID¹³¹](#),
T. Kupfer [ID⁴⁹](#), A. Kupich [ID³⁷](#), J. Kuppambatti [ID^{aj,41}](#), O. Kuprash [ID⁵⁴](#), H. Kurashige [ID⁸⁴](#),
L.L. Kurchaninov [ID^{156a}](#), Y.A. Kurochkin [ID³⁷](#), A. Kurova [ID³⁷](#), M. Kuze [ID¹⁵⁴](#), A.K. Kvam [ID¹⁰³](#),
J. Kvita [ID¹²²](#), T. Kwan [ID¹⁰⁴](#), K.W. Kwok [ID^{64a}](#), N.G. Kyriacou [ID¹⁰⁶](#), E. Kyriakis-Bitaros [ID¹⁶¹](#),
L.A.O. Laatu [ID¹⁰²](#), C. Lacasta [ID¹⁶⁵](#), F. Lacava [ID^{75a,75b}](#), H. Lacker [ID¹⁸](#), D. Lacour [ID¹²⁷](#), N.N. Lad [ID⁹⁶](#),
E. Ladygin [ID³⁸](#), B. Laforge [ID¹²⁷](#), S. Lafrasse [ID⁴](#), T. Lagouri [ID^{137e}](#), S. Lai [ID⁵⁵](#), I.K. Lakomic [ID^{85a}](#),
N. Lalloue [ID⁶⁰](#), H.P. Lam [ID^{64c}](#), J.E. Lambert [ID¹²⁰](#), S. Lammers [ID⁶⁸](#), E. Lampardaki [ID¹⁰](#), W. Lampl [ID⁷](#),
C. Lampoudis [ID^{152,e}](#), A.N. Lancaster [ID¹¹⁵](#), E. Lançon [ID²⁹](#), U. Landgraf [ID⁵⁴](#), M.P.J. Landon [ID⁹⁴](#),
C. Landraud [ID³⁶](#), V.S. Lang [ID⁵⁴](#), R.J. Langenberg [ID¹⁰³](#), R.R. Langstaff [ID^{156a}](#), A.J. Lankford [ID¹⁶⁰](#),
F. Lanni [ID³⁶](#), K. Lantzs [ID²⁴](#), A. Lanza [ID^{73a}](#), A. Lapertosa [ID^{57b,57a}](#), D. Laporte [ID¹²⁷](#), J.F. Laporte [ID¹³⁵](#),
T. Lari [ID^{71a}](#), F. Lasagni Manghi [ID^{23b}](#), M. Lassnig [ID³⁶](#), V. Latonova [ID¹³¹](#), S.L. Latorre [ID^{71b}](#), T.S. Lau [ID^{64a}](#),
A. Laudrain [ID¹⁰⁰](#), D.M. Laugier [ID¹⁰²](#), A. Laurier [ID³⁴](#), S.D. Lawlor [ID⁹⁵](#), Z. Lawrence [ID¹⁰¹](#), I. Lax [ID^{23a,23b}](#),
M. Lazzaroni [ID^{71a,71b}](#), B. Le [ID¹⁰¹](#), F. Le Goff [ID³⁶](#), X. Le [ID⁴⁴](#), P. Le Boulout [ID¹³⁵](#), B. Leban [ID⁹³](#), A. Lebedev [ID⁸¹](#),
M. LeBlanc [ID³⁶](#), D. Leboeuf [ID¹³⁵](#), C. Leboube [ID³⁶](#), T. LeCompte [ID⁶](#), F. Ledroit-Guillon [ID⁶⁰](#), A.C.A. Lee [ID⁹⁶](#),
C.A. Lee [ID²⁹](#), G.R. Lee [ID¹⁶](#), L. Lee [ID⁶¹](#), S.C. Lee [ID¹⁴⁸](#), S. Lee [ID^{47a,47b}](#), T.F. Lee [ID⁹²](#), L.L. Leeuw [ID^{33c}](#),
B. Lefebvre [ID^{156a}](#), H.P. Lefebvre [ID⁹⁵](#), M. Lefebvre [ID¹⁶⁷](#), M. Lefevre [ID¹³⁵](#), C. Leggett [ID^{17a}](#),
K. Lehmann [ID¹⁴²](#), G. Lehmann Miotto [ID³⁶](#), M. Leigh [ID⁵⁶](#), W.A. Leight [ID¹⁰³](#), U. Leis [ID¹¹⁰](#), A. Leisos [ID^{152,t}](#),
M.A.L. Leite [ID^{82c}](#), C.E. Leitgeb [ID⁴⁸](#), R. Leitner [ID¹³³](#), M.S. Lenckowski [ID^{156a}](#), K.J.C. Leney [ID⁴⁴](#),
T. Lenz [ID²⁴](#), S. Leone [ID^{74a}](#), C. Leonidopoulos [ID⁵²](#), A. Leopold [ID¹⁴⁴](#), T.J. Lepota [ID^{33g}](#), C. Leroy [ID¹⁰⁸](#),
R. Les [ID¹⁰⁷](#), C.G. Lester [ID³²](#), H.K. Leung [ID^{64a}](#), M. Levchenko [ID³⁷](#), J. Levêque [ID⁴](#), D. Levin [ID¹⁰⁶](#),
L.J. Levinson [ID¹⁷¹](#), M.P. Lewicki [ID⁸⁶](#), D.J. Lewis [ID⁴](#), A. Li [ID⁵](#), B. Li [ID^{62b}](#), C. Li [ID^{62a}](#), C. Li [ID^{62b}](#),
C-Q. Li [ID^{62c}](#), F. Li [ID^{62a}](#), H. Li [ID^{62a}](#), H. Li [ID^{62b}](#), H. Li [ID^{14c}](#), H. Li [ID^{62b}](#), J. Li [ID^{62c}](#), K. Li [ID¹³⁸](#),
L. Li [ID^{62c}](#), M. Li [ID^{14a,14d}](#), Q.Y. Li [ID^{62a}](#), S. Li [ID^{14a,14d}](#), S. Li [ID^{62d,62c,d}](#), T. Li [ID^{62b}](#), X. Li [ID¹⁰⁴](#), X. Li [ID⁴⁴](#),
Z. Li [ID^{62b}](#), Z. Li [ID¹²⁶](#), Z. Li [ID¹⁰⁴](#), Z. Li [ID⁹²](#), Z. Li [ID^{14a,14d}](#), F. Liang [ID⁴⁴](#), Z. Liang [ID^{14a}](#), S. Liaqat [ID³⁶](#),
M. Liberatore [ID⁴⁸](#), B. Liberti [ID^{76a}](#), K. Lie [ID^{64c}](#), J. Lieber Marin [ID^{82b}](#), H. Lien [ID⁶⁸](#), C. Lim [ID^{156a}](#),
K. Lin [ID¹⁰⁷](#), T.H. Lin [ID¹⁰⁰](#), R.A. Linck [ID⁶⁸](#), R.E. Lindley [ID⁷](#), J.H. Lindon [ID²](#), A. Linss [ID⁴⁸](#),
E. Lipeles [ID¹²⁸](#), A. Lipniacka [ID¹⁶](#), M.W. Lippert [ID¹¹⁰](#), A. Lister [ID¹⁶⁶](#), J.D. Little [ID⁴](#), B. Liu [ID^{14a}](#),
B.X. Liu [ID¹⁴²](#), C. Liu [ID⁴⁴](#), D. Liu [ID^{62d,62c}](#), H.B. Liu [ID²⁹](#), J.B. Liu [ID^{62a}](#), J.K.K. Liu [ID³²](#), K. Liu [ID^{62d,62c}](#),
M. Liu [ID^{62a}](#), M.Y. Liu [ID^{62a}](#), P. Liu [ID^{14a}](#), Q. Liu [ID^{62d,138,62c}](#), S. Liu [ID^{62a}](#), T. Liu [ID²⁹](#), X. Liu [ID^{62a}](#),

Y. Liu ^{14c,14d}, Y.L. Liu ¹⁰⁶, Y.W. Liu ^{62a}, D. Liubimtcev³⁸, M. Livan ^{73a,73b}, M.P. Liz Vargas^{137d},
 J. Llorente Merino ¹⁴², S.L. Lloyd ⁹⁴, E.M. Lobodzinska ⁴⁸, P. Loch ⁷, S. Loffredo ^{76a,76b},
 T. Lohse ¹⁸, K. Lohwasser ¹³⁹, C.A. Loiseau¹³⁵, M. Lokajicek ¹³¹, J.D. Long ¹⁶³, I. Longarini ¹⁶⁰,
 L. Longo ^{70a,70b}, R. Longo ¹⁶³, D. Lopez Mateos⁶¹, I. Lopez Paz ⁶⁷, A. Lopez Solis ⁴⁸,
 J. Lorenz ¹⁰⁹, N. Lorenzo Martinez ⁴, A.M. Lory ¹⁰⁹, P.J. Lösel¹⁰⁹, X. Lou ^{47a,47b}, X. Lou ^{14a,14d},
 A. Lounis ⁶⁶, J. Love ⁶, P.A. Love ⁹¹, J.J. Lozano Bahilo ¹⁶⁵, G. Lu ^{14a,14d}, M. Lu ⁸⁰,
 S. Lu ¹²⁸, Y.J. Lu ⁶⁵, H.J. Lubatti ¹³⁸, C. Luci ^{75a,75b}, F.L. Lucio Alves ^{14c}, A. Lucotte ⁶⁰,
 F. Luehring ⁶⁸, I. Luise ¹⁴⁵, O. Lukianchuk ⁶⁶, L. Luminari^{75b}, R. Lunadei ^{75b}, O. Lundberg ¹⁴⁴,
 B. Lund-Jensen ¹⁴⁴, N.A. Luongo ¹²³, N. Lupu¹⁵⁰, M.S. Lutz ¹⁵¹, R.J. Luz ⁶, D. Lynn ²⁹,
 H. Lyons⁹², R. Lysak ¹³¹, E. Lytken ⁹⁸, F. Lyu ^{14a}, V. Lyubushkin ³⁸, T. Lyubushkina ³⁸,
 M.M. Lyukova ¹⁴⁵, H. Ma ²⁹, L.L. Ma ^{62b}, Y. Ma ⁹⁶, D.M. Mac Donell ¹⁶⁷, G. Maccarrone ⁵³,
 J.C. MacDonald ¹³⁹, R. Madar ⁴⁰, W.F. Mader ⁵⁰, K. Madhoun ³⁴, J. Maeda ⁸⁴, T. Maeno ²⁹,
 M. Maerker ⁵⁰, H. Maguire ¹³⁹, D.J. Mahon ⁴¹, R. Maier¹¹⁰, A. Maio ^{130a,130b,130d}, K. Maj ^{85a},
 O. Majersky ^{28a}, S. Majewski ¹²³, N. Makovec ⁶⁶, V. Maksimovic ¹⁵, B. Malaescu ¹²⁷,
 J.A. Malaquin³⁶, Pa. Malecki ⁸⁶, V.P. Maleev ³⁷, F. Malek ⁶⁰, D. Malito ^{43b,43a}, U. Mallik ⁸⁰,
 C. Malone ³², S. Maltezos¹⁰, P. Maly¹³¹, S. Malyukov³⁸, J. Mamuzic ¹³, F.M. Manca^{71b},
 G. Mancini ⁵³, G. Manco ^{73a,73b}, J.P. Mandalia ⁹⁴, I. Mandić ⁹³, I. Mandjavidze¹³⁵,
 L. Manhaes de Andrade Filho ^{82a}, I.M. Maniatis ¹⁷¹, M. Manisha ¹³⁵, J. Manjarres Ramos ⁵⁰,
 D.C. Mankad ¹⁷¹, A. Mann ¹⁰⁹, A. Manousos ⁷⁹, S.M. Manson ^{156a}, B. Mansoulie ¹³⁵,
 I. Mantos ¹⁵², S. Manzoni ³⁶, E. Maragkou Didi¹⁵², A. Marantis ¹⁵², G. Marchiori ⁵,
 M. Marcisovsky ¹³¹, C. Marcon ^{71a,71b}, M. Marinescu ²⁰, M. Marjanovic ¹²⁰, E.J. Marshall ⁹¹,
 Z. Marshall ^{17a}, S. Marti-Garcia ¹⁶⁵, T.A. Martin ¹⁶⁹, V.J. Martin ⁵², B. Martin dit Latour ¹⁶,
 L. Martinelli ^{75a,75b}, M. Martinez ^{13,u}, P. Martinez Agullo ¹⁶⁵, V.I. Martinez Outschoorn ¹⁰³,
 P. Martinez Suarez ¹³, S. Martin-Haugh ¹³⁴, V.S. Martoiu ^{27b}, A.C. Martyniuk ⁹⁶, A. Marzin ³⁶,
 P. Mas¹³⁵, S.R. Maschek ¹¹⁰, D. Mascione ^{78a,78b}, L. Masetti ¹⁰⁰, T. Mashimo ¹⁵³, J. Masik ¹⁰¹,
 A.L. Maslennikov ³⁷, L. Massa ^{23b}, P. Massarotti ^{72a,72b}, N. Massol⁴, P. Mastrandrea ^{74a,74b},
 A. Mastroberardino ^{43b,43a}, T. Masubuchi ¹⁵³, D. Matakias²⁹, T. Mathisen ¹⁶², N. Matsuzawa¹⁵³,
 P. Mättig ²⁴, J. Maurer ^{27b}, B. Maček ⁹³, D.A. Maximov ³⁷, R. Mazini ¹⁴⁸, I. Maznas ^{152,e},
 M. Mazza ¹⁰⁷, S.M. Mazza ¹³⁶, C. Mc Ginn ²⁹, J.P. Mc Gowan ¹⁰⁴, S.P. Mc Kee ¹⁰⁶,
 C.C. McCracken ¹⁶⁶, E.F. McDonald ¹⁰⁵, A.E. McDougall ¹¹⁴, J.A. Mcfayden ¹⁴⁶,
 R.P. McGovern ¹²⁸, G. Mchedlidze ^{149b}, R.P. Mckenzie ^{33g}, T.C. McLachlan ⁴⁸,
 D.J. McLaughlin ⁹⁶, K.D. McLean ¹⁶⁷, S.J. McMahon ¹³⁴, P.C. McNamara ¹⁰⁵,
 C.M. Mcpartland ⁹², R.A. McPherson ^{167,x}, T. Megy ⁴⁰, I. Mehalev ¹⁵⁰, S. Mehlhase ¹⁰⁹,
 A. Mehta ⁹², B. Meirose ⁴⁵, D. Melini ¹⁵⁰, B.R. Mellado Garcia ^{33g}, A.H. Melo ⁵⁵, F. Meloni ⁴⁸,
 E.D. Mendes Gouveia ^{130a}, A.M. Mendes Jacques Da Costa ²⁰, S. Meneghini^{23a}, H.Y. Meng ¹⁵⁵,
 L. Meng ⁹¹, S. Menke ¹¹⁰, M. Menouni¹⁰², M. Mentink ³⁶, E. Meoni ^{43b,43a}, C. Merlassino ¹²⁶,
 L. Merola ^{72a,72b}, C. Meroni ^{71a}, G. Merz¹⁰⁶, O. Meshkov ³⁷, I. Mesolongitis ¹⁶¹, J. Metcalfe ⁶,
 A.S. Mete ⁶, C. Meyer ⁶⁸, J-P. Meyer ¹³⁵, P. Miao^{62a}, A. Miccoli^{70b}, S. Michal⁵⁶, M. Michetti ¹⁸,
 R.P. Middleton ¹³⁴, S. Miglioranzi ^{57a}, J. Migne³⁶, L. Mijović ⁵², G. Mikenberg ¹⁷¹,
 M. Mikestikova ¹³¹, M. Mikuž ⁹³, H. Mildner ¹³⁹, A. Milic ³⁶, C.D. Milke ⁴⁴, D.W. Miller ³⁹,
 L.S. Miller ³⁴, A. Milov ¹⁷¹, M. Milovanovic ⁴⁸, D.A. Milstead^{47a,47b}, T. Min^{14c},
 A.A. Minaenko ³⁷, Y. Minami ¹⁵³, M. Miñano Moya ¹⁰³, I.A. Minashvili ^{149b}, L. Mince ⁵⁹,
 A.I. Mincer ¹¹⁷, B. Mindur ^{85a}, M. Mineev ³⁸, Y. Mino ⁸⁷, L.M. Mir ¹³, M. Miralles Lopez ¹⁶⁵,
 M. Mironova ¹²⁶, M.C. Missio ¹¹³, T. Mitani ¹⁷⁰, A. Mitra ¹⁶⁹, V.A. Mitsou ¹⁶⁵, O. Miu ¹⁵⁵,
 P.S. Miyagawa ⁹⁴, Y. Miyazaki⁸⁹, A. Mizukami ⁸³, J.U. Mjörnmark ⁹⁸, T. Mkrtchyan ^{63a},
 G. Mladenovic ¹⁵, T. Mlinarevic ⁹⁶, M. Mlynarikova ³⁶, T. Moa ^{47a,47b}, S. Mobius ⁵⁵,
 K. Mochizuki ¹⁰⁸, P. Moder ⁴⁸, P. Mogg ¹⁰⁹, A.F. Mohammed ^{14a,14d}, S. Mohapatra ⁴¹,

G. Mokgatitswane ^{id33g}, S. Mokrenko³⁸, L. Moleri ^{id171}, E.J. Molina Gonzalez¹³⁵, A. Monay E Silva^{82a}, B. Mondal ^{id141}, S. Mondal ^{id132}, K. Mönig ^{id48}, E. Monnier ^{id102}, L. Monsonis Romero¹⁶⁵, J. Montejó Berlingen ^{id36}, M. Montella ^{id119}, M.M. Monti ^{id71b}, F. Monticelli ^{id90}, J.P.A. Moraga Jimenez⁵², N. Morange ^{id66}, A.L. Moreira De Carvalho ^{id130a}, M. Moreno Llácer ^{id165}, C. Moreno Martinez ^{id56}, P. Morettini ^{id57b}, S. Morgenstern ^{id169}, M. Morii ^{id61}, M. Morinaga ^{id153}, A.K. Morley ^{id36}, F. Morodei ^{id75a,75b}, L. Morvaj ^{id36}, P. Moschovakos ^{id36}, B. Moser ^{id36}, M. Mosidze^{149b}, T. Moskalets ^{id54}, P. Moskvitina ^{id113}, J. Moss ^{id31.o}, P. Moszkowicz ^{id85a}, E.J.W. Moyse ^{id103}, O. Mtintsilana ^{id33g}, S. Muanza ^{id102}, J. Mueller ^{id129}, R.S.P. Mueller¹⁰⁹, D. Muenstermann ^{id91}, R. Müller ^{id19}, G.A. Mullier ^{id162}, J.J. Mullin¹²⁸, A.E. Mulski ^{id61}, D.P. Mungo ^{id155}, J.L. Munoz Martinez ^{id13}, D. Munoz Perez ^{id165}, F.J. Munoz Sanchez ^{id101}, M. Mur¹³⁵, M. Murin ^{id101}, W.J. Murray ^{id169,134}, A. Murrone ^{id71a,71b}, J.M. Muse ^{id120}, M. Muškinja ^{id17a}, C. Mwewa ^{id29}, A.G. Myagkov ^{id37,a}, A.J. Myers ^{id8}, A.A. Myers¹²⁹, G. Myers ^{id68}, M. Myska ^{id132}, B.P. Nachman ^{id17a}, O. Nackenhorst ^{id49}, M. Naeem³⁶, A. Nag ^{id50}, K. Nagai ^{id126}, K. Nagano ^{id83}, J.L. Nagle ^{id29,ah}, E. Nagy ^{id102}, A.M. Nairz ^{id36}, Y. Nakahama ^{id83}, K. Nakamura ^{id83}, H. Nanjo ^{id124}, R. Narayan ^{id44}, E.A. Narayanan ^{id112}, J. Narevicius¹⁷¹, L.L. Narvaez Paredes^{137f}, I. Naryshkin ^{id37}, M. Naseri ^{id34}, C. Nass ^{id24}, M.D. Natsios¹⁰, G. Navarro ^{id22a}, J. Navarro-Gonzalez ^{id165}, R. Nayak ^{id151}, A. Nayaz ^{id18}, P.Y. Nechaeva ^{id37}, F. Nechansky ^{id48}, L. Nedic ^{id126}, T.J. Neep ^{id20}, A. Negri ^{id73a,73b}, M. Negrini ^{id23b}, C. Nellist ^{id113}, C. Nelson ^{id104}, K. Nelson ^{id106}, S. Nemecek ^{id131}, M. Nessi ^{id36,h}, M.S. Neubauer ^{id163}, F. Neuhaus ^{id100}, J. Neundorf ^{id48}, R. Newhouse ^{id166}, P.R. Newman ^{id20}, C.W. Ng ^{id129}, Y.S. Ng¹⁸, Y.W.Y. Ng ^{id48}, B. Ngair ^{id35e}, H.D.N. Nguyen ^{id108}, R.B. Nickerson ^{id126}, R. Nicolaidou ^{id135}, J. Nielsen ^{id136}, M. Niemeyer ^{id55}, N. Nikiforou ^{id36}, V. Nikolaenko ^{id37,a}, I. Nikolic-Audit ^{id127}, K. Nikolopoulos ^{id20}, M. Nila¹⁰⁷, P. Nilsson ^{id29}, I. Ninca ^{id48}, H.R. Nindhito ^{id56}, A. Nisati ^{id75a}, N. Nishu ^{id2}, R. Nisius ^{id110}, J-E. Nitschke ^{id50}, E.K. Nkadimeng ^{id33g}, S.J. Noacco Rosende ^{id90}, T. Nobe ^{id153}, D.L. Noel ^{id32}, J. Noel³⁶, Y. Noguchi ^{id87}, T. Nommensen ^{id147}, M.A. Nomura²⁹, M.B. Norfolk ^{id139}, R.R.B. Norisam ^{id96}, B.J. Norman ^{id34}, J. Novak ^{id93}, T. Novak ^{id48}, O. Novgorodova ^{id50}, L. Novotny ^{id132}, R. Novotny ^{id112}, L. Nozka ^{id122}, K. Ntekas ^{id160}, N.M.J. Nunes De Moura Junior ^{id82b}, E. Nurse⁹⁶, F.G. Oakham ^{id34,ae}, J. Ocariz ^{id127}, A. Ochi ^{id84}, I. Ochoa ^{id130a}, W. Ockenfels²⁴, R. Oehm¹⁰⁹, S. Oerdek ^{id162}, J.T. Offermann ^{id39}, A. Ogradnik ^{id85a}, A. Oh ^{id101}, C.C. Ohm ^{id144}, H. Oide ^{id83}, K. Oikonomou¹⁵², R. Oishi ^{id153}, M.L. Ojeda ^{id48}, Y. Okazaki ^{id87}, M.W. O'Keefe⁹², Y. Okumura ^{id153}, A. Olariu^{27b}, L.F. Oleiro Seabra ^{id130a}, S.A. Olivares Pino ^{id137e}, D. Oliveira Damazio ^{id29}, D. Oliveira Goncalves ^{id82a}, J.L. Oliver ^{id160}, M.J.R. Olsson ^{id160}, A. Olszewski ^{id86}, J. Olszowska ^{id86,*}, Ö.O. Öncel ^{id54}, D.C. O'Neil ^{id142}, A.P. O'Neill ^{id19}, A. Onofre ^{id130a,130e}, P.U.E. Onyisi ^{id11}, R. Openshaw^{156a}, M.J. Oreglia ^{id39}, G.E. Orellana ^{id90}, D. Orestano ^{id77a,77b}, N. Orlando ^{id13}, R.S. Orr ^{id155}, V. O'Shea ^{id59}, R. Ospanov ^{id62a}, M.S. Ostrega ^{id36}, G. Otero y Garzon ^{id30}, H. Otono ^{id89}, P.S. Ott ^{id63a}, G.J. Ottino ^{id17a}, M. Ouchrif ^{id35d}, J. Ouellette ^{id29,ah}, F. Ould-Saada ^{id125}, M. Owen ^{id59}, R.E. Owen ^{id134}, K.Y. Oyulmaz ^{id21a}, A. Ozbey ^{id21a}, V.E. Ozcan ^{id21a}, N. Ozturk ^{id8}, S. Ozturk ^{id21d}, J. Pacalt ^{id122}, H.A. Pacey ^{id32}, K. Pachal ^{id51}, A. Pacheco Pages ^{id13}, C. Padilla Aranda ^{id13}, G. Padovano ^{id75a,75b}, S. Pagan Griso ^{id17a}, G. Palacino ^{id68}, A. Palazzo ^{id70a,70b}, S. Palestini ^{id36}, M. Palka ^{id85b}, J. Pan ^{id174}, T. Pan ^{id64a}, C. Pancake¹⁴⁵, D.K. Panchal ^{id11}, C.E. Pandini ^{id114}, J.G. Panduro Vazquez ^{id95}, H. Pang ^{id14b}, P. Pangaud¹⁰², P. Pani ^{id48}, L. Panico^{72b}, G. Panizzo ^{id69a,69c}, L. Paolozzi ^{id56}, C. Papadatos ^{id108}, S. Parajuli ^{id44}, A. Paramonov ^{id6}, C. Paraskevopoulos ^{id10}, D. Paredes Hernandez ^{id64b}, T.H. Park ^{id155}, M.A. Parker ^{id32}, F. Parodi ^{id57b,57a}, E.W. Parrish ^{id115}, V.A. Parrish ^{id52}, J.A. Parsons ^{id41}, G. Paruzza^{77a,77b}, U. Parzefall ^{id54}, P. Paschalias¹⁵², B. Pascual Dias ^{id108}, L. Pascual Dominguez ^{id151}, V.R. Pascuzzi ^{id17a}, B. Pasmantirer¹⁷¹, F. Pasquali ^{id114}, E. Pasqualucci ^{id75a}, S. Passaggio ^{id57b}, F. Pastore ^{id95}, E. Pastori ^{id76b}, P. Pasuwan ^{id47a,47b}, P. Patel ^{id86}, J.R. Pater ^{id101}, T. Pauly ^{id36}, C.I. Pazos ^{id158}, J. Parkes ^{id143},

M. Pedersen ¹²⁵, R. Pedro ^{130a}, S.V. Peleganchuk ³⁷, A. Pelosi ^{75b}, O. Penc ³⁶, E.A. Pender ⁵²,
C. Peng ^{64b}, H. Peng ^{62a}, K.E. Penski ¹⁰⁹, M. Penzin ³⁷, M. Pepe ^{70a,70b}, B.S. Peralva ^{82d,82d},
M.C. Pereira ^{156a}, A.P. Pereira Peixoto ⁶⁰, L. Pereira Sanchez ^{47a,47b}, D.V. Perepelitsa ^{29,ah},
E. Perez Codina ^{156a}, F. Perez Gomez ³⁶, M. Perganti ¹⁰, L. Perini ^{71a,71b,*}, H. Pernegger ³⁶,
S. Perrella ³⁶, A. Perrevoort ¹¹³, O. Perrin ⁴⁰, G. Perrot ⁴, K. Peters ⁴⁸, R.F.Y. Peters ¹⁰¹,
B.A. Petersen ³⁶, T.C. Petersen ⁴², E. Petit ¹⁰², V. Petousis ¹³², C. Petridou ^{152,e},
M. Petrucci ^{75a,75b}, F. Petrucci ^{77b}, A. Petrukhin ¹⁴¹, M. Pettee ^{17a}, N.E. Pettersson ³⁶,
A. Petukhov ³⁷, K. Petukhova ¹³³, A. Peyaud ¹³⁵, R. Pezoa ^{137f}, L. Pezzotti ³⁶, G. Pezzullo ¹⁷⁴,
B. Pfeifer ⁵⁴, T.M. Pham ¹⁷², T. Pham ¹⁰⁵, P.W. Phillips ¹³⁴, M.W. Phipps ¹⁶³, G. Piacquadio ¹⁴⁵,
E. Pianori ^{17a}, F. Piazza ^{71a,71b}, R. Piegai ³⁰, D. Pietreanu ^{27b}, G. Pileggi ⁵³, A.D. Pilkington ¹⁰¹,
M. Pinamonti ^{69a,69c}, J.L. Pinfold ², B.C. Pinheiro Pereira ^{130a}, J.T. Pinnell ³⁶, R.A. Pinto Bustos ^{137a},
R.N. Pirea ^{27e}, Y. Piret ¹³⁵, M. Pirola ^{73b}, F. Piscitelli ^{76b}, C. Pitman Donaldson ⁹⁶, D.A. Pizzi ³⁴,
L. Pizzimento ^{76a,76b}, A. Pizzini ¹¹⁴, M.-A. Pleier ²⁹, V. Plesanovs ⁵⁴, V. Pleskot ¹³³, E. Plotnikova ³⁸,
A. Pluzhnikov ³⁸, G. Poddar ⁴, S. Podkladkin ¹¹⁰, R. Poettgen ⁹⁸, P. Poffenberger ¹⁶⁷, L. Poggioli ¹²⁷,
I. Pogrebnyak ¹⁰⁷, D. Pohl ²⁴, I. Pokharel ⁵⁵, S. Polacek ¹³³, G. Polesello ^{73a}, A. Poley ^{142,156a},
R. Polifka ¹³², A. Polini ^{23b}, E. Politis ¹⁶¹, C.S. Pollard ¹⁶⁹, Z.B. Pollock ¹¹⁹, V. Polychronakos ²⁹,
E. Pompa Pacchi ^{75a,75b}, D. Ponomarenko ³⁷, P. Ponsot ¹³⁵, L. Pontecorvo ³⁶, G. Pontoriere ^{72a,72b},
B. Ponzio ⁵³, S. Popa ^{27a}, G.A. Popeneciu ^{27d}, R.D. Porter ¹⁶⁷, D.M. Portillo Quintero ^{156a},
S. Pospisil ¹³², P. Postolache ^{27c}, K. Potamianos ¹²⁶, I.N. Potrap ³⁸, C.J. Potter ³², H. Potti ¹,
T. Poulsen ⁴⁸, J. Poveda ¹⁶⁵, M.E. Pozo Astigarraga ³⁶, A. Prades Ibanez ¹⁶⁵,
P.B. Pranav Bhagawath Prasad ¹⁶⁷, M.M. Prapa ⁴⁶, J. Pretel ⁵⁴, D. Price ¹⁰¹, M. Primavera ^{70a},
M.A. Principe Martin ⁹⁹, R. Privara ¹²², M.L. Proffitt ¹³⁸, N. Proklova ¹²⁸, K. Prokofiev ^{64c},
G. Prono ¹³⁵, G. Proto ^{76a,76b}, S. Protopopescu ²⁹, J. Proudfoot ⁶, M. Przybycien ^{85a},
J.E. Puddefoot ¹³⁹, D. Pudzha ³⁷, P. Puzo ⁶⁶, D. Pyatizbyantseva ³⁷, J. Qian ¹⁰⁶, W. Qian ¹³⁴,
D. Qichen ¹⁰¹, Y. Qin ¹⁰¹, T. Qiu ⁹⁴, A. Quadt ⁵⁵, M. Queitsch-Maitland ¹⁰¹, L. Quercia ³⁶,
G. Quetant ⁵⁶, G. Rabanal Bolanos ⁶¹, J. Rabel ¹²⁹, D. Rafanoharana ⁵⁴, F. Ragusa ^{71a,71b},
J.L. Rainbolt ³⁹, J.A. Raine ⁵⁶, S. Rajagopalan ²⁹, E. Ramakoti ³⁷, I.A. Ramirez-berend ³⁴,
K. Ran ^{48,14d}, N.P. Rapheeha ^{33g}, T. Rashid ⁶⁶, V. Raskina ¹²⁷, D.F. Rassloff ^{63a}, S. Rave ¹⁰⁰,
B. Ravina ⁵⁵, I. Ravinovich ¹⁷¹, M. Raymond ³⁶, A.L. Read ¹²⁵, N.P. Readioff ¹³⁹,
D.M. Rebuzzi ^{73a,73b}, G. Redlinger ²⁹, K. Reeves ⁴⁵, J.A. Reidelsturz ¹⁷³, D. Reikher ¹⁵¹,
A. Rej ¹⁴¹, C. Rembser ³⁶, A. Renardi ⁴⁸, M. Renda ^{27b}, M.B. Rendel ¹¹⁰, A. Renklioglu ^{21a},
F. Renner ⁴⁸, A.G. Rennie ⁵⁹, S. Resconi ^{71a}, M. Ressegotti ^{57b,57a}, E.D. Resseguie ^{17a},
S. Rettie ³⁶, J.G. Reyes Rivera ¹⁰⁷, B. Reynolds ¹¹⁹, E. Reynolds ^{17a}, M. Rezaei Estabragh ¹⁷³,
O.L. Rezanova ³⁷, P. Reznicek ¹³³, M. Riallot ¹³⁵, N. Ribaric ⁹¹, E. Ricci ^{78a,78b}, W. Richert ^{156a},
R. Richter ¹¹⁰, S. Richter ^{47a,47b}, E. Richter-Was ^{85b}, M. Ridel ¹²⁷, S. Ridouani ^{35d}, P. Rieck ¹¹⁷,
P. Riedler ³⁶, C.J. Riegel ¹⁷³, M. Rijssenbeek ¹⁴⁵, A. Rimoldi ^{73a,73b}, M. Rimoldi ⁴⁸,
L. Rinaldi ^{23b,23a}, T.T. Rinn ²⁹, M.P. Rinnagel ¹⁰⁹, G. Ripellino ¹⁴⁴, I. Riu ¹³, P. Rivadeneira ⁴⁸,
J.C. Rivera Vergara ¹⁶⁷, F. Rizatdinova ¹²¹, E. Rizvi ⁹⁴, C. Rizzi ⁵⁶, B.A. Roberts ¹⁶⁹,
B.R. Roberts ^{17a}, S.H. Robertson ^{104,x}, A. Robichaud-Veronneau ¹⁰⁴, M. Robin ⁴⁸,
D. Robinson ³², C.M. Robles Gajardo ^{137f}, M. Robles Manzano ¹⁰⁰, A. Robson ⁵⁹, A. Rocchi ^{76a,76b},
C. Roda ^{74a,74b}, S. Rodriguez Bosca ^{63a}, Y. Rodriguez Garcia ^{22a}, A. Rodriguez Rodriguez ⁵⁴,
A.M. Rodríguez Vera ^{156b}, S. Roe ³⁶, J.T. Roemer ¹⁶⁰, A.R. Roepe-Gier ¹²⁰, J. Roggel ¹⁷³,
O. Røhne ¹²⁵, A. Roich ¹⁷¹, R.A. Rojas ¹⁰³, B. Roland ⁵⁴, C.P.A. Roland ⁶⁸, J. Roloff ²⁹,
A. Romaniouk ³⁷, E. Romano ^{73a,73b}, M. Romano ^{23b}, A.C. Romero Hernandez ¹⁶³,
N. Rompotis ⁹², L. Roos ¹²⁷, S. Rosati ^{75a}, L. Roscilli ^{72a,72b}, L. Rose-Dulcina ³⁶, S. Röbl ¹⁰⁹,
B.J. Rosser ³⁹, C. Rossi ^{57a}, E. Rossi ⁴, E. Rossi ^{72a,72b}, F. Rossi ¹³⁵, L.P. Rossi ^{57b},
L. Rossini ⁴⁸, R. Rosten ¹¹⁹, M. Rotaru ^{27b}, B. Rottler ⁵⁴, C. Rougier ¹⁰², D. Rousseau ⁶⁶,

D. Rousso ³², A.R. Rovani ^{57a}, G. Rovelli ^{73a,73b}, A. Roy ¹⁶³, A. Rozanov ¹⁰², Y. Rozen ¹⁵⁰,
 X. Ruan ^{33g}, A. Rubio Jimenez ¹⁶⁵, A.J. Ruby ⁹², V.H. Ruelas Rivera ¹⁸, T.A. Ruggeri ¹,
 A. Ruggieri ^{75b}, D. Ruggieri ^{75b}, F. Rühr ⁵⁴, A. Ruiz-Martinez ¹⁶⁵, A. Rummler ³⁶, Z. Rurikova ⁵⁴,
 N.A. Rusakovich ³⁸, E.R. Ruscino ^{57a}, H.L. Russell ¹⁶⁷, J.P. Rutherford ⁷, O. Růžička ¹⁶⁴,
 K. Rybacki ⁹¹, M. Rybar ¹³³, E.B. Rye ¹²⁵, V. Ryjov ³⁶, A. Ryzhov ³⁷, J.A. Sabater Iglesias ⁵⁶,
 F.S. Sabatini ^{71b}, P. Sabatini ¹⁶⁵, L. Sabetta ^{75a,75b}, H.F-W. Sadrozinski ¹³⁶, F. Safai Tehrani ^{75a},
 B. Safarzadeh Samani ¹⁴⁶, M. Safdari ¹⁴³, S. Saha ¹⁰⁴, M. Sahinsoy ¹¹⁰, M. Saimpert ¹³⁵,
 M. Saito ¹⁵³, T. Saito ¹⁵³, M. Sajid ³⁶, D. Salamani ³⁶, G. Salamanna ^{77a,77b}, A. Salnikov ¹⁴³,
 F. Salomon ¹⁰², J. Salt ¹⁶⁵, A. Salvador Salas ¹³, D. Salvatore ^{43b,43a}, F. Salvatore ¹⁴⁶,
 A. Salzburger ³⁶, J. Samarati ³⁶, D. Sammel ⁵⁴, D. Sampsonidis ^{152,e}, D. Sampsonidou ^{62d,62c},
 J. Sánchez ¹⁶⁵, A. Sanchez Pineda ⁴, V. Sanchez Sebastian ¹⁶⁵, H. Sandaker ¹²⁵, C.O. Sander ⁴⁸,
 J.A. Sandesara ¹⁰³, M. Sandhoff ¹⁷³, C. Sandoval ^{22b}, D.P.C. Sankey ¹³⁴, B. Sanny ¹⁷³, T. Sano ⁸⁷,
 A. Sansoni ⁵³, L. Santi ^{75a,75b}, C. Santoni ⁴⁰, H. Santos ^{130a,130b}, S.N. Santpur ^{17a}, A. Santra ¹⁷¹,
 K.A. Saoucha ¹³⁹, J.G. Saraiva ^{130a,130d}, J. Sardain ⁷, O. Sasaki ⁸³, K. Sato ¹⁵⁷,
 T.P. Satterthwaite ³², R. Satzowski ¹⁰⁹, C. Sauer ^{63b}, F. Sauerburger ⁵⁴, E. Sauvan ⁴, R.M. Sauve ^{156a},
 P. Savard ^{155,ae}, R. Sawada ¹⁵³, C. Sawyer ¹³⁴, L. Sawyer ⁹⁷, I. Sayago Galvan ¹⁶⁵, C. Sbarra ^{23b},
 A. Sbrizzi ^{23b,23a}, C. Scagliotti ^{73b}, T. Scanlon ⁹⁶, J. Schaarschmidt ¹³⁸, P. Schacht ¹¹⁰,
 D. Schaefer ³⁹, U. Schäfer ¹⁰⁰, A.C. Schaffer ^{66,44}, D. Schaile ¹⁰⁹, O. Schaile ¹⁰⁹,
 R.D. Schamberger ¹⁴⁵, E. Schanet ¹⁰⁹, C. Scharf ¹⁸, M.M. Schefer ¹⁹, V.A. Schegelsky ³⁷,
 D. Scheirich ¹³³, F. Schenck ¹⁸, L. Scherino ³⁶, M. Schernau ¹⁶⁰, C. Scheulen ⁵⁵,
 C. Schiavi ^{57b,57a}, Z.M. Schillaci ²⁶, E.J. Schioppa ^{70a,70b}, M. Schioppa ^{43b,43a}, B. Schlag ¹⁰⁰,
 K.E. Schleicher ⁵⁴, S. Schlenker ³⁶, J. Schmeing ¹⁷³, M.A. Schmidt ¹⁷³, K. Schmieden ¹⁰⁰,
 C. Schmitt ¹⁰⁰, S. Schmitt ⁴⁸, R.M. Schnarr ³⁴, L. Schoeffel ¹³⁵, A. Schoening ^{63b}, P.G. Scholer ⁵⁴,
 E. Schopf ¹²⁶, A.L.S. Schorlemmer ⁴⁹, M. Schott ¹⁰⁰, J. Schovancova ³⁶, S. Schramm ⁵⁶,
 F. Schroeder ¹⁷³, H-C. Schultz-Coulon ^{63a}, J. Schumacher ¹¹³, M. Schumacher ⁵⁴,
 B.A. Schumm ¹³⁶, Ph. Schune ¹³⁵, H.R. Schwartz ¹³⁶, A. Schwartzman ¹⁴³, T.A. Schwarz ¹⁰⁶,
 Ph. Schwemling ¹³⁵, R. Schwienhorst ¹⁰⁷, A. Sciandra ¹³⁶, G. Sciolla ²⁶, A. Sciuccati ³⁶,
 G.J. Scott ⁷, F. Scuri ^{74a}, F. Scutti ¹⁰⁵, C.D. Sebastiani ⁹², C. Secord ¹⁶⁷, K. Sedlaczek ⁴⁹,
 P. Seema ¹⁸, S.C. Seidel ¹¹², A. Seiden ¹³⁶, B.D. Seidlitz ⁴¹, C. Seitz ⁴⁸, J.M. Seixas ^{82b},
 G. Sekhniaidze ^{72a}, S.J. Sekula ⁴⁴, L. Selem ⁴, A. Seletskiy ³⁸, N. Semprini-Cesari ^{23b,23a},
 S. Sen ⁵¹, D. Sengupta ⁵⁶, V. Senthilkumar ¹⁶⁵, L. Serin ⁶⁶, L. Serkin ^{69a,69b}, M. Serochkin ³⁸,
 M. Sessa ^{77a,77b}, H. Severini ¹²⁰, K.A. Sexton ²⁹, F. Sforza ^{57b,57a}, A. Sfyrla ⁵⁶, E. Shabalina ⁵⁵,
 E. Shafto ¹⁴⁵, R. Shaheen ¹⁴⁴, J.D. Shahinian ¹²⁸, O. Shaked ¹⁷¹, D. Shaked Renous ¹⁷¹,
 L.Y. Shan ^{14a}, M. Shapiro ^{17a}, A. Sharma ³⁶, A.S. Sharma ¹⁶⁶, P. Sharma ⁸⁰, S. Sharma ⁴⁸,
 P.B. Shatalov ³⁷, K. Shaw ¹⁴⁶, S.M. Shaw ¹⁰¹, Q. Shen ^{62c,5}, D.J. Sheppard ¹⁴², P.N. Sherpa ¹⁶,
 P. Sherwood ⁹⁶, L. Shi ⁹⁶, C.O. Shimmin ¹⁷⁴, Y. Shimogama ¹⁷⁰, J.D. Shinner ⁹⁵,
 I.P.J. Shipsey ¹²⁶, S. Shirabe ⁶⁰, M. Shiyakova ³⁸, J. Shlomi ¹⁷¹, M. Shoa ¹⁷¹, M.J. Shochet ³⁹,
 J. Shojaii ¹⁰⁵, D. Shooltz ¹⁰⁷, D.R. Shope ¹²⁵, S. Shrestha ^{119,ai}, E.M. Shrif ^{33g}, M.J. Shroff ¹⁶⁷,
 A. Shutov ³⁸, P. Sicho ¹³¹, A.M. Sickles ¹⁶³, E. Sideras Haddad ^{33g}, O. Sidiropoulou ³⁶,
 A. Sidoti ^{23b}, F. Siegert ⁵⁰, Dj. Sijacki ¹⁵, R. Sikora ^{85a}, F. Sili ⁹⁰, J.M. Silva ²⁰,
 M.V. Silva Oliveira ³⁶, S.B. Silverstein ^{47a}, S. Simion ⁶⁶, V.K. Simola ³⁶, R. Simoniello ³⁶,
 E.L. Simpson ⁵⁹, L.R. Simpson ¹⁰⁶, N.D. Simpson ⁹⁸, S. Simsek ^{21d}, S. Sindhu ⁵⁵, P. Sinervo ¹⁵⁵,
 S. Singh ³⁶, S. Singh ¹⁴², S. Singh ¹⁵⁵, S. Sinha ⁴⁸, S. Sinha ^{33g}, M. Sioli ^{23b,23a},
 W. Sippach ^{aj,41}, I. Siral ³⁶, S.Yu. Sivoklov ^{37,*}, M.J. Siyad ¹³⁴, J. Sjölin ^{47a,47b}, A. Skaf ⁵⁵,
 E. Skorda ⁹⁸, P. Skubic ¹²⁰, M. Slawinska ⁸⁶, K. Sliwa ³⁶, V. Smakhtin ¹⁷¹, B.H. Smart ¹³⁴,
 J. Smiesko ³⁶, S.Yu. Smirnov ³⁷, Y. Smirnov ³⁷, L.N. Smirnova ^{37,a}, O. Smirnova ⁹⁸,
 A.C. Smith ⁴¹, D.S. Smith ¹¹⁹, E.A. Smith ³⁹, H.A. Smith ¹²⁶, J.L. Smith ⁹², R. Smith ¹⁴³,

M. Smizanska ⁹¹, K. Smolek ¹³², A. Smykiewicz ⁸⁶, A.A. Snesarev ³⁷, H.L. Snoek ¹¹⁴,
S. Snyder ²⁹, R. Sobie ^{167,x}, A. Soffer ¹⁵¹, C.A. Solans Sanchez ³⁶, E.Yu. Soldatov ³⁷,
U. Soldevila ¹⁶⁵, M.A. Solis ⁷, F.S. Soliveres Riviere ¹⁰², A.A. Solodkov ³⁷, S. Solomon ⁵⁴,
A. Soloshenko ³⁸, K. Solovieva ⁵⁴, O.V. Solovyanov ⁴⁰, V. Solovyev ³⁷, R. Soluk ², P. Sommer ³⁶,
A. Sonay ¹³, W.Y. Song ^{156b}, J.M. Sonneveld ¹¹⁴, A. Sopcak ¹³², A.L. Sopio ⁹⁶,
F. Sopkova ^{28b}, J. Sorbe ¹³⁵, V. Sothilingam ^{63a}, S. Sottocornola ⁶⁸, R. Soualah ^{116b},
Z. Soumami ^{35e}, D. South ⁴⁸, D. Soyk ¹¹⁰, S. Spagnolo ^{70a,70b}, M. Spalla ¹¹⁰, F. Spanò ⁹⁵,
P. Speers ¹⁴², D. Sperlich ⁵⁴, G. Spigo ³⁶, M. Spina ¹⁴⁶, S. Spinali ⁹¹, D.P. Spiteri ⁵⁹,
R. Spiwoks ³⁶, M. Spousta ¹³³, E.J. Staats ³⁴, A. Stabile ^{71a,71b}, R.J. Staley ²⁰, R. Stamen ^{63a},
M. Stamenkovic ¹¹⁴, I. Stamoulos ¹⁶¹, A. Stampekis ²⁰, M. Standke ²⁴, E. Stanecka ⁸⁶,
M.V. Stange ⁵⁰, B. Stanislaus ^{17a}, M.M. Stanitzki ⁴⁸, M. Stankaityte ¹²⁶, B. Stapf ⁴⁸,
E.A. Starchenko ³⁷, G.H. Stark ¹³⁶, J. Stark ¹⁰², D.M. Starko ^{156b}, P. Staroba ¹³¹,
P. Starovoitov ^{63a}, S. Stärz ¹⁰⁴, R. Staszewski ⁸⁶, G. Stavropoulos ⁴⁶, J. Steentoft ¹⁶²,
P. Steinberg ²⁹, A.L. Steinhebel ¹²³, B. Stelzer ^{142,156a}, H.J. Stelzer ¹²⁹, O. Stelzer-Chilton ^{156a},
H. Stenzel ⁵⁸, T.J. Stevenson ¹⁴⁶, G.A. Stewart ³⁶, M.C. Stockton ³⁶, G. Stoicea ^{27b},
M. Stolarski ^{130a}, S. Stonjek ¹¹⁰, N. Stouras ¹⁶¹, A. Straessner ⁵⁰, J. Strandberg ¹⁴⁴,
S. Strandberg ^{47a,47b}, M. Strauss ¹²⁰, T. Strebler ¹⁰², V. Strickland ^{34,ae}, P. Strizenc ^{28b},
R. Ströhmer ¹⁶⁸, D.M. Strom ¹²³, L.R. Strom ⁴⁸, R. Stroynowski ⁴⁴, A. Strubig ^{47a,47b},
S.A. Stucci ²⁹, B. Stugu ¹⁶, J. Stupak ¹²⁰, J.T. Sturdy ¹⁶⁰, N.A. Styles ⁴⁸, D. Su ¹⁴³, S. Su ^{62a},
W. Su ^{62d,138,62c}, X. Su ^{62a,66}, K. Sugizaki ¹⁵³, V.V. Sulin ³⁷, M.J. Sullivan ⁹²,
D.M.S. Sultan ^{78a,78b}, L. Sultanaliyeva ³⁷, S. Sultansoy ^{3b}, T. Sumida ⁸⁷, Q. Sun ⁴⁴, S. Sun ¹⁰⁶,
S. Sun ¹⁷², O. Sunneborn Gudnadottir ¹⁶², M.R. Sutton ¹⁴⁶, M. Svatos ¹³¹, M. Swiatlowski ^{156a},
T. Swirski ¹⁶⁸, I. Sykora ^{28a}, M. Sykora ¹³³, T. Sykora ¹³³, D. Ta ¹⁰⁰, K. Tackmann ^{48,v},
A. Taffard ¹⁶⁰, R. Tafirout ^{156a}, J.S. Tafoya Vargas ⁶⁶, S. Taghvirad ¹³⁴, R.H.M. Taibah ¹²⁷,
R. Takashima ⁸⁸, K. Takeda ⁸⁴, E.P. Takeva ⁵², Y. Takubo ⁸³, M. Talby ¹⁰², A.A. Talyshev ³⁷,
K.C. Tam ^{64b}, N.M. Tamir ¹⁵¹, A. Tanaka ¹⁵³, J. Tanaka ¹⁵³, R. Tanaka ⁶⁶, M. Tanasini ^{57b,57a},
J. Tang ^{62c}, S. Tang ²⁹, Z. Tao ¹⁶⁶, S. Tapia Araya ^{137f}, S. Tapprogge ¹⁰⁰, B. Tar ¹¹⁹,
A. Tarek Abouelfadl Mohamed ¹⁰⁷, S. Tarem ¹⁵⁰, Z. Tarem ¹⁵⁰, K. Tariq ^{62b}, G. Tarna ^{102,27b},
G.F. Tartarelli ^{71a}, P. Tas ¹³³, M. Tasevsky ¹³¹, M. Tasevsky ¹³¹, E. Tassi ^{43b,43a}, A.C. Tate ¹⁶³,
G. Tateno ¹⁵³, Y. Tayalati ^{35e,w}, G.N. Taylor ¹⁰⁵, W. Taylor ^{156b}, H. Teagle ⁹², A.S. Tee ¹⁷²,
R. Teixeira De Lima ¹⁴³, P. Teixeira-Dias ⁹⁵, J.J. Teoh ¹⁵⁵, K. Terashi ¹⁵³, J. Terron ⁹⁹,
S. Terzo ¹³, M. Testa ⁵³, P. Teterin ³⁷, M. Teurnier ³⁶, R.J. Teuscher ^{155,x}, A. Thaler ⁷⁹,
O. Theiner ⁵⁶, N. Themistokleous ⁵², T. Thevenaux-Pelzer ¹⁸, O. Thielmann ¹⁷³, C.T. Thomas ⁸¹,
D.W. Thomas ⁹⁵, J.O. Thomas ⁴⁴, J.P. Thomas ²⁰, E.A. Thompson ⁴⁸, P.D. Thompson ²⁰,
E. Thomson ¹²⁸, E.J. Thorpe ⁹⁴, Y. Tian ⁵⁵, V. Tikhomirov ^{37,a}, Yu.A. Tikhonov ³⁷,
S. Timoshenko ³⁷, E.X.L. Ting ¹, P. Tipton ¹⁷⁴, S. Tisserant ¹⁰², S.H. Tlou ^{33g}, A. Thourji ⁴⁰,
J. Tobias ⁵⁴, K. Todome ^{23b,23a}, T. Todorov ⁴, S. Todorova-Nova ¹³³, S. Todt ⁵⁰, M. Togawa ⁸³,
J. Tojo ⁸⁹, S. Tokár ^{28a}, K. Tokushuku ⁸³, O. Toldaiev ⁶⁸, R. Tombs ³², M. Tomoto ^{83,111},
D. Tompkins ⁷, L. Tompkins ¹⁴³, K.W. Topolnicki ^{85b}, P. Tornambe ¹⁰³, E. Torrence ¹²³,
H. Torres ⁵⁰, E. Torró Pastor ¹⁶⁵, M. Toscani ³⁰, C. Tosciri ³⁹, M. Tost ¹¹, D.R. Tovey ¹³⁹,
A. Traeet ¹⁶, L. Tranchand ³⁶, I.S. Trandafir ^{27b}, F.F. Trantou ⁴⁶, P. Trattino ^{72b}, R. Travaglini ^{23a},
T. Trefzger ¹⁶⁸, A. Tricoli ²⁹, I.M. Trigger ^{156a}, S. Trincaz-Duvoid ¹²⁷, D.A. Trischuk ²⁶,
B. Trocme ⁶⁰, I. Troglazov ³⁸, A. Trofymov ⁶⁶, C. Troncon ^{71a}, G.T. Troska ⁴⁹, D. Trotta ^{72b},
M. Trovatelli ¹⁶⁷, M. Trovato ⁶, L. Truong ^{33c}, M. Trzebinski ⁸⁶, A. Trzupek ⁸⁶, F. Tsai ¹⁴⁵,
M. Tsai ¹⁰⁶, B.W.H. Tse ^{64a}, I. Tsiafis ¹⁵², A. Tsiamis ^{152,e}, P.V. Tsiarehshka ³⁷, S. Tsigaridas ^{156a},
A. Tsirigotis ^{152,t}, V. Tsiskaridze ¹⁴⁵, E.G. Tskhadadze ^{149a}, M. Tsopoulou ^{152,e}, Y. Tsujikawa ⁸⁷,
I.I. Tsukerman ³⁷, V. Tsulaia ^{17a}, S. Tsuno ⁸³, O. Tsur ¹⁵⁰, D. Tsybychev ¹⁴⁵, Y. Tu ^{64b},

A. Tudorache ^{id27b}, V. Tudorache ^{id27b}, A.N. Tuna ^{id36}, S. Turchikhin ^{id38}, P. Turco ^{43a}, I. Turk Cakir ^{id3a},
 R. Turra ^{id71a}, T. Turtuvshin ^{id38,y}, E.M. Tusi ^{76b}, P.M. Tuts ^{id41}, S. Tzamarias ^{id152,e}, P. Tzanis ^{id10},
 S. Tzanos ¹⁰, E. Tzovara ^{id100}, K. Uchida ¹⁵³, K.R. Ukah ^{id61}, F. Ukegawa ^{id157}, P.A. Ulloa Poblete ^{id137c},
 E.N. Umaka ^{id29}, G. Unal ^{id36}, M. Unal ^{id11}, A. Undrus ^{id29}, G. Unel ^{id160}, K. Uno ^{id153}, J. Urban ^{id28b},
 V. Urbasek ¹²², P. Urquijo ^{id105}, G. Usai ^{id8}, R. Ushioda ^{id154}, M. Usman ^{id108}, M. Usseglio ¹³⁵,
 Z. Uysal ^{id21b}, L. Vacavant ^{id102}, V. Vacek ^{id132}, Th. Vacher ¹³⁵, B. Vachon ^{id104}, K.O.H. Vadla ^{id125},
 T. Vafeiadis ^{id36}, A. Vaitkus ^{id96}, C. Valderanis ^{id109}, E. Valdes Santurio ^{id47a,47b}, M. Valente ^{id156a},
 S. Valentinetti ^{id23b,23a}, A. Valero ^{id165}, A. Vallier ^{id102}, J.A. Valls Ferrer ^{id165}, D.R. Van Arneman ^{id114},
 T.R. Van Daalen ^{id138}, P. Van Gemmeren ^{id6}, M.G. Van Overbeek ¹¹⁴, M. Van Rijnbach ^{id125,36},
 S. Van Stroud ^{id96}, I. Van Vulpen ^{id114}, M. Vanadia ^{id76a,76b}, W. Vandelli ^{id36}, M. Vandenbroucke ^{id135},
 E.R. Vandewall ^{id121}, D. Vannicola ^{id151}, L. Vannoli ^{id57b,57a}, A. Varga-Rehling ¹⁰⁹, R. Vari ^{id75a},
 E.W. Varnes ^{id7}, C. Varni ^{id17a}, T. Varol ^{id148}, D. Varouchas ^{id66}, L. Varriale ^{id165}, K.E. Varvell ^{id147},
 M.E. Vasile ^{id27b}, L. Vaslin ⁴⁰, G.A. Vasquez ^{id167}, F. Vazeille ^{id40}, T. Vazquez Schroeder ^{id36},
 A. Vdovin ¹⁵⁰, J. Veatch ^{id31}, V. Vecchio ^{id101}, M.J. Veen ^{id103}, I. Veliscek ^{id126}, L.M. Veloce ^{id155},
 F. Veloso ^{id130a,130c}, S. Veneziano ^{id75a}, A. Ventura ^{id70a,70b}, N. Venturi ³⁶, A. Verbytskyi ^{id110},
 F. Vercellati ^{73b}, M. Verducci ^{id74a,74b}, M.M. Vergain ³⁶, C. Vergis ^{id24}, M. Verissimo De Araujo ^{id82b},
 W. Verkerke ^{id114}, B. Verlaat ³⁶, J.C. Vermeulen ^{id114}, C. Vernieri ^{id143}, P.J. Verschuuren ^{id95},
 M. Vessella ^{id103}, M.C. Vetterli ^{id142,ae}, A. Vgenopoulos ^{id152,e}, N. Viaux Maira ^{id137f}, P. Vichoudis ^{id36},
 T. Vickey ^{id139}, O.E. Vickey Boeriu ^{id139}, G.H.A. Viehhauser ^{id126}, J. Vieira De Souza ^{id100},
 L. Vigani ^{id63b}, E. Vigeolas ¹⁰², M. Villa ^{id23b,23a}, M. Villaplana Perez ^{id165}, E.M. Villhauer ⁵²,
 E. Vilucchi ^{id53}, M.G. Vincter ^{id34}, M. Vinogradov ³⁸, G.S. Virdee ^{id20}, A. Vishwakarma ^{id52},
 C. Vittori ^{id23b,23a}, I. Vivarelli ^{id146}, S. Vlachos ^{id10}, V. Vladimirov ¹⁶⁹, E. Voevodina ^{id110}, F. Vogel ^{id109},
 S. Vogt ^{id110}, P. Vokac ^{id132}, J. Von Ahnen ^{id48}, E. Von Toerne ^{id24}, B. Vormwald ^{id36}, V. Vorobel ^{id133},
 K. Vorobev ^{id37}, M. Vos ^{id165}, K. Voss ^{id141}, J.H. Vosseveld ^{id92}, M. Vozak ^{id114}, L. Vozdecky ^{id94},
 N. Vranjes ^{id15}, M. Vranjes Milosavljevic ^{id15}, M. Vreeswijk ^{id114}, C. Vuillemin ¹³⁵, R. Vuillemer ^{id36},
 O. Vujinovic ^{id100}, I. Vukotic ^{id39}, S. Wada ^{id157}, C. Wagner ¹⁰³, W. Wagner ^{id173}, S. Wahdan ^{id173},
 H. Wahlberg ^{id90}, R. Wakasa ^{id157}, M. Wakida ^{id111}, V.M. Walbrecht ^{id110}, J. Walder ^{id134},
 R. Walker ^{id109}, R.W. Walker ^{id7}, W. Walkowiak ^{id141}, A.M. Wang ^{id61}, A.Z. Wang ^{id172}, C. Wang ^{id62a},
 C. Wang ^{id62c}, H. Wang ^{id17a}, J. Wang ^{id64a}, J. Wang ^{id62a}, J. Wang ^{id106}, Q. Wang ⁴¹, R.-J. Wang ^{id100},
 R. Wang ^{id61}, R. Wang ^{id6}, S.M. Wang ^{id148}, S. Wang ^{id62b}, T. Wang ^{id62a}, W.T. Wang ^{id80},
 X. Wang ^{id106}, X. Wang ^{id14c}, X. Wang ^{id62a}, X. Wang ^{id163}, X. Wang ^{id62c}, Y. Wang ^{id62d}, Y. Wang ^{id14c},
 Z. Wang ^{id106}, Z. Wang ^{id62d,51,62c}, Z. Wang ^{id106}, A. Warburton ^{id104}, R.J. Ward ^{id20}, N. Warrack ^{id59},
 A.T. Watson ^{id20}, H. Watson ^{id59}, M.F. Watson ^{id20}, G. Watts ^{id138}, B.M. Waugh ^{id96}, C. Weaverdyck ¹⁰⁶,
 A.F. Webb ^{id11}, C. Weber ^{id29}, H.A. Weber ^{id18}, J. Weber ⁴⁹, M. Weber ⁵⁶, M.S. Weber ^{id19},
 S.A. Weber ^{id34}, S.M. Weber ^{id63a}, C. Wei ^{id62a}, Y. Wei ^{id126}, A.R. Weidberg ^{id126}, J. Weingarten ^{id49},
 M. Weirich ^{id100}, C. Weiser ^{id54}, S.W. Welch ^{id121}, C.J. Wells ^{id48}, P.S. Wells ^{id36}, T. Wenaus ^{id29},
 B. Wendland ^{id49}, T. Wengler ^{id36}, N.S. Wenke ^{id110}, M. Wensing ^{id173}, N. Wermes ^{id24}, M. Wessels ^{id63a},
 K. Whalen ^{id123}, A.M. Wharton ^{id91}, A.S. White ^{id61}, A. White ^{id8}, M.J. White ^{id1}, D. Whiteson ^{id160},
 L. Wickremasinghe ^{id124}, W. Wiedenmann ^{id172}, C. Wiel ^{id50}, M. Wielers ^{id134}, C. Wiglesworth ^{id42},
 L.A.M. Wiik-Fuchs ^{id54}, D.J. Wilbern ¹²⁰, H.G. Wilkens ^{id36}, D.M. Williams ^{id41}, H.H. Williams ¹²⁸,
 S. Williams ^{id32}, S. Willocq ^{id103}, P.J. Windischhofer ^{id126}, I. Wingerter-Seez ^{id4}, F. Winklmeier ^{id123},
 B.T. Winter ^{id54}, J.K. Winter ^{id101}, M. Wittgen ¹⁴³, T. Wittig ⁴⁹, M. Wobisch ^{id97}, R. Wölker ^{id126},
 J. Wollrath ¹⁶⁰, K. Wolniewicz ²⁹, M.W. Wolter ^{id86}, H. Wolters ^{id130a,130c}, V.W.S. Wong ^{id166},
 A.F. Wongel ^{id48}, S.D. Worm ^{id48}, B.K. Wosiek ^{id86}, J. Wotschack ³⁶, A.S. Woyshville ^{id119},
 K.W. Woźniak ^{id86}, K. Wraight ^{id59}, J. Wu ^{id14a,14d}, M. Wu ^{id64a}, M. Wu ^{id113}, S.L. Wu ^{id172},
 W. Wu ^{id62c}, W. Wu ^{id106}, X. Wu ^{id56}, Y. Wu ^{id62a}, Z. Wu ^{id135,62a}, J. Wuerzinger ^{id126}, T.R. Wyatt ^{id101},
 B.M. Wynne ^{id52}, S. Xella ^{id42}, L. Xia ^{id14c}, M. Xia ^{id14b}, J. Xiang ^{id64c}, X. Xiao ^{id106}, M. Xie ^{id62a},

X. Xie ^{62a}, S. Xin ^{14a,14d}, J. Xiong ^{17a}, I. Xiotidis ¹⁴⁶, D. Xu ^{14a}, H. Xu ^{62a}, H. Xu ^{62a}, H. Xu ²⁹, L. Xu ^{62a}, R. Xu ¹²⁸, R. Xu ⁴¹, T. Xu ¹⁰⁶, W. Xu ¹⁰⁶, Y. Xu ^{14b}, Z. Xu ^{62b}, Z. Xu ^{14a}, B. Yabsley ¹⁴⁷, S. Yacoob ^{33a}, N. Yamaguchi ⁸⁹, Y. Yamaguchi ¹⁵⁴, S. Yamamoto ¹⁵³, H. Yamauchi ¹⁵⁷, T. Yamazaki ^{17a}, Y. Yamazaki ⁸⁴, J. Yan ^{62c}, S. Yan ¹²⁶, Z. Yan ²⁵, A. Yandyan ¹⁷⁵, H.J. Yang ^{62c,62d}, H.T. Yang ^{62a}, S. Yang ^{62a}, T. Yang ^{64c}, X. Yang ^{62a}, X. Yang ^{14a}, Y. Yang ⁴⁴, Z. Yang ^{62a,106}, L. Yao ¹²⁹, W-M. Yao ^{17a}, Y.C. Yap ⁴⁸, H. Ye ^{14c}, H. Ye ⁵⁵, J. Ye ⁴⁴, S. Ye ²⁹, X. Ye ^{62a}, Y. Yeh ⁹⁶, I. Yeletskikh ³⁸, B.K. Yeo ^{17a}, M.R. Yexley ⁹¹, S.C. Yildiz ¹⁶⁰, P. Yin ⁴¹, W. Yin ²⁹, K. Yorita ¹⁷⁰, S. Younas ^{27b}, C.J.S. Young ⁵⁴, C. Young ¹⁴³, Y. Yu ^{62a}, M. Yuan ¹⁰⁶, R. Yuan ^{62b,k}, L. Yue ⁹⁶, X. Yue ^{63a}, S. Yukhimchuk ³⁸, M. Zaazoua ^{35e}, B. Zabinski ⁸⁶, A. Zachariadou ¹⁶¹, H. Zaghia ¹³⁵, V. Zahradnik ¹³¹, E. Zaid ⁵², T. Zakareishvili ^{149b}, N. Zakharchuk ³⁴, S. Zambito ⁵⁶, J.A. Zamora Saa ^{137d,137b}, J. Zang ¹⁵³, D. Zanzi ⁵⁴, O. Zaplatilek ¹³², S.V. Zeiβner ⁴⁹, C. Zeitnitz ¹⁷³, J.C. Zeng ¹⁶³, D.T. Zenger Jr ²⁶, O. Zenin ³⁷, T. Ženiš ^{28a}, S. Zenz ⁹⁴, S. Zerradi ^{35a}, D. Zerwas ⁶⁶, M. Zhai ^{14a,14d}, B. Zhang ^{14c}, D.F. Zhang ¹³⁹, J. Zhang ^{62b}, J. Zhang ⁶, K. Zhang ^{14a,14d}, L. Zhang ^{14c}, P. Zhang ^{14a,14d}, R. Zhang ¹⁷², S. Zhang ¹⁰⁶, T. Zhang ¹⁵³, X. Zhang ^{62c}, X. Zhang ^{62b}, Y. Zhang ^{62c,5}, Z. Zhang ^{17a}, Z. Zhang ⁶⁶, H. Zhao ¹³⁸, P. Zhao ⁵¹, T. Zhao ^{62b}, X. Zhao ⁴⁴, Y. Zhao ¹³⁶, Z. Zhao ^{62a}, A. Zhemchugov ³⁸, X. Zheng ^{62a}, Z. Zheng ¹⁴³, E. Zhivun ²⁹, D. Zhong ¹⁶³, B. Zhou ¹⁰⁶, C. Zhou ¹⁷², H. Zhou ⁷, N. Zhou ^{62c}, S. Zhou ^{62a}, Y. Zhou ⁷, C.G. Zhu ^{62b}, C. Zhu ^{14a,14d}, H.L. Zhu ^{62a}, H. Zhu ^{14a}, J. Zhu ¹⁰⁶, Y. Zhu ^{62c}, Y. Zhu ^{62a}, X. Zhuang ^{14a}, K. Zhukov ³⁷, V. Zhulanov ³⁷, A. Zibell ¹⁶⁸, J. Zich ¹⁶⁴, N.I. Zimine ³⁸, J. Zimmermann ¹¹⁰, S. Zimmermann ⁵⁴, J. Zinsser ^{63b}, M. Ziolkowski ¹⁴¹, L. Živković ¹⁵, A. Zoccoli ^{23b,23a}, K. Zoch ⁵⁶, I. Zolkin ¹⁵¹, E. Zonca ¹³⁵, T.G. Zorbas ¹³⁹, O. Zormpa ⁴⁶, W. Zou ⁴¹, G. Zuk ¹²⁹, A. Zullo ^{75b}, L. Zwalinski ³⁶.

¹Department of Physics, University of Adelaide, Adelaide; Australia.

²Department of Physics, University of Alberta, Edmonton AB; Canada.

³(^a)Department of Physics, Ankara University, Ankara; (^b)Division of Physics, TOBB University of Economics and Technology, Ankara; Türkiye.

⁴LAPP, Univ. Savoie Mont Blanc, CNRS/IN2P3, Annecy; France.

⁵APC, Université Paris Cité, CNRS/IN2P3, Paris; France.

⁶High Energy Physics Division, Argonne National Laboratory, Argonne IL; United States of America.

⁷Department of Physics, University of Arizona, Tucson AZ; United States of America.

⁸Department of Physics, University of Texas at Arlington, Arlington TX; United States of America.

⁹Physics Department, National and Kapodistrian University of Athens, Athens; Greece.

¹⁰Physics Department, National Technical University of Athens, Zografou; Greece.

¹¹Department of Physics, University of Texas at Austin, Austin TX; United States of America.

¹²Institute of Physics, Azerbaijan Academy of Sciences, Baku; Azerbaijan.

¹³Institut de Física d'Altes Energies (IFAE), Barcelona Institute of Science and Technology, Barcelona; Spain.

¹⁴(^a)Institute of High Energy Physics, Chinese Academy of Sciences, Beijing; (^b)Physics Department, Tsinghua University, Beijing; (^c)Department of Physics, Nanjing University, Nanjing; (^d)University of Chinese Academy of Science (UCAS), Beijing; China.

¹⁵Institute of Physics, University of Belgrade, Belgrade; Serbia.

¹⁶Department for Physics and Technology, University of Bergen, Bergen; Norway.

¹⁷(^a)Physics Division, Lawrence Berkeley National Laboratory, Berkeley CA; (^b)University of California, Berkeley CA; United States of America.

¹⁸Institut für Physik, Humboldt Universität zu Berlin, Berlin; Germany.

- ¹⁹Albert Einstein Center for Fundamental Physics and Laboratory for High Energy Physics, University of Bern, Bern; Switzerland.
- ²⁰School of Physics and Astronomy, University of Birmingham, Birmingham; United Kingdom.
- ²¹(*a*) Department of Physics, Bogazici University, Istanbul; (*b*) Department of Physics Engineering, Gaziantep University, Gaziantep; (*c*) Department of Physics, Istanbul University, Istanbul; (*d*) Istinye University, Sariyer, Istanbul; Türkiye.
- ²²(*a*) Facultad de Ciencias y Centro de Investigaciones, Universidad Antonio Nariño, Bogotá; (*b*) Departamento de Física, Universidad Nacional de Colombia, Bogotá; Colombia.
- ²³(*a*) Dipartimento di Fisica e Astronomia A. Righi, Università di Bologna, Bologna; (*b*) INFN Sezione di Bologna; Italy.
- ²⁴Physikalisches Institut, Universität Bonn, Bonn; Germany.
- ²⁵Department of Physics, Boston University, Boston MA; United States of America.
- ²⁶Department of Physics, Brandeis University, Waltham MA; United States of America.
- ²⁷(*a*) Transilvania University of Brasov, Brasov; (*b*) Horia Hulubei National Institute of Physics and Nuclear Engineering, Bucharest; (*c*) Department of Physics, Alexandru Ioan Cuza University of Iasi, Iasi; (*d*) National Institute for Research and Development of Isotopic and Molecular Technologies, Physics Department, Cluj-Napoca; (*e*) University Politehnica Bucharest, Bucharest; (*f*) West University in Timisoara, Timisoara; (*g*) Faculty of Physics, University of Bucharest, Bucharest; Romania.
- ²⁸(*a*) Faculty of Mathematics, Physics and Informatics, Comenius University, Bratislava; (*b*) Department of Subnuclear Physics, Institute of Experimental Physics of the Slovak Academy of Sciences, Kosice; Slovak Republic.
- ²⁹Physics Department, Brookhaven National Laboratory, Upton NY; United States of America.
- ³⁰Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Departamento de Física, y CONICET, Instituto de Física de Buenos Aires (IFIBA), Buenos Aires; Argentina.
- ³¹California State University, CA; United States of America.
- ³²Cavendish Laboratory, University of Cambridge, Cambridge; United Kingdom.
- ³³(*a*) Department of Physics, University of Cape Town, Cape Town; (*b*) iThemba Labs, Western Cape; (*c*) Department of Mechanical Engineering Science, University of Johannesburg, Johannesburg; (*d*) National Institute of Physics, University of the Philippines Diliman (Philippines); (*e*) University of South Africa, Department of Physics, Pretoria; (*f*) University of Zululand, KwaDlangezwa; (*g*) School of Physics, University of the Witwatersrand, Johannesburg; South Africa.
- ³⁴Department of Physics, Carleton University, Ottawa ON; Canada.
- ³⁵(*a*) Faculté des Sciences Ain Chock, Réseau Universitaire de Physique des Hautes Energies - Université Hassan II, Casablanca; (*b*) Faculté des Sciences, Université Ibn-Tofail, Kénitra; (*c*) Faculté des Sciences Semlalia, Université Cadi Ayyad, LPHEA-Marrakech; (*d*) LPMR, Faculté des Sciences, Université Mohamed Premier, Oujda; (*e*) Faculté des sciences, Université Mohammed V, Rabat; (*f*) Institute of Applied Physics, Mohammed VI Polytechnic University, Ben Guerir; Morocco.
- ³⁶CERN, Geneva; Switzerland.
- ³⁷Affiliated with an institute covered by a cooperation agreement with CERN.
- ³⁸Affiliated with an international laboratory covered by a cooperation agreement with CERN.
- ³⁹Enrico Fermi Institute, University of Chicago, Chicago IL; United States of America.
- ⁴⁰LPC, Université Clermont Auvergne, CNRS/IN2P3, Clermont-Ferrand; France.
- ⁴¹Nevis Laboratory, Columbia University, Irvington NY; United States of America.
- ⁴²Niels Bohr Institute, University of Copenhagen, Copenhagen; Denmark.
- ⁴³(*a*) Dipartimento di Fisica, Università della Calabria, Rende; (*b*) INFN Gruppo Collegato di Cosenza, Laboratori Nazionali di Frascati; Italy.
- ⁴⁴Physics Department, Southern Methodist University, Dallas TX; United States of America.

- ⁴⁵Physics Department, University of Texas at Dallas, Richardson TX; United States of America.
- ⁴⁶National Centre for Scientific Research "Demokritos", Agia Paraskevi; Greece.
- ⁴⁷(^a) Department of Physics, Stockholm University; (^b) Oskar Klein Centre, Stockholm; Sweden.
- ⁴⁸Deutsches Elektronen-Synchrotron DESY, Hamburg and Zeuthen; Germany.
- ⁴⁹Fakultät Physik, Technische Universität Dortmund, Dortmund; Germany.
- ⁵⁰Institut für Kern- und Teilchenphysik, Technische Universität Dresden, Dresden; Germany.
- ⁵¹Department of Physics, Duke University, Durham NC; United States of America.
- ⁵²SUPA - School of Physics and Astronomy, University of Edinburgh, Edinburgh; United Kingdom.
- ⁵³INFN e Laboratori Nazionali di Frascati, Frascati; Italy.
- ⁵⁴Physikalisches Institut, Albert-Ludwigs-Universität Freiburg, Freiburg; Germany.
- ⁵⁵II. Physikalisches Institut, Georg-August-Universität Göttingen, Göttingen; Germany.
- ⁵⁶Département de Physique Nucléaire et Corpusculaire, Université de Genève, Genève; Switzerland.
- ⁵⁷(^a) Dipartimento di Fisica, Università di Genova, Genova; (^b) INFN Sezione di Genova; Italy.
- ⁵⁸II. Physikalisches Institut, Justus-Liebig-Universität Giessen, Giessen; Germany.
- ⁵⁹SUPA - School of Physics and Astronomy, University of Glasgow, Glasgow; United Kingdom.
- ⁶⁰LPSC, Université Grenoble Alpes, CNRS/IN2P3, Grenoble INP, Grenoble; France.
- ⁶¹Laboratory for Particle Physics and Cosmology, Harvard University, Cambridge MA; United States of America.
- ⁶²(^a) Department of Modern Physics and State Key Laboratory of Particle Detection and Electronics, University of Science and Technology of China, Hefei; (^b) Institute of Frontier and Interdisciplinary Science and Key Laboratory of Particle Physics and Particle Irradiation (MOE), Shandong University, Qingdao; (^c) School of Physics and Astronomy, Shanghai Jiao Tong University, Key Laboratory for Particle Astrophysics and Cosmology (MOE), SKLPPC, Shanghai; (^d) Tsung-Dao Lee Institute, Shanghai; China.
- ⁶³(^a) Kirchhoff-Institut für Physik, Ruprecht-Karls-Universität Heidelberg, Heidelberg; (^b) Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg, Heidelberg; Germany.
- ⁶⁴(^a) Department of Physics, Chinese University of Hong Kong, Shatin, N.T., Hong Kong; (^b) Department of Physics, University of Hong Kong, Hong Kong; (^c) Department of Physics and Institute for Advanced Study, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong; China.
- ⁶⁵Department of Physics, National Tsing Hua University, Hsinchu; Taiwan.
- ⁶⁶IJCLab, Université Paris-Saclay, CNRS/IN2P3, 91405, Orsay; France.
- ⁶⁷Centro Nacional de Microelectrónica (IMB-CNM-CSIC), Barcelona; Spain.
- ⁶⁸Department of Physics, Indiana University, Bloomington IN; United States of America.
- ⁶⁹(^a) INFN Gruppo Collegato di Udine, Sezione di Trieste, Udine; (^b) ICTP, Trieste; (^c) Dipartimento Politecnico di Ingegneria e Architettura, Università di Udine, Udine; Italy.
- ⁷⁰(^a) INFN Sezione di Lecce; (^b) Dipartimento di Matematica e Fisica, Università del Salento, Lecce; Italy.
- ⁷¹(^a) INFN Sezione di Milano; (^b) Dipartimento di Fisica, Università di Milano, Milano; Italy.
- ⁷²(^a) INFN Sezione di Napoli; (^b) Dipartimento di Fisica, Università di Napoli, Napoli; Italy.
- ⁷³(^a) INFN Sezione di Pavia; (^b) Dipartimento di Fisica, Università di Pavia, Pavia; Italy.
- ⁷⁴(^a) INFN Sezione di Pisa; (^b) Dipartimento di Fisica E. Fermi, Università di Pisa, Pisa; Italy.
- ⁷⁵(^a) INFN Sezione di Roma; (^b) Dipartimento di Fisica, Sapienza Università di Roma, Roma; Italy.
- ⁷⁶(^a) INFN Sezione di Roma Tor Vergata; (^b) Dipartimento di Fisica, Università di Roma Tor Vergata, Roma; Italy.
- ⁷⁷(^a) INFN Sezione di Roma Tre; (^b) Dipartimento di Matematica e Fisica, Università Roma Tre, Roma; Italy.
- ⁷⁸(^a) INFN-TIFPA; (^b) Università degli Studi di Trento, Trento; Italy.
- ⁷⁹Universität Innsbruck, Department of Astro and Particle Physics, Innsbruck; Austria.
- ⁸⁰University of Iowa, Iowa City IA; United States of America.

- ⁸¹Department of Physics and Astronomy, Iowa State University, Ames IA; United States of America.
- ⁸²(^a)Departamento de Engenharia Elétrica, Universidade Federal de Juiz de Fora (UFJF), Juiz de Fora; (^b)Universidade Federal do Rio De Janeiro COPPE/EE/IF, Rio de Janeiro; (^c)Instituto de Física, Universidade de São Paulo, São Paulo; (^d)Rio de Janeiro State University, Rio de Janeiro; Brazil.
- ⁸³KEK, High Energy Accelerator Research Organization, Tsukuba; Japan.
- ⁸⁴Graduate School of Science, Kobe University, Kobe; Japan.
- ⁸⁵(^a)AGH University of Science and Technology, Faculty of Physics and Applied Computer Science, Krakow; (^b)Marian Smoluchowski Institute of Physics, Jagiellonian University, Krakow; Poland.
- ⁸⁶Institute of Nuclear Physics Polish Academy of Sciences, Krakow; Poland.
- ⁸⁷Faculty of Science, Kyoto University, Kyoto; Japan.
- ⁸⁸Kyoto University of Education, Kyoto; Japan.
- ⁸⁹Research Center for Advanced Particle Physics and Department of Physics, Kyushu University, Fukuoka ; Japan.
- ⁹⁰Instituto de Física La Plata, Universidad Nacional de La Plata and CONICET, La Plata; Argentina.
- ⁹¹Physics Department, Lancaster University, Lancaster; United Kingdom.
- ⁹²Oliver Lodge Laboratory, University of Liverpool, Liverpool; United Kingdom.
- ⁹³Department of Experimental Particle Physics, Jožef Stefan Institute and Department of Physics, University of Ljubljana, Ljubljana; Slovenia.
- ⁹⁴School of Physics and Astronomy, Queen Mary University of London, London; United Kingdom.
- ⁹⁵Department of Physics, Royal Holloway University of London, Egham; United Kingdom.
- ⁹⁶Department of Physics and Astronomy, University College London, London; United Kingdom.
- ⁹⁷Louisiana Tech University, Ruston LA; United States of America.
- ⁹⁸Fysiska institutionen, Lunds universitet, Lund; Sweden.
- ⁹⁹Departamento de Física Teórica C-15 and CIAFF, Universidad Autónoma de Madrid, Madrid; Spain.
- ¹⁰⁰Institut für Physik, Universität Mainz, Mainz; Germany.
- ¹⁰¹School of Physics and Astronomy, University of Manchester, Manchester; United Kingdom.
- ¹⁰²CPPM, Aix-Marseille Université, CNRS/IN2P3, Marseille; France.
- ¹⁰³Department of Physics, University of Massachusetts, Amherst MA; United States of America.
- ¹⁰⁴Department of Physics, McGill University, Montreal QC; Canada.
- ¹⁰⁵School of Physics, University of Melbourne, Victoria; Australia.
- ¹⁰⁶Department of Physics, University of Michigan, Ann Arbor MI; United States of America.
- ¹⁰⁷Department of Physics and Astronomy, Michigan State University, East Lansing MI; United States of America.
- ¹⁰⁸Group of Particle Physics, University of Montreal, Montreal QC; Canada.
- ¹⁰⁹Fakultät für Physik, Ludwig-Maximilians-Universität München, München; Germany.
- ¹¹⁰Max-Planck-Institut für Physik (Werner-Heisenberg-Institut), München; Germany.
- ¹¹¹Graduate School of Science and Kobayashi-Maskawa Institute, Nagoya University, Nagoya; Japan.
- ¹¹²Department of Physics and Astronomy, University of New Mexico, Albuquerque NM; United States of America.
- ¹¹³Institute for Mathematics, Astrophysics and Particle Physics, Radboud University/Nikhef, Nijmegen; Netherlands.
- ¹¹⁴Nikhef National Institute for Subatomic Physics and University of Amsterdam, Amsterdam; Netherlands.
- ¹¹⁵Department of Physics, Northern Illinois University, DeKalb IL; United States of America.
- ¹¹⁶(^a)New York University Abu Dhabi, Abu Dhabi; (^b)University of Sharjah, Sharjah; United Arab Emirates.
- ¹¹⁷Department of Physics, New York University, New York NY; United States of America.

- ¹¹⁸Ochanomizu University, Otsuka, Bunkyo-ku, Tokyo; Japan.
- ¹¹⁹Ohio State University, Columbus OH; United States of America.
- ¹²⁰Homer L. Dodge Department of Physics and Astronomy, University of Oklahoma, Norman OK; United States of America.
- ¹²¹Department of Physics, Oklahoma State University, Stillwater OK; United States of America.
- ¹²²Palacký University, Joint Laboratory of Optics, Olomouc; Czech Republic.
- ¹²³Institute for Fundamental Science, University of Oregon, Eugene, OR; United States of America.
- ¹²⁴Graduate School of Science, Osaka University, Osaka; Japan.
- ¹²⁵Department of Physics, University of Oslo, Oslo; Norway.
- ¹²⁶Department of Physics, Oxford University, Oxford; United Kingdom.
- ¹²⁷LPNHE, Sorbonne Université, Université Paris Cité, CNRS/IN2P3, Paris; France.
- ¹²⁸Department of Physics, University of Pennsylvania, Philadelphia PA; United States of America.
- ¹²⁹Department of Physics and Astronomy, University of Pittsburgh, Pittsburgh PA; United States of America.
- ¹³⁰^(a)Laboratório de Instrumentação e Física Experimental de Partículas - LIP, Lisboa;^(b)Departamento de Física, Faculdade de Ciências, Universidade de Lisboa, Lisboa;^(c)Departamento de Física, Universidade de Coimbra, Coimbra;^(d)Centro de Física Nuclear da Universidade de Lisboa, Lisboa;^(e)Departamento de Física, Universidade do Minho, Braga;^(f)Departamento de Física Teórica y del Cosmos, Universidad de Granada, Granada (Spain);^(g)Departamento de Física, Instituto Superior Técnico, Universidade de Lisboa, Lisboa; Portugal.
- ¹³¹Institute of Physics of the Czech Academy of Sciences, Prague; Czech Republic.
- ¹³²Czech Technical University in Prague, Prague; Czech Republic.
- ¹³³Charles University, Faculty of Mathematics and Physics, Prague; Czech Republic.
- ¹³⁴Particle Physics Department, Rutherford Appleton Laboratory, Didcot; United Kingdom.
- ¹³⁵IRFU, CEA, Université Paris-Saclay, Gif-sur-Yvette; France.
- ¹³⁶Santa Cruz Institute for Particle Physics, University of California Santa Cruz, Santa Cruz CA; United States of America.
- ¹³⁷^(a)Departamento de Física, Pontificia Universidad Católica de Chile, Santiago;^(b)Millennium Institute for Subatomic physics at high energy frontier (SAPHIR), Santiago;^(c)Instituto de Investigación Multidisciplinario en Ciencia y Tecnología, y Departamento de Física, Universidad de La Serena;^(d)Universidad Andres Bello, Department of Physics, Santiago;^(e)Instituto de Alta Investigación, Universidad de Tarapacá, Arica;^(f)Departamento de Física, Universidad Técnica Federico Santa María, Valparaíso; Chile.
- ¹³⁸Department of Physics, University of Washington, Seattle WA; United States of America.
- ¹³⁹Department of Physics and Astronomy, University of Sheffield, Sheffield; United Kingdom.
- ¹⁴⁰Department of Physics, Shinshu University, Nagano; Japan.
- ¹⁴¹Department Physik, Universität Siegen, Siegen; Germany.
- ¹⁴²Department of Physics, Simon Fraser University, Burnaby BC; Canada.
- ¹⁴³SLAC National Accelerator Laboratory, Stanford CA; United States of America.
- ¹⁴⁴Department of Physics, Royal Institute of Technology, Stockholm; Sweden.
- ¹⁴⁵Departments of Physics and Astronomy, Stony Brook University, Stony Brook NY; United States of America.
- ¹⁴⁶Department of Physics and Astronomy, University of Sussex, Brighton; United Kingdom.
- ¹⁴⁷School of Physics, University of Sydney, Sydney; Australia.
- ¹⁴⁸Institute of Physics, Academia Sinica, Taipei; Taiwan.
- ¹⁴⁹^(a)E. Andronikashvili Institute of Physics, Iv. Javakhishvili Tbilisi State University, Tbilisi;^(b)High Energy Physics Institute, Tbilisi State University, Tbilisi;^(c)University of Georgia, Tbilisi; Georgia.

- ¹⁵⁰Department of Physics, Technion, Israel Institute of Technology, Haifa; Israel.
- ¹⁵¹Raymond and Beverly Sackler School of Physics and Astronomy, Tel Aviv University, Tel Aviv; Israel.
- ¹⁵²Department of Physics, Aristotle University of Thessaloniki, Thessaloniki; Greece.
- ¹⁵³International Center for Elementary Particle Physics and Department of Physics, University of Tokyo, Tokyo; Japan.
- ¹⁵⁴Department of Physics, Tokyo Institute of Technology, Tokyo; Japan.
- ¹⁵⁵Department of Physics, University of Toronto, Toronto ON; Canada.
- ¹⁵⁶(^a) TRIUMF, Vancouver BC; (^b) Department of Physics and Astronomy, York University, Toronto ON; Canada.
- ¹⁵⁷Division of Physics and Tomonaga Center for the History of the Universe, Faculty of Pure and Applied Sciences, University of Tsukuba, Tsukuba; Japan.
- ¹⁵⁸Department of Physics and Astronomy, Tufts University, Medford MA; United States of America.
- ¹⁵⁹United Arab Emirates University, Al Ain; United Arab Emirates.
- ¹⁶⁰Department of Physics and Astronomy, University of California Irvine, Irvine CA; United States of America.
- ¹⁶¹University of West Attica, Athens; Greece.
- ¹⁶²Department of Physics and Astronomy, University of Uppsala, Uppsala; Sweden.
- ¹⁶³Department of Physics, University of Illinois, Urbana IL; United States of America.
- ¹⁶⁴University of West Bohemia, Physics Department, Plzen; Czech Republic.
- ¹⁶⁵Instituto de Física Corpuscular (IFIC), Centro Mixto Universidad de Valencia - CSIC, Valencia; Spain.
- ¹⁶⁶Department of Physics, University of British Columbia, Vancouver BC; Canada.
- ¹⁶⁷Department of Physics and Astronomy, University of Victoria, Victoria BC; Canada.
- ¹⁶⁸Fakultät für Physik und Astronomie, Julius-Maximilians-Universität Würzburg, Würzburg; Germany.
- ¹⁶⁹Department of Physics, University of Warwick, Coventry; United Kingdom.
- ¹⁷⁰Waseda University, Tokyo; Japan.
- ¹⁷¹Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot; Israel.
- ¹⁷²Department of Physics, University of Wisconsin, Madison WI; United States of America.
- ¹⁷³Fakultät für Mathematik und Naturwissenschaften, Fachgruppe Physik, Bergische Universität Wuppertal, Wuppertal; Germany.
- ¹⁷⁴Department of Physics, Yale University, New Haven CT; United States of America.
- ¹⁷⁵Yerevan Physics Institute, Yerevan; Armenia.
- ^a Also Affiliated with an institute covered by a cooperation agreement with CERN.
- ^b Also at Borough of Manhattan Community College, City University of New York, New York NY; United States of America.
- ^c Also at Bruno Kessler Foundation, Trento; Italy.
- ^d Also at Center for High Energy Physics, Peking University; China.
- ^e Also at Center for Interdisciplinary Research and Innovation (CIRI-AUTH), Thessaloniki ; Greece.
- ^f Also at Centro Studi e Ricerche Enrico Fermi; Italy.
- ^g Also at CERN, Geneva; Switzerland.
- ^h Also at Département de Physique Nucléaire et Corpusculaire, Université de Genève, Genève; Switzerland.
- ⁱ Also at Departament de Física de la Universitat Autònoma de Barcelona, Barcelona; Spain.
- ^j Also at Department of Financial and Management Engineering, University of the Aegean, Chios; Greece.
- ^k Also at Department of Physics and Astronomy, Michigan State University, East Lansing MI; United States of America.
- ^l Also at Department of Physics and Astronomy, University of Louisville, Louisville, KY; United States of America.

- m* Also at Department of Physics, Ben Gurion University of the Negev, Beer Sheva; Israel.
- n* Also at Department of Physics, California State University, East Bay; United States of America.
- o* Also at Department of Physics, California State University, Sacramento; United States of America.
- p* Also at Department of Physics, King's College London, London; United Kingdom.
- q* Also at Department of Physics, University of Fribourg, Fribourg; Switzerland.
- r* Also at Department of Physics, University of Thessaly; Greece.
- s* Also at Department of Physics, Westmont College, Santa Barbara; United States of America.
- t* Also at Hellenic Open University, Patras; Greece.
- u* Also at Institutio Catalana de Recerca i Estudis Avancats, ICREA, Barcelona; Spain.
- v* Also at Institut für Experimentalphysik, Universität Hamburg, Hamburg; Germany.
- w* Also at Institute of Applied Physics, Mohammed VI Polytechnic University, Ben Guerir; Morocco.
- x* Also at Institute of Particle Physics (IPP); Canada.
- y* Also at Institute of Physics and Technology, Ulaanbaatar; Mongolia.
- z* Also at Institute of Physics, Azerbaijan Academy of Sciences, Baku; Azerbaijan.
- aa* Also at Institute of Theoretical Physics, Ilia State University, Tbilisi; Georgia.
- ab* Also at Lawrence Livermore National Laboratory, Livermore; United States of America.
- ac* Also at RWTH Aachen University, III. Physikalisches Institut A, Aachen; Germany.
- ad* Also at The Collaborative Innovation Center of Quantum Matter (CICQM), Beijing; China.
- ae* Also at TRIUMF, Vancouver BC; Canada.
- af* Also at Università di Napoli Parthenope, Napoli; Italy.
- ag* Also at University of Chinese Academy of Sciences (UCAS), Beijing; China.
- ah* Also at University of Colorado Boulder, Department of Physics, Colorado; United States of America.
- ai* Also at Washington College, Maryland; United States of America.
- aj* Also at Department of Electrical Engineering, Columbia University, New York NY; United States of America.
- * Deceased