

The ATLAS trigger system for LHC Run 3 and trigger performance in 2022



The ATLAS collaboration

E-mail: atlas.publications@cern.ch

ABSTRACT. The ATLAS trigger system is a crucial component of the ATLAS experiment at the LHC. It is responsible for selecting events in line with the ATLAS physics programme. This paper presents an overview of the changes to the trigger and data acquisition system during the second long shutdown of the LHC, and shows the performance of the trigger system and its components in the proton-proton collisions during the 2022 commissioning period as well as its expected performance in proton-proton and heavy-ion collisions for the remainder of the third LHC data-taking period (2022–2025).

KEYWORDS: Online farms and online filtering; Trigger algorithms; Trigger concepts and systems (hardware and software); Trigger detectors

ARXIV EPRINT: [2401.06630](https://arxiv.org/abs/2401.06630)

2024 JINST 19 P06029



Contents

1	Introduction	2
2	The Run-3 ATLAS detector	3
2.1	Overview of ATLAS trigger and data acquisition system	4
2.2	The ATLAS run structure	5
3	Changes to the Trigger/DAQ system for Run 3	6
3.1	Level-1 calorimeter trigger	6
3.2	Level-1 muon trigger	8
3.3	Level-1 topological trigger	10
3.4	HLT software	11
3.5	DAQ system	13
3.6	Validation strategy	15
4	Trigger menu	15
4.1	Baseline physics trigger menu for proton-proton collisions	18
4.2	Baseline physics trigger menu for heavy-ion collisions	20
5	High-level trigger reconstruction	22
5.1	Inner detector tracking	22
5.2	Calorimeter reconstruction	36
5.3	Tracking in the muon spectrometer	39
6	Physics triggers	41
6.1	Electrons and photons	41
6.2	Muons	46
6.3	Taus	47
6.4	Jets	51
6.5	Jets containing b -hadrons	55
6.6	Missing transverse momentum	59
6.7	B -physics and light states	61
6.8	Minimum-bias and forward signatures	63
6.9	Unconventional tracking signatures	67
7	Auxiliary triggers	73
7.1	Non-collision background	73
7.2	Zero-bias trigger	75
7.3	Triggers for tracking performance studies	76

7.4	Triggers for detector performance studies (calibration, noise)	77
7.5	Beam-spot algorithms	78
8	Trigger software performance	80
8.1	Trigger software performance scaling	80
8.2	HLT processing time	82
9	Conclusion	85
The ATLAS collaboration		93

1 Introduction

The trigger system [1, 2] is an essential component of the ATLAS experiment [3] as it is responsible for deciding whether or not to permanently record data from a given bunch-crossing interaction (also referred to as an *event*) for further study. It aims to minimise the rate of the recorded events while maintaining an excellent and unbiased efficiency for physics processes of interest. During the Run 1 (2009–2013) and Run 2 (2015–2018) operational periods at the Large Hadron Collider (LHC), the ATLAS trigger system collected data from proton-proton (pp) collisions at a centre-of-mass energy up to $\sqrt{s} = 13$ TeV, a peak instantaneous luminosity of $2.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and up to a mean of 60 interactions per bunch crossing (pile-up). As part of the LHC heavy-ion (HI) programme, lead-lead nuclei collisions (Pb+Pb) at a centre-of-mass energy up to $\sqrt{s} = 5$ TeV per nucleon and a peak instantaneous luminosity of $6.2 \times 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ were achieved, with other ion species and collision configurations used in addition.

During the second LHC Long Shutdown (LS2, 2019–2021), the ATLAS trigger system underwent a major upgrade. The purpose of this upgrade was to enhance the physics reach of the experiment for the ongoing operation in Run 3 (2022–2025). In Run 3 the centre-of-mass energy of the pp collisions is $\sqrt{s} = 13.6$ TeV with an instantaneous luminosity planned to be kept constant at its peak value, approximately $2.4 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ for a duration as long as ten hours per LHC fill and 60–70 interactions per bunch crossing.

The ATLAS Run-3 trigger system, its performance in the pp collisions during the 2022 commissioning period, including rates and efficiencies, as well as expected trigger algorithm modifications for HI data taking are described in this paper. After a brief introduction to the ATLAS detector in section 2, section 3 summarises the changes to the trigger and data acquisition during LS2. Section 4 gives overviews of trigger selections implemented for Run 3 followed by an introduction to the reconstruction algorithms used at the High Level Trigger (HLT) in section 5. The implementation of different triggers and their performance are discussed in sections 6 and 7. The trigger software performance is presented in section 8.

The results presented in this paper are based on the pp collision data recorded by the ATLAS experiment and the corresponding simulated Monte Carlo (MC) events. The rates are quoted for an instantaneous luminosity value of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The integrated luminosity used varies over the figures and is measured to be at most 31 fb^{-1} for the 2022 data set [4]. Due to the absence of HI physics data taking in 2022, there are no HI results to show in this paper.

Trigger efficiencies are calculated with respect to either reconstructed or true physics objects from MC simulation. Typical methods used to determine trigger efficiency in data [5] are either the tag-and-probe or bootstrap methods. The tag-and-probe method is based on a sample of events with a known resonance, such as Z boson or J/ψ , decaying to two charged leptons. The bootstrap method involves successive measurements of trigger efficiency with respect to unbiased or lower-threshold triggers. MC generator information about the physics objects can be used to determine trigger efficiency in the MC simulation.

2 The Run-3 ATLAS detector

The ATLAS detector at the LHC has undergone a substantial upgrade with improvements to various detector subsystems and their electronics, in order to enable the broad physics programme planned for the Run-3 data taking. The original configuration of the detector (as it was built for the start of Run 1 of the LHC) is described in ref. [6].

The ATLAS detector covers nearly the entire solid angle around the collision point.¹ It consists of an inner tracking detector surrounded by a superconducting solenoid, electromagnetic and hadron calorimeters, and a muon spectrometer incorporating three large superconducting air-core toroidal magnets.

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 (PIX) 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 [7]. It is followed by the silicon microstrip tracker (SCT), which usually provides eight measurements per track. These silicon detectors are complemented by the outermost of the three tracking subsystems, 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 (typically 30 in total) above a higher energy-deposit threshold corresponding to transition radiation. The TRT gas configuration has a significant impact on the particle identification. Due to a number of leaks in flexible active gas exhaust pipes that developed during Run 1 and Run 2, it became too costly to operate the entire detector with the baseline xenon-based gas mixture. For Run 3 an argon-based gas mixture is used in the entire barrel and within a few endcap wheels on one side of the detector [3]. Due to poor absorption of transition radiation photons by the argon gas, the particle identification function is significantly reduced in the barrel region. However, in combination with dE/dx measurements, it still contributes to the ATLAS electron identification, particularly at particle energies below 10 GeV. The particle identification performance of the endcaps is largely preserved.

ATLAS uses two sampling calorimeter technologies covering the pseudorapidity range $|\eta| < 4.9$. Within the region $|\eta| < 3.2$, electromagnetic (EM) calorimetry is provided by barrel and endcap high-granularity lead/liquid-argon (LAr) calorimeters, consisting of three layers with varying granularities. In addition, a LAr presampler covering $|\eta| < 1.8$ corrects for energy loss in material upstream

¹ATLAS uses a right-handed coordinate system with its origin at the nominal interaction point (IP) in the centre of the detector and the z -axis along the beam pipe. The x -axis points from the IP to the centre of the LHC ring, and the y -axis points upwards. Polar 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 $\eta = -\ln \tan(\theta/2)$. Angular distance is measured in units of $\Delta R \equiv \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$.

of the calorimeters. Hadron calorimetry is provided by the steel/scintillator-tile calorimeter (Tile calorimeter), segmented into three barrel structures within $|\eta| < 1.7$, and two copper/LAr hadron endcap calorimeters. The solid angle coverage is completed with forward copper/LAr and tungsten/LAr calorimeter modules optimised for electromagnetic and hadronic energy measurements respectively. During LS2, the LAr calorimeter electronics was augmented with a new digital trigger path [8] providing finer granularity inputs to the upgraded trigger system discussed in section 3.

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 T m across most of the detector. Three stations of precision chambers, each consisting of layers of monitored drift tubes (MDTs), cover the region $|\eta| < 2.7$. In the innermost station of the endcaps, $|\eta| > 1.3$, the detectors used in Runs 1 and 2, Small Wheels, have been replaced by the New Small Wheels (NSWs) [3]. The NSWs use two technologies: small-strip thin gap chamber (sTGC) and Micromegas (MM) detectors, both with high-rate tolerance and improved resolution. The muon trigger system covers the range $|\eta| < 2.4$ with resistive plate chambers (RPCs) in the barrel ($|\eta| < 1.05$), and thin gap chambers (TGCs) in the endcap regions ($1.05 < |\eta| < 2.4$). 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 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. Such maintenance work is expected to continue throughout Run 3 in periods with no data taking. The resulting impact on the muon trigger efficiency is described in section 3.

2.1 Overview of ATLAS trigger and data acquisition system

The selection and recording of events is handled by the Trigger and Data Acquisition (TDAQ) system shown in figure 1.

The first-level (Level-1, L1) trigger is mainly based on two independent systems which use custom electronics to trigger on reduced-granularity information from either the calorimeters (L1Calo) or the muon detectors (L1Muon). The L1 topological processor (L1Topo) system uses kinematic information from objects reconstructed in the L1Calo and L1Muon systems and applies topological selections. Changes to these systems for Run 3 are described in detail in section 3.

The L1 trigger decision is formed by the Central Trigger Processor (CTP), based on inputs received from the L1Calo trigger system, the L1Muon trigger system through the Muon-to-Central Trigger Processor Interface (MUCTPI) [9], the L1Topo system as well as several other subsystems. These subsystems are the Minimum Bias Trigger Scintillators (MBTS) [10], the ATLAS Forward Proton (AFP) [11] detector and ALFA detector [12] discussed further in section 6.8; the LUCID-2 Cherenkov Counter [13], used for the primary luminosity measurements [14] complemented by measurements using the inner detector and calorimeters; and the Zero Degree Calorimeter (ZDC) [15] which is installed for heavy ion data taking (see section 4.2). The CTP is also responsible for applying *dead time*, a mechanism to limit the number of close-by L1 accepts [16].

As detailed in section 4, events that satisfy the *trigger menu* requirements based on object type, threshold and multiplicity are accepted at a rate up to the maximum detector read-out rate of 100 kHz (down from the bunch crossing rate of about 40 MHz) at a fixed latency (detector read-out time window) below $2.5 \mu\text{s}$. Up to 512 distinct L1 trigger items may be configured in the CTP.

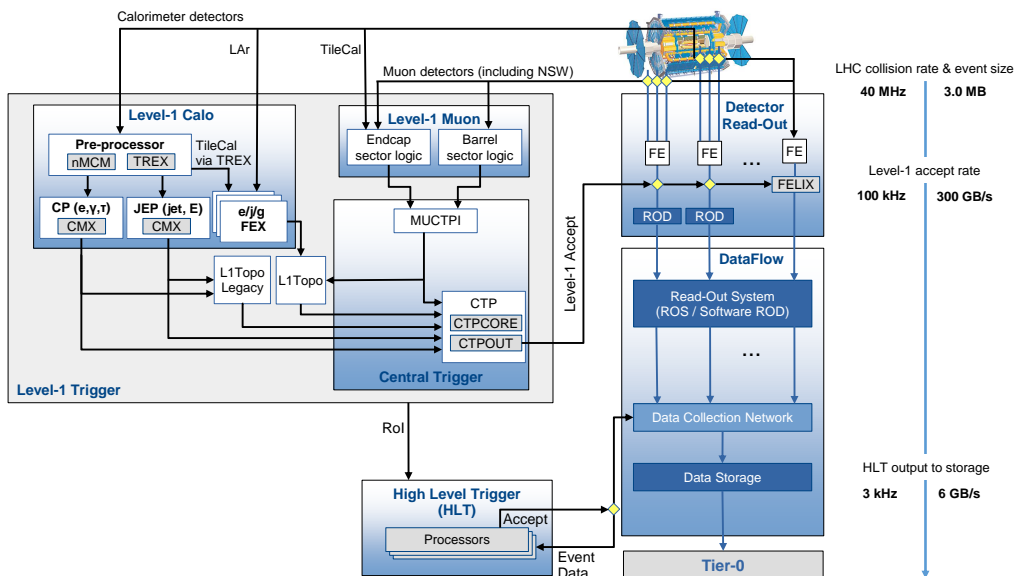


Figure 1. The ATLAS TDAQ system in Run 3 with emphasis on the components relevant for triggering as well as the detector read-out and data flow. Level-1 Calo, Level-1 Muon and the Central Trigger all send data to the Read-Out System (or FELIX/Software ROD), described in section 3.5, primarily for the purposes of offline validation and error checking. This is not shown in the diagram for simplicity. Abbreviations used are defined in sections 2.1 and 3.

If accepted by the L1 trigger, events are then sent to a software-based HLT. Here *online* algorithms reconstruct the event at progressively higher levels of detail than at L1, either in the full detector volume or in restricted Regions-of-Interest (RoIs), which are detector regions in which candidate trigger objects have been identified by the L1 trigger. The HLT software is incorporated in the same software framework as is used *offline* to reconstruct recorded events. For Run 3, this software framework was redesigned to support multi-threaded execution [17] as detailed in section 3.4. The physics output rate of the HLT during an ATLAS data-taking run is expected to be 3 kHz on average (see section 4 for more details).

The Data Acquisition (DAQ) system [18], described in section 3.5, transports data from custom subdetector electronics to offline processing, according to the decisions made by the trigger. Data are compressed prior to processing from the original event size of about 3 MB to below 2 MB. An extensive software suite [19] 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.

2.2 The ATLAS run structure

An ATLAS run is a period of data acquisition with a stable detector configuration. In the case of collecting data for physics analyses (*physics data taking*), it usually coincides with an LHC fill, which can last many hours. A unique number is assigned to every run at its beginning by the DAQ system. A run is divided into Luminosity Blocks (LB), defined as intervals of approximately constant instantaneous luminosity and stable detector conditions (including the trigger system and its configuration), with a nominal length of one minute. To define a data sample appropriate for physics

studies, quality criteria are applied to select LBs where conditions are acceptable for a particular analysis. The instantaneous luminosity in a given LB is multiplied by the LB duration to obtain the integrated luminosity delivered in that LB. From a data quality point of view, the LB represents the smallest quantity of data that can be declared good or bad for physics analysis. Further details on the LHC fill cycle, fill patterns and ATLAS run structure can be found in ref. [20].

3 Changes to the Trigger/DAQ system for Run 3

This section describes changes in the L1Calo, L1Muon and L1Topo systems as well as the HLT software and DAQ system implemented for the Run 3, including some initial performance data where available. More details about the hardware upgrades can be found in ref. [3].




Figure 2. The trigger granularity from each 0.1×0.1 trigger tower after the upgrade of the LAr Calorimeter electronics. Ten E_T values are provided from “1-4-4-1” longitudinal/transverse samples, each forming a SuperCell.

3.1 Level-1 calorimeter trigger

During LS2, the L1Calo trigger was upgraded to perform on-detector digitisation of transverse energies from the LAr calorimeters. With this change in Run 3, L1Calo can receive finer-granularity input data from the LAr calorimeter. While in Runs 1 and 2, these inputs consisted of *trigger towers* spanning 0.1×0.1 in η and ϕ , the electromagnetic calorimeter information is now provided in the form of *SuperCells* containing sums of four or eight calorimeter cells. Each trigger tower contains ten SuperCells as shown in figure 2. These trigger towers are distributed to Feature Extraction (FEX) processors. In this system, the digitisation and calibration of the LAr calorimeter data are now performed in the LAr calorimeter electronics. The Tile calorimeter data are still received in analog format, digitised by the pre-processors as shown in figure 1 and detailed below, and transmitted by the new Tile Rear EXTension (TREX) modules to both the FEXs and the Run-2 L1Calo system (hereafter referred to as the legacy system).

The upgraded L1Calo system includes new electromagnetic (eFEX) and jet (jFEX) feature extractors as well as a global feature extractor (gFEX). The full *SuperCell* granularity is available at the eFEX to reconstruct EM objects and hadronically decaying τ leptons, as well as shower-shape variables used for their identification [3]. Three programmable threshold parameters (commonly referred to as loose, medium and tight) are available for three shower-shape variables for the EM

triggers and one shower-shape variable for tau triggers. The resulting threshold pass bits for each shower shape variable are sent to the Run-3 L1Topo system and allow for the configuration of different levels of background rejection for these triggers. In the region with $|\eta| < 2.5$ jFEX receives towers with a granularity of 0.1×0.1 in η - ϕ space, which is a factor of two better in both dimensions than the inputs used for the Run-2 L1-jet-trigger system. In the region with $|\eta| > 2.5$ the granularity is similar to Run 2 with some improvements in the far forward region. The jFEX is used to reconstruct small-radius (small- R) jets with $R = 0.4$. Its performance is expected to be similar to that of the legacy L1Calo system for the single jet triggers, while improving the reconstruction for the nearby jets in the multi-jet triggers. In addition, jFEX brings new capability to reconstruct large-radius (large- R) jets with $R = 0.8$, hadronically decaying τ leptons in the range of $|\eta| \leq 2.5$, and electromagnetic objects in the forward region of $2.3 \leq |\eta| \leq 4.9$. The latter two triggers have an optional isolation requirement. Missing transverse momentum (E_T^{miss}) is computed by summing transverse energy values in slices of constant ϕ and then weighted by $\cos(\phi)$ and $\sin(\phi)$ in order to determine the x - and y -components, respectively. The gFEX has been designed with a coarser granularity than jFEX (similar to the Run-2 system) so that the data from the entire calorimeter can be processed on a single module, facilitating the identification of boosted objects and global observables. The gFEX identifies large- R jets within a “circular” 1.8×1.8 area ($R < 0.9$) and provides local pile-up density and substructure information. The gFEX computes E_T^{miss} by separating the transverse energy (E_T) sums into “hard” and “soft” where the hard term consists of the E_T sum of towers 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.

Both jFEX and gFEX have energy sum algorithms based on vector and scalar sums of E_T . The pile-up subtraction algorithms introduced in jFEX and gFEX improve L1 jet and energy sum triggers efficiency, rates and purity. More details on all these algorithms can be found in ref. [3]. Complementary information on the same types of objects (e.g. hadronically decaying τ leptons from eFEX and jFEX or E_T^{miss} from jFEX and gFEX) can be used in the Run-3 L1Topo selection to further improve their L1 trigger performance.

The legacy L1Calo system [2, 21] can be operated in parallel to the upgraded L1Calo system in Run 3. In the legacy system, analogue detector signals are digitised and calibrated in the multi-chip modules (nMCM) in the pre-processor system and sent in parallel to the Cluster Processors (CP) and the Jet/Energy-sum Processors (JEP). The CP system identifies electron, photon, and τ -lepton candidates above a programmable threshold, and the JEP system identifies jet candidates and produces global sums of total and missing transverse energy.

During the commissioning and validation of the upgraded L1Calo system in 2022, calorimeter-based physics triggers were provided by the legacy L1Calo system. In parallel, the electromagnetic objects for the new Run-3 L1Calo trigger were commissioned to be used from the start of 2023 data taking. After final tuning, they allow for a significant reduction in the L1 trigger rate of up to 10 kHz and an improved efficiency for electron triggers requiring isolation. Figure 3 shows the 2022 performance of the legacy trigger in comparison to the new one for the barrel region.² At the time of submission of this manuscript, the final optimisation of the Run-3 L1Calo trigger system was still ongoing.

²For Run 3, the L1 threshold nomenclature is changed so that the threshold value corresponds to a 50% efficiency point, not the efficiency plateau as was the case in Run 2. Thus L1 Run-3 triggers with thresholds of 14 GeV for a single muon and 8 GeV for a pair of muons have the same performance, respectively, as a 20 GeV single muon and a 10 GeV di-muon L1 triggers in Run 2.




Figure 3. (Left) Efficiency of the L1 single-object electromagnetic (EM) trigger and (right) the correlation of the L1Calo transverse energies measured by the legacy Cluster Processor system and the Run-3 eFEX. Both measurements are done in the inner EM barrel within $|\eta| < 0.8$. Isolated legacy trigger [5] and eFEX trigger with identification [3] have the same rate. Only statistical uncertainties are shown.

3.2 Level-1 muon trigger

The L1Muon trigger uses hits from the RPCs (in the barrel) and the TGCs (in the endcaps) to determine the deviation of the hit pattern from that of a muon with infinite momentum [22]. For Run 3 the L1Muon transverse momentum (p_T) thresholds have been redefined to improve the performance of both the high- p_T single muon and the low- p_T multi-muon triggers.

To reduce the rate of the low-momentum charged particles in the endcap regions, the L1Muon trigger in Run 2 applied coincidence requirements between the outer TGC station and either the inner TGC stations or the tile calorimeter. In Run 3, the replacement of the Small Wheels by the NSWs allows for a further rate reduction. Additional RPC modules have also been deployed in the inner barrel station for Run 3. These new Barrel Inner Small (BIS78) chambers are located in the transition $1.0 < \eta < 1.3$ region between the barrel and endcap on one side of the detector only and will be utilised to reduce the trigger rate of the endcap region which is not covered by inner endcap TGC stations. These new chambers are considered to be a pilot project for the inner barrel upgrade in preparation for Run 4 (2029–2032) when the second side of BIS78 will be installed.

The upgrade of the endcap trigger processor board, called Endcap Sector Logic, increases the number of TGC thresholds from 6 to 15 and supports the attaching of up to four flags to muon candidates. Figure 4 (left) shows the efficiency of single-muon triggers for the endcap region. The finer granularity of low- p_T TGC thresholds results in a better resolution of the invariant mass of two muons in the Run-3 L1Topo system, crucial for B -physics and light states triggers discussed in section 6.7. The rate of these low- p_T multi-muon triggers can be further reduced by using information on the charge of the muon candidate, which is provided by one of the new flags. As can be seen from figure 4 (right), the accuracy of muon charge identification is close to 100% at low p_T , although it gets worse with increasing p_T , due to the reduction of curvature of the muon path in the magnetic field. The other three flags indicate the following conditions: the muon candidate satisfies the coincidence in all three outer stations of the TGC (the Full station flag); the Inner Coincidence flag described in the previous paragraph (in 2022 only the inner endcap TGC stations in the region of $1.05 < |\eta| < 1.3$ were used to satisfy it, outside this region the flag was always set); the muon candidate does not




Figure 4. (Left) Efficiency of L1 muon triggers in the endcap region for various p_T thresholds. (Right) The efficiency of the charge identification for L1 muons with $p_T > 3$ GeV. Only statistical uncertainties are shown.

pass through regions of weak magnetic field, where the momentum resolution of such candidates is poor (the Hot-RoIs flag). All three flags were required by default for the 2022 single L1 muon triggers. The di-muon triggers use only the Full station flag.

Muon triggers in the barrel region are provided by three concentric layers of RPC doublets, with the trigger decision relying on coincidence logic. For the low- p_T thresholds, the coincidence of three out of four layers in the two inner doublets of the RPC, which was required in Run 2, was relaxed during 2022 by requiring at least one hit in each of the two doublets. This increased the trigger efficiency with a minimal impact on the rates. 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. The number of RPC thresholds remains the same as in Run 2, with a total of six thresholds: three low- p_T (2-station) and three high- p_T (3-station) thresholds. The firmware of the RPC Pad and of the Barrel Sector Logic boards is upgraded to flag the possibility of two muons in a single tower,³ which would otherwise be identified as only one candidate. By propagating this information to the HLT, the identification of close-by muons (e.g. from boosted J/ψ decays) is enhanced, as shown in figure 5 (left). Figure 5 (right) shows the L1Muon efficiency of various single muon triggers in the barrel region with respect to offline muon selection. Their maximum efficiencies are about 5% lower than at the end of Run 2, due to inefficiencies in the RPC detectors mainly caused by leaks in the gas distribution system. The lower L1 efficiency of the 3-station triggers is due to the additional hit requirements at the outer station and the detector coverage.

The MUCTPI is upgraded to provide full-granularity muon RoI information to the L1Topo system and to be able to interface with the new Endcap Sector Logic. A special SL2MuCTPI board [3] is introduced to interface the Barrel Sector Logic with the new MUCTPI. The MUCTPI receives muon candidate information from the barrel and endcap Sector Logic and calculates the overall muon candidate multiplicity for each p_T threshold, taking into account the possible overlap between trigger sectors in order to avoid double counting of muon candidates. Additional logic in the MUCTPI allows for the exclusion of candidates from the RPC feet areas, or to flag Barrel-Only and Endcap-Only muon candidates. The MUCTPI sends to the CTP information for up to 32 combinations of the p_T

³A tower is a region of approximately 0.2×0.2 in η - ϕ space comprising four possible RoI positions which are read out by a Pad Logic board.




Figure 5. (Left) The efficiency of trigger reconstruction with respect to offline muon selection as a function of distance ($\Delta R_{\mu\mu}$) between two close-by muons from J/ψ decay in the barrel region for a L1 trigger with two muon candidates of $p_T > 10$ GeV and a L1 trigger with a single Barrel-Only (BO) muon candidate of $p_T > 10$ GeV and close-by flag (M). (Right) Efficiency of L1 muon triggers for the barrel region for various p_T thresholds. Only statistical uncertainties are shown.

thresholds and flags, which is an increase from the six available in Run 2. Combinations with different thresholds in the barrel and endcap regions can be also implemented.

The commissioning of the L1Muon trigger proceeded systematically in several steps. It started with the commissioning of the upgraded endcap trigger system using test pulses injected into the on-detector electronics. This allowed for the detection of hardware problems and incorrect fibre connections, as well as the adjustment of the trigger timing and the validation of the trigger logic. Subsequently, the use of cosmic rays provided a full integration test with the barrel and endcap trigger systems with the new MUCTPI and the CTP. The commissioning of the upgraded L1Muon system (without the NSW) was finalised in 2022, while the commissioning of the NSW and the BIS78 RPC chambers is still ongoing at the time of writing.

3.3 Level-1 topological trigger

The addition of the L1Topo trigger in 2016 allowed for both the reduction of energy thresholds without an increase in trigger rates as well as a higher complexity of the algorithms available at L1 through topological selections. This resulted in a significant improvement of the background rejection and enhanced acceptance of physics signal events, despite the increase of luminosity during Run 2. A technical description of the Run-2 L1Topo system can be found in ref. [23].

The principle of operation of the L1Topo system remains the same between Run 2 and Run 3. The L1Topo system receives Trigger OBjects (TOBs) containing kinematic (e.g. E_T , position) and further qualifying information (e.g. flags) from the L1Calo and L1Muon systems and applies topological selections. The Run-3 upgrades of both the L1Calo and the L1Muon systems have an impact on the L1Topo, since they provide the input data directly for the topological selections. The outputs from the L1Calo and the MUCTPI are remapped in the Topo-Fiber Optic eXchange (TopoFOX) module before arriving at the L1Topo, contrary to the Run-2 system where the Common Merger Module (CMX) and MUCTPI2Topo [24] were needed for e.g. TOB sorting. In Run 3, the L1Topo hardware is upgraded with three new modules (TOPO1, TOPO2 and TOPO3) designed to deal with the new input format. Each module has two Field Programmable Gate Array (FPGA) processors on which to run the algorithms.

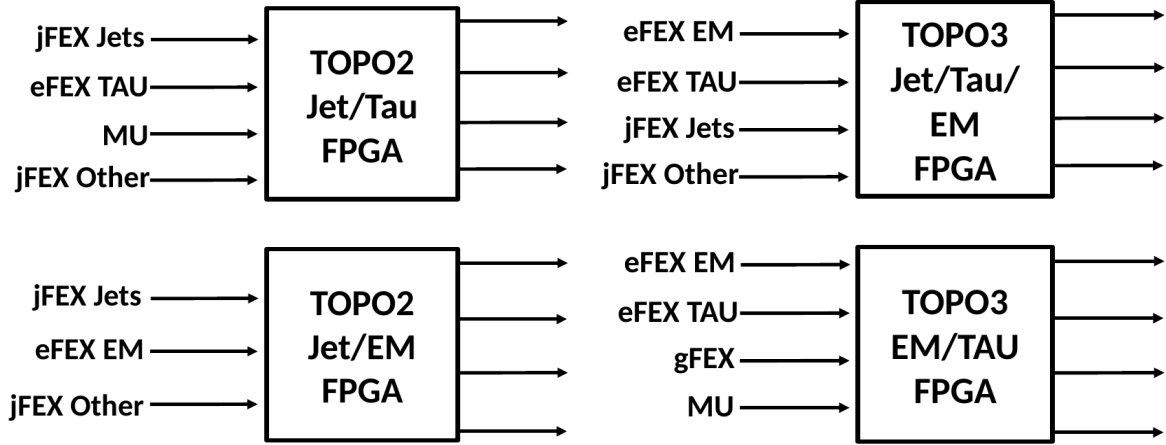


Figure 6. Distribution of L1 TOB inputs for TOPO2 and TOPO3 modules.

In Run 3, the selections based on the multiplicities of the L1Calo objects passing pre-defined energy thresholds are performed in the TOPO1 module. 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 number of available thresholds and multiplicity bits per L1Calo object type is configurable in L1Topo firmware and could be modified in the course of Run 3 according to the trigger menu requirements discussed in section 4.

The TOPO2 and TOPO3 modules are used for the topological selections as in the Run-2 L1Topo system [23]. Available inputs for these two L1Topo models, shown in figure 6, impose constraints on the possible combinations of L1Calo and L1Muon TOBs in topological algorithms. The L1Topo algorithms are distributed across the FPGAs as evenly as possible to achieve optimal resource usage. The configuration of algorithms with their parameters is stored in the trigger menu. The L1Topo algorithm decisions are transmitted to the CTP and are limited to 32 bits per FPGA.

The L1Topo legacy system consists of two Run-2 L1Topo modules with their original Run-2 firmware configuration, but with inputs only from the legacy L1Calo system and not from L1Muon. The decisions of only 22 out of 113 topological algorithms that were implemented in Run 2 [23] are sent to the CTP in Run 3. The L1Topo legacy system was used to collect physics events in 2022, while the upgraded system underwent commissioning. When the three new L1Topo modules are fully validated, the two L1Topo legacy modules will be disconnected.

3.4 HLT software

The ATLAS software framework Athena [25] was used in Runs 1 and 2 at all stages of the event data processing path, from detector simulation to event reconstruction and physics analysis (referred to as “offline”), as well as for the HLT selections (referred to as “online”). Athena is based on the inter-experiment framework Gaudi [26]. In Run 2, memory limitations of the simulation and reconstruction workflows led to the development of a multi-process event processing model (AthenaMP), which was adopted as an intermediate solution, also at the HLT. In this approach, the main process is forked after initialisation into a number of worker processes equal to the number of events which should be processed in parallel. Each worker processes events independently using a single thread, sharing read-only memory with other workers via the copy-on-write mechanism. Thanks to this mechanism,

a reduction of the overall memory requirements was possible, with respect to the instantiation of totally independent processes.

A typical online reconstruction sequence makes use of dedicated fast trigger algorithms to provide early background rejection, followed by more precise and CPU-intensive algorithms that are similar or identical to those used for offline reconstruction to make the final selection. Reconstruction algorithms process detector data to extract features and hypothesis algorithms test the selection hypothesis for all active trigger chains, where a chain consists of a L1 trigger item and a series of HLT algorithms organised into distinct steps that reconstruct physics objects and apply kinematic selections to them.

For Run 3, the ATLAS software framework was adapted to support multi-threaded execution (AthenaMT), based on the concurrent-processing version of Gaudi, named *GaudiHive* [27] which is itself based on Intel’s Threading Building Blocks (TBB) library [28]. HLT requirements were taken into account during this transition [29]. AthenaMT allows greater memory sharing across compute cores than is possible with AthenaMP and, consequently, greater flexibility and efficiency when running on hardware with limited memory per core. It is also a prerequisite for the eventual use of compute accelerators such as GPUs in online and offline data processing. All data processing steps (except event generation) will use AthenaMT in Run 4.

Three types of parallelism are included: inter-event parallelism (multiple events are processed in parallel), intra-event parallelism (multiple algorithms can run in parallel for a single event) and in-algorithm parallelism (algorithms can internally utilise multi-threading and vectorisation). The execution order of the job’s algorithms is determined by the input data required, and output data produced, by each algorithm. These data dependencies must be exposed by the developer to the AthenaMT scheduler via “data handles”, which are used by the scheduler to appropriately order the algorithms for execution. The execution depends on the configured number of threads (made available to execute algorithms) and event slots (made available to concurrently process events). Asynchronous or time-varying data (conditions data) whose lifetime can be longer than an event are handled by the framework with special conditions algorithms, which are scheduled appropriately by the framework so that the required conditions data are available when needed.

To take advantage of the major changes in AthenaMT, the ATLAS HLT framework was to a large extent redesigned and rewritten. While in the Run-2 HLT framework all trigger algorithms were implemented using an HLT-specific interface (with offline algorithms requiring a wrapper), the Run-3 framework requires no HLT-specific interfaces and takes full advantage of the AthenaMT scheduler and other AthenaMT components to control event execution in the HLT. These modifications eased online integration of the offline reconstruction developments due to the seamless integration between the offline and HLT frameworks.

The development of the AthenaMT framework was a common project of the offline and HLT groups from the start [29], incorporating the key principles of the HLT event selection in ATLAS that remain unchanged from Runs 1 and 2:

1. Where possible, trigger selections make use of partial-event data which are processed step-wise (from low to high granularities) using regional reconstruction inside RoIs corresponding to a part (or the whole) detector.
2. Event processing is terminated as soon as the event is known to have failed all active trigger selections.

These two key principles require some extensions to AthenaMT, namely the concept of *EventViews* to accommodate the regional reconstruction as well as the introduction of the control flow to allow for early rejection. *EventViews* allow for an algorithm to be executed multiple times in a single event on either partial or full event data, making use of already existing framework functionality to execute the reconstruction steps. Their execution is prepared by Input Maker algorithms that are scheduled to run before the reconstruction algorithms. Subsequent combinatorial hypothesis algorithms allow for topological selections to be defined between sets of candidate physics objects. HLT-specific filter algorithms reside between the steps of the step-wise HLT processing. They provide the possibility of early rejection by gating which steps are allowed to be executed on any given event.

The scheduling of the execution of algorithms in the HLT is augmented by a control-flow logic determined when the HLT configuration is prepared, applying an additional layer of steering logic to the event processing. This layer is defined on nested lists of algorithms and is in addition to the basic execution steering based on each algorithm's input and output data. Figure 7 shows an example of a data flow graph. The control flow in conjunction with the data dependencies ensure that for a given step, its components are executed in a fixed order: filter \rightarrow Input Maker \rightarrow reconstruction \rightarrow hypothesis algorithms. The additional control logic is required because each processing step can have multiple Filters: for example, the step 1 electron and muon paths in figure 7 each have two feeding Filters (electron, electron+muon) and (muon, electron+muon), respectively. These paths may activate in a given event due to one or both of these Filters passing. The control flow creates a graph representation of all possible execution paths at configuration time. The graph is built from a list of all physics selections configured to be executed and does not change during run time. However, each trigger chain corresponds to one path through the graph, and these chains can be individually enabled or disabled during run time or executed on only a fraction of events.

The HLT software configuration is stored in the trigger-specific Oracle database in JSON [30] format with blobs of entire JSON job configurations being uploaded to the database. This allows for a considerable simplification of the database schema compared to Run 2 [20] allowing further improvements to associated tools and metadata. Several tools allow developers easy inspection, modification or downloading of JSON files of database configuration data for debugging of problems and reproducing setups. The trigger configuration service in AthenaMT can be either run from the trigger database, from JSON files or from standard ATLAS python-based configurations.

In 2022 the HLT software was used in a mixed multi-processing/multi-threaded configuration with two event slots per worker during commissioning and low-luminosity data-taking periods. For the MC production the configuration was pure MT with typically eight event slots per grid job [31]. The HLT software performance in 2022 is further detailed in section 8.

3.5 DAQ system

The DAQ system, as shown in figure 1, supports the operation of the two trigger levels, beginning 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 systems. In Runs 1 and 2 the off-detector stage was performed in detector-specific custom hardware modules called Read-Out Drivers (RODs), typically sitting in VME crates. On receipt of a L1 accept, these RODs read out their data to the first common stage of the DAQ system, the Read-Out System (ROS). The ROS consists of a set of commodity server machines hosting custom-built I/O cards. These




Figure 7. Example of control flow for menu processing. The control flow graph is created during initialisation and the steps are executed based on the available data. If a filter passes, processing continues through the next steps until all filters fail where the processing of the following steps is stopped. If the last step is reached with a chain passing all of its steps the event will be accepted.

receive data from detector front-end electronics and store them in internal memory buffers, called Read-Out Buffers (ROBs). The Data Collection Manager (DCM) orchestrates the data flow from the ROS through the HLT Multi-Process Processing Units (HLT MPPUs). During processing the HLT requests data from the ROS as needed before either accepting an event for permanent storage or requesting its deletion from the ROS buffers. In 2022 a dedicated HLT computing farm of around 56k (112k with hyperthreading) cores was running up to 90k HLT MPPU selection applications. Every processing node hosts one DCM and one or a few HLT MPPU selection applications, depending on the chosen multi-process and multi-threading configuration. Individual events are assigned to processing nodes in the computing farm by a HLT SuperVisor (HLTSV) application according to available free event nodes on each machine.

New in Run 3 is the introduction of the Front-End Link eXchange (FELIX) read-out system and software, running on commodity servers (SW ROD). These new DAQ systems are integrated into the read-out path for those detector systems with new or upgraded front-end electronics (LAr digital trigger, L1Calo, L1Topo, NSW). The interface to the HLT is identical between ROS and SW ROD, with event data routed on-demand to HLT processing nodes in both cases. While the ATLAS data flow architecture remains unchanged for Run 3, the change to a multi-threaded event processing environment required changes in the interface between the HLT MPPU and the DCM, and in the event assignment from the HLTSV. Once HLT processing has been completed, accepted events are sent to a dedicated cluster of servers (known for historical reasons as Sub-Farm Outputs (SFOs)) for packing, compression,

and transfer to offline storage. To accommodate the increased average rate of physics data to be written out, the Run-3 physics bandwidth to permanent storage has been doubled to allow for up to 8 GB/s.

3.6 Validation strategy

The ATLAS software is being continuously developed to adapt to changing requirements and conditions, fix defects, improve resource usage and in the particular case of the software used for recording data, to further optimise the performance of the selection algorithms. As the software is evolving, its validation both in terms of the signature reconstruction performance (see sections 5 and 6) and the trigger menu functionality and execution speed (see section 4) is of great importance to ensure the reliability and predictability of its performance.

Validation jobs are run nightly using the ATLAS Release Tester (ART) system [32] to monitor the status of the trigger software. To facilitate the monitoring and tracking of the ART test results, an automated system has been developed. This system analyzes any log files produced by the tests, looking for errors and warnings. It is able to catch any newly introduced changes on a daily basis, including changes to the effectiveness of the software system with respect to time constraints and allocation of resources. Any newly introduced failures are flagged and followed up with the corresponding developer. By tracking the history of the number of events accepted by a given trigger, it is possible to also monitor the evolution of the trigger algorithm performance and detect unexpected degradation or changes immediately.

In addition to the daily validation, the trigger software is periodically validated using high statistics simulated samples (e.g. $t\bar{t}$ as well as other more specific physics signal processes) and collision data. For both types of validation, sets of histograms representing the performance of the HLT and signature reconstruction are produced and checked against a reference. The main purpose of validating with simulated samples is to check the reconstruction performance, including new developments and to validate new trigger chains that are foreseen to be used online in terms of selection efficiency. Before a new software version is deployed on the HLT farm, it is used to process an enhanced-bias data set [33]. The enhanced-bias data set consists of roughly one million events, and is collected with filters that ensure an even representation of events with different topologies and rarities. This data reprocessing is, due to its large statistics, executed on the LHC computing grid using the ATLAS production system [34]. Not only is the performance checked in terms of reconstruction efficiency and rates of trigger chains, but also the performance in terms of resource usage (such as run time, memory consumption, etc.) is checked to ensure smooth data taking. This is important in order to achieve an overall high data quality, and efficient data taking. Further information about the assessment of the data quality can be found in ref. [20].

4 Trigger menu

Events are selected to be recorded if they satisfy the conditions of one or more trigger chains. The list of trigger chains used for data taking is known as a trigger menu, which also includes prescales for each trigger chain.⁴ The trigger menu consists of *Physics* triggers, detailed in section 6, and

⁴To control the rate of accepted events and to manage CPU consumption at the HLT, a prescale value, or simply prescale, can be applied. For a prescale value of n , a trigger chain has a probability of $1/n$ to be activated in the event. By default they are randomly generated for every individual trigger. However, a mechanism of coherent prescale sets exists for defining

Auxiliary triggers, detailed in section 7. Triggers which use information from more than one object type are called *Combined*.

Physics triggers are used for physics analyses and can be subdivided into the following categories:

- *Primary* triggers, which cover all signatures relevant to the ATLAS physics programme and are typically unprescaled.
- *Support* triggers, which are used for efficiency and performance measurements, background estimates, or for monitoring. These are typically operated at a small rate (of the order of 0.5–1.0 Hz each) using prescale factors. About 15% of the HLT bandwidth in Run 2 was dedicated to support triggers.
- *Alternative* triggers, using alternative (sometimes experimental or new) reconstruction algorithms complementary to the primary or support selections, and often heavily overlapping with the primary triggers. They are often used, for example, as part of the commissioning process for future primary triggers.
- *Backup* triggers, with tighter selections and lower processing or output rate. They can replace the relevant primary triggers if their CPU usage or output rate becomes too high. These triggers require almost no additional computing resources or output rate as they select a subset of the primary triggered events.

Auxiliary triggers can in turn be subdivided into the following categories:

- *Calibration* triggers, used for detector calibrations.
- *Cosmic ray* triggers.
- *Beam-induced background* triggers, which are recorded in LHC bunches with single beam or no beam present.
- *Noise* triggers, which are collected by a *random* trigger at L1.⁵
- *Other* dedicated triggers and algorithms.

To facilitate further processing and analysis, accepted events are recorded into different data sets, called *streams*, which are designed to have minimal overlap. The trigger menu defines the streams to which an event is written, depending on the trigger chains that accepted the event. The five different types of data streams considered in the recording rate budget available at the HLT during nominal *pp* data taking are:

- *Physics* streams: collision events of interest for physics studies. The events contain full detector information and dominate in terms of processing, bandwidth and storage requirements. There are three physics streams for the *pp* data taking: the Main stream, the *B*-physics and light states

groups of triggers whose prescales are correlated. Individual prescale factors can be given to each chain at L1 or at the HLT, and can be any value greater than or equal to one. The value -1 is used to disable triggers.

⁵The random trigger item at L1 corresponds to the read-out from the detector of events chosen at random. They are always prescaled. L1 random triggers on filled bunches can be used to seed specific triggers to overcome potential inefficiencies at L1, while L1 random triggers on the unfilled LHC bunches are typically used for noise and background studies.

(BLS) stream and the Hadronic stream. Events in the Main stream are promptly reconstructed after completion of the first-pass calibration and data quality assessment, as described in ref. [35], while events in the other two, *Delayed*, streams are reconstructed when resources allow it.

- *Express* stream: a very small subset of the physics stream events for prompt monitoring, detector calibration, and first-pass data quality checks. It is fully reconstructed offline within a day of having been recorded.
- *Background* streams: background events of interest for physics and detector performance studies.
- *Debug* streams: events for which no trigger decision could be made are written to this stream. Typical reasons are crashes, timeouts in the HLT processing, and HLT data payloads exceeding set thresholds. These events need to be analysed and recovered separately to identify and fix possible problems in the TDAQ system.
- *Calibration* streams: events containing only partial detector information for calibration of specific subdetectors.
- *Trigger-Level Analysis (TLA)* streams: events sent to this stream contain only specific physics objects reconstructed by the HLT, and optionally only partial detector information. These data are used directly in the corresponding physics analysis (e.g. ref. [36]). The average TLA event size is 4.5 kB in 2022, which is approximately 0.3% of the size of a full ATLAS event.
- *Monitoring* streams: events to be sent to dedicated monitoring nodes for online analysis for, e.g., detector monitoring, but not recorded.

For special data-taking configurations, it is possible to introduce additional streams, such as, for example, the *Enhanced bias* physics data stream, which is used to record events for trigger rate predictions [20, 33]. With the exception of the debug streams, the streaming model is inclusive, which means that an event can be written to multiple streams.

The trigger menu composition and trigger thresholds are optimised for several luminosity ranges in order to maximise the physics output of the experiment and to fit within the rate and bandwidth constraints of the ATLAS detector, TDAQ system and offline computing. The effect of such optimisation on the HLT stream output for a typical run is shown in figure 8. For Run 3, the most relevant constraints are the maximum L1 rate of 100 kHz (unchanged with respect to Run 2) as defined by the ATLAS detector read-out capability, a target average HLT physics output rate of 3 kHz (1.2 kHz in Run 2) and CPU resources of the HLT farm, as detailed in section 8. Substantial increases in the HLT output rate were enabled by the expanded SFO capacity, providing improved coverage for physics, while also allowing space for sustained operation with very long periods (up to ten hours) of constant instantaneous luminosity close to its peak value, which can be seen in figure 8 before 6:30. Such periods make the average Run-3 HLT rate almost equal to the peak rate, in contrast to Run 2 for which the average was about 2/3 of the peak. To ensure an optimal trigger menu within the rate constraints for a given LHC luminosity, prescale factors can be applied to L1 and HLT triggers and these are changed during data taking in such a way that triggers may be disabled entirely or only executed for a certain fraction of events. Supporting triggers usually run at a constant rate. The small event size of triggers in the TLA stream allows it to be recorded at rates of the order of 1–10 kHz while taking up a




Figure 8. An example of (left) the rate and (right) bandwidth output by the HLT streams in a typical 2022 pp run. The total HLT event rate is lower than the sum of the stream rates, because the same events may be written to multiple streams. The peak bandwidth of the TLA stream is approximately 25 MB/s, and consequently barely visible in the right hand figure. The event size per stream can vary up to a factor of 300.

minor fraction of the total HLT bandwidth. This strategy is effective in avoiding high prescales at the HLT for low p_T TLA triggers. Some triggers are enabled only later in an LHC fill when the luminosity and pile-up have reduced and the required HLT farm resources are available. The effects of prescale changes towards the end of the run on the HLT output rate and bandwidth can be seen in figure 8 around 8:30, 11:00 and 14:30. Further flexibility is provided by bunch groups, which allow triggers to include specific requirements on the proton bunches in the LHC. These requirements include paired (colliding) bunch-crossings for physics triggers, empty or unpaired crossings for background studies or searches for long-lived particle decays, and dedicated bunch groups for detector calibration.

4.1 Baseline physics trigger menu for proton-proton collisions

The primary pp menu triggers cover all signatures relevant to the ATLAS physics programme including electrons, photons, muons, taus, jets and E_T^{miss} which are used for Standard Model (SM) precision measurements including decays of the Higgs, W and Z bosons, and searches for physics beyond the SM such as heavy particles, supersymmetry or exotic particles. A set of low- p_T di-muon and di-electron triggers is used to collect B -meson decays, which are essential for the B -physics programme of ATLAS.

The Run-3 trigger menu aims to maximise the physics impact of the Run-3 data set by exploiting the newly implemented detector features, more performant HLT hardware, and algorithmic advancements, while simultaneously maintaining a level of consistency with the Run-2 trigger menu to allow for combined analyses on both data sets. Trigger thresholds at L1 and HLT were generally kept the same as during Run 2, benefiting from improvements to reduce trigger rate. The trigger menu strategy remains focused on assigning the majority of the rate to inclusive triggers rather than analysis-specific triggers. In particular, the Run-3 trigger menu maintains the unprescaled isolated single-electron and single-muon trigger p_T thresholds around 25 GeV, as described in sections 6.1.2 and 6.2.1. Dedicated triggers are added for specific analyses that are not covered by inclusive triggers. The additional available HLT rate compared to Run 2 is dedicated to expanding the physics menu, in the physics and TLA streams, lowering trigger thresholds and including new triggers for previously unexplored phase space. The increase in the SFO capacity enables a larger recording bandwidth which is exploited

with the delayed streams. The breakdown of the approximate rates feeding these streams grouped by trigger signature, as detailed in section 6, is shown in table 1.

In 2022 the baseline physics trigger menu was based on the L1Calo and L1Topo legacy systems. This led to some limitations; in particular, HLT chains seeded from L1Muon topological triggers could not be run until the new L1Topo and L1Muon were commissioned. This mostly affected triggers for the ATLAS B -physics programme.

Table 1. Example breakdown of approximate total rates for physics triggers grouped by signature at luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and $\sqrt{s} = 13.6 \text{ TeV}$. Rates are quoted for the Main, Delayed and TLA streams subtracting off the contributions from the less inclusive streams.

Signature	Rate per stream [Hz]		
	Main	Delayed	TLA
Electron	270		
Photon	120		
Muon	290		
Tau	160		
Missing transverse momentum	140		
Unconventional Tracking	40		
B -physics and light states		240	
Jet	490	460	5000
Jet with b -hadrons	190	160	
Combined	240	50	830

The evolution of the average recording rates of the physics data streams in 2022 is shown in figure 9. The average Run-3 rates of the Main, Hadronic and BLS streams are designed to be 1.6 kHz, 1.1 kHz and 0.5 kHz, respectively. However, in 2022 they used only part of this allocation, which in the case of BLS was dictated by the commissioning schedule.




Figure 9. The average recording rate of physics data streams for the ATLAS pp physics runs taken in 2022. The average of all runs for these three streams is indicated as a dash-dotted line, and the average of the Main stream is indicated as a dashed line.

The triggers with the largest rate contributing to the delayed streams are summarised in table 2. The B -physics triggers are strongly limited by the rate of L1 triggers. In order to maximise the acceptance, L1 triggers with lower thresholds are enabled at the end of the fills, once the luminosity is significantly below its peak value. The rate of the BLS stream is kept approximately constant and thus the fraction of the BLS events in the recorded physics events increases significantly at lower luminosities.

Table 2. Summary of selected triggers in the delayed streams. The VBF di-jet trigger selects events characteristic of Vector Boson Fusion, requiring the presence of a pair of jets satisfying kinematic correlations and an invariant mass greater than 1000 GeV. All jets are within $|\eta| < 2.5$ unless otherwise specified. Rates are given at luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and $\sqrt{s} = 13.6 \text{ TeV}$.

Trigger	p_T threshold [GeV]	Rate [Hz]
VBF di-jet	70 ($ \eta < 3.2$), 50 ($ \eta < 4.9$)	270
Two jets, two b -jets ($\epsilon = 77\%$)	80, 55, 28, 20	160
Six jets	6×35	140
Five jets, one b -jet	$5 \times 35, 25$	50
B -physics di-muon	11, 6	40
$B \rightarrow K^* ee$	5, 5	170

The Run-3 TLA stream was expanded to include photon, muon and b -jet triggers, in addition to the jet triggers, that were used in Run 2 [36]. It is also now possible to record multiple trigger object collections at the same time, permitting the analysis of more complex final states. The content of each individual event in the TLA stream is determined by specialised physics chains targeting these objects and their combinations, as detailed in sections 6.1.4, 6.4.4 and 6.5.3.

The improvements to existing triggers and their Run-3 performance are detailed in sections 5 and 6. In particular, the tracking improvements [37] allowed for the introduction of triggers based on unconventional tracking which are listed in table 3 and detailed in section 6.9. The shown rates are not unique and can have overlap with other physics triggers, as is the case for the displaced tau trigger, for example.

4.2 Baseline physics trigger menu for heavy-ion collisions

The characteristics of HI collisions of Pb+Pb are largely different from pp collisions. Apart from any hard scattering of interest, each HI collision is composed of multiple simultaneous nucleon-nucleon interactions, which generate a sizable underlying event (UE) contribution dominated by soft particle production. As a result of this large UE contribution, the particle multiplicities and the total energy deposited into calorimeters in HI collisions are on average much larger than those in pp collisions, and they also vary significantly from event to event. The event-by-event variations correspond to the variations in the size and geometry of the overlap region of the two colliding Pb nuclei. The variations are characterised by the centrality of Pb+Pb collisions, which is strongly correlated to total transverse energy measured in the forward calorimeter, ΣE_T^{FCal} [38]. Centrality classes are defined by dividing the ΣE_T^{FCal} distribution of minimum-bias inelastic Pb+Pb collisions into percentiles. The 0–10% interval includes events with the largest ΣE_T^{FCal} , corresponding to the most central collisions (largest geometric overlap), while the 90–100% interval includes the most peripheral collisions (smallest

Table 3. Summary of selected unconventional tracking trigger rates at luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and $\sqrt{s} = 13.6 \text{ TeV}$.

Trigger	Rate [Hz]
Isolated high p_T track	1
Large dE/dx	5
Disappearing track	4
Hit-based displaced vertex E_T^{miss} -seeded	1
Hit-based displaced vertex jet-seeded	1
Emerging jet (jet-seeded)	10
Emerging jet (photon-seeded)	10
Displaced single-jet	15
Displaced di-jet	7
Displaced electron	13
Displaced muon	4

geometric overlap). In the trigger, centrality can hence be mapped to the total energy measured in the full calorimeter system or in its forward part, both of which are accessible at L1.

Besides the variations of particle multiplicity from event to event, there is also an azimuthal anisotropy of particle production present in each HI event. This is a result of the initial-state spatial anisotropy of the overlap region leading to sizable anisotropies of particle momenta in the final state.

The trigger menu for HI collisions must be designed to handle the large event-by-event variations and azimuthal anisotropy in each event. On the other hand, HI collisions are operated at a lower centre-of-mass energy per nucleon pair and a lower luminosity compared to pp collisions, so the hard scattering rate is lower in HI collisions. The main goal of the Run-3 HI trigger menu design is to keep p_T thresholds for unprescaled triggers of different signatures as low as possible while minimizing the sensitivity of all triggers to UE contributions. The unprescaled single electron and photon p_T thresholds are set at 15 GeV, while the unprescaled single muon trigger has a p_T threshold around 8 GeV. This strategy ensures the collection of the majority of events with leptonic W and Z boson decays, which are the main source of events for the study of electroweak processes. The lower p_T threshold for the single muon trigger allows for the selection of events with semi-leptonic decays of heavy-flavour hadrons. The unprescaled single jet p_T threshold is set to 85 GeV to collect events with jets for the study of QCD processes. In addition to single physics signature triggers, dedicated multi-object triggers are added: an unprescaled di-muon trigger with p_T thresholds of 4 GeV for both muons to collect events with di-muon decays of quarkonium states; an unprescaled muon and jet trigger with a muon p_T threshold of 4 GeV and a jet p_T threshold of 60 GeV to collect events with b -jets based on soft muons from B -hadron decays.

To reduce their sensitivity to UE contributions, calorimeter-based triggers — such as electron, photon and jet triggers — include a correction for the average energy contributed by the UE. The average UE energy is evaluated per calorimeter layer and cell following the iterative procedure used in offline reconstruction [39, 40]. For each interval of width $\Delta\eta = 0.1$, the UE energy is averaged over ϕ taking into account the azimuthal anisotropy of deposited energy due to the collective motion of particles created in HI collisions. This estimated average UE energy is then subtracted from the

cluster constituent cells. As a result, the cells in the RoI — which are used as inputs to clustering and identification algorithms — contain a reduced systematic bias due to the UE.

In cases where two nuclei do not interact hadronically due to a lack of geometric overlap, they can still interact electromagnetically, producing so-called ultraperipheral collisions (UPC). The large electric charge of Pb nuclei ($Z = 82$) results in a photon flux that is enhanced by a factor of Z^2 compared to pp collisions. Therefore, cross-sections for electromagnetic (photon-photon and photon-nucleus) interactions are significantly increased. The events produced in these processes are typically accompanied by small multiplicities of produced particles, as well as forward rapidity gaps defined by very low activity in the forward calorimeters.

The three physics streams for Pb+Pb data taking differ from those in pp collisions and aim to target different classes of events:

- *physics_HardProbes*: events produced by hard processes in inelastic Pb+Pb collisions and triggered with high- p_T electrons, photons, muons, jets and b -jets.
- *physics_UPC*: events produced by various processes in ultraperipheral Pb+Pb collisions. Triggers in this stream select events with low- p_T electrons, photons, muons and jets as well as specific track multiplicities, in combination with vetos on the total calorimeter energy, and/or deposits in the ZDC.
- *physics_PC/physics_CC*: minimally trigger-biased (minimum bias) inelastic Pb+Pb collisions separated into peripheral (PC) and central (CC) events. Triggers for these streams are seeded off L1 total energy triggers.

The calibration and performance streams largely remain the same.

The HI triggers with muons rely on standard reconstruction described in sections 6.2 and 6.7, and are not discussed further. More details on other signature-specific HI configurations can be found in sections 6.1.5, 6.4.5, 6.5.2 and 6.8.

5 High-level trigger reconstruction

Once an event is accepted at L1, it is processed by the HLT by making use of finer-granularity calorimeter information, precision measurements from the MS and tracking information from the ID. The following sections describe the main HLT algorithms used for ID tracking, calorimeter clustering and muon reconstruction.

5.1 Inner detector tracking

The sequence of track reconstruction steps is briefly summarised below. More detailed information can be found in ref. [41]. The track reconstruction starts with the data preparation in the ID, which reconstructs clusters and space points using information from the pixel and SCT data providers which fetch the raw detector data from the pixel and SCT read-out systems. This step makes use of the RoI mechanism which allows the HLT to request only the data from those silicon modules inside the RoI. The data preparation is then followed by the fast-tracking step, which runs a custom pattern recognition and a fast-track fit. Following this, a precision tracking step takes the tracks reconstructed by the fast tracking and refits them using the offline track reconstruction code. It does not use the

slower offline pattern recognition code. Additionally, these track candidates are extended into the TRT: using hits at larger radii leads to an improved track momentum resolution. The TRT data preparation is performed only for the precision tracking following the extension of the tracks into the TRT.

The Run-3 tracking software is enhanced with respect to Run 2 with several new developments:

- Restriction of the RoI width along the beamline for the track reconstruction for the muon trigger isolation requirement (section 5.1.2);
- Implementation of the bremsstrahlung recovery for the electron triggers (section 5.1.3);
- Track reconstruction in the entire ID volume (full scan) for the jet- and E_T^{miss} -based triggers, followed by the vertex finding with these tracks (section 5.1.4);
- Running of a preselection tracking stage in the b -jet trigger to allow the use of fast b -tagging algorithms prior to execution of the full scan tracking (section 5.1.5);
- Implementation of large radius tracking (LRT) [42] (section 5.1.6).

The HLT minimum-bias triggers (section 6.8) run the full offline pattern recognition and track reconstruction [43].

Except where stated, the standard selection⁶ of offline tracks [43] and objects [44–48] is used for performance studies in this subsection. Efficiencies are measured by taking the correlated ratio of the number of offline reference objects that have a matched trigger track to the total number of offline objects passing the selection. The methodology for tracking performance studies is described in detail in ref. [41].

5.1.1 Tau lepton tracking

The tau lepton trigger tracking runs as a two-stage process, essentially unchanged since Run 2 and described in more detail in section 6.3 and in ref. [41]. First, the fast tracking runs in a core RoI of 0.2×0.2 in η - ϕ space, fully extended along the beamline in the range $|z| < 225$ mm, to identify the leading p_T tracks from the tau decay. This is followed by running both the fast tracking and the precision tracking in a wider RoI (0.8×0.8 in η - ϕ space) to determine whether the tau candidates are isolated, and to accommodate the wider opening angle for tau candidates with three tracks associated with them (3-prong). This RoI is centred on the z position of the leading p_T track identified by the first stage and limited to $|z| < 10$ mm relative to this leading track. The efficiency of the tau lepton tracking is above 99%.

5.1.2 Muon tracking

The muon triggers use the standard tracking sequence, which is essentially unchanged since Run 2 [41] and has an efficiency better than 99% for tracks originating from the beamline with an almost negligible dependence of the efficiency on the pile-up multiplicity. Those triggers in which a muon track is required to be isolated with respect to other tracks from the interaction run the muon isolation (*muonIso*) tracking stage, after the standard muon fast and precision tracking, using a wider RoI

⁶With the addition that for the standard tight offline requirement on the number of holes (e.g. missing hits from active layers) [43], the tracks are required to contain at least one pixel hit.

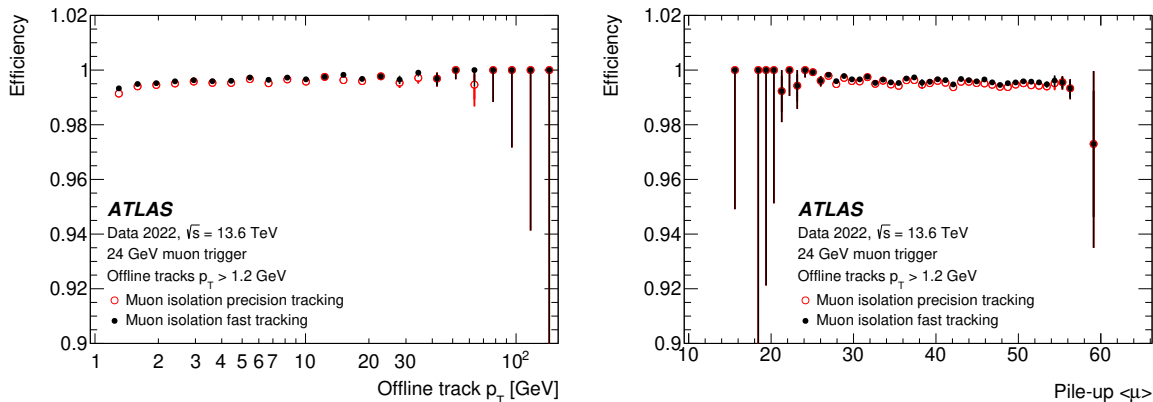


Figure 10. The tracking efficiency for non-muon tracks reconstructed in the *muonIso* RoI centred on the trigger muon selected by the 24 GeV muon performance trigger described in section 7.3, but without selection on the trigger tracks from the *muonIso* reconstruction. Efficiencies are shown for both the fast and precision tracking algorithms as a function of (left) the offline muon transverse momentum and (right) the average pile-up. Only statistical uncertainties are shown.

(0.7×0.7) in η - ϕ space. For Run 3, this wider RoI has a restricted full z -width of 20 mm centred on the z position at the beamline of the muon candidate identified in the first round of precision tracking and after the muon reconstruction. The restriction of the z -width reduces the execution time of the *muonIso* tracking, and is discussed in more detail in section 5.1.7.

The tracking efficiency in the widened RoI used for the muon isolation is shown in figure 10. It is evaluated with respect to offline tracks within the wider RoI once the muon track candidate is removed. The efficiency is greater than 99% across the full p_T range shown. There is a small dependence (less than 0.5%) of the efficiency on pile-up.

5.1.3 Electron tracking

The electron triggers run both the fast, and precision tracking in the same RoI, produced following the fast calorimeter reconstruction, as detailed in section 6.1. The dimensions of the RoI used for the tracking are 0.1×0.2 in η - ϕ space, with the η width being reduced to half of that used in Run 2, with no reduction of the physics performance.

The track reconstruction for electrons can be affected by bremsstrahlung which causes tracks, in particular at lower p_T , to deviate from the expected helical path through the tracking detectors. To improve the resolution of the electron track reconstruction in Run 3 the offline Gaussian sum filter (GSF) algorithm [49] is added to the electron track reconstruction sequence in the HLT.

Figure 11 shows the efficiency of the electron track reconstruction as a function of the transverse energy and the pile-up for offline electron candidates with E_T greater than 14 GeV. To minimise the contribution from non-electron tracks, the tag-and-probe method is used for this study [41]. The slow increase of the efficiency at low E_T is due to the removal of some seed candidates that share hits with other seeds in the fast-tracking stage, reducing the efficiency for bremsstrahlung candidates where the seeds may not lie so closely on a single helix. By construction, the GSF efficiency can not be higher than the precision tracking efficiency since only tracks from the precision tracking are used for the refit for the GSF tracks. The small apparent excess in the GSF tracking efficiency in

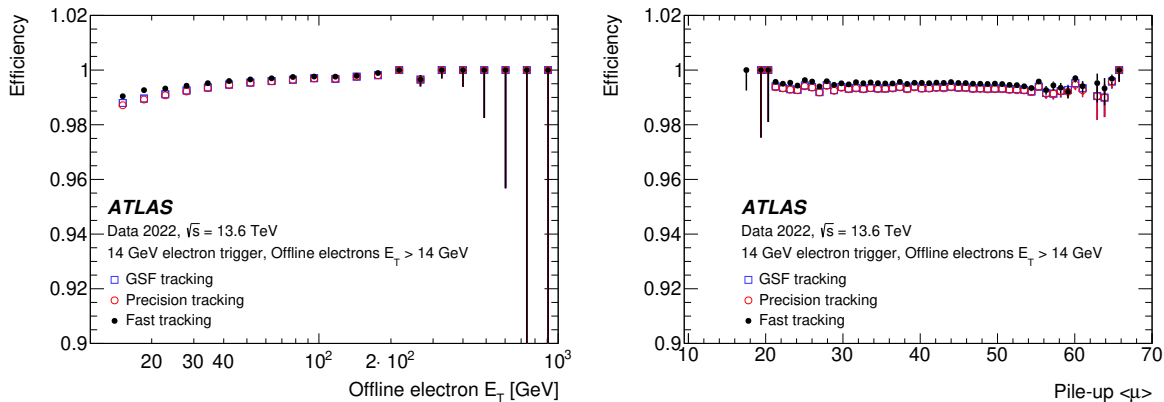


Figure 11. The ID tracking efficiency for electron candidates selected by the 14 GeV electron performance trigger described in section 7.3 which does not use ID tracking information. The efficiency with respect to offline electrons is shown as a function of (left) the offline electron transverse energy and (right) the average pile-up. Only statistical uncertainties are shown.

figure 11 (left) arises since the resolutions of the GSF tracks are better, and as such fewer tracks are excluded by failing the matching criteria. In addition, different performance triggers are used for the GSF and precision tracking studies, such that the event samples used for the calculation of efficiency are not exactly identical.

The effect of bremsstrahlung on the reconstruction of electron candidates is shown in figure 12: as the electron candidates are required to have $E_T > 14$ GeV, the track candidates with p_T below this value would have undergone bremsstrahlung. Because of this both the precision and fast tracking overestimate the $1/p_T$ with a very long tail to positive values shown in figure 12 (top left). This bias becomes progressively smaller at higher p_T as shown in figure 12 (top right). For the GSF tracking however, the distribution is more symmetric, and the bias is close to zero over the entire p_T range.

The resolution with respect to offline of the trigger electron $1/p_T$ is also seen in figure 12 (bottom) and clearly shows the resolution improvement from the GSF tracking with respect to the precision tracking. For reliable estimates of higher offline track transverse momentum, only candidates with $E_T/p_T > 0.8$ are used for the determination of the resolutions. The GSF tracking improves the $1/p_T$ resolution by nearly a factor of two. Similar improvement is observed for the resolutions of the azimuthal angle and transverse impact parameter, but not in the track pseudorapidity and z at the beamline as the latter variables are less sensitive to the bending of the track in the magnetic field.

5.1.4 Full scan tracking and vertex finding

To improve the trigger reconstruction of hadronic signatures (jets, b -jets, E_T^{miss} , etc.) with respect to the offline reconstruction, tracks from the entire ID volume are combined with the calorimeter topological clusters, described in section 5.2, to form particle flow objects (PFO) [50]. This approach improves the jet energy resolution at lower transverse momenta and better separates the hard interaction from pile-up.

The full scan tracking is executed only for the hadronic signatures. These algorithms are activated at a rate close to 14 kHz at $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ after the preselections discussed in sections 6.4–6.6. Even with the reduced input rate after the preselections and with all improvements described below, the full scan fast tracking accounts for 26% of the total event processing time of the HLT, as described




Figure 12. (Top left) The ID trigger residual in $1/p_T$ with respect to offline and (top right) the mean of the residual as a function of the offline track p_T . The resolution in the inverse transverse momentum as a function of offline track (bottom left) η and (bottom right) p_T . Distributions are shown for the fast, precision and GSF tracking algorithms for electron candidates selected with electron performance trigger with $E_T > 14$ GeV and passing the offline electron $E_T/p_T > 0.8$ requirement. Only statistical uncertainties are shown.

later in section 8. An optimisation of the processing time at the cost of a reduced tracking efficiency was thus crucial for the implementation of the full scan tracking for Run 3.

One of the most time-consuming aspects of the full scan tracking is the seed making. In this case the standard seed making [41] is reconfigured to use space-point triplets consisting of pixel-only or SCT-only hits. Additionally, machine learning techniques [51, 52] for the seed selection are adopted to further reduce the processing time. The assignment of SCT hits to the tracks is only performed during the track extrapolation into the SCT with a reduced window to search for these hits in the subsequent layers. The full scan tracking for the hadronic triggers is executed only once per event for the full ID volume, with the resulting tracks and vertices used by each trigger that requires them.

The full scan tracking efficiency for events recorded with a 45 GeV jet trigger is shown in figure 13. The efficiency is approximately 94% at 1 GeV, the threshold used in the pattern recognition, and reaches a plateau at approximately $p_T > 5$ GeV. A slight asymmetry of the efficiency between negative and positive η is due to the mean beam-spot position not being at $z = 0$. Since η for the tracks is always defined with respect to the z_0 position of the point of closest approach of the track to the beamline and not with respect to $z = 0$, two tracks with identical η but different z_0 pass through different parts of the detector. The overall efficiency is approximately 98.5%, but with a large pseudorapidity




Figure 13. The full scan track finding efficiency with respect to offline tracks versus the offline track (top left) p_T , (top right) η , and (bottom) the average pile-up. The efficiency is evaluated for the 45 GeV jet trigger. Only statistical uncertainties are shown.

dependence: it is only 90% for $2.4 < |\eta| < 2.5$. There is also a dependence on the pile-up, falling from 98.5% at $\langle\mu\rangle = 20$ to approximately 97.5% for $\langle\mu\rangle = 54$.

The offline vertex algorithm is also used for the vertex reconstruction in the trigger. It uses tracks from the full scan tracking to identify the likely primary interaction as well as any additional interactions. The vertex with the largest sum of the squared transverse momenta for the tracks assigned to it is chosen as the primary interaction vertex, for both offline and in the trigger. Only the primary offline vertex is considered as a reference for the trigger vertex efficiency study. The efficiency for reconstructing the primary interaction vertex in the trigger is shown in figure 14. For offline vertices with more than six constituent tracks, the trigger vertex efficiency is better than 99.5%. Similarly, the vertex finding efficiency is close to 100% for all z values and all pile-up multiplicities.

The resolution of the reconstructed z position of the vertex is shown in figure 15, illustrating the dependence on the number of tracks from which the vertex is constructed. The resolution is approximately 1 mm for very low track multiplicities and can have values as low as $40 \mu\text{m}$ at higher multiplicities. The resolution is between 40 and $50 \mu\text{m}$ over the full z range.

5.1.5 Tracking in the jets containing b -hadrons

To reduce the rate of full scan tracking for b -jet triggers in Run 3, a calorimeter-only jet-finding preselection step is implemented, followed by a b -tagging preselection stage running the fast tracking.




Figure 14. The trigger efficiency for finding the primary offline vertex, with respect to (top left) track multiplicity, (top right) z position, and (bottom) the average pile-up. The efficiency is evaluated for the 45 GeV jet trigger. Only statistical uncertainties are shown.




Figure 15. The resolution of the primary trigger vertex with respect to the primary offline vertex (left) track multiplicity and (right) z position. The efficiency is evaluated for the 45 GeV jet trigger. Only statistical uncertainties are shown.

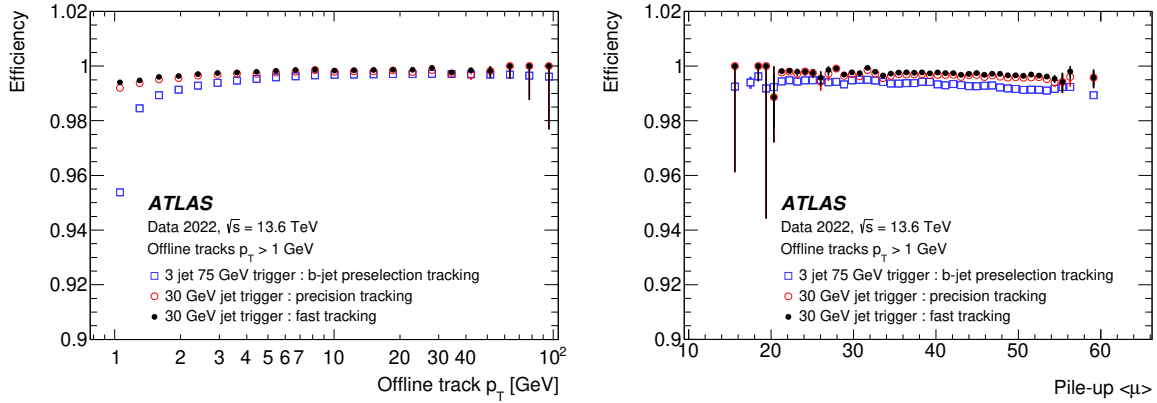


Figure 16. The performance of the tracking for the b -jet signature versus (left) the offline track p_T and (right) the average pile-up. Efficiencies with respect to offline tracks are shown for b -jet preselection stage as well as for the b -jet RoI fast and precision tracking stages. Only statistical uncertainties are shown.

This instance of the fast tracking uses the same optimisations as the full scan tracking but executes in a single composite RoI (super-RoI), which is constructed from all regions centred around $|\Delta\eta| < 0.4$, $|\Delta\phi| < 0.4$ with respect to the jet axes for all jets with $E_T > 20$ GeV identified in the jet preselection stage. This RoI is extended along the beamline by $|z| < 150$ mm. A fast b -tagging algorithm is then applied to the reconstructed tracks, as detailed in section 6.5 and ref. [53]. Shown in figure 16, for the b -tagging preselection stage the efficiency is only 95% at the 1 GeV threshold and does not reach the plateau until approximately 5 GeV. The preselection tracking efficiency also shows a slight reduction with increasing pile-up.

Following this preselection, after the rejection of events where no b -tag is found, the full scan tracking is executed in the full ID volume to determine the primary vertex position and reconstruct the tracks used for the particle flow (PFlow) jet reconstruction. The need to reduce the processing time for the full scan tracking, described in the previous section, necessarily compromises the efficiency to some degree. While this may be acceptable for the PFlow reconstruction, for which it is predominantly intended, this would not be desirable for the final b -tagging where the best possible efficiency for individual tracks is required. Consequently, as in Run 2, the standard RoI-based fast and precision track reconstruction is executed in a separate RoI for each PFlow jet, centred on the jet direction and with the RoI z -position determined from the primary vertex information. These precision tracks are then used for the full b -tagging, described in section 6.5. For Run 3, the track p_T requirement used in the pattern recognition for the b -jet tracking stage is reduced to 0.8 GeV from the 1 GeV used in Run 2. As shown in figure 16, the RoI-based fast and precision tracking results in higher efficiencies for the tracks from b -jets. The efficiency for the precision b -jet tracking is better than 99% over the full phase space, an improvement on the 84% in the 1.0–1.2 GeV range seen in Run 2 [41].

5.1.6 Large radius tracking

LRT is a new feature introduced in the HLT for Run 3. It uses the same algorithms as the previously described standard tracking, but with modified configurations in order to reconstruct tracks at large radii or impact parameter (d_0). These configurations are based on the offline LRT reconstruction [42], which was improved for Run 3 to reduce the number of fake tracks and processing time. As with the standard tracking, LRT is split into fast and precision track reconstruction steps, which can be

performed either inside RoIs or in the entire ID volume (full scan LRT). In the case of RoI-based LRT, e.g. for leptons, the LRT is run by itself. For the full scan LRT, the tracking uses the remaining hits after a standard tracking pass, as is done for the offline LRT.

The standard seeding step [41] is modified as follows: only hits from the SCT are used and the ordering of seeds by impact parameter is removed. Tracking is expanded to cover $|z_0| < 500$ mm and $|d_0| < 300$ mm, but with generally stricter requirements on track quality to reduce the number of fake tracks, such as requiring at least eight hits on the track. This limits tracks to originate before the first SCT layer in the barrel region, at a radius of approximately 300 mm. The same momentum threshold of $p_T > 1$ GeV as standard tracking is used. For LRT in electron and muon RoIs, the size of the RoI is expanded in both η and ϕ , compared to the tracking for prompt leptons, in order to accommodate tracks that do not point to the beamline. The size of the RoI in ϕ is a limiting factor in the efficiency to reconstruct tracks at large d_0 , and a trade-off has to be made between tracking acceptance and computing cost. In order to reduce the processing time for full scan LRT, tracking is restricted to $|d_0| > 2$ mm, tracks are not extended into the TRT, and track candidates are required to have $p_T > 1$ GeV (on top of the p_T requirement on the seeds).




Figure 17. Fast and precision LRT efficiencies for (top) electron and (bottom) muon triggers as a function of offline electron and muon (left) track p_T and (right) d_0 . The corresponding merged collections of standard and large radius offline electron and muon tracks are used. The efficiency for the additional GSF step for electrons is also included. Only statistical uncertainties are shown.

Figure 17 shows the performance of LRT in electron and muon trigger RoIs with respect to offline electrons and muons, respectively. No additional identification requirements are applied beyond the

reconstruction. Data were collected during 2022 using a trigger that does not apply any selection based on the tracking, described in section 7.3. The offline lepton tracks are required to have at least eight silicon hits, $|d_0| > 2$ mm, and $p_T > 10$ GeV. Leptons using both the standard offline track reconstruction and large radius offline track reconstruction are used as the reference. There is a small overlap where standard and LRT tracks matched to the same calorimeter cluster or muon spectrometer segment that is not removed here. Since there are very few SM processes that produce displaced tracks, the tracks in these plots originate from a combination of photon conversions, heavy-flavour hadron decays, long-lived neutral kaon decays, and combinations of hits resulting in fake tracks, which are not necessarily representative of high momentum leptons.

The efficiencies in data are similar to that of tracks in $t\bar{t}$ MC simulation, while the efficiency for leptons in long-lived particle (LLP) signal MC simulation is much higher. Figures 18 and 19 overlay the efficiency of electron GSF- and muon precision-tracking with respect to offline electrons or muons for these samples. Tracking in the trigger for signal-like tracks is much more efficient out to large impact parameter values and less dependent on the amount of pile-up.




Figure 18. The GSF tracking efficiency for the LRT electron trigger with respect to the merged standard and large radius offline electron track collections versus (top left) p_T , (top right) d_0 , (bottom left) η , and (bottom right) the average pile-up for data, semi-leptonic $t\bar{t}$ MC events, and simulated pair production of selectrons with a 1 ns lifetime and a mass of 100 GeV. Only statistical uncertainties are shown.

Data events collected in one run of 2022 were reprocessed to run the full scan trigger tracking configuration of LRT for every event. A collection of offline K_S^0 candidate vertices is produced by selecting both standard and LRT offline tracks with $p_T > 1$ GeV and opposite-charge that form a vertex

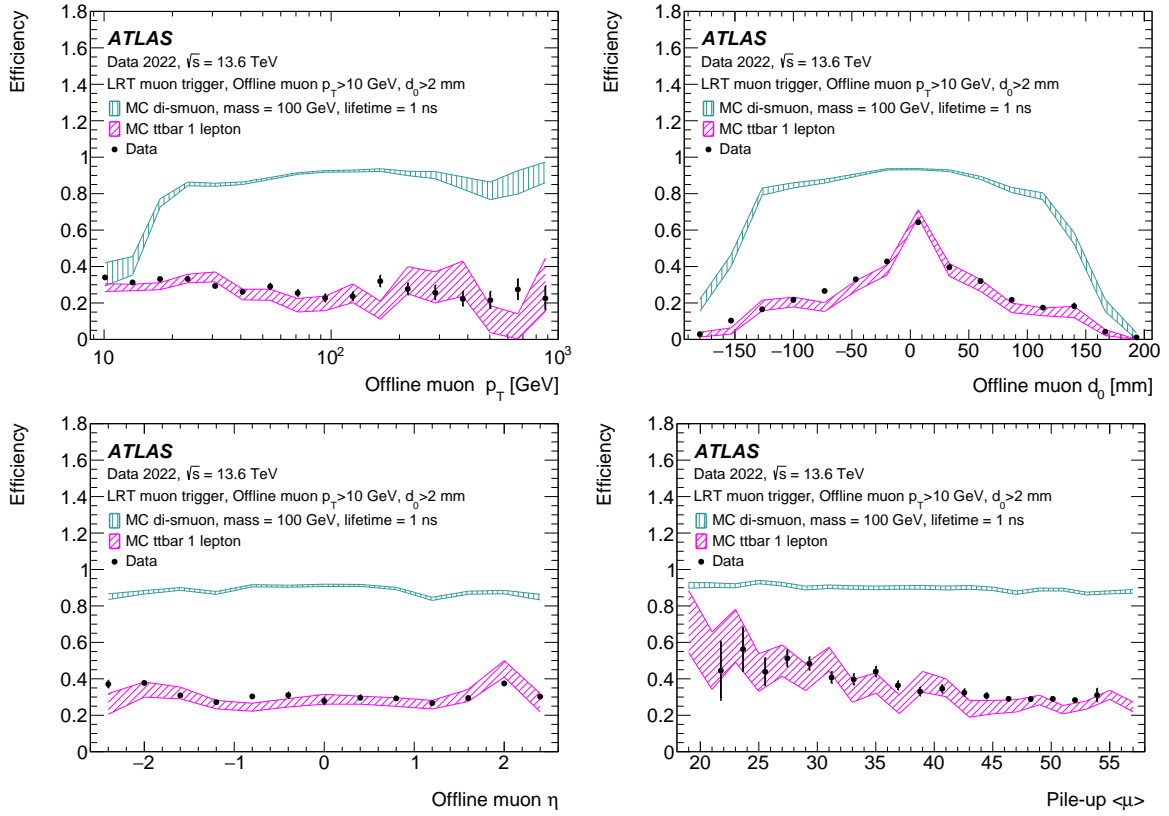


Figure 19. The precision tracking efficiency for the LRT muon trigger with respect to the merged standard and large radius offline muon track collections versus (top left) p_T , (top right) d_0 , (bottom left) η , and (bottom right) the average pile-up for data, semi-leptonic $t\bar{t}$ MC events, and simulated pair production of smuons with a 1 ns lifetime and a mass of 100 GeV. Only statistical uncertainties are shown.

with a mass within 25 MeV of 497 MeV [54]. The offline tracks associated to the K_S^0 candidates are then matched to online standard full scan tracks if they fall within $\Delta R < 0.05$ and $\Delta d_0 < 2.5$ mm of each other. Remaining offline tracks are used in the denominator to compute the online LRT efficiency. Figure 20 shows the efficiency with respect to these remaining offline tracks versus the p_T and d_0 of the offline track. The tracks from K_S^0 decays tend to have low p_T and the number of offline tracks falls off rapidly after a p_T of 5 GeV and a $|d_0|$ of 80 mm. Due to the steeply falling p_T spectrum of the K_S^0 decay products, low momentum particles with poor p_T resolution make up a significant proportion of the high p_T offline reference sample. Many of these are not reconstructed by the trigger leading to an apparent decrease in trigger efficiency at large p_T which is not representative of the actual efficiency for higher p_T particles. The bottom of figure 20 shows the efficiency of matching both tracks of the offline vertices to online tracks, including standard and LRT, versus the reconstructed radius of the K_S^0 vertex. The efficiency drops at radii corresponding to barrel layers of the Pixel detector, e.g. 88.5 mm and 122.5 mm.

5.1.7 Execution time

The ID track reconstruction in the HLT comprises 59% of the total event processing time in the trigger as detailed further in table 5 of section 8.2. The fast tracking is typically the most time-consuming




Figure 20. Efficiency of online full scan LRT with respect to offline tracks associated with K_S^0 candidates that have not been matched to standard online full scan tracks. This efficiency is shown as a function of (left) the p_T and (right) d_0 of the offline tracks. An additional requirement of $|d_0| > 5$ mm is used when computing the efficiency versus p_T to remove tracks with low displacement, which are expected to be reconstructed by the standard tracking. The bottom plot shows the efficiency, combining both standard and LRT, to reconstruct both tracks of the offline vertex versus the radius of the vertex. Only statistical uncertainties are shown.

algorithm due to the combinatorial nature of the pattern recognition stage which has a non-linear dependence on pile-up. The precision tracking timing has a much smaller dependence on pile-up since it depends only on the number of tracks passed to it from the fast tracking. The precision tracking also includes hits from the TRT, which are obtained by the extension of the tracks into the TRT before the final fit is performed.

The timing measurements presented here were recorded from the HLT farm, during a typical physics run from October 2022. Figure 21 shows the mean execution times for the ID trigger-related algorithms for the muon reconstruction as a function of the pile-up interaction multiplicity. Shown are the data preparation times for the silicon detectors prior to the fast tracking, and the execution times for the fast, and precision tracking themselves together with the combined silicon data preparation time and the total time for the full TRT extension used for the precision tracking. The combined pixel and SCT data preparation, at 12 ms at high pile-up is slightly faster than the full TRT extension, at approximately 13 ms. The slowest component from the silicon data preparation is the pixel data provider which fetches the pixel data from the pixel read-out. The slowest component overall, taking approximately 40 ms per RoI at high pile-up is the fast tracking, which exhibits a non-linear dependence




Figure 21. Execution times for the muon track reconstruction in the trigger. Shown are the times for (left) the pixel and SCT data preparation, and (right) the muon fast, and precision tracking, together with the combined data preparation for pixel and SCT, and the total time spent in the TRT extension, all as a function of the average pile-up. Solid lines show ad hoc fits to data to guide the eye. Only statistical uncertainties are shown.




Figure 22. Execution times for the muon isolation track reconstruction shown as a function of the average pile-up. (Left) Detailed breakdowns for the Pixel and SCT data preparation components. (Right) The muon isolation tracking, the combined silicon data preparation and TRT extension times. Solid lines show ad hoc fits to data to guide the eye. Only statistical uncertainties are shown.

on the pile-up multiplicity due to the combinatorial nature of the pattern recognition. The precision tracking takes only approximately 6 ms at high pile-up in the muon RoI.

The timing of the *muonIso* tracking, shown in figure 22, is significantly improved compared to Run 2 [41]. With the reduced z width, the muon isolation tracking algorithms take less than 25 ms in total per RoI at a pile-up of 52. The total combined data preparation for the pixel and SCT takes less than 5 ms, with the longest components being the pixel and SCT data providers, which fetch the data across the network, with the pixel and SCT clustering being somewhat faster. The TRT extension has a combined execution time of approximately 20 ms, in this case, longer than the other algorithms. In contrast, the standard muon fast and precision tracking with $|z| < 225$ mm takes around 40 ms per RoI with approximately 15 ms for the silicon data preparation. The fast tracking alone for the muon isolation in Run 2 typically exceeded 115 ms per RoI at similar pile-up [41].

The execution times for the electron-based algorithms are shown in figure 23. The small size of the electron RoI (0.1×0.2 in η - ϕ space) results in an execution time for the electron pixel and SCT




Figure 23. Execution times for the (left) TRT extension and (right) tracking and data preparation for the electron tracking, shown as a function of the average pile-up. Solid lines show ad hoc fits to data to guide the eye. Only statistical uncertainties are shown.

data preparation of less than 16 ms at a pile-up of 52. The electron tracking takes an additional 20 ms in total for the fast and precision tracking together. The new GSF tracking takes approximately 11 ms which is nearly three times slower than the precision tracking due to the more complicated nature of the GSF reconstruction. Since the GSF tracking is seeded by the tracks from the precision tracking, it has only a weak dependence on the pile-up. The extension of the electron tracks into the TRT and subsequent TRT data preparation and processing result in an additional 11 ms per RoI, as shown in figure 23 (left). The dominant contribution to the TRT processing is 4 ms at high pile-up from the TRT data provider, which includes the fetching of the TRT data from the ROS system.

The execution times for the full scan trigger track and vertex reconstruction are shown in figure 24. The combined data preparation for pixel and SCT is around 250 ms. Even with the optimisation to reduce the processing time, the full scan tracking still takes longer than 1.1 s per event at pile-up of 52, while the vertex finding with the full scan tracks takes only approximately 14 ms.

The execution times for the various tracking stages of b -jet triggers are shown in figure 25. The super-RoI tracking for the b -jet preselection stage takes 140 ms at a pile-up of 52, with the data preparation taking approximately 80 ms in total. The precision tracking is not executed for the preselection. The data preparation for the final RoI-based b -tagging stage is extremely fast (less than 4 ms) since only the data for those parts of the detector not already processed by the preselection tracking need to be reconstructed. The fast and precision tracking in the RoI for this final b -tagging stage takes less than 35 ms per jet RoI at high pile-up. As with the other signatures the fast tracking shows the largest dependence on the pile-up, although the longest contribution is the TRT extension at low pile-up, which shows a similar trend to the precision tracking. It is slightly exceeded by the fast tracking at higher pile-up.

Figures 26 and 27 show the execution times in the HLT for the muon and electron LRT, respectively. For muon LRT compared to standard muon tracking, the data preparation steps take approximately 1.5–3 times as long as standard tracking, except for the pixel data provider which is similar. The muon LRT fast tracking is 2–3 times faster than standard tracking, while the precision tracking is similar. Differences in the processing times of the LRT tracking with respect to the standard tracking are due to the special configuration, such as using larger RoIs or seeding using only the SCT hits. For electron LRT compared to standard electron tracking, the data preparation steps are a few milliseconds




Figure 24. Execution times for (top left) the pixel and SCT data preparation, (top right) fast tracking and total silicon data preparation time for the full scan tracking in jet triggers, and (bottom) for the vertex finding with the full scan tracks, all shown as a function of the average pile-up. Solid lines show ad hoc fits to data to guide the eye. Only statistical uncertainties are shown.

slower, the fast tracking is approximately 1.5–2 times faster, while the precision and GSF tracking are 1.5–2 times slower. The execution time of the track extension step is similar to standard tracking. While the other steps to include the TRT are 1.5–2 times slower.

5.2 Calorimeter reconstruction

The HLT Calorimeter (HLTCalo) software performs the translation of the raw data read out from the LAr and Tile calorimeters into the final software objects, the CaloCells. There are 187,652 CaloCells for the entire detector. CaloCells associate the geometrical information (η , ϕ and the longitudinal layer) of calorimeter cells with their energy, pulse peaking time, hardware gain and quality factor as detailed in ref. [55]. The calorimeter reconstruction algorithms use the CaloCells information to reconstruct clusters of energy for candidate electrons, photons, taus and jet objects as well as shower shape variables useful for particle identification. The HLTCalo software handles both the high rate of regional data requests, tens of kHz in multiple RoIs, and a similar rate of data processing requests which performs reconstruction in the entire volume of the calorimeters for jets and missing transverse energy determination (full scan). The HLTCalo software was adapted for the Run-3 AthenaMT framework, though still not optimally as discussed in section 8.1.

Two different clustering algorithms are used to reconstruct the clusters of energy deposited in the calorimeter: a sliding-window algorithm [2] used for the fast electron and photon reconstruction step and a topological-clustering (*topo-cluster*) algorithm [56]. The topo-cluster algorithm begins with a




Figure 25. Execution times for data preparation, tracking, and TRT extension for (top) the b -jet preselection tracking and (bottom) b -jet RoI as a function of the average pile-up. The detailed breakdowns of the silicon data preparation components are shown on the left. Solid lines show ad hoc fits to data to guide the eye. Only statistical uncertainties are shown.




Figure 26. The data preparation, TRT track extension, and tracking times for the muon large radius tracking algorithms as a function of the mean pile-up. The detailed breakdown of the pixel and SCT data preparation stages are shown on the left, and the tracking times and combined silicon data preparation time and total TRT extension on the right. Solid lines show ad hoc fits to data to guide the eye. Only statistical uncertainties are shown.




Figure 27. The TRT track extension, data preparation and tracking times for the electron large radius tracking algorithms as a function of the mean pile-up. The detailed breakdown of the TRT extension and data preparation stages are shown on the left, and the tracking times and combined silicon data preparation time and total TRT extension on the right. Solid lines show ad hoc fits to data to guide the eye. Only statistical uncertainties are shown.

seed-cell search and iteratively adds neighbouring cells to the cluster if their energies⁷ (E_{cell}) are above a given energy threshold that is a function of the root-mean-square of the expected electronics and pile-up noise (σ). The seed cells are first identified as those cells that have $|E_{\text{cell}}|/\sigma > 4$. All neighbouring cells with $|E_{\text{cell}}|/\sigma > 2$ are then added to the cluster and, finally, all the remaining neighbours to these cells with $|E_{\text{cell}}|/\sigma > 0$ are also added. These σ thresholds (4, 2, 0) can, in principle, be adapted but the numbers above are the same as used in the offline version of the topo-cluster algorithm. The clusters are grown by energy occupancy with no predefined shape. Given the number of searches (via look-up table helpers) of the cell neighbours and calculations of energy ratios, this algorithm is one of the most resource-consuming parts of the HLTCalo reconstruction.

One of the necessary conditions to calculate the energy in a given cell is related to pile-up-dependent effects. There is a small and measurable bunch-by-bunch energy shift which depends on the LHC luminosity profile (i.e. the structure of empty and filled bunches as well as their relative intensities, which is given by the profile of $\langle\mu\rangle$ as a function of bunch-crossing identification). The details of the offline correction procedure are provided in ref. [57]. The only difference for the online application is that this profile is updated only if the pile-up value $\langle\mu\rangle$ changes by more than 5%, which means that a few tens of LBs can pass before the update. This correction improves the resolution of many variables used for particle identification and energy estimation. The calculation of this correction was adapted to the new AthenaMT environment of Run 3 by separating it into a run-long component and an event-by-event component.

The measurements presented here were done in the HLT farm during a typical physics run. The E_T resolution of HLTCalo clusters obtained with different reconstruction algorithms are shown in figure 28. The peaks of all four distributions are consistent with zero within 1%. The residual differences are due to different reconstruction and calibration algorithms. The fast calorimeter reconstruction relies on a sliding window cluster algorithm, which allows for a good electron energy estimate, despite the very short time available (< 5 ms). It has an energy resolution of 4.3% and a longer low-energy

⁷The LAr calorimeter electronics are designed such that signals from pile-up in earlier bunch crossings appear as negative energy.




Figure 28. Transverse energy resolution in percent for HLTCalo clusters obtained with respect to calorimeter clusters reconstructed offline for 2022 data for the following algorithms: full scan topo-cluster for jets and E_T^{miss} , topo-cluster in RoI for taus, topo-cluster in RoIs for electrons (e) and photons (γ), fast sliding-window reconstruction for e/γ .

tail than the topo-cluster reconstruction algorithms, which are shapeless and pick up much more low-energy activity. The topo-clusters used for electrons, photons and taus are reconstructed in an RoI, as discussed in sections 6.1 and 6.3, and those used for jets and E_T^{miss} are reconstructed in the full calorimeter volume (details in sections 6.4 and 6.6). Topo-clusters can be calibrated at the EM scale, for which the energy of an isolated topo-cluster is the sum of its constituent cell energies. Alternatively, local hadronic calibration [58] can be applied as a cell-level correction to improve the mean response to hadronic showers, and is primarily used for tau reconstruction, detailed further in section 6.3, as well as for large-radius jet triggers, detailed in section 6.4. The E_T resolutions are 4.1%, 3.7% and 2.8% for e/γ , tau and full scan topo-clusters, respectively. The processing time of the main HLTCalo reconstruction algorithms per call is shown in figure 29 (left). It varies from a few milli-seconds running at 22 (14) kHz for topo-cluster reconstruction in e/γ (tau) RoIs to a few tens of milliseconds at 23 kHz for full scan topo-clustering. The latter also exhibits a known pile-up dependence as shown in figure 29 (right), unlike the other algorithms. The fast calorimeter reconstruction runs at a total call rate of 132 kHz⁸ and its timing distribution has a double-peaked structure seen in figure 29 (left). The lower peak in the processing time is due to a fraction of calls which do not require time-consuming data requests as data was cached once by the first RoI request and only the reconstruction is run.

5.3 Tracking in the muon spectrometer

The HLT muon reconstruction [22] consists of two steps: the first is fast and trigger specific, while the second is based on precision reconstruction. The precision reconstruction makes use of the same software as the offline muon reconstruction with some adaptations for online running. Muon candidates are reconstructed from combined tracks in the MS and the ID subdetectors. Most of the muon triggers are based on combined muon candidates. MS-only candidates are used to trigger particular topologies, such as LLPs where there might be no corresponding ID track.

⁸It runs at L1 EM rate of 46 kHz for about 2.9 RoIs per event on average.




Figure 29. (Left) Processing time per call for the following algorithms: topo-cluster in RoIs for e/γ , topo-cluster in RoIs for taus, fast sliding-window reconstruction for e/γ , full scan topo-cluster for jets and E_T^{miss} . (Right) Processing time as a function of the average pile-up for the full scan topo-cluster algorithm. A solid line shows an ad hoc fit to data to guide the eye.

In the fast reconstruction stage, each L1 muon candidate is refined by including the precision data from the MDT chambers in the RoI defined by the L1 candidate. A track fit is performed using the MDT drift times and positions as well as the sTGC and MM chambers in the endcap regions, and a p_T measurement is assigned using lookup tables, creating MS-only muon candidates. The MS-only muon track is back-extrapolated to the interaction point and combined with tracks reconstructed in the ID to form a combined muon candidate with a refined track parameter resolution.

In order to recover the efficiency in low p_T di-muon topologies, e.g. B -meson decays where multiple muons arrive close together in the MS, a new *inside-out* algorithm was developed for Run 3. The MS-only back-extrapolated tracks are used as an RoI to reconstruct the ID tracks, which are then extrapolated to the MS and used as seeds for the fast MS-only muon reconstruction. Collimated di-muon trigger candidates can hence be distinguished.

The precision stage follows the same strategy, combining tracks in the MS and the ID subdetectors to reconstruct the muon candidates. It starts from the refined RoIs identified in the fast reconstruction step to form muon candidates using information from the MS detectors. They are then extrapolated to the interaction point and combined with ID tracks. If the MS track cannot be matched to an ID track, combined muon candidates are searched for by extrapolating ID tracks to the MS.

The full scan mode is used to find additional muons that are not found by the RoI-based method mainly due to L1 inefficiencies (‘noL1’ triggers). In the full scan mode, muon candidates are first sought in all muon detectors. Then, RoIs are constructed around the found MS tracks and ID tracks are reconstructed within these RoIs. The same combination procedure as for the RoI-based method is used to construct combined full scan muons. Given the high CPU demand of the full scan reconstruction, it is only executed in multi-object triggers with at least one of the trigger objects found by an RoI-based algorithm.

The performance of the muon tracks at the HLT level compared to the offline muon reconstruction is illustrated in figure 30. The p_T resolution of the trigger muon candidates is measured with $Z \rightarrow \mu\mu$ candidate events using the tag-and-probe method [22]. Events are required to contain a pair of muons with opposite charge and an invariant mass within 10 GeV of the Z pole mass and the offline reconstructed muons are required to pass the *medium* identification requirements. The distributions of

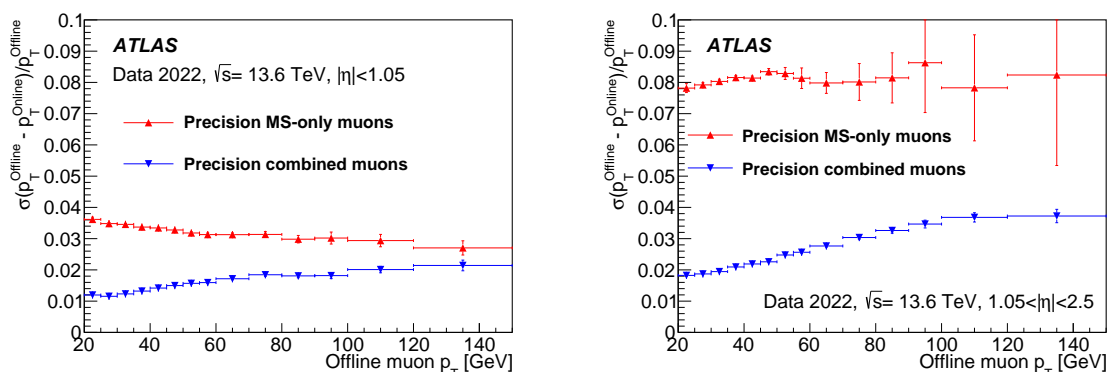


Figure 30. Width of the residuals for p_T as a function of the offline muon p_T for the precision MS-only and combined algorithms in (left) the barrel and (right) endcaps. Only statistical uncertainties are shown.

the relative residuals between online and offline track parameters are constructed in bins of offline muon p_T and the width of the distribution in each bin is obtained by means of a Gaussian fit. The resolution for combined trigger muons is better than the resolution for MS-only muons as expected, thanks to the higher precision of the ID track measurements, especially at low transverse momentum. The relative p_T resolution is about 3–4% (8%) and 1–2% (2–4%) in the barrel (endcap) region for precision MS-only and combined muons, respectively, and tends generally to degrade towards higher- p_T . The MS-only resolution in the endcap region is expected to improve after the NSW commissioning is completed.

The timing measurements presented in were recorded from the HLT farm, during a typical physics run from November 2022. The processing times per RoI are shown in figure 31 for the fast MS-only and fast combined algorithms. The call frequencies are 37 kHz and 23 kHz, respectively. Figure 32 shows the most CPU consuming steps of the precision reconstruction: segment-finding algorithm, MS-only track building, MS and combined muon candidate building and inside-out recovery. The call frequency of the segment-finding, track-building and MS muon candidate building algorithms is 11 kHz, while they are 2.7 kHz and 0.34 kHz for the combined muon candidate building and inside-out recovery, respectively. The large tails of the distribution of the fast combined and precision combined algorithms are related to the algorithms not being fully optimised for Run-3 conditions in 2022. The performance has since been optimised.

6 Physics triggers

The final event selection is based on trigger signatures, such as leptons, hadrons, and global event quantities such as missing transverse momentum. They are formed by placing different selection criteria on the various reconstructed objects. The selection criteria and performance of the various trigger signatures is described in this section, highlighting the differences with respect to the Run-2 selection and performance of primary triggers during 2022.

6.1 Electrons and photons

6.1.1 Electron and photon trigger reconstruction and selection

Electron and photon reconstruction at the HLT is performed on each EM RoI provided by L1, which satisfies the E_T requirement, and any other L1 selection requirements specified by the trigger menu (e.g.




Figure 31. Processing times per ROI for the fast MS-only and fast combined algorithms. The mean time of each algorithm is indicated in the legend. The last bin of the distribution includes the overflow events.




Figure 32. Processing times per ROI for the most CPU consuming precision steps: segment-finding and track-building algorithm (left), muon candidate building and inside-out recovery (right). The mean time of each algorithm is indicated in the legend. The last bin of the distributions includes the overflow events.

isolation) of at least one active HLT chain. In the HLT, fast calorimeter algorithms are executed first, allowing precision algorithms to run at a reduced rate later in the trigger sequence. The reconstruction of candidate electrons and photons uses the sliding-window algorithm with rectangular clustering windows of size $\Delta\eta \times \Delta\phi = 0.075 \times 0.175$ in the barrel and 0.125×0.125 in the endcaps. The fast calorimeter selection step has three implementations [5]. The default fast calorimeter selection step for electrons uses a neural-network-based *Ringer* algorithm [59], which uses as input energy sums of all the cells in 100 concentric rings centred around the most energetic cell in each calorimeter sampling layer. In Run 2 it was used only for triggering electrons with $E_T \geq 15$ GeV, but it is applied from $E_T \geq 5$ GeV in Run 3. The *Ringer* algorithm is optimised in two regions of E_T , between 5–15 GeV with $J/\psi \rightarrow ee$ MC samples and for $E_T > 15$ GeV with $Z \rightarrow ee$ MC samples. For electrons, it is optionally possible to use fast calorimeter selections which take as an input either only the cluster E_T (E_T -based) or the cluster E_T with three shower shape parameters (cut-based) [5]. The electron candidates are then required to have tracks within the ROI (obtained from the fast track reconstruction) matching the corresponding clusters [5]. In contrast, photon candidates are reconstructed using the calorimeter information only, with cut-based selection criteria that are the same as in Run 2 [5].

For Run 3, the HLT electron and photon precision reconstruction becomes closer to the offline reconstruction [44], due to the implementation of a new super-clustering algorithm, described below, and the use of the GSF tracking algorithm for electrons, described in section 5.1.3.

Electron and photon super-clusters are reconstructed in two stages. The first stage consists of finding the seed topo-cluster candidates in the same RoI as used for the tracking: 0.1×0.2 in η - ϕ space. These seed topo-cluster candidates form the basis of super-clusters. The second stage is the identification of topo-clusters near the seed candidates which are identified as satellite cluster candidates. They may emerge from bremsstrahlung radiation or topo-cluster splitting. For a cluster to become an electron super-cluster seed, it is required to have a minimum E_T of 1 GeV and match to a track with at least four hits in the silicon tracking detectors. For photon reconstruction, a cluster must have E_T greater than 1.5 GeV to qualify as a super-cluster seed. If a cluster meets the seed cluster requirements, the algorithm attempts to find satellite clusters, which represent secondary EM showers originating from the same initial electron or photon. A cluster is considered a satellite if it falls within a window of $\Delta\eta \times \Delta\phi = 0.075 \times 0.125$ around the seed cluster barycentre. The energy of the clusters is calibrated using a multivariate technique such that the response of the calorimeter layers is corrected in data and simulation [44].

At the precision step, the electron identification relies on a multivariate technique using a likelihood discriminant, while the photon identification is cut-based. These identifications, as well as isolation requirements which are applied to some triggers, remain unchanged with respect to their Run-2 configurations detailed in ref. [5].

6.1.2 Electron and photon trigger menu

Data taken with electron and photon triggers are used in a wide range of ATLAS physics analyses, from SM precision physics to searches for new physics. The various triggers cover the energy range between a few GeV and several TeV. The 2022 electron and photon trigger thresholds remain the same as in 2018 [5]. The minimum E_T thresholds for the isolated single electron and non-isolated photon triggers are 26 and 140 GeV with rates of 186 Hz and 46 Hz at luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, respectively. To increase the single electron trigger efficiency at E_T of 60 GeV and 140 GeV, triggers with no isolation requirements and looser identification are present at rates of 20 Hz and 2 Hz, respectively. Additional single electron triggers with E_T thresholds at 140 (300) GeV and the cut-based (E_T -based) fast calorimeter selection, instead of the default Ringer algorithm, run at 3 (6) Hz improving sensitivity to merged electrons coming from decays of boosted dibosons [5]. Triggering on low E_T electrons and photons is very challenging because of the high rates at low trigger thresholds. However, this can be mitigated by requiring the presence of multiple electrons or photons in the event, which helps to reduce the trigger thresholds with respect to single electron or photon triggers. The primary di-photon trigger is mainly designed for the efficient selection of events with Higgs boson candidates in the di-photon decay channel. It has a rate of 18 Hz for medium photon identification and trigger E_T thresholds of 35 and 25 GeV for the leading and subleading photons. A second symmetric di-photon trigger with loose identification and the E_T requirement of 50 GeV has a rate of 7 Hz. A di-electron trigger at lower E_T requirement of 17 GeV each and loose identification has a rate of 12 Hz.

6.1.3 Electron and photon trigger performance

The efficiency of the lowest unrescaled single-electron trigger with respect to offline electron candidates is shown in figure 33 for each step of the online reconstruction. This and the next measurements

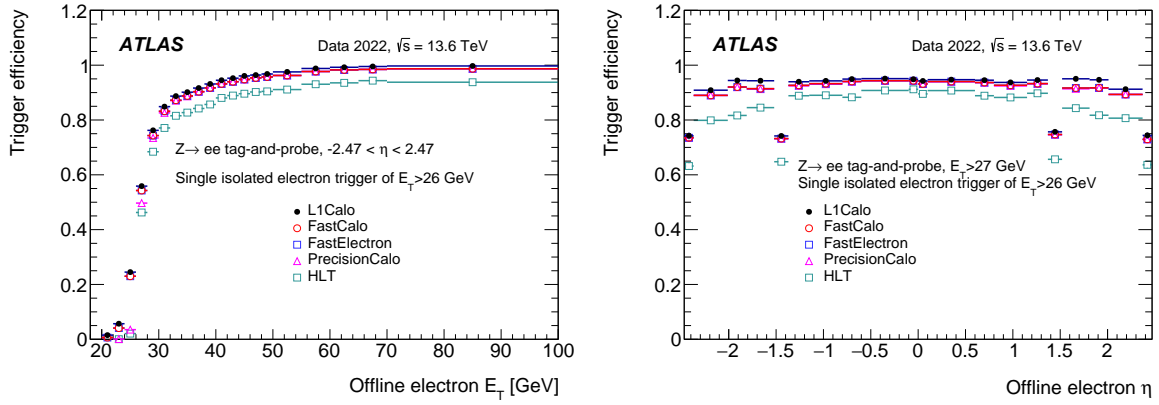


Figure 33. Efficiency of the primary isolated electron trigger with $E_T > 26$ GeV for all steps of the HLT electron reconstruction as a function of the offline electron (left) E_T and (right) η . Efficiency is given with respect to the offline electrons which satisfy the tight identification and isolation criteria. Only statistical uncertainties are shown.

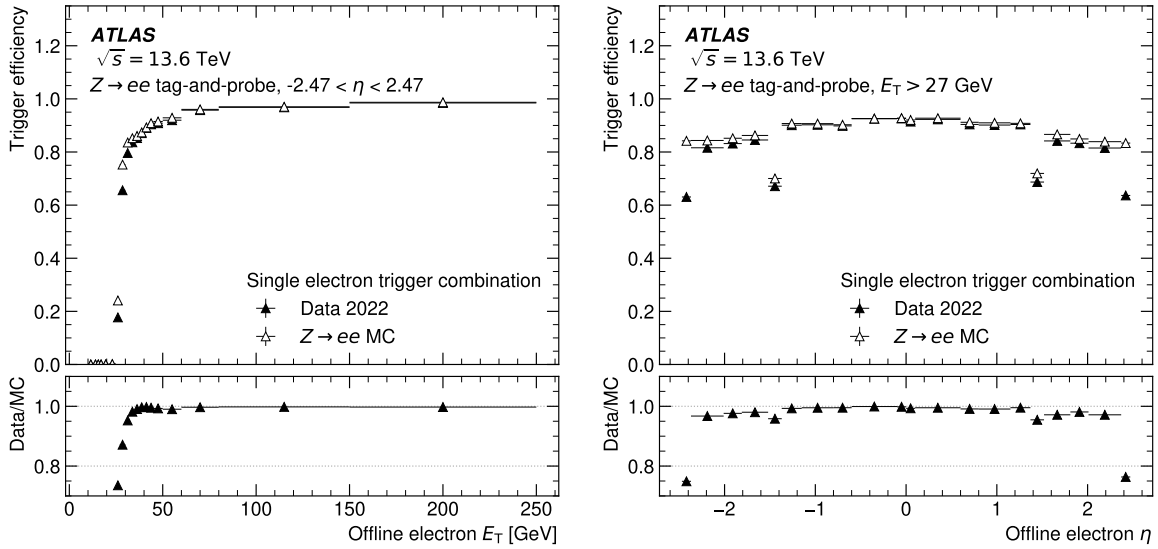


Figure 34. Combined efficiency of three primary single electron triggers as a function of the offline electron (left) E_T and (right) η for both MC simulation and 2022 data. Efficiency is given with respect to the offline electrons which satisfy the tight identification and isolation criteria. Only statistical uncertainties are shown.

are performed using the tag-and-probe method in $Z \rightarrow ee$ events [5]. The HLT inefficiency is due to differences in the online and offline electron identification and isolation [5] and a 1–2% lower efficiency of the trigger precision tracks, shown in figure 11. The efficiency of the combination of all primary single-electron triggers is shown in figure 34. Scale factors derived from the observed data/MC simulation differences are used to correct other MC simulation samples used in data analyses.

Figure 35 shows the efficiencies of the primary di-photon triggers extracted with the bootstrap method [5]. Photon trigger efficiencies are always very high, but while the 25 GeV trigger is fully efficient at 5 GeV above its threshold, higher E_T triggers are not, as can be seen in figure 35 (right). As for the electron triggers, scale factors are used to correct for the observed data/MC simulation differences.




Figure 35. Efficiency of the 25, 35 and 50 GeV legs of di-photon triggers as a function of the offline photon (left) E_T and (right) η for both MC simulation and 2022 data. The ratios of data to MC simulation efficiencies are also shown. Efficiency is given with respect to the offline photons reconstructed with the tight offline identification and isolation requirements; candidates in the calorimeter transition region $1.37 < |\eta| < 1.52$ are not considered. Legacy non-isolated L1 EM trigger with $E_T > 20$ GeV defined in ref. [5] is used as a seed. For (right), only offline candidates with E_T values 5 GeV above the corresponding trigger threshold are used. Only statistical uncertainties are shown.

6.1.4 Trigger level analysis with photons

Hadronically decaying low-mass resonances around the electroweak scale can be probed through composite trigger chains targeting both the final-state di-jet pair and an additional feature, such as an initial-state radiation photon. The Run-3 TLA workflow is upgraded to allow for the recording of HLT-reconstructed photons, either standalone or in association with other TLA-compatible signatures to enable such searches at rates higher than possible with the standard triggers. A new TLA trigger records all jets in events with at least one tight photon [5] with $E_T > 35$ GeV and three additional jets with $p_T > 25$ GeV reconstructed with the PFlow algorithm, described in section 6.4. The photon will, in most events, be reconstructed as a jet as well, making this effectively a trigger for a signature with a photon and two jets. A nominal HLT rate of such trigger is about 700 Hz at luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ at $\sqrt{s} = 13.6$ TeV. This TLA approach allows extended sensitivity of hadronic searches to resonance masses potentially as low as 100 GeV.

HLT photons saved to the TLA stream are decorated with additional variables computed within the HLT precision photon reconstruction and calorimeter isolation trigger sequences. Such variables are intended to enable the derivation of a custom residual photon energy calibration, as well as to constrain the contribution from non-prompt, non-isolated photons in the TLA data set.

Figure 36 shows the energy response of HLT photons spatially matched to offline photons in the E_T range of interest for Run-3 TLA searches targeting photons. HLT photons calibrated with the default 2022 calibration sequence are shown to be already within 2% of offline-calibrated photons across the probed E_T range.

6.1.5 Electron and photon triggers for heavy-ion collisions

Due to the compactness of electron and photon showers in the detector, the standard L1 EM RoIs can still be used in HI collisions to select electron and photon candidates with reasonable rate and purity. Typical E_T thresholds for L1 EM RoIs are 12 and 16 GeV. However, the large UE contribution present in HI collisions leads to significant distortions of the shower shapes and subsequent inefficiency




Figure 36. Energy response in TLA photons spatially matched to the leading offline photons. HLT-level photons are calibrated with the default sequence employed within the HLT reconstruction for the 2022 data taking, while reference offline photons are calibrated according to the Run-2 recommendations [44]. Events are required to satisfy a 25 GeV tight single-photon prescaled support trigger. Only statistical uncertainties are shown.

associated with the electron and photon identification requirements at HLT. To improve the electron and photon trigger performance, the UE subtraction procedure is applied at the HLT.

Two electron HLT sequences are designed as primary physics triggers in the HI menu: a cut-based trigger with an $E_T > 20$ GeV requirement and a likelihood-based trigger with an $E_T > 15$ GeV requirement. For both triggers the UE subtraction is performed at the HLT as described in section 4.2. For the likelihood-based electron identification [5], the standard pp probability density functions are evaluated using the UE-corrected variables. An advantage of the likelihood-based approach is the significant reduction in the output rate in comparison to the cut-based trigger for a given p_T threshold. The likelihood trigger has significantly better purity at the cost of a loss in trigger efficiency.

The primary photon HLT sequence in the HI trigger menu uses a cut-based photon identification with an $E_T > 20$ GeV requirement including the subtraction of the UE contribution. After the UE subtraction procedure, the photon trigger efficiency remains high across the full range of centralities.

6.2 Muons

The reconstruction of muon candidates at the HLT is described in detail in section 5.3. The criteria for selecting muon candidates are dependent on the algorithm used for their reconstruction. The MS-only chains select on the p_T of the muon candidate measured solely by the muon spectrometer, while the combined muon chains apply requirements on the matching between the ID and MS tracks and their combined p_T .

6.2.1 Muon trigger menu

Muon triggers cover a wide momenta range between a few GeV, for B -meson-decay studies, up to several TeV for new phenomena searches. The primary triggers in the muon trigger menu include single-muon triggers with and without isolation requirements, symmetric and asymmetric di-muon and

multi-muon triggers. The Run-3 muon trigger menu is similar to that used in Run 2 [22], accounting for the refinement of L1Muon thresholds, discussed in section 3.2.

The improved suppression of fake muons and the 2022 running conditions allowed for a lowered HLT threshold of the lowest-unprescaled isolated single-muon trigger by 2 GeV to $p_T > 24$ GeV maintaining its rate at about 200 Hz at a luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. As in Run 2, combined muon candidates must fulfil the following track isolation requirement: the scalar sum of the p_T values of tracks within a variable-size cone around the muon (excluding its own track) must be less than 7% of the muon p_T . The track isolation cone size for muons, ΔR , is given by the smaller of $\Delta R = 10 \text{ GeV}/p_T$ and $\Delta R = 0.3$. A non-isolated trigger with $p_T > 50$ GeV helps to increase the efficiency for high- p_T muons. Additionally, a trigger that selects only muons in the barrel region ($|\eta| < 1.05$) reconstructed using MS-only information is available at $p_T > 60$ GeV.

There were no changes to the multi-muon thresholds in 2022 with respect to Run 2. Two combined muon candidates are required with a p_T threshold of 14 GeV each at the HLT with a rate of 24 Hz at a luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. To avoid an efficiency loss due to the limited acceptance of L1, di-muon and tri-muon triggers seeded by the single L1 muon trigger are also present in the menu. These sub-leading muons, with p_T greater than 8 GeV and (4 GeV, 4 GeV) for di-muon and tri-muon triggers, respectively, are reconstructed by the muon full scan algorithm described in section 5.3. Low- p_T di-muon triggers are further discussed in section 6.7.

6.2.2 Muon trigger performance

The muon efficiencies are determined using the tag-and-probe method with $Z \rightarrow \mu\mu$ events. The efficiencies of the combination of the two lowest-unprescaled single-muon triggers are shown in figures 37 and 38. The HLT efficiency relative to L1 is close to 100% both in the barrel and in the endcaps. The L1 muon trigger efficiency is about 60% in the barrel and 80% in the endcap regions for offline medium muons [45] with $p_T > 25$ GeV.

The measured trigger efficiency in figure 37 is lower than that of the expected efficiency in MC. This is due to the L1 muon trigger in the barrel region being simulated with an optimistic lower-bound on chamber efficiency to allow for a reasonable MC efficiency for chambers whose efficiencies are later recovered during data taking. Scale factors derived from these observed differences are used to correct other MC simulation samples used in data analyses.

6.3 Taus

6.3.1 Tau trigger reconstruction and selection

The HLT tau trigger targets only hadronic decays of tau lepton candidates ($\tau_{\text{had-vis}}$). Events with taus decaying leptonically are recorded by electron and muon triggers described in sections 6.1 and 6.2, respectively. The HLT tau trigger reconstruction is subdivided in two steps: a calorimeter-only preselection and track reconstruction with an offline-like selection.

At the first step, the $\tau_{\text{had-vis}}$ candidate is reconstructed purely from calorimeter information. Calorimeter cells inside the RoI identified at L1 are retrieved and the topo-clustering algorithm, described in section 5.2, is executed. Thanks to the full detector granularity and the bunch-by-bunch pile-up corrections, the energies of these reconstructed topo-clusters are very close to the offline ones. These clusters are calibrated with the local hadron calibration (LC) [58] and their vectorial sum is used as a ‘jet seed’ for the reconstruction of the $\tau_{\text{had-vis}}$ candidate. The energy of the $\tau_{\text{had-vis}}$ candidate




Figure 37. Muon trigger efficiencies as a function of (left) p_T and (right) ϕ for $p_T > 25$ GeV of the offline medium muon for the combination of two single-muon triggers in the barrel region. The ratios of data to MC simulation efficiencies are also shown. Only statistical uncertainties are shown.




Figure 38. Muon trigger efficiencies as a function of (left) p_T and (right) ϕ for $p_T > 25$ GeV of the offline medium muon for the combination of two single-muon triggers in the endcap region. The ratios of data to MC simulation efficiencies are also shown. Only statistical uncertainties are shown.

is calculated from the LC clusters in a cone of $\Delta R < 0.2$ around the barycentre of the jet seed. A dedicated $\tau_{\text{had-vis}}$ energy calibration is applied to improve the precision of the energy measurement and follows the offline procedure [60]. Then, a selection on the minimum p_T of the $\tau_{\text{had-vis}}$ candidate is applied, and only the remaining candidates pass to the next step.

The second step of the HLT tau trigger first runs a fast-tracking algorithm, followed by a precise measurement of the tracks associated with the $\tau_{\text{had-vis}}$ candidate and a final $\tau_{\text{had-vis}}$ identification based on a Recurrent-Neural-Network (RNN) algorithm. The fast tracking algorithm is a trigger-specific pattern recognition algorithm that runs in two stages. In the first stage, the leading track is sought in a narrow $\Delta\eta \times \Delta\phi$ around the $\tau_{\text{had-vis}}$ candidate along the entire beamline. In the second stage, additional tracks associated with the $\tau_{\text{had-vis}}$ candidate are sought in a larger $\Delta\eta \times \Delta\phi$ region but in a narrow range around the leading track along the beamline. This strategy is CPU-efficient as it minimises the volume in which the pattern recognition algorithm is executed, as discussed in section 5.1.1.

A precision-tracking algorithm similar to the offline one is run using the tracks identified by the second step of the fast tracking as seeds to measure their properties more precisely. Using these

tracks as well as the calorimeter information, the input variables for the $\tau_{\text{had-vis}}$ identification are computed. Three sets of variables for zero-prong, one-prong and multi-prong $\tau_{\text{had-vis}}$ candidates are used depending on whether the number of precision tracks associated to the $\tau_{\text{had-vis}}$ in $\Delta R < 0.2$ is zero, one, or more than one, respectively. The architecture implementation of the RNN follows closely its offline counterpart, as described in ref. [60]. Finally, tau candidates are required to have up to three tracks within $\Delta R < 0.2$ and up to one track within $0.2 < \Delta R < 0.4$ around the $\tau_{\text{had-vis}}$. In addition, they are required to be identified by the RNN. Given the small increase in event rate and potential efficiency gain, the identification working point is loosened for the $\tau_{\text{had-vis}}$ $p_T > 280$ GeV, and both criteria (identification and number of tracks) are completely dropped when the $\tau_{\text{had-vis}}$ $p_T > 440$ GeV. A different RNN is used to trigger on LLP and trained for zero-, one- and multi-prong taus using a MC sample of LLP.

6.3.2 Tau trigger menu

The tau trigger menu selects a wide spectrum of final states that involve hadronically decaying tau leptons. There are four main categories of triggers: single-tau, di-tau, tau+X (X=light leptons, E_T^{miss}) and events for tag-and-probe performance studies. For each of these categories, the definition of the triggers varies based on (i) the identification requirement, (ii) the p_T threshold applied and, for the multi-object final states, (iii) the eventual presence of a topological cut on the angular distance between the two objects selected in the final state. Dedicated triggers for identifying tau leptons originating from the decay of LLP are included. The lowest threshold of the single-tau unrescaled triggers is 160 GeV and 180 GeV for the standard and LLP taus. The individual rates are about 40 Hz and 50 Hz, respectively, at a luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The overlap fraction of the higher threshold LLP trigger with the 160 GeV standard trigger is 90%. Lower threshold single-tau triggers with or without looser requirements are rescaled at L1 and/or the HLT and are used to support the primary triggers. There are two primary di-tau triggers, which run at rates of 25 Hz and 10 Hz, respectively. The first, main, trigger requires two tau candidates of $p_T > 35$ and 25 GeV and the second, LLP, trigger requires $p_T > 80$ and 60 GeV. LLP tau triggers are discussed further in section 6.9.

6.3.3 Tau trigger performance

The tau trigger efficiency is determined using a tag-and-probe analysis, which selects events in two complementary signal regions [2], one enriched in $Z \rightarrow \tau\tau$ events and another one enriched in dileptonic $t\bar{t} \rightarrow W(e\nu)bW(\tau\nu)b$ decay events. The events for the efficiency measurement are selected using single light lepton triggers as well as dedicated tag-and-probe trigger chains that use as ‘‘tag’’ a single light lepton trigger other than the tau. The combination of the two signal regions allow for the measurement of the tau trigger performance with sufficient statistics both at the low and high p_T regimes up to a tau p_T of about 200 GeV. The background events in these signal regions mostly come from events where a jet has been misidentified as a hadronic τ , and these are estimated using either data-driven methods or simulated events. In the $Z \rightarrow \tau\tau$ enriched region the background from misidentified jets come from QCD or W +jets processes, while in the $t\bar{t}$ enriched region the background comes from QCD or semi-leptonic $t\bar{t}$ decay events. Other minor background sources are due to diboson or $Z \rightarrow ee$ events where an electron is mis-reconstructed as a hadronic τ . Figure 39 shows the comparison between data and MC simulation of the tau p_T in the one-prong case for the $Z \rightarrow \tau\tau$ and $t\bar{t}$ selections obtained with the electron tag and tau probe trigger.




Figure 39. Distributions for data and MC simulation of the p_T of one-prong taus in (left) the $Z \rightarrow \tau\tau$ and (right) the $t\bar{t}$ signal regions (SRs). The data set was recorded with the $e - \tau$ tag-and-probe trigger. The bottom panels show the ratio between data and the signal-plus-background prediction. Only statistical uncertainties are shown.




Figure 40. Tau trigger efficiencies as a function of offline tau p_T for (left) the one-prong and (right) three-prong tau candidates. The legacy isolated L1 tau trigger with $E_T > 12$ GeV defined in ref. [2] is used as the seed. The measurement is performed using a combination of $Z \rightarrow \tau\tau$ and $t\bar{t}$ selections [2] in events recorded with the $e - \tau$ tag-and-probe trigger. The efficiency from the sum of the $Z \rightarrow \tau\tau$ and $t\bar{t}$ signal processes estimated with MC (data after subtracting the background) is shown with open (filled) markers. The bottom panels show the ratio of the efficiencies. Only statistical uncertainties are shown.

The tau trigger efficiency is calculated with respect to the reconstructed offline tau candidates and is separated between one-prong and three-prong tau cases. Figure 40 shows the resulting efficiency as a function of the offline tau p_T for one-prong and three-prong tau candidates using 2022 data recorded with a combined electron tag and tau probe trigger.

In the one-prong selection, the resulting efficiency of the tau trigger reaches a plateau slightly below 95% for tau candidates with $p_T > 40$ GeV, while it is about 90% for the three-prong case. The small differences between the data and MC signal prediction observed in figure 40 are used to derive scale factors to correct other MC simulation samples used in data analyses.

6.4 Jets

6.4.1 Jet trigger reconstruction

A detailed description of jet triggers in Run 2 can be found in ref. [2]. At the HLT, small- and large-radius jets are reconstructed using the anti- k_r algorithm [61] with a radius parameter of $R = 0.4$ and $R = 1.0$, respectively. While jet triggers in Run 2 employed only calorimeter topo-clusters as input, Run-3 triggers are able to rely on PFO thanks to the extended tracking capabilities offered by the HLT in Run 3.

Despite the large reductions in tracking time at the HLT, the CPU time usage of jet chains is still high due to full scan tracking, as discussed in section 5.1.4. For high- p_T jet chains such as single- and multi-jet chains, the large gap between L1 and HLT thresholds allows for an early reduction in event rate via the application of a calorimeter-only jet selection (*calorimeter preselection*) before any kind of tracking is executed. It is based on topo-cluster jets, described in section 5.2. Minimal HLT jet p_T requirements are adjusted to obtain the maximal CPU-time reduction which does not impact trigger efficiencies. The preselection stage is found to reduce the output rates of HLT jet trigger chains by less than 2%. Table 4 shows the impact of different preselections on the event rates of various jet triggers.

Table 4. Impact of calorimeter preselections on the event rates of jet triggers in which tracking reconstruction is performed. The percentage of events passing the preselection step is estimated with respect to the original L1 trigger rate from a reprocessing of Run-2 Enhanced-bias data [33]. The HLT selection column shows the lowest jet trigger threshold which this preselection is applied for.

L1 selection	Preselection	HLT selection	Events passing preselection
1 jet, $p_T > 100$ GeV	1 jet, $p_T > 225$ GeV	1 jet, $p_T > 275$ GeV	23%
3 jets, $p_T > 50$ GeV	4 jets, $p_T > 85$ GeV	4 jets, $p_T > 115$ GeV	15%
4 jets, $p_T > 15$ GeV	5 jets, $p_T > 50$ GeV	5 jets, $p_T > 70$ GeV	7%
4 jets, $p_T > 15$ GeV	6 jets, $p_T > 40$ GeV	6 jets, $p_T > 55$ GeV	6%

Even though this strategy works well for baseline jet triggers, it is not optimal for triggers whose HLT jet thresholds are very close to the L1 thresholds. One example of such triggers is the b -jet triggers, where a more complex preselection step detailed in sections 5.1.5 and 6.5.1 is included in the HLT reconstruction.

After the preselection step, fast full scan tracking is executed as detailed in section 5.1.4 and the resulting tracks are extrapolated to the calorimeter. The PFO formation starts by matching full scan tracks and topo-clusters taking into account the extrapolated track position and topo-cluster location. The topo-clusters which are not matched to any track are referred to as neutral PFOs. The matched topo-clusters have their energy removed according to the expected calorimeter energy deposited by the matched track. The tracks are considered as charged PFOs. Any topo-cluster surviving the energy subtraction procedure becomes a neutral PFO. Charged PFOs not matched to the primary vertex are discarded, which is the dominant means of pile-up suppression. Outside the geometrical

acceptance of the tracker, only the calorimeter information is available. Hence, in the forward region the topo-clusters alone are used as inputs to the PFlow jet reconstruction.

After reconstruction, small- R jets are calibrated through a procedure similar to that used offline [46]. A correction accounting for pile-up contamination is applied on an event-by-event basis to the jets. This is followed by another correction, compensating for the energy response of the calorimeter. Finally, a sequence of calibrations is applied to PFlow triggers to correct for the residual discrepancies between reconstructed and simulated jets and accounts, for example, for energy differences resulting from the different showering of quarks and gluons.

For Run 3, large- R jet triggers are also extended to use PFOs as inputs. In addition, the Constituent Subtraction [62] and SoftKiller [63] algorithms are applied to neutral PFOs to subtract energy contamination originating from pile-up interactions. After this step, Soft Drop grooming [64] with parameters $\beta = 1.0$ and $z_{\text{cut}} = 0.1$ is applied to mitigate the contamination from initial state radiation and underlying event. Finally, a calibration similar to small- R jets is used to correct for the energy response of the ATLAS calorimeter and adjust the reconstructed mass of the large- R jet.

6.4.2 Jet trigger menu

Jet triggers are used for a wide set of measurements, ranging from precision physics measurements to detector performance studies. The inclusive jet, di-jet and multi-jet production measurements rely on the events selected by small- R single- and multi-jet triggers. Events selected by these triggers are also employed for the calibrations of the jet energy scale and resolution, as well as for new physics searches such as supersymmetry. Another important class of jet triggers is represented by large- R triggers, selecting final states with boosted weak vector bosons (W/Z) or Higgs bosons. These triggers are generally employed by analyses searching for heavy resonances predicted, for example, by theories of extra-dimensions and the two-Higgs-doublet model.

The jet trigger chains are initiated by L1 algorithms selecting single jets, multi-jets, or, for very low- p_T HLT thresholds, a random trigger at L1. The primary, unrescaled jet chain applies a threshold of 420 GeV and has a rate of 42 Hz (at $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$). Small- R single jet chains with p_T thresholds below 420 GeV are prescaled to provide a complete jet p_T spectrum, with a constant rate per chain ranging from 1 to 4 Hz; large- R single jets with thresholds between 110 GeV and 460 GeV are also collected using prescaled triggers at a rate range between 2 and 3 Hz per chain. Trigger chains, which select events based on the scalar sum of the transverse momentum of all jets (H_T), contribute with a rate of 34 Hz. Topo-cluster-based jet triggers are also available in Run 3 as a backup for PFlow jet triggers.

6.4.3 Jet trigger performance

The efficiency of jet triggers is primarily a function of the p_T of the leading jet in the event, and is measured using the bootstrap method. In addition to the selection criteria specified in the following figures, the offline jets are required to satisfy a timing cut of less than 12.5 ns to minimise the contamination from out-of-time pile-up jets.

The efficiencies of L1 and HLT single jet triggers as a function of offline reconstructed jet p_T are shown in figures 41 and 42, respectively. The reference trigger selections reported on the legend have a looser L1 and HLT selection than the probed triggers to ensure an unbiased reference dataset. It must be noted that the jet-energy-scale correction is applied at the HLT but not at L1, meaning that the L1-jet thresholds are effectively 50–100% higher when considered at the jet-energy scale of




Figure 41. Efficiency of the legacy L1 single jet triggers in the (left) central and (right) forward regions. The efficiency is computed using the bootstrap method with respect to events taken by an independent trigger, shown in brackets, that is 100% efficient in the relevant region. Only statistical uncertainties are shown.




Figure 42. Efficiency of single jet trigger selection by the HLT in the (left) central and (right) forward regions. Central PFlow jet triggers exploit the particle flow reconstruction with calorimeter clusters and tracks, while forward PFlow jet triggers rely on the topo-cluster-based reconstruction to trigger events outside the ID acceptance ($|\eta| < 2.5$). The efficiency is computed using the bootstrap method with respect to events taken by an independent trigger, shown in brackets, that is 100% efficient in the relevant region. Only statistical uncertainties are shown.

the HLT. A steep rise in efficiency near the nominal thresholds reflects the compatibility of energy scales for jets reconstructed in the trigger and offline.

The efficiency of a multi-jet trigger requiring N jets depends mostly on the p_T of the N^{th} p_T -ranked jet. The relevant efficiency measurements are shown in figure 43 as a function of the offline reconstructed N^{th} jet p_T for the primary L1 and HLT multi-jet triggers. Lower efficiency at high p_T in the L1 four-jet trigger is due to the merging of the close-by jets at L1, which should be improved in the upgraded L1Calo system.

Efficiencies of the large- R jets as a function of the leading offline jet p_T are shown in figure 44 (top left), for both HLT and L1 trigger chains. Figures 44 (top right) and (bottom left) present the efficiencies as a function of the offline jet mass and p_T , respectively, for two main HLT chains with mass cuts. The L1 seed is shown not to have an impact on the efficiency of these triggers due to the large gap between L1 and HLT thresholds. The triggers shown are fully efficient at high jet p_T and mass values, and the steep rise in efficiency near the nominal thresholds reflects a good energy resolution.




Figure 43. Multi-jet trigger efficiency as a function of the offline reconstructed N^{th} jet p_T (left) at L1 and (right) in the HLT. The efficiency is computed using the bootstrap method with respect to events taken by an independent single jet trigger, shown in brackets, that is 100% efficient in the relevant region. Only statistical uncertainties are shown.




Figure 44. L1 and HLT efficiencies are shown for single large- R jet triggers as a function of (top left) the leading offline large- R jet p_T and (top right) highest mass offline large- R jet mass (m), (bottom left) large- R single jet triggers with a mass cut as a function of leading offline large- R jet p_T and (bottom right) H_T triggers as a function of offline H_T . The efficiency of triggers is computed using the bootstrap method with respect to events taken by an independent trigger, shown in brackets, that is 100% efficient in the relevant region. Only statistical uncertainties are shown.

The H_T trigger efficiency as a function of offline H_T is presented in figure 44 (bottom right) for both L1 and HLT triggers, which also rises sharply. The observed data/MC differences will be corrected later-on in physics analyses by dedicated scale factors applied to MC simulation samples.

6.4.4 Trigger level analysis with jets

HLT topo-cluster-based and PFlow jets with transverse momentum larger than 20 GeV are saved to the TLA stream described in section 4. To remain within the CPU constraints of the HLT farm, full scan tracking and PFO reconstruction are executed only if the event passes the calorimeter preselection step.

In 2022, jet chains populating the TLA stream were seeded by two L1Calo legacy triggers, L1 jet $p_T > 100$ GeV (which reaches 50% efficiency at around 160 GeV) and $H_T > 190$ GeV, with a total rate of approximately 4 kHz. This allows for the recording of orders of magnitude more events containing lower- p_T jets with respect to standard ATLAS jet triggers, extending the sensitivity of hadronic searches to resonance masses as low as 400 GeV [36].

As shown in figure 45, the response of trigger jets with the default calibration is already within 2% of offline jets across the momentum range of interest to TLA studies. For Run 3, a custom residual energy scale calibration is to be derived and applied at the analysis level in order to further improve the response and resolution of TLA jets. In order to derive such a calibration, an extensive set of jet variables, beyond the jet four-momentum, is saved in the TLA stream. This includes the energy fraction deposited in the EM and hadronic endcap calorimeters, the momentum density of soft radiation in the event, the number of primary vertices reconstructed at the HLT, the jet active area [46], and the number of constituents encompassing 90% of the jet transverse energy.

6.4.5 Jet triggers for heavy-ion collisions

Jets are direct probes of the quark-gluon plasma produced in HI collisions, and studying how they are modified in such collisions compared to pp collisions provides insight into their interactions with this QCD medium. The main challenge for triggering on jets in HI collisions is the presence of a large amount of UE. The HLT jet algorithm uses the HI UE subtraction procedure described in section 4.2 and selects events containing jets with transverse energies exceeding a threshold ranging from 60 to 85 GeV. Jets are reconstructed at the HLT across the entire calorimeter using the anti- k_t algorithm with distance parameter $R = 0.4$, from projective towers of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ formed from the summation of calorimeter cell energies with UE subtraction applied.

6.5 Jets containing b -hadrons

6.5.1 b -jet trigger reconstruction and selection

Triggers based on jets containing b -hadrons, so-called b -jet triggers, are designed to allow for highly efficient recording of fully hadronic events with predominantly heavy flavour content jets. The detailed description of the b -jet triggers in Run 2 can be found in ref. [65], the changes for Run 3 are discussed here. The basic inputs to b -tagging are reconstructed jets, reconstructed tracks and the position of the primary vertex. The jet reconstruction is described in section 6.4. The primary vertex finding and b -jet track reconstruction are discussed in sections 5.1.4 and 5.1.5, respectively.

The main b -jet identification steps are as follows:

- A fast b -tagging algorithm is run on super-RoIs and uses the tracks inside them to provide the first b -jet preselection. To reduce the rate of the full scan tracking for b -jet triggers, the b -jet




Figure 45. The response curve for HLT-level PFlow jets in the momentum range of interest for Run-3 TLA. The online-to-offline response is evaluated at the final energy scale in the calibration sequence of the jets. HLT-level jets are calibrated with the default sequence used at the HLT for the 2022 data taking. Offline jets are corrected with the calibration available during 2022 data taking. Events are required to pass the single jet $p_T > 15$ GeV (circle markers), $p_T > 25$ GeV (triangle markers) or $p_T > 35$ GeV (square markers) random-seeded triggers depending on the p_T of the leading reference jet. For each event, only HLT jets matched ($\Delta R < 0.3$) to the leading and sub-leading offline jets reconstructed within $|\eta| < 2.4$ are considered. The overlaid heat map refers to the number of events in the analyzed sample for each configuration of the reference jet momentum and response.

preselection stage is run with lower track p_T thresholds and with a super-RoI composed of wider regions around jets than were used for the vertex tracking in Run 2 [41, 65]. The details of the fast b -tagging algorithm used at this stage are provided in ref. [53].

- The final b -tagging algorithm, detailed below, uses precision tracking and primary vertex information.

The b -jet identification relies on the properties of b -hadrons: long lifetime (about 1.5 ps), hard fragmentation, a relatively large mass of about 5 GeV, and a displaced (secondary) vertex (SV) formed a few millimeters away from the primary vertex. Tracks associated with b -jets are characterised by larger transverse (d_0) and longitudinal (z_0) impact parameters. In addition, b -hadrons decay semileptonically, either promptly, or via a subsequent c -hadron decay, to electrons or muons. The branching ratio of these semileptonic decays is about 20% each and results in the presence of a low- p_T lepton close to the b -jet.

The flavour tagging identification is done in two steps [48]. Low-level algorithms reconstruct characteristic features of the heavy flavour jets based on track properties, such as impact parameters, or combine those tracks to reconstruct the SV. The outputs of these low-level algorithms are then combined into high-level algorithms, usually using some multivariate technique.

Low-level taggers used at the HLT are the IP2D algorithm, which utilises the signed transverse impact parameter significance of tracks, and the IP3D algorithm, which uses, in addition, the longitudinal impact parameter significance [48].

The secondary vertex algorithm, SV1, uses tracks associated with jets (after rejecting those compatible with K_s or Λ , photon conversions or interactions with detector material) to assign decay products from b - or c -hadrons to a single common SV. Several discriminating variables associated with the SV are then used as inputs to the high-level tagger. Finally, the JetFitter algorithm exploits the topology of weak b - and subsequent c -hadron decays inside a jet to reconstruct the full decay chain, recreating the approximate b -hadron path from the primary vertex via bottom and charm vertices.

Another low-level algorithm is the Deep Impact Parameter Sets (DIPS) [66], based on the Deep Sets architecture. It uses impact parameter information, accounting for correlations between the track features, among other variables. DIPS considers tracks in the jet as an unordered and variable-sized set, which is physically better motivated than the algorithm based on recurrent neural network [65] used in Run 2, given that the b -hadron decay products do not exhibit any intrinsic sequential ordering. The performance of the DIPS tagger is investigated using $t\bar{t}$ MC events and 2022 data collected with a calibration trigger with one electron, one muon and two jets. The final DIPS discriminant distribution is shown in figure 46 (left). In the case of SM $HH \rightarrow b\bar{b}b\bar{b}$, a key signature relying on b -jet triggers, the b -jet preselection step lowers the input rate to the remaining HLT by a factor of five at the cost of reducing the overall signal efficiency by roughly 2%. More details about DIPS, its training and usage in b -jet triggers are given in ref. [53].




Figure 46. (Left) Distribution of the b -jet discriminant score for the DIPS algorithm shown in data from selected runs and $t\bar{t}$ MC events collected with a calibration trigger with one electron, one muon and two jets. Only statistical uncertainties are shown. (Right) Expected trigger rates as a function of b -tagging efficiency using the DL1d and GN1 algorithms, while requiring at least four HLT PFlow jets, three of which are required to be above the b -tagging threshold. Rates are estimated with Run-3 Enhanced-Bias data, and the b -jet efficiencies are estimated for $t\bar{t}$ samples using PFlow jet. The relative errors on trigger rates are between 8 (7)% and 18 (20)% for DL1d (GN1).

The newly developed algorithm based on deep feed-forward Neural Networks, the so-called DL1 series [48], replaced boosted-decision-tree based taggers utilised in Run 2 [65]. The particular instance of the algorithm used in 2022 is called DL1d and takes as inputs the kinematic variables (p_T and η) of the jet as well as final discriminants from lower-level taggers (IPxD, SV1, JetFitter and DIPS).

From 2023 the b -jet trigger relies on a novel algorithm GN1, which is based on Graph Neural Networks (GNNs) [67]. Unlike the DIPS and DL1d, the GN1 utilises a single neural network taking

the tracks and jet information directly as inputs and is thus independent of other flavour tagging algorithms. The GN1 combines a GNN with the two auxiliary training objectives: the grouping of tracks originating from a common vertex and the prediction of the underlying physics process from which each track originated. This approach leads to a better understanding of the jet’s internal structure and, thus, a better algorithm performance.

The expected trigger rate as a function of the b -tagging efficiency using the DL1d and GN1 algorithms is shown in figure 46 (right). The expected light jet rejection as a function of the b -tagging efficiency for various b -taggers as well as their operating points are shown in figure 47. The GN1 tagger performance exceeds that of DL1d, which was the main tagger in 2022.




Figure 47. The light jet rejection as a function of the b -tagging efficiency (ROC curve) for the DIPS algorithm used on topo-cluster jets for preselection stage (dash-dotted line), the DIPS algorithm used on PFlow jets for TLA triggers (dashed line), the DL1d algorithm used as the final tagger in 2022 (solid line) and the GN1 algorithm used as the final tagger from 2023 (dotted line) obtained on a $t\bar{t}$ MC sample. Statistical uncertainties for each ROC curve are represented with shaded regions around the curves. The vertical dashed lines represent the operating points for b -tagging used at the HLT. The bottom panel displays the ratio of all the ROC curves with respect to the DL1d performance.

6.5.2 b -jet trigger menu

b -jet triggers are crucial for several precision measurements and searches for new particles such as resonant and non-resonant $HH \rightarrow b\bar{b}b\bar{b}$ [68, 69], VBF $H \rightarrow b\bar{b}$ [70, 71], $t\bar{t}H \rightarrow t\bar{t}b\bar{b}$ with both top quarks decaying hadronically [72], $b\phi \rightarrow bb\bar{b}$ [73], third generation squarks [74], low mass di- b -jet resonances [75], etc.

The b -jet trigger menu consists of the following physics triggers: single- b -jet triggers, multi- b -jet triggers and specialised b -jet triggers designed for specific physics processes which often occur in combination with other signatures. Single- b -jet triggers and multi- b -jet triggers are seeded from dedicated single-jet or multi-jet L1 objects. At the HLT various combinations of the p_T thresholds and b -tagging operating points are employed. In multi- b -jet triggers, between one and four jets are tagged, while other jets are not tagged. An asymmetric chain requiring four jets with p_T greater

than 80 GeV, 55 GeV, 28 GeV, and 20 GeV of which two are b -tagged (with efficiency of 77%) has a rate around 150 Hz, the lowest unprescaled single b -jet trigger ($p_T > 225$ GeV) around 50 Hz and one jet of $p_T > 150$ GeV plus two b -jets with $p_T > 55$ GeV and b -tagging efficiency of 70% around 10 Hz at luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

Flavour-tagging calibration triggers [65] include a single lepton (electron or muon) with one or two jets with $p_T > 20$ GeV, electron-plus-muon with one or two jets with $p_T > 20$ GeV, and muon-jet matched triggers, that are also used for physics studies in HI collisions. These latter triggers were seeded from either a single-muon or a muon-plus-jet trigger item at L1 in 2022. New chains seeded from L1 objects which use topological information, such as ΔR , are planned to be introduced later in Run 3. At the HLT, a muon is required to be matched to a jet by fulfilling a requirement of $\Delta R(\mu, \text{jet}) < 0.5$, with an additional requirement for pp collisions $\Delta z(\mu, \text{jet}) < 2$ mm, where the z -position of the jet is taken to be the primary vertex z -position.

6.5.3 Trigger level analysis with b -jets

The Run-3 TLA HLT infrastructure allows for the full outcome of flavour tagging algorithms to be written out in events triggered by TLA chains targeting b -jet signatures. In every TLA event, each HLT-level jet (for which flavour tagging is performed) is linked to the corresponding b -tagging object in which the values of low and high-level taggers are stored. Furthermore, the value of the DIPS tagger is added to every PFlow jet saved to the TLA stream, regardless of whether the triggering chain is configured to perform dedicated flavour tagging or not. This setup allows for the recording of orders of magnitude more events containing lower- p_T b -jets with respect to standard b -jet triggers, extending the sensitivity of low mass di- b -jet resonance searches to resonance masses down to 100 GeV.

6.6 Missing transverse momentum

6.6.1 E_T^{miss} trigger reconstruction

The missing transverse momentum (E_T^{miss}) trigger is used to identify events with particles that do not interact with the ID or calorimeters because of the absence of strong or electromagnetic interactions and that have lifetimes large enough to leave the detector without decaying into detectable particles. Examples of such processes include SM Z boson decays to neutrinos [76], as well as searches for beyond SM decays involving dark matter [77] or supersymmetric [78] particles.

The E_T^{miss} can be computed from the magnitude of the vector sum of constituents, c :

$$E_T^{\text{miss}} = \left| \sum_c \vec{p}_{T,c} \right|.$$

In addition to the E_T^{miss} , the scalar sum of the constituents is also computed. The main variations in how E_T^{miss} is computed in the trigger involve details of which constituents to utilise in the sum in the equation above.

The following HLT E_T^{miss} algorithms are available in Run 3:

- **cell**: the E_T^{miss} is formed from a sum over all calorimeter cells passing the selection $E_i > -5\sigma_i$, $|E_i| > 2\sigma_i$, where σ_i is the estimated noise in that cell. This is to protect against spurious large negative energy signals, discussed in section 5.2. The cell energy magnitude must be greater than $2\sigma_i$ to reduce the effect of noise from electronics and pile-up.

- `tcpufit`: the E_T^{miss} is formed using topo-clusters as inputs, which are combined in η - ϕ patches with dimensions of approximately 0.7×0.7 (roughly the size of an $R = 0.5$ jet). The energy contribution from pile-up to these patches is estimated by a fit over them, requiring that pile-up events have no true E_T^{miss} and are approximately evenly distributed over the calorimeter [79]. The estimated pile-up contribution to each patch is subtracted, and the remaining patch transverse momenta are summed to obtain the E_T^{miss} . The full algorithm is described in appendix A of ref. [79].
- `trkmht`: the E_T^{miss} is formed over all jets passing a Jet Vertex Tagger (JVT) selection [80] where applicable. The vector sum of the p_T of tracks from the primary vertex that are not associated with a passing jet defines a track soft term which is added into the E_T^{miss} calculation.
- `pfopufit`: the `pfopufit` algorithm uses the same techniques as `tcpufit`, but the input topo-clusters are modified to use PFO tracks [50]. As well as the improved momentum resolution of the PFOs, the vertex information provided by the charged PFOs is used to improve the categorisation of deposits into hard-scatter and pile-up.
- `mhtpufit`: the E_T^{miss} is formed from a sum over jets passing a JVT selection [80]. A similar technique to `tcpufit` is used to correct these jets for the impact of pile-up. Two variants are used: `mhtpufit_pf` uses PFOs and jets formed from them to estimate the pile-up contributions, `mhtpufit_em` uses jets formed from EM scale topo-clusters and the hadronic scale topo-clusters to estimate the pile-up contributions.
- `pfsum`: the E_T^{miss} is formed from a sum over the PFOs in the event. Two variants exist: in `pfsum_vssk` the PFOs have their energies reduced according to their Voronoi areas [81], whereas in `pfsum_cssk` the constituent subtraction method [62] is used instead. In both cases, the SoftKiller algorithm [63] is used to remove PFOs from low-energy areas of the calorimeter.

In 2022, the primary L1 item for E_T^{miss} was the L1Calo legacy trigger with $E_T^{\text{miss}} > 50$ GeV [21, 79]. Before any kind of tracking is executed, an early reduction of rate at HLT is achieved by a calorimeter preselection requirement that the cell-based E_T^{miss} is greater than 65 GeV. The choice of default Run-3 algorithm is made based on the performance in terms of background rejection versus signal efficiency and is discussed in the following paragraphs.

6.6.2 E_T^{miss} trigger performance and menu

Since muons are treated as invisible by the E_T^{miss} trigger algorithms described above, events with a boosted $Z \rightarrow \mu\mu$ can be used for the E_T^{miss} performance studies with $p_T(\mu\mu)$ serving as a proxy for E_T^{miss} . A background acceptance vs. trigger efficiency curve is shown in figure 48 (left) comparing the performance of a selection of E_T^{miss} algorithms considered for Run 3. The efficiencies are calculated as the fraction of events passing a given E_T^{miss} requirement for data events with an actual number of interactions per bunch crossing of at least 38, and a $Z \rightarrow \mu\mu$ event selection passing a single muon trigger, with $76 < M(\mu\mu) < 106$ GeV and $p_T(\mu\mu) > 175$ GeV. The background rate is obtained from an offline trigger reprocessing of the 2022 data collected with zero-bias triggers described in section 7.2. The new Run-3 default E_T^{miss} trigger is based on `pfopufit` algorithm, which shows the best ability to retain signal, whilst rejecting background, of all the algorithms listed in section 6.6.1.




Figure 48. (Left) Background acceptance vs. trigger efficiency curves for a selection of Run-3 E_T^{miss} trigger algorithms. (Right) Trigger efficiency as a function of $p_T(\mu\mu)$ for the primary L1 E_T^{miss} trigger and the full trigger chain, with specific thresholds listed. The trigger rates for the HLT chains shown are approximately the same rate (to within 10%). The data are presented in comparison to 2018 efficiency. Only statistical uncertainties are shown.

The Run-2 default `tcpufit` trigger is maintained as backup and for analyses where the primary vertex may not match the online primary vertex. These triggers had unrescaled HLT thresholds of 90 GeV and 115 GeV in 2022 and rates of about 70 Hz and 30 Hz at luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, respectively. Requiring the presence of additional objects in the event, allows for decreased thresholds for E_T^{miss} triggers, as discussed in section 6.9.

A signal efficiency curve is shown in figure 48 (right) comparing the performance of the primary L1 E_T^{miss} trigger alone to the default HLT E_T^{miss} trigger chains used in Run 3 and Run 2. The efficiencies are calculated as the fraction of events passing a given $p_T(\mu\mu)$ requirement for data events with a $Z \rightarrow \mu\mu$ event selection passing a single muon trigger, with $76 < M(\mu\mu) < 106 \text{ GeV}$. The 2018 L1 efficiency matches that of 2022 within a few percent, so it is not shown. The improved efficiency seen for the Run-3 `pfopufit`-based chain compared to the Run-2 chain reflects the `pfopufit` algorithm’s enhanced momentum resolution and hard-scatter/pile-up categorisation.

6.7 B-physics and light states

The trigger selection of events for B -physics analyses is primarily based on the identification of b -hadrons through decays with a muon pair in the final state. However, one or more muons or electrons could also be present in specific selections. Examples are decays with charmonium, $B \rightarrow (J/\psi \text{ or } \psi')X \rightarrow \mu\mu X$, rare decays $B_{(s)}^0 \rightarrow \mu\mu$, and $B \rightarrow \mu\mu X$ decays. Decays of prompt charmonium and bottomonium are also identified through their di-muon decays, and are therefore similar to B -meson decays, apart from the lack of measurable displacement from the pp interaction point. As the BLS topologies are significantly different from the majority of the ATLAS physics triggers, they are recorded to a separate BLS stream, as discussed in section 4.1. In 2022 this stream collected data with an average rate of approximately 240 Hz.

As B -mesons are light, the resulting muon momenta are rather soft. To control the input rate to the HLT, most of the B -physics triggers require two muons at L1. Their rate is substantially reduced compared to single-muon L1 triggers. Depending on the mass of the resonance this can result in the muons being within a single RoI or in separate RoIs. At the HLT, muons are reconstructed using the same algorithms as described in section 5.3 with the additional requirement that the combined muons have opposite charges and form a good vertex (where a fit is performed using the ID track parameters) within a certain invariant mass window. For example, the mass range of the $bJ\psi\mu\mu$ is between 2.5 and 4.3 GeV while $b\Upsilon\mu\mu$ is between 8 and 12 GeV. Examples of these mass spectra can be seen in figure 49.




Figure 49. Invariant mass distribution of offline-selected di-muon candidates passing the lowest thresholds of di-muon B -physics triggers. Triggers targeting different invariant mass ranges are illustrated with different colours, and the differing thresholds are shown with different shadings. No accounting for overlaps between triggers is made, and the distributions are shown overlaid, and not stacked.

Di-muon trigger rate restrictions at L1 define the lowest muon transverse momentum thresholds for primary B -physics triggers. HLT triggers based on two L1 muons passing a 3 GeV p_T threshold need to be prescaled for most data-taking luminosities. The prescales are adjusted to maximise the recorded number of low- p_T di-muon events while remaining within operational constraints of the ATLAS TDAQ system as discussed in section 4.1. Higher threshold HLT triggers seeded from two muons passing 5 GeV or 8 GeV thresholds at L1 are unprescaled when the L1 and HLT bandwidths allow it.

Additional primary and supporting triggers are also implemented. Triggers requiring three muons at L1 help to maintain the lowest muon p_T thresholds for certain event signatures with a likely presence of a third muon. For semileptonic decays, such as $B^0 \rightarrow \mu\mu K^{*0} (\rightarrow K^+\pi^-)$, searches for additional ID tracks and a combined vertex fit are performed, assuming a few exclusive decay hypotheses. For the start of Run 3 a so-called **bBmuX** selection is implemented to perform the partial reconstruction of the B -hadron decay final state $D^{*+}\mu^-X$ with the D^{*+} reconstructed through the cascade hadronic decay chain $D^{*+} \rightarrow \pi^+D^0 (\rightarrow K^-\pi^+)$.

To evaluate the B -meson selection at the HLT, supporting triggers without vertex or charge requirements are used. The mass spectra of various **bBmuX** trigger decisions are plotted and fitted using events from the **BphysDelayed** stream. The consistency of the mass peaks in figure 50 with PDG averages and previous measurements [82] demonstrate the validity of the HLT reconstruction and selection.




Figure 50. To validate the selection of the bBmumux the mass spectra of the selected (left) $B_d \rightarrow J/\psi K^*$ and (right) $B_s \rightarrow J/\psi \phi$ events from the BLS Delayed stream are plotted along with the spectra for a variety of bJpsimumu and bBmumux trigger decisions. Only statistical uncertainties are shown.

6.8 Minimum-bias and forward signatures

To study diffraction, soft-QCD and similar topics, highly efficient triggers are necessary to select pp inelastic, diffractive and scattering interactions with the least possible trigger bias. These data also serve as references for heavy-ion measurements. The minimum-bias triggers discussed below are enabled only during dedicated low- μ runs.

6.8.1 Minimum-bias triggers

Depending on the number of collisions per bunch crossing, μ , the minimum-bias triggers play various roles. At very low values of μ ($\mu \ll 1$) the triggers require one or two signals on any side of the Minimum Bias Trigger Scintillator (MBTS) detector to select actual collision events and to ensure that the data sample is not dominated by empty events with no pp interactions. The MBTS detector was replaced for Run 3, as discussed in ref. [3]. It is a two-armed large-area plastic scintillator with a very high light yield covering $2.0 < |\eta| < 4.0$. On each side, the scintillator is divided into two concentric rings, each consisting of 8 octants covering the full azimuth. Altogether, the MBTS deliver 32 signals to the CTP, which can be used independently or combined. Further event selection is possible at the HLT by requiring matching timing between MBTS signals or the reconstruction of tracks in the ID. As a side benefit, reconstructing the MBTS signals at the HLT allows for the monitoring of the timing and energies of individual channels.

At μ values close to unity the need for MBTS as a source for the hardware trigger is less relevant and the minimum-bias sample can be collected by a random trigger at L1 followed by the requirement of a track reconstructed at the HLT.

At $\mu \sim 1-3$, every bunch crossing contains an inelastic collision and the L1 random trigger is sufficient for minimum-bias triggering. In these conditions, the focus shifts to collecting events with high multiplicity tracks (HMT) or events with a high momentum track for analyses like Bose-Einstein correlations or azimuthal correlations similar to what is done in HI physics [83]. The possibility of using a trigger with the total (transverse) energy deposit in the calorimeter is also planned.

At the start of 2022 data taking no track selection was applied at the HLT. This allows for the comparison in performance of the minimum-bias tracking selection in offline versus the trigger. In

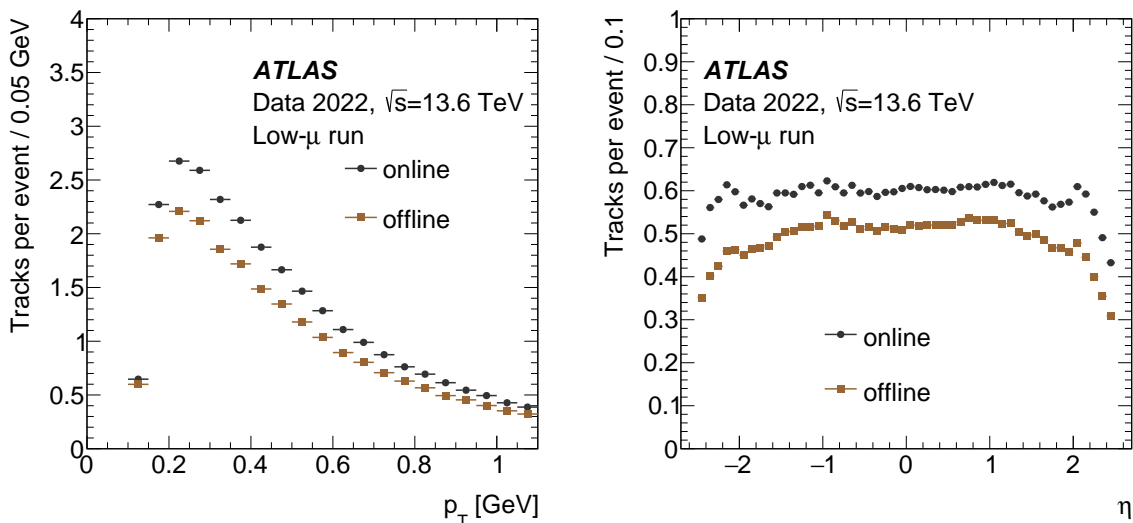


Figure 51. Comparison of (left) p_T spectra and (right) η of tracks reconstructed offline and passing minimum-bias quality selection and tracks reconstructed by the HLT online during low pile-up run in 2022.

Run 3 the track reconstruction for the minimum-bias trigger is based on offline algorithms. The minimum-bias online tracks are required to satisfy the minimum-bias track selection: $p_T > 0.1$ GeV and $|\eta| < 2.5$. In addition to the online selection above, the offline tracks are required to satisfy the standard minimum-bias analysis selection [84]: $|d_0| < 1$ mm, $|z_0 \sin \theta| < 1.5$ mm, a measurement in the IBL if it is expected, at least one pixel measurement and the number of measurements in the SCT above 2, 4, and 6 for tracks of p_T above 0.1, 0.3, and 0.4 GeV, respectively. Furthermore, the online and offline conditions are not exactly identical, e.g. the map of dead pixel and SCT modules is known only after data have been taken. A comparison of transverse momentum, p_T , and η for online tracks and offline tracks undergoing corresponding minimum-bias selections is shown in figure 51. Despite the looser online track selection, the purity of minimum-bias triggers is quite high.

Triggers with track reconstruction include a preselection step based on the count of pixel clusters and SCT space-points (coinciding pairs of hits on both sides of an SCT module). In particular, the single-track trigger preselection requires two-pixel clusters and three SCT space points. This preselection step reduces the input rate prior to execution of the tracking algorithms without any efficiency loss. Random-seeded triggers with only the preselection step and no further HLT requirements are used to collect unbiased samples for performance studies. Figure 52 shows the efficiency of HLT tracking selection and the L1 MBTS trigger as a function of the number of tracks reconstructed offline passing minimum-bias selection, as obtained with a data set collected with the random-seeded trigger at L1. The HLT tracking selection is nearly fully efficient if at least two tracks in the event are present (a required minimum multiplicity for the $p_T > 0.1$ GeV working point). The inefficiency of the MBTS triggers is attributed to disabled modules (one on each side) and only a partial geometric overlap between the MBTS and the inner tracking volume. For track p_T above 0.5 GeV, the track trigger is again fully efficient and the MBTS trigger performance mostly recovered. This is because the presence of at least one track of this momentum is usually correlated with additional activity resulting in signals in a few MBTS counters and thus corresponding triggers.

For HMT triggers the requirement on pixel clusters is removed and only the SCT space-points are counted and subject to a threshold requirement, optimised to have full efficiency at a given number of




Figure 52. (Left) Efficiency of triggering for events with at least two tracks of $p_T > 0.1$ GeV and passing minimum-bias quality track criteria (see text) by single track trigger and MBTS triggers requiring single (MBTS 1) or two hits (MBTS 2) in the scintillator. (Right) Efficiency of triggering for events with at least one track of $p_T > 0.5$ GeV passing minimum-bias track quality criteria by the same triggers as shown on the left panel. The efficiency is measured as a function of the number of offline tracks, N_{trk} , with respect to the fully efficient trigger requiring only two clusters in the pixels detector and three space-points in SCT seeded from random bunch crossing at L1. Efficiencies of the MBTS single and two hit triggers are measured with respect to the single track trigger. Only statistical uncertainties are shown. These triggers were active during low pile-up runs in 2022.

tracks. For HMT triggers at moderate values of μ , an additional pile-up mitigation strategy is required. This is because the counting of all tracks in the event instead of only tracks belonging to the highest multiplicity vertex impacts the purity of this type of trigger. In addition, performing tracking in the full volume of the ID becomes more time-consuming with rising μ , and so any preselection is beneficial. An algorithm to approximate the vertex position along z and the count of tracks originating from it was developed and optimised for expected data-taking conditions. It uses the triplets of clusters from the Pixel detector that are used to construct a linear extrapolation to the luminous region. Extrapolated positions are histogrammed along z with the binning optimised so that the coincidental combinations form a negligible background and real vertices form a well-pronounced peak. The threshold cut is applied on the count in the peak in the step preceding the tracking step. Because of the availability of the vertex z position, the tracks in HMT triggers with pile-up suppression are counted only if they are within 10 mm along z from the approximate vertex position. The impact of the additional pile-up suppression procedure on the efficiency of HMT triggers is shown in figure 53: here again the offline tracks are required to pass a minimum-bias selection and the triggers are seeded from the random L1 trigger. At moderate μ values, these chains are planned to be seeded by a trigger that sums up all energy in the calorimeters, which was not yet commissioned at the time of the 2022 low- μ runs.

6.8.2 Forward triggers

In addition to the minimum-bias triggers above, a diverse set of triggers aims to record events with elastic, diffractive or central-exclusive interactions. Their usage depends on the data-taking conditions




Figure 53. Efficiencies of high multiplicity triggers targeting selection of events with 40, 60, and 80 tracks as a function of number of tracks passing minimum-bias selection quality, N_{trk} . For each requirement the performance of the pile-up suppression (PU SUPP) variant of the trigger, designed to be robust in higher values of pile-up, is shown. Only statistical uncertainties are shown. These triggers were active during low pile-up runs in 2022.

such as pile-up or LHC beam optics [85]. The latter is usually defined by the value of the betatron function at the collision point, β^* [85]. During the high- β^* runs the focus is on elastic scattering events triggered solely by the ALFA detector, although there are also some triggers combining ALFA detector information with that of other ATLAS subdetectors which target the soft diffractive events. A detailed description of ALFA triggers can be found in ref. [12]. During low- β^* (‘standard’) runs triggers are based on signals from various stations of the AFP detectors with or without coincidence with standard signatures from the central detector. A detailed description of AFP can be found in refs. [3, 11]. The AFP can deliver trigger signals from two of its detector systems: the Silicon Trackers (SiT) and the Time of Flight (ToF) detectors. SiT trigger signals are expected to be more efficient (98–99% per station) than the signals from the ToF system (about 80% per side), but are also known to suffer from a 400 ns dead time after each hit [86], causing the efficiency to significantly decrease for the later bunches in a bunch train when the pile-up exceeds one. As the SiT trigger dead time depends on the beam intensity and train structure, the following logic is used:

- During high- μ runs, when the presence of protons in AFP is expected to be in every second bunch crossing,⁹ the trigger items are based on ToF,
- during low- μ runs, when the probability of observing a proton in consecutive bunch crossings is small, the trigger items are based on SiT.

During high- μ data taking, the physics programme using proton tagging is focused on the measurements of exclusive and two-photon exchange processes, typically in conjunction with high- p_T

⁹In the first approximation this depends on pile-up. E.g. for $\mu = 50$ and a probability of registering pile-up proton originating from a single pp interaction of 2%, the chance of observing a proton in AFP is $1 - (1 - 0.02)^{50} \approx 64\%$.

object(s) produced in the ATLAS detector to keep the rate low. An example of such a process is exclusive jet production when un-prescaled jets with p_T of about 150 GeV are required. In order to trigger such events, the presence of a jet with a minimum threshold of 50 GeV and a proton(s) in AFP (ToF trigger) is required at L1. At the HLT, matching is required between the kinematics of the centrally reconstructed di-jet system and the scattered protons as reconstructed using the SiT. Depending on the settings of the algorithm selection criteria, the efficiency is expected to be 60–85%. In addition, the match between the di-jet vertex and the vertex z location reconstructed using the ToF data can be applied. Such trigger chains may obtain a rate reduction of a factor of 100 when compared to a nominal jet trigger with a similar p_T threshold.

The composition of the trigger menu for low- μ runs depends on the exact data-taking conditions. The number of colliding bunches, pile-up and beam optics [85] play a key role in the expected event rate. Studies of soft diffractive processes are usually realised using the least biased triggers. For AFP, this translates into the requirement of a proton to be detected in both stations on the same side of ATLAS (single tagged events) or all four AFP stations (both ATLAS sides; double-tagged events). AFP triggers can be combined with other ATLAS trigger objects. For example, for diffractive charm meson production, in addition to the presence of protons in AFP, a track with a certain transverse momentum is additionally required.

6.9 Unconventional tracking signatures

The search for LLPs is an important part of the Run-3 physics programme, as they appear in many motivated scenarios of phenomena beyond the SM. LLPs which decay within the ID volume or pass entirely through it result in a variety of unconventional tracking signatures in the detector. Several new triggers presented below were developed to target such unique signatures: long-lived charged particles, displaced jets, jets with displaced tracks and displaced leptons. They make use of both standard tracking and LRT, described in section 5.1. The use of tracking for these signatures leads to lower background rates, which allows for lower particle momentum or E_T^{miss} requirements in the trigger, resulting in large gains in sensitivity for Run-3 searches. Rates for the triggers discussed below are given in table 3.

6.9.1 Long-lived charged particles that partially or fully traverse the inner detector

Three new triggers are developed targeting long-lived charged particles that partially or fully traverse the ID. They make use of standard prompt full scan tracking executed after a calorimeter E_T^{miss} or jet preselection is applied to reduce the full scan tracking rate, as described in section 5.1.4. In Run 2, these searches generally relied on the E_T^{miss} trigger [87–89] to select events, which resulted in a low acceptance for models not producing large E_T^{miss} .

Isolated high p_T track. Heavy charged particles (such as charginos, sleptons, and R -hadrons) with sufficiently long lifetimes can leave an isolated, high- p_T track in the ID. A new isolated high- p_T track trigger is introduced for Run 3 to increase sensitivity to LLP signatures with low E_T^{miss} in the event. It requires a `tcput` E_T^{miss} threshold of 80 GeV, 30 GeV below the lowest unpre-scaled E_T^{miss} trigger, to reduce the amount of CPU-time spent on full scan tracking. The trigger uses tracks from the fast tracking step to select events with at least one isolated track with $p_T > 120$ GeV that also passes additional track quality requirements. For the track to be isolated, the scalar sum of the track p_T within $\Delta R < 0.3$ must be less than 10 GeV. Figure 54 shows the expected performance of the isolated




Figure 54. Isolated high- p_T track trigger expected efficiency using Run-3 MC simulation of SUSY di-stau events with a mass of 600 GeV and a lifetime of 10 ns vs. (left) the cell and (right) the tcpufit algorithm E_T^{miss} . The efficiency of the new trigger (circles) is compared with that of a logical “or” of the lowest unprescaled E_T^{miss} triggers based on tcpufit and pfopufit algorithms described in section 6.6 (squares). Both of these are overlaid on the signal sample E_T^{miss} distribution of fraction of events per bin (gray histogram), and include the L1 trigger efficiency. Only statistical uncertainties are shown.

high- p_T track trigger with respect to the HLT E_T^{miss} calculated with cell and tcpufit algorithms, compared to the “or” of the lowest unprescaled E_T^{miss} triggers based on tcpufit and pfopufit algorithms. All of these algorithms are described in section 6.6. SUSY di-stau LLP MC simulation for a stau mass of 600 GeV and a lifetime of 10 ns is used to evaluate the efficiency. The new trigger increases the acceptance to signal-like events at lower values of E_T^{miss} .

Large dE/dx triggers. New heavy, charged particles, mentioned above, may also leave large energy deposits in the ID silicon layers compared to what is expected from a minimally ionizing particle. The measurement of these large ionisation energy losses per unit pathlength, dE/dx , in the Pixel detector is a handle to identify tracks as signal candidates. A new trigger targeting long-lived, heavy, charged particles [88] uses the dE/dx measurement capabilities of the ID for trigger decisions. The same E_T^{miss} triggers, as used for the isolated track trigger above, are required before running the full scan tracking. The dE/dx trigger selects events with at least one track with $p_T > 50$ GeV, an average $dE/dx > 1.7$ MeV/cm, at least two hits with $dE/dx > 1.7$ MeV/cm, track $|d_0| < 2.5$ mm, and $|\eta| < 2.5$. Figure 55 shows the online dE/dx distribution of tracks with $p_T > 10$ GeV in data collected during 2022 with E_T^{miss} , single-jet, and multi-jet triggers where full scan tracking was run. It also shows the correlation of the online dE/dx measurement to the offline measurement.

Disappearing track triggers. Charged particles with slightly shorter lifetimes than the signatures above can decay part way through the ID and leave a short tracklet of a few hits. These tracklets are referred to as disappearing tracks when the charged particle decays into invisible and low- p_T particles that are not reconstructed. As in the previous two cases, the disappearing track trigger makes use of full scan tracking executed after the same E_T^{miss} trigger requirements, which are lower than the lowest unprescaled E_T^{miss} trigger. The fast tracking algorithm is modified to save tracklets with four hits in the inner layers of the ID that fail to become tracks. In order to reduce the large background to this signature, a Boosted Decision Tree (BDT) based on the track parameters, quality of fit, and number of hits in the Pixel and SCT detectors is used to separate signal-like tracklets from

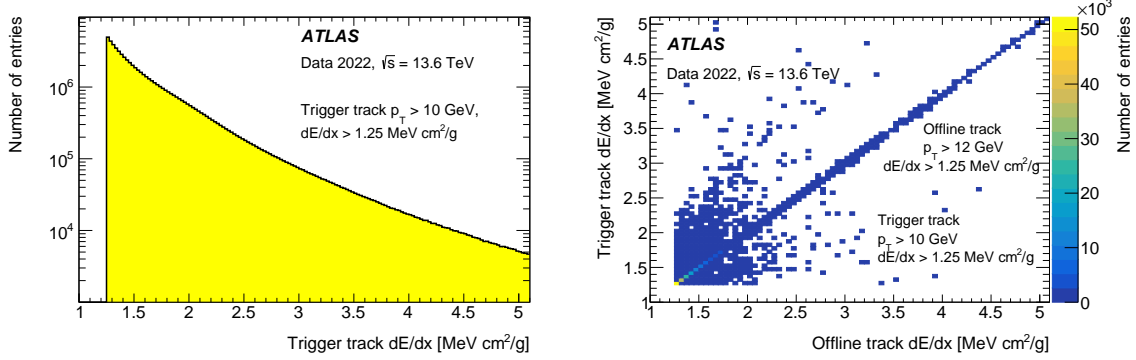


Figure 55. (Left) Distribution of dE/dx for tracks reconstructed in the trigger used for monitoring the dE/dx -based trigger. (Right) Correlation of online vs. offline dE/dx for offline tracks with $p_T > 12$ GeV.




Figure 56. Efficiency as a function of the chargino-pair system transverse momentum for the HLT E_T^{miss} trigger (dotted line) and a logical OR with the disappearing track trigger (solid line) with respect to the L1 $E_T^{\text{miss}} > 50$ GeV trigger. Both triggers are based on `tcpufit` E_T^{miss} algorithm. MC simulation with Run-2 conditions is used. The events are taken from models with a chargino lifetime of 1 ns and the chargino mass of either 91 GeV or in the range 200–1000 GeV with 100 GeV spacing. Only events with $E_T^{\text{miss}} > 60$ GeV and at least one chargino decaying between the pixel and SCT detectors (13–30 cm) in the central region, $|\eta| < 1.8$, are considered.

background. The disappearing track trigger selects events with at least one tracklet with $p_T > 20$ GeV that passes a stringent requirement on the BDT score. Figure 56 shows the expected performance of the trigger compared to that of the E_T^{miss} trigger with a 110 GeV threshold, which was used in Run 2, for a model with long-lived charginos [90]. The acceptance of the trigger to events with chargino momentum below 150 GeV is greatly improved.

6.9.2 Long-lived particle decays into jets

A second set of three new triggers have been developed targeting long-lived particle decays into jets, which may be displaced themselves or contain displaced tracks. The triggers make use of full scan tracking run after a E_T^{miss} or jet requirement and in some cases use LRT as an additional handle to select events.

Hit-based displaced vertex triggers. Neutral LLPs may travel some distance into the detector before decaying, resulting in a displaced vertex (DV) or jet. Previous searches relied on a variety of triggers

looking for other objects in the final state. Two hit-based DV triggers make use of hits not associated to tracks after the standard full scan tracking is performed in events passing the L1 jet and E_T^{miss} triggers, respectively. Jets with a large number of remaining hits on the outer layers of the ID, and few on the inner layers, are indicative of DVs with displaced tracks that are not reconstructed. Using this as an extra requirement allows for lower E_T^{miss} and jet thresholds compared to the lowest unrescaled triggers. A BDT trained on the fraction of hits in the layers of the ID is used to implement this selection on jets. The first trigger seeded by the same 80 GeV E_T^{miss} threshold as above selects events containing jets with $p_T > 200$ GeV and $|\eta| < 1$. Figure 57 shows its expected performance versus the number of pile-up interactions using MC simulation of a heavy Higgs boson decaying into two long-lived scalars, each subsequently decaying into two b -jets compared to a background process of $t\bar{t}$ with an all-hadronic final state. The signal efficiency is around 70%, compared to a background efficiency of less than 5%, for high numbers of additional interactions. The algorithm is tuned such that there is no strong dependence of the signal efficiency on pile-up. The second trigger seeded by a L1 jet with $p_T > 100$ GeV selects events containing jets with $p_T > 260$ GeV and $|\eta| < 1$. It requires a jet with $p_T > 180$ GeV at the HLT calorimeter preselection step before running the tracking.

Another hit-based trigger under development is seeded by a L1 trigger with $E_T^{\text{miss}} > 50$ GeV and runs only the fast tracking step of LRT in a composite RoI around the jets. Tracks with $p_T > 2$ GeV are then used to build vertices using a modified version of the offline secondary-vertexing algorithm [91]. The algorithm is optimised to be faster by requiring track pairs to be consistent with an approximate vertex position and by reducing the combinatorics of clustering by binning the vertex positions.




Figure 57. Expected performance of the hit-based displaced vertex trigger seeded by a E_T^{miss} threshold of 80 GeV versus the average pile-up. MC simulation is used for signal and background processes. The signal process is a heavy Higgs boson ($m_H = 1$ TeV) decaying into two scalars ($m_S = 50$ GeV) with a proper decay length of 9 mm, which subsequently decay into two b -jets each. The background process is $t\bar{t}$ production with an all-hadronic final state. Only statistical uncertainties are shown.

Emerging jet triggers. In the models of new phenomena with a dark sector, decays of particles into hadronically interacting particles may result in emerging jets with a large fraction of events possessing a DV and tracks with large d_0 . The fraction of momentum associated to prompt tracks in the jet relative to the total jet momentum is used to identify those with a large fraction of displaced tracks, i.e. a low momentum fraction associated to prompt tracks. This trigger is run in events that contain a L1 jet with $p_T > 100$ GeV and an HLT jet with $p_T > 200$ GeV in order to reduce the rate of full scan tracking. The p_T of tracks within $\Delta R < 1.2$ of the leading large-R jet axis are summed. Jets are

only considered if within $|\eta| < 1.8$ and the tracks must have $p_T > 2$ GeV, $|\eta| < 2.4$, $|d_0| < 8$ mm, and $|d_0| < 2.5 \cdot \sigma_{d_0}$ mm, where the d_0 resolution (σ_{d_0}) is computed from Z +jets events in data. The ratio of the sum of track p_T over the jet p_T is computed, and events are selected if this ratio, PTF, is less than 0.08. This extra requirement allows for a reduction of more than a factor of two of the jet p_T threshold as compared to that of the lowest unrescaled trigger (420 GeV in 2022).

A second trigger of this type is seeded by a 45 GeV photon and selects events with two large-R jets with $|\eta| < 2.0$, $p_T > 55$ GeV, and $\text{PTF} < 0.1$. Figure 58 (left) shows that the emerging jet trigger is efficient down to much lower jet p_T than the single large-R jet trigger. A model of a 1.5 TeV Z' decaying into two 20 GeV dark pions with a proper decay length of 50 mm is used [92]. The overall efficiency depends on the acceptance of the PTF requirement.




Figure 58. (Left) Efficiency of the new emerging jet trigger compared to a single large-radius jet trigger for a Z' decaying into two dark pions with a proper decay length ($c\tau$) of 50 mm. The efficiency is calculated as a function of the p_T of the leading offline anti- k_t $R=1.0$ jet, reclustered from $R=0.4$ topocluster jets. (Right) Efficiency of the LRT-based displaced jets triggers plotted against the leading offline anti- k_t $R=0.4$ jet p_T for a model of exotic decays of the Higgs boson into long-lived pseudoscalars a with mass 55 GeV and proper decay length of 100 mm. Both studies use MC simulation with Run-2 conditions. Only statistical uncertainties are shown.

Displaced jet triggers. Similar to the previous case, neutral LLPs decaying into quarks or gluons may result in displaced jets. Two new LRT-based triggers select such events with a jet from initial-state radiation and at least one or two displaced jets, but with a lower p_T threshold than the lowest unrescaled single-jet trigger at 420 GeV. Tracks with $p_T > 1$ GeV and $|\eta| < 2.4$ within $\Delta R < 0.4$ of the jet are counted as prompt ($n_{\text{prompt}}^{\text{trk}}$) or displaced ($n_{\text{disp}}^{\text{trk}}$) if they have $|d_0| < \text{or} > 3$ mm, respectively. After the standard tracking is run for the events that contain a L1 jet with $p_T > 100$ GeV and an HLT jet with $p_T > 180$ GeV, jets are preselected by requiring $n_{\text{prompt}}^{\text{trk}} \leq 2$. LRT is then run on the remaining hits in RoIs around the three leading jets passing the preselection. In addition to the leading jet $p_T > 180$ GeV requirement, the displaced single-jet trigger selects events with a $p_T > 140$ GeV jet with $n_{\text{disp}}^{\text{trk}} \geq 3$ and $n_{\text{prompt}}^{\text{trk}} \leq 1$. Similarly, the displaced di-jet trigger selects events with two $p_T > 50$ GeV jets where at least one must satisfy $n_{\text{disp}}^{\text{trk}} \geq 3$ and $n_{\text{prompt}}^{\text{trk}} \leq 2$, and the other may satisfy only $n_{\text{disp}}^{\text{trk}} \geq 0$ if $n_{\text{prompt}}^{\text{trk}} \leq 1$. Figure 58 (right) shows the expected performance of the displaced jets triggers based on MC simulation of a model with exotic Higgs decay into long-lived pseudoscalars a in the jet p_T region below the primary single jet trigger. Details of the model can be found in ref. [93].

6.9.3 Long-lived particle decays into SM leptons

LLPs may also decay into SM leptons resulting in displaced electrons, muons, and taus. Three new sets of triggers target these decays using LRT, and also standard tracking in the case of taus.

To select such events in Run 2, searches generally relied on triggers without tracking information such as photon triggers to select displaced electrons and MS-only triggers to select muons [94]. These triggers had high p_T thresholds of 50–120 GeV, requiring two objects for the 50 GeV threshold, and restrictions in η in the case of muons. New triggers for displaced electrons and muons directly trigger on these signatures allowing for lower thresholds with respect to those used in Run 2. The displaced electron trigger runs LRT in RoIs in events passing the same L1 threshold as the primary prompt electron chain, described in section 6.1. It selects events that have an electron with $p_T > 30$ GeV and $|d_0| > 3$ mm that passes a loose likelihood electron identification [5] without the use of d_0 or requirements on the number of hits in the Pixel detector. The displaced muon trigger runs LRT in RoIs in events passing the same L1 threshold as the primary prompt muon chain, described in section 6.2. It selects events that have a muon with $p_T > 20$ GeV and $|d_0| > 2$ mm. Figures 59 and 60 show the expected efficiency of the displaced electron and muon HLT triggers with respect to their L1 seeds in terms of the offline reconstructed lepton transverse impact parameter (d_0) and the production radius of offline reconstructed electrons and muons, respectively, using MC simulation of LLP di-stau production. MC simulation of pair production of staus with a 1 ns lifetime is used, and stau masses of 100–500 GeV in 100 GeV steps are merged. The offline electrons are required to pass the same loose likelihood identification as used in the trigger and have $|d_0| > 3$ mm. The offline muons are required to pass the medium working point described in ref. [45] without a cut on the number of pixel hits and have $|d_0| > 2$ mm. The acceptance of the LRT-based triggers extends out to the first layer of the SCT at 300 mm, where the layout of the detector no longer allows for eight hits on the track in most regions. The standard tracking runs out to a $|d_0|$ of 5 and 10 mm for electrons and muons respectively, limiting the acceptance of the standard prompt lepton triggers. The larger value for muon tracking increases the acceptance of B -meson decays. Combining the standard and LRT light lepton triggers provides continuous acceptance from small to large displacements.

Two dedicated triggers for displaced hadronically decaying taus are in development for Run 3. One is based on standard tracking and the other on LRT. A search during Run 2 with this signature was not performed, but would have needed to use jet or E_T^{miss} triggers with high thresholds. The standard tau identification RNN [95] is retrained to use standard tracking to identify displaced taus. Standard input samples are replaced by representative signal samples with displaced tau content, which additionally include association of large radius tracks in the tau reconstruction. MC samples with lifetimes of 0.01–100 ns are combined for the training. The first tau LLP trigger selects events containing a tau with $p_T > 200$ GeV passing the medium working point of this displaced tau identification. This identification is also used in multi-object triggers looking for a tau+X, which allows for lower thresholds on the tau object. Primary L1 di-tau, single muon, single isolated electron and E_T^{miss} triggers are used to seed them. The second tau LLP trigger runs LRT in an RoI around the calorimeter tau seed and uses the same RNN-based tau identification as trained for the previous trigger. Events are selected if they contain at least one tau passing the displaced tau identification. The trigger is expected to have a p_T requirement slightly lower than the lowest unrescaled single tau trigger, 160 GeV. Figure 61 shows the expected performance of these triggers compared to the prompt tau trigger based on an MC simulation of 100 GeV, 1 ns staus. The efficiency is computed with respect to standard and large




Figure 59. Displaced (left) electron and (right) muon trigger efficiencies with respect to their L1 seeds versus the offline lepton d_0 (LRT, open triangles). The efficiencies for the isolated primary single electron and muon triggers (Standard), described in sections 6.1.2 and 6.2.1, are shown as inverted triangles. MC simulation samples with Run-3 conditions are used. Only statistical uncertainties are shown.




Figure 60. Displaced (left) electron and (right) muon trigger efficiencies with respect to their L1 seeds versus the offline reconstructed lepton production radius (LRT, open triangles). The efficiencies for the isolated primary single electron and muon triggers (Standard), described in sections 6.1.2 and 6.2.1, are shown as inverted triangles. A logical OR between the LRT and standard triggers (Combination) is marked with open circles. MC simulation samples with Run-3 conditions are used. Only statistical uncertainties are shown.

radius offline tau tracks that are truth-matched to signal tau decays. Hence, multiple tracks from the same generated tau may be included in the computation. New LLP and LLP LRT tau triggers significantly improve sensitivity to the displaced tau signals.

7 Auxiliary triggers

7.1 Non-collision background

Non-collision backgrounds comprise detector signals which do not originate from the collisions of paired bunches at the ATLAS interaction point. They are categorised into beam-induced and




Figure 61. Displaced tau trigger efficiency as a function of truth signal tau (left) p_T and (right) decay radius of the tau. MC simulation with Run-3 conditions of staus with a mass of 100 GeV and a 1 ns lifetime is used. The efficiencies for the prompt tau trigger (inverted triangle), LLP tau trigger using standard tracking (square), and LLP tau trigger using LRT (triangle) are overlaid by the logical OR of all three (circle). The gray histogram shows the distribution of truth taus for this model. Only statistical uncertainties are shown.

cosmic-induced backgrounds. Both beam-induced and cosmic-induced backgrounds are non-negligible background sources in searches for new phenomena targeting delayed or displaced detector signatures, like those discussed in section 6.9. A number of dedicated triggers are used to select non-collision backgrounds for study.

7.1.1 Cosmic ray triggers

The cosmic-induced background results from the impact of extremely energetic cosmic muons on the detector, which induce signals in the muon spectrometers or energy deposits in the calorimeters. These events are collected either during LHC data taking (with L1 triggers in the empty bunch crossings which do not have pp collisions) or during periods with no LHC data taking (e.g. for detector commissioning as described in section 3.2). The cosmic-muon triggers during data taking do not run any muon reconstruction at the HLT and they only run for dedicated cosmic data taking when there is no beam. Further details on studies of the ATLAS detector performance with cosmic-ray muons can be found in ref. [96].

The `IDCosmic` stream is designed to have a total rate of about 20 Hz. It consists of two single L1 muon triggers with p_T thresholds of 3 GeV and 8 GeV at approximate rates of 8 Hz and 5 Hz respectively, the latter being unprescaled. In the same stream, a dedicated TRT-based L1 trigger for cosmics data taking uses a fast read-out path for groups of channels in the detector [97], running at a typical rate of around 10 Hz.

The `CosmicCalo` stream aims for a target rate of 5 Hz and includes legacy L1Calo EM triggers with $E_T > 3$ GeV and 7 GeV, a tau lepton trigger with $E_T > 8$ GeV and jet triggers with $E_T > 12$ GeV and 30 GeV for $|\eta| < 3.1$ and $E_T > 30$ GeV for $3.1 < |\eta| < 4.9$.

7.1.2 Beam-induced background triggers

Beam-induced background (BIB) [98] originates from (1) the inelastic interactions of protons with residual gas molecules upstream and nearby the detector, producing showers of secondary particles (beam-gas background), (2) protons with high transverse amplitude, or (3) from protons deflected

in beam-gas scattering hitting the tertiary collimators resulting in background/secondary particles entering the detector (beam halo). The online monitoring of BIB is essential to track live information about the beam conditions, which is also provided to the LHC. Further studies of BIB are essential to understand its origin and composition, and maintain and develop adequate monitoring.

Two single jet triggers with a prescaled L1 threshold of 12 GeV and an unprescaled L1 threshold of 50 GeV are used to record BIB events in bunch crossings where a proton bunch is present in only one or in neither of the beams. In addition, triggers based on hits in the Beam Conditions Monitor (BCM) [99] are used to record BIB events. The BCM consists of two stations of detectors located symmetrically around the interaction point at $z = \pm 184$ cm and $r = 55$ mm ($\eta \approx 4.2$). Each station has four modules of two diamond sensors read out in parallel. The implementation of the BCM hit-based trigger is based on a coincidence of an early hit on one side (A or C side) and one hit in-time with the bunch crossing on the other side (C or A side) in unpaired bunch crossings, where only one of the two beams is filled with a proton bunch. In order to be more independent of the presence of unpaired bunch crossings and to improve the purity of the selected events, a new logic is implemented for Run 3 which relies on two early hits on the same side (2A or 2C). This allows for the triggering on paired bunch crossings while maintaining a similar rate to the AC/CA counterparts, especially when triggering on the first colliding pair of bunches in a train which helps to mitigate against the impact of afterglow (increased cavern backgrounds following pp collisions at high μ). Furthermore, this allows for a measurement of the composition and fraction of BIB in paired bunch crossings, which was found to be different compared to unpaired bunch crossings. The contribution of ghost collisions¹⁰ is eliminated and thus the purity of the recorded BIB events is increased. A new bunch group with first paired bunch crossings in a train, containing all paired bunch crossings following a gap of at least 29 empty bunch crossings is defined. In order to evaluate the performance of the new triggers, they are defined in different flavours, triggering on empty, unpaired, and the aforementioned first paired bunch crossings in a train. Figure 62 shows the L1 rate of the old and new BCM triggers as a function of a bunch crossing identification (BCID) value. For each trigger the per-BCID rate is shown for five different bunch groups, paired colliding bunches, the first colliding bunches in a train, empty bunch crossings, unpaired bunch crossings with either beam 1 or beam 2 filled with a proton bunch. As expected the rate of the 2A/2C variants is much lower than the AC/CA variants in BCIDs with paired colliding bunches. Further, the rate in unpaired BCIDs is similar for both variations.

7.2 Zero-bias trigger

The zero-bias data are collected using a dedicated trigger, which fires one LHC turn after a L1 EM trigger with $E_T > 15$ GeV fires. This approach allows to collect data which is unbiased with respect to the activity in the event and at the same time proportional to the luminosity in each bunch crossing, which can not be achieved with random triggers. Such triggers are used for detector studies as well as for the MC simulation overlay method [100], which relies on the zero-bias data to account for the pile-up background, cavern background, and detector noise. The zero-bias trigger, prescaled to a constant rate of about 10 Hz, was based on the legacy L1Calo EM trigger in 2022. A new Run-3 zero-bias algorithm is implemented in the new L1Topo system to be used following decommissioning

¹⁰Ghost collisions occur between protons from a filled bunch, which typically has $> 10^{11}$ protons, and protons in an unfilled (empty) bunch, which has $< 10^8$ protons, due to diffusion from filled bunches which takes place at the interaction point.




Figure 62. Per-BCID L1 rate of BCM hit-based triggers based on coincidence of (top left) AC and (top right) 2A, and (bottom left) CA and (bottom right) 2C. The per-BCID rates are shown for paired colliding bunches, the first colliding bunches in a train, empty bunch crossings (orange), unpaired bunch crossings with either beam 1 or beam 2 filled with a proton bunch. The BCID range shown is restricted to the first 350 BCIDs.

of the L1Calo legacy system, because the data connections between the new L1Calo system and CTP prevent the use of the legacy zero-bias trigger.

7.3 Triggers for tracking performance studies

There are dedicated triggers in the trigger menu to study lepton tracking performance, described in section 5.1. These `idperf` chains run the same selection as the corresponding triggers under study but do not apply any requirements on tracking. This means that they are unbiased with respect to tracking and can be used for efficiency measurements. For instance, if an electron candidate is formed from a track and a cluster, the `idperf` electron chain does not make any selection on the electron candidate, as this would include requirements on track-cluster matching to define the candidate, as described in section 6.1.

Two types of `idperf` triggers are employed in the trigger menu: di-lepton triggers and single lepton triggers. The di-lepton `idperf` triggers target Z boson or J/ψ decays to leptons in order to study the performance with the tag-and-probe method. These chains are based on primary single or di-object electron and muon triggers, described in sections 6.1 and 6.2, and can run unrescaled at about 15–30 Hz not adding any unique rate. These `idperf` triggers do not apply any requirements on tracking on the subleading trigger leg which can be an electron, muon or tau candidate. The single object

idperf triggers have thresholds of 5, 14, 20, 26 and 30 GeV for electrons, 6, 20, 24 and 40 GeV for muons and 25, 35, 80 and 160 GeV for taus. These are heavily prescaled to record events at 0.5–1 Hz.

There is at most one reference object (one track) per RoI from the leptonic triggers. The purity of the muon triggers is quite high. For the tau trigger, it is very low, and for the electron trigger, lower still. There are very few real electrons in the sample accumulated by single electron idperf triggers, so the dilepton idperf triggers are crucial to evaluate electron tracking performance.

7.4 Triggers for detector performance studies (calibration, noise)

Triggers used for detector performance studies often run at very high rates. To compensate for this, only partial detector information is recorded through a strategy called Partial Event Building (PEB), which has the potential to significantly reduce the event size and thus overall bandwidth. These triggers operate either with a fixed subset of ROBs which are to be included in events recorded to the PEB stream, or a list of ROBs can be derived dynamically on an event-by-event basis based on a set of RoIs. A combination of both static and dynamic ROB lists is also possible. These triggers are used for the calibration of muon subdetectors, low- p_T offline muon calibration and LAr calorimeter performance studies.

For example, a dedicated calibration trigger, called *a laser trigger*, is active during collisions runs to monitor the stability of the Tile calorimeter channels. The LHC abort gap of about $3 \mu\text{s}$ is used to send and register laser pulses that are arbitrated, timed and controlled with respect to the LHC signals [101]. Laser pulses are sent to the Tile calorimeter at the increased rate of 12 Hz for Run 3 and recorded for an analysis. Laser events are used to track possible fast gain changes in the photomultiplier tubes and monitor the timing calibration.

Random events with PEB information for Pixel and SCT detectors are recorded in the unfilled LHC bunches into dedicated streams to identify noisy channels which need to be masked^u in these subdetectors. These triggers typically run at about 10 Hz. A pixel with noise hit occupancy above 5×10^{-4} is classified as a noisy channel and masked in DAQ.

7.4.1 LArNoiseBurst algorithm

The LAr calorimeter has a small but luminosity-dependent probability of generating large noise signals involving a considerable number of cells, as discussed in detail in ref. [102]. While this happens with milli-Hertz frequency, high spurious values of energy are provided by these cells and it is therefore important to veto affected events from the data quality monitoring and downstream of this from potentially affected physics analyses. Veto windows occur across a much shorter timescale than that of a single luminosity block. The detection and registration of LAr noise bursts is done via a LArNoiseBurst algorithm which runs in the HLT.

Noise bursts generate electrical pulses which do not correspond to physics objects and thus these have a bad quality of fit when energy reconstruction is attempted. Offline tools examine the quality factor of the cell energy reconstruction procedure, counting the number of cells with a bad quality factor in a given event. Given that many other factors could interfere with the cell quality factor (e.g. intense pile-up in a particular event), a minimum of two events that are strongly correlated in time is required to declare a veto interval.

^uNot used for the track reconstruction.

In the HLT, the offline algorithm is run for every accepted event that is written to certain streams, including the *Main* stream, the TLA, delayed streams and streams recording events with cosmic signatures in the calorimeter. Events that are considered bad by the offline tool have their time stamp published online via an information messaging system. A dedicated application then picks the time stamps, performs a coincidence operation between them (a maximum time difference of $250\ \mu\text{s}$ is required) and declares the interval to a database used by the offline reconstruction. The veto is then applied to reject events in the post-run analysis. The computing resources taken by this algorithm are negligible due to very low rate of the Noise Bursts and large overlap with other calorimeter triggers.

7.5 Beam-spot algorithms

A dedicated system allows the HLT to continuously measure the position, size and orientation of the luminous region (also known as beam spot) at the ATLAS interaction point [103, 104]. Beam-spot parameters can change from fill to fill, and some, in particular the transverse position and size, show significant variation over the course of a fill as well. Knowledge of the current beam-spot parameters is crucial for several HLT algorithms, most notably for the selection of events with b -jets, and for HLT tracking itself. The parameters are also transmitted online to the LHC status display [105].

The beam-spot parameters are continuously monitored and archived, determined bunch-by-bunch as well as fill average, and sampled over different intervals from one to many minutes depending on the required statistics. One dedicated set that is used by the HLT is only updated whenever significant deviations from the currently used values are detected. The process of obtaining the parameters, updating them and feeding them to the algorithms as condition parameters is referred to as beam-spot calibration.

The calibration and update process is not trivial as it consists of multiple steps involving several cooperating sub-systems to synchronise the beam-spot parameters across the distributed HLT farm at the same time. This requires orchestrating a series of steps, primarily through the CTP:

- HLT algorithms extract tracking and vertexing information and publish their distributions in the form of histograms,
- histograms from all individual HLT instances are aggregated by a monitoring infrastructure referred to as the Gatherer,
- on each new luminosity block the merged histograms are processed and new estimates of the beam-spot parameters are calculated by an external application called the BeamspotTool,
- when the BeamspotTool determines that the new estimate constitutes a significant change with respect to the previous one, it sends these new parameters to the CTP process,
- the CTP process writes new beam-spot parameters to the conditions database with a validity interval starting with the next luminosity block and then notifies the HLT processes of the pending update via its event fragment,
- individual HLT processes read the new beam-spot parameters from the conditions database when they receive events from the next luminosity block, for which the CoralServer and CoralProxy [106] infrastructure provides scalable access to the conditions database from the HLT farm.

The online beam-spot calibration received significant improvements in preparation for Run 3.

One long-standing issue with the online beam-spot calibration during Run 2 was a small (5–15 μm) systematic difference observed in the transverse beam position with respect to the offline calibration. During LS2 this was tracked down to an incorrect transverse position at the coordinate system origin being used for track clustering, which caused a systematic bias towards the origin. The issue was resolved by using the current estimate of the transverse beam location for track clustering. Reprocessing of the Run-2 data with the fixed vertex finder showed the online calibration matching the offline measurement with excellent precision. In 2022, there was still an observable difference in transverse position between offline and online due to different alignment constants, which makes a direct comparison difficult.

The original calibration used for Run 1 and Run 2 is based on optimised vertex finding and fitting algorithms. The Run-3 calibration introduces an additional algorithm based only on track information that does not involve vertexing. The luminous region is narrower than the typical vertexing uncertainty, and an accurate estimation of the transverse widths requires precise knowledge of the vertexing resolution. The resolution itself depends on detector and trigger conditions, and is therefore evaluated in real time through a split-vertex method that was introduced in Run 1. However, the smaller beam-spot sizes of Run 2 required larger resolution corrections and the vertex method has shown limitations in case of low statistics. To study alternatives, a new method was developed for Run 3. This method utilises information from reconstructed tracks only and determines beam-spot parameters by a likelihood fit to the observed set of d_0 and ϕ parameters of tracks. In the ATLAS online environment the tracking information is local to each of the HLT processing nodes, and it is not possible to collect all track data at one location to perform fitting. To support fitting of all available tracking data at a single location, the likelihood function is approximated to a sufficiently small set of additive terms. These terms are calculated from the local set of tracks on each HLT node and merged for the final fit by utilizing the Gatherer infrastructure. The new method demonstrated reasonable performance with the reprocessed Run-2 data and a shorter ramp-up period between the start of data collection and calibration availability compared to the vertex-based method. For Run 3 both vertex-based and track-based algorithms are in use to further study their performance. During the initial 2022 data-taking period the track-based algorithm was used preferentially due to its better robustness, lower demand on statistics and earlier availability of the beam spot. For high-luminosity running the vertex algorithm is still the preferred method for its accuracy over a wider range in beam sizes.

The multi-step calibration process outlined above introduces a delay in availability of the calibrated beam-spot position for HLT algorithms, with a typical latency of several luminosity blocks. This is not an issue for the HLT during a run since the variations are small on that timescale. However, at the start of each new fill this requires a short bootstrap period, during which certain HLT algorithms, such as those used for b -tagging, have to be held off until a first measurement of the beam spot has succeeded. This results in a loss of a few luminosity blocks of data for the individual triggers: on average 4.5 luminosity blocks were missing beam-spot calibration at the start of each data-taking run in Run 2, which corresponds to approximately one percent of a duration of a typical run.

Improvements in the Gatherer infrastructure helped to reduce the calibration bootstrap delay. The calibration depends on histograms produced by the HLT and merged by the Gatherer. Improvements in handling of histograms for each of the luminosity blocks in the Gatherer have reduced propagation delays for those types of histograms. This results in a faster availability of those histograms for the BeamspotTool application and helps in shortening the inherent ramp-up delay, which has reduced to an average of 2.5 luminosity blocks in Run 3.

With the higher bunch intensities achieved by the injector upgrades, the LHC had for the first time begun β^* -levelling in order to limit pile-up in 2022. This involves a large variation in spot sizes, in turn demanding more frequent updates of the HLT during the levelling phase between β^* of 60 cm to 30 cm. In 2023 the levelling range is extended to be between 120 cm and 30 cm.

Another development begun in 2022 that can be applied in 2023 is to constrain HLT tracking to a 3σ window around the longitudinal beam position which is projected to save a significant fraction of CPU time in the HLT. This will involve yet another bootstrap procedure to settle on the beam spot at the start of each run.

Figure 63 shows the time evolution of the vertical beam-spot position as measured, and subsequently applied, by the HLT over the course of one LHC fill during November 2022 data taking. An update of all parameters is performed whenever the position changes by more than 10% of the width, or the width changes by more than 10%, or the uncertainty on any of the parameters decreases by more than half.

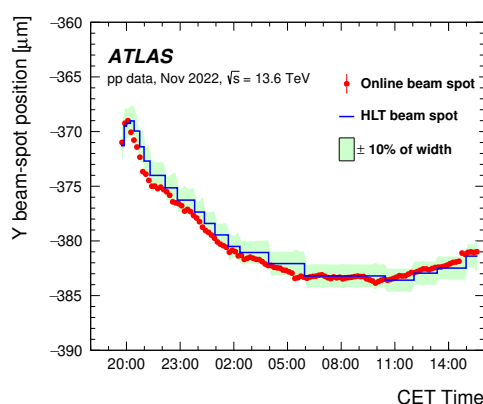


Figure 63. Time evolution of the vertical beam-spot position over the course of one LHC fill. The dots represent the monitored beam-spot position measured in luminosity blocks. The line represents the position currently used by the HLT algorithms which is measured as an average over a sliding window of a series of luminosity blocks. The band indicates a range corresponding to 10% of the vertical beam-spot width that was measured at the same time.

8 Trigger software performance

8.1 Trigger software performance scaling

For Run 3 the same process forking architecture of the HLTMPPU was kept as in Run 2. However, due to the multithreaded event selection, each worker process can now contain multiple event slots, which are processed in parallel. In Run 2, each worker process handled one event, and the selection progressed sequentially. This allowed the HLTSV to assign new events directly to idle worker processes. In Run 3, the AthenaMT scheduler decides when an event slot is freed and when a new event can be read from the input source. The HLTSV, therefore, assigns events to an event queue for each HLTMPPU from which new events are pulled by the AthenaMT scheduler for each worker process.

In the Run-3 system, the number of forked worker processes per HLTMPPU and the number of parallel event slots in each worker process are freely configurable. This allows for an adiabatic transition of a pure multi-process-based event selection like in Run 2 (i.e. many worker processes with




Figure 64. (Left) Application throughput in events/s, (middle) CPU usage (CPU time divided by wall time) in percent and (right) memory usage in GB as a function of the number of events processed in parallel for AthenaMT executing trigger selection algorithms. Blue squares represent a multi-processing approach. Pink circles represent a multi-threading approach with a single process using a number of threads equal to the number of events requested to process in parallel. Threads are not bound to events; instead, a pool of a number of threads (N_{thread}) is used to process an equivalent sized pool of events ($N_{\text{event}} = N_{\text{thread}}$). Green diamonds and orange triangles represent a hybrid approach where a number of processes, $N_{\text{process}} = N_{\text{event}}/N_{\text{thread}}$, forked after initialisation use a fixed number of threads each ($N_{\text{thread}} = 4$ and 8, respectively) to process an equivalent number of events in parallel. Differences between these approaches are discussed in section 3.4.

one event slot each) to a pure multithreaded configuration with one worker process and many parallel event slots. Event throughput measurements are used to decide what configuration will be used.

Figure 64 shows the trigger software performance scaling in terms of application throughput in events/s (left), CPU usage (middle), and memory usage (right) as a function of the number of events processed in parallel for AthenaMT executing trigger selection algorithms. The measurements were performed in a standalone local environment using a machine identical to those used in the ATLAS HLT computing farm during data taking. It is a dual processor machine with 128 GB RAM using a NUMA memory architecture and two AMD EPYC 7302 CPUs, where each CPU has 16 real cores with two hyper-threads per core, giving the total number of 64 threads. The data sample contains a mix of events representative of the real HLT input data and trigger selection configuration identical to one used during data taking. Four ways of achieving the parallelism are presented. Data were taken in 2022 using a pure multi-process configuration with 48 forks as event throughput is the most critical metric for the HLT. Different constraints apply, however, to Grid [31] processing of MC simulated events. Here, available memory per core is significantly more restricted than on the dedicated HLT machines, and these tasks execute in Run 3 in a pure multi-threaded configuration, the trigger simulation included. A typical Grid site in 2022 ran the multi-threaded configuration with eight event slots. As ATLAS’ transition to multi-threading was new for Run 3, a number of components still make use of mutex locking to provide safe shared access to certain common resources. In aggregate, these bottlenecks severely limit event throughput at very high levels of multi-threaded execution, as seen in figure 64. Work continues to refactor these components to minimise blocking behaviour in future software releases for both Run 3 and Run 4. In addition, the hybrid modes of operation continue to be studied as significant memory savings are made even in hybrid modes which use a small number of multi-threading event slots, hence minimising losses due to resource contention.

Figure 65 shows the calorimeter processing time dependency on the number of threads obtained offline by concurrently processing various numbers of data events with different thread configurations.




Figure 65. Mean execution time per call of (top left) the fast calorimeter reconstruction, (top right) the unpacking of the full scan calorimeter cells and (bottom) the full scan topo-clustering algorithm as a function of the number of threads for various numbers of concurrently processed events.

In the Run-3 AthenaMT framework, the same block of cells cannot be requested by two different algorithms (e.g. two overlapping RoIs or a full scan request at the same time as any RoI) as it may cause simultaneous reading and writing to a given memory area. To avoid this, locks are added to the HLTCalo data preparation service, leading to a non-linear dependency between the processing time and the number of parallel processing threads allocated for intra-event processing as shown in figure 65. Unpacking of the calorimeter cells and topological clustering algorithms do not show any dependence on the number of threads, contrary to fast reconstruction, which exhibits turn-on-like dependence on the number of threads the shape of which depends on the number of concurrently processed events. Optimisations are planned to either employ a more fine-grained locking or through an updated processing model. The time of the full scan calorimeter cell unpacking per call is slightly extended (mostly due to locks associated to RoIs) with the increased number of threads, but more linear scaling performance is achieved through increasing the number of inter-event parallel processing slots. This does not affect the HLTCalo algorithm’s functional performance, such as cell and cluster parameters or reconstruction efficiency.

8.2 HLT processing time

The overall utilisation of the HLT is dependent on the mean time taken to process each event and the incoming event rate from L1. The mean time to process an event is a product of the trigger menu, its associated luminosity-dependent prescale sets, and of the collisions being supplied by the LHC. Here the processing time is most heavily dependent on the mean number of pile-up interactions. An




Figure 66. (Left) Mean HLT wall time as a function of the average pile-up throughout a run. A vertical line marks the instantaneous luminosity at which additional trigger selections are enabled. Error bars denote the Gaussian width of the underlying per event measurements. (Right) HLT processing time distribution per event for an instantaneous luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ at pile-up of 51. Shown are the wall time spent per event as well as the sum of the algorithms only processing time. The difference between these two distributions is due to the time spent in framework operations.

example of the mean HLT processing time per event at different instantaneous luminosity values throughout a run is presented in figure 66 (left). The mean HLT processing time decreases with decreasing average pile-up due to a reduction in event complexity. The menu is adjusted based on the current rate of events in order to fully use available resources while retaining a constant 5000 slot safety margin (about 6% of the final 2022 farm) to protect against losing slots due to power glitches, technical problems, etc. The slope of the distribution in figure 66 (left) becomes less steep at lower μ values where additional trigger selections are enabled. Additional event processing restrictions came from the delayed delivery of replacement ROS servers. In 2022 the ROS (like in Run 2) was able to supply full detector information at half of the maximum L1 rate (50 kHz). However, it is possible to reach the L1 rate of 100 kHz for a subset of detectors by re-configuring the read-out, for example by increasing the number of ROS servers assigned to it.

In addition to the online monitoring of high-level HLT processing statistics, offline tools, such as Cost Monitoring [33], are available to investigate the detailed performance of the HLT (e.g. the average processing times of individual algorithms or whole trigger chains). Data required for such studies is saved to a dedicated calibration stream which enables recording of this data for all events, not just those accepted. This stream contains information on algorithms' execution times and on any data requests they make to the ROSes. By default, only the first 250 LBs are monitored in this way during physics data taking. After the run has finished, the data from the stream undergo additional post-processing and the performance details are available on a dedicated website.

The Cost Monitoring used during Runs 1 and 2 had to be adapted to the multi-threaded framework for Run 3. While the following figures and tables present data taken using the MP configuration in 2022, some changes and challenges which become more relevant in more MT-like configuration are pointed out in the following. Upgraded monitoring includes a new MT-compatible algorithm gathering the data, a redesigned post-processing framework, and additional monitors (including thread monitoring). Given that the new HLT runs under the AthenaMT scheduler, the analysis of the results is no longer as simple as in a sequential framework, therefore multiple metrics were prepared to provide an overview of the event processing time, for example, an event wall time or the sum of the

execution times of all algorithms. The wall time includes the framework operations/delays between scheduled algorithms. A comparison between those values can be seen in figure 66 (right), which shows the wall time which includes both algorithm execution time and time spent on framework operations (including algorithms scheduling, data traffic) as well as the total time of just the algorithms. These data were recorded in a MP-mode, where the scheduler was only tasked with processing a single event at a time using a single CPU core.

In the figure, three peaks can be identified, representing fast (approximately 30 ms), medium (approximately 300 ms), and slow (approximately 2 s) events. The last type of the event is the rarest due to the early rejection mechanism. For faster events, when only fast algorithms are executed and most of the trigger chains are rejected, the impact of steering operations from the underlying framework is observable as an overhead of about 10 ms. For events with execution times much longer than 100 ms, when time-consuming algorithms, for example, tracking algorithms, are executed, it is negligible. Table 5 shows an example of collected Cost Monitoring metrics for a given period of time during a run, broken down into the main types of object reconstruction and the framework. The Cost Monitoring data are used in an iterative process to identify areas where optimisation of code and strategy can yield the greatest impact with zero or negligible physics impact. The difference between the wall time and the total time of all algorithms per event can further diverge in different ways when running in MT mode. For the MP case, the total event wall time is always greater than the algorithm time, but with sufficient intra-event parallelism in an MT configuration, the total event wall time can be smaller than the algorithm time, as some of the event processing may occur simultaneously on multiple cores.

Table 5. Example of collected Cost Monitoring metrics for 50 LBs for an instantaneous luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ at pile-up of 51, showing how the total time is distributed between the main types of object reconstruction and the framework software. The fractional time consumption of algorithms is calculated based on the sum of the execution times of all algorithms.

	Total time [%]
ID reconstruction	59
Muon reconstruction	14
Calorimeter reconstruction	11
Combined reconstruction and hypothesis algorithms	8
Trigger infrastructure	2
Other	6

At the peak instantaneous luminosity during the run, the HLT processing time of one event is approximately 600 ms (590 ms algorithm time), compared to 500 ms achieved during Run 2. Out of this number about 59% of the total event time is spent on the ID tracking, despite the improvements discussed in section 5.1. This is due to the expanded use of the full scan tracking for the hadronic signatures in Run 3. The HLT farm computing capacity has been increased for Run 3 by routine replacement of old servers, benefiting from the sustained industry trend of increasing processor performance. This has enabled the additional HLT processing for Run 3 that is described in this paper to support the physics goals of the experiment. During Run 2, the HLT farm consisted of processors with a performance of 22.8 HS06/core [107] (0.8 MHS06 total farm performance). For Run 3, 60% of the farm was upgraded to 36.2 HS06/core (1.7 MHS06 total farm performance) for early 2022, rising to 100% by November 2022 (2.0 MHS06 available for the remainder of Run 3 starting from 2023).

9 Conclusion

A large number of trigger upgrades and developments for the ATLAS experiment were made during the second long shutdown of the LHC in preparation for the Run-3 data taking. A summary of the various updates as well as the first Run-3 performance studies can be found in this paper.

Many changes in the L1 trigger system improve both its rejection of background events and acceptance for interesting physics processes. Upgrades in the L1 calorimeter trigger increased the granularity of information used by the trigger and enable it to run more sophisticated algorithms to identify physics objects, and to calculate missing transverse momentum in the event with higher precision. The L1 muon system was enhanced through the addition of New Small Wheels, new resistive plate chambers in the barrel/endcap transition region as well as upgraded electronics. The topological trigger, installed during Run 2, has also undergone a hardware upgrade. Its refined kinematic measurements of muons and other objects as well as new flexibility to define multiplicity triggers allow for more sophisticated event selections at L1.

The underlying framework of the HLT was completely rewritten in order to execute trigger algorithms within the multi-threaded software framework AthenaMT, sharing common features between the trigger software and the software used for offline reconstruction. These changes help to optimise the use of the HLT computing resources, both in terms of computing power and memory consumption. New-and-improved selection algorithms and strategies further improve the reconstruction of objects at the trigger level, in the case of some of the hadronic signatures, at the expense of higher CPU time needs due to the expanded use of full scan tracking.

While maintaining a level of consistency with the Run-2 trigger menu, the Run-3 trigger menu sets out to exploit the newly implemented detector features, more performant HLT hardware, and algorithmic advancements. Trigger thresholds at L1 and HLT were generally kept the same as during Run 2, benefiting from improvements to reduce trigger rate. The additional available HLT rate compared to Run 2 is dedicated to expanding the physics menu in the physics and TLA streams, lowering trigger thresholds and including new triggers for previously unexplored phase space, which make extensive use of the various inner detector tracking algorithms in the HLT.

The ATLAS trigger system was successfully (re-)commissioned with the first data acquired at 13.6 TeV, with some final commissioning steps for L1Calo and L1Muon to be completed during Run 3. First performance studies of the different trigger signatures and trigger efficiencies with respect to the offline quantities are presented using the 13.6 TeV proton-proton collision data collected during 2022.

Acknowledgments

We thank CERN for the very successful operation of the LHC and its injectors, as well as the support staff at CERN and at our institutions worldwide without whom ATLAS could not be operated efficiently.

The crucial computing support from all WLCG partners is acknowledged gratefully, in particular from CERN, the ATLAS Tier-1 facilities at TRIUMF/SFU (Canada), NDGF (Denmark, Norway, Sweden), CC-IN2P3 (France), KIT/GridKA (Germany), INFN-CNAF (Italy), NL-T1 (The Netherlands), PIC (Spain), RAL (U.K.) and BNL (U.S.A.), the Tier-2 facilities worldwide and large non-WLCG resource providers. Major contributors of computing resources are listed in ref. [108].

We gratefully 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, The 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, Taipei; TENMAK, Türkiye; STFC, United Kingdom; DOE and NSF, United States of America.

Individual groups and members have received support from BCKDF, CANARIE, CRC and DRAC, Canada; CERN-CZ, PRIMUS 21/SCI/017 and UNCE SCI/013, Czech Republic; COST, ERC, ERDF, Horizon 2020, ICSC-NextGenerationEU 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.

In addition, individual members wish to acknowledge support from CERN: European Organization for Nuclear Research (CERN P.J.A.S.); Chile: Agencia Nacional de Investigación y Desarrollo (FONDECYT 1190886, FONDECYT 1210400, FONDECYT 1230812, FONDECYT 1230987); China: National Natural Science Foundation of China (NSFC — 12175119, NSFC 12275265, NSFC-12075060); Czech Republic: PRIMUS Research Programme (PRIMUS/21/SCI/017); European Union: H2020 European Research Council (ERC — 101002463), European Research Council (ERC — 948254), Horizon 2020 Framework Programme (MUCCA — CHIST-ERA-19-XAI-00), European Union, Future Artificial Intelligence Research (FAIR-NextGenerationEU PE00000013), Italian Center for High Performance Computing, Big Data and Quantum Computing (ICSC, NextGenerationEU), Marie Skłodowska-Curie Actions (EU H2020 MSC IF GRANT NO 101033496); France: Agence Nationale de la Recherche (ANR-20-CE31-0013, ANR-21-CE31-0013, ANR-21-CE31-0022), Investissements d’Avenir IDEX (ANR-11-LABX-0012), Investissements d’Avenir Labex (ANR-11-LABX-0012); Germany: Baden-Württemberg Stiftung (BW Stiftung-Postdoc Eliteprogramme), Deutsche Forschungsgemeinschaft (DFG — 469666862, DFG — CR 312/5-1); Italy: Istituto Nazionale di Fisica Nucleare (FELLINI G.A. n. 754496, ICSC, NextGenerationEU); Japan: Japan Society for the Promotion of Science (JSPS KAKENHI JP21H05085, JSPS KAKENHI JP22H01227, JSPS KAKENHI JP22H04944, JSPS KAKENHI JP22KK0227); Netherlands: Netherlands Organisation for Scientific Research (NWO Veni 2020 — VI.Veni.202.179); Norway: Research Council of Norway (RCN-314472); Poland: Polish National Agency for Academic Exchange (PPN/PPO/2020/1/00002/U/00001), Polish National Science Centre (NCN 2021/42/E/ST2/00350, NCN OPUS nr 2022/47/B/ST2/03059, NCN UMO-2019/34/E/ST2/00393, UMO-2020/37/B/ST2/01043, UMO-2021/40/C/ST2/00187); Slovenia: Slovenian Research Agency (ARIS grant J1-3010); Spain: BBVA Foundation (LEO22-1-603), Generalitat Valenciana (Artemisa, FEDER, IDIFEDER/2018/048), La Caixa Banking Foundation (LCF/BQ/PI20/11760025), Ministry of Science and Innovation (MCIN & NextGenEU PCI2022-135018-2, MICIN & FEDER PID2021-125273NB, RYC2019-028510-I, RYC2020-030254-I, RYC2021-031273-I, RYC2022-038164-I), PROMETEO and GenT Programmes Generalitat Valenciana (CIDEAGENT/2019/023, CIDEAGENT/2019/027); Sweden: Swedish Research Council (VR

2018-00482, VR 2022-03845, VR 2022-04683, VR grant 2021-03651), Knut and Alice Wallenberg Foundation (KAW 2017.0100, KAW 2018.0157, KAW 2018.0458, KAW 2019.0447); Switzerland: Swiss National Science Foundation (SNSF — PCEFP2_194658); United Kingdom: Leverhulme Trust (Leverhulme Trust RPG-2020-004); United States of America: U.S. Department of Energy (ECA DE-AC02-76SF00515), Neubauer Family Foundation.

References

- [1] ATLAS collaboration, *Performance of the ATLAS Trigger System in 2010*, *Eur. Phys. J. C* **72** (2012) 1849 [[arXiv:1110.1530](#)].
- [2] ATLAS collaboration, *Performance of the ATLAS Trigger System in 2015*, *Eur. Phys. J. C* **77** (2017) 317 [[arXiv:1611.09661](#)].
- [3] ATLAS collaboration, *The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3*, *2024 JINST* **19** P05063 [[arXiv:2305.16623](#)].
- [4] ATLAS collaboration, *Preliminary analysis of the luminosity calibration of the ATLAS 13.6 TeV data recorded in 2022*, *ATL-DAPR-PUB-2023-001* (2023).
- [5] ATLAS collaboration, *Performance of electron and photon triggers in ATLAS during LHC Run 2*, *Eur. Phys. J. C* **80** (2020) 47 [[arXiv:1909.00761](#)].
- [6] ATLAS collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *2008 JINST* **3** S08003.
- [7] B. Abbott et al., *Production and Integration of the ATLAS Insertable B-Layer*, *2018 JINST* **13** T05008 [[arXiv:1803.00844](#)].
- [8] G. Aad et al., *The Phase-I trigger readout electronics upgrade of the ATLAS Liquid Argon calorimeters*, *2022 JINST* **17** P05024 [[arXiv:2202.07384](#)].
- [9] A. Armbruster et al., *The ATLAS Muon to Central Trigger Processor Interface Upgrade for the Run 3 of the LHC*, in the proceedings of the *2017 IEEE Nuclear Science Symposium and Medical Imaging Conference*, Atlanta, GA, U.S.A., 21–28 October 2017, pp. 1–5 [[DOI:10.1109/NSSMIC.2017.8532707](#)].
- [10] A. Sidoti, *Minimum Bias Trigger Scintillators in ATLAS Run II*, *2014 JINST* **9** C10020.
- [11] L. Adamczyk et al., *Technical Design Report for the ATLAS Forward Proton Detector*, *CERN-LHCC-2015-009* (2015) [ATLAS-TDR-024].
- [12] S. Abdel Khalek et al., *The ALFA Roman Pot Detectors of ATLAS*, *2016 JINST* **11** P11013 [[arXiv:1609.00249](#)].
- [13] G. Avoni et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, *2018 JINST* **13** P07017.
- [14] ATLAS collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, *Eur. Phys. J. C* **83** (2023) 982 [[arXiv:2212.09379](#)].
- [15] ATLAS collaboration, *Zero degree calorimeters for ATLAS*, *CERN-LHCC-2007-001* (2007).
- [16] H. Bertelsen et al., *Operation of the upgraded ATLAS Central Trigger Processor during the LHC Run 2*, *2016 JINST* **11** C02020.
- [17] C. Leggett et al., *AthenaMT: upgrading the ATLAS software framework for the many-core world with multi-threading*, *J. Phys. Conf. Ser.* **898** (2017) 042009.

- [18] ATLAS TDAQ collaboration, *The ATLAS Data Acquisition and High Level Trigger system*, 2016 *JINST* **11** P06008.
- [19] ATLAS collaboration, *The ATLAS collaboration Software and Firmware*, ATL-SOFT-PUB-2021-001 (2021).
- [20] ATLAS collaboration, *Operation of the ATLAS trigger system in Run 2*, 2020 *JINST* **15** P10004 [[arXiv:2007.12539](#)].
- [21] ATLAS collaboration, *Performance of the upgraded PreProcessor of the ATLAS Level-1 Calorimeter Trigger*, 2020 *JINST* **15** P11016 [[arXiv:2005.04179](#)].
- [22] ATLAS collaboration, *Performance of the ATLAS muon triggers in Run 2*, 2020 *JINST* **15** P09015 [[arXiv:2004.13447](#)].
- [23] ATLAS collaboration, *Performance of the ATLAS Level-1 topological trigger in Run 2*, *Eur. Phys. J. C* **82** (2022) 7 [[arXiv:2105.01416](#)].
- [24] ATLAS collaboration, *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System*, CERN-LHCC-2013-018 (2013) [ATLAS-TDR-023].
- [25] ATLAS collaboration, *Athena*, version 22.0.1, Zenodo (2019) [[DOI:10.5281/zenodo.2641997](#)].
- [26] G. Barrand et al., *GAUDI — A software architecture and framework for building HEP data processing applications*, *Comput. Phys. Commun.* **140** (2001) 45.
- [27] M. Clemencic, B. Hegner, P. Mato and D. Piparo, *Introducing concurrency in the Gaudi data processing framework*, *J. Phys. Conf. Ser.* **513** (2014) 022013.
- [28] UXL Foundation, *oneAPI Threading Building Blocks (oneTBB)*, <https://github.com/oneapi-src/oneTBB>.
- [29] ATLAS collaboration, *ATLAS Future Framework Requirements Group Report*, ATL-SOFT-PUB-2016-001 (2016).
- [30] F. Pezoa, J.L. Reutter, F. Suarez, M. Ugarte and D. Vrgoč, *Foundations of JSON Schema*, in the proceedings of the *25th International Conference on World Wide Web (WWW '16)*, Montréal, QC, Canada, 11–15 April 2016, pp. 263–273 [[DOI:10.1145/2872427.2883029](#)].
- [31] K. Bos et al., *LHC computing Grid: Technical Design Report. Version 1.06 (20 Jun 2005)*, CERN-LHCC-2005-024 (2005).
- [32] T.C. Donszelmann, W. Lampl and G.A. Stewart, *ART ATLAS Release Tester using the Grid*, *Eur. Phys. J. Web Conf.* **245** (2020) 05015.
- [33] ATLAS collaboration, *Trigger monitoring and rate predictions using Enhanced Bias data from the ATLAS Detector at the LHC*, ATL-DAQ-PUB-2016-002 (2016).
- [34] M. Borodin et al., *Scaling up ATLAS production system for the LHC Run 2 and beyond: project ProdSys2*, *J. Phys. Conf. Ser.* **664** (2015) 062005.
- [35] ATLAS collaboration, *ATLAS data quality operations and performance for 2015–2018 data-taking*, 2020 *JINST* **15** P04003 [[arXiv:1911.04632](#)].
- [36] ATLAS collaboration, *Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. Lett.* **121** (2018) 081801 [[arXiv:1804.03496](#)].
- [37] ATLAS collaboration, *Software Performance of the ATLAS Track Reconstruction for LHC Run 3*, ATL-PHYS-PUB-2021-012 (2021).
- [38] ATLAS collaboration, *Measurement of longitudinal flow decorrelations in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ and 5.02 TeV with the ATLAS detector*, *Eur. Phys. J. C* **78** (2018) 142 [[arXiv:1709.02301](#)].

- [39] ATLAS collaboration, *Measurements of azimuthal anisotropies of jet production in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector*, *Phys. Rev. C* **105** (2022) 064903 [[arXiv:2111.06606](#)].
- [40] ATLAS collaboration, *Measurement of the nuclear modification factor for inclusive jets in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector*, *Phys. Lett. B* **790** (2019) 108 [[arXiv:1805.05635](#)].
- [41] ATLAS collaboration, *The ATLAS inner detector trigger performance in pp collisions at 13 TeV during LHC Run 2*, *Eur. Phys. J. C* **82** (2022) 206 [[arXiv:2107.02485](#)].
- [42] ATLAS collaboration, *Performance of the reconstruction of large impact parameter tracks in the inner detector of ATLAS*, *Eur. Phys. J. C* **83** (2023) 1081 [[arXiv:2304.12867](#)].
- [43] ATLAS collaboration, *Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2*, *Eur. Phys. J. C* **77** (2017) 673 [[arXiv:1704.07983](#)].
- [44] ATLAS collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data*, *2019 JINST* **14** P12006 [[arXiv:1908.00005](#)].
- [45] ATLAS collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** (2021) 578 [[arXiv:2012.00578](#)].
- [46] ATLAS collaboration, *Jet energy scale and resolution measured in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **81** (2021) 689 [[arXiv:2007.02645](#)].
- [47] ATLAS collaboration, *E_T^{miss} performance in the ATLAS detector using 2015–2016 LHC p-p collisions*, *ATLAS-CONF-2018-023* (2018).
- [48] ATLAS collaboration, *ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset*, *Eur. Phys. J. C* **83** (2023) 681 [[arXiv:2211.16345](#)].
- [49] ATLAS collaboration, *Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung*, *ATLAS-CONF-2012-047* (2012).
- [50] ATLAS collaboration, *Jet reconstruction and performance using particle flow with the ATLAS Detector*, *Eur. Phys. J. C* **77** (2017) 466 [[arXiv:1703.10485](#)].
- [51] N. Lad, *Machine learning predictor for ‘measurement-to-track’ association for the ATLAS Inner Detector trigger*, *J. Phys. Conf. Ser.* **2438** (2023) 012101.
- [52] ATLAS collaboration, *The ATLAS Inner Detector tracking trigger at 13 TeV in LHC Run-2 and new developments on standard and unconventional tracking signatures for the upcoming LHC Run-3*, *ATL-DAQ-PROC-2022-005* (2022).
- [53] ATLAS collaboration, *Fast b-tagging at the high-level trigger of the ATLAS experiment in LHC Run 3*, *2023 JINST* **18** P11006 [[arXiv:2306.09738](#)].
- [54] ATLAS collaboration, *Kshort and Λ production in pp interactions at $\sqrt{s} = 0.9$ and 7 TeV measured with the ATLAS detector at the LHC*, *Phys. Rev. D* **85** (2012) 012001 [[arXiv:1111.1297](#)].
- [55] D. Oliveira Damazio, *ATLAS High Level Calorimeter Trigger Software Performance for Cosmic Ray Events*, *ATL-COM-DAQ-2009-045* (2009).
- [56] ATLAS collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, *Eur. Phys. J. C* **77** (2017) 490 [[arXiv:1603.02934](#)].
- [57] ATLAS collaboration, *Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton-proton collision data*, *2019 JINST* **14** P03017 [[arXiv:1812.03848](#)].
- [58] T. Barillari et al., *Local Hadronic Calibration*, *ATL-COM-LARG-2008-006* (2008).

- [59] J.M. Seixas, L. Caloba, M. Souza, A. Braga and A. Rodrigues, *Neural second level trigger system based on calorimetry*, *Comput. Phys. Commun.* **95** (1996) 143.
- [60] ATLAS collaboration, *Reconstruction, Identification, and Calibration of hadronically decaying tau leptons with the ATLAS detector for the LHC Run 3 and reprocessed Run 2 data*, [ATL-PHYS-PUB-2022-044](#) (2022).
- [61] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_r jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)].
- [62] P. Berta, L. Masetti, D.W. Miller and M. Spusta, *Pileup and Underlying Event Mitigation with Iterative Constituent Subtraction*, *JHEP* **08** (2019) 175 [[arXiv:1905.03470](#)].
- [63] M. Cacciari, G.P. Salam and G. Soyez, *SoftKiller, a particle-level pileup removal method*, *Eur. Phys. J. C* **75** (2015) 59 [[arXiv:1407.0408](#)].
- [64] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft Drop*, *JHEP* **05** (2014) 146 [[arXiv:1402.2657](#)].
- [65] ATLAS collaboration, *Configuration and performance of the ATLAS b-jet triggers in Run 2*, *Eur. Phys. J. C* **81** (2021) 1087 [[arXiv:2106.03584](#)].
- [66] ATLAS collaboration, *Deep Sets based Neural Networks for Impact Parameter Flavour Tagging in ATLAS*, [ATL-PHYS-PUB-2020-014](#) (2020).
- [67] ATLAS collaboration, *Graph Neural Network Jet Flavour Tagging with the ATLAS Detector*, [ATL-PHYS-PUB-2022-027](#) (2022).
- [68] ATLAS collaboration, *Search for resonant pair production of Higgs bosons in the $b\bar{b}b\bar{b}$ final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **105** (2022) 092002 [[arXiv:2202.07288](#)].
- [69] ATLAS collaboration, *Search for nonresonant pair production of Higgs bosons in the $b\bar{b}b\bar{b}$ final state in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **108** (2023) 052003 [[arXiv:2301.03212](#)].
- [70] ATLAS collaboration, *Measurements of Higgs bosons decaying to bottom quarks from vector boson fusion production with the ATLAS experiment at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** (2021) 537 [[arXiv:2011.08280](#)].
- [71] ATLAS collaboration, *Search for Higgs boson production in association with a high-energy photon via vector-boson fusion with decay into bottom quark pairs at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **03** (2021) 268 [[arXiv:2010.13651](#)].
- [72] ATLAS collaboration, *Search for the Standard Model Higgs boson decaying into $b\bar{b}$ produced in association with top quarks decaying hadronically in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *JHEP* **05** (2016) 160 [[arXiv:1604.03812](#)].
- [73] ATLAS collaboration, *Search for heavy neutral Higgs bosons produced in association with b-quarks and decaying into b-quarks at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **102** (2020) 032004 [[arXiv:1907.02749](#)].
- [74] ATLAS collaboration, *Search for single vector-like B quark production and decay via $B \rightarrow bH(b\bar{b})$ in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **11** (2023) 168 [[arXiv:2308.02595](#)].
- [75] ATLAS collaboration, *Search for resonances in the mass distribution of jet pairs with one or two jets identified as b-jets in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **98** (2018) 032016 [[arXiv:1805.09299](#)].
- [76] ATLAS collaboration, *Observation of $H \rightarrow b\bar{b}$ decays and VH production with the ATLAS detector*, *Phys. Lett. B* **786** (2018) 59 [[arXiv:1808.08238](#)].

- [77] ATLAS collaboration, *Search for invisible Higgs-boson decays in events with vector-boson fusion signatures using 139fb^{-1} of proton-proton data recorded by the ATLAS experiment*, *JHEP* **08** (2022) 104 [[arXiv:2202.07953](#)].
- [78] ATLAS collaboration, *Search for direct production of electroweakinos in final states with one lepton, missing transverse momentum and a Higgs boson decaying into two b -jets in pp collisions at $\sqrt{s} = 13\text{ TeV}$ with the ATLAS detector*, *Eur. Phys. J. C* **80** (2020) 691 [[arXiv:1909.09226](#)].
- [79] ATLAS collaboration, *Performance of the missing transverse momentum triggers for the ATLAS detector during Run-2 data taking*, *JHEP* **08** (2020) 080 [[arXiv:2005.09554](#)].
- [80] ATLAS collaboration, *Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8\text{ TeV}$ using the ATLAS detector*, *Eur. Phys. J. C* **76** (2016) 581 [[arXiv:1510.03823](#)].
- [81] ATLAS collaboration, *Constituent-level pile-up mitigation techniques in ATLAS*, *ATLAS-CONF-2017-065* (2017).
- [82] PARTICLE DATA Group, *Review of Particle Physics*, *Prog. Theor. Exp. Phys.* **2022** (2022) 083C01.
- [83] ATLAS collaboration, *Measurement of long-range multiparticle azimuthal correlations with the subevent cumulant method in pp and $p + Pb$ collisions with the ATLAS detector at the CERN Large Hadron Collider*, *Phys. Rev. C* **97** (2018) 024904 [[arXiv:1708.03559](#)].
- [84] ATLAS collaboration, *Charged-particle distributions in $\sqrt{s} = 13\text{ TeV}$ pp interactions measured with the ATLAS detector at the LHC*, *Phys. Lett. B* **758** (2016) 67 [[arXiv:1602.01633](#)].
- [85] M. Trzebiński, *Machine Optics Studies for the LHC Measurements*, *Proc. SPIE* **9290** (2014) 929026 [[arXiv:1408.1836](#)].
- [86] J. Lange et al., *Beam tests of an integrated prototype of the ATLAS Forward Proton detector*, *2016 JINST* **11** P09005 [[arXiv:1608.01485](#)].
- [87] ATLAS collaboration, *Search for heavy charged long-lived particles in the ATLAS detector in 36.1fb^{-1} of proton-proton collision data at $\sqrt{s} = 13\text{ TeV}$* , *Phys. Rev. D* **99** (2019) 092007 [[arXiv:1902.01636](#)].
- [88] ATLAS collaboration, *Search for heavy, long-lived, charged particles with large ionisation energy loss in pp collisions at $\sqrt{s} = 13\text{ TeV}$ using the ATLAS experiment and the full Run 2 dataset*, *JHEP* **06** (2023) 158 [[arXiv:2205.06013](#)].
- [89] ATLAS collaboration, *Search for long-lived charginos based on a disappearing-track signature in pp collisions at $\sqrt{s} = 13\text{ TeV}$ with the ATLAS detector*, *JHEP* **06** (2018) 022 [[arXiv:1712.02118](#)].
- [90] ATLAS collaboration, *Search for long-lived charginos based on a disappearing-track signature using 136fb^{-1} of pp collisions at $\sqrt{s} = 13\text{ TeV}$ with the ATLAS detector*, *Eur. Phys. J. C* **82** (2022) 606 [[arXiv:2201.02472](#)].
- [91] ATLAS collaboration, *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*, *ATL-PHYS-PUB-2019-013* (2019).
- [92] D. Linthorne and D. Stolarski, *Triggering on emerging jets*, *Phys. Rev. D* **104** (2021) 035019 [[arXiv:2103.08620](#)].
- [93] ATLAS collaboration, *Search for exotic decays of the Higgs boson into long-lived particles in pp collisions at $\sqrt{s} = 13\text{ TeV}$ using displaced vertices in the ATLAS inner detector*, *JHEP* **11** (2021) 229 [[arXiv:2107.06092](#)].
- [94] ATLAS collaboration, *Search for Displaced Leptons in $\sqrt{s} = 13\text{ TeV}$ pp Collisions with the ATLAS Detector*, *Phys. Rev. Lett.* **127** (2021) 051802 [[arXiv:2011.07812](#)].
- [95] ATLAS collaboration, *Identification of hadronic tau lepton decays using neural networks in the ATLAS experiment*, *ATL-PHYS-PUB-2019-033* (2019).

- [96] ATLAS collaboration, *Studies of the performance of the ATLAS detector using cosmic-ray muons*, *Eur. Phys. J. C* **71** (2011) 1593 [[arXiv:1011.6665](#)].
- [97] S. Fratina et al., *The TRT Fast-OR Trigger*, [ATL-INDET-PUB-2009-002](#) (2009).
- [98] ATLAS collaboration, *Beam-induced and cosmic-ray backgrounds observed in the ATLAS detector during the LHC 2012 proton-proton running period*, *2016 JINST* **11** P05013 [[arXiv:1603.09202](#)].
- [99] V. Cindro et al., *The ATLAS beam conditions monitor*, *2008 JINST* **3** P02004.
- [100] ATLAS collaboration, *ATLAS simulation using real data: Embedding and overlay*, *J. Phys. Conf. Ser.* **898** (2017) 042004.
- [101] ATLAS TILE CALORIMETER SYSTEM collaboration, *The Laser calibration of the ATLAS Tile Calorimeter during the LHC run 1*, *2016 JINST* **11** T10005 [[arXiv:1608.02791](#)].
- [102] ATLAS collaboration, *Monitoring and data quality assessment of the ATLAS liquid argon calorimeter*, *2014 JINST* **9** P07024 [[arXiv:1405.3768](#)].
- [103] ATLAS collaboration, *Online Determination of the LHC Luminous Region with the ATLAS High-Level Trigger*, *Phys. Procedia* **37** (2012) 2080.
- [104] F. Winklmeier, R. Bartoldus, J. Cogan, A. Salnikov and E. Strauss, *A System for Monitoring and Tracking the LHC Beam Spot within the ATLAS High Level Trigger*, in the proceedings of the *2012 18th IEEE-NPSS Real Time Conference*, Berkeley, CA, U.S.A., 9–15 June 2012, pp. 1–7 [[DOI:10.1109/RTC.2012.6418175](#)].
- [105] <https://op-webtools.web.cern.ch/vistar/vistars.php?usr=LHCCONFIG>.
- [106] A. Valassi, R. Bartoldus, A. Kalkhof, A. Salnikov and M. Wache, *CORAL server and CORAL server proxy: Scalable access to relational databases from CORAL applications*, *J. Phys. Conf. Ser.* **331** (2011) 042025.
- [107] M. Michelotto et al., *A comparison of HEP code with SPEC benchmarks on multi-core worker nodes*, *J. Phys. Conf. Ser.* **219** (2010) 052009.
- [108] ATLAS collaboration, *ATLAS Computing Acknowledgements*, [ATL-SOFT-PUB-2023-001](#) (2023).

The ATLAS collaboration

G. Aad ¹⁰², E. Aakvaag ¹⁶, B. Abbott ¹²⁰, K. Abeling ⁵⁵, N.J. Abicht ⁴⁹, S.H. Abidi ²⁹,
 A. Aboulhorma ^{35e}, H. Abramowicz ¹⁵¹, H. Abreu ¹⁵⁰, Y. Abulaiti ¹¹⁷, B.S. Acharya ^{69a,69b,m},
 C. Adam Bourdarios ⁴, L. Adamczyk ^{86a}, S.V. Addepalli ²⁶, M.J. Addison ¹⁰¹, J. Adelman ¹¹⁵,
 A. Adiguzel ^{21c}, T. Adye ¹³⁴, A.A. Affolder ¹³⁶, Y. Afik ³⁹, M.N. Agaras ¹³, J. Agarwala ^{73a,73b},
 A. Aggarwal ¹⁰⁰, C. Agheorghiesei ^{27c}, A. Ahmad ³⁶, F. Ahmadov ^{38,z}, W.S. Ahmed ¹⁰⁴, S. Ahuja ⁹⁵,
 X. Ai ^{62e}, G. Aielli ^{76a,76b}, A. Aikot ¹⁶³, M. Ait Tamlihat ^{35e}, B. Aitbenchikh ^{35a}, I. Aizenberg ¹⁶⁹,
 M. Akbiyik ¹⁰⁰, T.P.A. Åkesson ⁹⁸, A.V. Akimov ³⁷, D. Akiyama ¹⁶⁸, N.N. Akolkar ²⁴, S. Aktas ^{21a},
 K. Al Khoury ⁴¹, G.L. Alberghi ^{23b}, J. Albert ¹⁶⁵, P. Albicocco ⁵³, G.L. Albouy ⁶⁰,
 S. Alderweireldt ⁵², Z.L. Alegria ¹²¹, M. Aleksa ³⁶, I.N. Aleksandrov ³⁸, C. Alexa ^{27b},
 T. Alexopoulos ¹⁰, F. Alfonsi ^{23b}, M. Algren ⁵⁶, M. Alhroob ¹²⁰, B. Ali ¹³², H.M.J. Ali ⁹¹,
 S. Ali ¹⁴⁸, S.W. Alibocus ⁹², M. Aliev ^{33c}, G. Alimonti ^{71a}, W. Alkakh ⁵⁵, C. Allaire ⁶⁶,
 B.M.M. Allbrooke ¹⁴⁶, J.F. Allen ⁵², C.A. Allendes Flores ^{137f}, P.P. Allport ²⁰, A. Aloisio ^{72a,72b},
 F. Alonso ⁹⁰, C. Alpigiani ¹³⁸, M. Alvarez Estevez ⁹⁹, A. Alvarez Fernandez ¹⁰⁰, M. Alves Cardoso ⁵⁶,
 M.G. Alvigi ^{72a,72b}, M. Aly ¹⁰¹, Y. Amaral Coutinho ^{83b}, A. Ambler ¹⁰⁴, C. Amelung ³⁶,
 M. Amerl ¹⁰¹, C.G. Ames ¹⁰⁹, D. Amidei ¹⁰⁶, S.P. Amor Dos Santos ^{130a}, K.R. Amos ¹⁶³,
 V. Ananiev ¹²⁵, C. Anastopoulos ¹³⁹, T. Andeen ¹¹, J.K. Anders ³⁶, S.Y. Andreato ^{47a,47b},
 A. Andreazza ^{71a,71b}, S. Angelidakis ⁹, A. Angerami ^{41,ac}, A.V. Anisenkov ³⁷, A. Annovi ^{74a},
 C. Antel ⁵⁶, M.T. Anthony ¹³⁹, E. Antipov ¹⁴⁵, M. Antonelli ⁵³, F. Anulli ^{75a}, M. Aoki ⁸⁴,
 T. Aoki ¹⁵³, J.A. Aparisi Pozo ¹⁶³, M.A. Aparo ¹⁴⁶, L. Aperio Bella ⁴⁸, C. Appelt ¹⁸, A. Apyan ²⁶,
 S.J. Arbiol Val ⁸⁷, C. Arcangeletti ⁵³, A.T.H. Arce ⁵¹, E. Arena ⁹², J-F. Arguin ¹⁰⁸,
 S. Argyropoulos ⁵⁴, J.-H. Arling ⁴⁸, O. Arnaez ⁴, H. Arnold ¹¹⁴, G. Artoni ^{75a,75b}, H. Asada ¹¹¹,
 K. Asai ¹¹⁸, S. Asai ¹⁵³, N.A. Asbah ⁶¹, K. Assamagan ²⁹, R. Astalos ^{28a}, S. Atashi ¹⁵⁹,
 R.J. Atkin ^{33a}, M. Atkinson ¹⁶², H. Atmani ^{35f}, P.A. Atlasiddha ¹²⁸, K. Augsten ¹³²,
 S. Auricchio ^{72a,72b}, A.D. Aurio ²⁰, V.A. Austrup ¹⁰¹, G. Avolio ³⁶, K. Axiotis ⁵⁶, G. Azuelos ^{108,ag},
 D. Babal ^{28b}, H. Bachacou ¹³⁵, K. Bachas ^{152,q}, A. Bachi ³⁴, F. Backman ^{47a,47b}, A. Badea ³⁹,
 T.M. Baer ¹⁰⁶, P. Bagnaia ^{75a,75b}, M. Bahmani ¹⁸, D. Bahner ⁵⁴, A.J. Bailey ¹⁶³, V.R. Bailey ¹⁶²,
 J.T. Baines ¹³⁴, L. Baines ⁹⁴, O.K. Baker ¹⁷², E. Bakos ¹⁵, D. Bakshi Gupta ⁸, V. Balakrishnan ¹²⁰,
 R. Balasubramanian ¹¹⁴, E.M. Baldin ³⁷, P. Balek ^{86a}, E. Ballabene ^{23b,23a}, F. Balli ¹³⁵,
 L.M. Baltes ^{63a}, W.K. Balunas ³², J. Balz ¹⁰⁰, E. Banas ⁸⁷, M. Bandieramonte ¹²⁹,
 A. Bandyopadhyay ²⁴, S. Bansal ²⁴, L. Barak ¹⁵¹, M. Barakat ⁴⁸, E.L. Barberio ¹⁰⁵,
 D. Barberis ^{57b,57a}, M. Barbero ¹⁰², M.Z. Barel ¹¹⁴, K.N. Barends ^{33a}, T. Barillari ¹¹⁰,
 M.-S. Barisits ³⁶, T. Barklow ¹⁴³, P. Baron ¹²², D.A. Baron Moreno ¹⁰¹, A. Baroncelli ^{62a},
 G. Barone ²⁹, A.J. Barr ¹²⁶, J.D. Barr ⁹⁶, 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. Basan ¹⁰⁰, M. Baselga ⁴⁹, A. Bassalat ^{66,b}, M.J. Basso ^{156a},
 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}, M. Bauer ³⁶, P. Bauer ²⁴, L.T. Bazzano Hurrell ³⁰, J.B. Beacham ⁵¹,
 T. Beau ¹²⁷, J.Y. Beaucamp ⁹⁰, P.H. Beauchemin ¹⁵⁸, P. Bechtel ²⁴, H.P. Beck ^{19,p}, K. Becker ¹⁶⁷,
 A.J. Beddall ⁸², V.A. Bednyakov ³⁸, C.P. Bee ¹⁴⁵, L.J. Beamster ¹⁵, T.A. Beeraman ³⁶,
 M. Begalli ^{83d}, M. Begel ²⁹, A. Behara ¹⁴⁵, J.K. Behr ⁴⁸, J.F. Beirer ³⁶, F. Beisiegel ²⁴,
 M. Belfkir ^{116b}, G. Bella ¹⁵¹, L. Bellagamba ^{23b}, A. Bellerive ³⁴, P. Bellos ²⁰, K. Beloborodov ³⁷,
 D. Benckekroun ^{35a}, F. Bendebba ^{35a}, Y. Benhammou ¹⁵¹, K.C. Benkendorfer ⁶¹, L. Beresford ⁴⁸,

M. Beretta ⁵³, E. Bergeaas Kuutmann ¹⁶¹, N. Berger ⁴, B. Bergmann ¹³², J. Beringer ^{17a},
 G. Bernardi ⁵, C. Bernius ¹⁴³, F.U. Bernlochner ²⁴, F. Bernon ^{36,102}, A. Berrocal Guardia ¹³,
 T. Berry ⁹⁵, P. Berta ¹³³, A. Berthold ⁵⁰, I.A. Bertram ⁹¹, S. Bethke ¹¹⁰, A. Betti ^{75a,75b},
 A.J. Bevan ⁹⁴, N.K. Bhalla ⁵⁴, M. Bhamjee ^{33c}, S. Bhatta ¹⁴⁵, D.S. Bhattacharya ¹⁶⁶, P. Bhattarai ¹⁴³,
 K.D. Bhide ⁵⁴, V.S. Bhopatkar ¹²¹, R.M. Bianchi ¹²⁹, G. Bianco ^{23b,23a}, O. Biebel ¹⁰⁹, R. Bielski ¹²³,
 M. Biglietti ^{77a}, C.S. Billingsley ⁴⁴, M. Bindi ⁵⁵, A. Bingul ^{21b}, C. Bini ^{75a,75b}, A. Biondini ⁹²,
 C.J. Birch-sykes ¹⁰¹, G.A. Bird ³², M. Birman ¹⁶⁹, M. Biros ¹³³, S. Biryukov ¹⁴⁶, T. Bisanz ⁴⁹,
 E. Bisceglie ^{43b,43a}, J.P. Biswal ¹³⁴, D. Biswas ¹⁴¹, K. Bjørke ¹²⁵, I. Bloch ⁴⁸, A. Blue ⁵⁹,
 U. Blumenschein ⁹⁴, J. Blumenthal ¹⁰⁰, V.S. Bobrovnikov ³⁷, M. Boehler ⁵⁴, B. Boehm ¹⁶⁶,
 D. Bogavac ³⁶, A.G. Bogdanchikov ³⁷, C. Bohm ^{47a}, V. Boisvert ⁹⁵, P. Bokan ³⁶, T. Bold ^{86a},
 M. Bomben ⁵, M. Bona ⁹⁴, M. Boonekamp ¹³⁵, C.D. Booth ⁹⁵, A.G. Borbély ⁵⁹, I.S. Bordulev ³⁷,
 H.M. Borecka-Bielska ¹⁰⁸, G. Borissov ⁹¹, D. Bortoletto ¹²⁶, D. Boscherini ^{23b}, M. Bosman ¹³,
 J.D. Bossio Sola ³⁶, K. Bouaouda ^{35a}, N. Bouchhar ¹⁶³, J. Boudreau ¹²⁹, E.V. Bouhova-Thacker ⁹¹,
 D. Boumediene ⁴⁰, R. Bouquet ¹⁶⁵, A. Boveia ¹¹⁹, J. Boyd ³⁶, D. Boye ²⁹, I.R. Boyko ³⁸,
 J. Bracinik ²⁰, N. Brahimy ^{62d}, E.D. Brandani ¹²⁹, G. Brandt ¹⁷¹, O. Brandt ³², F. Braren ⁴⁸,
 B. Brau ¹⁰³, J.E. Brau ¹²³, R. Brenner ¹⁶⁹, L. Brenner ¹¹⁴, R. Brenner ¹⁶¹, S. Bressler ¹⁶⁹,
 D. Britton ⁵⁹, D. Britzger ¹¹⁰, I. Brock ²⁴, G. Brooijmans ⁴¹, E. Brost ²⁹, L.M. Brown ¹⁶⁵,
 L.E. Bruce ⁶¹, T.L. Bruckler ¹²⁶, P.A. Bruckman de Renstrom ⁸⁷, B. Brüers ⁴⁸, A. Bruni ^{23b},
 G. Bruni ^{23b}, M. Bruschi ^{23b}, N. Bruscino ^{75a,75b}, T. Buanes ¹⁶, Q. Buat ¹³⁸, D. Buchin ¹¹⁰,
 A.G. Buckley ⁵⁹, O. Bulekov ³⁷, B.A. Bullard ¹⁴³, S. Burdin ⁹², C.D. Burgard ⁴⁹, A.M. Burger ³⁶,
 B. Burghgrave ⁸, O. Burlayenko ⁵⁴, 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, S. Cabrera Urbán ¹⁶³, L. Cadamuro ⁶⁶, D. Caforio ⁵⁸,
 H. Cai ¹²⁹, Y. Cai ^{14a,14e}, Y. Cai ^{14c}, V.M.M. Cairo ³⁶, O. Cakir ^{3a}, N. Calace ³⁶, P. Calafiura ^{17a},
 G. Calderini ¹²⁷, P. Calfayan ⁶⁸, G. Callea ⁵⁹, L.P. Caloba ^{83b}, D. Calvet ⁴⁰, S. Calvet ⁴⁰,
 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 ³⁶, C. Camincher ¹⁶⁵, M. Campanelli ⁹⁶, A. Camplani ⁴²,
 V. Canale ^{72a,72b}, J. Cantero ¹⁶³, Y. Cao ¹⁶², F. Capocasa ²⁶, M. Capua ^{43b,43a}, A. Carbone ^{71a,71b},
 R. Cardarelli ^{76a}, J.C.J. Cardenas ⁸, F. Cardillo ¹⁶³, G. Carducci ^{43b,43a}, T. Carli ³⁶, G. Carlino ^{72a},
 J.I. Carlotto ¹³, B.T. Carlson ^{129,r}, E.M. Carlson ^{165,156a}, L. Carminati ^{71a,71b}, A. Carnelli ¹³⁵,
 M. Carnesale ^{75a,75b}, S. Caron ¹¹³, E. Carquin ^{137f}, S. Carrá ^{71a}, G. Carratta ^{23b,23a},
 A.M. Carroll ¹²³, T.M. Carter ⁵², M.P. Casado ^{13,i}, M. Caspar ⁴⁸, F.L. Castillo ⁴,
 L. Castillo Garcia ¹³, V. Castillo Gimenez ¹⁶³, N.F. Castro ^{130a,130e}, A. Catinaccio ³⁶, J.R. Catmore ¹²⁵,
 T. Cavaliere ⁴, V. Cavaliere ²⁹, N. Cavalli ^{23b,23a}, Y.C. Cekmecelioglu ⁴⁸, E. Celebi ^{21a}, S. Cella ³⁶,
 F. Celli ¹²⁶, M.S. Centonze ^{70a,70b}, V. Cepaitis ⁵⁶, K. Cerny ¹²², A.S. Cerqueira ^{83a}, A. Cerri ¹⁴⁶,
 L. Cerrito ^{76a,76b}, F. Cerutti ^{17a}, B. Cervato ¹⁴¹, A. Cervelli ^{23b}, G. Cesarini ⁵³, S.A. Cetin ⁸²,
 D. Chakraborty ¹¹⁵, J. Chan ^{17a}, W.Y. Chan ¹⁵³, J.D. Chapman ³², E. Chapon ¹³⁵,
 B. Chargeishvili ^{149b}, D.G. Charlton ²⁰, M. Chatterjee ¹⁹, C. Chauhan ¹³³, Y. Che ^{14c}, S. Chekanov ⁶,
 S.V. Chekulaev ^{156a}, G.A. Chelkov ^{38,a}, A. Chen ¹⁰⁶, B. Chen ¹⁵¹, B. Chen ¹⁶⁵, H. Chen ^{14c},
 H. Chen ²⁹, J. Chen ^{62c}, J. Chen ¹⁴², M. Chen ¹²⁶, S. Chen ¹⁵³, S.J. Chen ^{14c}, X. Chen ^{62c,135},
 X. Chen ^{14b,af}, Y. Chen ^{62a}, C.L. Cheng ¹⁷⁰, 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, V. Chiarella ⁵³, G. Chiarelli ^{74a}, N. Chiedde ¹⁰², G. Chiodini ^{70a}, A.S. Chisholm ²⁰,

A. Chitan ^{27b}, M. Chitishvili ¹⁶³, M.V. Chizhov ³⁸, K. Choi ¹¹, Y. Chou ¹³⁸, E.Y.S. Chow ¹¹³,
 K.L. Chu ¹⁶⁹, M.C. Chu ^{64a}, X. Chu ^{14a,14e}, J. Chudoba ¹³¹, J.J. Chwastowski ⁸⁷, D. Cierci ¹¹⁰,
 K.M. Ciesla ^{86a}, V. Cindro ⁹³, A. Ciocio ^{17a}, F. Cirotto ^{72a,72b}, Z.H. Citron ^{169,k}, M. Citterio ^{71a},
 D.A. Ciubotaru ^{27b}, A. Clark ⁵⁶, P.J. Clark ⁵², C. Clarry ¹⁵⁵, J.M. Clavijo Columbie ⁴⁸, S.E. Clawson ⁴⁸,
 C. Clement ^{47a,47b}, J. Clercx ⁴⁸, Y. Coadou ¹⁰², M. Cobal ^{69a,69c}, A. Coccaro ^{57b},
 R.F. Coelho Barrue ^{130a}, R. Coelho Lopes De Sa ¹⁰³, S. Coelli ^{71a}, B. Cole ⁴¹, J. Collot ⁶⁰,
 P. Conde Muiño ^{130a,130g}, M.P. Connell ^{33c}, S.H. Connell ^{33c}, I.A. Connelly ⁵⁹, E.I. Conroy ¹²⁶,
 F. Conventi ^{72a,ah}, H.G. Cooke ²⁰, A.M. Cooper-Sarkar ¹²⁶, A. Cordeiro Oudot Choi ¹²⁷,
 L.D. Corpe ⁴⁰, M. Corradi ^{75a,75b}, F. Corriveau ^{104,x}, A. Cortes-Gonzalez ¹⁸, M.J. Costa ¹⁶³,
 F. Costanza ⁴, D. Costanzo ¹³⁹, B.M. Cote ¹¹⁹, G. Cowan ⁹⁵, K. Cranmer ¹⁷⁰, D. Cremonini ^{23b,23a},
 S. Crépe-Renaudin ⁶⁰, F. Crescioli ¹²⁷, M. Cristinziani ¹⁴¹, M. Cristoforetti ^{78a,78b}, V. Croft ¹¹⁴,
 J.E. Crosby ¹²¹, G. Crosetti ^{43b,43a}, A. Cueto ⁹⁹, T. Cuhadar Donszelmann ¹⁵⁹, H. Cui ^{14a,14e},
 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 ^{57b,57a}, J.V. Da Fonseca Pinto ^{83b}, C. Da Via ¹⁰¹, W. Dabrowski ^{86a},
 T. Dado ⁴⁹, S. Dahbi ^{33g}, T. Dai ¹⁰⁶, D. Dal Santo ¹⁹, C. Dallapiccola ¹⁰³, M. Dam ⁴², G. D'amen ²⁹,
 V. D'Amico ¹⁰⁹, J. Damp ¹⁰⁰, J.R. Dandoy ³⁴, M. Danninger ¹⁴², V. Dao ³⁶, G. Darbo ^{57b},
 S. Darmora ⁶, S.J. Das ^{29,ai}, S. D'Auria ^{71a,71b}, C. David ^{33a}, T. Davidek ¹³³, B. Davis-Purcell ³⁴,
 I. Dawson ⁹⁴, H.A. Day-hall ¹³², K. De ⁸, R. De Asmundis ^{72a}, N. De Biase ⁴⁸, S. De Castro ^{23b,23a},
 N. De Groot ¹¹³, P. de Jong ¹¹⁴, H. De la Torre ¹¹⁵, A. De Maria ^{14c}, A. De Salvo ^{75a},
 U. De Sanctis ^{76a,76b}, F. De Santis ^{70a,70b}, A. De Santo ¹⁴⁶, J.B. De Vivie De Regie ⁶⁰, D.V. Dedovich ³⁸,
 J. Degens ¹¹⁴, A.M. Deiana ⁴⁴, F. Del Corso ^{23b,23a}, J. Del Peso ⁹⁹, F. Del Rio ^{63a}, L. Delagrangé ¹²⁷,
 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 ¹⁷², M. Demichev ³⁸,
 S.P. Denisov ³⁷, L. D'Eramo ⁴⁰, D. Derendarz ⁸⁷, F. Derue ¹²⁷, P. Dervan ⁹², K. Desch ²⁴,
 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 ³⁶, A. Di Luca ^{78a,78b}, B. Di Micco ^{77a,77b},
 R. Di Nardo ^{77a,77b}, M. Diamantopoulou ³⁴, F.A. Dias ¹¹⁴, T. Dias Do Vale ¹⁴², M.A. Diaz ^{137a,137b},
 F.G. Diaz Capriles ²⁴, M. Didenko ¹⁶³, E.B. Diehl ¹⁰⁶, L. Diehl ⁵⁴, S. Díez Cornell ⁴⁸,
 C. Díez Pardos ¹⁴¹, C. Dimitriadi ^{161,24}, A. Dimitrievska ^{17a}, J. Dingfelder ²⁴, I-M. Dinu ^{27b},
 S.J. Dittmeier ^{63b}, F. Dittus ³⁶, F. Djama ¹⁰², T. Djobava ^{149b}, C. Doglioni ^{101,98}, A. Dohnalova ^{28a},
 J. Dolejsi ¹³³, Z. Dolezal ¹³³, K.M. Dona ³⁹, M. Donadelli ^{83c}, B. Dong ¹⁰⁷, J. Donini ⁴⁰,
 A. D'Onofrio ^{72a,72b}, M. D'Onofrio ⁹², J. Dopke ¹³⁴, A. Doria ^{72a}, N. Dos Santos Fernandes ^{130a},
 P. Dougan ¹⁰¹, M.T. Dova ⁹⁰, A.T. Doyle ⁵⁹, M.A. Draguet ¹²⁶, E. Dreyer ¹⁶⁹, I. Drivas-koulouris ¹⁰,
 M. Drnevich ¹¹⁷, M. Drozdova ⁵⁶, D. Du ^{62a}, T.A. du Pree ¹¹⁴, F. Dubinin ³⁷, M. Dubovsky ^{28a},
 E. Duchovni ¹⁶⁹, G. Duckeck ¹⁰⁹, O.A. Ducu ^{27b}, D. Duda ⁵², A. Dudarev ³⁶, E.R. Duden ²⁶,
 M. D'uffizi ¹⁰¹, L. Duflot ⁶⁶, M. Dührssen ³⁶, A.E. Dumitriu ^{27b}, M. Dunford ^{63a}, S. Dungs ⁴⁹,
 K. Dunne ^{47a,47b}, A. Duperrin ¹⁰², H. Duran Yildiz ^{3a}, M. Düren ⁵⁸, A. Durglishvili ^{149b},
 B.L. Dwyer ¹¹⁵, G.I. Dyckes ^{17a}, M. Dyndal ^{86a}, B.S. Dziedzic ⁸⁷, Z.O. Earnshaw ¹⁴⁶,
 G.H. Eberwein ¹²⁶, B. Eckerova ^{28a}, S. Eggebrecht ⁵⁵, E. Egidio Purcino De Souza ¹²⁷, L.F. Ehrke ⁵⁶,
 G. Eigen ¹⁶, K. Einsweiler ^{17a}, T. Ekelof ¹⁶¹, P.A. Ekman ⁹⁸, S. El Farkh ^{35b}, Y. El Ghazali ^{35b},
 H. El Jarrari ³⁶, A. El Moussaouy ¹⁰⁸, V. Ellajosyula ¹⁶¹, M. Ellert ¹⁶¹, F. Ellinghaus ¹⁷¹, N. Ellis ³⁶,
 J. Elmsheuser ²⁹, M. Elsing ³⁶, D. Emeliyanov ¹³⁴, Y. Enari ¹⁵³, I. Ene ^{17a}, S. Epari ¹³,
 P.A. Erland ⁸⁷, M. Errenst ¹⁷¹, M. Escalier ⁶⁶, C. Escobar ¹⁶³, E. Etzion ¹⁵¹, G. Evans ^{130a},

H. Evans ⁶⁸, L.S. Evans ⁹⁵, M.O. Evans ¹⁴⁶, A. Ezhilov ³⁷, S. Ezzarqtouni ^{35a}, F. Fabbri ⁵⁹, L. Fabbri ^{23b,23a}, G. Facini ⁹⁶, V. Fadeyev ¹³⁶, R.M. Fakhruddinov ³⁷, D. Fakoudis ¹⁰⁰, S. Falciano ^{75a}, L.F. Falda Ulhoa Coelho ³⁶, P.J. Falke ²⁴, J. Faltova ¹³³, C. Fan ¹⁶², Y. Fan ^{14a}, Y. Fang ^{14a,14e}, M. Fanti ^{71a,71b}, M. Faraj ^{69a,69b}, Z. Farazpay ⁹⁷, A. Farbin ⁸, A. Farilla ^{77a}, T. Farooque ¹⁰⁷, S.M. Farrington ⁵², F. Fassi ^{35e}, D. Fassouliotis ⁹, M. Faucci Giannelli ^{76a,76b}, W.J. Fawcett ³², L. Fayard ⁶⁶, P. Federic ¹³³, P. Federicova ¹³¹, O.L. Fedin ^{37,a}, M. Feickert ¹⁷⁰, L. Feligioni ¹⁰², D.E. Fellers ¹²³, C. Feng ^{62b}, M. Feng ^{14b}, Z. Feng ¹¹⁴, M.J. Fenton ¹⁵⁹, L. Ferencz ⁴⁸, R.A.M. Ferguson ⁹¹, S.I. Fernandez Luengo ^{137f}, P. Fernandez Martinez ¹³, M.J.V. Fernoux ¹⁰², J. Ferrando ⁹¹, A. Ferrari ¹⁶¹, P. Ferrari ^{114,113}, R. Ferrari ^{73a}, D. Ferrere ⁵⁶, C. Ferretti ¹⁰⁶, F. Fiedler ¹⁰⁰, P. Fiedler ¹³², A. Filipčič ⁹³, E.K. Filmer ¹, F. Filthaut ¹¹³, M.C.N. Fiolhais ^{130a,130c,c}, L. Fiorini ¹⁶³, W.C. Fisher ¹⁰⁷, T. Fitschen ¹⁰¹, P.M. Fitzhugh ¹³⁵, I. Fleck ¹⁴¹, P. Fleischmann ¹⁰⁶, T. Flick ¹⁷¹, M. Flores ^{33d,ad}, L.R. Flores Castillo ^{64a}, L. Flores Sanz De Acedo ³⁶, F.M. Follega ^{78a,78b}, N. Fomin ¹⁶, J.H. Foo ¹⁵⁵, A. Formica ¹³⁵, A.C. Forti ¹⁰¹, E. Fortin ³⁶, A.W. Fortman ^{17a}, M.G. Foti ^{17a}, L. Fountas ^{9,j}, D. Fournier ⁶⁶, H. Fox ⁹¹, P. Francavilla ^{74a,74b}, S. Francescato ⁶¹, S. Franchellucci ⁵⁶, M. Franchini ^{23b,23a}, S. Franchino ^{63a}, D. Francis ³⁶, L. Franco ¹¹³, V. Franco Lima ³⁶, L. Franconi ⁴⁸, M. Franklin ⁶¹, G. Frattari ²⁶, A.C. Freegard ⁹⁴, W.S. Freund ^{83b}, Y.Y. Frid ¹⁵¹, J. Friend ⁵⁹, N. Fritzsche ⁵⁰, A. Froch ⁵⁴, D. Froidevaux ³⁶, J.A. Frost ¹²⁶, Y. Fu ^{62a}, S. Fuenzalida Garrido ^{137f}, M. Fujimoto ¹⁰², K.Y. Fung ^{64a}, E. Furtado De Simas Filho ^{83b}, M. Furukawa ¹⁵³, J. Fuster ¹⁶³, A. Gabrielli ^{23b,23a}, A. Gabrielli ¹⁵⁵, P. Gadow ³⁶, G. Gagliardi ^{57b,57a}, L.G. Gagnon ^{17a}, S. Galantzan ¹⁵¹, E.J. Gallas ¹²⁶, B.J. Gallop ¹³⁴, K.K. Gan ¹¹⁹, S. Ganguly ¹⁵³, Y. Gao ⁵², F.M. Garay Walls ^{137a,137b}, B. Garcia ²⁹, C. García ¹⁶³, A. Garcia Alonso ¹¹⁴, A.G. Garcia Caffaro ¹⁷², J.E. García Navarro ¹⁶³, M. Garcia-Sciveres ^{17a}, G.L. Gardner ¹²⁸, R.W. Gardner ³⁹, N. Garelli ¹⁵⁸, D. Garg ⁸⁰, R.B. Garg ^{143,n}, J.M. Gargan ⁵², C.A. Garner ¹⁵⁵, C.M. Garvey ^{33a}, P. Gaspar ^{83b}, V.K. Gassmann ¹⁵⁸, G. Gaudio ^{73a}, V. Gautam ¹³, P. Gauzzi ^{75a,75b}, I.L. Gavrilenko ³⁷, A. Gavrilyuk ³⁷, C. Gay ¹⁶⁴, G. Gaycken ⁴⁸, E.N. Gazis ¹⁰, A.A. Geanta ^{27b}, C.M. Gee ¹³⁶, A. Gekow ¹¹⁹, C. Gemme ^{57b}, M.H. Genest ⁶⁰, A.D. Gentry ¹¹², S. George ⁹⁵, W.F. George ²⁰, T. Gerialis ⁴⁶, P. Gessinger-Befurt ³⁶, M.E. Geyik ¹⁷¹, M. Ghani ¹⁶⁷, M. Ghneimat ¹⁴¹, K. Ghorbanian ⁹⁴, A. Ghosal ¹⁴¹, A. Ghosh ¹⁵⁹, A. Ghosh ⁷, B. Giacobbe ^{23b}, S. Giagu ^{75a,75b}, T. Giani ¹¹⁴, P. Giannetti ^{74a}, A. Giannini ^{62a}, S.M. Gibson ⁹⁵, M. Gignac ¹³⁶, D.T. Gil ^{86b}, A.K. Gilbert ^{86a}, B.J. Gilbert ⁴¹, D. Gillberg ³⁴, G. Gilles ¹¹⁴, L. Ginabat ¹²⁷, D.M. Gingrich ^{2,ag}, M.P. Giordani ^{69a,69c}, P.F. Giraud ¹³⁵, G. Giugliarelli ^{69a,69c}, D. Giugni ^{71a}, F. Giuli ³⁶, I. Gkialas ^{9,j}, L.K. Gladilin ³⁷, C. Glasman ⁹⁹, G.R. Gledhill ¹²³, G. Glemža ⁴⁸, M. Glisic ¹²³, I. Gnesi ^{43b,f}, Y. Go ²⁹, M. Goblirsch-Kolb ³⁶, B. Gocke ⁴⁹, D. Godin ¹⁰⁸, B. Gokturk ^{21a}, S. Goldfarb ¹⁰⁵, T. Golling ⁵⁶, M.G.D. Gololo ^{33g}, D. Golubkov ³⁷, J.P. Gombas ¹⁰⁷, A. Gomes ^{130a,130b}, G. Gomes Da Silva ¹⁴¹, A.J. Gomez Delegido ¹⁶³, R. Gonçalo ^{130a,130c}, L. Gonella ²⁰, A. Gongadze ^{149c}, F. Gonnella ²⁰, J.L. Gonski ⁴¹, R.Y. González Andana ⁵², S. González de la Hoz ¹⁶³, R. Gonzalez Lopez ⁹², C. Gonzalez Renteria ^{17a}, M.V. Gonzalez Rodrigues ⁴⁸, R. Gonzalez Suarez ¹⁶¹, S. Gonzalez-Sevilla ⁵⁶, G.R. Gonzalvo Rodriguez ¹⁶³, L. Goossens ³⁶, B. Gorini ³⁶, E. Gorini ^{70a,70b}, A. Gorišek ⁹³, T.C. Gosart ¹²⁸, A.T. Goshaw ⁵¹, 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}, I. Grabowska-Bold ^{86a}, K. Graham ³⁴, E. Gramstad ¹²⁵, S. Grancagnolo ^{70a,70b}, C.M. Grant ^{1,135}, P.M. Gravila ^{27f}, F.G. Gravili ^{70a,70b}, H.M. Gray ^{17a}, M. Greco ^{70a,70b}, C. Grefe ²⁴, I.M. Gregor ⁴⁸, P. Grenier ¹⁴³,

S.G. Grewe¹¹⁰, C. Grieco¹³, A.A. Grillo¹³⁶, K. Grimm³¹, S. Grinstein^{13,t}, J.-F. Grivaz⁶⁶,
 E. Gross¹⁶⁹, J. Grosse-Knetter⁵⁵, J.C. Grundy¹²⁶, L. Guan¹⁰⁶, W. Guan²⁹, C. Gubbels¹⁶⁴,
 J.G.R. Guerrero Rojas¹⁶³, G. Guerrieri^{69a,69c}, F. Guescini¹¹⁰, D. Guest¹⁸, R. Gugel¹⁰⁰,
 J.A.M. Guhit¹⁰⁶, A. Guida¹⁸, E. Guilloton^{167,134}, S. Guindon³⁶, F. Guo^{14a,14e}, J. Guo^{62c},
 L. Guo⁴⁸, Y. Guo¹⁰⁶, R. Gupta⁴⁸, R. Gupta¹²⁹, S. Gurbuz²⁴, S.S. Gurdasani⁵⁴, G. Gustavino³⁶,
 M. Guth⁵⁶, P. Gutierrez¹²⁰, L.F. Gutierrez Zagazeta¹²⁸, M. Gutsche⁵⁰, C. Gutschow⁹⁶,
 C. Gwenlan¹²⁶, C.B. Gwilliam⁹², E.S. Haaland¹²⁵, A. Haas¹¹⁷, M. Habedank⁴⁸, C. Haber^{17a},
 H.K. Hadavand⁸, A. Hadeef⁵⁰, S. Hadzic¹¹⁰, A.I. Hagan⁹¹, J.J. Hahn¹⁴¹, E.H. Haines⁹⁶,
 M. Haleem¹⁶⁶, J. Haley¹²¹, J.J. Hall¹³⁹, G.D. Hallewell¹⁰², L. Halser¹⁹, K. Hamano¹⁶⁵,
 M. Hamer²⁴, G.N. Hamity⁵², E.J. Hampshire⁹⁵, J. Han^{62b}, K. Han^{62a}, L. Han^{14c}, L. Han^{62a},
 S. Han^{17a}, Y.F. Han¹⁵⁵, K. Hanagaki⁸⁴, M. Hance¹³⁶, D.A. Hangal⁴¹, H. Hanif¹⁴²,
 M.D. Hank¹²⁸, J.B. Hansen⁴², P.H. Hansen⁴², K. Hara¹⁵⁷, D. Harada⁵⁶, T. Harenberg¹⁷¹,
 S. Harkusha³⁷, M.L. Harris¹⁰³, Y.T. Harris¹²⁶, J. Harrison¹³, N.M. Harrison¹¹⁹, P.F. Harrison¹⁶⁷,
 N.M. Hartman¹¹⁰, N.M. Hartmann¹⁰⁹, Y. Hasegawa¹⁴⁰, R. Hauser¹⁰⁷, C.M. Hawkes²⁰,
 R.J. Hawkins³⁶, Y. Hayashi¹⁵³, S. Hayashida¹¹¹, D. Hayden¹⁰⁷, C. Hayes¹⁰⁶, R.L. Hayes¹¹⁴,
 C.P. Hays¹²⁶, J.M. Hays⁹⁴, H.S. Hayward⁹², F. He^{62a}, M. He^{14a,14e}, Y. He¹⁵⁴, Y. He⁴⁸,
 Y. He⁹⁶, N.B. Heatley⁹⁴, V. Hedberg⁹⁸, A.L. Heggelund¹²⁵, N.D. Hehir^{94,*}, C. Heidegger⁵⁴,
 K.K. Heidegger⁵⁴, W.D. Heidorn⁸¹, J. Heilman³⁴, S. Heim⁴⁸, T. Heim^{17a}, J.G. Heinlein¹²⁸,
 J.J. Heinrich¹²³, L. Heinrich^{110,ae}, J. Hejbal¹³¹, A. Held¹⁷⁰, S. Hellesund¹⁶, C.M. Helling¹⁶⁴,
 S. Hellman^{47a,47b}, R.C.W. Henderson⁹¹, L. Henkelmann³², A.M. Henriques Correia³⁶, H. Herde⁹⁸,
 Y. Hernández Jiménez¹⁴⁵, L.M. Herrmann²⁴, T. Herrmann⁵⁰, G. Herten⁵⁴, R. Hertenberger¹⁰⁹,
 L. Hervas³⁶, M.E. Hesping¹⁰⁰, N.P. Hessey^{156a}, E. Hill¹⁵⁵, S.J. Hillier²⁰, J.R. Hinds¹⁰⁷,
 F. Hinterkeuser²⁴, M. Hirose¹²⁴, S. Hirose¹⁵⁷, D. Hirschbuehl¹⁷¹, T.G. Hitchings¹⁰¹, B. Hiti⁹³,
 J. Hobbs¹⁴⁵, R. Hobincu^{27e}, N. Hod¹⁶⁹, M.C. Hodgkinson¹³⁹, B.H. Hodgkinson³², A. Hoecker³⁶,
 D.D. Hofer¹⁰⁶, J. Hofer⁴⁸, T. Holm²⁴, M. Holzbock¹¹⁰, L.B.A.H. Hommels³², B.P. Honan¹⁰¹,
 J. Hong^{62c}, T.M. Hong¹²⁹, B.H. Hooberman¹⁶², W.H. Hopkins⁶, Y. Horii¹¹¹, S. Hou¹⁴⁸,
 A.S. Howard⁹³, J. Howarth⁵⁹, J. Hoya⁶, M. Hrabovsky¹²², A. Hrynevich⁴⁸, T. Hryn'ova⁴,
 P.J. Hsu⁶⁵, S.-C. Hsu¹³⁸, Q. Hu^{62a}, Y.F. Hu^{14a,14e}, S. Huang^{64b}, X. Huang^{14c},
 X. Huang^{14a,14e}, Y. Huang¹³⁹, Y. Huang^{14a}, Z. Huang¹⁰¹, Z. Hubacek¹³², M. Huebner²⁴,
 F. Huegging²⁴, T.B. Huffman¹²⁶, C.A. Hugli⁴⁸, M. Huhtinen³⁶, S.K. Huiberts¹⁶, R. Hulskens¹⁰⁴,
 N. Huseynov¹², J. Huston¹⁰⁷, J. Huth⁶¹, R. Hyneman¹⁴³, G. Iacobucci⁵⁶, G. Iakovidis²⁹,
 I. Ibragimov¹⁴¹, L. Iconomidou-Fayard⁶⁶, J.P. Iddon³⁶, P. Iengo^{72a,72b}, R. Iguchi¹⁵³, T. Iizawa¹²⁶,
 Y. Ikegami⁸⁴, N. Ilic¹⁵⁵, H. Imam^{35a}, M. Ince Lezki⁵⁶, T. Ingebretsen Carlson^{47a,47b},
 G. Introzzi^{73a,73b}, M. Iodice^{77a}, V. Ippolito^{75a,75b}, R.K. Irwin⁹², M. Ishino¹⁵³, W. Islam¹⁷⁰,
 C. Issever^{18,48}, S. Istin^{21a,ak}, H. Ito¹⁶⁸, R. Iuppa^{78a,78b}, A. Ivina¹⁶⁹, J.M. Izen⁴⁵, V. Izzo^{72a},
 P. Jacka^{131,132}, P. Jackson¹, B.P. Jaeger¹⁴², C.S. Jagfeld¹⁰⁹, G. Jain^{156a}, P. Jain⁵⁴, K. Jakobs⁵⁴,
 T. Jakoubek¹⁶⁹, J. Jamieson⁵⁹, K.W. Janas^{86a}, M. Javurkova¹⁰³, L. Jeanty¹²³, J. Jejelava^{149a,aa},
 P. Jenni^{54,g}, C.E. Jessiman³⁴, C. Jia^{62b}, J. Jia¹⁴⁵, X. Jia⁶¹, X. Jia^{14a,14e}, Z. Jia^{14c}, S. Jiggins⁴⁸,
 J. Jimenez Pena¹³, S. Jin^{14c}, A. Jinaru^{27b}, O. Jinnouchi¹⁵⁴, P. Johansson¹³⁹, K.A. Johns⁷,
 J.W. Johnson¹³⁶, D.M. Jones³², E. Jones⁴⁸, P. Jones³², R.W.L. Jones⁹¹, T.J. Jones⁹²,
 H.L. Joos^{55,36}, R. Joshi¹¹⁹, J. Jovicevic¹⁵, X. Ju^{17a}, J.J. Junggeburth¹⁰³, T. Junkermann^{63a},
 A. Juste Rozas^{13,t}, M.K. Juzek⁸⁷, S. Kabana^{137e}, A. Kaczmarska⁸⁷, M. Kado¹¹⁰, H. Kagan¹¹⁹,
 M. Kagan¹⁴³, A. Kahn⁴¹, A. Kahn¹²⁸, C. Kahra¹⁰⁰, T. Kaji¹⁵³, E. Kajomovitz¹⁵⁰, N. Kakati¹⁶⁹,

I. Kalaitzidou ⁵⁴, C.W. Kalderon ²⁹, A. Kamenshchikov ¹⁵⁵, N.J. Kang ¹³⁶, D. Kar ^{33g}, K. Karava ¹²⁶, M.J. Kareem ^{156b}, E. Karentzos ⁵⁴, I. Karkanias ¹⁵², O. Karkout ¹¹⁴, S.N. Karpov ³⁸, Z.M. Karpova ³⁸, V. Kartvelishvili ⁹¹, A.N. Karyukhin ³⁷, E. Kasimi ¹⁵², J. Katzy ⁴⁸, S. Kaur ³⁴, K. Kawade ¹⁴⁰, M.P. Kawale ¹²⁰, C. Kawamoto ⁸⁸, T. Kawamoto ^{62a}, E.F. Kay ³⁶, F.I. Kaya ¹⁵⁸, S. Kazakos ¹⁰⁷, V.F. Kazanin ³⁷, Y. Ke ¹⁴⁵, J.M. Keaveney ^{33a}, R. Keeler ¹⁶⁵, G.V. Kehris ⁶¹, J.S. Keller ³⁴, A.S. Kelly ⁹⁶, J.J. Kempster ¹⁴⁶, P.D. Kennedy ¹⁰⁰, O. Kepka ¹³¹, B.P. Kerridge ¹³⁴, S. Kersten ¹⁷¹, B.P. Kerševan ⁹³, S. Keshri ⁶⁶, L. Keszeghova ^{28a}, S. Ketabchi Haghighat ¹⁵⁵, R.A. Khan ¹²⁹, A. Khanov ¹²¹, A.G. Kharlamov ³⁷, T. Kharlamova ³⁷, E.E. Khoda ¹³⁸, M. Kholodenko ³⁷, T.J. Khoo ¹⁸, G. Khorauli ¹⁶⁶, J. Khubua ^{149b}, Y.A.R. Khwaira ⁶⁶, B. Kibirige ^{33g}, A. Kilgallon ¹²³, D.W. Kim ^{47a,47b}, Y.K. Kim ³⁹, N. Kimura ⁹⁶, M.K. Kingston ⁵⁵, A. Kirchhoff ⁵⁵, C. Kirfel ²⁴, F. Kirfel ²⁴, J. Kirk ¹³⁴, A.E. Kiryunin ¹¹⁰, C. Kitsaki ¹⁰, O. Kivernyk ²⁴, M. Klassen ^{63a}, C. Klein ³⁴, L. Klein ¹⁶⁶, M.H. Klein ⁴⁴, S.B. Klein ⁵⁶, U. Klein ⁹², P. Klimek ³⁶, A. Klimentov ²⁹, T. Klioutchnikova ³⁶, P. Kluit ¹¹⁴, S. Kluth ¹¹⁰, E. Kneringer ⁷⁹, T.M. Knight ¹⁵⁵, A. Knue ⁴⁹, R. Kobayashi ⁸⁸, D. Kobylanskii ¹⁶⁹, S.F. Koch ¹²⁶, M. Kocian ¹⁴³, P. Kodyš ¹³³, D.M. Koeck ¹²³, P.T. Koenig ²⁴, T. Koffas ³⁴, O. Kolay ⁵⁰, I. Koletsou ⁴, T. Komarek ¹²², K. Köneke ⁵⁴, A.X.Y. Kong ¹, T. Kono ¹¹⁸, N. Konstantinidis ⁹⁶, P. Kontaxakis ⁵⁶, B. Konya ⁹⁸, R. Kopeliansky ⁶⁸, S. Koperny ^{86a}, K. Korcyl ⁸⁷, K. Kordas ^{152,e}, A. Korn ⁹⁶, S. Korn ⁵⁵, I. Korolkov ¹³, N. Korotkova ³⁷, B. Kortman ¹¹⁴, O. Kortner ¹¹⁰, S. Kortner ¹¹⁰, W.H. Kostecka ¹¹⁵, V.V. Kostyukhin ¹⁴¹, A. Kotsokechagia ¹³⁵, A. Kotwal ⁵¹, A. Koulouris ³⁶, A. Kourkumeli-Charalampidi ^{73a,73b}, C. Kourkumelis ⁹, E. Kourlitis ^{110,ae}, O. Kovanda ¹²³, R. Kowalewski ¹⁶⁵, W. Kozanecki ¹³⁵, A.S. Kozhin ³⁷, V.A. Kramarenko ³⁷, G. Kramberger ⁹³, P. Kramer ¹⁰⁰, M.W. Krasny ¹²⁷, A. Krasznahorkay ³⁶, J.W. Kraus ¹⁷¹, J.A. Kremer ⁴⁸, T. Kresse ⁵⁰, J. Kretschmar ⁹², K. Kreul ¹⁸, P. Krieger ¹⁵⁵, S. Krishnamurthy ¹⁰³, M. Krivos ¹³³, K. Krizka ²⁰, K. Kroeninger ⁴⁹, H. Kroha ¹¹⁰, J. Kroll ¹³¹, J. Kroll ¹²⁸, K.S. Krowpman ¹⁰⁷, U. Kruchonak ³⁸, H. Krüger ²⁴, N. Krumnack ⁸¹, M.C. Kruse ⁵¹, O. Kuchinskaja ³⁷, S. Kuday ^{3a}, S. Kuehn ³⁶, R. Kuesters ⁵⁴, T. Kuhl ⁴⁸, V. Kukhtin ³⁸, Y. Kulchitsky ^{37,a}, S. Kuleshov ^{137d,137b}, M. Kumar ^{33g}, N. Kumari ⁴⁸, P. Kumari ^{156b}, A. Kupco ¹³¹, T. Kupfer ⁴⁹, A. Kupich ³⁷, O. Kuprash ⁵⁴, H. Kurashige ⁸⁵, L.L. Kurchaninov ^{156a}, O. Kurdysh ⁶⁶, Y.A. Kurochkin ³⁷, A. Kurova ³⁷, M. Kuze ¹⁵⁴, A.K. Kvam ¹⁰³, J. Kvita ¹²², T. Kwan ¹⁰⁴, N.G. Kyriacou ¹⁰⁶, L.A.O. Laatu ¹⁰², C. Lacasta ¹⁶³, F. Lacava ^{75a,75b}, H. Lacker ¹⁸, D. Lacour ¹²⁷, N.N. Lad ⁹⁶, E. Ladygin ³⁸, B. Laforge ¹²⁷, T. Lagouri ^{27b}, F.Z. Lahbabi ^{35a}, S. Lai ⁵⁵, I.K. Lakomic ^{86a}, N. Lalloue ⁶⁰, J.E. Lambert ¹⁶⁵, S. Lammers ⁶⁸, W. Lampl ⁷, C. Lampoudis ^{152,e}, A.N. Lancaster ¹¹⁵, E. Lançon ²⁹, U. Landgraf ⁵⁴, M.P.J. Landon ⁹⁴, V.S. Lang ⁵⁴, O.K.B. Langrekken ¹²⁵, A.J. Lankford ¹⁵⁹, F. Lanni ³⁶, K. Lantzsch ²⁴, A. Lanza ^{73a}, A. Lapertosa ^{57b,57a}, J.F. Laporte ¹³⁵, T. Lari ^{71a}, F. Lasagni Manghi ^{23b}, M. Lassnig ³⁶, V. Latonova ¹³¹, A. Laudrain ¹⁰⁰, A. Laurier ¹⁵⁰, S.D. Lawlor ¹³⁹, Z. Lawrence ¹⁰¹, R. Lazaridou ¹⁶⁷, M. Lazzaroni ^{71a,71b}, B. Le ¹⁰¹, E.M. Le Boulicaut ⁵¹, B. Leban ⁹³, A. Lebedev ⁸¹, M. LeBlanc ¹⁰¹, F. Ledroit-Guillon ⁶⁰, A.C.A. Lee ⁹⁶, S.C. Lee ¹⁴⁸, S. Lee ^{47a,47b}, T.F. Lee ⁹², L.L. Leeuw ^{33c}, H.P. Lefebvre ⁹⁵, M. Lefebvre ¹⁶⁵, C. Leggett ^{17a}, G. Lehmann Miotto ³⁶, M. Leigh ⁵⁶, W.A. Leight ¹⁰³, W. Leinonen ¹¹³, A. Leisos ^{152,s}, M.A.L. Leite ^{83c}, C.E. Leitgeb ¹⁸, R. Leitner ¹³³, K.J.C. Leney ⁴⁴, T. Lenz ²⁴, S. Leone ^{74a}, C. Leonidopoulos ⁵², A. Leopold ¹⁴⁴, C. Leroy ¹⁰⁸, R. Les ¹⁰⁷, C.G. Lester ³², M. Levchenko ³⁷, J. Levêque ⁴, L.J. Levinson ¹⁶⁹, G. Levrini ^{23b,23a}, M.P. Lewicki ⁸⁷, D.J. Lewis ⁴, A. Li ⁵, B. Li ^{62b}, C. Li ^{62a}, C.-Q. Li ¹¹⁰, H. Li ^{62a}, H. Li ^{62b}, H. Li ^{14c}, H. Li ^{14b}, H. Li ^{62b}, J. Li ^{62c}, K. Li ¹³⁸,

L. Li ^{62c}, M. Li ^{14a,14e}, Q.Y. Li ^{62a}, S. Li ^{14a,14e}, S. Li ^{62d,62c,d}, T. Li ⁵, X. Li ¹⁰⁴, Z. Li ¹²⁶,
 Z. Li ¹⁰⁴, Z. Li ^{14a,14e}, S. Liang ^{14a,14e}, Z. Liang ^{14a}, M. Liberatore ¹³⁵, B. Liberti ^{76a}, K. Lie ^{64c},
 J. Lieber Marin ^{83b}, H. Lien ⁶⁸, K. Lin ¹⁰⁷, R.E. Lindley ⁷, J.H. Lindon ², E. Lipeles ¹²⁸,
 A. Lipniacka ¹⁶, A. Lister ¹⁶⁴, J.D. Little ⁴, B. Liu ^{14a}, B.X. Liu ¹⁴², D. Liu ^{62d,62c}, J.B. Liu ^{62a},
 J.K.K. Liu ³², K. Liu ^{62d,62c}, M. Liu ^{62a}, M.Y. Liu ^{62a}, P. Liu ^{14a}, Q. Liu ^{62d,138,62c}, X. Liu ^{62a},
 X. Liu ^{62b}, Y. Liu ^{14d,14e}, Y.L. Liu ^{62b}, Y.W. Liu ^{62a}, J. Llorente Merino ¹⁴², S.L. Lloyd ⁹⁴,
 E.M. Lobodzinska ⁴⁸, P. Loch ⁷, T. Lohse ¹⁸, K. Lohwasser ¹³⁹, E. Loiacono ⁴⁸, M. Lokajicek ^{131,*},
 J.D. Lomas ²⁰, J.D. Long ¹⁶², I. Longarini ¹⁵⁹, L. Longo ^{70a,70b}, R. Longo ¹⁶², I. Lopez Paz ⁶⁷,
 A. Lopez Solis ⁴⁸, N. Lorenzo Martinez ⁴, A.M. Lory ¹⁰⁹, G. Löschcke Centeno ¹⁴⁶, O. Loseva ³⁷,
 X. Lou ^{47a,47b}, X. Lou ^{14a,14e}, A. Lounis ⁶⁶, P.A. Love ⁹¹, G. Lu ^{14a,14e}, M. Lu ⁸⁰, S. Lu ¹²⁸,
 Y.J. Lu ⁶⁵, H.J. Lubatti ¹³⁸, C. Luci ^{75a,75b}, F.L. Lucio Alves ^{14c}, F. Luehring ⁶⁸, I. Luise ¹⁴⁵,
 O. Lukianchuk ⁶⁶, O. Lundberg ¹⁴⁴, B. Lund-Jensen ^{144,*}, N.A. Luongo ⁶, M.S. Lutz ³⁶, A.B. Lux ²⁵,
 D. Lynn ²⁹, R. Lysak ¹³¹, E. Lytken ⁹⁸, V. Lyubushkin ³⁸, T. Lyubushkina ³⁸, M.M. Lyukova ¹⁴⁵,
 H. Ma ²⁹, K. Ma ^{62a}, L.L. Ma ^{62b}, W. Ma ^{62a}, Y. Ma ¹²¹, D.M. Mac Donnell ¹⁶⁵, G. Maccarrone ⁵³,
 J.C. MacDonald ¹⁰⁰, P.C. Machado De Abreu Farias ^{83b}, R. Madar ⁴⁰, W.F. Mader ⁵⁰, T. Madula ⁹⁶,
 J. Maeda ⁸⁵, T. Maeno ²⁹, H. Maguire ¹³⁹, V. Maiboroda ¹³⁵, A. Maio ^{130a,130b,130d}, K. Maj ^{86a},
 O. Majersky ⁴⁸, S. Majewski ¹²³, N. Makovec ⁶⁶, V. Maksimovic ¹⁵, B. Malaescu ¹²⁷, Pa. Malecki ⁸⁷,
 V.P. Maleev ³⁷, F. Malek ^{60,o}, M. Mali ⁹³, D. Malito ⁹⁵, U. Mallik ⁸⁰, S. Maltezos ¹⁰, S. Malyukov ³⁸,
 J. Mamuzic ¹³, G. Mancini ⁵³, M.N. Mancini ²⁶, G. Manco ^{73a,73b}, J.P. Mandalia ⁹⁴, I. Mandić ⁹³,
 L. Manhaes de Andrade Filho ^{83a}, I.M. Maniatis ¹⁶⁹, J. Manjarres Ramos ^{102,ab}, D.C. Mankad ¹⁶⁹,
 A. Mann ¹⁰⁹, S. Manzoni ³⁶, L. Mao ^{62c}, X. Mapekula ^{33c}, A. Marantis ^{152,s}, G. Marchiori ⁵,
 M. Marcisovsky ¹³¹, C. Marcon ^{71a}, M. Marinescu ²⁰, S. Marium ⁴⁸, M. Marjanovic ¹²⁰,
 E.J. Marshall ⁹¹, Z. Marshall ^{17a}, S. Marti-Garcia ¹⁶³, T.A. Martin ¹⁶⁷, V.J. Martin ⁵²,
 B. Martin dit Latour ¹⁶, L. Martinelli ^{75a,75b}, M. Martinez ^{13,t}, P. Martinez Agullo ¹⁶³,
 V.I. Martinez Outschoorn ¹⁰³, P. Martinez Suarez ¹³, S. Martin-Haugh ¹³⁴, V.S. Martoiu ^{27b},
 A.C. Martyniuk ⁹⁶, A. Marzin ³⁶, D. Mascione ^{78a,78b}, L. Masetti ¹⁰⁰, T. Mashimo ¹⁵³, J. Masik ¹⁰¹,
 A.L. Maslennikov ³⁷, P. Massarotti ^{72a,72b}, P. Mastrandrea ^{74a,74b}, A. Mastroberardino ^{43b,43a},
 T. Masubuchi ¹⁵³, T. Mathisen ¹⁶¹, J. Matousek ¹³³, N. Matsuzawa ¹⁵³, J. Maurer ^{27b}, B. Maček ⁹³,
 D.A. Maximov ³⁷, R. Mazini ¹⁴⁸, I. Maznas ¹¹⁵, M. Mazza ¹⁰⁷, S.M. Mazza ¹³⁶, E. Mazzeo ^{71a,71b},
 C. Mc Ginn ²⁹, J.P. Mc Gowan ¹⁰⁴, S.P. Mc Kee ¹⁰⁶, C.C. McCracken ¹⁶⁴, 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, S.J. McMahon ¹³⁴, C.M. Mcpartland ⁹², R.A. McPherson ^{165,x},
 S. Mehlhase ¹⁰⁹, A. Mehta ⁹², D. Melini ¹⁶³, B.R. Mellado Garcia ^{33g}, A.H. Melo ⁵⁵, F. Meloni ⁴⁸,
 A.M. Mendes Jacques Da Costa ¹⁰¹, H.Y. Meng ¹⁵⁵, L. Meng ⁹¹, S. Menke ¹¹⁰, M. Mentink ³⁶,
 E. Meoni ^{43b,43a}, G. Mercado ¹¹⁵, C. Merlassino ^{69a,69c}, L. Merola ^{72a,72b}, C. Meroni ^{71a,71b},
 J. Metcalfe ⁶, A.S. Mete ⁶, C. Meyer ⁶⁸, J-P. Meyer ¹³⁵, R.P. Middleton ¹³⁴, L. Mijović ⁵²,
 G. Mikenberg ¹⁶⁹, M. Mikestikova ¹³¹, M. Mikuž ⁹³, H. Mildner ¹⁰⁰, A. Milic ³⁶, D.W. Miller ³⁹,
 E.H. Miller ¹⁴³, L.S. Miller ³⁴, A. Milov ¹⁶⁹, D.A. Milstead ^{47a,47b}, T. Min ^{14c}, A.A. Minaenko ³⁷,
 I.A. Minashvili ^{149b}, L. Mince ⁵⁹, A.I. Mincer ¹¹⁷, B. Mindur ^{86a}, M. Mineev ³⁸, Y. Mino ⁸⁸,
 L.M. Mir ¹³, M. Miralles Lopez ⁵⁹, M. Mironova ^{17a}, A. Mishima ¹⁵³, M.C. Missio ¹¹³, A. Mitra ¹⁶⁷,
 V.A. Mitsou ¹⁶³, Y. Mitsumori ¹¹¹, O. Miu ¹⁵⁵, P.S. Miyagawa ⁹⁴, T. Mkrtchyan ^{63a},
 M. Mlinarevic ⁹⁶, T. Mlinarevic ⁹⁶, M. Mlynarikova ³⁶, S. Mobius ¹⁹, P. Mogg ¹⁰⁹,
 M.H. Mohamed Farook ¹¹², A.F. Mohammed ^{14a,14e}, S. Mohapatra ⁴¹, G. Mokgatitswane ^{33g},

L. Molero ¹⁶⁹, B. Mondal ¹⁴¹, S. Mondal ¹³², K. Mönig ⁴⁸, E. Monnier ¹⁰², L. Monsonis Romero ¹⁶³, J. Montejo Berlingen ¹³, M. Montella ¹¹⁹, F. Montekali ^{77a,77b}, F. Monticelli ⁹⁰, S. Monzani ^{69a,69c}, N. Morange ⁶⁶, A.L. Moreira De Carvalho ^{130a}, M. Moreno Llácer ¹⁶³, C. Moreno Martinez ⁵⁶, P. Moretini ^{57b}, S. Morgenstern ³⁶, M. Morii ⁶¹, M. Morinaga ¹⁵³, F. Morodei ^{75a,75b}, L. Morvaj ³⁶, P. Moschovakos ³⁶, B. Moser ³⁶, M. Mosidze ^{149b}, T. Moskalets ⁵⁴, P. Moskvitina ¹¹³, J. Moss ^{31,l}, E.J.W. Moyse ¹⁰³, O. Mtintsilana ^{33g}, S. Muanza ¹⁰², J. Mueller ¹²⁹, D. Muenstermann ⁹¹, R. Müller ¹⁹, G.A. Mullier ¹⁶¹, A.J. Mullin ³², J.J. Mullin ¹²⁸, D.P. Mungo ¹⁵⁵, D. Munoz Perez ¹⁶³, F.J. Munoz Sanchez ¹⁰¹, M. Murin ¹⁰¹, W.J. Murray ^{167,134}, M. Muškinja ^{17a}, C. Mwewa ²⁹, A.G. Myagkov ^{37,a}, A.J. Myers ⁸, G. Myers ⁶⁸, M. Myska ¹³², B.P. Nachman ^{17a}, O. Nackenhorst ⁴⁹, K. Nagai ¹²⁶, K. Nagano ⁸⁴, J.L. Nagle ^{29,ai}, E. Nagy ¹⁰², A.M. Nairz ³⁶, Y. Nakahama ⁸⁴, K. Nakamura ⁸⁴, T. Nakamura ⁸⁵, K. Nakkalil ⁵, H. Nanjo ¹²⁴, R. Narayan ⁴⁴, E.A. Narayanan ¹¹², I. Naryshkin ³⁷, M. Naseri ³⁴, S. Nasri ^{116b}, C. Nass ²⁴, G. Navarro ^{22a}, J. Navarro-Gonzalez ¹⁶³, R. Nayak ¹⁵¹, A. Nayaz ¹⁸, P.Y. Nechaeva ³⁷, F. Nechansky ⁴⁸, L. Nedic ¹²⁶, T.J. Neep ²⁰, A. Negri ^{73a,73b}, M. Negrini ^{23b}, C. Nellist ¹¹⁴, C. Nelson ¹⁰⁴, K. Nelson ¹⁰⁶, S. Nemecek ¹³¹, M. Nessi ^{36,h}, M.S. Neubauer ¹⁶², F. Neuhaus ¹⁰⁰, J. Neundorf ⁴⁸, R. Newhouse ¹⁶⁴, P.R. Newman ²⁰, C.W. Ng ¹²⁹, Y.W.Y. Ng ⁴⁸, B. Ngair ^{116a}, H.D.N. Nguyen ¹⁰⁸, R.B. Nickerson ¹²⁶, R. Nicolaidou ¹³⁵, J. Nielsen ¹³⁶, M. Niemeyer ⁵⁵, J. Niermann ^{55,36}, N. Nikiforou ³⁶, V. Nikolaenko ^{37,a}, I. Nikolic-Audit ¹²⁷, K. Nikolopoulos ²⁰, P. Nilsson ²⁹, I. Ninca ⁴⁸, H.R. Nindhito ⁵⁶, G. Ninio ¹⁵¹, A. Nisati ^{75a}, N. Nishu ², R. Nisius ¹¹⁰, J-E. Nitschke ⁵⁰, E.K. Nkadimeng ^{33g}, T. Nobe ¹⁵³, D.L. Noel ³², T. Nommensen ¹⁴⁷, M.B. Norfolk ¹³⁹, R.R.B. Norisam ⁹⁶, B.J. Norman ³⁴, M. Noury ^{35a}, J. Novak ⁹³, T. Novak ⁴⁸, L. Novotny ¹³², R. Novotny ¹¹², L. Nozka ¹²², K. Ntekas ¹⁵⁹, N.M.J. Nunes De Moura Junior ^{83b}, E. Nurse ⁹⁶, J. Ocariz ¹²⁷, A. Ochi ⁸⁵, I. Ochoa ^{130a}, S. Oerdek ^{48,u}, J.T. Offermann ³⁹, A. Ogrodnik ¹³³, A. Oh ¹⁰¹, C.C. Ohm ¹⁴⁴, H. Oide ⁸⁴, R. Oishi ¹⁵³, M.L. Ojeda ⁴⁸, Y. Okumura ¹⁵³, L.F. Oleiro Seabra ^{130a}, S.A. Olivares Pino ^{137d}, D. Oliveira Damazio ²⁹, D. Oliveira Goncalves ^{83a}, J.L. Oliver ¹⁵⁹, Ö.O. Öncel ⁵⁴, A.P. O'Neill ¹⁹, A. Onofre ^{130a,130e}, P.U.E. Onyisi ¹¹, M.J. Oreglia ³⁹, G.E. Orellana ⁹⁰, D. Orestano ^{77a,77b}, N. Orlando ¹³, R.S. Orr ¹⁵⁵, V. O'Shea ⁵⁹, L.M. Osojnak ¹²⁸, R. Ospanov ^{62a}, G. Otero y Garzon ³⁰, H. Otono ⁸⁹, P.S. Ott ^{63a}, G.J. Ottino ^{17a}, M. Ouchrif ^{35d}, F. Ould-Saada ¹²⁵, M. Owen ⁵⁹, R.E. Owen ¹³⁴, K.Y. Oyulmaz ^{21a}, V.E. Ozcan ^{21a}, F. Ozturk ⁸⁷, N. Ozturk ⁸, S. Ozturk ⁸², H.A. Pacey ¹²⁶, A. Pacheco Pages ¹³, C. Padilla Aranda ¹³, G. Padovano ^{75a,75b}, S. Pagan Griso ^{17a}, G. Palacino ⁶⁸, A. Palazzo ^{70a,70b}, J. Pampel ²⁴, J. Pan ¹⁷², T. Pan ^{64a}, D.K. Panchal ¹¹, C.E. Pandini ¹¹⁴, J.G. Panduro Vazquez ⁹⁵, H.D. Pandya ¹, H. Pang ^{14b}, P. Pani ⁴⁸, G. Panizzo ^{69a,69c}, L. Paolozzi ⁵⁶, S. Parajuli ¹⁶², A. Paramonov ⁶, C. Paraskevopoulos ⁵³, D. Paredes Hernandez ^{64b}, A. Pareti ^{73a,73b}, K.R. Park ⁴¹, T.H. Park ¹⁵⁵, M.A. Parker ³², F. Parodi ^{57b,57a}, E.W. Parrish ¹¹⁵, V.A. Parrish ⁵², J.A. Parsons ⁴¹, U. Parzefall ⁵⁴, B. Pascual Dias ¹⁰⁸, L. Pascual Dominguez ¹⁵¹, E. Pasqualucci ^{75a}, S. Passaggio ^{57b}, F. Pastore ⁹⁵, P. Patel ⁸⁷, U.M. Patel ⁵¹, J.R. Pater ¹⁰¹, T. Pauly ³⁶, C.I. Pazos ¹⁵⁸, J. Pearkes ¹⁴³, M. Pedersen ¹²⁵, R. Pedro ^{130a}, S.V. Peleganchuk ³⁷, O. Penc ³⁶, E.A. Pender ⁵², G.D. Penn ¹⁷², K.E. Penski ¹⁰⁹, M. Penzin ³⁷, B.S. Peralva ^{83d}, A.P. Pereira Peixoto ⁶⁰, L. Pereira Sanchez ^{47a,47b}, D.V. Perepelitsa ^{29,ai}, E. Perez Codina ^{156a}, M. Perganti ¹⁰, H. Pernegger ³⁶, O. Perrin ⁴⁰, 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}, A. Petrukhin ¹⁴¹, M. Pettee ^{17a}, N.E. Pettersson ³⁶, A. Petukhov ³⁷, K. Petukhova ¹³³, R. Pezoa ^{137f}, L. Pezzotti ³⁶, G. Pezzullo ¹⁷², T.M. Pham ¹⁷⁰, T. Pham ¹⁰⁵, P.W. Phillips ¹³⁴, G. Piacquadio ¹⁴⁵, E. Pianori ^{17a}, F. Piazza ¹²³,

R. Piegaia ³⁰, D. Pietreanu ^{27b}, A.D. Pilkington ¹⁰¹, M. Pinamonti ^{69a,69c}, J.L. Pinfeld ², B.C. Pinheiro Pereira ^{130a}, A.E. Pinto Pinoargote ^{100,135}, L. Pintucci ^{69a,69c}, K.M. Piper ¹⁴⁶, A. Pirttikoski ⁵⁶, D.A. Pizzi ³⁴, L. Pizzimento ^{64b}, A. Pizzini ¹¹⁴, M.-A. Pleier ²⁹, V. Plesanovs ⁵⁴, V. Pleskot ¹³³, E. Plotnikova ³⁸, G. Poddar ⁴, R. Poettgen ⁹⁸, L. Poggioli ¹²⁷, I. Pokharel ⁵⁵, S. Polacek ¹³³, G. Polesello ^{73a}, A. Poley ^{142,156a}, A. Polini ^{23b}, C.S. Pollard ¹⁶⁷, Z.B. Pollock ¹¹⁹, E. Pompa Pacchi ^{75a,75b}, D. Ponomarenko ¹¹³, L. Pontecorvo ³⁶, S. Popa ^{27a}, G.A. Popeneciu ^{27d}, A. Poreba ³⁶, D.M. Portillo Quintero ^{156a}, S. Pospisil ¹³², M.A. Postill ¹³⁹, P. Postolache ^{27c}, K. Potamianos ¹⁶⁷, P.A. Potepa ^{86a}, I.N. Potrap ³⁸, C.J. Potter ³², H. Potti ¹, T. Poulsen ⁴⁸, J. Poveda ¹⁶³, M.E. Pozo Astigarraga ³⁶, A. Prades Ibanez ¹⁶³, J. Pretel ⁵⁴, D. Price ¹⁰¹, M. Primavera ^{70a}, M.A. Principe Martin ⁹⁹, R. Privara ¹²², T. Procter ⁵⁹, M.L. Proffitt ¹³⁸, N. Proklova ¹²⁸, K. Prokofiev ^{64c}, G. Proto ¹¹⁰, J. Proudfoot ⁶, M. Przybycien ^{86a}, W.W. Przygoda ^{86b}, A. Psallidas ⁴⁶, J.E. Puddefoot ¹³⁹, D. Pudzha ³⁷, D. Pyatiizbyantseva ³⁷, J. Qian ¹⁰⁶, D. Qichen ¹⁰¹, Y. Qin ¹⁰¹, T. Qiu ⁵², A. Quadt ⁵⁵, M. Queitsch-Maitland ¹⁰¹, G. Quetant ⁵⁶, R.P. Quinn ¹⁶⁴, G. Rabanal Bolanos ⁶¹, 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 ^{48,14e}, N.P. Rapheeha ^{33g}, H. Rasheed ^{27b}, V. Raskina ¹²⁷, D.F. Rassloff ^{63a}, A. Rastogi ^{17a}, S. Rave ¹⁰⁰, B. Ravina ⁵⁵, I. Ravinovich ¹⁶⁹, M. Raymond ³⁶, A.L. Read ¹²⁵, N.P. Readioff ¹³⁹, D.M. Rebuzzi ^{73a,73b}, G. Redlinger ²⁹, A.S. Reed ¹¹⁰, K. Reeves ²⁶, J.A. Reidelsturz ¹⁷¹, D. Reikher ¹⁵¹, A. Rej ⁴⁹, C. Rembser ³⁶, M. Renda ^{27b}, M.B. Rendel ¹¹⁰, F. Renner ⁴⁸, A.G. Rennie ¹⁵⁹, A.L. Rescia ⁴⁸, S. Resconi ^{71a}, M. Ressegotti ^{57b,57a}, S. Rettie ³⁶, J.G. Reyes Rivera ¹⁰⁷, E. Reynolds ^{17a}, O.L. Rezanova ³⁷, P. Reznicek ¹³³, H. Riani ^{35d}, N. Ribaric ⁹¹, E. Ricci ^{78a,78b}, R. Richter ¹¹⁰, S. Richter ^{47a,47b}, E. Richter-Was ^{86b}, M. Ridel ¹²⁷, S. Ridouani ^{35d}, P. Rieck ¹¹⁷, P. Riedler ³⁶, E.M. Riefel ^{47a,47b}, J.O. Rieger ¹¹⁴, M. Rijssenbeek ¹⁴⁵, M. Rimoldi ³⁶, L. Rinaldi ^{23b,23a}, T.T. Rinn ²⁹, M.P. Rinnagel ¹⁰⁹, G. Ripellino ¹⁶¹, I. Riu ¹³, J.C. Rivera Vergara ¹⁶⁵, F. Rizatdinova ¹²¹, E. Rizvi ⁹⁴, B.A. Roberts ¹⁶⁷, B.R. Roberts ^{17a}, S.H. Robertson ^{104,x}, 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 ¹⁵⁹, A.R. Roepe-Gier ¹³⁶, J. Roggel ¹⁷¹, O. Røhne ¹²⁵, R.A. Rojas ¹⁰³, C.P.A. Roland ¹²⁷, J. Roloff ²⁹, A. Romaniouk ³⁷, E. Romano ^{73a,73b}, M. Romano ^{23b}, A.C. Romero Hernandez ¹⁶², N. Rompotis ⁹², L. Roos ¹²⁷, S. Rosati ^{75a}, B.J. Rosser ³⁹, E. Rossi ¹²⁶, E. Rossi ^{72a,72b}, L.P. Rossi ^{57b}, L. Rossini ⁵⁴, R. Rosten ¹¹⁹, M. Rotaru ^{27b}, B. Rottler ⁵⁴, C. Rougier ^{102,ab}, D. Rousseau ⁶⁶, D. Rousso ³², A. Roy ¹⁶², S. Roy-Garand ¹⁵⁵, A. Rozanov ¹⁰², Z.M.A. Rozario ⁵⁹, Y. Rozen ¹⁵⁰, A. Rubio Jimenez ¹⁶³, A.J. Ruby ⁹², V.H. Ruelas Rivera ¹⁸, T.A. Ruggeri ¹, A. Ruggiero ¹²⁶, A. Ruiz-Martinez ¹⁶³, A. Rummler ³⁶, Z. Rurikova ⁵⁴, N.A. Rusakovich ³⁸, H.L. Russell ¹⁶⁵, G. Russo ^{75a,75b}, J.P. Rutherford ⁷, S. Rutherford Colmenares ³², K. Rybacki ⁹¹, M. Rybar ¹³³, E.B. Rye ¹²⁵, A. Ryzhov ⁴⁴, J.A. Sabater Iglesias ⁵⁶, P. Sabatini ¹⁶³, H.F.W. Sadrozinski ¹³⁶, F. Safai Tehrani ^{75a}, B. Safarzadeh Samani ¹³⁴, M. Safdari ¹⁴³, S. Saha ¹⁶⁵, M. Sahinsoy ¹¹⁰, A. Saibel ¹⁶³, M. Saimpert ¹³⁵, M. Saito ¹⁵³, T. Saito ¹⁵³, D. Salamani ³⁶, A. Salnikov ¹⁴³, J. Salt ¹⁶³, A. Salvador Salas ¹⁵¹, D. Salvatore ^{43b,43a}, F. Salvatore ¹⁴⁶, A. Salzburger ³⁶, D. Sammel ⁵⁴, D. Sampsonidis ^{152,e}, D. Sampsonidou ¹²³, J. Sánchez ¹⁶³, V. Sanchez Sebastian ¹⁶³, H. Sandaker ¹²⁵, C.O. Sander ⁴⁸, J.A. Sandesara ¹⁰³, M. Sandhoff ¹⁷¹, C. Sandoval ^{22b}, D.P.C. Sankey ¹³⁴, T. Sano ⁸⁸, A. Sansoni ⁵³, L. Santi ^{75a,75b}, C. Santoni ⁴⁰, H. Santos ^{130a,130b},

A. Santra ¹⁶⁹, K.A. Saoucha ¹⁶⁰, J.G. Saraiva ^{130a,130d}, J. Sardain ⁷, A. Sarkar ⁵⁶, O. Sasaki ⁸⁴,
 K. Sato ¹⁵⁷, C. Sauer ^{63b}, F. Sauerburger ⁵⁴, E. Sauvan ⁴, P. Savard ^{155,ag}, R. Sawada ¹⁵³,
 C. Sawyer ¹³⁴, L. Sawyer ⁹⁷, I. Sayago Galvan ¹⁶³, C. Sbarra ^{23b}, A. Sbrizzi ^{23b,23a}, T. Scanlon ⁹⁶,
 J. Schaarschmidt ¹³⁸, U. Schäfer ¹⁰⁰, A.C. Schaffer ^{66,44}, D. Schaile ¹⁰⁹, R.D. Schamberger ¹⁴⁵,
 C. Scharf ¹⁸, M.M. Schefer ¹⁹, V.A. Schegelsky ³⁷, D. Scheirich ¹³³, F. Schenck ¹⁸, M. Schernau ¹⁵⁹,
 C. Scheulen ⁵⁵, C. Schiavi ^{57b,57a}, E.J. Schioppa ^{70a,70b}, M. Schioppa ^{43b,43a}, B. Schlag ^{143,n},
 K.E. Schleicher ⁵⁴, S. Schlenker ³⁶, J. Schmeing ¹⁷¹, M.A. Schmidt ¹⁷¹, K. Schmieden ¹⁰⁰,
 C. Schmitt ¹⁰⁰, N. Schmitt ¹⁰⁰, S. Schmitt ⁴⁸, L. Schoeffel ¹³⁵, A. Schoening ^{63b}, P.G. Scholer ⁵⁴,
 E. Schopf ¹²⁶, M. Schott ¹⁰⁰, J. Schovancova ³⁶, S. Schramm ⁵⁶, T. Schroer ⁵⁶,
 H-C. Schultz-Coulon ^{63a}, M. Schumacher ⁵⁴, B.A. Schumm ¹³⁶, Ph. Schune ¹³⁵, A.J. Schuy ¹³⁸,
 H.R. Schwartz ¹³⁶, A. Schwartzman ¹⁴³, T.A. Schwarz ¹⁰⁶, Ph. Schwemling ¹³⁵, R. Schwienhorst ¹⁰⁷,
 A. Sciandra ¹³⁶, G. Sciolla ²⁶, F. Scuri ^{74a}, C.D. Sebastiani ⁹², K. Sedlaczek ¹¹⁵, P. Seema ¹⁸,
 S.C. Seidel ¹¹², A. Seiden ¹³⁶, B.D. Seidlitz ⁴¹, C. Seitz ⁴⁸, J.M. Seixas ^{83b}, G. Sekhniaidze ^{72a},
 L. Selem ⁶⁰, N. Semprini-Cesari ^{23b,23a}, D. Sengupta ⁵⁶, V. Senthilkumar ¹⁶³, L. Serin ⁶⁶,
 L. Serkin ^{69a,69b}, M. Sessa ^{76a,76b}, H. Severini ¹²⁰, F. Sforza ^{57b,57a}, A. Sfyrla ⁵⁶, Q. Sha ^{14a},
 E. Shabalina ⁵⁵, R. Shaheen ¹⁴⁴, J.D. Shahinian ¹²⁸, D. Shaked Renous ¹⁶⁹, L.Y. Shan ^{14a},
 M. Shapiro ^{17a}, A. Sharma ³⁶, A.S. Sharma ¹⁶⁴, P. Sharma ⁸⁰, P.B. Shatalov ³⁷, K. Shaw ¹⁴⁶,
 S.M. Shaw ¹⁰¹, A. Shcherbakova ³⁷, Q. Shen ^{62c,5}, D.J. Sheppard ¹⁴², P. Sherwood ⁹⁶, L. Shi ⁹⁶,
 X. Shi ^{14a}, C.O. Shimmin ¹⁷², J.D. Shinner ⁹⁵, I.P.J. Shipsey ¹²⁶, S. Shirabe ⁸⁹, M. Shiyakova ^{38,v},
 J. Shlomi ¹⁶⁹, M.J. Shochet ³⁹, J. Shojaii ¹⁰⁵, D.R. Shope ¹²⁵, B. Shrestha ¹²⁰, S. Shrestha ^{119,aj},
 E.M. Shrif ^{33g}, M.J. Shroff ¹⁶⁵, P. Sicho ¹³¹, A.M. Sickles ¹⁶², E. Sideras Haddad ^{33g}, A. Sidoti ^{23b},
 F. Siegert ⁵⁰, Dj. Sijacki ¹⁵, F. Sili ⁹⁰, J.M. Silva ²⁰, M.V. Silva Oliveira ²⁹, S.B. Silverstein ^{47a},
 S. Simion ⁶⁶, R. Simoniello ³⁶, E.L. Simpson ⁵⁹, H. Simpson ¹⁴⁶, L.R. Simpson ¹⁰⁶, N.D. Simpson ⁹⁸,
 S. Simsek ⁸², S. Sindhu ⁵⁵, P. Sinervo ¹⁵⁵, S. Singh ¹⁵⁵, S. Sinha ⁴⁸, S. Sinha ¹⁰¹, M. Sioli ^{23b,23a},
 I. Siral ³⁶, E. Sitnikova ⁴⁸, J. Sjölin ^{47a,47b}, A. Skaf ⁵⁵, E. Skorda ²⁰, P. Skubic ¹²⁰, M. Slawinska ⁸⁷,
 V. Smakhtin ¹⁶⁹, B.H. Smart ¹³⁴, S.Yu. Smirnov ³⁷, Y. Smirnov ³⁷, L.N. Smirnova ^{37,a}, O. Smirnova ⁹⁸,
 A.C. Smith ⁴¹, E.A. Smith ³⁹, H.A. Smith ¹²⁶, J.L. Smith ⁹², R. Smith ¹⁴³, M. Smizanska ⁹¹,
 K. Smolek ¹³², A.A. Snesev ³⁷, S.R. Snider ¹⁵⁵, H.L. Snoek ¹¹⁴, S. Snyder ²⁹, R. Sobie ^{165,x},
 A. Soffer ¹⁵¹, C.A. Solans Sanchez ³⁶, E.Yu. Soldatov ³⁷, U. Soldevila ¹⁶³, A.A. Solodkov ³⁷,
 S. Solomon ²⁶, A. Soloshenko ³⁸, K. Solovieva ⁵⁴, O.V. Solovyanov ⁴⁰, V. Solovyev ³⁷, P. Sommer ³⁶,
 A. Sonay ¹³, W.Y. Song ^{156b}, A. Sopczak ¹³², A.L. Sapiro ⁹⁶, F. Sopkova ^{28b}, J.D. Sorenson ¹¹²,
 I.R. Sotarriva Alvarez ¹⁵⁴, V. Sothilingam ^{63a}, O.J. Soto Sandoval ^{137c,137b}, S. Sottocornola ⁶⁸,
 R. Soualah ¹⁶⁰, Z. Soumami ^{35e}, D. South ⁴⁸, N. Soybelman ¹⁶⁹, S. Spagnolo ^{70a,70b}, M. Spalla ¹¹⁰,
 D. Sperlich ⁵⁴, G. Spigo ³⁶, S. Spinali ⁹¹, D.P. Spiteri ⁵⁹, M. Spousta ¹³³, E.J. Staats ³⁴,
 R. Stamen ^{63a}, A. Stampekis ²⁰, M. Standke ²⁴, E. Stanecka ⁸⁷, M.V. Stange ⁵⁰, B. Stanislaus ^{17a},
 M.M. Stanitzki ⁴⁸, B. Stapf ⁴⁸, E.A. Starchenko ³⁷, G.H. Stark ¹³⁶, J. Stark ^{102,ab}, P. Staroba ¹³¹,
 P. Starovoitov ^{63a}, S. Stärz ¹⁰⁴, R. Staszewski ⁸⁷, G. Stavropoulos ⁴⁶, J. Steentoft ¹⁶¹, P. Steinberg ²⁹,
 B. Stelzer ^{142,156a}, H.J. Stelzer ¹²⁹, O. Stelzer-Chilton ^{156a}, H. Stenzel ⁵⁸, T.J. Stevenson ¹⁴⁶,
 G.A. Stewart ³⁶, J.R. Stewart ¹²¹, M.C. Stockton ³⁶, G. Stoicea ^{27b}, M. Stolarski ^{130a}, S. Stonjek ¹¹⁰,
 A. Straessner ⁵⁰, J. Strandberg ¹⁴⁴, S. Strandberg ^{47a,47b}, M. Stratmann ¹⁷¹, M. Strauss ¹²⁰,
 T. Strebler ¹⁰², P. Strizenec ^{28b}, R. Ströhmer ¹⁶⁶, D.M. Strom ¹²³, R. Stroynowski ⁴⁴,
 A. Strubig ^{47a,47b}, S.A. Stucci ²⁹, B. Stugu ¹⁶, J. Stupak ¹²⁰, N.A. Styles ⁴⁸, D. Su ¹⁴³, S. Su ^{62a},
 W. Su ^{62d}, X. Su ^{62a,66}, K. Sugizaki ¹⁵³, V.V. Sulin ³⁷, M.J. Sullivan ⁹², D.M.S. Sultan ^{78a,78b},

L. Sultanalievya [137](#), S. Sultansoy [3b](#), T. Sumida [88](#), S. Sun [106](#), S. Sun [170](#),
 O. Sunneborn Gudnadottir [161](#), N. Sur [102](#), M.R. Sutton [146](#), H. Suzuki [157](#), M. Svatos [131](#),
 M. Swiatlowski [156a](#), T. Swirski [166](#), I. Sykora [28a](#), M. Sykora [133](#), T. Sykora [133](#), D. Ta [100](#),
 K. Tackmann [48,u](#), A. Taffard [159](#), R. Tafirout [156a](#), J.S. Tafoya Vargas [66](#), Y. Takubo [84](#), M. Talby [102](#),
 A.A. Talyshv [37](#), K.C. Tam [64b](#), N.M. Tamir [151](#), A. Tanaka [153](#), J. Tanaka [153](#), R. Tanaka [66](#),
 M. Tanasini [57b,57a](#), Z. Tao [164](#), S. Tapia Araya [137f](#), S. Tapprogge [100](#),
 A. Tarek Abouelfadl Mohamed [107](#), S. Tarem [150](#), K. Tariq [14a](#), G. Tarna [102,27b](#), G.F. Tartarelli [71a](#),
 P. Tas [133](#), M. Tasevsky [131](#), E. Tassi [43b,43a](#), A.C. Tate [162](#), G. Tateno [153](#), Y. Tayalati [35e,w](#),
 G.N. Taylor [105](#), W. Taylor [156b](#), A.S. Tee [170](#), 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](#), R.J. Teuscher [155,x](#), A. Thaler [79](#),
 O. Theiner [56](#), N. Themistokleous [52](#), T. Theveneaux-Pelzer [102](#), O. Thielmann [171](#), D.W. Thomas [95](#),
 J.P. Thomas [20](#), E.A. Thompson [17a](#), P.D. Thompson [20](#), E. Thomson [128](#), Y. Tian [55](#),
 V. Tikhomirov [37,a](#), Yu.A. Tikhonov [37](#), S. Timoshenko [37](#), D. Timoshyn [133](#), E.X.L. Ting [1](#),
 P. Tipton [172](#), S.H. Tlou [33g](#), A. Tnourji [40](#), K. Todome [154](#), S. Todorova-Nova [133](#), S. Todt [50](#),
 M. Togawa [84](#), J. Tojo [89](#), S. Tokár [28a](#), K. Tokushuku [84](#), O. Toldaiev [68](#), R. Tombs [32](#),
 M. Tomoto [84,111](#), L. Tompkins [143,n](#), K.W. Topolnicki [86b](#), E. Torrence [123](#), H. Torres [102,ab](#),
 E. Torró Pastor [163](#), M. Toscani [30](#), C. Tosciri [39](#), M. Tost [11](#), D.R. Tovey [139](#), A. Traet [16](#),
 I.S. Trandafir [27b](#), T. Trefzger [166](#), A. Tricoli [29](#), I.M. Trigger [156a](#), S. Trincaz-Duvoid [127](#),
 D.A. Trischuk [26](#), B. Trocmé [60](#), L. Truong [33c](#), M. Trzebinski [87](#), A. Trzupek [87](#), F. Tsai [145](#),
 M. Tsai [106](#), A. Tsiamis [152,e](#), P.V. Tsiarehka [37](#), S. Tsigaridas [156a](#), A. Tsirigotis [152,s](#),
 V. Tsiskaridze [155](#), E.G. Tskhadadze [149a](#), M. Tsopoulou [152,e](#), Y. Tsujikawa [88](#), I.I. Tsukerman [37](#),
 V. Tsulaia [17a](#), S. Tsuno [84](#), K. Tsurii [118](#), D. Tsybychev [145](#), Y. Tu [64b](#), A. Tudorache [27b](#),
 V. Tudorache [27b](#), A.N. Tuna [61](#), S. Turchikhin [57b,57a](#), I. Turk Cakir [3a](#), R. Turra [71a](#),
 T. Turtuvshin [38,y](#), P.M. Tuts [41](#), S. Tzamarias [152,e](#), P. Tzanis [10](#), E. Tzovara [100](#), F. Ukegawa [157](#),
 P.A. Ulloa Poblete [137c,137b](#), E.N. Umaka [29](#), G. Unal [36](#), M. Unal [11](#), A. Undrus [29](#), G. Unel [159](#),
 J. Urban [28b](#), P. Urquijo [105](#), P. Urrejola [137a](#), G. Usai [8](#), R. Ushioda [154](#), M. Usman [108](#), Z. Uysal [82](#),
 V. Vacek [132](#), B. Vachon [104](#), K.O.H. Vadla [125](#), T. Vafeiadis [36](#), A. Vaitkus [96](#), C. Valderanis [109](#),
 E. Valdes Santurio [47a,47b](#), M. Valente [156a](#), S. Valentinetti [23b,23a](#), A. Valero [163](#),
 E. Valiente Moreno [163](#), A. Vallier [102,ab](#), J.A. Valls Ferrer [163](#), D.R. Van Arneman [114](#),
 T.R. Van Daalen [138](#), A. Van Der Graaf [49](#), P. Van Gemmeren [6](#), M. Van Rijnbach [125,36](#),
 S. Van Stroud [96](#), I. Van Vulpen [114](#), M. Vanadia [76a,76b](#), W. Vandelli [36](#), E.R. Vandewall [121](#),
 D. Vannicola [151](#), L. Vannoli [57b,57a](#), R. Vari [75a](#), E.W. Varnes [7](#), C. Varni [17b](#), T. Varol [148](#),
 D. Varouchas [66](#), L. Varriale [163](#), K.E. Varvell [147](#), M.E. Vasile [27b](#), L. Vaslin [84](#), G.A. Vasquez [165](#),
 A. Vasyukov [38](#), R. Vavricka [100](#), F. Vazeille [40](#), T. Vazquez Schroeder [36](#), J. Veatch [31](#), V. Vecchio [101](#),
 M.J. Veen [103](#), I. Velisek [126](#), L.M. Veloce [155](#), F. Veloso [130a,130c](#), S. Veneziano [75a](#),
 A. Ventura [70a,70b](#), S. Ventura Gonzalez [135](#), A. Verbitskyi [110](#), M. Verducci [74a,74b](#), C. Vergis [24](#),
 M. Verissimo De Araujo [83b](#), W. Verkerke [114](#), J.C. Vermeulen [114](#), C. Vernieri [143](#), M. Vessella [103](#),
 M.C. Vetterli [142,ag](#), A. Vgenopoulos [152,e](#), N. Viaux Maira [137f](#), T. Vickey [139](#), O.E. Vickey Boeriu [139](#),
 G.H.A. Viehhauser [126](#), L. Vigani [63b](#), M. Villa [23b,23a](#), M. Villaplana Perez [163](#), E.M. Villhauer [52](#),
 E. Vilucchi [53](#), M.G. Vincter [34](#), G.S. Virdee [20](#), A. Vishwakarma [52](#), A. Visible [114](#), C. Vittori [36](#),
 I. Vivarelli [146](#), E. Voevodina [110](#), F. Vogel [109](#), J.C. Voigt [50](#), P. Vokac [132](#), Yu. Volkotrub [86a](#),
 J. Von Ahnen [48](#), E. Von Toerne [24](#), B. Vormwald [36](#), V. Vorobel [133](#), K. Vorobev [37](#), M. Vos [163](#),
 K. Voss [141](#), M. Vozak [114](#), L. Vozdecky [94](#), N. Vranjes [15](#), M. Vranjes Milosavljevic [15](#),

M. Vreeswijk¹¹⁴, N.K. Vu^{62d,62c}, R. Vuillermet³⁶, O. Vujanovic¹⁰⁰, I. Vukotic³⁹, S. Wada¹⁵⁷, C. Wagner¹⁰³, J.M. Wagner^{17a}, W. Wagner¹⁷¹, S. Wahdan¹⁷¹, H. Wahlberg⁹⁰, M. Wakida¹¹¹, J. Walder¹³⁴, R. Walker¹⁰⁹, W. Walkowiak¹⁴¹, A. Wall¹²⁸, T. Wamorkar⁶, A.Z. Wang¹³⁶, C. Wang¹⁰⁰, C. Wang¹¹, H. Wang^{17a}, J. Wang^{64c}, R.-J. Wang¹⁰⁰, R. Wang⁶¹, R. Wang⁶, S.M. Wang¹⁴⁸, S. Wang^{62b}, T. Wang^{62a}, W.T. Wang⁸⁰, W. Wang^{14a}, X. Wang^{14c}, X. Wang¹⁶², X. Wang^{62c}, Y. Wang^{62d}, Y. Wang^{14c}, Z. Wang¹⁰⁶, Z. Wang^{62d,51,62c}, Z. Wang¹⁰⁶, A. Warburton¹⁰⁴, R.J. Ward²⁰, N. Warrack⁵⁹, S. Waterhouse⁹⁵, A.T. Watson²⁰, H. Watson⁵⁹, M.F. Watson²⁰, E. Watton^{59,134}, G. Watts¹³⁸, B.M. Waugh⁹⁶, C. Weber²⁹, H.A. Weber¹⁸, M.S. Weber¹⁹, S.M. Weber^{63a}, C. Wei^{62a}, Y. Wei¹²⁶, A.R. Weidberg¹²⁶, E.J. Weik¹¹⁷, J. Weingarten⁴⁹, M. Weirich¹⁰⁰, C. Weiser⁵⁴, C.J. Wells⁴⁸, T. Wenaus²⁹, B. Wendland⁴⁹, T. Wengler³⁶, N.S. Wenke¹¹⁰, N. Wermes²⁴, M. Wessels^{63a}, A.M. Wharton⁹¹, A.S. White⁶¹, A. White⁸, M.J. White¹, D. Whiteson¹⁵⁹, L. Wickremasinghe¹²⁴, W. Wiedenmann¹⁷⁰, M. Wielers¹³⁴, C. Wiglesworth⁴², D.J. Wilbern¹²⁰, H.G. Wilkens³⁶, D.M. Williams⁴¹, H.H. Williams¹²⁸, S. Williams³², S. Willocq¹⁰³, B.J. Wilson¹⁰¹, P.J. Windischhofer³⁹, F.I. Winkel³⁰, F. Winklmeier¹²³, B.T. Winter⁵⁴, J.K. Winter¹⁰¹, M. Wittgen¹⁴³, M. Wobisch⁹⁷, Z. Wolfs¹¹⁴, J. Wollrath¹⁵⁹, M.W. Wolter⁸⁷, H. Wolters^{130a,130c}, E.L. Woodward⁴¹, S.D. Worm⁴⁸, B.K. Wosiek⁸⁷, K.W. Woźniak⁸⁷, S. Wozniowski⁵⁵, K. Wraight⁵⁹, C. Wu²⁰, J. Wu^{14a,14e}, M. Wu^{64a}, M. Wu¹¹³, S.L. Wu¹⁷⁰, X. Wu⁵⁶, Y. Wu^{62a}, Z. Wu¹³⁵, J. Wuerzinger^{110,ae}, T.R. Wyatt¹⁰¹, B.M. Wynne⁵², S. Xella⁴², L. Xia^{14c}, M. Xia^{14b}, J. Xiang^{64c}, M. Xie^{62a}, X. Xie^{62a}, S. Xin^{14a,14e}, A. Xiong¹²³, J. Xiong^{17a}, D. Xu^{14a}, H. Xu^{62a}, L. Xu^{62a}, R. Xu¹²⁸, T. Xu¹⁰⁶, Y. Xu^{14b}, Z. Xu⁵², Z. Xu^{14c}, B. Yabsley¹⁴⁷, S. Yacoob^{33a}, Y. Yamaguchi¹⁵⁴, E. Yamashita¹⁵³, T. Yamashita⁸⁵, H. Yamauchi¹⁵⁷, T. Yamazaki^{17a}, Y. Yamazaki⁸⁵, J. Yan^{62c}, S. Yan¹²⁶, Z. Yan²⁵, H.J. Yang^{62c,62d}, H.T. Yang^{62a}, S. Yang^{62a}, T. Yang^{64c}, X. Yang³⁶, X. Yang^{14a}, Y. Yang⁴⁴, Y. Yang^{62a}, Z. Yang^{62a}, W.-M. Yao^{17a}, H. Ye^{14c}, H. Ye⁵⁵, J. Ye^{14a}, S. Ye²⁹, X. Ye^{62a}, Y. Yeh⁹⁶, I. Yeletsikh³⁸, B.K. Yeo^{17b}, M.R. Yexley⁹⁶, P. Yin⁴¹, K. Yorita¹⁶⁸, S. Younas^{27b}, C.J.S. Young³⁶, C. Young¹⁴³, C. Yu^{14a,14e}, Y. Yu^{62a}, M. Yuan¹⁰⁶, R. Yuan^{62b}, L. Yue⁹⁶, M. Zaazoua^{62a}, B. Zabinski⁸⁷, E. Zaid⁵², Z.K. Zak⁸⁷, T. Zakareishvili¹⁶³, N. Zakharchuk³⁴, S. Zambito⁵⁶, J.A. Zamora Saa^{137d,137b}, J. Zang¹⁵³, D. Zanzi⁵⁴, O. Zaplatilek¹³², C. Zeitnitz¹⁷¹, H. Zeng^{14a}, J.C. Zeng¹⁶², D.T. Zenger Jr²⁶, O. Zenin³⁷, T. Ženiš^{28a}, S. Zenz⁹⁴, S. Zerradi^{35a}, D. Zerwas⁶⁶, M. Zhai^{14a,14e}, D.F. Zhang¹³⁹, J. Zhang^{62b}, J. Zhang⁶, K. Zhang^{14a,14e}, L. Zhang^{14c}, P. Zhang^{14a,14e}, R. Zhang¹⁷⁰, S. Zhang¹⁰⁶, S. Zhang⁴⁴, T. Zhang¹⁵³, X. Zhang^{62c}, X. Zhang^{62b}, Y. Zhang^{62c,5}, Y. Zhang⁹⁶, Y. Zhang^{14c}, Z. Zhang^{17a}, Z. Zhang⁶⁶, H. Zhao¹³⁸, T. Zhao^{62b}, Y. Zhao¹³⁶, Z. Zhao^{62a}, A. Zhemchugov³⁸, J. Zheng^{14c}, K. Zheng¹⁶², X. Zheng^{62a}, Z. Zheng¹⁴³, D. Zhong¹⁶², B. Zhou¹⁰⁶, H. Zhou⁷, N. Zhou^{62c}, Y. Zhou^{14c}, Y. Zhou⁷, C.G. Zhu^{62b}, J. Zhu¹⁰⁶, Y. Zhu^{62c}, Y. Zhu^{62a}, X. Zhuang^{14a}, K. Zhukov³⁷, N.I. Zimine³⁸, J. Zinsser^{63b}, M. Ziolkowski¹⁴¹, L. Živković¹⁵, A. Zoccoli^{23b,23a}, K. Zoch⁶¹, T.G. Zorbas¹³⁹, O. Zormpa⁴⁶, W. Zou⁴¹, L. Zwalinski³⁶

¹ Department of Physics, University of Adelaide, Adelaide, Australia

² Department 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

⁴ LAPP, Université Savoie Mont Blanc, CNRS/IN2P3, Annecy, France

⁵ APC, Université Paris Cité, CNRS/IN2P3, Paris, France

⁶ High Energy Physics Division, Argonne National Laboratory, Argonne IL, United States of America

- ⁷ Department of Physics, University of Arizona, Tucson AZ, United States of America
- ⁸ Department of Physics, University of Texas at Arlington, Arlington TX, United States of America
- ⁹ Physics Department, National and Kapodistrian University of Athens, Athens, Greece
- ¹⁰ Physics Department, National Technical University of Athens, Zografou, Greece
- ¹¹ Department of Physics, University of Texas at Austin, Austin TX, United States of America
- ¹² Institute of Physics, Azerbaijan Academy of Sciences, Baku, Azerbaijan
- ¹³ Institut 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) School of Science, Shenzhen Campus of Sun Yat-sen University; ^(e) University of Chinese Academy of Science (UCAS), Beijing, China
- ¹⁵ Institute of Physics, University of Belgrade, Belgrade, Serbia
- ¹⁶ Department 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
- ¹⁸ Institut für Physik, Humboldt Universität zu Berlin, Berlin, Germany
- ¹⁹ Albert Einstein Center for Fundamental Physics and Laboratory for High Energy Physics, University of Bern, Bern, Switzerland
- ²⁰ School 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, 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
- ²⁴ Physikalisches Institut, Universität Bonn, Bonn, Germany
- ²⁵ Department of Physics, Boston University, Boston MA, United States of America
- ²⁶ Department 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) National University of Science and Technology Politehnica, 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
- ²⁹ Physics Department, Brookhaven National Laboratory, Upton NY, United States of America
- ³⁰ Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Departamento de Física, y CONICET, Instituto de Física de Buenos Aires (IFIBA), Buenos Aires, Argentina
- ³¹ California State University, CA, United States of America
- ³² Cavendish 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
- ³⁴ Department 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
- ³⁶ CERN, Geneva, Switzerland
- ³⁷ Affiliated with an institute covered by a cooperation agreement with CERN
- ³⁸ Affiliated with an international laboratory covered by a cooperation agreement with CERN
- ³⁹ Enrico Fermi Institute, University of Chicago, Chicago IL, United States of America
- ⁴⁰ LPC, Université Clermont Auvergne, CNRS/IN2P3, Clermont-Ferrand, France

- ⁴¹ *Nevis Laboratory, Columbia University, Irvington NY, United States of America*
- ⁴² *Niels 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*
- ⁴⁴ *Physics Department, Southern Methodist University, Dallas TX, United States of America*
- ⁴⁵ *Physics Department, University of Texas at Dallas, Richardson TX, United States of America*
- ⁴⁶ *National Centre for Scientific Research “Demokritos”, Agia Paraskevi, Greece*
- ⁴⁷ ^(a) *Department of Physics, Stockholm University;* ^(b) *Oskar Klein Centre, Stockholm, Sweden*
- ⁴⁸ *Deutsches Elektronen-Synchrotron DESY, Hamburg and Zeuthen, Germany*
- ⁴⁹ *Fakultät Physik, Technische Universität Dortmund, Dortmund, Germany*
- ⁵⁰ *Institut für Kern- und Teilchenphysik, Technische Universität Dresden, Dresden, Germany*
- ⁵¹ *Department of Physics, Duke University, Durham NC, United States of America*
- ⁵² *SUPA - School of Physics and Astronomy, University of Edinburgh, Edinburgh, United Kingdom*
- ⁵³ *INFN e Laboratori Nazionali di Frascati, Frascati, Italy*
- ⁵⁴ *Physikalisches Institut, Albert-Ludwigs-Universität Freiburg, Freiburg, Germany*
- ⁵⁵ *II. Physikalisches Institut, Georg-August-Universität Göttingen, Göttingen, Germany*
- ⁵⁶ *Département 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*
- ⁵⁸ *II. Physikalisches Institut, Justus-Liebig-Universität Giessen, Giessen, Germany*
- ⁵⁹ *SUPA - School of Physics and Astronomy, University of Glasgow, Glasgow, United Kingdom*
- ⁶⁰ *LPSC, Université Grenoble Alpes, CNRS/IN2P3, Grenoble INP, Grenoble, France*
- ⁶¹ *Laboratory 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;* ^(e) *School of Physics and Microelectronics, Zhengzhou University, 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*
- ⁶⁵ *Department of Physics, National Tsing Hua University, Hsinchu, Taiwan*
- ⁶⁶ *IJCLab, Université Paris-Saclay, CNRS/IN2P3, 91405, Orsay, France*
- ⁶⁷ *Centro Nacional de Microelectrónica (IMB-CNM-CSIC), Barcelona, Spain*
- ⁶⁸ *Department 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*
- ⁷⁹ *Universität Innsbruck, Department of Astro and Particle Physics, Innsbruck, Austria*
- ⁸⁰ *University of Iowa, Iowa City IA, United States of America*
- ⁸¹ *Department of Physics and Astronomy, Iowa State University, Ames IA, United States of America*
- ⁸² *Istinye University, Sariyer, Istanbul, Türkiye*
- ⁸³ ^(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*

- 84 KEK, High Energy Accelerator Research Organization, Tsukuba, Japan
- 85 Graduate School of Science, Kobe University, Kobe, Japan
- 86 ^(a) AGH University of Krakow, Faculty of Physics and Applied Computer Science, Krakow; ^(b) Marian Smoluchowski Institute of Physics, Jagiellonian University, Krakow, Poland
- 87 Institute of Nuclear Physics Polish Academy of Sciences, Krakow, Poland
- 88 Faculty of Science, Kyoto University, Kyoto, Japan
- 89 Research Center for Advanced Particle Physics and Department of Physics, Kyushu University, Fukuoka, Japan
- 90 Instituto de Física La Plata, Universidad Nacional de La Plata and CONICET, La Plata, Argentina
- 91 Physics Department, Lancaster University, Lancaster, United Kingdom
- 92 Oliver Lodge Laboratory, University of Liverpool, Liverpool, United Kingdom
- 93 Department of Experimental Particle Physics, Jožef Stefan Institute and Department of Physics, University of Ljubljana, Ljubljana, Slovenia
- 94 School of Physics and Astronomy, Queen Mary University of London, London, United Kingdom
- 95 Department of Physics, Royal Holloway University of London, Egham, United Kingdom
- 96 Department of Physics and Astronomy, University College London, London, United Kingdom
- 97 Louisiana Tech University, Ruston LA, United States of America
- 98 Fysiska institutionen, Lunds universitet, Lund, Sweden
- 99 Departamento de Física Teórica C-15 and CIAFF, Universidad Autónoma de Madrid, Madrid, Spain
- 100 Institut für Physik, Universität Mainz, Mainz, Germany
- 101 School of Physics and Astronomy, University of Manchester, Manchester, United Kingdom
- 102 CPPM, Aix-Marseille Université, CNRS/IN2P3, Marseille, France
- 103 Department of Physics, University of Massachusetts, Amherst MA, United States of America
- 104 Department of Physics, McGill University, Montreal QC, Canada
- 105 School of Physics, University of Melbourne, Victoria, Australia
- 106 Department of Physics, University of Michigan, Ann Arbor MI, United States of America
- 107 Department of Physics and Astronomy, Michigan State University, East Lansing MI, United States of America
- 108 Group of Particle Physics, University of Montreal, Montreal QC, Canada
- 109 Fakultät für Physik, Ludwig-Maximilians-Universität München, München, Germany
- 110 Max-Planck-Institut für Physik (Werner-Heisenberg-Institut), München, Germany
- 111 Graduate School of Science and Kobayashi-Maskawa Institute, Nagoya University, Nagoya, Japan
- 112 Department of Physics and Astronomy, University of New Mexico, Albuquerque NM, United States of America
- 113 Institute for Mathematics, Astrophysics and Particle Physics, Radboud University/Nikhef, Nijmegen, The Netherlands
- 114 Nikhef National Institute for Subatomic Physics and University of Amsterdam, Amsterdam, The Netherlands
- 115 Department of Physics, Northern Illinois University, DeKalb IL, United States of America
- 116 ^(a) New York University Abu Dhabi, Abu Dhabi; ^(b) United Arab Emirates University, Al Ain, United Arab Emirates
- 117 Department of Physics, New York University, New York NY, United States of America
- 118 Ochanomizu University, Otsuka, Bunkyo-ku, Tokyo, Japan
- 119 Ohio State University, Columbus OH, United States of America
- 120 Homer L. Dodge Department of Physics and Astronomy, University of Oklahoma, Norman OK, United States of America
- 121 Department of Physics, Oklahoma State University, Stillwater OK, United States of America
- 122 Palacký University, Joint Laboratory of Optics, Olomouc, Czech Republic
- 123 Institute for Fundamental Science, University of Oregon, Eugene, OR, United States of America
- 124 Graduate School of Science, Osaka University, Osaka, Japan
- 125 Department of Physics, University of Oslo, Oslo, Norway
- 126 Department of Physics, Oxford University, Oxford, United Kingdom
- 127 LPNHE, Sorbonne Université, Université Paris Cité, CNRS/IN2P3, Paris, France
- 128 Department of Physics, University of Pennsylvania, Philadelphia PA, United States of America
- 129 Department 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

- 131 *Institute of Physics of the Czech Academy of Sciences, Prague, Czech Republic*
- 132 *Czech Technical University in Prague, Prague, Czech Republic*
- 133 *Charles University, Faculty of Mathematics and Physics, Prague, Czech Republic*
- 134 *Particle Physics Department, Rutherford Appleton Laboratory, Didcot, United Kingdom*
- 135 *IRFU, CEA, Université Paris-Saclay, Gif-sur-Yvette, France*
- 136 *Santa 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*
- 138 *Department of Physics, University of Washington, Seattle WA, United States of America*
- 139 *Department of Physics and Astronomy, University of Sheffield, Sheffield, United Kingdom*
- 140 *Department of Physics, Shinshu University, Nagano, Japan*
- 141 *Department Physik, Universität Siegen, Siegen, Germany*
- 142 *Department of Physics, Simon Fraser University, Burnaby BC, Canada*
- 143 *SLAC National Accelerator Laboratory, Stanford CA, United States of America*
- 144 *Department of Physics, Royal Institute of Technology, Stockholm, Sweden*
- 145 *Departments of Physics and Astronomy, Stony Brook University, Stony Brook NY, United States of America*
- 146 *Department of Physics and Astronomy, University of Sussex, Brighton, United Kingdom*
- 147 *School of Physics, University of Sydney, Sydney, Australia*
- 148 *Institute 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*
- 150 *Department of Physics, Technion, Israel Institute of Technology, Haifa, Israel*
- 151 *Raymond and Beverly Sackler School of Physics and Astronomy, Tel Aviv University, Tel Aviv, Israel*
- 152 *Department of Physics, Aristotle University of Thessaloniki, Thessaloniki, Greece*
- 153 *International Center for Elementary Particle Physics and Department of Physics, University of Tokyo, Tokyo, Japan*
- 154 *Department of Physics, Tokyo Institute of Technology, Tokyo, Japan*
- 155 *Department of Physics, University of Toronto, Toronto ON, Canada*
- 156 ^(a) *TRIUMF, Vancouver BC;* ^(b) *Department of Physics and Astronomy, York University, Toronto ON, Canada*
- 157 *Division of Physics and Tomonaga Center for the History of the Universe, Faculty of Pure and Applied Sciences, University of Tsukuba, Tsukuba, Japan*
- 158 *Department of Physics and Astronomy, Tufts University, Medford MA, United States of America*
- 159 *Department of Physics and Astronomy, University of California Irvine, Irvine CA, United States of America*
- 160 *University of Sharjah, Sharjah, United Arab Emirates*
- 161 *Department of Physics and Astronomy, University of Uppsala, Uppsala, Sweden*
- 162 *Department of Physics, University of Illinois, Urbana IL, United States of America*
- 163 *Instituto de Física Corpuscular (IFIC), Centro Mixto Universidad de Valencia - CSIC, Valencia, Spain*
- 164 *Department of Physics, University of British Columbia, Vancouver BC, Canada*
- 165 *Department of Physics and Astronomy, University of Victoria, Victoria BC, Canada*
- 166 *Fakultät für Physik und Astronomie, Julius-Maximilians-Universität Würzburg, Würzburg, Germany*
- 167 *Department of Physics, University of Warwick, Coventry, United Kingdom*
- 168 *Waseda University, Tokyo, Japan*
- 169 *Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot, Israel*
- 170 *Department of Physics, University of Wisconsin, Madison WI, United States of America*
- 171 *Fakultät für Mathematik und Naturwissenschaften, Fachgruppe Physik, Bergische Universität Wuppertal, Wuppertal, Germany*
- 172 *Department of Physics, Yale University, New Haven CT, United States of America*

^a Also Affiliated with an institute covered by a cooperation agreement with CERN

^b Also at An-Najah National University, Nablus, Palestine

^c Also at Borough of Manhattan Community College, City University of New York, New York NY, United States of America

^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, Ben Gurion University of the Negev, Beer Sheva, Israel
- ^l Also at Department of Physics, California State University, Sacramento, United States of America
- ^m Also at Department of Physics, King's College London, London, United Kingdom
- ⁿ Also at Department of Physics, Stanford University, Stanford CA, United States of America
- ^o Also at Department of Physics, Stellenbosch University, South Africa
- ^p Also at Department of Physics, University of Fribourg, Fribourg, Switzerland
- ^q Also at Department of Physics, University of Thessaly, Greece
- ^r Also at Department of Physics, Westmont College, Santa Barbara, United States of America
- ^s Also at Hellenic Open University, Patras, Greece
- ^t Also at Institutio Catalana de Recerca i Estudis Avancats, ICREA, Barcelona, Spain
- ^u Also at Institut für Experimentalphysik, Universität Hamburg, Hamburg, Germany
- ^v Also at Institute for Nuclear Research and Nuclear Energy (INRNE) of the Bulgarian Academy of Sciences, Sofia, Bulgaria
- ^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, Mongolian Academy of Sciences, 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 L2IT, Université de Toulouse, CNRS/IN2P3, UPS, Toulouse, France
- ^{ac} Also at Lawrence Livermore National Laboratory, Livermore, United States of America
- ^{ad} Also at National Institute of Physics, University of the Philippines Diliman (Philippines), Philippines
- ^{ae} Also at Technical University of Munich, Munich, Germany
- ^{af} Also at The Collaborative Innovation Center of Quantum Matter (CICQM), Beijing, China
- ^{ag} Also at TRIUMF, Vancouver BC, Canada
- ^{ah} Also at Università di Napoli Parthenope, Napoli, Italy
- ^{ai} Also at University of Colorado Boulder, Department of Physics, Colorado, United States of America
- ^{aj} Also at Washington College, Chestertown, MD, United States of America
- ^{ak} Also at Yeditepe University, Physics Department, Istanbul, Türkiye
- * Deceased