



Development of secondary vertex reconstruction using tracks in the ATLAS Muon Spectrometer for long-lived particle decays

The ATLAS Collaboration

Many promising beyond the Standard Model theories predict invisible long-lived particles that travel macroscopic distances before decaying into visible Standard Model particles with a common spatial origin. Fast and accurate reconstruction of secondary vertices therefore plays a central role in ATLAS long-lived particle searches. While significant advancements have been made in efficiently reconstructing Inner Detector tracks originating far from the primary interaction point, traditional secondary vertexing algorithms still remain limited by the ID track reconstruction acceptance, significantly reducing sensitivity to phase space where the bulk of expected decays lie beyond the Pixel detector. This note describes a novel secondary vertexing technique that leverages standalone muon tracks reconstructed exclusively in the ATLAS Muon Spectrometer, which demonstrates the ability to efficiently reconstruct ultra-displaced dimuon vertices up to 8 meters from the IP. Its performance is characterized on a range of simulated beyond-the-Standard Model signatures and in Run 3 data through a tag-and-probe approach using $J/\psi \rightarrow \mu\mu$ decays.

1 Introduction

Searches for long-lived particles represent some of the most compelling opportunities to discover physics beyond the Standard Model (SM) at the Large Hadron Collider (LHC). Many beyond the SM (BSM) scenarios, such as heavy neutral lepton (HNL) and hidden sector models, predict new invisible states with long proper lifetimes τ that travel macroscopic distances before decaying into SM particles. If these particles are charged, the trajectories can be reconstructed and used to find the decay position. The resulting displaced vertices provide a distinctive experimental signature largely free of SM backgrounds. However, traditional search strategies have historically been optimized under the assumption that new BSM particles decay promptly, close to the primary proton-proton (pp) interaction point, and signals left by such displaced decay products in the detectors have originally been left unexploited. Thus, dedicated reconstruction techniques are required to target these unconventional topologies, introducing new challenges into the conventional analysis paradigms.

Within the ATLAS experiment, charged particle trajectories are traditionally reconstructed through measurements provided by the ATLAS Inner Detector (ID). Energy deposits (hits) recorded in individual detector elements of the ID are used to construct a parameterized description of the trajectory (track) for an individual particle as it travels outward from the central collision point. To efficiently process the ATLAS dataset within computational resource constraints, candidate tracks in the primary reconstruction pass (described in Section 4) are required to have a transverse impact parameter, d_0 , lower than 5 mm. Here, d_0 quantifies the displacement perpendicular to the beam axis, while the longitudinal impact parameter, z_0 , measures the displacement along it. Reconstruction of displaced decays is therefore handled by an additional large-radius tracking (LRT) pass [1], which reprocesses hits left unassigned after the primary pass with loosened pointing requirements, allowing track candidates displaced up to 300 mm in the transverse plane and 500 mm along the beam direction. Event reconstruction proceeds through a sequence of dedicated algorithms that together form the ATLAS reconstruction chain, with track reconstruction constituting its first stage.

To prepare for the more challenging LHC Run 3 data-taking conditions, both the primary and LRT tracking passes were significantly reoptimized, dramatically reducing the computational overhead by up to a factor of two for the primary pass and up to ten for the LRT pass while also considerably improving overall candidate purity [1, 2]. These parallel optimizations allowed LRT reconstruction to be fully integrated into the standard Run 3 ATLAS event reconstruction chain, making the common data streams and central SM background Monte-Carlo (MC) simulation productions fully available to track-based LLP searches for both Run 3 and reprocessed Run 2 data.

The integration of LRT into the primary reconstruction chain represents a monumental achievement, enabling a rich and successful expansion of the ATLAS LLP search program [3–5]. However, by construction, LRT performance remains limited by the finite extent of the silicon detectors. Beyond the outermost layers of the Pixel detector, LLP decay products fail to leave sufficient silicon hits to form reconstructed tracks, and standard ATLAS secondary vertexing algorithms relying on ID track information thus rapidly lose efficiency beyond 300 mm from the IP, limiting the ability to reconstruct decays from particle τ beyond around 10 ns. As a result, ATLAS displaced vertex searches have limited sensitivity to the wide class of LLP scenarios predicting decays outside the ID volume, even if trackless searches for displaced decays exist. There are several such searches like these for long-lived hidden sector particles and axion-like particles resulting in displaced hadronic jets in the ATLAS calorimeter [6], and displaced hadronic vertices in the muon spectrometer (MS) [7]. However the focus of these searches were hadronic decays, and relied on clusters of energy deposits for reconstruction in the calorimeter or an abundance

of soft tracks in the chambers of the MS [8]. Similarly, a search for dark photons decaying to displaced lepton jets in the MS [9] also relied upon identification of the jet-like structure instead of reconstructing a muon vertex. There was also previously a search using the topology of non-collimated muons without ID tracks [10], however no attempt was made to fully reconstruct a vertex from the two muons. To be able to reconstruct a clean muonic MS-based vertex will open a new chapter for the ATLAS LLP search program.

To tackle this challenge, a new secondary vertexing algorithm has been developed that leverages purely information from the MS, which provides an alternative source of tracking information over a radial extent of several meters, far exceeding the total ID acceptance for displaced dimuon signatures. Standard ATLAS muon reconstruction algorithms first reconstruct MS track candidates through linking compatible segments built from precision chamber hits, which are then systematically combined with a corresponding compatible ID track to form a Combined muon (CB) candidate. Remaining high-quality MS tracks not ultimately matched to an ID track are classified as StandAlone (SA) and remain available for downstream analysis use. As ID track and CB muon reconstruction are both extremely efficient, SA muons within the $|\eta| < 2.5$ range of ID acceptance are very rare from SM processes while occurring naturally at a high rate from LLP decays that take place beyond the ID volume. This observation motivates the development of a dedicated secondary vertexing framework based on SA muons, enabling reconstruction of displaced dimuon vertices in the MS independent of ID information.

This note is organized as follows: a brief description of the ATLAS detector is given in Section 2. Section 3 describes the MC simulation and data samples used to evaluate the performance. Section 4 reviews the formalism of ATLAS ID and MS track reconstruction alongside standard secondary vertex reconstruction techniques. Section 5 describes the the SA muon vertexing algorithm. Performance results on simulation and data are presented in Sections 6 and 7, respectively, and concluding remarks are made in Section 8.

2 The ATLAS detector

The ATLAS detector [11] at the Large Hadron Collider (LHC) covers nearly the entire solid angle around the collision point.¹ It consists of an inner tracking detector surrounded by a thin superconducting solenoid, electromagnetic and hadronic calorimeters, and a muon spectrometer incorporating three large superconducting air-core toroidal magnets.

The ID 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 comprises 4 concentric barrel layers and 3 endcap disks on each side. It covers the collision region and typically provides four measurements per track, the first hit normally being in the insertable B-layer (IBL), which was installed before Run 2 [12, 13]. It is followed by the silicon microstrip SemiConductor Tracker (SCT), which is comprised of four barrel layers and 9 planar endcap disks per side, with pairs of modules mounted at a small stereo angle allowing for two measurements per traversed module layer. These silicon detectors are complemented by the transition radiation tracker (TRT), which enables radially extended track reconstruction up to $|\eta| = 2.0$. The TRT

¹ ATLAS uses a right-handed coordinate system with its origin at the nominal interaction point (IP) in the center of the detector and the z -axis along the beam pipe. The x -axis points from the IP to the center 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 $\eta = -\ln \tan(\theta/2)$ and is equal to the rapidity $y = \frac{1}{2} \ln \left(\frac{E+p_z}{E-p_z} \right)$ in the relativistic limit. Angular distance is measured in units of $\Delta R \equiv \sqrt{(\Delta y)^2 + (\Delta \phi)^2}$.

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 calorimeter system [14–16] covers the pseudorapidity range $|\eta| < 4.9$. Within the region $|\eta| < 3.2$, electromagnetic calorimetry is provided by barrel and endcap high-granularity lead/liquid-argon (LAr) calorimeters, with an additional thin LAr presampler covering $|\eta| < 1.8$ to correct for energy loss in material upstream of the calorimeters. Hadronic calorimetry is provided by the steel/scintillator-tile calorimeter, segmented into three barrel structures within $|\eta| < 1.7$, and two copper/LAr hadronic 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.

The muon spectrometer (MS) [17], as detailed in Figure 1, comprises separate trigger and high-precision tracking chambers measuring the deflection of muons in a magnetic field generated by the superconducting air-core toroidal magnets. The field integral of the toroids ranges between 2.0 and 6.0 Tm across most of the detector. Three layers of precision chambers, each consisting of layers of monitored drift tubes (MDTs), cover the region $|\eta| < 2.7$, except in the innermost layer of the endcap region, where layers of small-strip thin-gap chambers and Micromegas chambers both provide precision tracking in the region $1.3 < |\eta| < 2.7$, making up the New Small Wheel (NSW). The muon trigger system covers the range $|\eta| < 2.4$ with resistive-plate chambers in the barrel, thin-gap chambers in the endcap regions, and the aforementioned small-strip thin-gap chambers and Micromegas chambers in the innermost layer of the endcap.

The luminosity is measured mainly by the LUCID-2 [19] detector that records Cherenkov light produced in the quartz windows of photomultipliers located close to the beam pipe.

Events are selected by the first-level trigger system implemented in custom hardware, followed by selections made by algorithms implemented in software in the high-level trigger [20]. The first-level trigger accepts events from the 40 MHz bunch crossings at a rate below 100 kHz, which the high-level trigger further reduces in order to record complete events to disk at about 1 kHz.

A software suite [21] is used in data simulation, in the reconstruction and analysis of real and simulated data, in detector operations, and in the trigger and data acquisition systems of the experiment.

3 Data and simulated samples

To analyze the performance of the SA muon vertexing algorithm on BSM physics of interest, several representative BSM scenarios are studied using Monte-Carlo (MC) simulations. These scenarios are both characterized by the production of long-lived neutral particles which subsequently decay into a dimuon pair, with proper lifetimes resulting in the majority of decays appearing outside of ID acceptance. These benchmark scenarios are chosen to sample an inclusive range of event topologies that impact the reconstructed vertex properties and kinematics.

The first class of signals considered is a displaced HNL, which allows the study of the performance in boosted, highly collimated decay topologies common to the low mass regime. This signal is motivated by theories which provide an explanation to the origin of the neutrino mass through introducing a theoretical massive right-chiral counterpart to the SM left-chiral neutrino [22–24]. The small mixing between SM neutrinos and the HNL (N) results in a long proper lifetime for the BSM particle. Here, the signal is produced through a W boson decay into a muon and neutrino which mixes with the N , which subsequently




Figure 1: Cross-section of the ATLAS detector specifying the positions of the different subsystems of the MS. A potential muon track and potential fake muon are shown, as well as the direction of the magnetic field in the endcap toroid. Taken from [18].

decays into a dimuon pair and a muon neutrino via and off-shell W or Z boson, as shown in Figure 2(a). The mass of the N is 3 GeV, with mean proper decay lengths of 100 and 1000 mm. Samples are generated at next-to-leading-order (NLO) accuracy in QCD using MADGRAPH5_AMC@NLO 3.3.1 [25] and the SM HEAVYN NLO model libraries [26]. The NNPDF3.0NLO [27] parton distribution functions (PDFs) are used. Events are interfaced to PYTHIA 8.307[28] for parton showering, hadronization, and underlying event modelling, using the A14 tune. The leptonic decays of the HNL are handled by MADSPIN [29].

The second class of signals considered is a displaced dark photon (Z_d) produced in the context of the Hidden Abelian Higgs Model (HAHM) [30]. In this model, an additional $U(1)_d$ gauge symmetry is introduced in the hidden sector, which kinematically mixes with the SM hypercharge field through a small mixing parameter. The resulting Z_d can be produced through exotic Higgs boson decays via $H \rightarrow ZZ_d$, with subsequent decays into a dimuon pair as shown in Figure 2(b). The proper lifetime of the Z_d depends on the mixing strength, and can be sufficiently long for the decay to occur outside the ID volume. In contrast to the HNL benchmark, this signature probes the performance in the high-mass, large opening-angle regime, characteristic of more energetic decays with well-separated muon trajectories. Two LLP masses, 20 and 50 GeV, are considered, both with the mixing parameter set to enforce a proper lifetime of 10 ns. Samples are generated using MADGRAPH5_AMC@NLO 3.3.1 at leading-order (LO). The production is simulated for gluon-gluon fusion at 13.6 TeV using the NNPDF2.3LO PDFs [31] and the A14 tune [32].

To evaluate the effectiveness of the algorithm on genuine data, simulated $b\bar{b} \rightarrow J/\psi$ events are used




Figure 2: Feynman diagrams for the benchmark LLP models considered in this note. Figure (a) shows the HNL production and decay processes with muon-only couplings. The equivalent process with a virtual Z boson instead of a virtual W boson current decays is also considered. Figure (b) shows the HAHM ZZ_d production.

in a tag-and-probe approach where their yields are compared to data. These samples are generated using PYTHIA 8B, a PYTHIA 8 based wrapper adapted to allow forced gluon decays to b-quarks, using the CTEQ6L1 set of parton distribution functions [33] and the A14 set of tuned parameters.

A data sample recorded during 2024 is used to compare the SA muon vertexing performance between simulation and data. These are selected from the “B-physics delayed” stream, a specialized stream separate from the main physics dataset where data is collected using dedicated low p_T triggers tuned for b -physics analyses. This stream is not reconstructed promptly at the Tier-0 facility during data-taking and is instead saved in raw form for delayed processing when computational resources become available [34]. The total number of recorded events is 39,502,447, during LHC fill number 10084.

For all MC samples, detector response is simulated with a full GEANT4-based model of the ATLAS detector, including additional pp collisions in the same bunch crossing (pileup) that are overlaid with the signal process. Samples are generated with a flat pileup range between 0 and 100. In the case of comparison with data events, they are reweighted to match the pileup distribution with a mean value of 64.

4 Inner detector track and vertex reconstruction

Charged particle trajectories in ATLAS are parameterized by a five-component state vector $(d_0, z_0, \phi, \theta, q/p)$, defined relative to a reference axis aligned with the global z -direction of the detector and centered on the beamspot. The beamspot denotes the luminous region, i.e. the time-averaged centroid and spatial profile of pp interactions. This parameterization corresponds to the *perigee representation*, in which the track parameters are expressed at the point of closest approach of the reconstructed helix to the beamline, the nominal trajectory of the colliding proton beams through the detector. Here, d_0 and z_0 denote the transverse and longitudinal impact parameters with respect to the beamline, ϕ and θ represent the azimuthal and polar angles, and q/p corresponds to the inverse momentum multiplied by the charge. The objective of track reconstruction is to identify groups of detector measurements originating from a single charged particle and

to determine the optimal estimates of its trajectory parameters. Each such set of associated hits, together with its fitted track parameters and covariance matrix, defines a reconstructed track object.

The ATLAS ID track reconstruction for both prompt and displaced tracks in Run 2 and Run 3 has been extensively documented in other sources [1, 2, 35]. The following section provides a brief overview for the reader.

4.1 Primary and secondary pass

Charged-particle reconstruction in the ATLAS ID [2] begins with a data preparation phase in which neighboring readout channels from the Pixel and SCT subsystems are merged into clusters representing individual particle crossings. Each pixel cluster or pair of one-dimensional SCT clusters from the two sides of a module is then converted into a three-dimensional space point, with position uncertainties derived from the detector geometry and the sensor pitch. For SCT space points, the missing coordinate is reconstructed from the stereo angle between the two strip orientations. Tracks are reconstructed for charged particles with $p_T > 500$ MeV and $|\eta| < 2.5$; these limits are set respectively by the curvature of low-momentum particles in the magnetic field and by the instrumented acceptance of the ID.

The primary reconstruction sequence, often referred to as the inside-out pass, is optimized for tracks originating near the primary interaction point. It begins by forming triplets of compatible space points (“seeds”) in the silicon detectors that could originate from a common charged particle. Using the estimated seed trajectory, “search roads” are defined through successive detector layers, identifying additional clusters consistent with the expected path. Track candidates are assembled and refined with a combinatorial Kalman filter [36], which updates the track parameters and covariance as compatible clusters are incorporated.

To resolve ambiguities among overlapping candidates and suppress fakes, a dedicated ambiguity solver evaluates all candidates according to multiple quality metrics. Tracks sharing many hits with higher-scoring candidates are rejected, while limited shared hits are allowed to preserve efficiency in dense environments such as the cores of energetic jets. Clusters consistent with multiple crossings are shared amongst tracks using a neural-network-based algorithm that estimates the most probable particle multiplicity and positions [35].

Following ambiguity resolution, the surviving candidates are refitted with a global χ^2 minimization to obtain high-precision track parameters. Tracks failing convergence are discarded. An optional extension into the TRT is then attempted by associating drift circles near the extrapolated trajectory and refitting the full track including the TRT measurements, improving momentum resolution. To guard against spurious extensions, only tracks with a sufficient fraction of “precision” TRT hits, where the fitted trajectory agrees with the measured drift radius within 2.5 times the drift-radius uncertainty, are retained. For tracks with at least 15 TRT hits, a precision hit fraction exceeding 30% is required for the extension to be accepted. The probability of a successful extension depends on occupancy but typically exceeds 90% for muons traversing the full ID.

Because the inside-out configuration targets prompt tracks, an additional secondary pass—implemented as a back-tracking procedure—is applied to recover slightly displaced tracks such as those from photon conversions using hits left unassigned in the primary pass. Regions of interest are first identified by energy deposits in the electromagnetic calorimeter. Within those regions, TRT segments are built from groups of compatible drift circles and then used to seed short silicon tracks based on two space points in the Pixel

and SCT detectors. These are extended using the same combinatorial Kalman filter and subjected to an independent ambiguity resolution and refit sequence, including possible TRT extensions.

4.2 Large radius pass

Since the primary and secondary tracking passes were designed to be efficient and highly pure for prompt and slightly displaced decays, they are less than effective at reconstructing tracks from more displaced decays. Thus another secondary Large Radius Tracking (LRT) pass was developed to be complementary to the previous tracking passes [1]. LRT uses unassigned hits that have been left over from the previous passes, and proceeds with dramatically loosened requirements on the transverse impact parameter, d_0 , from 5 mm to 300 mm, and on the longitudinal impact parameter, z_0 , from 200 mm to 500 mm. Additionally, while the seeding for the primary tracking pass may be either from the pixel or SCT detectors, the LRT uses SCT-only seeding, due to the displaced topology of the tracks from displaced decays which typically do not result in three pixel hits. While the impact parameter requirements are extremely relaxed, the minimum p_T threshold for the tracks is increased from 0.5 to 1 GeV for the LRT pass to reduce the background from soft particles such as those from interactions with the detector material. The search roads are also narrowed to reduce the number of poor-quality tracks reconstructed. The LRT track collection provides a significant sensitivity increase for displaced charged particle searches, but has significantly limited efficiency beyond 300 mm in the transverse plane due to the eight minimum silicon hits required to form a track candidate. These requirements are nevertheless essential to maintain an acceptable level of purity and reconstruction latency.

4.3 ID secondary vertex reconstruction

Traditional secondary vertexing algorithms in ATLAS, including those commonly used for LLPs [37], primarily use the VKALVRT[38] fitting algorithm, which provides a general framework for fitting the common vertex position of a set of reconstructed charged particle tracks under arbitrary geometric and kinematic constraints. VKALVRT implements a full three-dimensional least-squares minimization of the track-to-vertex residuals using an adaptive iterative procedure that accounts for the helical track model, material effects, and magnetic field inhomogeneities.

Input tracks are provided in the standard ATLAS perigee representation accompanied by their full 5×5 covariance matrices. The algorithm linearizes the track equations around an initial vertex seed, typically derived from the fast intersection of two or more tracks assuming a perfect helix in a constant magnetic field, and iteratively refines both the track and vertex parameters until convergence. The minimization employs analytic derivatives of the track model with respect to the vertex coordinates and particle momentum components, allowing fast evaluation of the residuals and covariance propagation. The fit is terminated once all parameters converge within predefined tolerances or the maximum iteration count is reached.

Multiple constraints can be applied during the fit, including mass requirements for resonant decays, enforcing consistency between the secondary vertex and the PV, or imposing common-vertex requirements across multiple ("cascade") decay chains. During each iteration, tracks inconsistent with the current vertex hypothesis can be excluded or down-weighted according to configurable thresholds, improving robustness against outliers and fake tracks. The resulting covariance matrix of the fitted vertex is obtained from the inverse of the final derivative matrix, providing accurate uncertainty estimates on the fitted vertex position and reconstructed invariant mass.




Figure 3: Simplified overview of the primary outside-in combined tracking chain and the secondary inside-out chain used in ATLAS muon reconstruction. The primary chain begins in the MS, where track candidates are built and then combined with ID tracks. The secondary chain starts from ID tracks and attempts to recover muons missed by the primary chain by extending these tracks through the calorimeters and the MS. Recovered muons can be identified using calorimeter- or segment-tagging, or promoted to inside-out combined candidates if sufficient compatible MS measurements are found. [2]

In standard ATLAS reconstruction workflows, VKALVRT serves as the core vertex-fitting engine for secondary vertex algorithms, not only in LLP searches but also for photon conversions, hadronic interactions, and heavy-flavor decays. Its modular implementation allows downstream algorithms to adapt their configuration to different reconstruction contexts, although it has not previously been applied to tracks from the Muon Spectrometer.

5 Muon reconstruction

An overview of the muon reconstruction workflow [39] is shown in Figure 3. Two main track-finding chains are employed. The primary one, referred to as the outside-in chain, starts from tracks reconstructed in the MS and attempts to combine them with tracks in the ID. There is also a secondary chain aiming to recover very low-momentum muons that starts from ID tracks and searches for hits and segments in the MS; however, since this chain is not relevant here it will not be further described.

The first step in the primary chain is standalone muon reconstruction, which builds tracks using only MS hits. The curvature of the reconstructed track in the toroidal magnetic field provides a measurement of the muon momentum. Because the magnetic field also affects the drift of ionization electrons, Lorentz angle corrections must be applied to the drift paths.

The MS experiences significant background radiation, particularly in the high- η region, leading to additional hits in the millions of readout channels. To maintain a muon track reconstruction efficiency of $\geq 99\%$ while keeping hit combinatorics manageable, advanced filtering techniques are applied early on in the reconstruction to suppress background activity.

Standalone reconstruction proceeds through the following steps:

- **Data preparation:** The raw detector data are calibrated. For the Micromegas (MM) and small-strip Thin Gap Chamber (sTGC) detectors, adjacent channels with signals above threshold are clustered.

- **Pattern finding:** A coarse pattern recognition based on a Hough transform is applied to retain high signal efficiency while rejecting large amounts of background. A weak pointing constraint toward the interaction point further improves background rejection.
- **Segment fitting:** For each muon station, the hits selected in the pattern-finding step are fitted with straight lines to form segments. In the New Small Wheel (NSW), segments are built from both sTGC and Micromegas hits, with an interaction-point constraint applied in both the bending and non-bending directions. In the other stations, the precision coordinate is measured by the Monitored Drift Tubes (MDTs), while the Resistive Plate Chambers (RPCs) and Thin Gap Chambers (TGCs) provide the non-precision coordinate. In the inner barrel and outer endcap stations, only MDTs are available, allowing the reconstruction of 2D segments only.
- **Track finding and fitting:** Compatible muon segments are combined to form track seeds. Different strategies are used to accommodate the varying detector coverage. A global track fit is then performed, accounting for energy loss, multiple scattering, magnetic field variations, and alignment effects. Finally, an ambiguity resolution step is applied before building the final MS track collection.

The reconstructed MS track is then extrapolated toward the beam pipe using the beamspot as a constraint. These tracks are referred to as Muon Spectrometer Only Extrapolated (MSOE) tracks. If an ID track is found to be compatible with an MSOE track, the hits from both are refitted in a common fit to form a combined muon track. The resulting muon objects are assigned an author attribute that indicates the type of track used. Combined tracks are identified as Combined Muons (CB), while MSOE tracks without matching ID tracks are labelled as Muon Standalone (MuSA), referred to here as SA muons.

6 Standalone muon vertexing

6.1 Seed muon selection

SA Muons are first preselected for the formation of seed pairs. As the overwhelming majority of events from background SM processes contain one or fewer SA muons, the selection is designed to be as inclusive as possible to allow for efficient reconstruction of dimuon vertices across a wide range of highly displaced signatures without a significant increase in SM contamination.

As the MS is instrumented up to $|\eta| = 2.7$, SA muons outside of the ID acceptance range of $|\eta| < 2.5$ but before $|\eta| < 2.7$ can result from promptly decaying processes. To focus on candidates that fall inside the ID acceptance—and therefore could in principle acquire an ID track if one were reconstructed—seed muon candidates are required to satisfy $|\eta| < 2.5$. To ensure high-quality tracks, two additional cleaning requirements are also implemented. To effectively capture low-mass topologies, no additional minimum p_T selection is imposed beyond the nominal thresholds that arise naturally from the detector geometry and magnetic field, which correspond to about 2.5 GeV once segment reconstruction and calorimeter energy-loss effects are considered. However, candidates are required to have a p_T below 13 TeV to remove rare cases where the muon momentum determination fails. Similarly, candidates are also required to be within regions of the detector where the expected magnitude of the magnetic field integral is above 0.1 T·m. Muons failing this requirement largely originate at high η behind the endcap toroid, and as such this selection both suppresses contributions from beam-induced background and prevents rare instances of undefined behavior in the vertex fit procedure.

ATLAS muon reconstruction is performed separately for prompt track collections and those reconstructed using LRT tracks, resulting in two distinct collections that are subsequently combined to produce an inclusive muon collection. In some cases, high quality MS tracks not matched to an ID track during the prompt pass and classified as SA muons are “recovered” as full CB muons in the secondary LRT muon pass. As the BSM signatures of interest are characterized by the explicit lack of an ID track, a dedicated overlap-removal step removes SA muons from the seeding pool if their MS tracks are also used in the LRT muon collection. Any remaining SA muons passing the above criteria are included in the seeding pool.

6.2 Muon extrapolation

Widely-used secondary vertex fitting algorithms, such as VKALVRT [38], are designed and developed under the implicit assumption that track parameters are expressed at the point of closest approach to the beamline, which is the default representation of ID tracks in the ATLAS software frameworks. Conversely, SA muon track parameters are defined at the inner surface of the Muon Spectrometer, i.e. the interface between the calorimeters and the MS. These representations, on their own, are purely geometrical – they do not retain information about individual measurements used in their creation and instead both parameterize a description of their reconstructed trajectory in coordinates most generally useful for their respective downstream applications. During the standard CB muon reconstruction sequence, the ID and MS tracks are matched and refitted at this MS inner surface, where the extrapolation uncertainties are smallest. The extrapolation of the MS track towards the beamline is performed using a common tool used throughout ATLAS reconstruction [40], which provides a unified framework for the numerical propagation of charged particle parameters and their full covariance through the magnetic field and the intervening detector and support material volumes, used throughout the ATLAS reconstruction workflow. The tool combines field integration, geometry navigation, and material-effects modeling – including energy loss, multiple scattering, and bremsstrahlung – within a modular Runge-Kutta-Nystrom propagation scheme [41]. This approach ensures consistent covariance transport between detector subsystems and allows precise determination of the predicted track state and associated uncertainties at the perigee surface. The resulting parameterization at the perigee surface is the MSOE track.

By default, this associated MSOE track remains available for downstream analysis use, but the extrapolation assumes a priori that the charged particle producing the MS track originated close to the primary vertex (PV) and incorporates a soft constraint including the PV location as an additional “pseudomeasurement”, rendering it unsuitable for highly displaced signatures. However, without loss of generality, this procedure can be simply repeated on demand without the constraint, yielding an “emulated” track object that is parametrically equivalent to a real ID track reconstructed without pointing requirements or the geometric constraints of the ID detector volume. These extrapolated tracks can then be used as inputs to a subsequent secondary vertex fit, with an associated loss of precision and resolution following from the single-detector information and material propagation contributions to the parameter covariances.

If two or more SA muons remain in the seeding pool, such that a seed pair satisfying the above selection criteria can be formed, this custom re-extrapolation is attempted for each, and any successful result saved for further processing. To ensure physical quantities and convergence in the vertex fit procedure, extrapolated SA muons resulting in $|d_0| > 8$ m or $|z_0| > 10$ m are removed from consideration.

6.3 Vertex finding

The vertex finding begins by forming all possible combinatorial pairs of extrapolated SA muon tracks remaining in the collection. For each pair, a dedicated vertex fit is considered using a modified implementation of the VKALVRT algorithm with increased iteration count, relaxed tolerances for fit convergence, and updated boundary volumes consistent with the MS dimensions. This algorithm minimizes the χ^2 of the vertex, penalizing large differences of the extrapolated track parameters with respect to the regressed vertex position weighted by the respective parameter covariances. Both tracks in the fit are assigned equal weights, and the minimization occurs via a Kalman filter with the incrementing step corresponding to the addition of a track, after an initial fast fit approximation estimates a coarse vertex position. After all pairs have been attempted, any successful vertex fit with a χ^2 greater than 50 is rejected.

The overall rarity of SA muons from background SM processes results in a very small number of total events where more than one SA muon vertex candidate is successfully found. However, to ensure a unique track-to-vertex assignment, events with more than one identified vertex candidate undergo an additional overlap removal procedure. Vertex candidates are sorted by χ^2_{vtx} , and beginning from the best fit (lowest χ^2) vertex, contributing muons are subsequently removed from each other vertex in the collection, until no vertices remain that share a single SA muon. This procedure improves correct vertex identification and purity for LLP signatures expecting more than one displaced dimuon vertex while otherwise having minimal impact.

Remaining SA muon vertices passing all requirements are saved along with the constituent extrapolated tracks used in their formation. These tracks are then subsequently propagated to the associated secondary vertex location and their perigee parameters re-expressed to obtain a more accurate estimate of the geometric relationship between the dimuon pairs. A diagram of SA muon vertex reconstruction is shown in Figure 4.

7 Performance on LLP simulation

7.1 Truth-matching in simulation

To evaluate the performance of the vertexing algorithm, it is necessary to match a generated particle in simulation to a reconstructed SA muon. The relatively low occupancy and high purity of reconstructed SA muons permit the use of a simplified matching criterion, where muons are considered truth matched to a decay if they are reconstructed from an MS segment associated with one of the LLP decay products. In the context of this work, a vertex is only considered truth matched if *both* reconstructed muons are truth matched to the target LLP decay, ensuring that combinatorial background vertices formed from unrelated SA tracks are excluded from the efficiency and resolution studies. For the efficiency measurements below, this matching is evaluated only for truth LLP decays that satisfy the reconstructible selection defined in this subsection.




Figure 4: Schematic diagram of SA muon vertex reconstruction.

7.2 Reconstruction efficiency

The total vertex reconstruction efficiency is defined as the probability for a truth LLP decay producing two muons to yield a reconstructed SA muon vertex that is truth matched according to the criteria above, evaluated for each model described in Section 3. The efficiency is factorized into independent terms to isolate the effects of MS tracking and vertexing performance and provide a holistic evaluation of the performance in the various stages. This treatment is adapted from the approach used in Ref. [37]. Unless otherwise stated, all quantities below are evaluated after applying the reconstructible truth-decay selection introduced in this subsection.

To correctly reconstruct a secondary vertex, both constituent tracks from the decay must be reconstructed as SA muons. Thus, generated signal particles must first be *reconstructible*. A reconstructible truth decay vertex in the context of SA muon reconstruction is defined as one satisfying the following set of fiducial requirements:

- The transverse distance from the origin must be less than 8 m.
- The z position must satisfy $|z| < 10$ m.
- Both muons in the decay chain must satisfy $p_T > 1$ GeV.

The first two requirements ensure that the LLP decays before the second MS station, while the final requirement ensures that both muons have enough momentum to be reconstructed by the muon reconstruction algorithms. The *acceptance* ($\mathcal{A}$) is defined as the ratio of truth LLP decays with at least two reconstructed SA muons to the number of reconstructible truth decays.

The *total efficiency* (ϵ^{total}) corresponds to the ratio of reconstructed vertices to the number of reconstructible truth decays in a given fiducial volume. To decouple the vertexing efficiency from the SA muon reconstruction efficiency, the *algorithmic efficiency* (ϵ^{alg}) is defined as the ratio of reconstructed vertices to the number of truth LLP decays with at least two reconstructed SA muons, corresponding to the sample counted in $\mathcal{A}$. This quantity isolates the performance of the vertexing stage once suitable SA muon tracks are available. The algorithmic efficiency can be further broken down into the *seeding efficiency* (ϵ^{seed}) and *core efficiency* (ϵ^{core}). The seeding efficiency is defined as the ratio of the number of truth LLP decays with at least two seed-selected tracks passing the requirements in Section 6 to the number of truth LLP decays contributing to the acceptance, and the core efficiency then follows as the ratio of reconstructed vertices to the number of truth LLP decays with at least two seed-selected tracks. The total efficiency is then given by

$$\epsilon^{\text{total}} = \mathcal{A} \cdot \epsilon^{\text{alg}} = \mathcal{A} \cdot \epsilon^{\text{seed}} \cdot \epsilon^{\text{core}} \quad (1)$$

Thus, the total efficiency is broken down into constituent components, with the acceptance term accounting for the SA muon reconstruction efficiency, the seed term for the SA muon selection efficiency, and the core term for the pure SA muon vertex reconstruction efficiency.

The total reconstruction efficiency and the component terms are shown in Figure 5 as a function of the truth transverse decay length L_{xy} , defined as the transverse displacement between the production and decay points. The evolution of the efficiency with increasing decay length depends strongly on the overall decay topology is driven by a number of complex factors, with up to 45% efficiency for the HNL signatures and 32% for the Z_d . For all signatures, acceptance is minimal before 300 mm, where ID tracks are reconstructed with high efficiency leaving few available opportunities for SA muon formation – these decays instead lie within the acceptance of traditional vertexing algorithms. The small accepted population below 300 mm mostly comprises decays more forward than $|\eta| = 2.5$, which lie outside of ID acceptance even before 300 mm, and decays where one product would be reconstructed with a d_0 beyond 300 mm, resulting in a failure in LRT track formation. Both scenarios are more common for the high mass Z_d decays, which have more forward topologies and larger overall opening angles, contributing to higher average expected d_0 values. Similarly, acceptance for the Z_d scenario rises sharply after 300 mm and then slowly decays at higher L_{xy} due to implicit loose pointing requirements built into muon segment formation – extremely large d_0 values from wide decays close to the MS entrance reduce the likelihood that an MS track will be formed. In contrast, light and highly boosted HNL decays remain within the approximate pointing acceptance for all lifetimes. The acceptance increases moderately with L_{xy} as decays occur closer to the MS entrance and lose less energy in the calorimeter material, but then drops sharply around 7 m once decays occur too deep inside the MS volume to form a track.

Seeding efficiency remains around 90% for all signatures beyond 300 mm – the small fraction passing acceptance before 300 mm are further removed largely by the $|\eta| < 2.5$ selection requirement. The core vertexing efficiency remains approximately flat around 85% for the HNL topologies, while the Z_d has similar efficiency until 1000 mm and then drops to approximately 50%. While large opening-angle topologies are generally easier for the vertex-fitting algorithm itself, the initial extrapolation to the perigee surface is more likely to break down at large angles. This coincides with the region where the extrapolator switches propagation strategy in the dense calorimeter material beyond the TRT outer radius. The 100 mm HNL case seems slightly more efficient than the 1000 mm case, but this is simply due to differences in the average z displacement.




Figure 5: Shown are the (a) acceptance, (b) seed, (c) core and (d) total vertex reconstruction efficiency as a function of truth LLP decay radius for the four benchmark models.

7.3 Position resolution

To evaluate the vertex position resolution, the residual of the reconstructed position versus the truth LLP decay position is evaluated. Figure 6 shows the full distribution of residuals across all LLP decay radii. The distributions are approximately symmetric – in the case of the 100 mm lifetime HNL, a slight asymmetry is observed as there is a slightly higher population of events at the very positive residual tail. This is largely due to slight accumulated biases in the extrapolation and fitting procedure – for the light collimated signature it is more likely to overestimate distance from the IP, leading to a slight overestimation of opening angle and subsequently higher masses.

The position resolution is dominated by the MDT measurement resolution and extrapolation uncertainties, as well as the decay topology features such as the opening angle. For very collimated tracks, such as those from HNL decays, the combination of correlated track-parameter uncertainties and reduced geometric leverage from the small opening angle creates a challenging environment for the vertex fit, leading to weaker position resolution. In contrast, while efficiency is reduced by the MS track reconstruction requirements, the large opening angle of the higher mass Z_d leads to much better resolution. Resolution for each LLP




Figure 6: Residuals between the reconstructed vertex position and the decay position of the simulated particle, shown in the (a) transverse (L_{xy}) and (b) longitudinal (z) directions for the four benchmark models.

scenario per bin of L_{xy} is shown in Figure 7. The resolution is evaluated as the root-mean-square of the transverse and longitudinal vertex position residuals in each L_{xy} bin, after removing the upper and lower 5% of values, in order to reduce the impact of extreme outliers and provide a more representative measure of performance.




Figure 7: The resolutions (a) σ_{xy} and (b) σ_z as a function of the truth LLP decay length L_{xy}^{truth} for the four benchmark models.

The transverse resolution remains roughly flat across L_{xy} , between 200-500 mm for the HNL and 20-50 mm for the Z_d , until approximately 5 m from the IP. For decay radii beyond 5 m, decays begin to occur within the MS volume itself and thus the number and quality of MS measurements per track begins to degrade. Resolution is more coarse along z , with increases at lowest truth L_{xy} for the boosted HNL signatures due to the near-collinear dimuon topology. When both muons are nearly parallel to the beamline, the resulting shallow angle of intersection leaves the vertex position more poorly constrained in the longitudinal direction.

8 Mass resolution

The mass resolution is evaluated through a direct examination of the reconstructed dimuon-vertex mass, m_{DV} , shown in Figure 8 for the four benchmark models. Performance is overwhelmingly dominated by the opening angle of the decay topology, with weaker performance for the collimated HNL signature and stronger performance for the high mass Z_d . The mass distribution for the HNL is naturally smeared due to the missing presence of the neutrino in the decay mode – however, the resolution for the 3 GeV benchmark signals is approximately 500 MeV. The small shift between the 100 mm and 1000 mm distributions reflects slight biases in the extrapolation routine for the different decay lengths. Resolution for the Z_d is proportionally finer compared to the simulated mass value, between 400 to 600 MeV for the 20 and 50 GeV signals, respectively.




Figure 8: Reconstructed dimuon-vertex mass (m_{DV}) distributions for the (a) HNL and (b) Z_d benchmark models.

9 Performance on data

The inherent strengths of the SA muon vertexing, by necessity, introduce unique challenges to characterization of its performance on real data. There are no known SM signatures which produce highly displaced dimuon vertices – while the algorithm would in principle be sensitive to average decay lengths characteristic of K_S^0 or K_L^0 , there are no exclusive dimuon decay channels available.² A creative approach is therefore required to design a suitable proxy.

This work adopts a tag-and-probe strategy based on prompt $J/\psi \rightarrow \mu^+ \mu^-$ decays, providing a high-purity source of dimuons in data with well-understood kinematic properties. To identify tag candidate events, the combined prompt and LRT muon collection is passed through an established secondary vertexing algorithm also based on VKALVRT using the ID track information associated with each CB muon [42].

² The MS is theoretically sensitive to very late appearing pions – however, their behavior and resulting detector response is not well characterized. Existing muon reconstruction algorithms have relatively poor efficiency at reconstructing these tracks and significantly reduced accuracy at determining their charge and momentum. Future work could continue to explore SA muon vertexing under a pion hypothesis, but would likely require dedicated tuning at the “muon” reconstruction level.

This algorithm attempts two track vertex fits between each possible opposite-sign muon pair and requires a reconstructed vertex mass between 2.9 and 3.3 GeV alongside other basic quality criteria.

Once suitable tag vertex candidates have been identified, the individual MS tracks used in forming each constituent CB muon are retrieved and passed directly as seeding input to the “probe” SA muon vertexing algorithm. Because the initial MS track reconstruction is agnostic to whether a given candidate is subsequently matched to an ID track (CB) or remains unmatched (SA), this procedure is effectively identical in performance to running SA muon vertexing on genuine StandAlone muons. In this way, the probe construction isolates the MS contribution to the vertexing without relying on any ID information, providing a direct, data-driven validation of the SA muon vertex reconstruction chain under realistic detector conditions within the ID acceptance.

This approach is inherently conservative – as a light resonance produced within b -hadron decays, $J/\psi \rightarrow \mu^+\mu^-$ decays are both low p_T and highly collimated, well-known challenging features for secondary vertex reconstruction. The vertex position and mass resolutions naturally degrade as the opening angle between decay products shrinks, due to increasingly correlated track parameter uncertainties and reduced geometric leverage within the vertex fit. Additionally, as a prompt signature, the SA track extrapolations used in the probe construction must extend all the way back to the interaction point. These MS→IP extrapolations always traverse the entire calorimeter and magnetic field volume, sampling substantially more material and path length than the shorter MS→vertex extrapolations relevant for the displaced LLP decays, which typically originate a few hundred millimetres from the IP. This longer extrapolation amplifies uncertainties from multiple scattering, energy loss, and magnetic-field modelling, resulting in systematically larger spatial and momentum uncertainties than those expected for the LLP topologies of interest.

The yields in simulation for both distributions are normalized such that the number of J/ψ candidates with two ID tracks agrees with the data. To suppress non-signal and combinatorial background contributions, tag vertices are required to have $p_T^{\text{vtx}} > 5$ GeV, $\chi^2_{\text{vtx}} < 2$, and the opening angle ΔR between the track pair between 0.3 and 0.6. Events where the tag does not pass this selection are removed from consideration. Figure 9 shows the reconstructed mass distributions for both the tag ID track vertices and the probe SA muon vertices. The agreement between data and simulation in the tag is approximately flat, with some deviation in the tails due to increased combinatorial background contributions in data and differences in modeling of final state radiation. In the probe, the ratio remains approximately flat across the mass range, with a coarser resolution as expected. The flat ratio demonstrates that the efficiency and overall performance is well modeled between data and simulation, agreeing within 20% in the central J/ψ mass region. The slight shift towards higher masses in data versus simulation is due to small differences in uncalibrated muon momentum.

10 Conclusion

The ATLAS experiment has a diverse and thriving long-lived particle search program, recently strengthened by major improvements in large-radius tracking. Nevertheless, important gaps remain, in particular the limitations of vertex reconstruction outside the Inner Detector volume. To address this, a novel vertexing technique using Muon Spectrometer track information without an associated ID track has been developed and demonstrated in this note. This approach shows the capacity to clearly reconstruct well-known resonances such as $J/\psi \rightarrow \mu^+\mu^-$, as well as the efficacy to reconstruct displaced decays of two representative LLP signal models across different mass regimes, yielding distinct vertex topologies. These




Figure 9: Reconstructed m_{DV} of J/ψ candidate vertices using “tag” ID tracks (a) and “probe” SA muons (b). Distributions are fit to a double Gaussian with a third-order polynomial to model the combinatorial background contribution, overlaid.

results highlight the significant potential of this new vertexing method to open previously inaccessible regions of phase space for the ATLAS LLP program. Looking ahead, several developments could further enhance this strategy in future work. Integrating the standalone-muon vertexing with prompt and LRT-based reconstructions in a combined workflow and extending the approach to multitrack topologies or muon plus pion final states would strengthen its physics reach for a wider array of LLP scenarios. Additionally, applications beyond LLP searches, such as studies of material interactions in calorimeter material or cosmic ray muon analyses could also benefit from this technique with additional development. Finally, this technique can also be explored for the Phase II upgrade of the ATLAS detector for the High-Luminosity LHC, where improved muon system instrumentation and tracking capabilities may further enhance its performance and physics potential.

References

- [1] 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 [hep-ex] (cit. on pp. 2, 7, 8).
- [2] ATLAS Collaboration, *Software Performance of the ATLAS Track Reconstruction for LHC Run 3*, *Comput. Softw. Big Sci.* **8** (2024) 9, arXiv: 2308.09471 [hep-ex] (cit. on pp. 2, 7, 9).
- [3] ATLAS Collaboration, *Search for Light Long-Lived Particles in pp Collisions at $\sqrt{s} = 13$ TeV Using Displaced Vertices in the ATLAS Inner Detector*, *Phys. Rev. Lett.* **133** (2024) 161803, arXiv: 2403.15332 [hep-ex] (cit. on p. 2).
- [4] ATLAS Collaboration, *Search for heavy neutral leptons in decays of W bosons using leptonic and semi-leptonic displaced vertices in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *JHEP* **07** (2025) 196, arXiv: 2503.16213 [hep-ex] (cit. on p. 2).

- [5] ATLAS Collaboration, *Search for emerging jets in pp collisions at $\sqrt{s} = 13.6$ TeV with the ATLAS experiment*, *Rept. Prog. Phys.* **88** (2025) 097801, arXiv: [2505.02429 \[hep-ex\]](#) (cit. on p. 2).
- [6] ATLAS Collaboration, *Search for neutral long-lived particles that decay into displaced jets in the ATLAS calorimeter in association with leptons or jets using pp collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **11** (2024) 036, arXiv: [2407.09183 \[hep-ex\]](#) (cit. on p. 2).
- [7] ATLAS Collaboration, *Search for events with one displaced vertex from long-lived neutral particles decaying into hadronic jets in the ATLAS muon spectrometer in pp collisions at $\sqrt{s} = 13$ TeV*, (2025), arXiv: [2503.20445 \[hep-ex\]](#) (cit. on p. 2).
- [8] ATLAS Collaboration, *Standalone vertex finding in the ATLAS muon spectrometer*, *JINST* **9** (2014) P02001, arXiv: [1311.7070 \[physics.ins-det\]](#) (cit. on p. 3).
- [9] ATLAS Collaboration, *Search for light long-lived neutral particles from Higgs boson decays via vector-boson-fusion production from pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **84** (2024) 719, arXiv: [2311.18298 \[hep-ex\]](#) (cit. on p. 3).
- [10] ATLAS Collaboration, *Search for long-lived particles in final states with displaced dimuon vertices in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **99** (2019) 012001, arXiv: [1808.03057 \[hep-ex\]](#) (cit. on p. 3).
- [11] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003 (cit. on p. 3).
- [12] ATLAS Collaboration, *ATLAS Insertable B-Layer: Technical Design Report*, ATLAS-TDR-19; CERN-LHCC-2010-013, 2010, URL: <https://cds.cern.ch/record/1291633> (cit. on p. 3), Addendum: ATLAS-TDR-19-ADD-1; CERN-LHCC-2012-009, 2012, URL: <https://cds.cern.ch/record/1451888>.
- [13] B. Abbott et al., *Production and integration of the ATLAS Insertable B-Layer*, *JINST* **13** (2018) T05008, arXiv: [1803.00844 \[physics.ins-det\]](#) (cit. on p. 3).
- [14] ATLAS Collaboration, *ATLAS Calorimeter Performance: Technical Design Report*, ATLAS-TDR-1; CERN-LHCC-96-040, 1996, URL: <https://cds.cern.ch/record/331059> (cit. on p. 4).
- [15] ATLAS Collaboration, *ATLAS Liquid Argon Calorimeter: Technical Design Report*, ATLAS-TDR-2; CERN-LHCC-96-041, 1996, URL: <https://cds.cern.ch/record/331061> (cit. on p. 4).
- [16] ATLAS Collaboration, *ATLAS Tile Calorimeter: Technical Design Report*, ATLAS-TDR-3; CERN-LHCC-96-042, 1996, URL: <https://cds.cern.ch/record/331062> (cit. on p. 4).
- [17] ATLAS Collaboration, *ATLAS Muon Spectrometer: Technical Design Report*, ATLAS-TDR-10; CERN-LHCC-97-022, CERN, 1997, URL: <https://cds.cern.ch/record/331068> (cit. on p. 4).
- [18] S. Akatsuka, *The Phase-1 Upgrade of the ATLAS Level-1 Endcap Muon Trigger*, (2018), eprint: [1806.09234 \(physics.ins-det\)](#), URL: <https://arxiv.org/abs/1806.09234> (cit. on p. 5).
- [19] G. Avoni et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, *JINST* **13** (2018) P07017 (cit. on p. 4).

- [20] ATLAS Collaboration, *Performance of the ATLAS trigger system in 2015*, *Eur. Phys. J. C* **77** (2017) 317, arXiv: [1611.09661 \[hep-ex\]](#) (cit. on p. 4).
- [21] ATLAS Collaboration, *Software and computing for Run 3 of the ATLAS experiment at the LHC*, *Eur. Phys. J. C* **85** (2025) 234, arXiv: [2404.06335 \[hep-ex\]](#) (cit. on p. 4).
- [22] R. N. Mohapatra and G. Senjanovic, *Neutrino Mass and Spontaneous Parity Nonconservation*, *Phys. Rev. Lett.* **44** (1980) 912 (cit. on p. 4).
- [23] J. Schechter and J. W. F. Valle, *Neutrino Masses in $SU(2) \otimes U(1)$ Theories*, *Phys. Rev. D* **22** (1980) 2227 (cit. on p. 4).
- [24] J. Schechter and J. W. F. Valle, *Neutrino Decay and Spontaneous Violation of Lepton Number*, *Phys. Rev. D* **25** (1982) 774 (cit. on p. 4).
- [25] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079, arXiv: [1405.0301 \[hep-ph\]](#) (cit. on p. 5).
- [26] R. Ruiz, *Quantitative study on helicity inversion in Majorana neutrino decays at the LHC*, *Phys. Rev. D* **103** (1 2021) 015022, URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.015022> (cit. on p. 5).
- [27] NNPDF Collaboration, R. D. Ball et al., *Parton distributions for the LHC run II*, *JHEP* **04** (2015) 040, arXiv: [1410.8849 \[hep-ph\]](#) (cit. on p. 5).
- [28] C. Bierlich et al., *A comprehensive guide to the physics and usage of PYTHIA 8.3*, *SciPost Phys. Codebases* (2022) 8, arXiv: [2203.11601 \[hep-ph\]](#) (cit. on p. 5).
- [29] P. Artoisenet, R. Frederix, O. Mattelaer and R. Rietkerk, *Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations*, *JHEP* **03** (2013) 015, arXiv: [1212.3460 \[hep-ph\]](#) (cit. on p. 5).
- [30] S. Gopalakrishna, S. Jung and J. D. Wells, *Higgs boson decays to four fermions through an Abelian hidden sector*, *Phys. Rev. D* **78** (5 2008) 055002, URL: <https://link.aps.org/doi/10.1103/PhysRevD.78.055002> (cit. on p. 5).
- [31] NNPDF Collaboration, R. D. Ball et al., *Parton distributions with LHC data*, *Nucl. Phys. B* **867** (2013) 244, arXiv: [1207.1303 \[hep-ph\]](#) (cit. on p. 5).
- [32] ATLAS Collaboration, *ATLAS Pythia 8 tunes to 7 TeV data*, ATL-PHYS-PUB-2014-021, 2014, URL: <https://cds.cern.ch/record/1966419> (cit. on p. 5).
- [33] J. Pumplin et al., *New Generation of Parton Distributions with Uncertainties from Global QCD Analysis*, *JHEP* **07** (2002) 012, arXiv: [hep-ph/0201195](#) (cit. on p. 6).
- [34] ATLAS Collaboration, *The ATLAS trigger system for LHC Run 3 and trigger performance in 2022*, *JINST* **19** (2024) P06029, arXiv: [2401.06630 \[hep-ex\]](#) (cit. on p. 6).
- [35] 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 \[hep-ex\]](#) (cit. on p. 7).

- [36] R. Frühwirth, *Application of Kalman filtering to track and vertex fitting*, [Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment](#) **262** (1987) 444, ISSN: 0168-9002, URL: <https://www.sciencedirect.com/science/article/pii/0168900287908874> (cit. on p. 7).
- [37] 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, URL: <https://cds.cern.ch/record/2669425> (cit. on pp. 8, 13).
- [38] V. Kostyukhin, *VKalVrt - package for vertex reconstruction in ATLAS.*, ATL-PHYS-2003-031, 2003, URL: <https://cds.cern.ch/record/685551> (cit. on pp. 8, 11).
- [39] 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 \[hep-ex\]](#) (cit. on p. 9).
- [40] A. Salzburger, *The ATLAS Track Extrapolation Package*, tech. rep., All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-SOFT-PUB-2007-005>: CERN, 2007, URL: <https://cds.cern.ch/record/1038100> (cit. on p. 11).
- [41] E. Lund, L. Bugge, I. Gavrilenko and A. Strandlie, *Track parameter propagation through the application of a new adaptive Runge-Kutta-Nystroem method in the ATLAS experiment*, [JINST](#) **4** (2009) P04001 (cit. on p. 11).
- [42] ATLAS Collaboration, *K_S^0 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 \[hep-ex\]](#) (cit. on p. 17).