RUPRECHT-KARLS-UNIVERSITÄT HEIDELBERG

Tobias Heintz

Back to the Future: Probing Ultra-Slow Long-Lived Particles in Consecutive Event Pairs using the ATLAS Detector

Master Thesis

HD-KIP-25-107

KIRCHHOFF-INSTITUT FÜR PHYSIK

Department of Physics and Astronomy
Heidelberg University

Master Thesis
in Physics
submitted by

Tobias Heintz
born in
Schweinfurt (Germany)

2025

Back to the Future: Probing Ultra-Slow Long-Lived Particles in Consecutive Event Pairs using the ATLAS Detector

This Master Thesis has been carried out by Tobias Heintz at the
Kirchhoff Institute for Physics in Heidelberg
under the supervision of
Prof. Dr. Hans-Christian Schultz-Coulon
and
Prof. Dr. Oleg Brandt
Cavendish Laboratory, University of Cambridge

Back to the Future: Probing Ultra-Slow Long-Lived Particles in Consecutive Event Pairs using the ATLAS Detector

This thesis explores a novel approach to identify Long-Lived Particles (LLPs) whose decays occur across two consecutive pp bunch-crossings (BCs), reducing the capability of probing them due to limitations of standard trigger and reconstruction strategies.

A dedicated trigger is utilised that correlates missing transverse momentum from a slow LLP with a $> 25\,\text{ns}$ delayed, azimuthally aligned displaced jet. Validation studies in pp collision data affirm trigger performance and inform future improvements. An extended reconstruction framework was developed for out-of-time decays, which are not considered in the standard ATLAS digitisation applying the LHC BC structure. Reconstructed quantities are compared to expectations based on MC truth, observing good closure. Disagreements between reconstructed and truth jet energies are shown to be a feature of standard jet calibrations.

A toy model inspired by Inelastic Dipole Dark Matter benchmarks the trigger performance. Truth-level feasibility studies give sub-percent signal efficiencies and Run-2 yields of $\mathcal{O}$(10-100), found to be limited by a small fraction of slow LLPs and trigger-shadowing from initial-state radiations. A data-driven background projection predicts $\mathcal{O}(1600)$ events, leading to cross-section limits of $\mathcal{O}(0.1\,\text{pb})$. These are weaker than existing ATLAS constraints, calling for an improved benchmark model with a larger fraction of slow LLPs.

Zurück in die Zukunft: Untersuchung extrem langsamer langlebiger Teilchen in aufeinanderfolgenden Ereignispaaren mit dem ATLAS-Detektor

Diese Arbeit untersucht einen innovativen Ansatz zur Identifikation langlebiger Teilchen (LLPs), deren Zerfälle sich über zwei aufeinanderfolgende pp-Bunch-Crossings (BCs) erstrecken, wodurch die Möglichkeit ihrer Untersuchung aufgrund der Einschränkungen von Standard-Trigger- und Rekonstruktionsstrategien reduziert wird.

Ein speziell entwickelter Trigger korreliert den fehlenden transversalen Impuls eines langsamen LLPs mit einem um $> 25\,\text{ns}$ verzögerten, azimutal abgestimmten Jet mit versetztem Ursprung. Validierungsstudien in pp-Kollisionsdaten bestätigen die Triggerleistung und liefern Ansatzpunkte für künftige Verbesserungen. Zudem wurde ein erweitertes Rekonstruktionsframework für zeitlich verzetzte Zerfälle entwickelt, die in der Standard-ATLAS-Digitisierung unter Berücksichtigung der LHC-BC-Struktur nicht enthalten sind. Rekonstruierte Größen werden mit Erwartungen aus MC-Truth verglichen und zeigen gute Übereinstimmung. Abweichungen zwischen rekonstruierter und wahrer Jetenergie werden als Eigenschaft der Standard-Jetkalibrierungen identifiziert.

Ein Toy-Modell inspiriert von Inelastischer-Dipol Dunkler Materie dient zur Bewertung der Triggerleistung. Truth-Level-Machbarkeitsstudien ergeben Signal-Effizienzen im Subprozentbereich und Run-2-Erträge von $\mathcal{O}$(10-100), begrenzt durch den geringen Anteil langsamer LLPs und Trigger-Abschattung durch QCD-Anfangszustrahlung (ISR). Eine datengetriebene Hintergrundprojektion sagt $\mathcal{O}(1600)$ Ereignisse voraus und führt zu Wirkungsquerschnittsgrenzen von $\mathcal{O}(0.1\,\text{pb})$, die schwächer sind als bestehende ATLAS-Beschränkungen, und deshalb ein verbessertes Benchmark-Modell mit höherem Anteil langsamer LLPs nahelegen.

Erklärung

Ich versichere, dass ich diese Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Heidelberg, den 23.12.2025 Tobias Stephan Heintz

Author's Contribution

The Back-to-the-Future (B2F) analysis is carried out by a small team within the ATLAS Collaboration and benefits from its infrastructure in multiple essential ways. The development and operation of the proposed trigger, the successful data taking, and the provision and maintenance of the central software framework are all indispensable. Equally crucial is the exchange of expertise with specialists in trigger operations, the trigger menu coordination, calorimetry, jet calibration, tracking & vertexing, and Long-Lived Particle searches, as well as with software experts. Furthermore, theory colleagues contributed to shaping and validating the physics strategy.

To avoid confusions about the contributions made by the author of this thesis, this section explicitly lists them. Here, it is noteworthy that not the entirety of the work conducted by the author during their Master's project is reported in this thesis, since the novel and unique strategy of the B2F project involved a lot of trial and error, making significant parts of their work exploratory in nature and leading to a refinement of the final strategy.

The author was significantly involved in detailed feasibility studies, the development and improvement of the B2F trigger, the development of the reconstruction and analysis frameworks and an analysis strategy, with their main contributions as follows:

1. Overall Strategy (Chapter 5)

   The author sharpened the analysis motivation and strategy, including exploration of the relevant phase-space, background-reduction, and feasibility studies. They also promoted the analysis concept through a lightning talk at the annual ATLAS UK conference [1] and a presentation in a parallel session at the DPG spring meeting [2].

2. Phenomenology and Sensitivity Studies (Chapter 6)

   The author was significantly involved in defining a benchmark model and its validation in close collaboration with theorists André Lessa and José Zurita.
   The author designed an efficient Monte Carlo (MC) production chain based on `HTCondor DAGMan`, and set up the necessary `MadGraph` and `Pythia` configurations as well as the MLM merging and matching settings, based on the UFO card provided by André Lessa, which encodes the Feynman rules of the BSM model in a standardised output format. Furthermore, the author implemented a toy MC technique for lifetime reweighting, making the MC generation more efficient.
   The overall MC production chain was validated by the author, including comparisons of Pythia-only, MG+Pythia, and MG+MLM+Pythia simulations, both within the Athena framework and stand-alone, as well as validation of the lifetime reweighting procedure.
   The author also developed a user-specific truth-level framework, which included the implementation of particle propagation, Level-1 (L1) trigger algorithms, jet clustering, and related components. The framework was successfully validated against an alternative Delphes card provided by André Lessa.
   Detailed studies were performed on the decay kinematics of the benchmark model, and on the B2F trigger and selection efficiencies. The author conceived and carried out the background estimate and identified the key background-reduction mechanism enabling a search strategy requiring only a single displaced jet — unlike existing calorimeter-based searches, which

require two. This insight motivated a redefinition of the benchmark model. The upper limit setting based on the signal and background yields, as well as the re-interpretation of the Calorimeter Energy Ratio (CalRatio) and Muon Spectrometer Displaced Vertex (MSVtx) analyses were carried out by the author and validated against the results of the alternative Delphes framework, showing closure.
The author also contributed to discussions on potential publication strategies.

3. Trigger Development, Validation & Optimisation (Section 7.1)

   The author validated the trigger performance in pp collision data, identified discrepancies, uncovered a bug of the B2F Higher-Level Trigger (HLT) chain, traced it to misbehaving pile-up removal and EMF calculation, and initiated discussions with the CalRatio trigger experts. An alternative HLT chain was proposed by the author as an interim solution.
   The author estimated HLT rates for the background-enriched control region trigger together with Ben Kerridge, and performed validation studies due to a misconfiguration in the 2025 trigger menu.
   Contributions were also made to a more general L1-correlation trigger concept. The author additionally initiated discussions about the Run-4 trigger developments.

4. Development and Validation of Simulation and Reconstruction Framework in Athena (Sections 4.3, 7.2-7.4)

   The author was significantly involved in implementing an extension to the Athena framework allowing out-of-time decays to be processed in the ATLAS simulation, digitisation, and reconstruction frameworks.
   Based on an idea of John Chapman, the author implemented a user-developed Geant4 *Extensions* and *UserAction* to write delayed particles into a separate file during the simulation, and handle their simulation, digitisation and reconstruction separately. However, this approach was abandoned by the author after validation studies revealed a conceptual limitation, namely loosing significant contributions of the out-of-time part.
   Instead, they conceived a custom `HitWrapper` tool as alternative approach which makes a copy of the out-of-time part and subsequently shifts it by one bunch-crossing, bringing it back into the "standard" time-window such that it can be processed within the standard Athena framework. The author co-developed this tool with John Chapman, with a focus on validation and improvements.
   Furthermore, a concept for properly simulating the trigger decision was proposed by the author, although the actual implementation in Athena lies beyond the scope of the thesis.
   A user-specific event filter was implemented and the MC production chain was set up using the corresponding Athena scripts; the MC production was validated by the author, and detailed studies were carried out on jet calibration, calorimeter timing, and cluster-level behaviour.
   The author developed the reconstruction-level analysis framework for processing pp collision data as well as signal simulations, including grid integration; moreover, they defined preliminary event selection criteria for the signal selection.

Acknowledgments

First and foremost, I would like to express my sincere gratitude to my supervisors, Hans-Christian Schultz-Coulon and Oleg Brandt, for giving me the opportunity to pursue my Master's project in cooperation with the Heidelberg and Cambridge ATLAS groups. Over the past year, I have grown not only as a physicist, but also personally. I am particularly grateful for their tireless support and encouragement regarding my future plans within ATLAS.

Special thanks go to everyone who shared the challenges of B2F analysis with me, which is certainly one of the most demanding studies I have undertaken so far. I am especially grateful to Oleg for his continuous support and mentorship; I learned a great deal from his experience, and I deeply admire his optimism and willingness to pursue unconventional ideas. I also thank Gareth and Jack for their guidance and patience, as well as André and José for the enjoyable collaboration and for introducing me to aspects of theoretical physics. I am thankful to John for his technical contributions, without which significant parts of this work would not have been possible, and to Jeremy for his input on the sensitivity studies.

I would like to thank the HEP groups at the Kirchhoff Institute and the Cavendish Laboratory for providing such a welcoming and stimulating environment. I am grateful for the support, interesting discussions, and many enjoyable activities. I also thank Chris and Oleg for making my trip to CERN possible; it was a valuable and inspiring experience, and I look forward to returning in the future. Finally, I thank Tina for giving me the opportunity to engage in teaching and outreach within the prestigious Cambridge system, which was both highly enjoyable and a great source of inspiration.

I am deeply grateful to my family and friends for their unwavering support and understanding, particularly during the more challenging phases of this work.

Contents

1. Introduction

The elementary constituents of matter and how they make up the universe have inspired mankind for centuries. Atoms have been believed to be the fundamental constituents of nature for a long time[1], but with the discovery of the electron by J.J. Thomson at the end of the 19th century, the first sub-atomic particle was found. Rutherford furthermore showed that a very dense positively charged nucleus exists inside atoms besides the electrons. In contrast to the electron, the nucleus is not fundamental, but consists of protons and neutrons, which themselves are made up of quarks and gluons, as described in the so-called parton model. Together with the electroweak unification, this made up the first formulation of the Standard Model (SM) of particle physics as we know it today.

1: The term atom originates from the Greek word ατομος meaning indivisible.

Many measurements, including precise tests of higher-order corrections, showed outstanding agreement with the theoretical predictions, underscoring the robustness of this theoretical framework. Moreover, it turned out to have a very high predictive power – predicting not only the Higgs boson but further imposing *indirect* constraints on the free parameters of the SM model in advance of their direct measurement.

Despite this great success, open questions remain beyond the reach of the SM, the so-called (BSM) sector. Prominent examples include the nature of its underlying parameters, the origin of dark matter and energy, as well as the matter-antimatter asymmetry. None of these has been successfully addressed to date, although several theoretical extensions of the SM have been probed by a variety of experiments.

The afore-mentioned agreement between measurements and predictions sets rather high constraints on the energy scale of new phenomena. Therefore, a promising approach is testing the high-energy frontier, for example with the collider experiments at the Large Hadron Collider (LHC) — the most energetic particle collider to date, reaching the terra-electronvolt scale.

One strategy that is currently gaining popularity focuses on phenomena that, due to their exotic signatures, may be of limited exploration power in classical searches, as conventional trigger and reconstruction strategies are less efficient, requiring dedicated search strategies. An example of this are particles with rather high lifetimes and decay lengths, which are commonly motivated by their unconventional signatures in the detector and are used

to probe a variety of theoretical extensions to the SM, in which Long-Lived Particles (LLPs) often occur naturally.

A currently under-explored LLP signature comes from slow-moving particles with sufficiently high lifetimes such that they decay with a significant time delay and a spatial displacement still inside the active detector volume. Conventional triggers would be of limited use for out-of-time decays, but a dedicated LLP trigger that correlates different detector signals over time could be used to detect out-of-time decaying LLPs.

In the case of electrically neutral LLPs, the expected signature would be missing transverse momentum (at a time $t = t_0 \sim 0$) followed by the decay signature (at $t \sim t_0 + t_{\text{delay}}$) pointing in the same azimuthal direction, for instance a displaced jet.

This is a very novel and unique approach in LLP searches, and the presented analysis is the first ATLAS LLP analysis ever to target a decay structure ranging over a time period of up to $t_{\text{delay}} \sim \mathcal{O}(50\,\text{ns})$, although ATLAS has been running and analysing data since 2010. Given the novel and unique approach, detailed feasibility studies need to be performed as a proof-of-concept and to physically motivate this unique strategy.

Therefore, this thesis presents feasibility studies for this dedicated LLP trigger, both from an experimental and theoretical perspective. A toy model inspired by Inelastic Dipole Dark Matter (IDDM) is utilised to benchmark the sensitivity of the trigger for this unconventional feature of the new physics; expected signal yields and backgrounds are extrapolated, and the resulting exclusion limits are compared to other BSM searches. For the experimental feasibility studies, the detector response is additionally taken into account, focusing on the reconstruction of out-of-time LLP decays, calibration of displaced jets, and the calorimeter timing. Also, the recorded pp collision data is used to validate the trigger performance.

The thesis is structured as follows. After a short summary of the theoretical framework of the SM in Chapter 2, the experimental environment at the LHC and the ATLAS detector is described in Chapter 3. The simulation of proton-proton collisions is outlined in Chapter 4 with a focus on simulating exotic LLPs. The main part of this thesis is divided into the phenomenological and experimental studies, presented in Chapter 6 and 7, respectively. Chapter 8 serves as a summary of the research undertaken and results attained in the course of this thesis. The supporting material is provided in the subsequent appendices.

2. The Standard Model of Particle Physics and Beyond

The following chapter introduces the theoretical framework of particle physics, the Standard Model (SM), as it is currently established. Moreover, an overview of searches and theories beyond this current formulation is presented, the so-called BSM sector, focussing on exotic Long-Lived Particles (LLPs).

2.1. The Standard Model of Particle Physics

The SM of particle physics [3] is a well-established and precisely-tested theoretical framework describing the *elementary particles and their interactions* via the fundamental forces, such as electromagnetism, strong and weak nuclear interactions. Since these interactions take place at the smallest measurable length scales, and the particles usually have relativistic properties, it is crucial to combine special relativity and quantum mechanics. This is achieved with the introduction of so-called quantum fields, Ψ, which span the entire space-time continuum and whose excitations are interpreted as particles.

[3]: Thomson (2013), *Modern Particle Physics*

Just as the laws of physics do not depend on the reference frame and therefore remain invariant under rotations, they should also remain invariant under phase transformations of the fields:

$$\Psi(x) \rightarrow \Psi'(x) = \Psi(x) \times \exp(i\alpha)\,. \tag{2.1}$$

Since the phase transformation might depend on the space-time, $\alpha \equiv \alpha(x)$, its derivative might not vanish, $\partial_\mu \alpha \not\equiv 0$, which could break the symmetry of the Lagrangian. The symmetry is restored by adding extra gauge fields, interpreted as the force carrying mediator particles, compensating for the additional terms [4].

[4]: Schmüser (2011), *Feynman-Graphen und Eichtheorien für Experimentalphysiker*

The SM contains, as shown in Figure 2.1, the following:

- The elementary matter particles (spin 1/2 fermions);
- The force carrying exchange particles (spin 1 vector bosons);
- The Higgs boson (spin 0 scalar) introducing a gauge-invariant way of generating particle masses.

The gauge *bosons* are the exchange particles of the force fields, such as the photon for the electromagnetic force, the gluons for the strong interaction, and the Z and $W^\pm$ bosons for the weak force.

The photon is massless and electrically neutral; its coupling is proportional to the electrical charge, Q, which forbids photon-photon self-interactions. This is significantly different for gluons

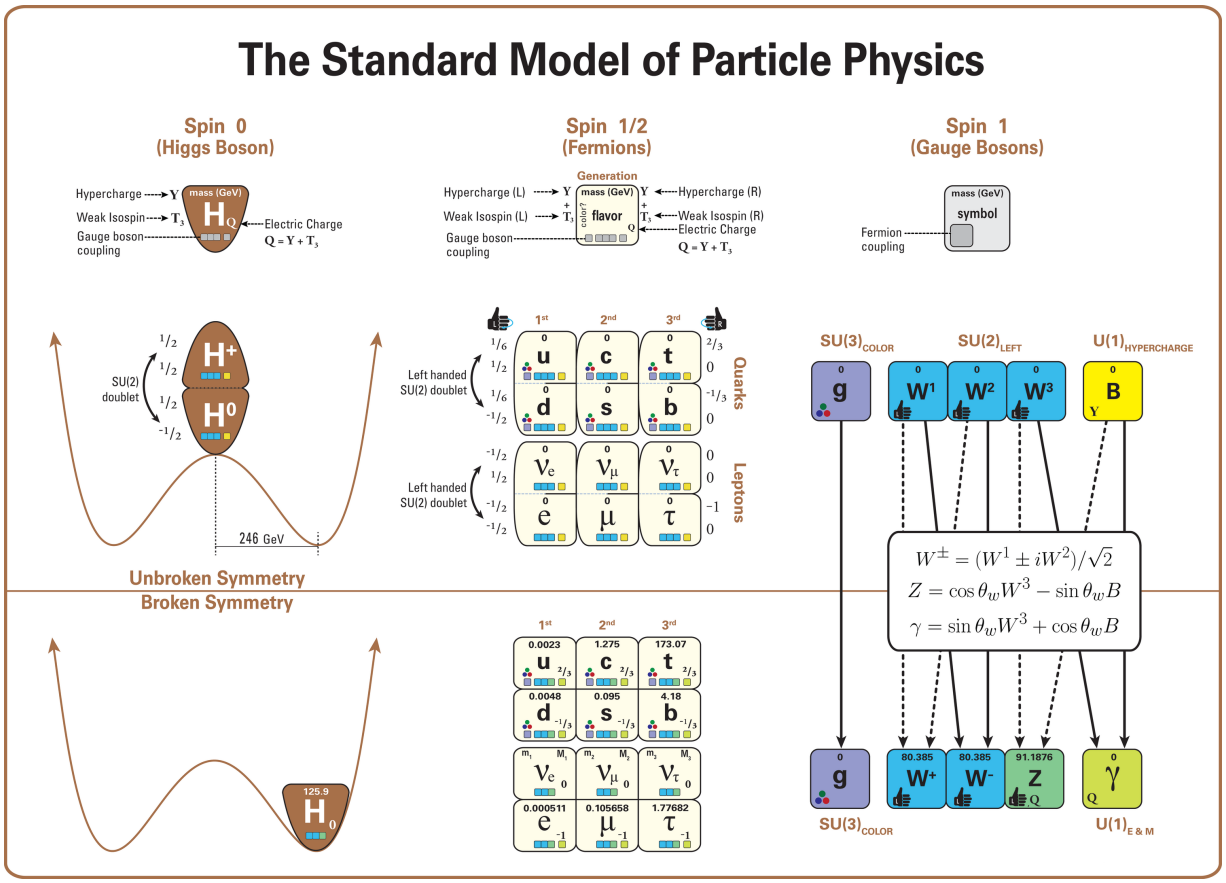

Figure 2.1.: Overview of the SM of particle physics. The SM is a quantum field theory, with its gauge structure shown in the upper part. The SM consists of elementary fermions (middle), gauge bosons as exchange particles of the force fields (right) and the Higgs boson (left), with the latter introducing a gauge-invariant way to introduce particle masses via a spontaneous symmetry breaking (lower part). Taken from [5].

which are colour-charged themselves, allowing for gluon-gluon self-interactions. That is, the strong force is described by a so-called *non-Abelian* gauge theory [6], with significant consequences like colour confinement and non-perturbative behaviour at low energies. Similarly, the non-vanishing weak isospin of the $W^{\pm}$ bosons, $T_3 = \pm 1$, allows for W self-interactions[2].

[6]: Peskin et al. (1995), *An Introduction to quantum field theory*

2: Further self-interactions of W, Z, and γ arise through the electroweak unification (see below).

The elementary *fermions* can be further classified into quarks and leptons, each of which exists in three generations; each generation is a doublet of up- and down-type quarks or electron- and neutrino-type leptons respectively:

Generation	1st	2nd	3rd
Quarks:	$\begin{pmatrix} u \\ d \end{pmatrix}$	$\begin{pmatrix} c \\ s \end{pmatrix}$	$\begin{pmatrix} t \\ b \end{pmatrix}$
Leptons:	$\begin{pmatrix} e \\ \nu_e \end{pmatrix}$	$\begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix}$	$\begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}$

The fermion masses increase with each generation[3], but the quantum numbers do not change throughout the individual generations; in contrast, quarks and leptons are distinct by their quantum numbers — especially, their electrical, colour, and weak charges differ, regarding which interactions they participate in: leptons are charged under the weak and/or electric force, and quarks carry additionally colour charges. That is, leptons interact via electromagnetism and/or the weak force, which are unified in the Weinberg-Salam-Glashow theory [7–9], whereas quarks additionally interact through the strong force, as described in quantum chromodynamics (QCD).

3: Though, no fixed mass hierarchy is known to-date for the neutrinos.

[7]: Weinberg (1967), *A Model of Leptons*
[8]: Salam (1968), *Weak and Electromagnetic Interactions*
[9]: Glashow (1961), *Partial-symmetries of weak interactions*

Electroweak Sector

From the charged current weak interaction (*e.g.* the $\beta^{\pm}$ decay) it is known that electromagnetism and the weak force are connected; in fact, they can be unified into one force at high energies, as present shortly after the Big Bang. The unification is subject to the symmetry described by the

$$\mathrm{SU(2)_L} \times \mathrm{U(1)_Y}$$

group. Here, $\mathrm{SU(2)_L}$, coming along with three[4] gauge fields, $\mathbf{W}_\mu = \left(W_\mu^{(1)}, W_\mu^{(2)}, W_\mu^{(3)}\right)$, describes rotations within the left handed doublets. On the other hand, $\mathrm{U(1)_Y}$, coming along with one gauge field, B_μ, covers rotations in the hypercharge space, $Y = Q - T_3$, mixing the electric charge, Q, and the weak isospin, T_3.

4: In general, a SU(N) symmetry comes with $N^2 - 1$ generators. To conserve local gauge invariance under the SU(N) transformation, the same number of gauge fields is required and is added to the covariant derivative.

Before electroweak (EW) symmetry breaking, the fermions and gauge bosons, as well as their interactions, are described by the EW Lagrangian:

$$\mathscr{L}_{\mathrm{EW}} = \underbrace{-\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \sum_\psi \bar{\psi} i \not{D} \psi}_{\mathscr{L}_{\mathrm{free}} + \mathscr{L}_{\mathrm{int}}} + \underbrace{\mathscr{L}_{\mathrm{Higgs}} + \mathscr{L}_{\mathrm{Yukawa}}}_{\text{Higgs Sector}} . \tag{2.2}$$

Here, the $\psi = u(p)\exp\left(ipx\right)$ terms describe the fermions[5] and the $F_{\mu\nu}F^{\mu\nu} = \mathbf{W}_{\mu\nu}\mathbf{W}^{\mu\nu} + B_{\mu\nu}B^{\mu\nu}$ are the kinematic terms of the gauge bosons:

5: Where $u(p)$ are the Dirac spinors, while anti-fermions are described by $\psi = v(p)\exp\left(-ipx\right)$ [3].

$$W_{\mu\nu}^{(i)} = \partial_\mu W_\nu^{(i)} - \partial_\nu W_\mu^{(i)} + g \sum_{j,k} \varepsilon_{ijk} W_\mu^{(j)} W_\nu^{(k)}, \tag{2.3}$$

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu. \tag{2.4}$$

It is notable that the latter part in $W_{\mu\nu}$ is causing the non-Abelian gauge structure, as visualised in Figure 2.2.

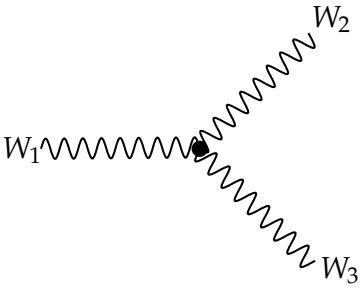

(a) Triple $W_1W_2W_3$ coupling, $-\frac{g}{2}\varepsilon_{jkl}\left(\partial^\mu W_j^\nu - \partial^\nu W_j^\mu\right)W_\mu^k W_\nu^l$.

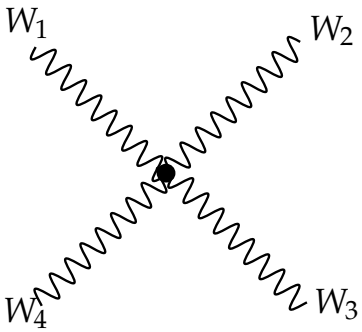

(b) Quartic $W_1W_2W_3W_4$ coupling, $-\frac{g^2}{4}\varepsilon_{jkl}\varepsilon_{jmn}W_k^\mu W_l^\nu W_\mu^m W_\nu^n$.

Figure 2.2.: Triple and quartic gauge couplings in the EW sector of the SM.

[10]: Higgs (1964), *Broken Symmetries and the Masses of Gauge Bosons*
[11]: Englert et al. (1964), *Broken Symmetry and the Mass of Gauge Vector Mesons*
[12]: Guralnik et al. (1964), *Global Conservation Laws and Massless Particles*

The interactions between the gauge fields and the fermions are encapsulated in the second part of the co-variant derivative

$$\not{D} = \gamma^\mu D_\mu = \gamma^\mu \left(\partial_\mu - ieA_\mu\right) \tag{2.5}$$

with $eA_\mu = \frac{g}{2}\mathbf{TW}_\mu + \frac{g'}{2}YB_\mu$, where $g = e/\sin\theta_W$ and $g' = e/\cos\theta_W$ are the $SU(2)_L$ and $U(1)_Y$ couplings, respectively, and $\cos\theta_W = m_W/m_Z$ is the weak mixing angle.

This theory is perfectly self-consistent, but has a significant limitation: it is known from experiment that the weak force treats the left-handed and right-handed part of $\psi \equiv \psi_L + \psi_R$ differently (also known as the $V-A$ structure). More specifically, EW interactions under the $SU(2)_L$ symmetry only rotate within the left-handed doublets, but leave the right-handed singlets un-transformed. This forbids adding an ad-hoc mass term $m\psi\bar{\psi}$ as it would break gauge symmetry.

This challenge is overcome by extending the EW unification with the Higgs formalism [10–12]. Within this formalism, the $SU(2)_L \times U(1)_Y$ break their symmetry, forming the massless photon, and the massive $W^\pm$, and Z Bosons:

$$W_\mu^\pm = \frac{1}{\sqrt{2}}\left(W_\mu^{(1)} \mp iW_\mu^{(2)}\right), \tag{2.6}$$

$$Z_\mu = -B_\mu \sin\theta_W + W_\mu^{(3)}\cos\theta_W, \text{and} \tag{2.7}$$

$$A_\mu = +B_\mu \cos\theta_W + W_\mu^{(3)}\sin\theta_W. \tag{2.8}$$

Higgs Formalism

Gauge-invariant mass terms can be added via the Yukawa interaction:

$$\mathcal{L}_{\text{Yukawa}} = \sum_\psi c_f \left(\bar{\psi}_L\psi_R\Phi + \Phi^\dagger\bar{\psi}_R\psi_L\right), \tag{2.9}$$

with the corresponding Yukawa couplings c_f being proportional to the fermion masses, and ψ_L being the left-handed fermion-doublets and ψ_R the right-handed fermion-singlets. Here, another SU(2) doublet of complex scalar fields is ad-hoc introduced:

$$\Phi = \begin{pmatrix}\phi^+\\ \phi^0\end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix}\phi_1 + i\phi_2\\ \phi_3 + i\phi_4\end{pmatrix}, \tag{2.10}$$

coming with the corresponding Lagrangian

$$\mathcal{L}_{\text{Higgs}} = \left(D_\mu\Phi\right)^\dagger\left(D^\mu\Phi\right) - V(\Phi). \tag{2.11}$$

The Higgs-Potential

$$V(\Phi) = -\mu^2 \left|\Phi^\dagger \Phi\right| + \lambda \left|\Phi^\dagger \Phi\right|^2 \sim |\Phi|^2 + |\Phi|^4 \tag{2.12}$$

has — because of $0 < \mu^2, \lambda \in \mathbb{R}$ — the shape of a Mexican hat, as visualised in Figure 2.3.

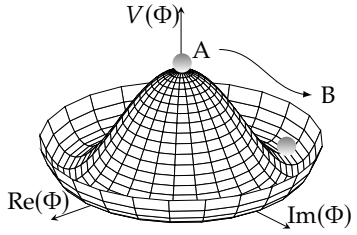

Figure 2.3.: Visualisation of the Higgs potential with the shape of a Mexican hat. Source code modified from [13].

In this general form, these expressions are gauge-invariant. While the potential has infinitely many minima, only one is implemented by nature. This is known as spontaneous symmetry breaking, meaning that the overall theory obeys a gauge symmetry, but the particular choice of minimum

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} =: \begin{pmatrix} 0 \\ \phi \end{pmatrix} \tag{2.13}$$

does not. Perturbations of ϕ around its minimum are described by an expansion

$$\phi \approx \frac{1}{\sqrt{2}} \left(v + \eta(x) + i\xi(x)\right), \tag{2.14}$$

with perturbations in the azimuthal direction through the massless Goldstone bosons[6] $\xi(x)$ and perturbations in the radial direction through the field of the massive gauge boson $\eta(x)$.

6: As the Higgs formalism is embedded into the $\mathrm{SU(2)_L \times U(1)_Y}$ symmetry: $\xi = \xi_0 + \xi_1 + \xi_2 + \xi_3$ with $\xi_i = \left\{B_\mu, W_\mu^{(1)}, W_\mu^{(2)}, W_\mu^{(3)}\right\}_i$.

In the unitary gauge [14] this simplifies to $\xi(x) = 0$ and $\eta(x) = h(x)$, which gives rise to a fermion-Higgs interaction term $c_f/\sqrt{2} \times h\, \bar{\psi}\psi$ in eq. 2.9. The remaining $c_f/\sqrt{2} \times v\, \bar{\psi}\psi$ terms are interpreted as the mass terms $m_f \bar{\psi}\psi$ with $m_f = \frac{c_f v}{\sqrt{2}}$.

[14]: Lee (1972), *Higher-Order Corrections to Leptonic Processes and the Renormalization of Weinberg's Theory of Weak Interactions in the Unitary Gauge*

That is, the Higgs formalism introduces a gauge-invariant way to generate the fermion masses via their interactions with the non-vanishing vacuum expectation value v of the Higgs potential.

It is notable that v also dictates the mass scale of the Higgs boson and of the gauge bosons, as can be seen from the expansion of ϕ in equation 2.11; furthermore, this also gives rise to the interaction between the Higgs boson and the gauge bosons as well as the self-interaction of the Higgs boson.

Quantum Chromodynamics

The strong interaction between quarks and gluons is described byquantum chromodynamics (QCD), a non-Abelian $\mathrm{SU(3)_C}$ symmetry coming with eight generators, $\lambda_{(A)}$, and gluon fields $G_\mu^{(A)}$.

The QCD Lagrangian[7] reads

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4}\sum_{A=1}^{8} G^{(A)}_{\mu\nu} G^{\mu\nu}_{(A)} + \sum_{\psi}\sum_{a,b} \bar{\psi}_{(a)} \left(i\not{D}_{ab} - m\delta_{ab}\right) \psi_{(b)}, \tag{2.15}$$

with $a, b \in \{\text{r,g,b}\}$ considering the colour flow of the quarks ψ, and $A = 1, \ldots, 8$ describing the gluon fields $G^{(A)}_{\mu}$:

$$G^{(A)}_{\mu\nu} = \partial_\mu G^{(A)}_{\nu} - \partial_\nu G^{(A)}_{\mu} - g_{\text{s}} \sum_{B,C} f_{ABC} G^{(B)}_{\mu} G^{(C)}_{\nu}, \tag{2.16}$$

where the last term is causing the non-Abelian gauge structure of QCD with f_{ABC} being the structure constants of SU(3) and g_{s} being the bare coupling constant of the strong interaction. Moreover, the co-variant derivative is given by

$$\not{D}_{ab} = \gamma^\mu \left(D_\mu\right)_{ab} = \gamma^\mu \left(\partial_\mu \delta_{ab} + i g_{\text{s}} \left(\sum_{A} \lambda_{(A)} G^{(A)}_{\mu}\right)_{ab}\right), \tag{2.17}$$

with the second part describing the interactions between gluons and quarks, where λ are the Gell-Mann matrices.

As the LHC collides the constituents of protons (*i.e.* quarks and gluons), QCD plays a critical role in LHC physics; therefore, the most relevant points are summarised in the following paragraphs.

Due to the gluon–gluon self-interaction, additional diagrams contribute to the vacuum polarisation, causing the strong coupling constant to increase towards lower energy scales, Q^2. As a consequence, collisions at the LHC can be separated into a hard-scattering regime, characterised by high Q^2 and moderate coupling where perturbation theory is applicable, and a hadronisation regime at low Q^2, where the coupling becomes large and perturbative calculations break down. This non-perturbative regime is characterised by parton showering and high jet activity, leading to large overall QCD backgrounds. Therefore, higher-order contributions must be taken into account to obtain reliable predictions. Furthermore, renormalisation is essential to ensure finite results over a broad range of energy scales [15].

Moreover, the initial state partons are characterised by the parton distribution functions which correspond to the probability of a parton having a certain momentum fraction x of the colliding protons. However, these functions are not known from first principle, and rely on measurements.

7: In contrast to the EW Lagrangian, the quark mass terms in the QCD Lagrangian are gauge-invariant.

[15]: Dodd et al. (2020), *The ideas of particle physics*

2.2. Physics Beyond the Standard Model

The SM has proven to be consistent with measurements in a broad range of energies: from low-energy quantum effects to high-energy electroweak precision measurements. Moreover, it not only successfully predicted the Higgs boson and its properties (except for its mass) [16, 17], but it also allowed constraints to be placed upon some of the free SM-parameters by measuring higher-order contributions [18]. The outstanding agreement of the measurements with the SM sets rather high constraints on the mass scale of physics beyond it, $\Lambda_{\mathrm{UV}} \sim \mathcal{O}(\mathrm{TeV})$. Notwithstanding this success, outstanding questions require the existence of physics beyond the current formulation of the SM [19]:

- Despite being favoured by cosmological measurements, dark matter and dark energy are currently not described in the framework of the SM. Also, the observed matter-antimatter asymmetry in the universe cannot be explained by the SM;
- The measurement of neutrino oscillation proving that neutrinos are massive, requires an extension of the SM structure, for example by adding right-handed neutrino singlets;
- The large amount of free parameters in the SM and their scales are currently not understood. For instance, the masses of the elementary particles range over a rather large scale, from neutrino masses on the sub-eV level to the mass of the top quark, $m_t = 173\,\mathrm{GeV}$[8]. Also, the value of the QCD vacuum angle $\theta \lesssim 10^{-10}$ requires either "fine tuning" of this parameter or a BSM explanation enforcing a vanishing QCD vacuum, also known as strong CP problem [20, 21];
- From a theoretical perspective, there are strong motivations to describe all elementary forces in a grand unified theory, as present under conditions as in the Big Bang. However, the current framework fails to accommodate gravity.

[16]: ATLAS Collaboration (2012), *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*
[17]: CMS Collaboration (2012), *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*
[18]: Bagger et al. (2006), *Precision electroweak measurements on the Z resonance*
[19]: Cirigliano et al. (2013), *Low energy probes of physics beyond the standard model*

8: Interestingly, this means that the Yukawa coupling of the top quark is very close to unity

$$c_t = \sqrt{2} m_t / v \sim 0.99.$$

[20]: Peccei (2008), *The Strong CP Problem and Axions*
[21]: Navas et al. (2024), *Review of particle physics*

These questions have been addressed in a variety of experiments in the last decades, among others in the LHC physics programme. For instance, the searches at the ATLAS experiment range from supersymmetric searches, dark matter searches to exotic searches. Despite the high energies accessible at the LHC, the huge dataset, and the variety of BSM searches, no significant excess has been observed yet; this naturally gives rise to uncenventional searches for signatures that might be of limited sensitivity in classical searches due to the limitations of conventional trigger and reconstruction strategies. One such approach is the search for Long-Lived Particles (LLPs), as detailed in the following section.

2.3. Exotic Long-Lived Particle Searches

An unstable particle, $|\psi_i\rangle \sim \exp\left(-ipx\right) = \exp\left(-i\left(Et - \mathbf{px}\right)\right)$, might decay through a perturbation $\mathscr{H}'$ to a particle state $|\psi_f\rangle$. Given the statistical nature of perturbations, the decay probability is an exponential[9], $\exp(-\Gamma t)$, with the decay rate Γ given by Fermi's golden rule [22]:

$$\Gamma = \frac{2\pi}{\hbar} \underbrace{\left|\langle\psi_f|\mathscr{H}'|\psi_i\rangle\right|^2}_{\sim\text{couplings}} \times \underbrace{\rho(E_f)}_{\text{phase-space}} \,. \tag{2.18}$$

The mean lifetime $\tau_0 = \Gamma^{-1}$ of a particle species is the time it takes on average between production and decay of an individual particle of this species. The lifetimes of SM particles span a wide range from $\tau_0 \sim 3 \times 10^{-25}\,\text{s}$ (Z boson) to $\tau_0 \sim 14\,\text{min}$ (neutrons) and even stable particles (*e.g.* protons). Based on the so-called proper lifetime $c\tau_0$, particles can be classified into:

- Promptly decaying particles;
- Particles decaying displaced from their production point;
- Detector-stable particles.

So-called Long-Lived Particles (LLPs) are stable over a non-negligible length scale of a typical HEP detector, and decay with a significant displacement. More precisely, their mean free path

$$L = \beta\gamma c\tau_0 \tag{2.19}$$

is of the order of the detector dimensions, typically $\mathcal{O}(\text{m})$.

Historically, the discovery of particles with relatively high lifetimes, such as the muon ($c\tau_0 \sim \mathcal{O}(10^2\,\text{m})$) and the charged pion ($c\tau_0 \sim \mathcal{O}(10\,\text{m})$), played a pivotal role in the emergence of particle physics as a distinct discipline, since their macroscopic decay signatures allowed these particles to be among the first particle discoveries (studying cosmic rays) that did not fit into the framework of atomic physics. Today, LLP searches gain popularity again, not only because of their unique decay signatures but also because they naturally arise in a variety of BSM extensions, such as SUSY, hidden sector theories, and extensions of the Higgs sector, addressing the hierarchy problem, providing DM candidates and more.

Phenomenological Motivation

A variety of well-motivated BSM extensions predicts LLPs, not because the long lifetimes solve a specific BSM question but the underlying dynamics require suppressed couplings or compressed mass spectra. The resulting couplings and phase-space may, as

9: Starting with a particle i with wavefunction ψ_i at $t = 0$, the probability to find i at a later time $t = t'$ is given by:

$$\begin{aligned} P(t = t') &= \left|\langle\psi_i(t = 0)|\psi_i(t = t')\rangle\right|^2 \\ &= \left|\langle e^{i\mathbf{px}}|e^{-iEt'}e^{i\mathbf{px}}\rangle\right|^2 \\ &= \left|\exp\left(-iEt'\right)\right|^2 \\ &= \exp\left(2\,\text{Im}(E)t'\right) \\ &= \exp\left(-\Gamma t'\right), \end{aligned}$$

where $E \to E - i\Gamma/2$, resulting in the typical Breit-Wigner resonance of unstable particles.

[22]: Dirac (1927), *The quantum theory of the emission and absorption of radiation*

indicated in Equation 2.18, give rise to macroscopic lifetimes [23, 24]:

- The phase-space could be suppressed by a mass compression of the initial and final state particles[10], with some supersymmetric theories being examples for this scenario, in which a small mass splitting is predicted between the lightest neutralino and chargino. The mechanism of mass compression is also partially — but sub-dominantly — responsible for the long lifetime in the presented LLP search, as detailed in Section 6.1;
- Small couplings could arise from approximate symmetries, small mixing angles, or heavy (off-shell) mediators that effectively reduce the couplings[11]; in many hidden sector theories the coupling between SM and BSM particles is exclusively via a BSM mediator mixing with a SM mediator by a small mixing angle; this mediator might additionally be rather heavy and/or couple only weakly to either the SM or the BSM sector, causing a long lifetime of the mediator. A small effective coupling is also the main lifetime handle in the presented LLP search, where a dark sector particles couples to the SM sector via a new BSM mediator mixing with the SM Higgs.

[23]: Knapen et al. (2022), *A guide to hunting long-lived particles at the LHC*
[24]: Genest (2022), *Searching for long-lived BSM particles at the LHC*

10: Similarly, the SM neutron obtains a high lifetime due to the mass compression in the $n \rightarrow pe^{-}\bar{\nu}_e$ decay channel: $m_n \approx m_p + m_e$.

11: Corresponding SM examples are the CKM suppressed K_L decay, and the heavy W and Z propagators in the weak interaction reducing its coupling strength causing the long lifetimes of the muon and pion.

Experimental Motivation

From an experimental perspective, LLPs are attractive as well, since they leave unconventional signatures in the detector which could typically be well separated from SM signatures, reducing the backgrounds. Some typical signatures of LLPs are visualised in Figure 2.4, mainly displaced objects (such as vertices, leptons, jets, photons) but also disappearing tracks, or heavy stable charged particles (HSCPs) resulting in stable tracks throughout the whole detector. This list can be extended by anomalous ionization, out-of-time energy deposits in the calorimeters, emerging jets, kinked tracks, and so on.

It is notable that these signatures are likely missed by conventional searches focussing on prompt signatures, since standard trigger and reconstruction strategies are usually only of limited use for LLP searches; therefore, dedicated trigger and reconstruction algorithms are developed for LLP searches, such as:

- Searches for LLPs decaying in the ATLAS calorimeter facilitate the unusual energy ratio of different calorimeter layers in so-called Calorimeter Energy Ratio (CalRatio) triggers [26];

[26]: ATLAS Collaboration (2022), *Search for neutral long-lived particles in pp collisions at $\sqrt{s}$ = 13 TeV that decay into displaced hadronic jets in the ATLAS calorimeter*

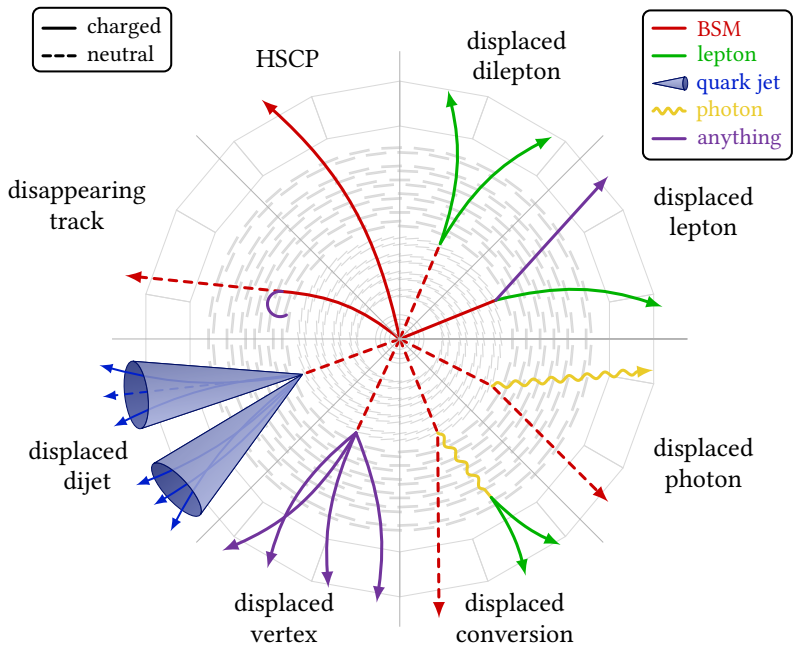

Figure 2.4.: Typical detector signatures of LLPs at LHC experiments. Source code taken from [25]

- Muon Spectrometer Displaced Vertex (MSVtx) searches use a specific reconstruction algorithm for decay vertices in the ATLAS muon spectrometer [27].

Dedicated LLP detectors, such as FASER [28], or proposed experiments like SHiP [29] and ANUBIS [30], offer sensitivity to an even larger phase-space.

[27]: ATLAS Collaboration (2014), *Standalone vertex finding in the ATLAS muon spectrometer*

[28]: Abreu et al. (2024), *The FASER detector*

[29]: Ahdida et al. (2022), *The SHiP experiment at the proposed CERN SPS Beam Dump Facility*

[30]: Bauer et al. (2025), *ANUBIS: Proposal to search for long-lived neutral particles in CERN service shafts*

3. The ATLAS Experiment at the Large Hadron Collider

One promising method to explore new physics beyond the SM is probing the high energy frontier. An experimentally attractive approach for this are hadron colliders, also referred to as discovery machines. Hadrons are compositions of partons, such as the proton $|p\rangle = |uud\rangle$; since the colliding partons might have a momentum-fraction between $0 < x < 1$ of the hadron, the effective centre-of-mass energy takes a broad range of values. In addition, higher energies can be reached due to the higher hadron mass significantly reducing synchrotron radiation. On the other hand, hadron colliders also come with a drawback, namely the omnipresent QCD background, which not only causes radiation damage to the particle detectors, but also needs to be dealt with at the level of on-line and off-line data processing.

The most energetic hadron collider to date is the Large Hadron Collider (LHC) [31] at the European Organization for Nuclear Research (CERN). By studying the collision products with large-scale particle detector experiments, the physical process can be reconstructed, which allows the known SM processes to be tested or new physics to be searched for. A broad physics spectrum is covered by the general-purpose detectors such as ATLAS [32] and CMS [33].

This chapter provides a description of the LHC and the ATLAS experiment, focusing on the calorimetry, trigger systems and readout systems of the detector; furthermore, the reconstruction and calibration of hadronic jets are detailed. This chapter is partially based on [34] and the references therein.

[31]: CERN (2025), *The Large Hadron Collider*

[32]: ATLAS Experiment (2025), *CERN accelerating science*

[33]: CMS Experiment (2025), *CERN accelerating science*

[34]: Heintz (2023), *Improvement of the template fit method for the fake photon background estimation at ATLAS*

3.1. The Large Hadron Collider

The Large Hadron Collider (LHC) is a circular particle collider with a circumference of 27 km. It contains two underground rings (between 45 m and 175 m below the surface) for the two colliding hadron beams, primarily used for proton beams, but also beams of heavy ions, such as lead, oxygen, neon, or xenon. For the remainder of this thesis, only pp collisions are considered.

The schematic structure of the proton beams is visualised in Figure 3.1: each beam contains $N_\mathrm{b} = 2808$ bunches of protons, each consisting of $N_p = \mathcal{O}\left(10^{11}\right)$ protons itself. These bunches are brought to collision at four interaction points where the main LHC experiments ALICE, ATLAS, CMS, and LHCb are located.

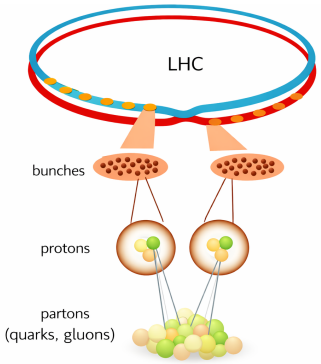

Figure 3.1.: Schematic sketch of the LHC beam structure at different length scales. Sketch based on [35].

The superconducting RF cavities [36] are operating with a frequency of roughly $f = 40\,\text{MHz}$, which means that proton bunches are colliding every 25 ns. The vast majority of the protons of the colliding bunches pass the IPs without any interaction, and about $\mathcal{O}(10-100)$ protons collide per bunch-crossing (BC). However, most of these collisions produce low-energy QCD processes (also known as *minimum-bias* events), while the highly energetic processes of interest typically occur every few BCs. Hypothetical BSM processes are expected to occur even more rarely.

[36]: Boussard et al. (1999), *The LHC Superconducting RF System*

The number of collisions depends on the instantaneous *Luminosity*[12], $\mathscr{L}$, and the cross-section[13], σ:

$$\text{Rate} = \mathscr{L} \times \sigma = \frac{N_p^2 \times f \times N_\text{b}}{4\pi \times \sigma_x \times \sigma_y} \times \sigma. \tag{3.1}$$

12: The instantaneous luminosity represents the particle flux multiplied by the number of collision centres.

13: The cross-section is a measure of the probability for the reaction between two colliding particles.

To maximise the event rate, the beams are heavily focused at the interaction points [37], such that the beam spread in x and y direction, $\sigma_{x,y}$, is minimised. Additionally, since the relevant cross-section depends on the beam energy, $\sigma \equiv \sigma(\sqrt{s})$, higher beam energies may also increase the event rate.

[37]: LHC Machine Outreach (2025), *Collisions and Collision Rates*

High beam energies can be reached due to the strong superconducting magnets which are cooled down with helium to below 2 K. The design centre-of-mass energy and luminosity are $\sqrt{s} = 14\,\text{TeV}$ and $10^{34}\,\text{cm}^{-2}\text{s}^{-1}$, respectively. The LHC has been delivering pp collisions since 2008, with different luminosities and centre-of-mass energies through the different run periods. The ongoing Run-3 is operating with $\sqrt{s} = 13.6\,\text{TeV}$ and instantaneous luminosities up to $2 \times 10^{34}\,\text{cm}^{-2}\text{s}^{-1}$. During the so-called high-luminosity phase (after the LHC upgrades during Shutdown-3), the LHC is expected to deliver up to $\mathscr{L} = 5 \times 10^{34}\,\text{cm}^{-2}\text{s}^{-1}$, which corresponds to an integrated luminosity of about $250\,\text{fb}^{-1}$ per year [38].

[38]: ATLAS Collaboration (2024), *The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3*

3.2. The ATLAS Detector

3.2.1. General Detector Layout

The ATLAS detector [39] was primarily built to discover the Higgs boson, to make SM (precision) measurements, and for BSM searches.

[39]: ATLAS Collaboration (2008), *The ATLAS Experiment at the CERN Large Hadron Collider*

Therefore, good tracking and reconstruction, calorimetry, and muon identification are required; moreover, a large coverage in polar and azimuthal angles is needed to measure missing transverse momentum in the transverse detector plane. Finally, high measurement granularity and efficient trigger systems are essential to handle the high event rates and tackle the QCD backgrounds.

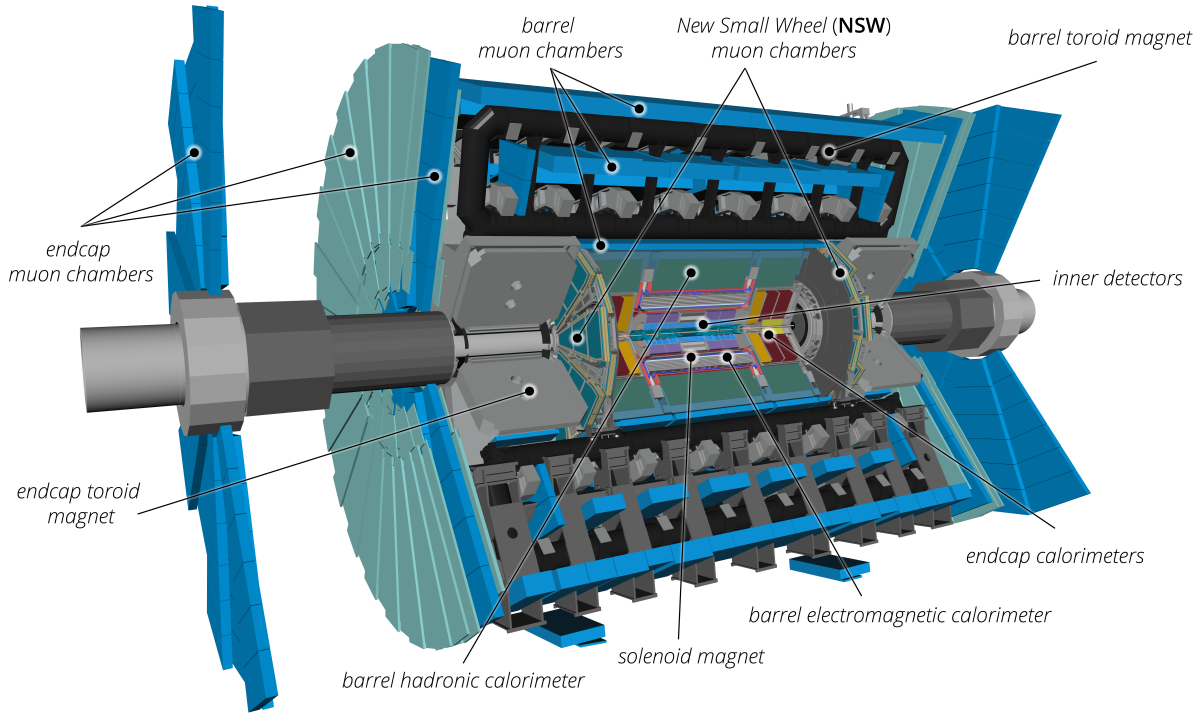

Figure 3.2.: Overview of the ATLAS detector in its Run-3 configuration. The ATLAS detector is a general purpose detector with several sub-detectors attached concentrically (onion-like) around the interaction point; these include the inner tracking detectors, the magnets, the electromagnetic and hadronic calorimeters, as well as the muon spectrometer. Taken from [38].

These features are combined in the ATLAS geometry (as depicted in Figure 3.2) by arranging different layers of sub-detectors concentrically around the IP, making ATLAS a cylindrical detector[14] with an overall size of about 25 m in diameter and 44 m in length.

3.2.2. Overview of the Run-3 Configuration

The *inner detectors* provide high granularity silicon detectors (pixel and strips) for charged particle tracking and vertexing, which allows the high track density to be tackled; moreover the central *solenoid magnet* facilitates the measurement of their momenta based on the track curvature throughout the magnetic field.

The high-precision *barrel* and *end-cap calorimeters* absorb the (full) energy of electromagnetically and hadronically interacting particles, and allow for high precision energy and position measurements of charged and neutral particles, except for neutrinos and muons. The high coverage up to $\eta = 4.9$ allows to measure missing momentum in the transverse plane. More details on calorimetry are provided in Section 3.2.3.

Muon identification and reconstruction is based on their deflection in the field of the *toroid magnets*; they are measured in *tracking*

14: By convention, cylindrical coordinates (r, φ) are used in the plane transverse to the beam (z direction) with origin $(r = 0 = z)$ in the interaction point. Quantities in the transverse plane are typically labelled with index "T", *e.g.*

$$\mathbf{p}_\mathrm{T} = (p_x, p_y), p_\mathrm{T} = |\mathbf{p}_\mathrm{T}| = \sqrt{p_x^2 + p_y^2}. \tag{3.2}$$

For hadron colliders it is also convenient to express the polar angle θ by the pseudorapidity

$$\eta = -\log\left[\tan\left(\frac{\theta}{2}\right)\right], \tag{3.3}$$

and to measure angular distances in terms of

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\varphi)^2}. \tag{3.4}$$

The ATLAS coordinate system is depicted in Figure 3.3.

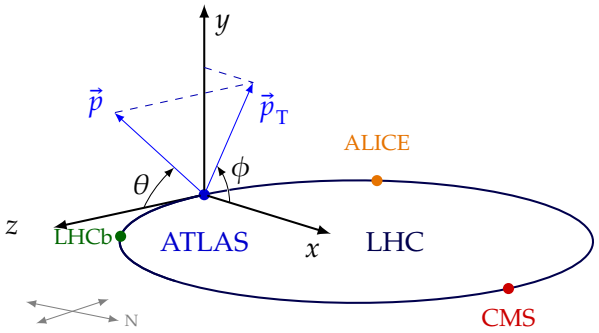

Figure 3.3.: The coordinate convention used in ATLAS is presented, introducing the cartesian axes, as well as the coordinates in the cylindrical/spherical system. Source code modified from [40].

chambers, which are the outermost layer of the ATLAS detector, since muons are the only charged particle species penetrating the calorimeters.

3.2.3. Calorimetry

Unlike "classical" calorimeters, the ones used in high energy physics typically do not measure the energy of a certain reaction by the temperature change (*i.e.* released heat), but by collecting the ionisation charge or scintillating light created in a cascade of secondary particles initiated by electromagnetic and hadronic interactions of a particle with the detector material. Ideally, the energy is fully[15] absorbed in the calorimeter when all such secondaries are stopped, in which case the energy of the incoming particle is proportional to the number of the secondaries. Due to the high granularity, the calorimeters also enable the measurement of jet properties, such as directionality and timing.

15: In reality, a fraction of the energy is lost due to invisible particles in hadronic showers, secondary particles in the absorber material, leakage into the inner detectors and muon spectrometer.

To contain the full shower, calorimeters need to be thick in terms of the radiation and interaction length, which are typical length scales of electromagnetic and hadronic showers, respectively [41].

[41]: Kolanoski et al. (2020), *Particle Detectors*

To keep calorimeters compact and affordable, a common approach in calorimetry is to combine dense absorber material with active detectors (so-called *sampling calorimeters*) as sketched in Figure 3.4. This is also chosen for all ATLAS calorimeters, as shown in Figure 3.5; to account for the distinct nature of electromagnetic and hadronic showers, the different calorimeters use various absorber and detector materials. Based on the detector technology, the ATLAS calorimeter can be classified into Liquid Argon (LAr) and scintillating (tile) calorimeters.

Liquid Argon Calorimeters

LAr is used for calorimeters relatively close to the beam pipe, *i.e.* the ECal barrel, the end-cap as well as forward calorimeters, since LAr is relatively radiation-hard compared to other technologies.

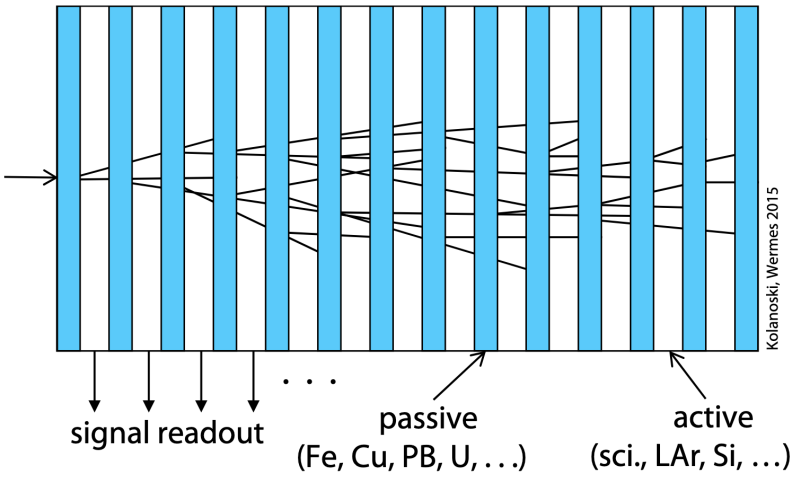

Figure 3.4.: Schematic of sampling calorimeters. Schematic sketch of a sampling calorimeter. Taken from [41].

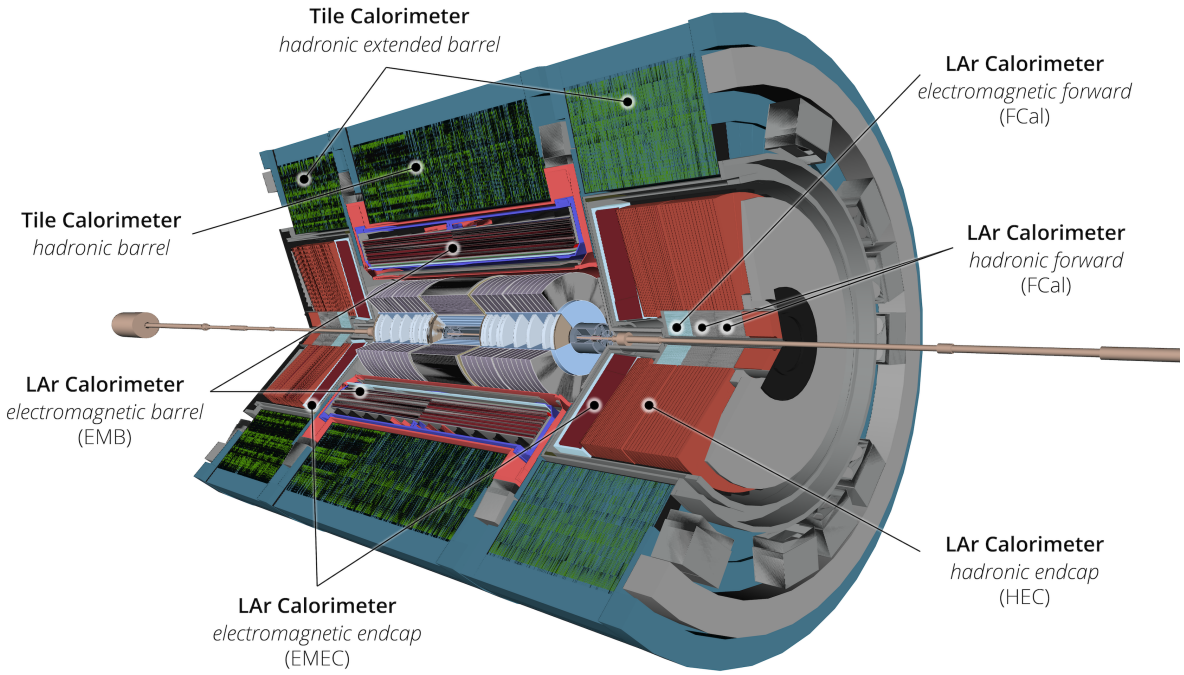

Figure 3.5.: Overview of the ATLAS calorimeters in the Run-3 configuration. The ATLAS calorimeter contains a scintillating tile barrel calorimeter and several LAr calorimeters, such as the electromagnetic barrel, the electromagnetic and hadronic end-cap, as well as the forward calorimeters. Taken from [38].

The LAr calorimeters [42] are operated as ionisation chambers, as sketched in Figure 3.6; that is, the cascade of secondary particles deposit their energy by ionising the Argon, creating free charges which drift towards the absorber plates across which a high voltage is applied.

This, however, comes with the price of rather high drift times of $\mathcal{O}(\mu s)$, much higher than the collision rate of 25 ns, which means that energy deposits of particles from previous pp collisions might contribute to the signal (so-called *out-of-time pile-up*).

[42]: Viehhauser et al. (2024), *Detectors in particle physics: A modern introduction*

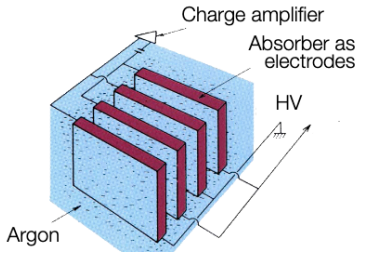

Figure 3.6.: Schematic sketch of LAr calorimeters. Taken from [43].

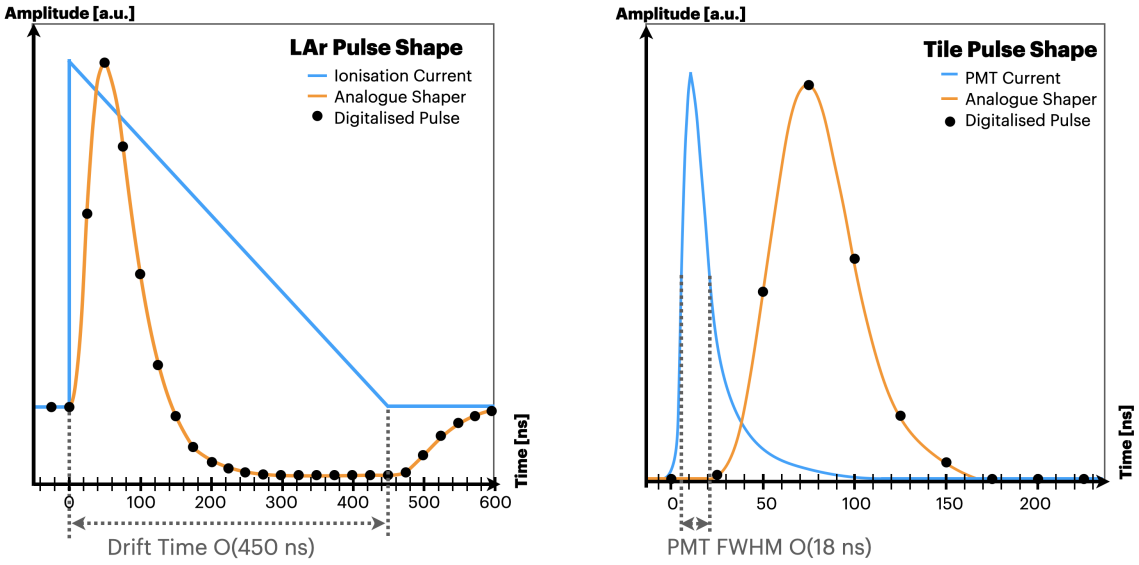

(a) LAr calorimeters.

(b) Scintillating tile calorimeters.

Figure 3.7.: Calorimeter pulses as function of time. The detector signals are shown as well as the shaped pulses. **(a)** The ionisation current is shown, as well as the sharper output pulse after the amplifier and the bipolar shaper. The pulse still has a rise time of more than one BC. Further steps in the processing are the digitalisation, and applying digital filters, calibration and corrections. Pulse shape extracted from [39]. **(b)** The PMT output as well as the broadened pulse and the digitised samples are shown. Pulse shape extracted from [47].

To tackle the out-of-time pile-up, and to assign the energy deposit to the correct BC, a two-staged filter system is used:

- **Bi-polar shaper:** Despite the relatively long duration of the overall signal, the ionisation current is rising relatively quickly. This can be used to fasten the LAr signal processing by integrating this current over a (much) shorter period τ than the drift time [44]. The integration is done by an electric $RC - CR^2$ circuit, where the latter is the integrator. Figure 3.7a shows the ionisation current as well as the output of the pulse shaper as a function of time.
- **Digital filter:** After the analogue shaping, the signal is digitised and the $i = -2, \ldots, 1$ samples, S_i, centred around the peak, $i = 0$, are stored and considered for the digital filtering, where optimal filters, a_i, b_i, are used to sharpen the signal while minimising the effects of noise, using the auto-correlation matrix of noise [45]. This allows for determining the amplitude A and time t of the pulse[16]:

$$A = \sum_i a_i S_i \quad \text{and} \quad t = \frac{1}{A} \times \sum_i b_i S_i \tag{3.5}$$

On average, the in-time and out-of-time pile-up contributions cancel each other since the area of the negative-amplitude pulse part is equal to the area of the positive-amplitude pulse part due to charge conservation.

[44]: Aubert et al. (1991), *Performance of a liquid argon electromagnetic calorimeter with an "accordion" geometry*

[45]: Loch (2024), *Signal reconstruction in the ATLAS calorimeter: from particles to topological cell clusters*

16: Here, the time of the pulse is defined as the time difference between the peak position and the closest digitised sample S_i [46].

Tile Calorimeters

Because of the nature of hadronic showers, which penetrate inherently deeper into material, the HCal needs to be denser or thicker in terms of the interaction length. To facilitate a higher active volume while keeping costs in a reasonable range, it is advantageous to use a technology based on converting the shower energy into photons, since photons have a longer range in transparent materials with respect to the ionisation charges in the LAr calorimeters [42].

Therefore, the hadronic barrel calorimeters use plastic scintillators (formed as tiles) as active material, readout by photomultipliers (PMTs), as illustrated in Figure 3.8.

The hadronic cascades interacts with the plastic scintillator and emits light, generating a pulse in the PMT, as shown in Figure 3.7b. After amplifying the signal, an electronic pulse shaper is used to rise the full-width at the half-maximum of the PMT pulse to $\mathcal{O}(50\,\text{ns})$ to accommodate the LHC frequency and allow for a digitalisation of the peak and its half maxima [49]. To account for out-of-time pile-up, optimal filters (as described in Equation 3.5) are used in seven subsequent readout windows, $i = -3, \ldots, 3$, centred around the peak.

[42]: Viehhauser et al. (2024), *Detectors in particle physics: A modern introduction*

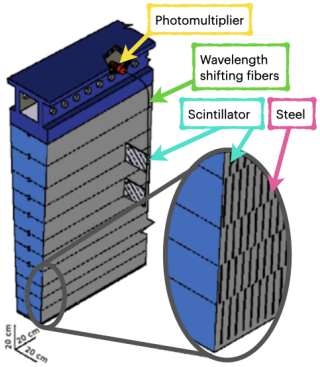

Figure 3.8.: Overview of the ATLAS tile calorimeter. Modified from [48].

[49]: Jen-La Plante et al. (2010), *Pulse shapes for signal reconstruction in the ATLAS Tile Calorimeter*

3.3. Trigger and Data Acquisition

As shown in Figure 3.9, the total rate of pp scattering at the LHC is, for a luminosity of $10^{33}\,\text{cm}^{-2}\text{s}^{-1}$, about $10^8\,/\text{s}$, which is far beyond the technical capabilities of the data acquisition system. However, the total rate is also more than 5 orders of magnitude higher than the rate of the interesting SM processes, such as the electroweak and Higgs processes, and even more for hypothetical BSM processes.

Therefore, an on-line decision on which events should be permanently stored for off-line analyses, is needed. Here, it is crucial to reject as many events produced by background processes as possible, such that the rates can be handled by the trigger and data acquisition system, while maximising the signal acceptance.

The Trigger and Data Acquisition (TDAQ) is currently implemented in a two-staged sequence in ATLAS.

The Level-1 (L1) uses reduced granularity information from hardware systems — in particular, the calorimeters (L1Calo) and muon chambers (L1Muon). Moreover, the Level-1 Topological Trigger (L1Topo) allows for topological combinations of multiple L1 objects. Here, reduced granularity means that several readout channels are grouped together to so-called *supercells*, of which "only" ten exist per 0.1×0.1 block in $\eta \times \varphi$ (formerly known as a trigger tower).

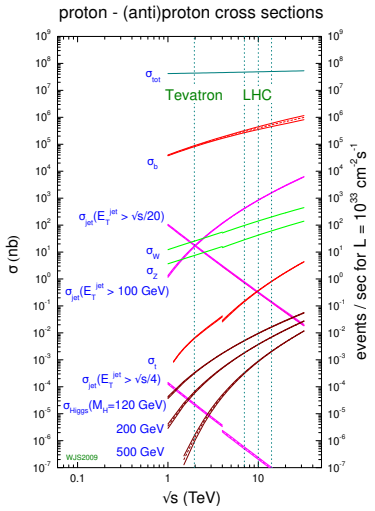

Figure 3.9.: The proton-(anti)proton cross-sections are presented as a function of the centre-of-mass energy. Taken from [50].

These supercells are processed with different feature extractors to find candidates of electrons, photons, hadronically decaying taus, jets of different radii ($R = 0.4, 1.0$), and to obtain event level quantities such as total and missing transverse energies.

The central trigger processor (CTP) makes a decision on whether to accept or reject a particular event at L1. An L1-Accept (L1A) implies further processing at the Higher-Level Trigger (HLT) and, temporarily, the storage of data, where the rate of L1As is limited to about 100 kHz. The CTP also needs to take into account the *deadtime* of the L1 trigger: after an L1A has been sent to the HLT, accepting further events in the next 2.5 µs is not possible due to TDAQ being busy with processing the L1A.

Once an event has been sent to the HLT, on-line algorithms in the Athena framework [51] are used in the region of interest for higher level event reconstruction. This allows to reject further background events, and to reduce the rate to about 3 kHz.

[51]: ATLAS Collaboration (2019), *Athena*

In case an event is also accepted by the HLT, the data is eventually readout and stored permanently for off-line analyses.

3.4. Reconstruction of Physics Objects

Physics objects, such as electrons, photons, taus, jets, and muons are reconstructed based on the detector signals, which is, to first order, done according to the typical interaction of these objects within the sub-detectors, as visualised in Figure 3.10. This allows to construct a list of particle *candidates*, from which the list of reconstructed particles is selected with higher-level methods and variables, such as the shower shape variables [52].

Because of the distinct nature of the different physics objects, individual algorithms are used for their reconstruction. The most relevant objects for this analysis are jets and missing transverse momentum, MET, which are reconstructed as described below. Descriptions of the reconstruction of the other objects can be found in [52] and the references therein.

[52]: Weber (2023), *Measurement of Vector Boson plus Jet Production Cross Sections and Dark Matter Interpretation with the ATLAS Detector*

Topological Jets

Hadronic and electromagnetic particle showers are reconstructed as jets based on topological combinations of energy deposits in neighbouring cells in the calorimeters (so-called topo-clusters), which are input for the jet clustering — typically done with the anti-k_t algorithm [54].

[54]: Cacciari et al. (2008), *The anti-ktjet clustering algorithm*

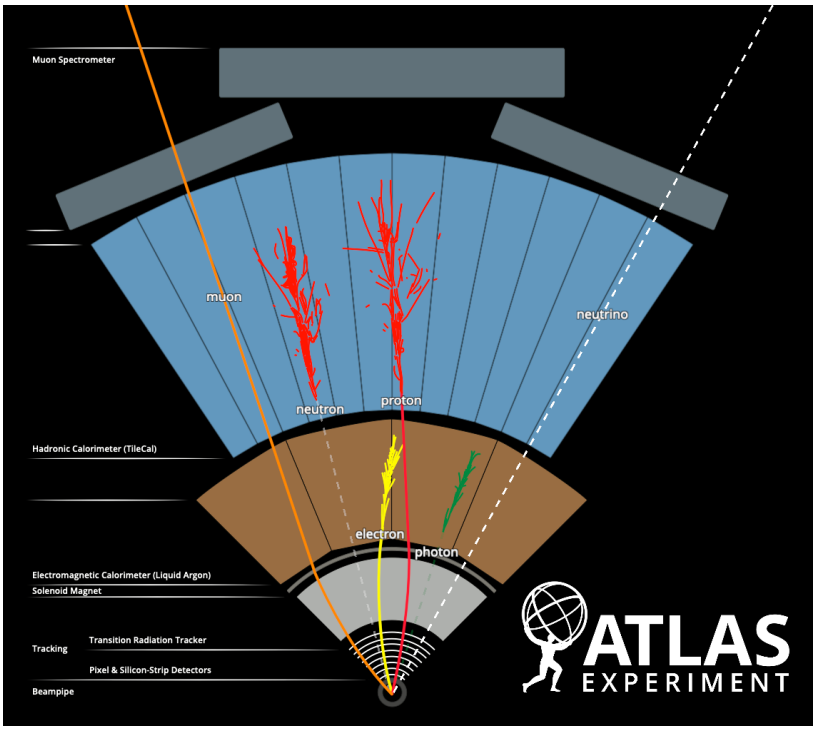

Figure 3.10.: Particle reconstruction and identification in the ATLAS sub-detectors. The particle species are reconstructed and identified based on their typical signatures in the sub-detectors. Taken from [53].

Subsequently, the jet energy is calibrated and jet vertex tagging helps to identify the amount of energy a jet is obtaining from the hard scatter event.

In principle, charged particle tracking can also be used for improving the jet energy resolution (especially for low energetic particles) in the so-called particle-flow algorithm [55], but given that the signal jets in the presented analysis are not expected to contain particles with tracks in the inner detector, only so-called topological jets are discussed here, whereas details on the particle flow algorithm can be found in [55].

The topo-clusters [34, 45, 56, 57] are formed with the *4-2-0 algorithm* based on the cell-energy significance

$$\zeta_{\text{cell}} = \frac{E^{\text{EM}}_{\text{cell}}}{\sigma_{\text{noise,cell}}}, \tag{3.6}$$

where the cell noise, $\sigma_{\text{noise,cell}}$, is the root-mean-square of the cell energies in randomly triggered events, with contributions coming from noise of the electronics and from pile-up. The cell energies, $E^{\text{EM}}_{\text{cell}}$, are measured at the so-called "EM" scale, *i.e.* without correcting for the calorimeter non-compensation.

In the first step, calorimeter cells with significance, $|\zeta_{\text{cell}}| > 4$ are used as seeds for the clustering. Iteratively, neighbouring cells with $|\zeta_{\text{cell}}| > 2$ are matched to the cluster — both in the η and φ directions as well as among the different calorimeter layers. Once the iteration terminates, all neighbouring cells with $|\zeta_{\text{cell}}| > 0$

[55]: ATLAS Collaboration (2017), *Jet reconstruction and performance using particle flow with the ATLAS Detector*

[34]: Heintz (2023), *Improvement of the template fit method for the fake photon background estimation at ATLAS*
[45]: Loch (2024), *Signal reconstruction in the ATLAS calorimeter: from particles to topological cell clusters*
[56]: ATLAS Collaboration (n.d.), *Jet energy measurement with the ATLAS detector in proton-proton collisions at* $\sqrt{s} = 7$ TeV
[57]: ATLAS Collaboration (2017), *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*

are added to the envelope of the topo-cluster. Negative energies (mainly from out-of-time pile-up) are considered in all three steps of the cluster formation, which improves the noise suppression, since positive and negative noise energies cancel each other on average.

Since the cluster size is not restricted in this algorithm, a procedure to potentially split the initial clusters into several ones is needed; to do so, the cluster formation as explained above is repeated with choosing different seeds, namely all local maxima inside the initial cluster. Here, a local maximum is defined as a cell with $E_{\text{cell}}^{\text{EM}} \geq 500\,\text{MeV}$ and with all neighbouring cells being less energetic.

In Run-3, the seeding requirement, $|\zeta_{\text{cell}}| > 4$, has been extended [58], by also requiring the cell time, as calculated in eq. 3.5, to be within ±12.5 ns. Cells failing this timing cut are also excluded in the iterative addition of neighbouring cells to clusters with other seeds. This timing cut is used in the off-line reconstruction[17], in order to reduce the contribution out-of-time pile-up. To keep sensitivity for out-of-time signals from exotic LLPs, an upper limit of this timing cut is implemented, which means that cells with positive energy and $|\zeta_{\text{cell}}| > 20$ are not affected by the timing cut:

$$\left\{\begin{matrix}|\zeta_{\text{cell}}| > 4 \\ |t_{\text{cell}}| < 12.5\,\text{ns}\end{matrix}\right\} \quad \text{or} \quad \left\{\begin{matrix}|\zeta_{\text{cell}}| > 20 \\ E_{\text{cell}} > 0\end{matrix}\right\}. \tag{3.7}$$

[58]: ATLAS Collaboration (2024), *Improving topological cluster reconstruction using calorimeter cell timing in ATLAS*

17: Currently, the timing cut is *not* used at the trigger level.

Each topo-cluster gets assigned the summed energies of all included cells, and the η, φ direction is calculated as an energy-weighted average of the contained cells:

$$E_{\text{clust}}^{\text{EM}} = \sum_{\text{cells}} w_{\text{cell},i}^{\text{geo}} \times E_{\text{cell},i}^{\text{EM}}, \tag{3.8}$$

$$\eta_{\text{clust}}\left(\varphi_{\text{clust}}\right) = \frac{\sum_{\text{cells}} w_{\text{cell},i}^{\text{geo}} \times \left|E_{\text{cell},i}^{\text{EM}}\right| \times \eta_{\text{cell},i}\left(\varphi_{\text{cell},i}\right)}{\sum_{\text{cells}} w_{\text{cell},i}^{\text{geo}} \times \left|E_{\text{cell},i}^{\text{EM}}\right|}, \tag{3.9}$$

where $w_{\text{cell}}^{\text{geo}}$ is a geometrical weight accounting for the splitting of clusters.

Also, the time of the cluster is calculated based on the cell time of cells with significance greater than 2:

$$t_{\text{clust}} = \frac{\sum_{\text{cells}|\zeta>2} \left(w_{\text{cell},i}^{\text{geo}} \times E_{\text{cell},i}^{\text{EM}}\right)^2 \times t_{\text{cell},i}}{\sum_{\text{cells}|\zeta>2} \left(w_{\text{cell},i}^{\text{geo}} \times E_{\text{cell},i}^{\text{EM}}\right)^2}. \tag{3.10}$$

Topo-clusters with positive energy, $E_{\text{clust}}^{\text{EM}} > 0$, are then clustered to

jets [59]. Jets are useful objects, since they can be measured in the detector, but they can also be theoretically calculated, as described in Chapter 4, and be related to the final state particles. Different algorithms can be used for the jet clustering; in general, these algorithms should be stable against soft and collinear radiations, which is also known as infrared and collinear safe [60].

[59]: Cacciari (2018), *Standard Model: Jet Physics*

[60]: Buckley et al. (2021), *Practical collider physics*

One way of achieving this is combining the topo-clusters iteratively utilising a distance measure, such as

$$d_{ij} = \min\left(\frac{1}{p_{\mathrm{T},i}^2}, \frac{1}{p_{\mathrm{T},j}^2}\right) \times \frac{\Delta R_{ij}^2}{R_{\mathrm{jet}}^2}, \quad d_{i\mathrm{B}} = \frac{1}{p_{\mathrm{T},i}^2}, \tag{3.11}$$

which is used in the so-called anti-k_t [54] algorithm. Here, d_{ij} is a measure for the distance between two topo-clusters i and j, and $d_{i\mathrm{B}}$ for the distance between a topo-cluster i and the beam (B). R_{jet} scales the jet size, and needs to be sufficiently high to contain the full shower, but it also must not be too large, to avoid contributions from other objects [61].

[54]: Cacciari et al. (2008), *The anti-ktjet clustering algorithm*

[61]: Balunas (2025), *Making the Most of the Mess: Jets and MET at ATLAS*

Iteratively, topo-clusters are combined with topo-clusters with smallest distance d_{ij} or $d_{i\mathrm{B}}$, which is terminated once no further topo-clusters are left. The individual topo-clusters that are combined during the jet clustering are also called jet *constituents*. Instead of using detector signals (such as topo-clusters), the same algorithm as described above can be applied on four momenta of truth (*i.e.* MC) particles.

An efficient version of this algorithm is based on the FastJet lemma, which reduces the number of required d_{ij} and $d_{i\mathrm{B}}$ calculations from $\mathcal{O}(N^3)$ to $\mathcal{O}(N \log N)$ by relating the distance d_{ij} to the geometrical distance, and efficiently searching the (nearest) neighbour(s), as implemented in the FastJet framework [62].

[62]: Cacciari et al. (2012), *FastJet User Manual*

The jet energy at the EM scale, E_j^{EM}, is simply the sum of the cluster energies, given by Equation 3.8. However, due to non-compensation, MC-data differences *etc.*, the jet energy needs to be calibrated before it can be interpreted in the context of final state particle properties. The calibration

$$E_j^{\mathrm{cal}} = \frac{E_j^{\mathrm{EM}} - \mathcal{O}(N_{\mathrm{PV}})}{\mathcal{F}(E_j^{\mathrm{EM}}, N_{\mathrm{PV}})} \tag{3.12}$$

includes pile-up corrections, calibration of the detector response, and *in-situ* corrections to account for differences between data and MC. Further jet selection criteria, and tagging, such as flavour tagging, vertex tagging, are done at the analysis level.

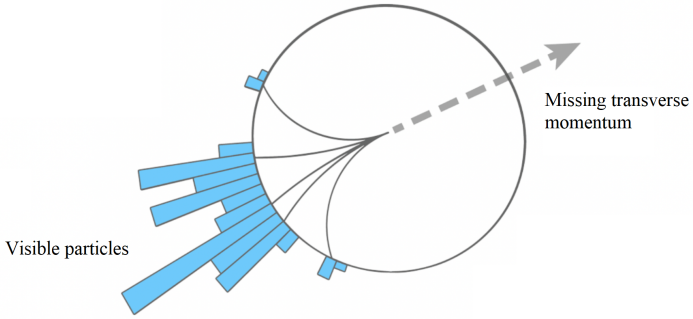

Figure 3.11.: Sketch of the energy balance in the transverse plane. Taken from [63].

Missing Transverse Momentum

Particles that do not interact via the electromagnetic or strong forces, such as neutrinos, or hypothetical BSM particles, do not interact with the detector. Therefore, a direct detection of these particles is not feasible. Still, their presence can be inferred by utilising the fact that the initial state partons' momenta are fully in the longitudinal direction (*i.e.* parallel to the beam axis); in contrast, there is no momentum in the transverse direction. Due to energy conservation, the same balance exists for the final state. That is, the sum of the transverse momenta of all final state particles vanishes. Splitting final state particles into visible and invisible ones yields

$$0 = \sum_i \mathbf{p}_{\mathrm{T},i}^{\mathrm{visible}} + \underbrace{\sum_j \mathbf{p}_{\mathrm{T},j}^{\mathrm{invisible}}}_{=:\,\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}} \Rightarrow \mathbf{p}_{\mathrm{T}}^{\mathrm{miss}} = -\sum_i \mathbf{p}_{\mathrm{T},i}^{\mathrm{visible}}, \tag{3.13}$$

as sketched in Figure 3.11.

The missing transverse momentum[18] can be explicitly written down as

$$\begin{aligned}\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}} = -\Bigg(&\sum_{\text{jets}} \mathbf{p}_{\mathrm{T}}^{j} + \sum_{\text{electrons}} \mathbf{p}_{\mathrm{T}}^{e} + \sum_{\text{muons}} \mathbf{p}_{\mathrm{T}}^{\mu} + \dots \\ &\dots + \sum_{\text{taus}} \mathbf{p}_{\mathrm{T}}^{\tau} + \sum_{\text{photons}} \mathbf{p}_{\mathrm{T}}^{\gamma} + \sum_{\text{TST/CST}} \mathbf{p}_{\mathrm{T}}^{\text{track/cluster}}\Bigg).\end{aligned} \tag{3.15}$$

Here, the last term represents the so-called soft term, which accounts for visible particles that do not enter reconstructed objects; it could either be calculated based on tracks (TST) or clusters (CST) that are not used for object reconstruction, where the former is the default choice.

18: Within ATLAS, missing transverse momentum and missing transverse energy (MET) are commonly used as synonyms:

$$\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}} \equiv \mathbf{E}_{\mathrm{T}}^{\mathrm{miss}} \equiv \not{\mathbf{E}}_{\mathrm{T}}. \tag{3.14}$$

4. Simulation of Proton-Proton Collisions

Simulations are crucial for understanding the detector response, optimising on- and off-line algorithms, such as reconstruction, they help optimising the physics analyses, and allow to benchmark new BSM physics phenomena by comparing theoretical predictions to the recorded data.

This chapter describes the basics of how to simulate an event, from the pp collision to the detector readout. The simulation chain is visualised and compared to the data processing workflow in Figure 4.1; it contains the event generation, as well as the propagation of the stable particles through the ATLAS detector, their interaction with the detector material, overlay of pile-up events, and the detector readout (digitisation).

Besides the "standard" simulation steps, this chapter also introduces a user-specific workflow for simulating exotic BSM signatures with out-of-time decaying LLPs.

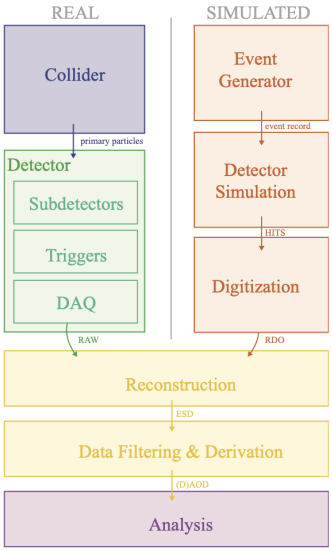

Figure 4.1.: Overview of the workflow in modern HEP analyses. Taken from [60].

4.1. Event Generation

Due to the nature of QCD, pp collisions at the LHC can be split into

- The hard process, *i.e.* QCD processes with rather high momentum transfer $Q \gg \Lambda_{\text{QCD}} \approx 220\,\text{MeV}$, where the strong coupling constant α_s is sufficiently small, such that it can be described by perturbative calculations;
- The parton shower (PS), *i.e.* the radiation of additional gluons and gluon-parton splitting, which also includes low Q processes, where perturbation theory breaks down, and showering models are used instead;
- The hadronisation process, where the partons form colourless hadrons, and unstable hadrons further decay into stable ones.

This is drawn schematically in Figure 4.2, and detailed in the following sections, greatly inspired by [64, 65].

[64]: Mitov (2024), *QCD (in the LHC precision era)*
[65]: Webber (2011), *Parton shower Monte Carlo event generators*

Hard Scattering Process

In proton collisions at the hard scale, $Q \gg \Lambda_{\text{QCD}}$, actually not the protons themselves (visualised as the red blobs in Figure 4.2) collide, but some of the partons inside them (black blob).

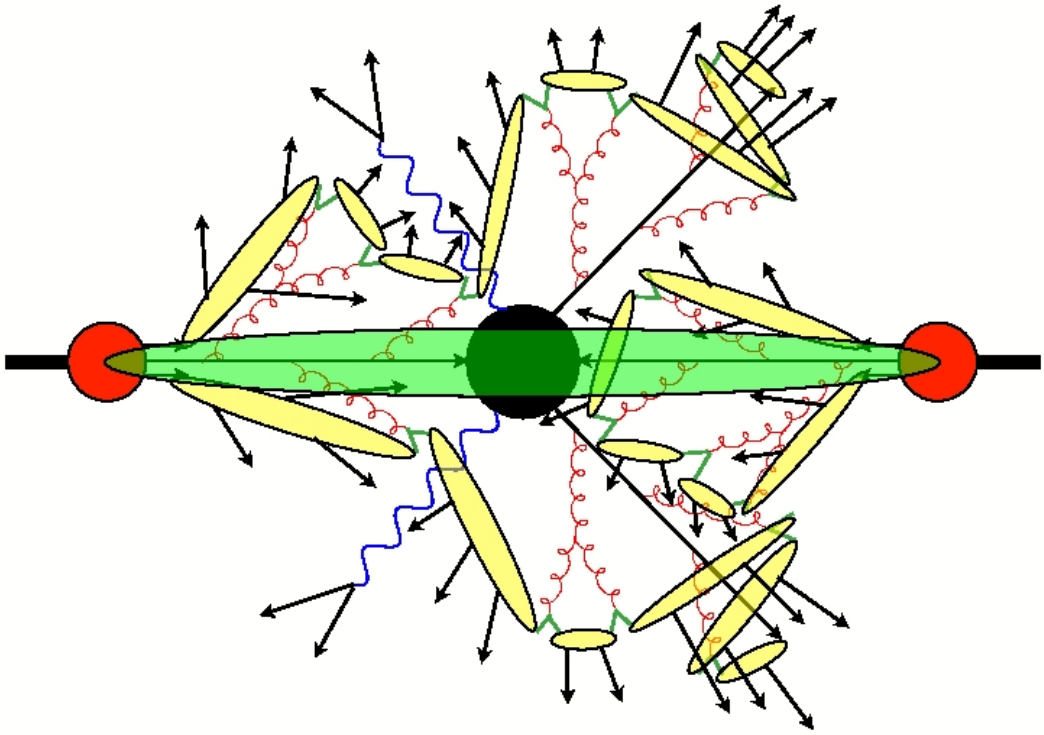

Figure 4.2.: Overview of the MC event generation. Schematic overview of components in the MC event generation chain. These include the hard scattering, the PS, the hadronisation, and the underlying event. Taken from [65].

The parton-level cross-section for $i + j \rightarrow 1 + ... + n$ with initial state partons i, j having momenta $p_{1,2}$, which can be expressed as a fraction $x_{1,2}$ of the colliding protons' momenta, and the final state momenta $q_1, ..., q_n$ reads

$$\mathrm{d}\sigma_{ij\rightarrow 1+...+n} = \frac{1}{2\hat{s}} \left|M_{ij}(p_1, p_2 \rightarrow q_1, ..., q_n)\right|^2 \mathrm{d}\Phi_n. \tag{4.1}$$

Here, $\hat{s} = x_1 x_2 \times s$ is the effective center-of-mass energy of the colliding partons, M_{ij} is the spin- and colour-averaged matrix element, calculated perturbatively at a fixed order in α_s, and

$$\mathrm{d}\Phi_n = (2\pi)^4 \, \delta^{(4)} \left(p_1 + p_2 - \sum_{l=1}^{n} q_l \right) \times \prod_{l=1}^{n-1} \frac{\mathrm{d}^3\mathbf{q}_l}{(2\pi)^3 \times 2E_l} \tag{4.2}$$

is the phase-space of the final state momenta.

This allows to express the $pp \rightarrow 1 + ... + n$ cross-section by summing up the parton-level cross-section for the different partons, $i, j \in \{g, u, \bar{u}, ...\}$, convoluted with the corresponding parton

density functions[19] $f_{i,j}$:

$$\mathrm{d}\sigma_{pp\to 1+\ldots+n} = \sum_{i,j} \int_0^1 \mathrm{d}x_1 \int_0^1 \mathrm{d}x_2 \, f_i(x_1) f_j(x_2) \, \mathrm{d}\sigma_{ij\to 1+\ldots+n}. \tag{4.3}$$

19: The parton density functions are derived from data, and scaled to different energy scales, Q, by the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) evolution equations [60].

By parametrising the final state momenta $\mathbf{q}_1, \ldots, \mathbf{q}_n$ as independent variables $z_1, \ldots, z_{3n}$ with $0 \leq z_k \leq 1$, $\forall k \in [1, 3n]$ (*e.g.* normalised energies, and cosine of angles), the scalar products $\mathbf{p}_{1,2} \cdot \mathbf{q}_l$ and $\mathbf{q}_l \cdot \mathbf{q}_m$ in the matrix element can be rewritten by the z_k, and the phase-space can be exchanged by the corresponding Jacobian. The cross-section then becomes

$$\begin{aligned} \mathrm{d}\sigma_{pp\to 1+\ldots+n} &= \int_0^1 \mathrm{d}x_1 \int_0^1 \mathrm{d}x_2 \left(\prod_{k=1}^{3(n-1)} \int_0^1 \mathrm{d}z_k \right) w(x_1, x_2, z_1, \ldots \\ &\qquad \ldots, z_{3(n-1)}), \\ &= \int_\Omega w(\vec{x}) \, \mathrm{d}x^{3n-1}, \quad \vec{x} = (x_1, x_2, z_1, \ldots, z_{3(n-1)}), \end{aligned} \tag{4.4}$$

where $\Omega = [0,1]^{3n-1}$ is the unit hypercube and $w(\vec{x})$ is a weight function absorbing the sum over i, j, the parton density functions, the matrix element, and the Jacobian.

Doing these integrations analytically is often not practical because of the high number of $3n-1$ dimensions, but using Monte Carlo (MC) integration techniques has proven to be an efficient alternative; that is, by evaluating the weight function individually for a sufficient number N of $\vec{x}_{(i)}$, the cross-section can be approximated as the average of the weights, because

$$\mathrm{d}\sigma_{pp\to 1+\ldots+n} = \int_\Omega w(\vec{x}) \, \mathrm{d}x^{3(n-1)} = \lim_{N\to\infty} \left(\frac{1}{N} \sum_{i=1}^{N} w\left(\vec{x}_{(i)}\right) \right). \tag{4.5}$$

The values for $x_{1,2}, z_k$ are either chosen randomly or – more efficiently – according to probability density functions approximating $w(\vec{x})$, *e.g.* using the VEGAS algorithm [66].

[66]: Peter Lepage (1978), *A new algorithm for adaptive multidimensional integration*

That is, the cross-section integration is simplified to evaluating the parton density functions and the squared matrix element for a sufficiently large amount of $\vec{x}_{(i)}$. A set of unweighted events following the physical cross-section is sampled by accepting/rejecting the events based on their weight. These methods are implemented in event generators, such as MadGraph [67].

[67]: Alwall et al. (2014), *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*

Parton Showering and Hadronisation

Partons from the initial and final state of the hard process may radiate gluons[20], as visulaised by the orange lines in Figure 4.2. These gluons radiate further gluons or produce $q\bar{q}$ pairs, and so on, resulting in a parton shower (PS). Complicated event topologies with high jet multiplicities and jet sub-structures are omnipresent in pp colliders, and make it impractical to calculate the PS with perturbation theory. Instead, the PS is modelled recursively using parton branching probabilities, and is applied on top of the hard-process, *e.g.* using PYTHIA8 [68]. The branching process typically evolves from a higher Q parton to partons with lower Q (*i.e.* increasing coupling α_s), until perturbation theory eventually breaks down at $Q \approx \mathcal{O}(1\,\text{GeV})$, causing the partons to start hadronising, *i.e.* forming colourless hadrons; subsequently, unstable hadrons decay further to stable hadrons. The hadronisation process is described by models such as the *string model*, as visualised by the yellow ellipses in Figure 4.2.

In addition to radiations from the partons involved in the hard process (as discussed above), also the remaining partons of the colliding protons will re-organise to colourless hadrons, since the proton remnant is not colourless anymore after the hard scattering. Analogously to the showering discussed above, this will result in additional hadrons in the final state, known as the underlying event, which is indicated as the green band in Figure 4.2.

20: Likewise an accelerated electromagnetic charge radiates photons.

[68]: Bierlich et al. (2022), *A comprehensive guide to the physics and usage of PYTHIA 8.3*

Multi-Jet Merging and Matching with Parton Shower

QCD multi-jets with N_{jets} up to $\mathcal{O}(10)$ are omnipresent at the LHC, and they are also expected to be the dominant background in the presented analysis. Therefore, using leading order (LO) perturbation theory and a subsequent PS, might not be sufficient for the signal simulation. In particular, hard ISR jets might affect the studied signature and cannot be neglected.

There are two common approaches of simulating multiple hard partons/jets, while still properly simulating the internal jet structure:

- Calculating the hard scattering at next-to-leading order (NLO) perturbation theory (which includes real and virtual corrections 2$^{\text{nd}}$ order), and subsequently simulating the PS;
- Generating a sample with LO perturbation theory, generating another sample with LO perturbation theory, but additional (real) partons/jets at the matrix level[21]. After merging them

21: To visualise this concept, a simple example would be producing $Z \to \mu^+\mu^-$ and merging it with $Z \to \mu^+\mu^- + 1\,\text{jet}$, where "jet" refers to a gluon or a quark.

according to their production cross-section, the PS is simulated.

Both approaches mentioned above require a matching process during the showering, to avoid double counting events where additional hard jets emerge from a PS.

It depends on the situation which approach is preferred, since both come with their advantages; the main difference is that the NLO calculations include real *and* virtual corrections, allowing for proper cancellation of divergences, and therefore to describe NLO corrections also at low Q and for collinear jets. In contrast, the LO merging only adds real corrections, and needs to avoid additional low Q jets and jets that are collinear. The LO merging, on the other hand, allows to accurately describe the kinematics across several hard jets at matrix element precision. Since the presented analysis will mainly be affected by additional (hard) ISR jets, the LO merging + PS approach is preferred.

In particular, samples with different jet multiplicities, $pp \rightarrow$ BSM + 0 jets and $pp \rightarrow$ BSM + 1 jet, are generated in MadGraph, and combined using the MLM merging scheme [69]. The two processes are mixed according to their production cross-sections, and subsequently, the PS is overlaid to also model the internal jet structure.

[69]: Lavesson et al. (2008), *Merging parton showers and matrix elements—back to basics*

Although quite limited to date, it is also possible to merge NLO samples with other NLO samples with additional jets at matrix level, which would be more accurate compared to the two methods described above [70].

[70]: Lönnblad (2002), *Correcting the Colour-Dipole Cascade Model with Fixed Order Matrix Elements*

4.2. Detector Response

The detector response and readout are simulated within the Athena framework [51, 71], which interfaces to a Geant4-based [72–74] simulation of the ATLAS detector geometry, using databases providing a detailed description of the layout and run-dependent operating conditions; the detector-stable particles from the event generation are taken as input for the following simulation chain:

- *Simulation*: propagation of the particles through the ATLAS detector, and simulation of their deflection in the magnetic field and of their interaction with the detector material;
- *Digitisation*: transformation of the deposited energies (including PU overlay) into pulses of currents/voltages;
- *Reconstruction* and simulation of *trigger decision* are identical to data processing, as described in Sections 3.3 and 3.4.

[51]: ATLAS Collaboration (2019), *Athena*
[71]: ATLAS Collaboration (2010), *The ATLAS Simulation Infrastructure*
[72]: Allison et al. (2016), *Recent developments in Geant4*
[73]: Allison et al. (2006), *Geant4 developments and applications*
[74]: Agostinelli et al. (2003), *Geant4—a simulation toolkit*

Due to the high CPU time consumption of the simulation, the samples are typically provided centrally for the whole ATLAS collaboration. However, due to the unusual signature with significant out-of-time contributions, a user-specific step is inserted between simulation and digitisation, which takes special care of the out-of-time hits.

This section describes the "standard" simulation and digitisation, based on [71], followed by an outline of how to account for the out-of-time part in Section 4.3.

[71]: ATLAS Collaboration (2010), *The ATLAS Simulation Infrastructure*

Simulation

The generated particles are propagated through the ATLAS detector by Geant4. Here, the deflection of charged particles in the magnetic fields of the ATLAS detector is done numerically. Moreover, hadronic and electromagnetic interactions between the particle and the detector material (active and passive) are simulated using models such as

- Quark–gluon string precompound model for inelastic reactions of high energy particles with nuclei;
- Bertini model for simulating the intranuclear cascade of hadronic showers;
- Multiple Coulomb scattering.

The energy deposits within the active material are stored as so-called *hits*, which contain the magnitude of the energy deposition, as well as the combination of time and position in form of the readout time,

$$\text{Hit} = (E, t_{\text{readout}}) \,, \tag{4.6}$$

$$t_{\text{readout}} = t - r/c; \quad r = |\mathbf{r}| = \sqrt{x^2 + y^2 + z^2}, \tag{4.7}$$

where $(t, \mathbf{r})$ is the lab frame four-vector of the energy deposit.

An extensive list of additional actions is performed during the simulation, as detailed in [71], including "standard" actions such as vertex smearing, but also optional and user-specific actions, such as decaying LLPs if they are considered as stable in the event generation, or stopping the simulation of particles meeting certain criteria.

[71]: ATLAS Collaboration (2010), *The ATLAS Simulation Infrastructure*

In the last step of event processing, a hit-merging is done to reduce the size of the hit file — in particular for hits in the calorimeter.

Subdetector	Simulation window [ns]	No. Bunch crossings (25 ns bunch spacing)
BCM	$-50, +25$	4
Pixel trackers	$-50, +25$	4
SCT	$-50, +25$	4
TRT	$-50, +50$	5
LAr calorimeter	$-801, +126$	38
Tile calorimeter	$-200, +200$	17
Muon chambers	$-1000, +700$	69

Table 4.1.: Active time window of sub-detectors for pile-up overlay. Due to the wide readout windows of some sub-detectors, pile-up is overlaid over multiple BCs. Taken from [71].

Digitisation

The hits from the simulation are converted into so-called *digits* for each readout cell, which, depending on the sub-detector, either indicate whether the cell has exceeded a certain current or voltage threshold in a particular time window, or includes the signal shape (*e.g.* pulses in the calorimeter). While the readout time of *hits* may take any value, *digits* are only considered for the detector readout if they are within a particular time window and exceed a particular voltage or current threshold.

To mimic the LHC environment realisticly, the signal event (*i.e.* the hits from the hard scatter event, generated and simulated as explained above) need to be overlaid with hits from pile-up events, *i.e.* simultaneous low energy inelastic pp collisions; moreover, additional hits come from beam-induced backgrounds (BIBs), such as beam halo/gas events. The number of pile-up events is sampled based on the run-dependent distribution of simultaneous pp collisions in data.

Since the readout spans several BCs for some sub-detectors, it is crucial to also account for hits from previous BCs. This is achieved by superimposing multiple pile-up/background samples shifted by various multiples of the BC window (25 ns) to cover the entire readout window, as summarized in Table 4.1.

Thereafter, the detector signal is created by sub-detector specific algorithms, basically by propagating charges from the active detector to the readout electronics and collecting the charges there. Also, detector noise and cross-talk effects are simulated, and run-specific conditions can be set, such as dead channels.

The final detector readout has an output format very similar to the raw data format. In fact, the simulated output could easily be converted to the data format and vice versa. That allows to process the trigger decision (L1 and HLT), as well as the reconstruction of

physics objects, such as jets, with the same algorithms, as described in Sections 3.3 and 3.4.

4.3. Simulating Exotic Long-Lived Particles

Usually, the event generation is point-like and therefore includes neither particle propagation nor the detector layout. This is perfectly fine for SM processes, where the Q scale is sufficiently high such that the length scale $\lambda \sim 1/Q$ is small enough to be negligible. However, special care is required for BSM processes that include LLPs with decay length at detector-scale (*i.e.* several centimeters or more).

In general, the LLPs can be treated as stable at the generator level, and their decay can be implemented with an Geant4 extension during the simulation of the detector response. However, if the LLP does not interact with the detector, its decay can still be described in the event generation, and the decay time can be added to the output file of the generation; in this manner, the decay time and position are properly propagated to the simulation of the detector response [75].

[75]: Marshall et al. (2025), *Internal Communication*

The situation becomes more complex when the LLP production and decay structure split between two BCs, $N-1$ and N; while the standard propagation of the detector-stable particles through and their interaction with the detector does process hits regardless of their time, the digitisation step needs to be performed separately for each BC, $N-1$ and N, because the BC structure is applied during this step.

The basic idea how to account for this, is to duplicate the output of the simulation, which contains all hits from the on-time *and* out-of-time parts, and subsequently shift all hits in the duplicated version by one BC by subtracting 25 ns from each hit, as illustrated in Figure 4.3:

$$\underbrace{\{E_i, t_{\text{readout},i} = t_i - |\mathbf{r}_i|/c\}_{\text{hits},i}}_{\equiv \text{HITS1}} \rightarrow \underbrace{\{E_i, t_{\text{readout},i} - 25\,\text{ns}\}_{\text{hits},i}}_{\equiv \text{HITS2}} . \tag{4.8}$$

This can be implemented as a `HitTimeWrapperTool` running in Athena between the main simulation job and the digitisation job, in a so-called merging step [76].

[76]: Chapman (2025), *Private Communication*

The two hit files, are then separately put through digitisation, pile-up overlay, trigger simulation, and reconstruction, which allows to derive two separate physics samples representing the on-time ($N-1$) and out-of-time (N) parts of the process.

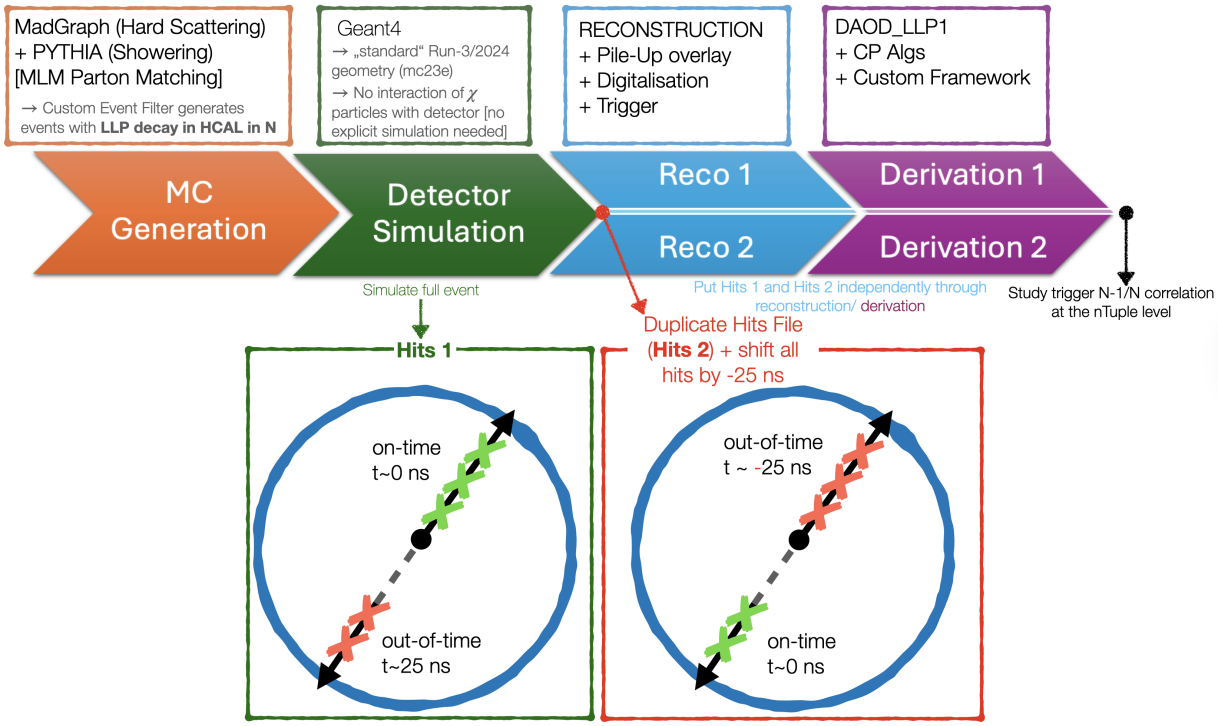

Figure 4.3.: Overview of the workflow of processing out-of-time decaying LLPs. The on-time and out-of-time parts are reconstructed individually. The on-time part is processed with the standard workflow inside Athena. For the out-of-time part, all Geant4 hits are shifted by one BC before processing them with the standard Athena reconstruction framework.

The trigger simulation is analysis-specific; details on the trigger strategy in the presented analysis are provided in Section 5.2. Here, the main idea of the trigger is to correlate a missing transverse momentum in $N-1$, $\text{MET}_{N-1} > 40\,\text{GeV}$, with a jet (also $> 40\,\text{GeV}$) in N with $\Delta\varphi(\text{MET}_{N-1}, \text{jet}_N) < 1.0$ and apply the CalRatio HLT chain [26] on top of the BC N.

The idea for simulating the L1 trigger correlation properly in Athena is to implement the trigger simulation of the $N-1$ part (L1 in $N-1$) and the trigger simulation of the N part (L1 in N and HLT in N) as two separate triggers:

$N-1$ trigger: `HLT_noalg_L1jXE40`,
N trigger: `HLT_j20_calratio_..._L1jJ40`,

both of which are individually tested in either parts. After deriving the physics samples, it is checked whether the $N-1$ trigger fires in the on-time part ($N-1$) and whether the N trigger fires in the out-of-time part (N). The L1 correlation between $N-1$ and N is then evaluated by comparing the φ of the L1 MET in $N-1$ with the φ of the triggering L1 jet(s) in N. In addition, it is required that no other triggers fire in the $N-1$ part, ensuring that N does *not* lie within the L1 deadtime.

[26]: ATLAS Collaboration (2022), *Search for neutral long-lived particles in pp collisions at $\sqrt{s}$ = 13 TeV that decay into displaced hadronic jets in the ATLAS calorimeter*

5. Dedicated Triggers for Delayed Long-Lived Particles

As the most energetic hadron collider to date, the LHC allows heavy BSM particles to be targetted, albeit with some limitations:

- The detector size sets an upper limit on the accessible decay length, $L = \beta c\gamma\tau_0 \lesssim \mathcal{O}(10\,\text{m})$; that is, for relativistic particles with $\beta\gamma = p/m > 1$, the proper lifetime $c\tau_0$ is usually also limited to $\mathcal{O}(10\,\text{m})$;
- Compressed hadronic decay signatures are limited by the QCD background dominating the low-energy signals, typically below $\mathcal{O}(100\,\text{GeV})$.

To address these challenges and to extend sensitivity to this difficult phase space, a dedicated LLP trigger — referred to in the following as *Back-to-the-Future Trigger (B2FT)* — is utilised. It targets electrically neutral, slow-moving LLPs that decay with a significant displacement in both space and time into kinematically compressed final states.

The trigger exploits correlations between two subsequent BCs: missing transverse energy in a BC $N-1$, where the LLP is produced but has not yet decayed, followed by a jet in the same azimuthal direction from its out-of-time decay in the subsequent BC N, as illustrated in Figure 5.1.

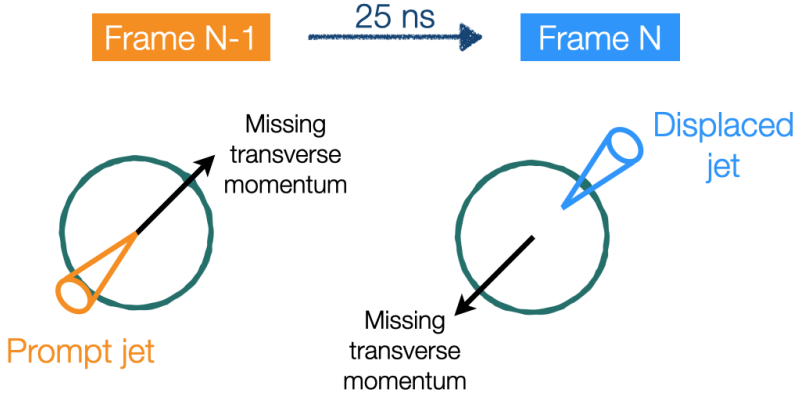

Figure 5.1.: Sketch of the LLP decay structure in the transverse detector plane. The trigger fires on a topological correlation of a MET in BC $N-1$ followed by a displaced and delayed jet in BC N.

By linking consecutive BCs, this strategy enhances sensitivity to long lifetimes while strongly suppressing QCD backgrounds, enabling the exploration of more compressed final states.

5.1. Target Phase-Space

This trigger strategy targets electrically neutral, slow-moving LLPs, whose production and decay might be split over several BCs. Based

on the the rough calorimeter-dimensions, $L \in (2, 5)$ m, and the BC readout window, $t_{\text{readout}} \in (-3, 15)$ ns and $t_{\text{readout}} \in (22, 40)$ ns for BCs $N-1$ and N, respectively[22], the required boost to arrive the HCal in the corresponding BC can be approximated:

$$t_{\text{readout}} = t_{\text{lab}} - \frac{L_{\text{lab}}}{c} = \frac{L_{\text{lab}}}{\beta c} - \frac{L_{\text{lab}}}{c},$$
$$\Rightarrow \beta = \frac{L_{\text{lab}}}{c t_{\text{readout}} + L_{\text{lab}}},$$

which can be evaluated at the maxima of the (rough) HCal dimensions and the boundaries of the time window:

$$\beta_{\text{min}} = \min\left\{\frac{L_{\text{min}}}{c t_{\text{max}} + L_{\text{min}}}, \frac{L_{\text{max}}}{c t_{\text{max}} + L_{\text{max}}}\right\},$$
$$\beta_{\text{max}} = \max\left\{\frac{L_{\text{min}}}{c t_{\text{min}} + L_{\text{min}}}, \frac{L_{\text{max}}}{c t_{\text{min}} + L_{\text{max}}}\right\}.$$

The LLP boost needs to be in $(0.1, 0.4)$ to arrive the HCal in BC N, while it should fall in $(0.3, 1.0)$ for BC $N-1$.

Here, the LLP boost is driven by the LLP production mechanism, in particular by the mass compression and the boost of the parent particle itself, which might have a longitudinal boost, β_z, and, when it recoils against ISR, also a boost in the transverse direction, β_{T}.

Additionally to arriving at the HCal one BC delayed, the LLP also needs to decay there; here, the probability for the LLP to decay at a time t_{lab} is described by an exponential decay probability with mean proper lifetime τ_0:

$$P(t) = \exp\left(-\frac{t_{\text{lab}}}{\gamma\tau_0}\right) = \exp\left(-\frac{t_{\text{readout}} + L_{\text{lab}}/c}{\gamma\tau_0}\right), \tag{5.1}$$
$$= \exp\left(-\frac{t_{\text{readout}}}{\gamma\tau_0}\right) \times \mathrm{e}^{-\beta}. \tag{5.2}$$

Integrated over the readout window, from t_{min} to t_{max}, and normalised, this can be evaluated as a two-dimensional function of boost (γ/β) and mean lifetime (τ_0):

$$\mathcal{P}(\gamma, \tau_0) := \frac{\int_{t_{\text{min}}}^{t_{\text{max}}} P(t)\,\mathrm{d}t}{\int_0^\infty P(t)\,\mathrm{d}t} = \exp\left(-\frac{t_{\text{min}}}{\gamma\tau_0}\right) - \exp\left(-\frac{t_{\text{max}}}{\gamma\tau_0}\right), \tag{5.3}$$

which is visualised for BCs $N-1$ and N in Figure 5.2.

The probability for an out-of-time decay in the HCal peaks at about 20% at a lifetime of $\tau_0 \sim 33$ ns. For the same lifetime, the probability for an on-time decay inside the HCal at about 40%. Ultimately, the boost distribution of the particular benchmark is the driving factor

22: Given the BC frequency of 25 ns, the time window of a BC N could be defined as $(N \times 25\,\text{ns}) \pm 12.5\,\text{ns}$. However, in analogy to the CalRatio search, this analysis defines the readout window slightly asymmetric: $(N \times 25\,\text{ns})^{+15}_{-3}$ ns. This choice of time window allows to reduce the effect of BIB events that might arrive early because BIB events typically do not travel through the interaction point; simultaneously, this choice of time window is still sensitive to delayed signals at higher times [26].

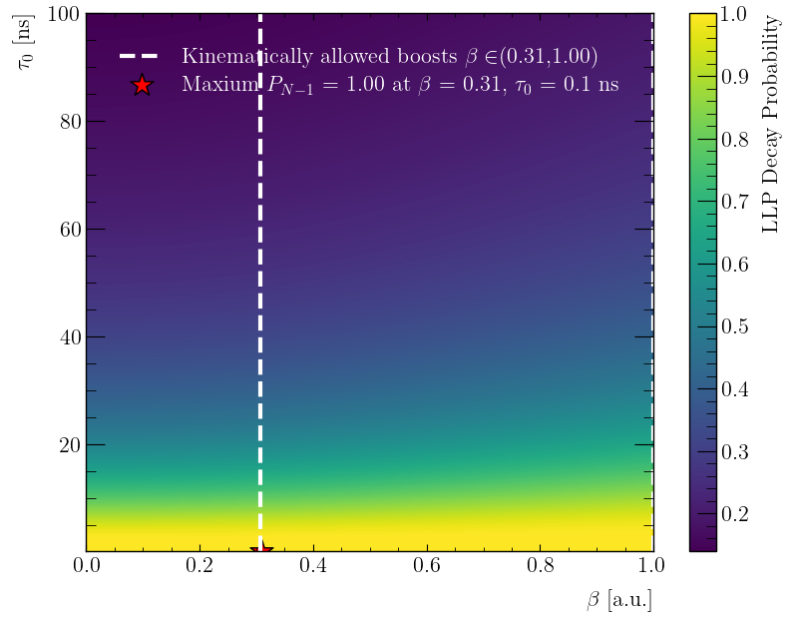

(a) On-time decay.

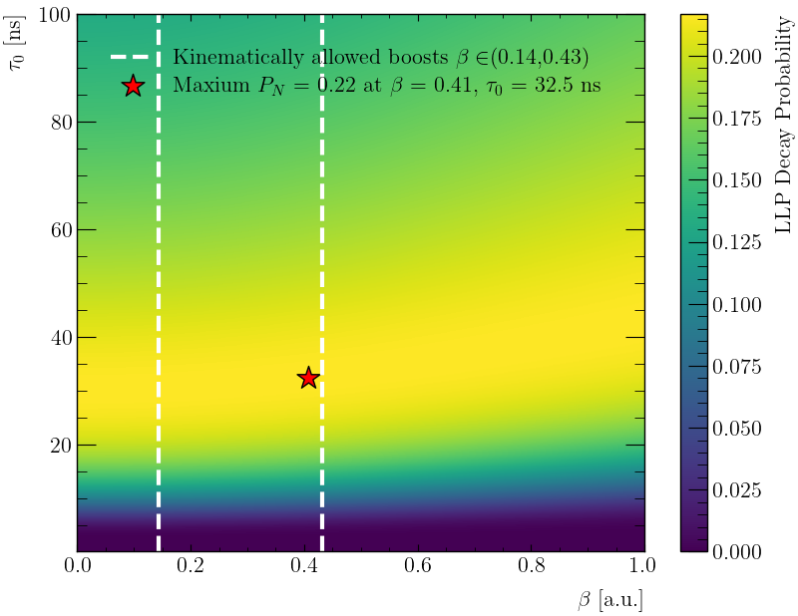

(b) Out-of-time decay.

Figure 5.2.: LLP decay probability in BCs $N-1$ and N. The decay probability of the LLP is presented for the two-dimensional $\tau \times \beta$ plane for BCs $N-1$ and N. The kinematic range for decays taking place inside the HCal is visualised with the white dashed lines.

which readout window is more sensitive. For example, assuming a boost distribution of a delta peak at 1.0 calls for using $N-1$, while a boost distribution of a delta peak at 0.1 for using N. In reality, the boost distribution is typically rather broad, and the $S/\sqrt{B}$ needs to be evaluated for the different readout windows to make a valid conclusion about which window is more sensitive.

5.2. Trigger Strategy

The B2FT facilitates the L1 trigger by studying the six most energetic small radius jets ($R = 0.4$) in a given BC N, using the jet feature extractor (jFex). Subsequently, the L1Topo checks for jets with $E_T > 40\,\text{GeV}$ if they are aligned with MET $> 40\,\text{GeV}$ in BC $N-1$, within $\Delta\varphi(\text{jet}_N, \text{MET}_{N-1}) < 1.0$. If such a jet/MET correlation is found at the L1, the CTP accepts the event N and sends it to the HLT.

[26]: ATLAS Collaboration (2022), *Search for neutral long-lived particles in pp collisions at $\sqrt{s}$ = 13 TeV that decay into displaced hadronic jets in the ATLAS calorimeter*

At the HLT, the CalRatio chain [26] is used, which targets hadronic jets with a low fraction of energy deposits in the ECal, caused by LLPs decaying at the outer edge of the ECal or somewhere in the HCal. It is worth mentioning that the HLT only accesses the displaced and delayed LLP decay in N, but not the $N-1$ event.

First, anti-k_t jets are reconstructed based on the topo-clusters as described in Section 3.4 using a jet radius of $R = 0.4$. Since the standard jet cleaning vetos jets with small energy fraction in the ECal[23], a dedicated algorithm is used instead — called CalRatio jet cleaning [26]. Here, the f_{EM} requirements of the usual jet cleaning are exchanged by a condition to reject jets with an energy fraction of more than 85% deposited in *one* calorimeter layer and with absolute value of negative energy higher than 4 GeV [26].

23: For example a jet with $f_{\text{EM}} < 0.05$ is rejected by the standard jet cleaning if either $f_{\text{ch}} < 0.05$ and $|\eta| < 2$ or if $|\eta| \geq 2$.

For firing the HLT, at least one jet passing the cleaning as described above is required to have

$$p_T > 20\,\text{GeV}, \quad |\eta| < 2.5, \quad \log_{10}(E_{\text{Had}}/E_{\text{EM}}) > 1.2 \tag{5.4}$$

where the latter cut ensures that the fraction of energy deposited in the ECal is smaller than $\sim 6\%$ of the energy deposited in the HCal.

Full scan tracking is performed in 0.8×0.8 in $\eta \times \varphi$ around these jets, to further veto jets having a track with $p_T > 2\,\text{GeV}$ within $\Delta R(\text{jet}, \text{track}) < 0.2$ to the jet axis. This cut is supposed to separate SM jets which statistically develop late (thus passing the EMF cut), but are not coming from neutral (LLP) decays in the calorimeter.

In the last step of the HLT hypothesis testing, beam induced backgrounds are removed as detailed in [26].

[26]: ATLAS Collaboration (2022), *Search for neutral long-lived particles in pp collisions at $\sqrt{s}$ = 13 TeV that decay into displaced hadronic jets in the ATLAS calorimeter*

In case of a successful HLT decision, the event N is readout and stored permanently, whereas the $N-1$ part can not be readout. That is, the correlation information between consecutive BCs is exclusively available at the L1 trigger and cannot be checked off-line. This not only means that the correlation information is exclusively available in the presented trigger strategy, but also that the cut values in $N-1$ (*i.e.* MET $> 40\,\text{GeV}$ and $\Delta\varphi < 1.0$) cannot be trivially

tightened at the off-line level to optimise the signal-to-background separation but needs to be fixed in the trigger firmware.

5.3. Background Suppression

Probing displaced and delayed jets in the calorimeters requires to tackle two sources of backgrounds, namely SM multi-jets and non-collision backgrounds.

- Although SM jets are *usually* neither delayed nor displaced, the amount of jets that statistically develop late in the calorimeter and therefore look signal-like, is still expected to be the dominating background source due to the overall high cross-section of SM multi-jets at the LHC;
- The non-collision backgrounds are not coming from the pp collisions but from other sources such as cosmic rays or BIB. Cosmic muons travelling through the upper hemisphere of the calorimeter and showering in the lower one, could potentially cause a displaced and delayed shower there. BIB includes collisions of protons with residual gas in the beam pipe or collisions of protons from the outer part of the beam with the focusing equipment, such as collimators [77]. Given the rather small rates of these processes and the additional MET correlation, the non-collision backgrounds are expected to be sub-dominant compared to the SM multi-jet background.

[77]: ATLAS Collaboration (2024), *Beam-induced backgrounds measured in the ATLAS detector during local gas injection into the LHC beam vacuum*

Since these backgrounds do in general not have a correlated MET in the previous BC, requiring the MET correlation at the L1 trigger helps to reduce the overall backgrounds. Nonetheless, the previous BC could have accidentally (*i.e.* by pure chance) an energetic MET, pointing in a correlated φ direction compared to the background event in N mimicking a displaced jet; the probability for this to happen is

$$P_{N-1 \wedge N} = P_{N-1} \times P_N \times \rho, \tag{5.5}$$

where P_N is the probability for a background event in BC N mimicking a displaced jet, P_{N-1} is the likelihood for a random event $N-1$ to have an energetic MET $> 40\,\text{GeV}$, and ρ is the correlation of the MET and the displaced jet in azimuthal direction, $\Delta\varphi(\text{jet}_N, \text{MET}_{N-1}) < 1.0$.

The requirement for a correlated MET in the previous BC reduces the background compared to not requiring this correlation by a

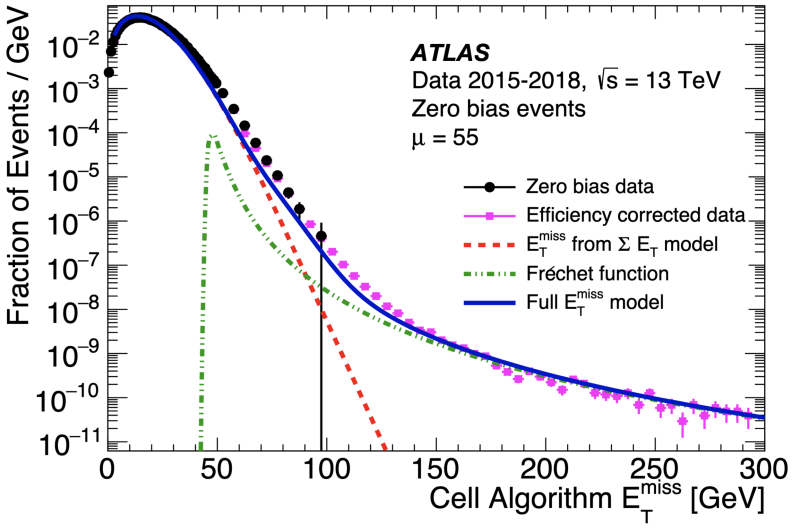

Figure 5.3.: Zero-Bias MET distribution in Run-2 pp data. The MET distribution (at the HLT) is presented for events fired with zero-bias triggers in Run-2 pp data. It shows that the MET distribution in randomly-triggered events is steeply falling, allowing to use the MET correlation as a strong signal-to-background discriminator. Taken from [78].

factor of

$$\frac{B_{\text{Corr}}}{B_{\text{noCorr}}} = \frac{P_{N-1} \times P_N \times \rho}{P_N} = P_{N-1} \times \rho, \tag{5.6}$$

which can be estimated using zero-bias data; as shown in Figure 5.3, the MET distribution in zero-bias data is steeply falling, and the probability for a random event to have an energetic MET above 40 GeV can be calculated as

$$P_{N-1} = \frac{\int_{40}^{\infty} f_{\text{ZB}}(\not{E}_{\text{T}})\, \mathrm{d}\not{E}_{\text{T}}}{\int_{0}^{\infty} f_{\text{ZB}}(\not{E}_{\text{T}})\, \mathrm{d}\not{E}_{\text{T}}} \approx 0.029. \tag{5.7}$$

Here, $f_{\text{ZB}}(\not{E}_{\text{T}})$ is the zero-bias MET distribution taken from the data points in Figure 5.3, which assumes Run-2 conditions, $\langle\mu\rangle = 55$, and facilitates the HLT cell algorithm MET distribution[24].

24: This particular HLT MET algorithm is actually quite close to the L1 MET algorithm, except for the finer granularity at the HLT level.

Similarly, the correlation factor, coming from the $\Delta\varphi < 1.0$ requirement, can be approximated by

$$\rho = \frac{\int_{0}^{1} f_{\text{ZB}}(\Delta\varphi)\, \mathrm{d}\Delta\varphi}{\int_{0}^{\pi} f_{\text{ZB}}(\Delta\varphi)\, \mathrm{d}\Delta\varphi} \approx \frac{1}{\pi}, \tag{5.8}$$

assuming $\Delta\varphi$ between the MET in $N-1$ and the jet in N to be uncorrelated for SM backgrounds, *i.e.* flat.

Combining both factors results in

$$P_{N-1} \times \rho = 0.029 \times 1/\pi \approx 0.0092, \tag{5.9}$$

which is a quite substantial background suppression by more than two orders of magnitude.

6. Phenomenology and Sensitivity Studies

To benchmark the experimental signature of MET in BC $N-1$ followed by a jet in the same azimuthal direction in BC N, a toy model is required to

- Contain an electrically neutral LLP with sufficiently high lifetime, covering a broad range of boosts, especially also small boosts;
- Provide a mechanism for the LLP to recoil against something visible (on-time), to observe the MET in $N-1$;
- Avoid the risk of conventional triggers being fired on the individual BCs, especially the decay in $N-1$ potentially causing N to lay within the L1 deadtime.

The simplest models combining these features produce two neutral particles pairwise in $N-1$ (back-to-back), one of which decays on-time in $N-1$ and the other one out-of-time in N, where the decay products are rather low-energy hadronic objects that are individually missed by conventional jet/MET triggers but which can be triggered on at a reasonable rate due to their $N-1$ and N correlation. The Inelastic Dipole Dark Matter (IDDM) is found to be a quite flexible model fulfilling the above-mentioned requirements.

This chapter presents a feasibility study of the B2FT from a phenomenological perspective exploring the IDDM as a benchmark model for this signature. First, the theoretical framework of IDDM is introduced in Section 6.1, followed by a list of the studied MC samples and an outline of the technical framework in Section 6.2. The event kinematics are studied in Section 6.3, and the expected signal and background yields are detailed in Sections 6.4 and 6.5, respectively. The resulting exclusion limits are compared to existing searches in Section 6.6. Section 6.7 concludes with an outlook on potential alternative models.

6.1. Inelastic Dipole Dark Matter

The Inelastic Dipole Dark Matter (IDDM) is a hidden-sector extension to the SM introducing *dark matter* Dirac fermion mass states $\chi_i, i \in \{0, 1, ..., n\}$ and the corresponding anti-particles $\bar{\chi}_i$. As illustrated in Figure 6.1, the χ_i states have a mass hierarchy $m_0 < m_1 < ... < m_n$, where the $i = 0$ state is the ground state (stable), and the $i = 1, ..., n$ states are excited states (instable).

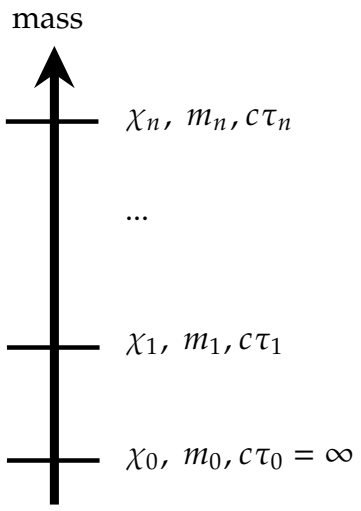

Figure 6.1.: The IDDM model introduces a stable ground state, χ_0, and several excited states χ_i with increasing mass m_i. The excited states may be long-lived or decay promptly.

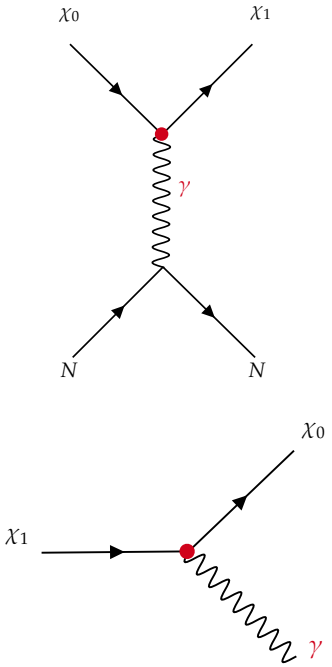

Figure 6.2.: Production and decay vertices in the original version of the IDDM via inelastic scattering and electromagnetic dipole operators.

As shown in the upper part of Figure 6.2, the ground state could scatter *inelastically* into an excited state, χ_i, with a small mass splitting between the ground state and the excited state, $\Delta_{i0} = m_i - m_0 \ll m_0$:

$$\chi_0 + N \rightarrow \chi_i + N, \tag{6.1}$$

where the coupling between the dark sector and SM particles is described by higher-order electromagnetic *dipole* operators. Subsequently, the excited state may decay via the same operators back into the ground state, $\chi_i \rightarrow \chi_0 + \gamma$, as illustrated in the lower part of Figure 6.2. Given the mass compression, $\Delta_{i0} \ll m_0$, and/or a small dipole operator coupling, the excited state may be be long-lived.

The IDDM as described above was originally developed for the DAMA/LIBRA experiment [79], and expanded for a FASER search [80, 81]; it is ultimately extended for the presented analysis [82], as detailed in the following sections, in particular with the following changes:

- The dark sector is linked to the SM via a so-called portal, which allows the IDDM fermions to be produced pairwise at the LHC, $pp \rightarrow \chi_i \chi_j$;
- The subsequent $\chi_{i,j} \rightarrow \chi_0 + X$ decay is modelled via effective higher-dimension operators coupling to the SM Higgs, while the electromagnetic dipole operator couplings are set to zero.

The relevant processes are visualised in Figure 6.3 and the full Feynman diagram is shown in Figure 6.4.

[79]: Smith et al. (2001), *Inelastic dark matter*

[80]: Jodłowski (2023), *Looking forward to inelastic DM with electromagnetic form factors at FASER and beam dump experiments*

[81]: Dienes et al. (2023), *Extending the discovery potential for inelastic-dipole dark matter with FASER*

[82]: Lessa et al. (2024), *Back to the Future Trigger*

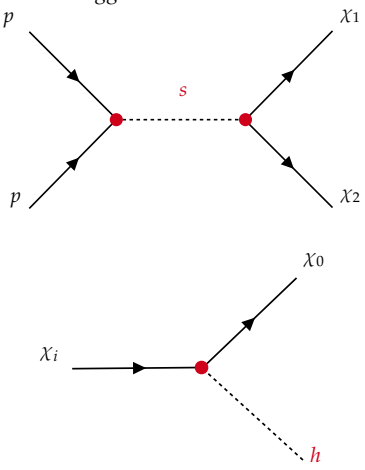

Figure 6.3.: Production and decay vertices in the modified version of the IDDM via scalar portal and Higgs operators.

Production Mechanism and Cross-Section

A scalar portal is used to produce the IDDM fermions pairwise via a BSM mediator by introducing a real scalar singlet φ in analogy to the Higgs mechanism; φ has a vacuum expectation value $v_D/\sqrt{2}$, and mixes with the Higgs doublet $\Phi = (\phi^+, \phi^0)$ into the SM Higgs h and the BSM scalar mediator s:

$$h = \left(\sqrt{2}\phi^0 - v\right)\cos\alpha - \left(\sqrt{2}\varphi - v_D\right)\sin\alpha, \tag{6.2}$$

$$s = \left(\sqrt{2}\phi^0 - v\right)\sin\alpha + \left(\sqrt{2}\varphi - v_D\right)\cos\alpha, \tag{6.3}$$

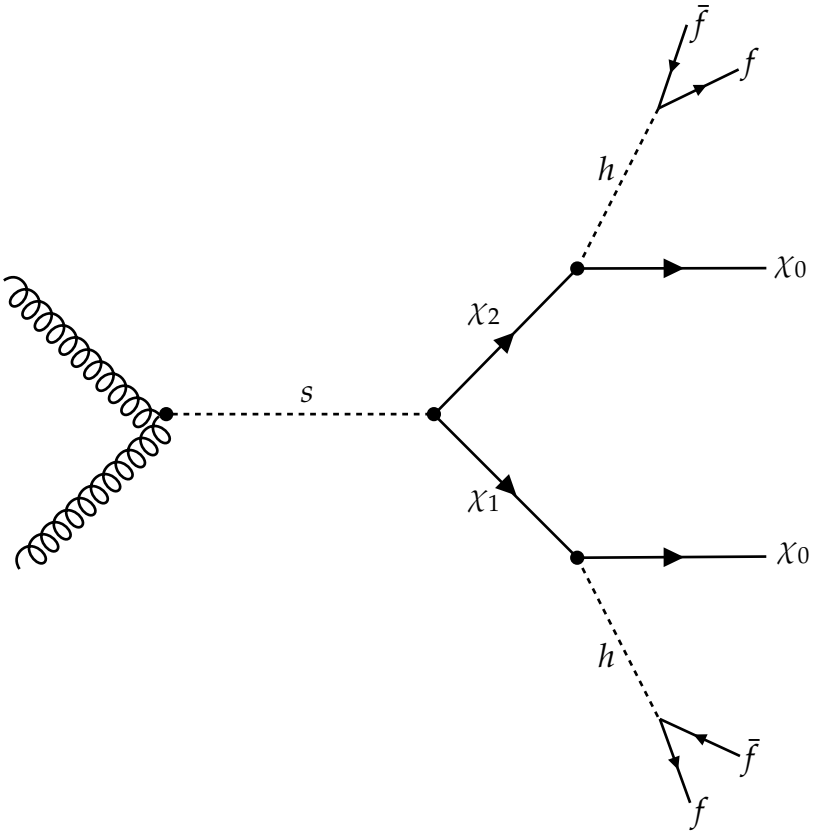

Figure 6.4.: Inelastic Dipole Dark Matter. The Feynman diagram of the studied IDDM toy model is shown. A pair of IDDM fermions is produced via a scalar portal (where the scalar BSM mediator is produced via a gluon-gluon induced top loop). Both IDDM fermions decay into the ground state IDDM fermion under the emission of a SM Higgs boson. The individual IDDM-Higgs couplings may be tuned, resulting in long-lived IDDM fermions. In particular, the model with $pp \to \chi_1(\chi_0 f\bar{f})\chi_2(\chi_0 f\bar{f})$ is studied in this analysis, where the χ_2 decays promptly, and the χ_1 is long-lived.

where $\sin\alpha \lesssim \mathcal{O}(0.2)$ is the mixing angle. Also, the Higgs potential is extended accordingly to

$$V(\Phi,\varphi) = \underbrace{-\mu_1|\Phi|^2 + \lambda_1|\Phi|^4}_{V_\Phi\ (c.f.\ \text{eq. 2.12})} \underbrace{-\mu_2|\varphi|^2 + \lambda_2|\varphi|^4}_{V_\varphi} + \underbrace{\lambda_3|\Phi|^2|\varphi|^2}_{V_{\Phi\varphi}}, \tag{6.4}$$

where the mass parameters, μ_i, and the quartic couplings, λ_i, can be re-parametrised by $m_h, m_s, v, v_\mathrm{D},$ and $\sin\alpha$.

The production mechanism for the BSM mediator is described effectively via a gluon-gluon induced top loop:

$$\mathscr{L}_{gg\varphi} = \frac{g_s^2}{48\pi^2 v}\left[\cos\alpha \times F\left(\frac{m_h^2}{m_t^2}\right) \times G^{\mu\nu}G_{\mu\nu}h + \ldots \right.$$
$$\left. \ldots + \sin\alpha \times F\left(\frac{m_s^2}{m_t^2}\right) \times G^{\mu\nu}G_{\mu\nu}s\right], \tag{6.5}$$

where $F(x)$ is the effective loop function [82]. The cross-section for $pp \to s + X$ is shown for $\sin\alpha = 0.2$ in Figure 6.5 for LO and higher-order ($\mathrm{N^3LO}$) calculations.

[82]: Lessa et al. (2024), *Back to the Future Trigger*

The scalar might subsequently decay into a pair of gluons, Higgs bosons, EW bosons (Z, W), SM fermions, or IDDM fermions. As seen in the Feynman rules in Table 6.1, the SM decay channels are

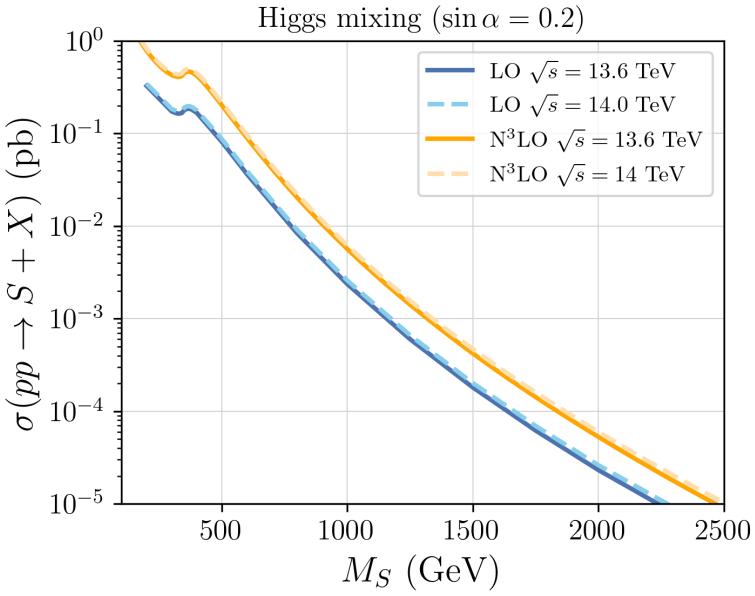

Figure 6.5.: Production cross-section of the scalar BSM mediator. The scalar production cross-section is presented for a mixing angle of $\sin\alpha = 0.2$. Taken from [82].

suppressed by $\sin\alpha$, whereas the IDDM decay channel $s \to \chi_i\chi_j$ is proportional to $\cos\alpha$.

The presented studies focus exclusively on the $s \to \chi_i\chi_j$ decay channel, with the aim of obtaining limits on the cross-section times the corresponding branching fraction.

IDDM Lagrangian and Feynman Rules

The Lagrangian of the modified IDDM model reads

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \mathcal{L}_{\varphi} + \mathcal{L}_{\chi} + \mathcal{L}_{\Phi\chi}, \tag{6.6}$$

were an additional $\mathcal{L}_{\text{dipole}}$ term could be added to allow for the original $\chi_i \to \chi_j + \gamma$ interaction[25]. More details on the individual components of the Lagrangian are provided in the following list and the resulting Feynman rules are summarised in Table 6.1.

- The SM Lagrangian, $\mathcal{L}_{\text{SM}}$, has components from the EW sector, eq. 2.2, QCD, eq. 2.15, the Higgs sector, eq. 2.11, and the Yukawa couplings, eq. 2.9:

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{EW}} + \mathcal{L}_{\text{QCD}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}}. \tag{6.9}$$

- The Lagrangian of the scalar, $\mathcal{L}_{\varphi}$, includes the kinematic term of φ and the extension of the $\Phi - \varphi$ potential, as defined in eq. 6.4:

$$\mathcal{L}_{\varphi} = \left(\partial^{\mu}\varphi\right)^2 - V_{\varphi} - V_{\Phi\varphi}. \tag{6.10}$$

25: The interaction between the IDDM fermions and photons would be described by an effective Lagrangian

$$\mathcal{L}_{\text{dipole}} = \frac{(C_{\gamma\chi\chi})_{ij}}{\Lambda}\chi_i\sigma^{\mu\nu}\chi_j F_{\mu\nu}, \tag{6.7}$$

with $\sigma^{\mu\nu} \sim [\gamma^{\mu}, \gamma^{\nu}]$, and $(C_{\gamma\chi\chi})_{ij}$ being the effective $\chi_i\chi_j\gamma$ couplings which are assumed to vanish in the presented analysis. The resulting vertex form for $\chi_i\chi_j A^{\mu}$ would read

$$-i\frac{1}{\Lambda}(C_{\gamma\chi\chi})_{ij}(\gamma^{\mu}\not{p} - \not{p}\gamma^{\mu}). \tag{6.8}$$

- The Lagrangian of the IDDM fermions, $\mathscr{L}_\chi$, includes their kinematic term, their mass terms, and their Yukawa-couplings to the scalar singlet:

$$\mathscr{L}_\chi = i\bar{\chi}_i \partial\!\!\!/ \chi_i - \tilde{M}_{ij}\chi_i\chi_j - \left(y_\chi\right)_{ij}\chi_i\chi_j\varphi. \tag{6.11}$$

  Here,

$$\tilde{M}_{ij} = m_i\delta_{ij} - \frac{v_D}{\sqrt{2}}(y_\chi)_{ij} - \frac{v^2}{2\Lambda}(C_{\Phi\chi\chi})_{ij} \tag{6.12}$$

  is the fermion mass matrix with m_i being the physical masses of the χ_i. Furthermore, $(y_\chi)_{ij}$ are Yukawa-like dimension-4 $\chi_i\chi_j\varphi$ couplings and $(C_{\Phi\chi\chi})_{ij}$ are effective dimension-6 $\chi_i\chi_j\Phi$ couplings, which are non-renormalizable and would come from some high scale (UV) physics.
- The interactions between the IDDM fermions and the SM Higgs are described by an effective Lagrangian

$$\mathscr{L}_{\Phi\chi} = \frac{\left(C_{\Phi\chi\chi}\right)_{ij}}{\Lambda}\chi_i\chi_j|\Phi|^2\,. \tag{6.13}$$

Interaction	**Vertex Term**
$\chi_i\chi_j h$	$\frac{i}{\sqrt{2}}(y_\chi)_{ij}\sin\alpha - i\frac{v}{\Lambda}(C_{\Phi\chi\chi})_{ij}\cos\alpha$
$s\chi_i\chi_j$	$-\frac{i}{\sqrt{2}}(y_\chi)_{ij}\cos\alpha - i\frac{v}{\Lambda}(C_{\Phi\chi\chi})_{ij}\sin\alpha$
shh	$-i\frac{m_s^2}{2v}\left(1+2\frac{m_h^2}{m_s^2}\right)\left(\cos\alpha + 2\frac{v}{v_D}\sin\alpha\right)\sin(2\alpha)$
$sf\bar{f}$	$-i\frac{m_f}{v}\sin\alpha$
$sW_\mu^- W_\nu^+$	$2ig^{\mu\nu}\frac{m_W^2}{v}\sin\alpha$
$sZ_\mu Z_\nu$	$2ig^{\mu\nu}\frac{m_Z^2}{v}\sin\alpha$
$sG_\mu G_\nu$	$i\frac{g_s^2}{12\pi^2 v}\sin\alpha F(m_s^2/m_t^2)(p_1^\mu p_2^\nu - g^{\mu\nu}p_1\cdot p_2)$

Table 6.1.: Feynman rules for the relevant interactions in the IDDM model.

Toy Models and Parametrisation

The IDDM model has the following free parameters:

- Mass of scalar particle, m_s;
- Masses of IDDM fermions, m_i for $i \in \{0, \ldots, n\}$;
- Yukawa-like dimension-4 $\chi_i\chi_j\varphi$ couplings, $(y_\chi)_{ij}$ for $0 \leq j \leq i \leq n$;
- Effective dimension-6 $\chi_i\chi_j\Phi$ couplings, $(C_{\Phi\chi\chi})_{ij}$;
- Mixing angle between the BSM scalar mediator s and the SM Higgs h, $\sin\alpha$;
- The ultra-violet cut-off, Λ_{UV};
- The vacuum expectation value of φ, v_{D}.

In total, this corresponds to $5 + n + 2 \times \frac{(n+1)(n+2)}{2}$ parameters, *i.e.*

12 (19) parameters for $n = 1\,(2)$, where n is the number of excited χ states. The amount of parameters can be reduced by studying toy models particularly interesting for the B2FT strategy, for which specific sets of couplings are chosen. Different toy models have been studied in this analysis, one of which is presented in this thesis, where the dark sector contains three IDDM fermions, $\chi_{0,1,2}$ with the following couplings:

$$C_{\Phi\chi\chi} = \begin{pmatrix} 0 & (C_{\Phi\chi\chi})_{10} & 0 \\ (C_{\Phi\chi\chi})_{10} & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \; y_\chi = \begin{pmatrix} 0 & 0 & (y_\chi)_{20} \\ 0 & 0 & (y_\chi)_{21} \\ (y_\chi)_{20} & (y_\chi)_{21} & 0 \end{pmatrix}. \tag{6.14}$$

This allows for a $\chi_1 + \chi_2$ production via $(y_\chi)_{21}$; the χ_1 solely decays via the effective $(C_{\Phi\chi\chi})_{10}$ operator, whereas the χ_2 decays promptly via the Yukawa-like $(y_\chi)_{20}$. That is, one prompt particle (χ_2) is produced alongside an LLP (χ_1) that may decay out-of-time; here the χ_1 lifetime could be tuned by the $(C_{\Phi\chi\chi})_{10}/\Lambda$ coupling and the mass splitting $\Delta_{10} = m_{\chi_1} - m_{\chi_0}$.

This toy model also allows for the $\chi_2\chi_0$ production via $(y_\chi)_{20}$, which would, however, result in a low energetic mono-jet plus MET signature with jet energies $\lesssim 100\,\text{GeV}$ in the studied phase-space.

It also allows for the $\chi_2 \to \chi_1 h$ decay via $(y_\chi)_{21}$, which would effectively result in a pair of LLPs and a prompt jet: $pp \to \chi_1 + \chi_2 \to \chi_1 + \chi_1 h$; this decay channel is kinematically suppressed if the masses are compressed, $m_2 \sim m_1 \gg m_0$.

Besides the Yukawa-couplings which are chosen to provide the desired signature, the remaining parameters can be re-parametrised to provide a set of physical parameters:

$$\{m_s, \sin\alpha\} \to \{m_s, \sigma_{\text{prod}}\}, \tag{6.15}$$

$$\{m_0, m_1, m_2\} \to \{\beta_1^*, \Delta_{21}, \Delta_{10}\}, \tag{6.16}$$

$$\{(C_{\Phi\chi\chi})_{10}, \Lambda_{\text{UV}}\} \to c\tau_{\chi_1}, \tag{6.17}$$

where

$$\beta_1^* = \frac{\sqrt{\left[m_s^2 - (m_1 + m_2)^2\right] \times \left[m_s^2 - (m_2 - m_1)^2\right]}}{m_s^2 - m_2^2 + m_1^2} \tag{6.18}$$

is the boost of the χ_1 in the rest frame of the scalar particle, and $\Delta_{ij} = m_i - m_j$ are the mass compressions; while Δ_{21} is required to be quite small, $\mathcal{O}(10\,\text{GeV})$, to avoid the $\chi_2 \to \chi_1 h$ decay, Δ_{10} is somewhat related to the shower energy of the LLP decay, and should take values around $\mathcal{O}(50 - 100\,\text{GeV})$.

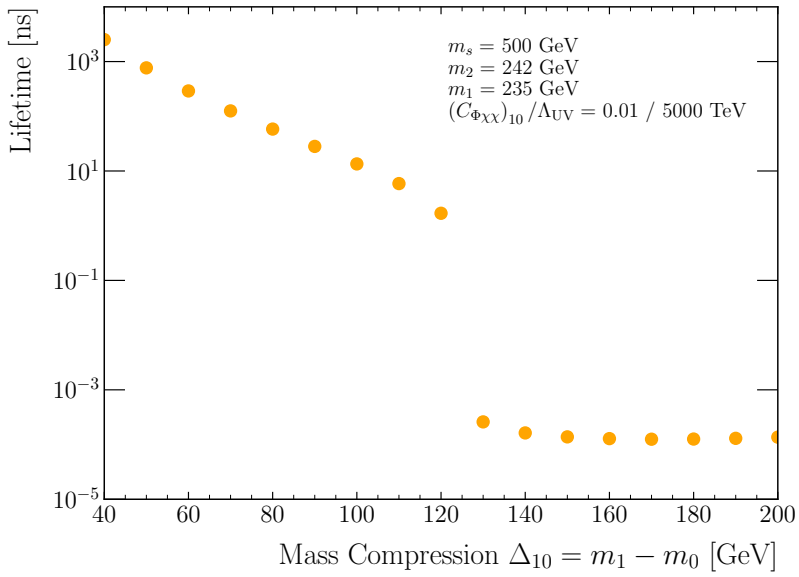

Figure 6.6.: Lifetime studies. Lifetime of the LLP for different mass compressions for a fixed (but arbitrary) coupling. A sharp cut is observed at the $m_h = 125\,\text{GeV}$ threshold, above which the phase-space is approximately constant, and below of which the phase-space is increasingly suppressed for decreasing mass compressions. The overall curve could be arbitrarily shifted to higher or lower lifetimes by choosing different $(C_{\Phi\chi\chi})_{10}/\Lambda_{UV}$ couplings.

Lifetime Studies

For a fixed coupling $(C_{\Phi\chi\chi})_{10}/\Lambda_{UV}$, the LLP lifetime mainly depends on the available phase-space for the $\chi_1 \to \chi_0 + h$ decay, which can be suppressed by a high mass compression, $\Delta_{10} = m_1 - m_0 \ll m_h$. The lifetime is studied as a function of Δ_{10} in Figure 6.6 for an arbitrary but fixed coupling. It exactly shows the behaviour of increasing lifetime with decreasing Δ_{10}, with a sharp cut at the m_h threshold, since $\Delta_{10} \gtrsim m_h$ allows for an on-shell production while $\Delta_{10} < m_h$ occurs suppressed via an off-shell Higgs. Therefore, the available phase-space is approximately constant for $\Delta_{10} > m_h$, and decreases significantly for $\Delta_{10} \ll m_h$. The overall curve may be arbitrary shifted to higher or lower lifetimes by choosing different $(C_{\Phi\chi\chi})/\Lambda_{UV}$ couplings, since $\tau \sim \left[\frac{(C_{\Phi\chi\chi})_{10}}{\Lambda_{UV}}\right]^{-2}$.

6.2. Monte Carlo Samples and Framework

A so-called unified Feynman output (UFO) card has been implemented for the model described in Section 6.1 [83], enabling the generation of $pp \to \chi_1\chi_2$ events with `MadGraph5_aMC@NLO v3.6.3`. Here, different jet multiplicities, $pp \to \chi_1\chi_2 + 0\,\text{jets}$ and $pp \to \chi_1\chi_2 + 1\,\text{jet}$, are generated, merged with the MLM scheme, and matched with the subsequent parton showering in `Pythia8.315`, as detailed in Chapter 4.

The background studies are mainly done data-driven based on an existing ATLAS Run-2 search [26], but di-jet simulations generated with Pythia and processed with Delphes are also used [83].

[83]: Lessa (2025), *back2futureTrigger*

[26]: ATLAS Collaboration (2022), *Search for neutral long-lived particles in pp collisions at $\sqrt{s}$ = 13 TeV that decay into displaced hadronic jets in the ATLAS calorimeter*

[83]: Lessa (2025), *back2futureTrigger*

Grid

Signal samples are generated for different points in the parameter-space grid, in particular for different scalar particle masses, χ_1 boosts, and mass splitting:

- m_s / GeV $\in \{300, 500, 1000, 2000\}$;
- $\beta_1^* \in \{0.1, 0.25, 0.5, 0.75\}$;
- Δ_{10}/ GeV $\in \{50, 100\}$;
- Δ_{21}/ GeV $\in \{7\}$[26].

26: No differences are observed for $\Delta_{21} \in (3, 50)$ GeV, and therefore only one Δ_{21} value is generated for the full grid.

The values of $\sin\alpha$ and v_D are fixed to 0.2 and 1 TeV, respectively; the Yukawa couplings are set to $(y_\chi)_{21} = 1.5$ and $(y_\chi)_{20} = 0.1$, while $(C_{\Phi\chi\chi})_{10}/\Lambda_{UV}$ is scaled to $c\tau_0 = 20$ m. Here, it is sufficient to generate one lifetime per grid point, which can be reweighted to different LLP lifetimes as detailed below.

Lifetime Reweighting

The signal yield can be obtained as function of the lifetime τ by generating only one reference sample with lifetime τ_{ref} and applying event-level weights depending on τ, τ_{ref}, and the decay time in the rest frame of the LLP(s), t_i, in this event [84]:

[84]: Burzynski (2025), *Lifetime reweighting*

$$w_{event} = \prod_{LLPs} w_i, \quad w_i = \frac{\tau_{ref}}{\tau} \times \exp\left[-\left(\frac{1}{\tau} - \frac{1}{\tau_{ref}}\right) \times t_i\right]. \quad (6.19)$$

Reweighting over a large τ range might be limited by the lack of statistics in the reference sample: in particular, reweighting from a smaller to a higher lifetime $\tau > \tau_{ref}$ might suffer by little statistics for high t_i in the reference sample, as sketched in Figure 6.7b; in general, reweighting from a higher to a smaller lifetime $\tau < \tau_{ref}$ is more stable due to the exponential distribution of t_i in the reference sample ensuring also sufficient population at low t_i, as drawn in Figure 6.7a.

The lifetime reweighting is validated, as presented exemplarily for one grid point in Figure 6.8, which shows the limitations of reweighting from lower to higher lifetimes; it also shows that reweighting from $c\tau_{ref} = \mathcal{O}(10\,\text{m})$ to lifetimes in the range $c\tau \in (1, 1000)$ m yields reasonable precision[27].

27: This procedure is also validated by comparing efficiencies and upper limits of generated and toy samples, observing closure.

MC Truth Framework

The generation chain explained above is done within the Athena framework, and the output is processed with a user-specific framework. Here, it is essential to perform the feasibility studies on the MC truth-level, since no existing framework at the detector level

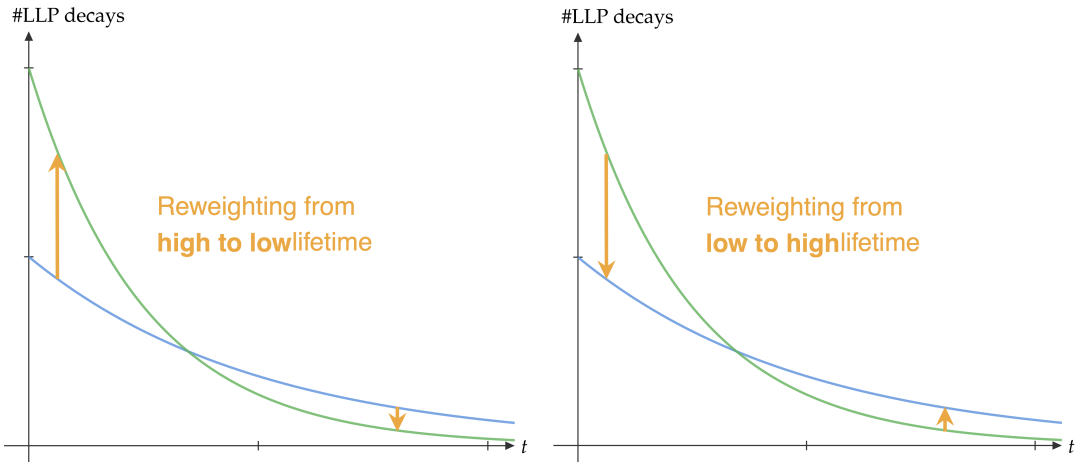

(a) Lifetime reweighting from high to lower lifetime. **(b)** Lifetime reweighting from low to higher lifetime.

Figure 6.7.: Schematic drawing of lifetime reweighting. The drawing presents the limitations of lacking statistics in the reference sample for two cases, namely reweighting from lower to higher lifetimes and vice versa. In particular the reweighting from a reference sample with low lifetime is limited by the lack of high decay times in the reference sample, biasing the reweighted sample.

can process multiple BCs, for which an extension to the Athena framework is needed; a fruitful sensitivity study at the MC truth-level not only serves as clear motivation for the development of such a framework, but also allows for its validation.

An alternative production chain with stand-alone MadGraph and Pythia, followed by a user-specific Delphes card [83] is used to validate the truth-level studies, showing closure between the two frameworks.

[83]: Lessa (2025), *back2futureTrigger*

The technicalities as well as the validation of the MC truth framework are detailed in Appendix A; here, the basic idea is to use all detector stable truth particles from the Pythia showering, propagate these particles to the calorimeter, and assign them to the corresponding BC based on their readout time in the calorimeter. From these particle collections, different objects are built, such as

- Trigger towers;
- L1 trigger small radius jets and MET;
- HLT/off-line level anti-k_t jets (with different jet radii);
- Jet level quantities, such as energy fractions in ECal and HCal;
- Tracks are reconstructed based on charged particles in the inner detector.

These objects are used to simulate the B2FT decision as realistically as possible on the MC truth-level, as described in the event selection in the following section.

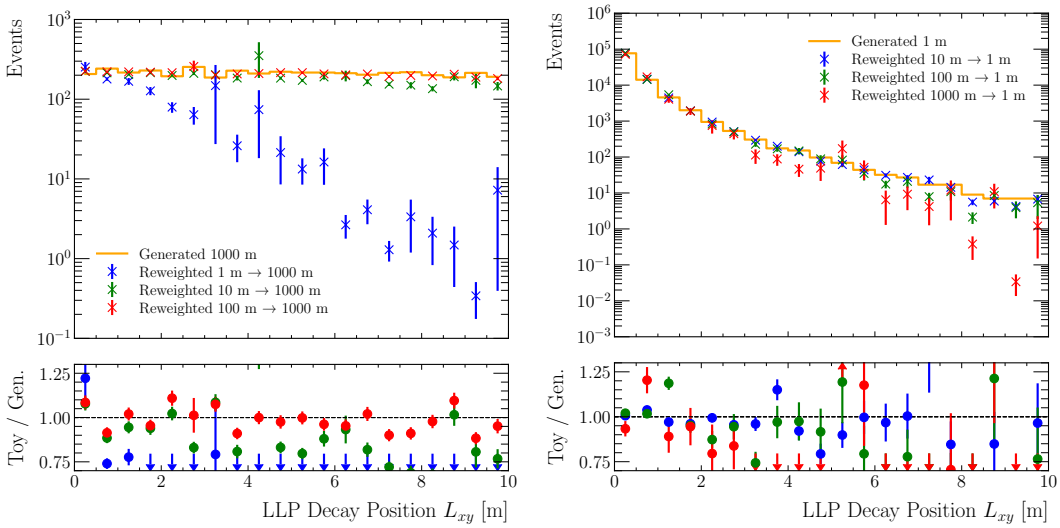

(a) Reweighting lower lifetimes to higher lifetimes, showing lack of statistics in low $c\tau$ samples.

(b) Reweighting higher lifetimes to lower lifetimes, showing more stable behaviour.

Figure 6.8.: MC lifetime reweighting. Generated and reweighted samples are compared, exemplarily for the LLP decay position in the transverse plane, L_{xy}. As expected, a limitation for the reweighting over large $c\tau$ ranges is observed, in particular for reweighting lower lifetimes to higher ones. It was decided to generate reference samples with $c\tau = 20\,\text{m}$ and reweight these samples to several lifetimes in the range $(1, 1000)\,\text{m}$. Systematical uncertainties are omitted at this stage of the analysis, since comparing the efficiencies and upper limits of generated and toy samples shows reasonable agreement.

6.3. Event Selection and Event Kinematics

The event selection used in the sensitivity studies includes the trigger decision, a pre-selection and a preliminary event selection. The overall event selection is optimised in signal-to-background discrimination after a successful feasibility study and development of a technical framework to process out of time decays in Athena.

The preliminary selection is greatly inspired by a previous CalRatio analysis [26]; in fact, it is crucial to follow their selection closely for the phenomenological feasibility studies to allow reuse of their background studies in the sensitivity studies. The event selection is summarised in Table 6.2, and details are provided below.

[26]: ATLAS Collaboration (2022), *Search for neutral long-lived particles in pp collisions at $\sqrt{s}$ = 13 TeV that decay into displaced hadronic jets in the ATLAS calorimeter*

Trigger Level

Simplified algorithms are used to reconstruct L1 objects and emulate the L1Topo correlation. Among the six leading jets in N, at least one is required to exceed $E_\text{T} > 40\,\text{GeV}$ and be within $\Delta\varphi(\text{MET}_{N-1}, \text{jet}_N) < 1.0$. The MET in $N-1$ is required to be above 40 GeV but simultaneously remain below 100 GeV; here, the upper cut on MET ensures to avoid conventional MET triggers to fire on $N-1$ and to shadow N.

The jet multiplicity and jet E_T distributions in BC N are shown in Figure 6.9 before applying the cuts.

Table 6.2.: Simplified event selection for the sensitivity studies. Here, ΔR_{min} is a measure how trackless a jet is, while λ_{ECal} and λ_{HCal} are the remaining hadronic interaction length from the jet origin to the calorimeter exit, which are utilised to approximate the energy fractions in the different calorimeters.

Stage	Selection
L1 Trigger ($N-1$, N)	L1 MET: $40 < \text{MET}_{N-1}/\,\text{GeV} < 100$ L1 small radius jet (N): $E_{\text{T}} > 40\,\text{GeV}$ and $\lvert\eta\rvert < 3.2$ $\Delta\varphi(\text{MET}_{N-1}, \text{jet}_N) < 1.0$ for one of the six leading jets
HLT Trigger (N)	Anti-k_t jet with $p_{\text{T}} > 20\,\text{GeV}$ and $\lvert\eta\rvert < 2.5$ $\Delta R_{\text{min}}(\text{jet, closest track} > 2\,\text{GeV}) > 0.2$ $\lambda_{\text{ECal}} < 1, \lambda_{\text{HCal}} > 1$
Pre-Selection (N)	Anti-k_t jet with $p_{\text{T}} > 40\,\text{GeV}$ and $\lvert\eta\rvert < 2.5$ $\Delta R_{\text{min}}(\text{jet, closest Track} > 2\,\text{GeV}) > 0.5$
Event Selection (N)	$-3 < t_{\text{jet}}/\,\text{ns} < 15$ (with $t = 0$ at BC N) $\lvert\eta_{\text{jet}}\rvert \notin [1.45, 1.55]$

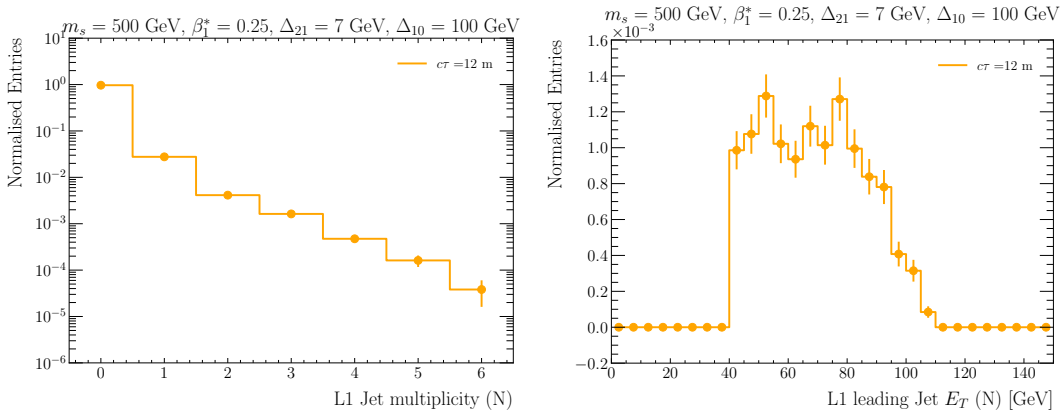

(a) L1 small radius jet multiplicity, where the last bin is an overflow bin, which contains all events with 6 or more jets.

(b) Leading L1 small radius jet E_{T} distribution for those events that contain at least one jet with transverse energy above 40 GeV.

Figure 6.9.: Relevant L1 jet distributions for one exemplary grid point. The relevant L1 small radius jet distributions are presented for one exemplary grid point. A huge fraction of the events does *not* contain a single jet in BC N due to the LLP decay time and boost making them decay outside the second BC and/or beyond the HCal. The distributions are normalised to the sum of weights from the lifetime reweighting.

Here, it is notable that a large fraction of events contains no jets in BC N, which can be attributed to:

- The decay probability of the LLP follows an exponential curve, causing a huge fraction of the LLPs decay outside the readout window of BC N, as shown in Figure 6.10a;
- The LLP needs sufficiently small boost of order $\mathcal{O}(0.1 - 0.4)$ such that the out-of-time decay may take place inside the HCal. However, as shown in Figure 6.10b, the boost is quite large, despite the low β_1^* value, which is the LLP boost in the rest frame of the scalar particle.

The combination of these features result in a quite low fraction of LLP decays in N inside the HCal, as shown in Figure 6.11.

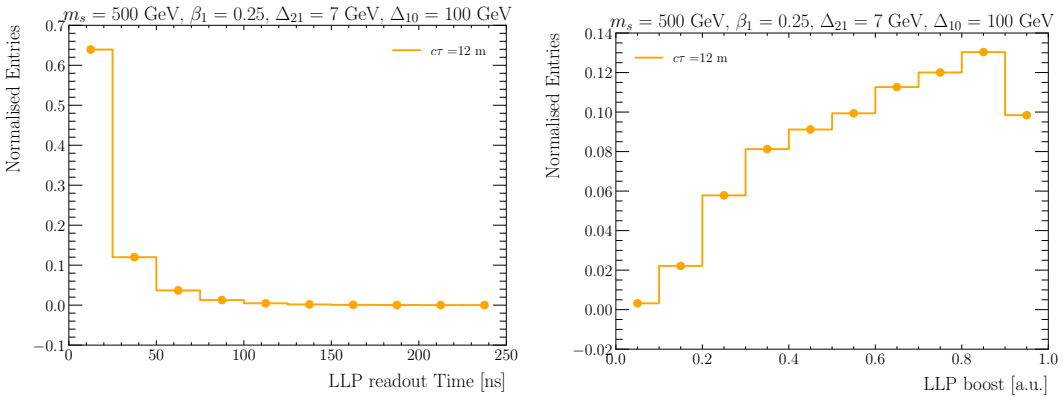

(a) LLP readout time, $t_{\text{readout}} = t_{\text{decay}} - L_{\text{decay}}/c$.

(b) LLP boost in the lab frame, $\beta = p/E$.

Figure 6.10.: Parton level distributions of BSM particles. The LLP readout time in the calorimeter as well as the LLP boost are presented for one exemplary grid point. A huge fraction of the LLPs decays outside the readout window of the second BC, and only about 15% fall within the N readout window. Moreover, a huge fraction of the LLPs has rather high lab-frame boost that are not in the required $\mathcal{O}(0.1 - 0.4)$ range, whereas only about 15% of the events contain an LLP with boost in this interval. Since boost and lifetime are independent, the combination of these two effects results in $\mathcal{O}(2\%)$ of LLP decays in the HCal in N, what agrees with the L1 jet multiplicity presented for the same grid point in Figure 6.9a, and also agrees with the 2.4% in the cutflow as shown in Table 6.3.

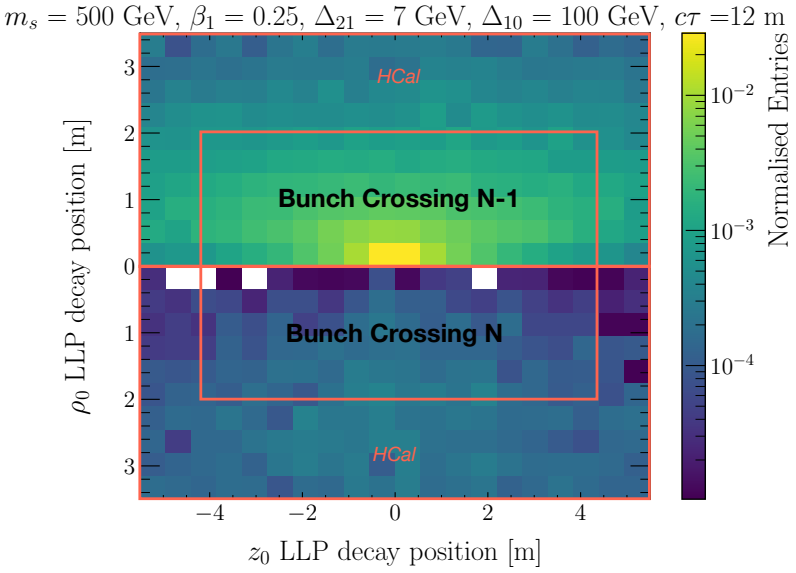

Figure 6.11.: LLP decay position. The decay position in the $L_{xy} \times L_z$ plane are presented for LLP decays in BC $N-1$ (upper part) and for BC N (lower part), showing that only a small fraction of LLPs decays in N inside the HCal.

Moreover, the MET distribution in $N-1$, as well as the $\Delta\varphi$ distance between the MET and the closest jet above 40 GeV are presented in Figure 6.12; the MET and its φ direction are also affected by additional ISR, causing a tail towards higher MET values and also broadening of the $\Delta\varphi$ distribution.

At the HLT stage, a displaced and trackless jet is required with $p_T > 20\,\text{GeV}$ and $|\eta| < 2.5$. Here, the degree of displacement is measured by the fraction of energy deposited in the ECal with respect to the total deposited energy,

$$\text{EMF} := \frac{E_{\text{ECal}}}{E_{\text{ECal}} + E_{\text{HCal}}}, \tag{6.20}$$

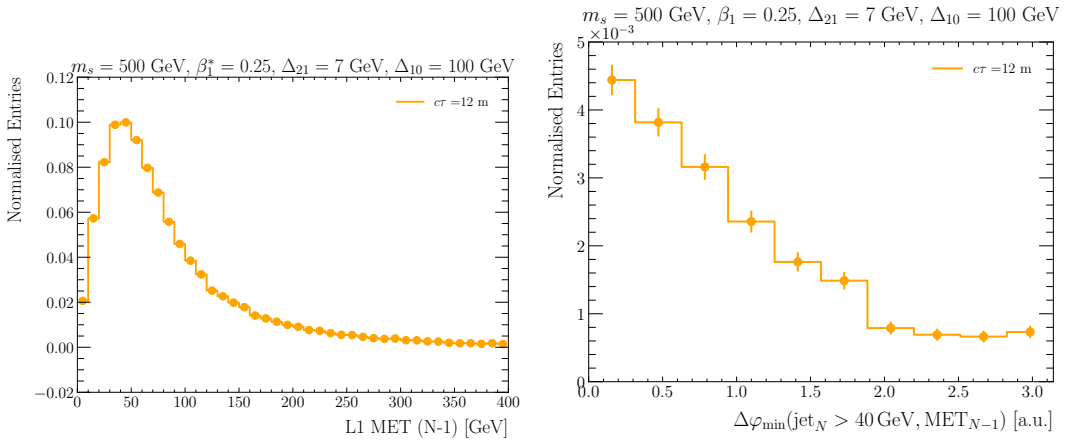

(a) L1 MET distribution in BC $N-1$.

(b) Azimuthal distance $\Delta\varphi(\text{MET}_{N-1}, \text{closest jet}_N)$.

Figure 6.12.: L1 MET and correlation distribution for one exemplary grid point. The MET in $N-1$ is required to be between 40 GeV and 100 GeV, and the $\Delta\varphi_{\text{min}}$ correlation between the MET in $N-1$ and the closest jet in N with $E_T > 40$ GeV must be smaller than 1.0.

which is required to be below 6.4% for displaced jets. Furthermore, the distance in ΔR between the jet and the closest track above 2 GeV,

$$\Delta R_{\text{min}} := \Delta R(\text{jet, closest track} > 2\text{ GeV}) \tag{6.21}$$

is a measure for how trackless a jet is, which is required to be larger than 0.2.

Anti-k_t truth jets are clustered with the FastJet framework using the detector stable truth particles extrapolated to their calorimeter entry. The resulting jet multiplicity and energy distributions are shown in Figure 6.13 after applying the p_{T} and η cuts.

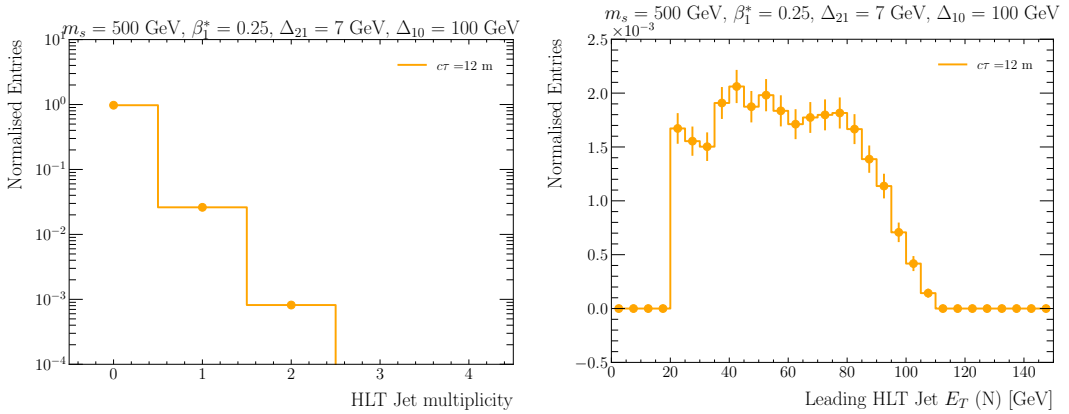

(a) Anti-k_t jet multiplicity.

(b) Leading anti-k_t jet E_{T} distribution.

Figure 6.13.: Relevant anti-k_t jet distributions for one exemplary grid point. The jet multiplicity and leading jet E_{T} distributions are presented for the HLT/off-line level for one exemplary grid point. The jet multiplicities are smaller compared to the L1, since p_{T} and $|\eta|$ cuts of 20 GeV and 2.5, respectively, are applied in the anti-k_t clustering, whereas the L1 jets span up to $|\eta| \leq 3.2$ and "only" noise cuts are applied which are of order 2 GeV per trigger tower.

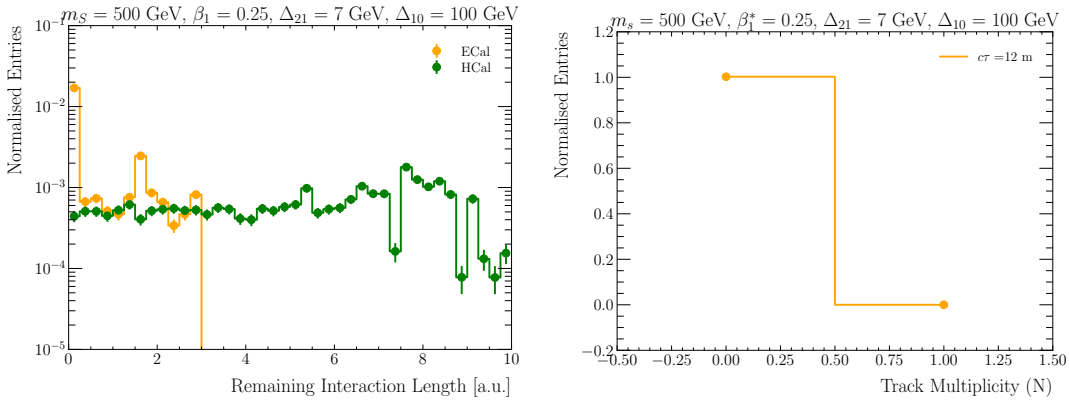

(a) Energy fraction deposited in the ECal and HCal. **(b)** HLT track multiplicity in BC N.

Figure 6.14.: Relevant HLT distributions. The energy fraction in the ECal compared to the total energy deposit, as well as the distance in ΔR between the jet and its closest track with $p_T > 2$ GeV are used to discriminate displaced jets from SM jets. **(a)** The remaining interaction length is required to be below 1 in the ECal and above 1 in the HCal, such that a sufficiently small amount of energy is deposited in the ECal and a sufficiently large amount of energy is deposited in the HCal. The peak at 0 in λ_{ECal} corresponds to those jets originating in the HCal that do not pass material in the ECal; these jets also only pass a part of the material budget in the HCal, and correspond therefore to the smaller $\lambda_{\text{HCal}} \lesssim 7.5$. On the other hand, those jets with $\lambda_{\text{ECal}} > 0$ correspond to jets that originate before the HCal; therefore, they need to pass the full HCal material, and hence correspond to the larger $\lambda_{\text{HCal}} \gtrsim 7.5$ values. **(b)** Since there are *exactly* zero tracks in N (at the MC truth-level, similar at the MC reco level), the $\Delta R_{\text{min}} > 0.2$ cut is trivially passed. This is observed for *all* grid points.

The EMF and ΔR_{min} cuts are approximated as follows:

- The energy fraction in the ECal, EMF, is approximated as a fiducial cut based on the jet origin[28] and the material budget in the ATLAS detector as detailed in Appendix A. This allows to estimate the "remaining hadronic interaction length", λ_{ECal}, from the jet origin to the ECal exit, which is shown for an exemplary grid point in Figure 6.14a; this distance is required to be below one, $\lambda_{\text{ECal}} < 1$, such that the shower develops only minimally inside the ECal and the energy deposit in the ECal is therefore reasonably low compared to the HCal. Simultaneously, the remaining interaction length in the HCal is required to be sufficiently large, $\lambda_{\text{HCal}} > 1$, such that sufficient energy is deposited in the HCal. Effectively, these two cuts ensure that the jet originates either at the outer edge of the ECal or inside the HCal (excluding the outer edge of the HCal).
- Tracks are reconstructed based on charged particles in the inner detector. Here, is notable that the number of tracks in BC N is exactly zero, as shown for an exemplary grid point in Figure 6.14b; this reflects the fact that the signal MC only contains prompt tracks (*e.g.* from ISR), but no delayed tracks. Overlaying pile-up might add some tracks to N, but these tracks are likely not associated to the primary vertex. In the absence of tracks the ΔR_{min} value is undefined (∞) and the ΔR_{min} cut is trivially passed.

28: Here, the jet origin is calculated as the energy weighted average of the jet constituents' production vertices.

Higher Level Selection

The pre-selection as well as the event selection in BC N include slightly tighter cuts on the jet p_{T}, η, and timing, as well as on ΔR_{min}. The overall selection is inspired by a previous CalRatio analysis [26], since there are significant parallels between the two analyses, especially sharing a similar HLT trigger. That being said, the two analyses are also distinct in some points, in particular, the presented analysis targets

[26]: ATLAS Collaboration (2022), *Search for neutral long-lived particles in pp collisions at $\sqrt{s}$ = 13 TeV that decay into displaced hadronic jets in the ATLAS calorimeter*

- Smaller LLP boosts, potentially causing more spherical LLP decays requiring larger jet radii;
- More compressed final states, resulting in lower jet energies;
- One displaced jet, while CalRatio targets two of them.

To address these differences, softer p_{T} cuts are used (40 GeV instead of 80 GeV), and the remaining cuts are adapted to one displaced jet instead of two, *e.g.* by requiring a lower summed EMF value. Moreover, different jet radii are explored to check for backwards travelling particles in the shower; however, since no differences are observed, usual $R = 0.4$ jets are used.

6.4. Trigger Efficiencies and Expected Signal Yields

The cutflow as described in Section 6.3 is applied to the signal samples listed in Section 6.2. The fraction of events passing the cuts are the selection efficiencies, ε, which are presented two-dimensionally in Figure 6.15.

Overall, the signal efficiencies are *below the percent level*, and they depend on the chosen mass of the scalar particle, the LLP lifetime and boost as well as the mass compression in the final state. In particular, the following dependencies are observed:

- **LLP lifetime:** The dependence of the efficiency on the LLP lifetime is mainly driven by the fiducial constraints of the HCal on the proper decay length $L = c\beta\gamma\tau$; that is, $\varepsilon(c\tau)$ peaks at $c\tau \sim \mathcal{O}(10\,\mathrm{m})$[29]. In addition, the delayed decay in N makes the lifetime dependence highly asymmetric towards longer lifetimes with reduced sensitivity at short lifetimes.
- **LLP boost:** The LLP lab frame boost, $\beta = p/E$, heavily affects the efficiency. However, β itself cannot be directly controlled by tuning the masses and couplings, but only the boost in the rest frame of the scalar particle, β^*, which is mainly driven by the mass compression of the scalar particle and the IDDM particles, as defined in eq. 6.18. The LLPs might get additional

29: Which aligns with the calculated 32.5 ns (*i.e.*, $c\tau \sim 9.75\,\mathrm{m}$) in Figure 5.2.

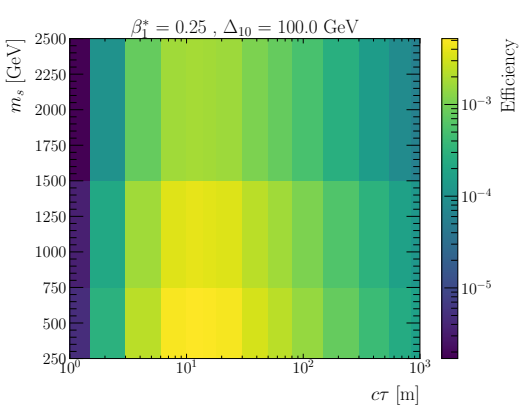

(a) Efficiencies in $m_s \times c\tau$ plane.

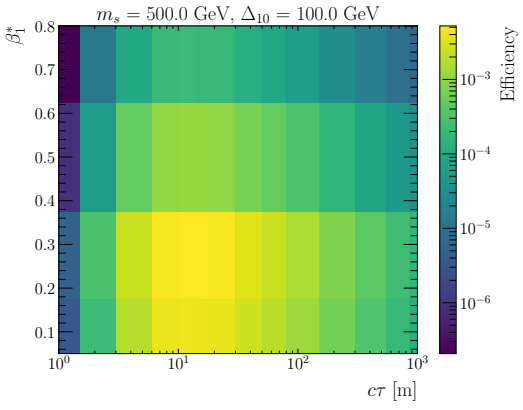

(b) Efficiencies in $\beta_1^* \times c\tau$ plane.

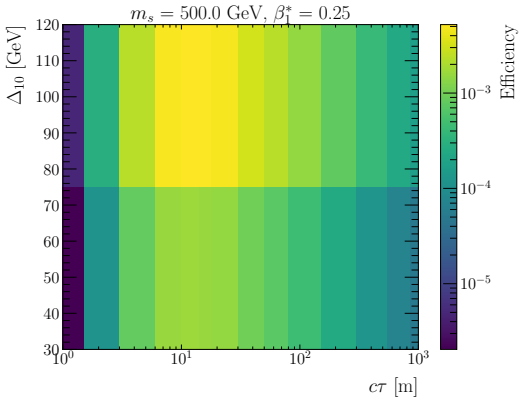

(c) Efficiencies in $\Delta_{10} \times c\tau$ plane.

Figure 6.15.: Two-dimensional selection efficiencies. The efficiencies are presented as function of lifetime and the decay kinematics.

boost from the scalar particle itself, affected by ISR and by the intrinsic boost of the scalar in longitudinal direction.

- **Scalar particle mass:** The mass of the scalar particle roughly sets the energy scale of the collision, $\sqrt{\hat{s}}$, and therefore is the dominant driver of ISR: the heavier the scalar particle, the harder potential ISR, increasing the LLP boost and causing the event to be shadowed by a conventional MET trigger firing in $N-1$. Both of these effects reduce the efficiencies with increasing m_s.
- **Mass compression:** The mass splitting between the final state particles, in particular $\Delta_{10} = m_1 - m_0$, is related to the energy of the jet formed by the $f\bar{f}$ pair, and therefore affects the efficiencies according to the jet E_T and MET cuts.

The cutflow for a sample where the efficiency peaks in the generated grid is presented in Table 6.3, showing the two main bottlenecks:

- Displaced jet inside the HCal in BC N: this requires a LLP with sufficiently high decay time *and* sufficiently small boost, which is strongly limited by ISR and the intrinsic boost of the scalar particle in longitudinal direction, both of which

Table 6.3.: The cutflow is presented for one exemplary grid point. Here, the sum of weights after each cut, the per-cut efficiencies, the cumulative efficiencies, and the expected signal yields are presented. For the latter, the Run-2 luminosity is assumed.

$m_s = 500\,\text{GeV}$, $\beta_1^* = 0.25$, $\Delta_{21} = 7\,\text{GeV}$, $\Delta_{10} = 100\,\text{GeV}$, $c\tau = 12\,\text{m}$

Cut	Events	ε_{cut}	$\varepsilon_{\text{cumulative}}$	# Signal
Sum of weights	1.00×10^5	—	—	—
$40\,\text{GeV} < \text{L1 MET}\,(N-1) < 100\,\text{GeV}$	4.42×10^4	0.442	0.442	1055
L1 Jet $(N) > 40$ GeV	1.06×10^3	2.40×10^{-2}	1.06×10^{-2}	25.3
$\Delta\varphi(\text{MET}_{N-1}, \text{jet}_N) < 1.0$	6.79×10^2	0.641	6.79×10^{-3}	16.2
Anti-k_t jet, $p_\text{T} > 20\,\text{GeV}$, $\lvert\eta\rvert < 2.5$	6.79×10^2	1.00	6.79×10^{-3}	16.2
$\lambda_{\text{ECal}} > 1.0$ and $\lambda_{\text{HCal}} > 1.0$	5.33×10^2	0.785	5.33×10^{-3}	12.7
$\Delta R_{\text{min}} > 0.2$	5.33×10^2	0.785	5.33×10^{-3}	12.7

can only be hardly avoided by tuning of the masses and couplings.

- Shadowing effect: in addition to the rise in boost, hard ISR also causes conventional MET triggers being fired in $N-1$, which is the second largest bottleneck.

Despite the relatively low signal efficiencies, the proposed strategy remains motivated, as the MET correlation also leads to a significant background reduction, as discussed in Section 5.3. However, the rather high boosts are quite concerning, because this is the feature driving the motivation to correlate two subsequent BCs. In the absence of low boosts, other on-time analyses only accessing one BCs are likely also sensitive to the signal.

The signal yield, S, is obtained by multiplying the efficiencies with the production cross-section, σ, and the Luminosity, $\mathscr{L}$:

$$S = \varepsilon \times \sigma \times \mathscr{L}. \tag{6.22}$$

Assuming the Run-2 luminosity, $\mathscr{L} = 140\,\text{fb}^{-1}$ [85], the signal yields peak in the $\mathcal{O}(10-100)$ range, as exemplarily shown for one grid point in Figure 6.16.

[85]: ATLAS Collaboration (2023), *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*

Here, the numerical values of the parameters such as the Yukawa-like couplings, $(y_\chi)_{ij}$, the mixing angle, $\sin\alpha$, as well as the vacuum expectation value of the dark sector, v_D, enter the cross-section, and make the signal yields to some extend arbitrary. To avoid this dependence, it is preferred to derive exclusion limits on the production cross-section times the branching fraction.

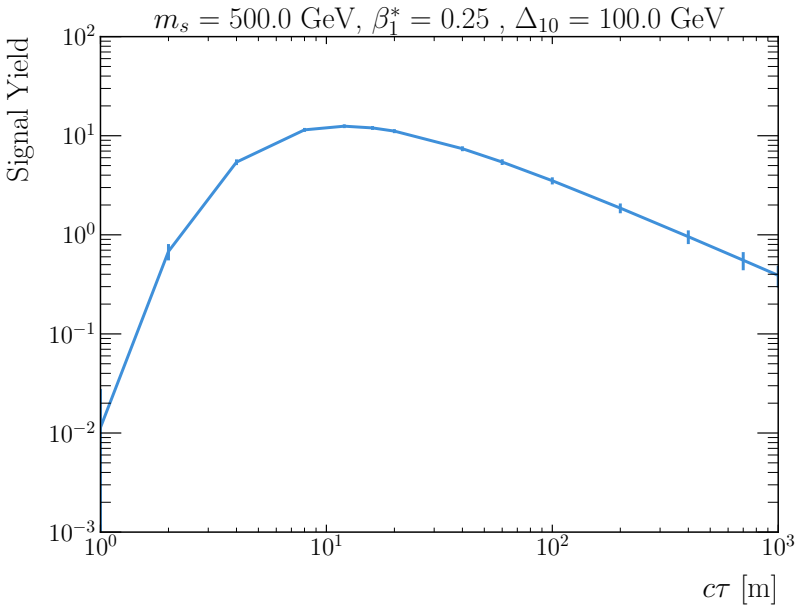

Figure 6.16.: Signal yields as function of lifetime. The expected signal yields are presented as a function of the proper lifetime assuming the Run-2 luminosity. The cross-section includes the kinematics (mass of scalar) and the phase-space suppression by the final state compression.

6.5. Expected Backgrounds and Exclusion Limits

The dominating background comes from SM multi-jets that develop late by statistics. This would require at least a full simulation using Geant4 if possible in MC at all. A data-driven alternative based on the studies of a previous CalRatio search turns out to be more efficient as a first estimate; since the B2F and CalRatio analyses have many parallels, especially sharing a similar HLT trigger, a back-to-the-envelope scaling of the CalRatio data rates allows to obtain the expected B2F data rates, which — assuming that the CalRatio data is background-only — can be used as a proxy for the background in the B2F analysis.

Here, it is important to understand that the aim is not to be super precise, but to provide an order-of-magnitude estimate to motivate the overall analysis strategy.

The mentioned data rates for the CalRatio cutflow are presented in Table 6.4, and include a pre-selection, event cleaning, low-E_T selection, all of which are similar to the preliminary B2F event selection introduced in Section 6.3, as well as an additional ABCD plane [86]; a couple of scaling factors are required to account for the differences in the two analyses:

- **Preselection:** The two analyses use different L1 triggers, similar[30] HLT chains and slightly different pre-selections. As summarised in Table 6.5, both analyses require a displaced jet above 40 GeV, with the difference that the B2F analysis additionally requires a correlated MET in the *previous* BC, whereas the CalRatio analysis requires an additional jet above 40 GeV in the *same* BC.

[86]: Buttinger (2018), *Background estimation with the ABCD method featuring the TROOFIT toolkit*

30: In principle, the HLT is identical, but in practice a slight difference is found due to the pile-up removal algorithm applied in the B2F HLT chain, as detailed in Section 7.1. Moreover, slightly different p_T thresholds are used: 20 GeV in the B2FT, and 30 GeV in the CalRatio trigger.

As detailed in Section 5.3, the MET correlation reduces the background by a factor of $1/0.0092 \sim 108$. In contrast, the additional jet in the CalRatio selection reduces the background by a factor of 5, where the numerical factor is derived using di-jet MC simulations [82]. That is, the background of the B2F analysis at the pre-selection stage corresponds to the background of the CalRatio analysis reduced by a factor of $108/5 \sim 22$.

[82]: Lessa et al. (2024), *Back to the Future Trigger*

- **Event cleaning:** The cut on the low-E_{T} boosted decision tree (BDT) score applied in the CalRatio selection cannot be used in B2F selection, since it uses two displaced jets as input. Also other cuts need to be scaled to represent the background rejection for one jet only. Here, the background rejection efficiencies of CalRatio are scaled by their square root, to obtain the according background rejection efficiencies for one jet only, where the square root relation comes from

$$\mathcal{P}(j_1 \lesseqgtr \text{cut} \wedge j_2 \lesseqgtr \text{cut}) \approx \left[\mathcal{P}(j \lesseqgtr \text{cut})\right]^2 . \tag{6.23}$$

- **Low-E_{T} selection:** The B2F analysis does not apply the $H_{\mathrm{T}}^{\mathrm{miss}}/H_{\mathrm{T}}$ and p_{T} cuts required in the CalRatio selection since the B2F analysis targets the low-energy phase-space with additional MET in the final state. Also, the cut on the energy fraction in the ECal (summed over all jets) needs to be scaled to account for the background rejection with one displaced jet only. Similarly, the cut on the neural network (NN) signal score *product* of both jets cannot be used directly in the B2F analysis, instead a cut on the leading jet NN signal score can be made; from Figure 6.17, an optimal cut value for signal-to-background discrimination is found at NN signal score > 0.95.
- **ABCD plane:** Given that the BDT score is taken as one of the axes of the CalRatio ABCD plane, it cannot be scaled to a corresponding ABCD plane for the B2F analysis, and is therefore skipped.

The background estimate is detailed in Appendix B resulting in about 1650 background events, which is quite large compared to the signal yields $\mathcal{O}(10-100)$; this gives quite bad signal-to-background discrimination, $S/B \sim \mathcal{O}(10^{-2})$, as well as quite low signal significance, $S/\sqrt{B} \sim \mathcal{O}(10^{-1})$, as shown for one exemplary grid point in Figure 6.18.

Since the signal yields depend on parameters, such as Yukawa-like couplings, mixing angle, and vacuum expectation value of the dark sector, the presented S/B and $S/\sqrt{B}$ curves might improve after fine-tuning these parameters.

The results indicate that a discovery or exclusion of the IDDM

Table 6.4.: Measured data rates in the CalRatio search. The CalRatio data rates are given for several steps throughout their analysis strategy. Numbers are presented for the SM multi-jet background, as well as the BIBs. Taken from [26].

	Low-E_{T} selection:	Main data	BIB
Preselection:	trigger, 2 clean jets	40 743 867	2 200 854
	$\sum \Delta R_{\mathrm{min}} > 0.5$	28 248 024	1 399 351
Event cleaning:	low-E_{T} BDT > 0.05	1 288 596	44 035
	Trigger matching	1 138 961	36 266
	$-3 < t < 15\mathrm{ns}$	1 123 239	35 245
	$\log_{10}(E_{\mathrm{H}}/E_{\mathrm{EM}}) > -1.5$ for $\mathrm{jet}^{\mathrm{sig}_1}$, $\mathrm{jet}^{\mathrm{sig}_2}$, $\mathrm{jet}^{\mathrm{bib}_1}$, $\mathrm{jet}^{\mathrm{bib}_2}$	1 038 019	33 100
	$\lvert\eta\rvert \notin [1.45, 1.55]$ for $\mathrm{jet}^{\mathrm{sig}_1}$, $\mathrm{jet}^{\mathrm{sig}_2}$	976 805	31 292
Low-E_{T} selection:	$H_{\mathrm{T}}^{\mathrm{miss}}/H_{\mathrm{T}} < 0.6$	965 748	30 712
	$p_{\mathrm{T}}(\mathrm{jet}^{\mathrm{sig}_{1\iota}}) > 80\,\mathrm{GeV}$	315 530	10 048
	$p_{\mathrm{T}}(\mathrm{jet}^{\mathrm{sig}_{2\iota}}) > 80\,\mathrm{GeV}$	73 484	2810
	$\sum_{\mathrm{jet}^{\mathrm{sig}_{1\iota}},\mathrm{jet}^{\mathrm{sig}_{2\iota}}} \log_{10}(E_{\mathrm{H}}/E_{\mathrm{EM}}) > 2$	3375	93
	Low-E_{T} NN product > 0.7	307	10
Region A:		23	0
Region B:		3	0
Region C:		220	7
Region D:		61	3

Table 6.5.: Rough cutflow of CalRatio and B2F at trigger and pre-selection level. The cuts at trigger and pre-selection level are compared, highlighting the difference of correlated MET in the previous BC for the B2FT *v.s.* additional jet in the same BC for the CalRatio selection.

Level	B2FT	CalRatio Trigger
L1	1 jet above 40 GeV, correlated MET	1 displaced jet above 30 GeV
HLT	1 displaced jet above 20 GeV	1 displaced jet above 30 GeV
Pre-Selection	1 jet above 40 GeV, $\Delta R_{\mathrm{min}} < 0.5$	2 jets above 40 GeV, $\Delta R_{\mathrm{min}} < 0.5$
Summary	1 displaced jet above 40 GeV correlated MET $\Delta R_{\mathrm{min}} < 0.5$	1 displaced jet above 40 GeV 1 additional jet above 40 GeV $\Delta R_{\mathrm{min}} < 0.5$

model might be quite challenging using the (preliminary) B2F analysis strategy. Nonetheless, upper limits can be obtained for BSM processes with comparable kinematics to the IDDM model, to show which cross-sections can be excluded with a 95 % confidence using the presented search strategy.

Limits on the production cross-section times the branching fraction can be derived using the CL_S method [87], which is a frequentist-motivated hypothesis test, with signal and background expec-

[87]: Read (2000), *Modified frequentist analysis of search results (the CL_s method)*

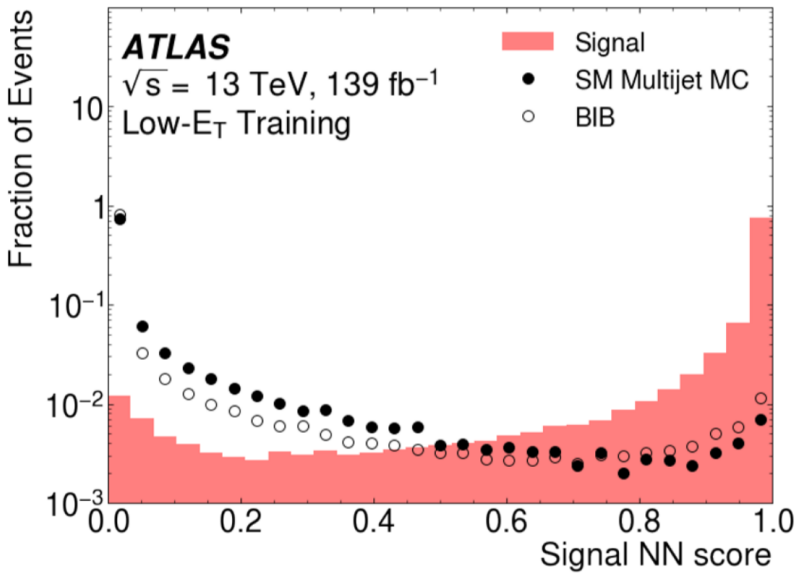

Figure 6.17.: CalRatio jet tagger NN signal score. The per-jet NN tagger from the CalRatio analysis is re-used for the background projection of the B2F analysis. The jet tagger has several calorimeter-based, tracker-based, and muon-spectrometer based input variables, and returns three scores, namely a signal score, a SM multi-jet score and a BIB score, indicating how likely a given jet is classified as signal-like, SM-like, or BIB-like. The scores range from zero to one, and all three scores add up to unity. For the background estimate, a cut on the signal score at 0.95 returns the best signal-to-background discrimination, and is therefore proposed for the B2F analysis. Taken from [26].

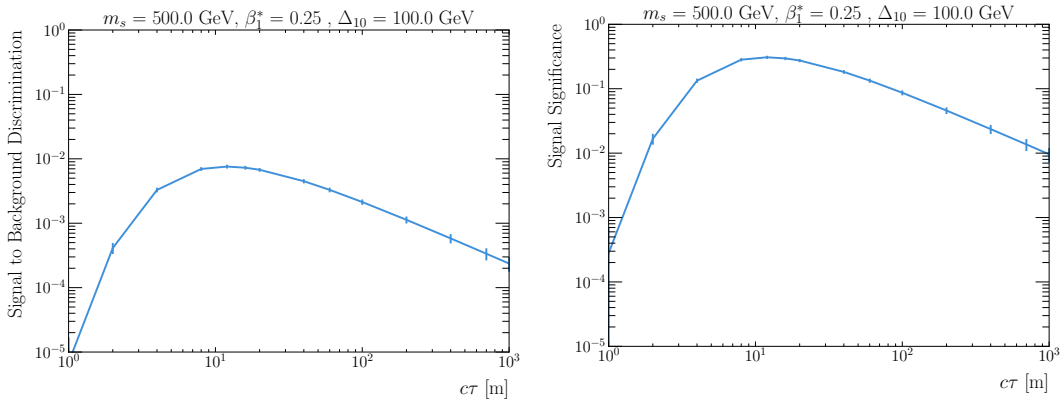

(a) Signal-to-background discrimination, S/B. **(b)** Signal significance, $S/\sqrt{B}$.

Figure 6.18.: Expected signal-to-background discrimination and signal significance. The expected signal-to-background discrimination as well as signal significance are presented for one exemplary grid point, showing quite low signal-to-background discrimination and signal significances.

tations entering a single-bin counting experiment. The test is performed per grid point with the `pyhf` package [88, 89] using the $\tilde{q}_\mu$ test statistic based on the profile likelihood ratio [90]. A global background uncertainty of 5 %, motivated by the CalRatio studies, is implemented as a Gaussian-constrained nuisance parameter. The resulting 95 % CL upper limits are presented in Figure 6.19 as a function of the scalar particle mass and final-state compression.

Overall the expected limits are rather weak with $\sigma_{\text{UL}} \times \mathcal{B} \sim \mathcal{O}(0.1\,\text{pb})$ in the optimal phase-space. In fact, the upper limits are — as expected from the weak signal-to-background discrimination, $S/B < 1$ — higher than the IDDM production cross-sections and even of the same magnitude as the overall $pp \longrightarrow s + X$ cross-sections as shown in Figure 6.5.

Another notable feature of the limits is their rather long tail towards

[88]: Heinrich et al. (n.d.), *pyhf: v0.7.6*
[89]: Heinrich et al. (2021), *pyhf: pure-Python implementation of HistFactory statistical models*
[90]: Cowan et al. (2011), *Asymptotic formulae for likelihood-based tests of new physics*

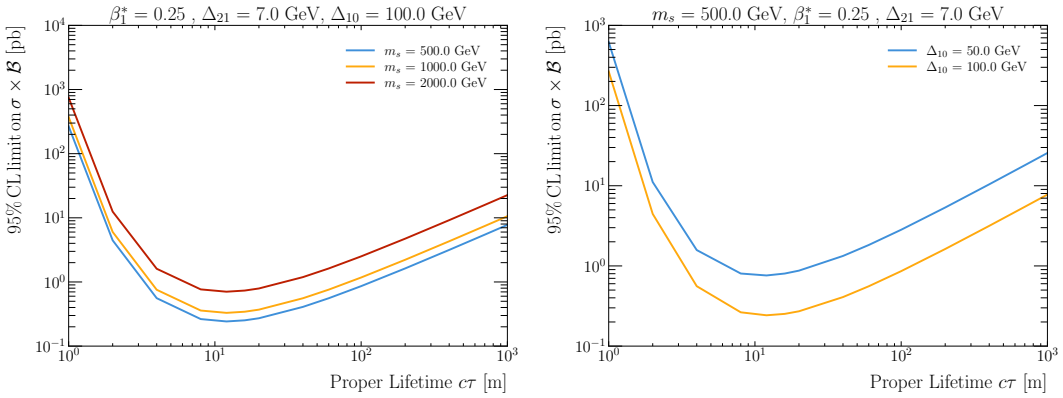

(a) Expected limits for various m_s samples.

(b) Expected limits for various Δ_{10} samples.

Figure 6.19.: Expected limits on production cross-section times branching fraction. The expected upper limits on $\sigma \times \mathcal{B}$ are presented as a function of the proper lifetime for several m_s and Δ_{10} samples. The dependences on m_s and Δ_{10} are inherited from the efficiencies as discussed above. The overall scale of the upper limits is quite weak, and, in fact, even higher than the IDDM production cross-section. Still, the presented limits serve as upper limits for BSM processes with comparable kinematics.

the higher lifetimes coming from the fact that two BCs are accessed, probing higher lifetimes. The dependences of the limits on the grid phase-space are inherited from the signal efficiencies as discussed above.

6.6. Comparison to Other BSM Searches

Given the high scalar particle masses, $m_s \gtrsim \mathcal{O}(100\,\text{GeV})$, only searches at the LHC, in particular at ATLAS and CMS, are potentially sensitive to the studied phase-space [82].

[82]: Lessa et al. (2024), *Back to the Future Trigger*

Searches for displaced vertices in the tracking detectors are generally sensitive to shorter lifetimes, and there are currently no other *calorimeter* searches sensitive to *one* displaced jet in the studied E_T range, due to the high QCD background at the trigger and analysis level. The presented B2F strategy reduces this background by more than two orders of magnitude, and therefore potentially capable of studying this regime. However, there is also the Muon Spectrometer Displaced Vertex (MSVtx) search that is sensitive to models with one displaced vertex. Naturally, the MSVtx search is also sensitive to higher lifetimes, and therefore potentially competitive.

In fact, the limits observed in the MSVtx search (for a different — but similar — benchmark model) are better by more than one order of magnitude compared to the expected limits derived with the preliminary B2F selection described above. The comparison is presented in Figure 6.20 for a low- and a high- energy case, where the effective centre-of-mass energy is controlled by the

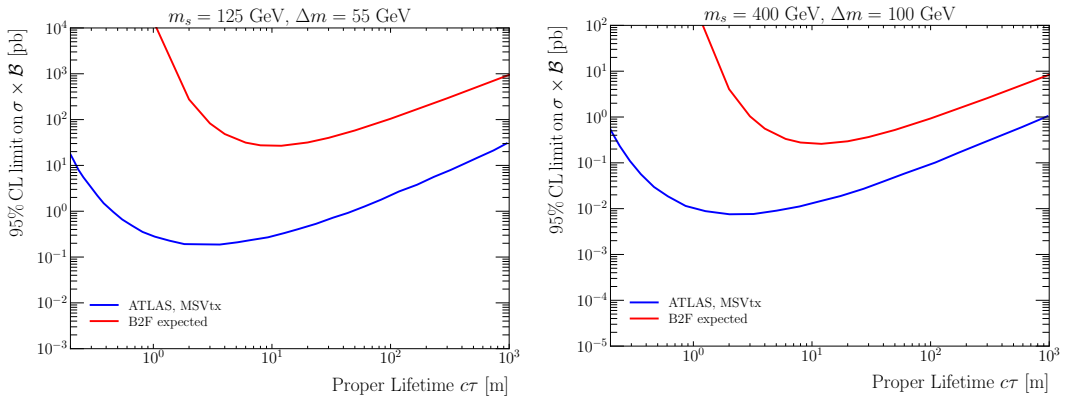

(a) Low E_T scenario.

(b) High E_T scenario.

Figure 6.20.: Comparison of expected limits to observed limits of other searches. The expected upper limits with the preliminary B2F analysis strategy are compared to observed limits of other LLP searches, in particular a seach for displaced vertices in the muon spectrometer of the ATLAS detector [91]. Here, Δm is defined slightly different for the B2F and MSVtx analyses since they use distinct benchmark models; in both cases Δm is supposed to be a measure of the energy available in the final state: for the IDDM in the B2F analysis $\Delta m = \Delta_{10} \sim E_{f\bar{f}}$ is chosen, while for the MSVtx search $\Delta m = m_s \sim E_{f\bar{f}}$.

scalar particle mass, $\sqrt{\hat{s}} \sim m_s$, and the jet energy by the final state compression, $E_j \sim \Delta m$.

These comparisons should, however, be taken with a grain of salt, since the MSVtx search contains a fully developed analysis strategy with optimised signal-to-background discrimination, whereas the presented search only applies a preliminary event selection inspired by a CalRatio search. To develop a full B2F analysis strategy, a simulation and reconstruction framework is required, as introduced in Chapter 7.

Moreover, two different benchmark models are compared and for a fair comparison a re-interpretation of the MSVtx analysis strategy with the IDDM model is required. Here, the main differences are that the IDDM model contains MET in the final state (from the χ_0 particles) whereas the model used in the MSVtx search[31] only contains visible particles in the final state. The additional MET would likely affect the MSVtx event selection.

31: The MSVtx search is benchmarked with a $pp \to \phi \to s(f\bar{f}) + s(f\bar{f})$ model, where both ϕ and s are scalars, and the s are long-lived.

Since the BSM sector also couples to the SM in these kind of models, *e.g.* through $gg \to s$ in the IDDM model, the SM decay channel, *e.g.* $s \to gg$ or $s \to f\bar{f}$, can not be avoided. That is, resonance searches could also probe for the same BSM extensions.

However, the SM decay channels are proportional to $\sin\alpha$, whereas the BSM process $s \to \chi_1\chi_2$ is proportional to $\cos\alpha$. Since $gg \to s$ is also proportional to $\sin\alpha$, the process $pp \to s \to$ SM is somewhat proportional to $\sin^2\alpha$ and $pp \to s \to \chi_1\chi_2$ to $\sin\alpha \times \cos\alpha$. That is, with careful tuning of the $\sin\alpha$ mixing angle, resonance searches can be disregarded as potential competition; additional

care is needed in this case such that the $\chi_1\chi_2$ cross-section is still sufficiently high.

6.7. Conclusions and Outlook

The proposed trigger strategy is benchmarked with a toy model inspired by the Inelastic Dipole Dark Matter (IDDM) using a simplified analysis strategy at the MC truth-level. Signal efficiencies are derived, which peak below the percent level, corresponding to signal yields of $\mathcal{O}(10 - 100)$ assuming the Run-2 luminosity. The expected backgrounds are extracted from a previous CalRatio search, resulting in $\mathcal{O}(1600)$ expected background events.

Signal yields and background rates are used to calculate upper limits on the production cross-section. Here, the limits are quite weak, $\mathcal{O}(0.1\,\text{pb})$, mainly due to two reasons:

- Low signal efficiencies: too high LLP boosts and ISR induced shadowing effects;
- High background rates: no full analysis strategy with signal-to-background optimised cuts.

In particular the first point is quite concerning since the model is already quite streamlined to the desired signature by tuning the couplings and masses. Therefore, another benchmark model is likely required to properly address the boost issue. Possible models known to provide slow particles are heavy ionising particles, potentially even being "stopped" in the calorimeter, sitting there until they decay. These models were not considered in the presented feasibility and sensitivity studies, since they do not necessarily provide the desired MET-jet correlation, because of the LLPs' electrical charge.

To address the ISR shadowing effect, looking into the the next BC instead of the previous BC could be considered; however, this is technically quite challenging in the current design of the ATLAS trigger [92].

[92]: Gugel (2025), *Private Communication*

Given the promising background suppression it seems appropriate to further investigate on the proposed trigger strategy despite the pending physical motivation given the weak limits. In fact, a more involved analysis strategy needs to be developed to make a fair comparison to other searches. For this, the development of a framework to process out-of-time decays in the internal ATLAS framework is required, as detailed in the next chapter.

7. Search for Delayed LLP Decays with the ATLAS Detector

The presented search aims to compare the SM-subtracted data fired by the B2FT to BSM predictions, *e.g.* to the IDDM model as introduced in Chapter 6. Therefore, this chapter is two-folded: the first part focuses on the trigger development and validation for successful data-taking, while the second part addresses precise BSM predictions through detector simulations, including the reconstruction of out-of-time decays in the ATLAS detector, and discusses the calibration of displaced jets and calorimeter timing.

7.1. Development, Validation, and Optimisation of *Back-to-the-Future* Triggers

The B2FT as introduced in Chapter 5 (referred to as *main* B2FT in the following) is implemented in the ATLAS trigger system utilising L1Topo. It roughly recorded 240 fb^{-1} of pp collision data in 2024 and 2025 combined.

In addition, a background-enriched control-region (BG-CR) trigger with relaxed $\Delta\varphi(\text{MET}_{N-1}, \text{jet}_N)$ correlation is implemented for background studies; while the main trigger requires $\Delta\varphi < 1.0$, the BG-CR trigger accepts all events with MET > 40 GeV in $N-1$ followed by a jet with $E_\text{T} > 40$ GeV regardless of the $\Delta\varphi$ correlation between these objects. Given that it is a control region trigger, a pre-scale is applied with a factor of 280 being conservatively high.

The main B2FT is running with a rate around 50 Hz at L1 and below 1 Hz at the HLT, while the BG-CR trigger is running with rather low rates of about 0.5 Hz at L1, and below 10 mHz at the HLT.

The main B2FT was observed to misbehave at the beginning of the 2025 data taking [92, 93]; in particular, a quite high rate (peaking at 4 kHz at L1) was observed, and, as shown in Figure 7.1a, a disagreement in the jet p_T (and multiplicity) distributions for 2024 and 2025 data.

The malfunctioning was caused by a misconfiguration of the L1Topo firmware; here, the delay option was dropped, and instead jets and MET were correlated in the *same* BC, resulting in a higher rate and a different event topology, making about 25 fb^{-1} of the 2025 data set un-useable for the B2F analysis[32].

[92]: Gugel (2025), *Private Communication*
[93]: Kerridge (2025), *Private Communication*

32: In particular the runs between 499248 and 500603.

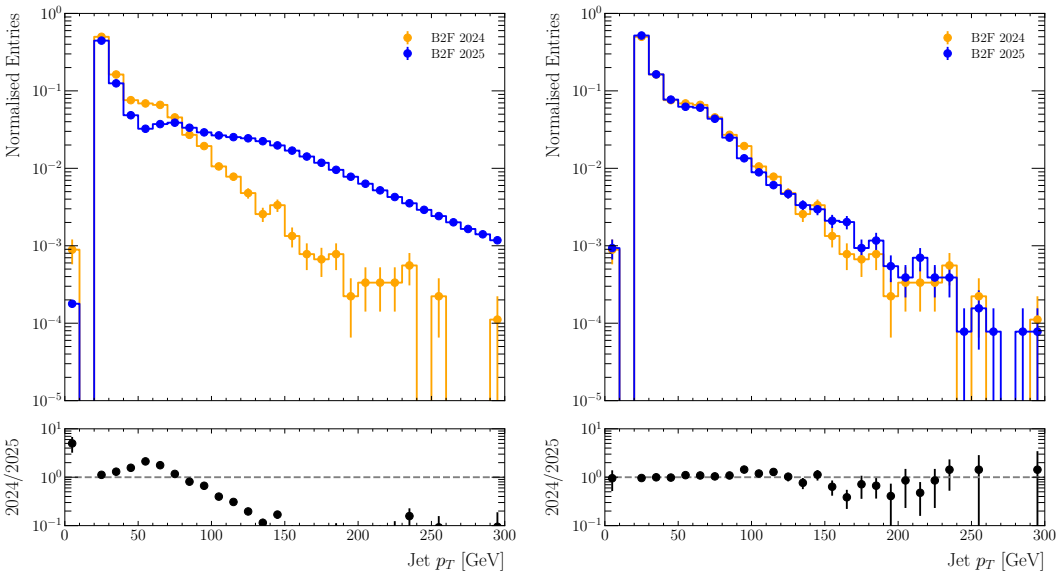

(a) Comparison of Data-Run `476218` (2024) to `499586` (2025, misconfigured L1Topo firmware).

(b) Comparison of Data-Run `476218` (2024) to `502782` (2025, fixed L1Topo firmware).

Figure 7.1.: Transverse momentum distribution of jets fired by the B2FT in the 2024 and 2025 data taking. A discrepancy between the jet p_T in 2024 and 2025 data is observed for runs 499248 through 500603 caused by a misconfiguration of the L1Topo firmware. Instead of correlating jets and MET of adjacent BCs, correlations were triggered on the same BC (due to the firmware misconfiguration) resulting in a quite different event topology with higher jet multiplicities and jet p_T.

The firmware configuration was fixed after the 2025 heavy ion runs, after which the B2FT rates normalised to the 2024 level, and the jet multiplicity and jet p_T distributions agree reasonably, as shown in Figure 7.1b.

The B2FT is also validated against the CalRatio trigger, as shown in Figure 7.2. Here, a significant difference in the energy fraction in the ECal is observed, which is unexpected since the two triggers share a similar HLT, in particular with the same cut on the energy fraction in the ECal, EMF $< 6.4\%$.

It is found that there is an additional algorithm in the HLT chain used for the B2FT, which is, however, not included in the HLT chain of the CalRatio triggers. This algorithm attempts to remove pile-up contributions (which could potentially affect the EMF of displaced LLP decays) and to re-calculate the EMF at the trigger level. However, this algorithms seems to misbehave in such that only pile-up contributions in the ECal are removed, but not in the HCal, which artificially decreases the EMF (at the trigger level), and hence lowers the threshold to pass the HLT, as sketched in Figure 7.3.

Since the B2FT applies the re-calculation algorithm at the trigger level before applying the EMF $< 6.4\%$ cut, jets with larger off-line EMF pass the HLT compared to the CalRatio trigger which does not apply the re-calculation algorithm; this effectively results in

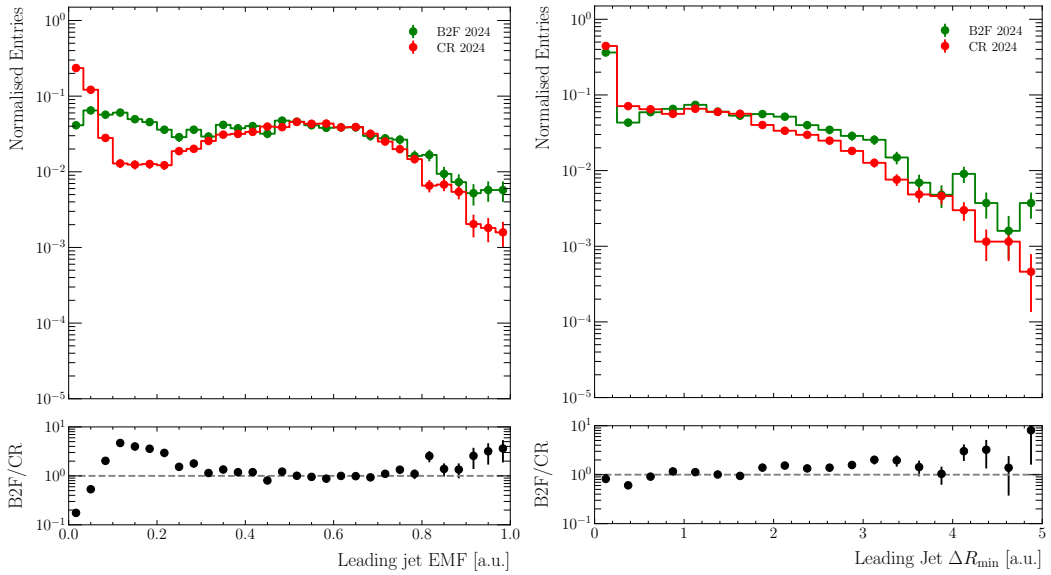

(a) Lowest-EMF (=: leading) jet EMF. **(b)** Leading jet ΔR_{min}(jet, track > 2 GeV).

Figure 7.2.: Comparison of relevant jet distributions fired by the B2FT and CalRatio trigger. The jet distributions relevant at the HLT are compared for the B2F and CalRatio triggers in 2024 data (Run 486706). Because the B2F and CalRatio triggers share a similar HLT, comparable EMF and ΔR_{min}(jet, closest track > 2 GeV) are expected — there might be small differences because of the different event topologies due to distinct L1 triggers. However, a significant discrepancy in the off-line EMF distribution is observed below EMF $\lesssim 0.24$ caused by a misbehaving algorithm for pile-up removal in the B2F HLT. While the CalRatio trigger fires on jets with EMF < 6.4%, the B2F HLT trigger uses jets with EMF < 0.24 and removes pile-up contributions from them, re-calculates their EMF at the trigger level, and accepts events with a jet with re-calculated EMF < 6.4%. However, the pile-up removal is only removing pile-up contributions in the ECal but not in the HCal, artificially lowering the re-calculated EMF, and hence lowering the threshold to pass the B2F HLT. This results in a degraded signal-to-background discrimination in the B2FT, and an optimisation of the cut value would be needed — in particular the 6.4% cut would likely needed be increased.

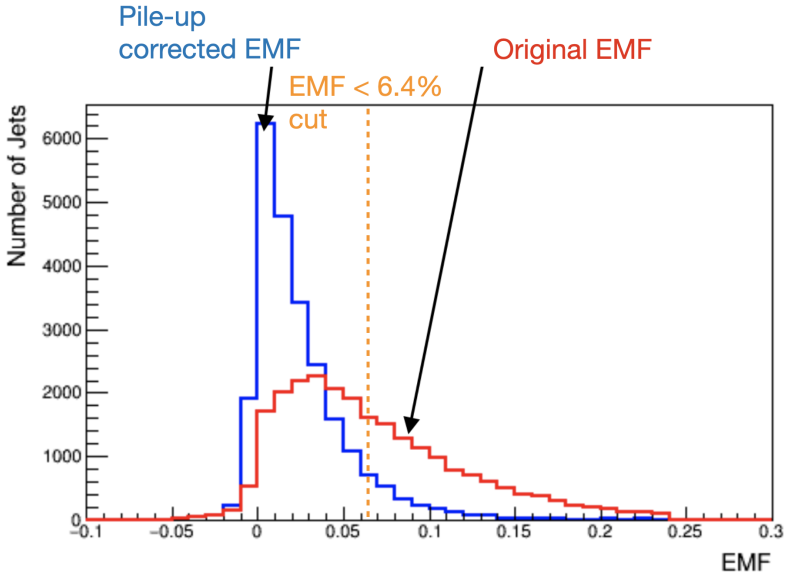

Figure 7.3.: Sketch of EMF re-calculation. This sketch shows the original EMF distribution (as it is also stored for standard off-line analysis) and the pile-up-corrected EMF distribution for an arbitrary signal MC sample. In general, pile-up may increase the EMF of displaced decaying LLPs, making signal jets less likely to pass the EMF < 6.4% cut. However, since pile-up contributions are only corrected for in the ECal but not in the HCal, the re-calculated EMF distribution is over-corrected making all jets more likely to pass the EMF cut, degrading the signal-to-background discrimination power of the EMF variable. Sketch based on [94].

a broader offline EMF distribution after the B2F trigger selection compared to CalRatio.

This not only makes signal jets more likely to pass the EMF cut at the trigger level, but it also degrades the signal-to-background

discrimination power of the EMF variable. The differences in the EMF between the trigger and off-line levels would require an optimisation of the EMF cuts at the trigger level, which, however, has not been addressed yet; the trigger cuts are currently under investigation by the CalRatio team, and for the meantime, the B2FT is adapted to schedule the standard CalRatio HLT chain without the pile-up removal and EMF re-calculation [95].

[95]: Ovsiannikova et al. (2025), *Private Communication*

The differences between the trigger and off-line levels also explain the limited performance of the naive trigger matching approach at the off-line level, where the fraction of events fired by the B2FT is studied that contains at least one jet with EMF $< 6.4\%$ and $\Delta R_{\text{min}} > 0.2$. Here, it is observed that only $\mathcal{O}(10\%)$ of the events at the off-line level do contain at least one such jet which potentially fired the B2FT, while it looks slightly better for the CalRatio trigger, where $\mathcal{O}(40\%)$ of the events contain at least one such jet.

7.2. Simulation and Reconstruction Framework

To compare the BSM predictions to the data, the simulation of the detector response and the detector readout are needed, which is quite non-standard, given the novel analysis strategy. While the simulation of the detector response in Geant4 does not apply the BC structure, this is done in the subsequent digitalisation step in the Athena framework. To avoid loosing the out-of-time part, a dedicated simulation and reconstruction strategy is required for the out-of-time part.

A rather model-independent approach[33] for this is introduced in Section 4.3 and is detailed in the remaining part of this chapter. Here, the MC samples are generated as detailed in Section 6.2, and are thereafter put through Geant4 to simulate the detector response. To avoid a waste of computational resources, an event filter is used to only simulate events where the LLP decays out-of-time and inside the HCal.

33: For the case that another model is found to perform better compared to the IDDM model.

The on-time part is processed with the standard ATLAS framework, while the out-of-time part is reconstructed separately by shifting the Geant4 hits by one BC before putting it through the standard reconstruction in the Athena framework. Subsequently, pile-up is overlaid individually on-top of the two parts during the digitisation step, but this has been omitted in the samples used in the presented studies for clarity/illustrational reasons.

After the reconstruction step, usually the trigger decision is simulated. However, a full simulation of the B2F trigger decision is not yet available within the Athena framework.

The result of the event generation, simulation and reconstruction (without the trigger simulation) is exemplarily shown for one event in Figure 7.4, which clearly illustrates the correlation between the on-time and out-of-time parts:

- In the on-time part (Figure 7.4a), the prompt jet and the MET from the LLP recoil against each other, back-to-back in the transverse plane;
- In the out-of-time part (Figure 7.4b), the delayed jet is perfectly pointing in the same direction as the on-time MET in the transverse detector plane, and is also recoiling back-to-back with the prompt jet in the longitudinal plane.

The simulation and reconstruction framework is validated against the truth-level framework as introduced in Section 6.2. Here, the results for the on-time part are not explicitly presented, since the on-time part is processed with the standard ATLAS framework and shows perfect closure. Also for the out-of-time part processed with the user-specific `HitTimeWrapperTool`, excellent agreement is observed in η and φ as shown in Figure 7.5, where ΔR_{min}(reco jet, closest truth jet) shows a sharp peak at 0 demonstrating that the reconstructed and truth-level jets are perfectly matched in η and φ. Also other important distributions, such as the jet timing and the energy *fraction* in the ECal agree reasonably well as shown in Figure 7.6 — however, these quantities can only be *approximated* at the truth-level.

There is one important exception, since the energy distributions (E, E_{T}, p_{T}, ...) show significant deviations of $\mathcal{O}(20-30\%)$ between the reconstruction and truth-levels, as shown in Figure 7.7; this non-closure is caused by a non-optimal energy calibration of displaced jets, since their topology is quite different to SM-like prompt jets which are used for deriving the calibration as detailed below [96].

[96]: Dandoy et al. (2025), *Internal Communication*

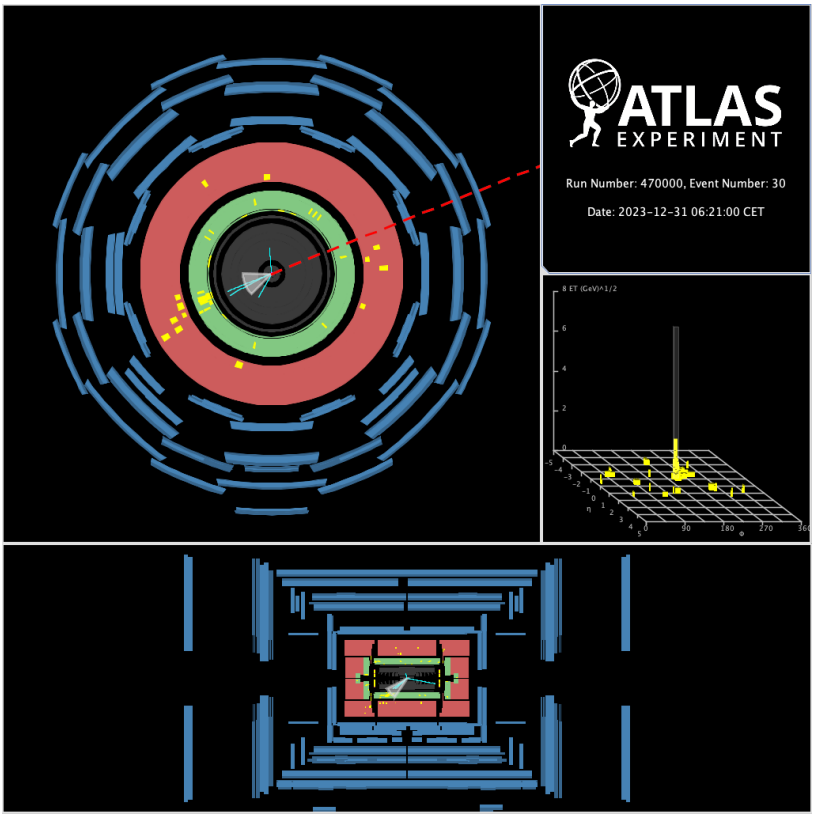

(a) On-time event display.

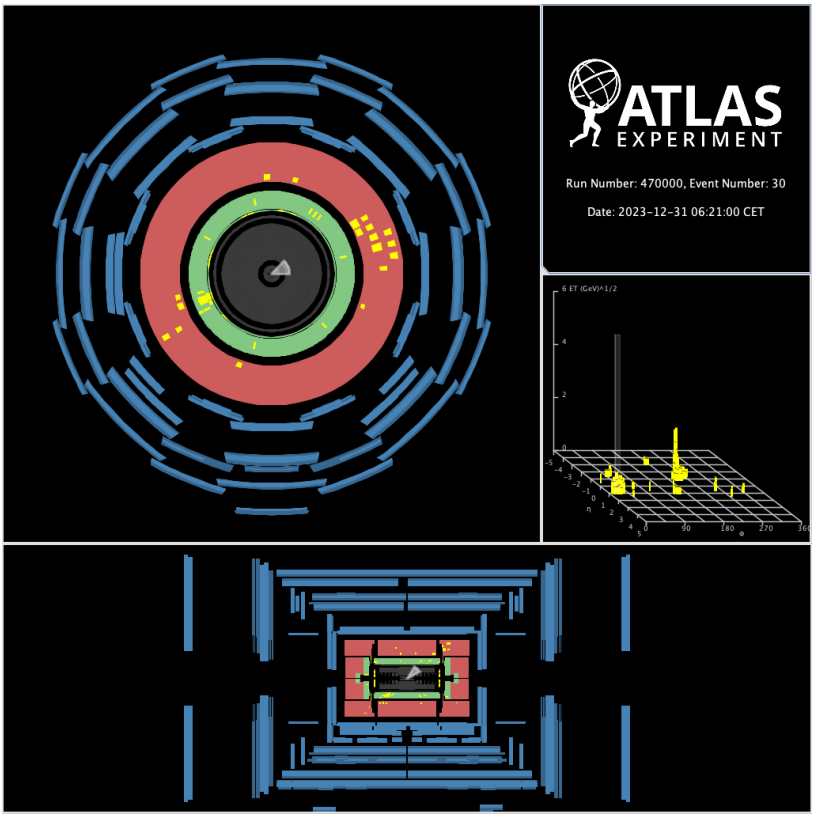

(b) Out-of-time event display.

Figure 7.4.: Event display of the signal MC reconstruction. The simulation and reconstruction is presented for one exemplary event which shows a particular illustrative correlation between the on-time and out-of-time part, since it does not contain hard ISR in the on-time part. The prompt jet and the missing transverse momentum from the LLP recoil back-to-back in the transverse plane, and the delayed jet is perfectly pointing towards the same azimuthal direction compared to the on-time MET, and is additionally recoiling against the prompt jet back-to-back in the longitudinal plane. Furthermore, the event display visualises the distinct topologies of the prompt and delayed jets, with latter originating displaced in the HCal with only little energy deposits in the ECal. Another feature of the ATLAS calorimeter is visible in the out-of-time display, which shows that the on-time energy deposits are still present in the out-of-time part, which comes from the rather broad calorimeter pulses.

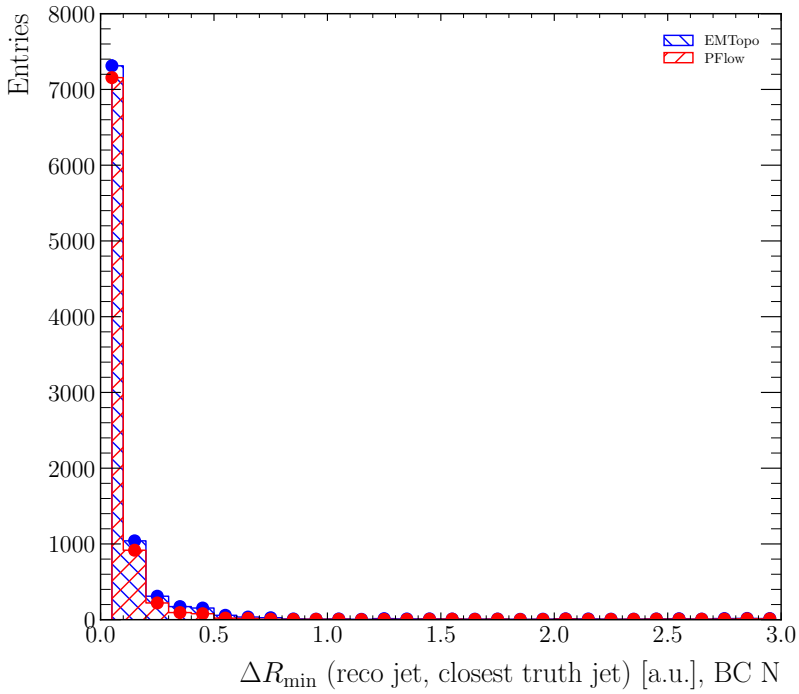

Figure 7.5.: Matching of reconstructed and truth-level jets. The reconstructed EMTopo and EMPFlow jets are matched to their closest truth-level anti-k_t jet based on η and φ. The sharp peak at 0 indicates that the reconstructed jets align perfectly with the truth-level jets in η and φ direction, showcasing that the dedicated reconstruction of the out-of-time part is yielding reasonable results.

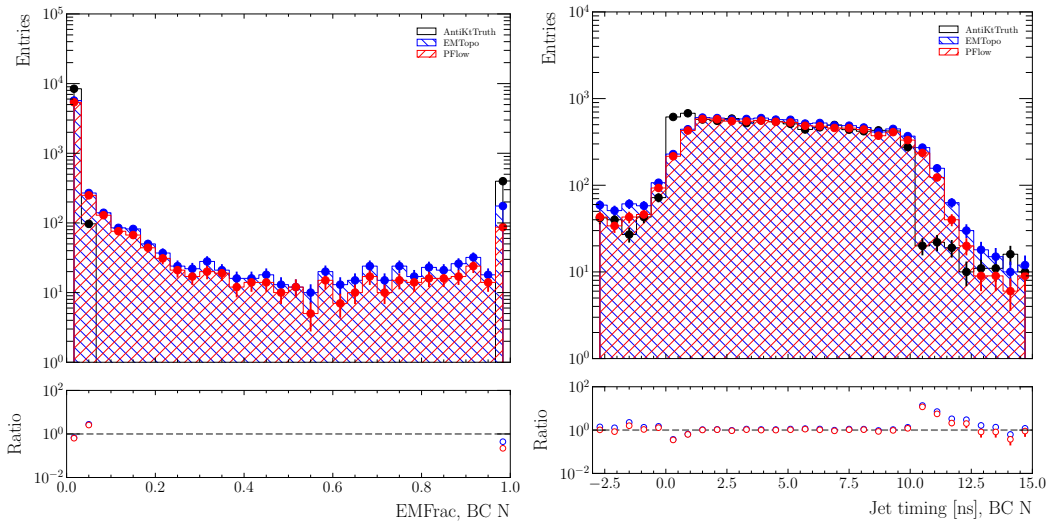

(a) Jet EMF. **(b)** Jet timing.

Figure 7.6.: Comparison of relevant distributions for reconstructed and truth-level jets. The jet EMF and timing distributions are compared for the different reconstructed jet collections, as well as truth-level anti-k_t jets. Overall, good agreement is observed between EMPFlow and EMTopo jets. The EMF and the timing of a jet can only be *approximated* at the truth-level based on the jet origin calculated as energy weighted average of the jet constituents' production vertices. In particular the EMF value at the truth-level is quite binary, either (*i*) 0 indicating that the jet originates in the HCal, or (*ii*) 1 indicating that the jet originates in the tracker or the inner part of the ECal, or (*iii*) 6.4% indicating that the jet originates at the outer edge of the ECal, or (*iv*) undefined for jets originating beyond the HCal.

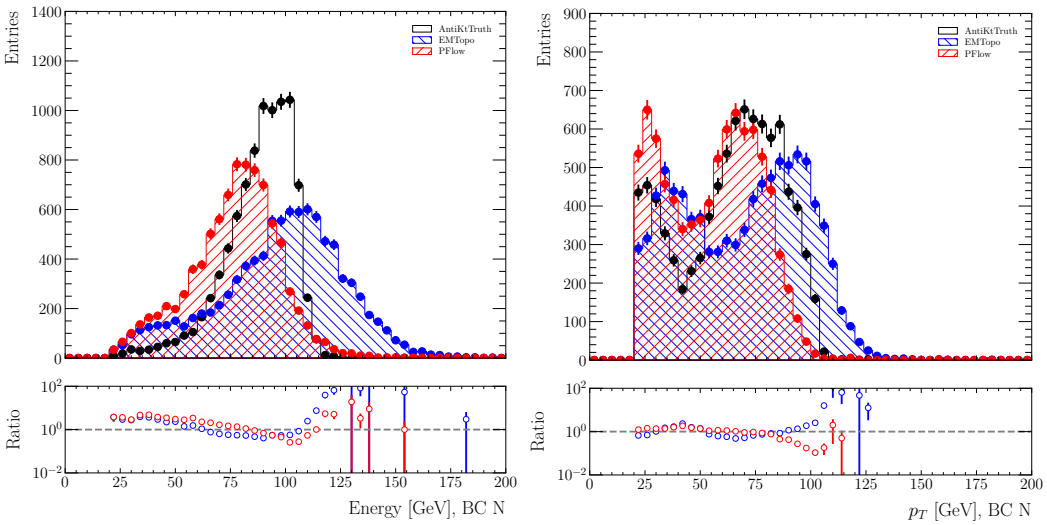

(a) Jet energy distribution.

(b) Jet transverse momentum distribution.

Figure 7.7.: Energy mismatch between reconstructed and truth-level jets. The jet energy and momentum distributions show a significant difference of $\mathcal{O}(20-30\%)$ between reconstructed and truth-level jets, caused by a non-optimal energy calibration of displaced jets. Moreover, the EMTopo and EMPFlow jet collections show quite different behaviour due to their distinct nature and calibration schemes, as discussed in more detail in the following section.

7.3. Energy Calibration of Displaced Jets

Although the energy scale is of less concern for BSM searches but the data-MC agreement is much more important, detailed jet energy calibration studies are performed, since the above-mentioned miscalibration is likely not a feature of the time delay but the spatial displacement of the jets, and hence affecting searches for displaced jets in general.

Given the trackless nature of displaced jets originating from LLP decays inside the calorimeter, topological jets clustered at the EM scale (EMTopo) are generally found to behave better for LLP searches compared to particle-flow jets clustered at the EM scale (EMPFlow)[34] [97]. This is because EMTopo jets use only topological calorimeter clusters as input for the jet clustering, while EMPFlow jets additionally rely on tracks to improve the jet energy resolution.

In principle, the uncalibrated jet energies (jet constituent level) aligns for EMTopo and EMPFlow jets in case of trackless jets, since the two clustering algorithms are equivalent in the absence of tracks; however, subsequently individual jet energy calibration schemes are applied to the different jet collections, causing distinct behaviour in the calibrated energies: EMTopo jets are generally much more affected by the jet calibration compared to EMPFlow jets, since the calorimeter clusters require a high correction due to the non-compensation of the calorimeter, while EMPFlow jets

34: Here, "EM" is a slightly misleading naming convention in this context, since "EM" does not directly refer to the EM scale, but rather to the fact that the jet clustering is happening at the EM scale; hadronic calibrations are subsequently applied at the jet level.

[97]: Burzynski (2025), *Private Communication*

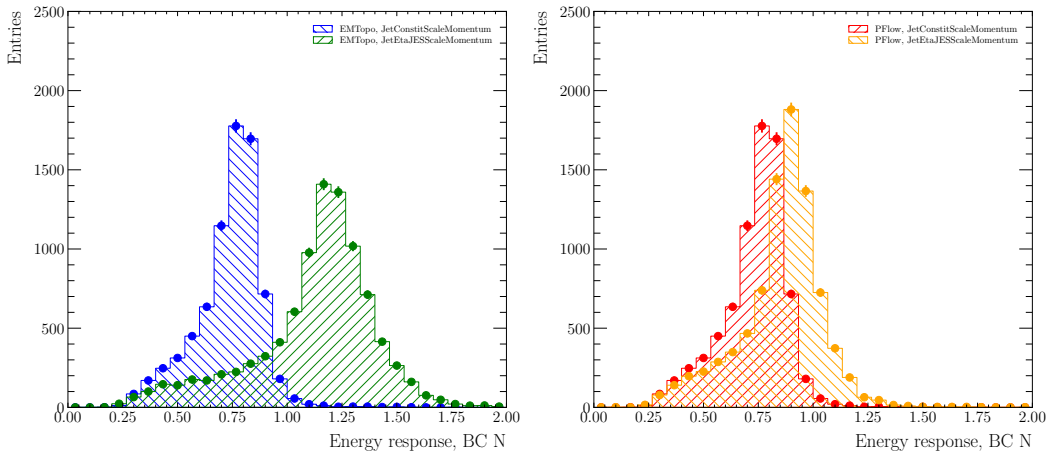

(a) EMTopo jet collection.

(b) EMPFlow jet collection.

Figure 7.8.: Energy response of displaced jets before and after the calibration. The energy response of displaced jets originating from LLP decays inside the HCal are presented for EMTopo and EMPFlow jets before ("jet constituent scale") and after the MCJes calibration. The uncalibrated energies align for EMTopo and EMPFlow jets, but do *not* account for non-compensation of the calorimeter, resulting in energy response shifted to values < 1. The MCJes calibration scheme is different for EMTopo and EMPFlow jets, resulting in different behaviour: the corrections applied on EMTopo jets are generally larger compared to EMPFlow jets which have significant contributions from tracks not requiring calibration. An *over*-correction is observed for EMTopo jets.

have significant contributions from tracks which do not require such a correction [98].

[98]: Bartels (2025), *Private Communication*

The energy scale is calibrated in the so-called MCJes step, which is a MC derived correction such that the energy response, $\mathscr{R} = E_{\text{reco}}/E_{\text{truth}}$, is a Gaussian distribution centred around 1. However, a significant shift is observed for the energy response of both EMPFlow and EMTopo jets at the MCJes level, as shown in Figure 7.8, indicating issues with the calibration of the jet energy *scale*. The response at the uncalibrated EM scale aligns well for the EMTopo and EMPFlow jets, as shown in Figure 7.9, confirming that it is not a reconstruction but rather a calibration issue.

Here, the issue is that the corresponding MCJes calibration factors are derived from (prompt) SM jets, and *no* specific calibration scheme is available for displaced jets, since it is not trivial to obtain generic calibrations for displaced jets, which would strongly depend on the jet origin (*i.e.*, on the LLP decay position and thus its lifetime).

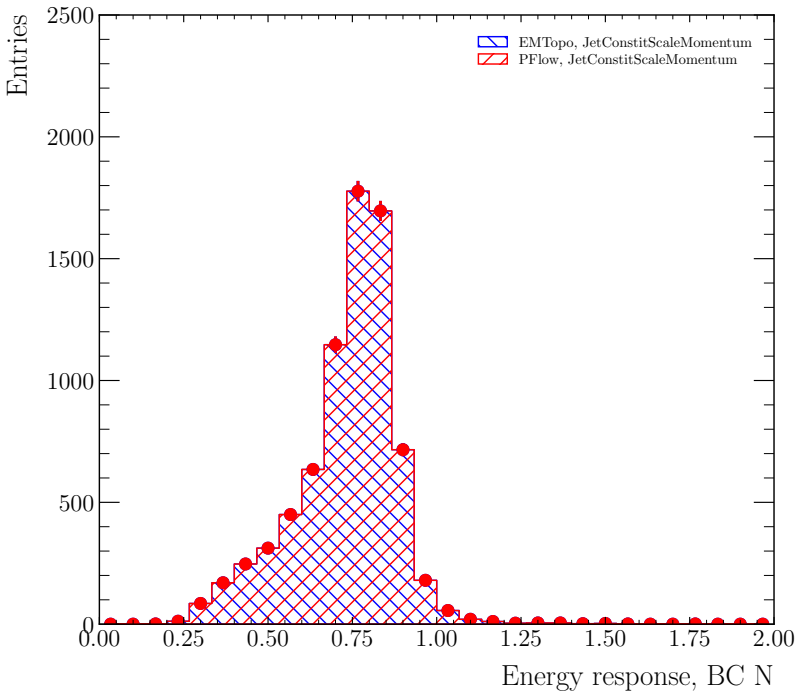

Figure 7.9.: Energy response at the EM scale. The energy response is presented for EMTopo and EMPFlow jets at the EM scale (uncalibrated). Perfect agreement is observed between the two jet collections, since the out-of-time jets are trackless, in which both jet clustering algorithms are equivalent. The non-compensation of the calorimeter is not accounted for at the EM scale, resulting in energy responses below 1.

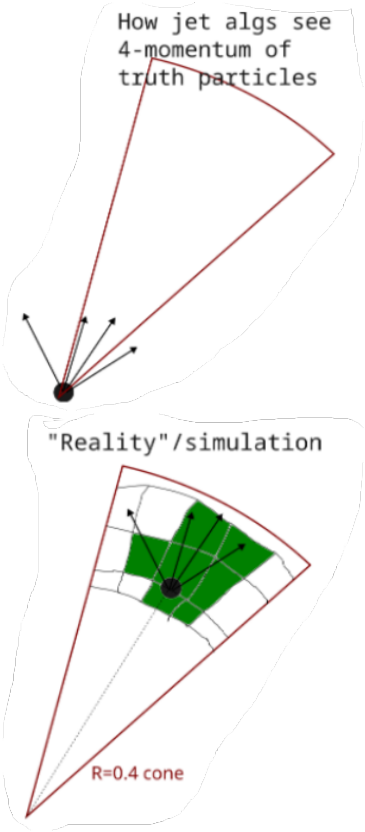

Figure 7.10.: Displaced jets in the MCJes calibration. Different treatment of displaced jets at the reconstruction and truth-levels. Taken from [96].

The application of these SM calibrations on displaced jets could explain the *over*-calibration of EMTopo jets in multiple ways:

- The calibration assumes that jets originate promptly from the interaction point. Under this assumption, a part of the QCD radiation associated with the initiating parton could fall outside the jet cone, as illustrated in the upper part of Figure 7.10. These out-of-cone losses are accounted for as part of the MCJes calibration, which applies correction factors to scale the reconstructed and truth jet energies.
  For displaced jets, however, the prompt-assumption no longer holds. Since the jet originates away from the interaction point, radiations that would be outside the jet cone in the prompt case are instead clustered into the jet, as shown in the lower part of Figure 7.10. Therefore, *no* out-of-cone energy losses are present for displaced jets. However, since the the calibrations are derived from SM simulations, the out-of-cone correction factors are still applied as part of the MCJes calibration, leading to an over-correction of displaced jets.
- The untypical low energy fraction in the ECal (EMF) and high energy fraction in the HCal (TileFrac) might cause an overcorrection of the calorimeter non-compensation, since the displaced jets are assumed to be quite hadronic given the low EMF. This might be the correct procedure for displaced SM jets (which are indeed typically hadronic), but is certainly not true for LLP-induced displaced jets, which have a usual electromagnetic and hadronic contribution.

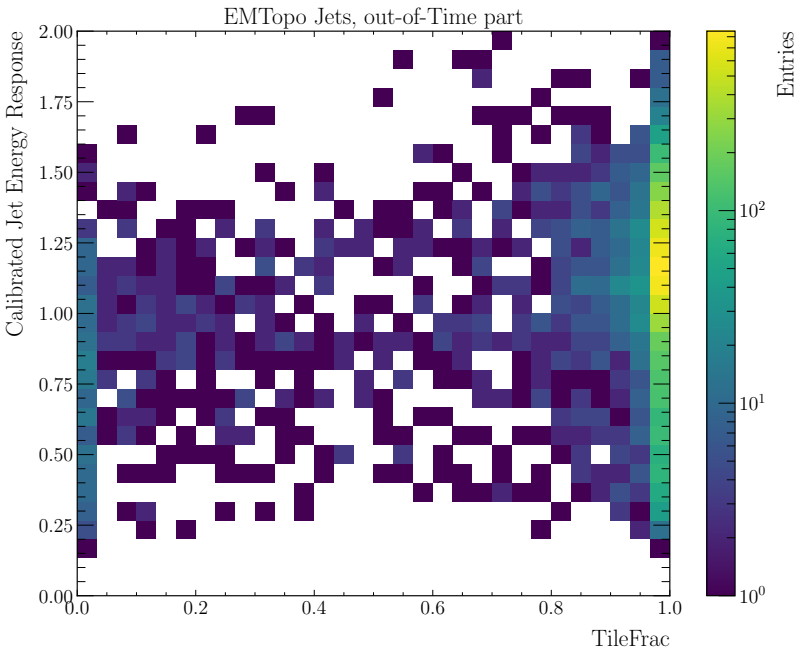

Figure 7.11.: Jet energy response *v.s.* TileFrac. The jet energy response is presented for EMTopo jets at the MCJes scale as a function of the energy fraction in the HCal. The energy response is best for low TileFrac (which is the usual SM case), and degrades with increasing TileFrac, particularly resulting in an over-correction.

The jet energy response is shown as a function of the energy fraction in the HCal in Figure 7.11, demonstrating that the EMTopo calibration performs best at low TileFrac (which is the usual SM case), but develops an over-calibration with increasing TileFrac.

- Displaced jets have unusually low energy deposits in front of the ECal, and unusually high energy leakage behind the HCal, which also might be incorrectly addressed in the jet energy calibration. However, since the number of associated muon segments is zero for most jets, as shown in Figure 7.12, the leakage behind the HCal is expected to have a sub-dominant role only.

As mentioned above, the actual energy scale is of less interest for BSM searches, but a good data-MC agreement is more important. Therefore, the calibration scheme is not required to represent the energy scale perfectly accurate, but it is required to be consistent in data and MC. Either the EM scale (uncalibrated) or the ATLAS recommended SM-calibrations could be used *consistently* in data *and* MC. Here, the EMTopo calibrations are preferred in the latter case, since they do not depend on tracks, and are more suitable for LLP searches.

The latter approach has the advantage that the calibration is appropriate for the SM multi-jet background, and already comes with an uncertainty for the backgrounds according to the standard ATLAS recommendations. However, it is not obvious how to handle the uncertainties for the signal MC. Here, it is proposed to use the standard uncertainties with a conservative factor, *e.g.* to double the uncertainties [96].

[96]: Dandoy et al. (2025), *Internal Communication*

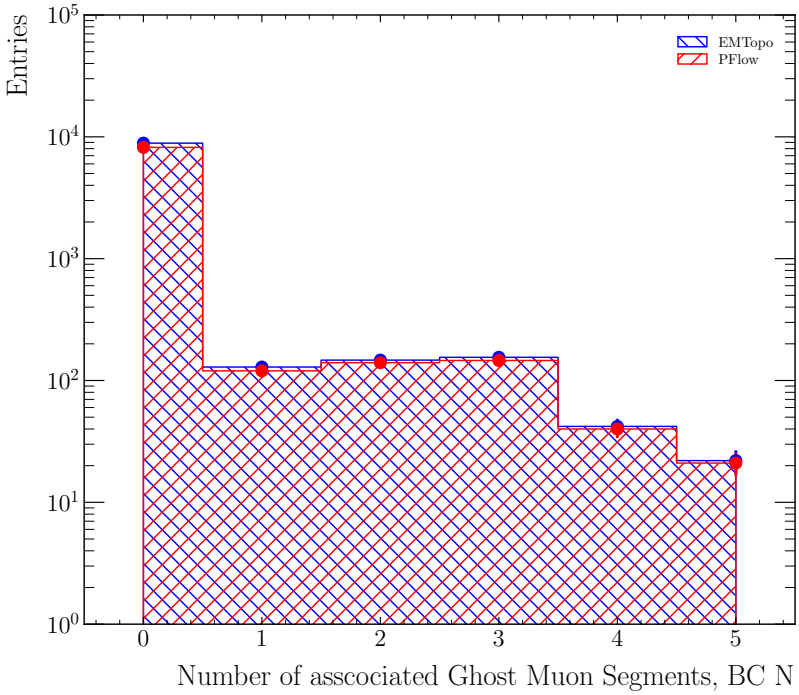

Figure 7.12.: Number of associated muon segments. The number of associated muon segments is presented for EMTopo and EMPFlow jets. Since most jets do not contain associated muon segments, the impact of atypically high energy leakage beyond the HCal is expected to be of subdominant importance for the energy (mis)calibration.

7.4. Calorimeter Timing

The event display (Figure 7.4) also illustrates another feature of the ATLAS calorimeter, namely that the energy deposits of the prompt jet are still present in the out-of-time part, since the shaped calorimeter pulses span a broad time window across multiple BCs. Therefore, energy deposits from $N-1$ contributing to N cannot be avoided, but using optimal filter (on-line) allows to reduce their contributions: while the filtered response remains about 1.0 for an in-time pulse, a pulse shifted by one BC still contributes with a weight of about 0.5, and pulses two and three BCs earlier would contribute with negative weights about -0.3 and -0.5, respectively [99].

[99]: Strizenec et al. (2025), *Private Communication*

This behaviour is particularly problematic when the delayed jet and a prompt jet point in similar η and φ directions, leading to energy deposits in the same calorimeter cells that cannot be separated; here, the prompt energy deposits would increase the energy of the delayed jet (with a weight of 0.5).

Matching reconstructed jets in BC N with their closest reconstructed jet in $N-1$ shows, however, that this has negligible impact, since there are hardly any jets in $N-1$ and N pointing in the same direction, as presented in Figure 7.13.

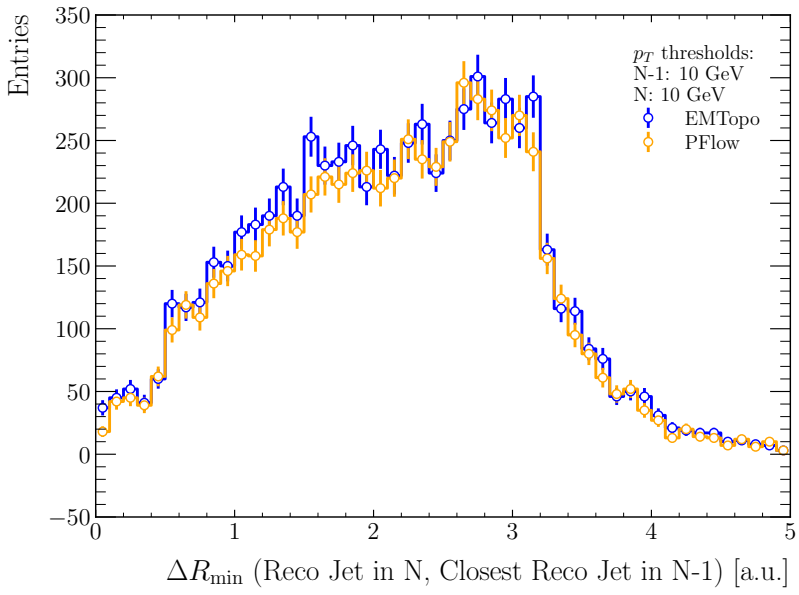

Figure 7.13.: Matching of reconstructed jets in N and $N-1$. Reconstructed jets with $p_{\text{T}} > 10\,\text{GeV}$ in BC N are matched to their closest reconstructed jet with $p_{\text{T}} > 10\,\text{GeV}$ in BC $N-1$, based on η and φ. Overall, most out-of-time jets are well separated from in-time jets, showing that out-of-time contributions from the on-time jets (due to their broad calorimeter pulse) are negligible.

7.5. Conclusions and Outlook

The proposed trigger is implemented in the ATLAS Level-1 Topological Trigger (L1Topo), and is running since 2024 corresponding to roughly $240\,\text{fb}^{-1}$ of pp collision data at 13.6 TeV as of the end of the 2025 data taking. Additionally, a background-enriched control region trigger is running since 2025, corresponding to roughly $100\,\text{fb}^{-1}$.

A misconfiguration in the 2025 L1Topo firmware led to a misoperation of the B2F triggers in early 2025 data taking, resulting in a loss of approximately $25\,\text{fb}^{-1}$ of the B2F dataset. After fixing the misconfiguration, the 2025 trigger rates and event topologies agree with the 2024 data taking.

In further validation studies, the HLT part of the trigger is found to misbehave, in such that it artificially decreases the fraction of energy deposited in the ECal with respect to the HCal; since exactly this energy fraction is the main handle for triggering on displaced jets, this degrades the signal-to-background discrimination for the B2FT. The cause of this problem lies in a malfunction of the pile-up removal and EMF recalculation at trigger level, which are currently re-optimised. For the meantime, an alternative HLT trigger without these algorithms is proposed to improve the signal-to-background discrimination for the final Run-3 data taking year, 2026, and also to lower the trigger rates.

It is under investigation whether this kind of L1-correlation trigger is technically feasible in the Run-4 trigger design, in which L1Topo is replaced by the L0 global trigger. Given that the Run-4 trigger design does not work sequential, the correlation of subsequent BCs is not straight-forward anymore, but might still be possible for

simple correlations such as MET (or H_T) from the previous frame with objects in the current one, such as jets, muons and so on.

Similar correlation information could also be used as background handle in other LLP searches, *e.g.* MET followed by correlated muons in the subsequent BC for the Muon Spectrometer Displaced Vertex (MSVtx) search. In particular LLP searches with slow LLPs could be considered, such as heavy ionising particles, potentially even being stopped in the calorimeter.

This concept can also be used in a more general manner by scheduling heavyly pre-scaled L1-correlation triggers as a "tag" for LLP searches to lower their backgrounds. Here, the idea is that the L1-correlation triggers are sensitive to correlations throughout several BCs, but are running in *parasitic* mode, *i.e.* they are *not* supposed to fire a readout, but are rather supposed to serve as a tag in events that are fired by usual LLP triggers. For example, a MSVtx trigger could fire an event, in which case the L1 decision of the heavily pre-scaled correlation trigger could be checked at the analysis level [100].

[100]: Gargan (2025), *L1-Tagged LLP Searches*

Concretely for the B2F analysis, the next steps are the development of a proper trigger simulation, and the definition of a full analysis strategy, including machine-learning based cuts to optimise the signal-to-background discrimination. Likely this will include a jet tagger NN, an event level BDT, and a data driven background estimation using an ABCD plane.

8. Summary

The physics programme at the Large Hadron Collider (LHC) has great potential to uncover some of the open questions beyond the Standard Model of particle physics (BSM). Studying the terra-electronvolt scale with the ATLAS and CMS detectors has, however, not resulted in the observation of a significant BSM excess yet. This naturally gives rise to the question whether the new physics might look different to what would be expected from the SM. Due to limitations of conventional trigger and reconstruction strategies, classical searches would offer only limited capability of probing this kind of new physics. Instead, dedicated searches are needed to address these unconventional BSM signatures, such as searches for Long-Lived Particles (LLPs), which might be split over multiple bunch-crossings (BCs) at the LHC in case of small Lorentz boosts.

A novel approach is presented in this thesis that utilises the Level-1 Topological Trigger of ATLAS to search for electrically neutral LLPs with a decay structure over two adjacent BCs. For this, a dedicated trigger — referred to as Back-to-the-Future Trigger (B2FT) — correlates missing transverse momentum in a certain BC coming from a slowly moving LLP with a displaced jet pointing in the same azimuthal direction in the next BC caused by the delayed LLP decay.

The B2FT performance is validated in pp collision data for the 2025 data taking, uncovering a misconfiguration in the L1Topo firmware. Moreover, validation studies using so-called Calorimeter Energy Ratio (CalRatio) triggers identify an issue with the pile-up removal in the B2F HLT chain, leading to the choice of an alternative HLT chain to improve the signal-to-background discrimination of the B2FT for the 2026 data taking.

Also the simulation of the detector response is quite non-standard for out-of-time decays, since the digitisation in the standard ATLAS reconstruction framework applies the LHC BC structure, making the out-of-time part being lost. Therefore, a dedicated reconstruction framework was developed that separately reconstructs the out-of-time part of the delayed LLP decay.

This framework is validated by comparing reconstructed quantities to the expectations from MC truth, showing perfect closure in observables such as η and φ, but a significant energy miscalibration for displaced jets.

The root of the energy miscalibration is found to be the application of a calibration scheme derived from SM jets, since user-specific calibrations for displaced jets would heavily depend on the LLP

decay position. The distinct topology of displaced jets, such as the absence of radiations outside the jet cone and the untypically low energy fraction in the electromagnetic calorimeter (ECal), explains the non-closure in energy response.

However, the agreement of reconstructed and truth energy is actually of less importance for BSM searches, while the main focus lays on a good data-MC agreement. Therefore, either using EM scale energies (uncalibrated) or SM-calibrated energies *consistently* in data and MC might be appropriate; here, it is proposed to use recommended SM-calibrations for topological jets with a conservative factor on the uncertainties.

Out-of-time contributions from the prompt jet to the delayed jet (due to the wide calorimeter pulses) are shown to have negligible impact on the energy response.

A toy model inspired by Inelastic Dipole Dark Matter (IDDM) was developed to benchmark the trigger performance; sensitivity studies are performed based on a preliminary analysis strategy at the Monte Carlo (MC) truth-level.

Trigger efficiencies and signal yields are derived, where the former peak below the percent level and the latter lie in the range of $\mathcal{O}(10 - 100)$ assuming the Run-2 luminosity. The main reason for the small efficiencies is found to be the rather low fraction of slow LLPs, caused by an increase of the boost by initial-state radiation (ISR). Moreover, the ISR also causes the second dominant bottleneck, namely the shadowing of events, where hard ISR fires a conventional trigger in the first BC, making the subsequent BC lay in the L1 deadtime.

A data-driven background projection is presented based on a previous search for LLPs in the HCal of the ATLAS detector, resulting in $\mathcal{O}(1600)$ background events, with an uncertainty of 5%.

Given the expected signal yields and background rates, upper limits are obtained on the production cross-section of hypothetical BSM extensions with comparable kinematics to the IDDM. The limits are rather weak, peaking at $\mathcal{O}(0.1\,\text{pb})$, and are beaten by the observed limits of a search for displaced vertices in the ATLAS muon spectrometer by more than one order of magnitude.

To make a fair comparison, a full analysis strategy is needed, including simulation of detector response and readout, as well as cuts to optimise the signal-to-background discrimination — potentially based on machine learning techniques. Also, it is advisable to check other benchmark models with slow boost LLPs, which are less affected by ISR.

Appendix

A. Monte Carlo Truth Framework

We try to mimic the overall workflow as closely as possible, in particular at the trigger level.

1. The hardware-level trigger decision, based on calorimeter energy deposits in the BCs $N-1$ and N.
2. Additionally for BC N, the software-based high-level trigger decision.
3. An off-line event selection, which is flexible, i.e. can still be modified, whereas the triggers in 1. and 2. are already implemented and taking data.

However, given that no trustworthy reconstruction of two subsequent bunch crossings exist, and even if we add such an option into Athena, it must be validated using realistic expectations at the truth-level. Also, the output of a full simulation of the detector response with Athena typically does not contain objects at the trigger level, but only objects at the off-line level.

For both reasons, it is crucial to understand our MC simulations thoroughly at the truth-level, which is implemented standalone[35], as explained in the following sections.

35: Actually, two different frameworks were used for this, validating each other, showing closure.

Still it is noticeable that some parts such as the full shower development, can only be *approximated* at the truth-level, and full simulation would be required to simulate the full shower.

A.1. Particle Propagation, Trigger Tower Assignment, and Readout-Time

The hardware-based trigger exploits energy deposits in the calorimeters (ECAL and HCAL). To make a trigger decision within the available time, the hardware-based trigger uses a finite granularity. The calorimeter cells are summed to so-called trigger towers (TTs), which are typically[36] (0.1×0.1) in $(\eta \times \varphi)$, with only two layers (ECAL and HCAL) per TT[37].

36: The TT size in η is slightly larger in the forward region, $|\eta| > 2.5$.

37: For simplicity, we use the legacy TT layout rather the more advanced super-cells which have more layers and a finer granularity.

All detector-stable truth particles that are visible in the calorimeters[38], and that are produced either in the trackers, ECAL, or HCAL, are assigned to the corresponding TT based on their (η, φ) coordinates *in the calorimeter*; i.e. **(i)** particles produced *inside the tracker*, at $(\vec{r}_0, t_0)$ with $\rho_0 = \sqrt{x_0^2 + y_0^2} < \rho_{\text{Calo}} \approx 1.4\,\text{m}$ and $|z_0| < z_{\text{Calo}} \approx 3.7\,\text{m}$, need to be extrapolated to their calorimeter entry. For this, we use a simplified ATLAS calorimeter geometry

38: i.e. all particles except neutrinos, muons, and the BSM particles.

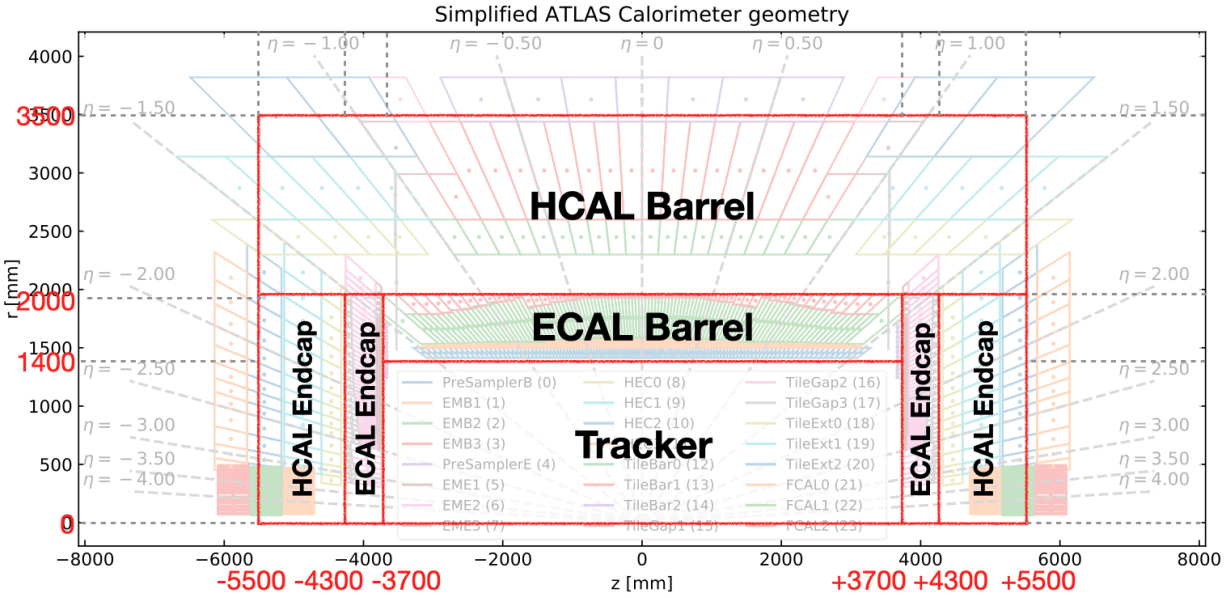

Figure A.1.: Simplified ATLAS geometry. Simplified ATLAS geometry in the (ρ, z) plane. Figure modified from [101].

as shown in Figure A.1, and approximate the calorimeter entry linearly as visualised in Figure A.2:

$$\vec{r}(t) = \vec{r}_0 + \vec{v} \cdot t \quad \text{with } \vec{v} = \vec{p}/E, \tag{A.1}$$

with t referring to the time of travel starting at $t = 0$ from $\vec{r}_0$. The calorimeter entry is obtained by solving $\rho(t) = \sqrt{x^2(t) + y^2(t)} = \rho_{\text{Calo}}$ and $|z(t)| = z_{\text{Calo}}$. These equations could give up to 3 solutions, from which the lowest positive real solution is the travel time (=: t_1) from the production vertex, $\vec{r}_0$, to the calorimeter entry, $\vec{r}_{\text{entry}} = \vec{r}(t_1)$. Subsequently, the η and φ coordinates could be calculated by

$$\eta = 1/2 \times \log\left(\frac{|\vec{r}_{\text{entry}}| + z_{\text{entry}}}{|\vec{r}_{\text{entry}}| - z_{\text{entry}}}\right) \quad \text{and} \quad \varphi = \text{atan2}(y_{\text{entry}}, x_{\text{entry}}). \tag{A.2}$$

For **(ii)** particles produced *within the ECAL or HCAL*, no extrapolation is needed, and the (η, φ) coordinates can directly be obtained by substituting the $\vec{r}_{\text{entry}}$ coordinates in (A.2) by the production vertex position, $\vec{r}_0$.

Particles produced within the trackers or the ECAL ($\rho_0 < 2\,\text{m}$ and $|z_0| < 4.3\,\text{m}$) are assigned to the corresponding ECAL TTs, whereas particles produced in the HCAL are assigned to the corresponding HCAL TTs. Particles produced outside the calorimeters ($\rho_0 > 3.5\,\text{m}$ or $|z_0| > 5.5\,\text{m}$) are not assigned to any TT[39].

The hardware based calorimeter trigger uses typically energy deposits within a certain BC N, but in our case the energy deposits of the previous BC $N-1$ are additionally considered for the trigger

39: That is, spherical particle decays outside the active detector volume with partially backwards travelling decay products are *not* considered for simplicity.

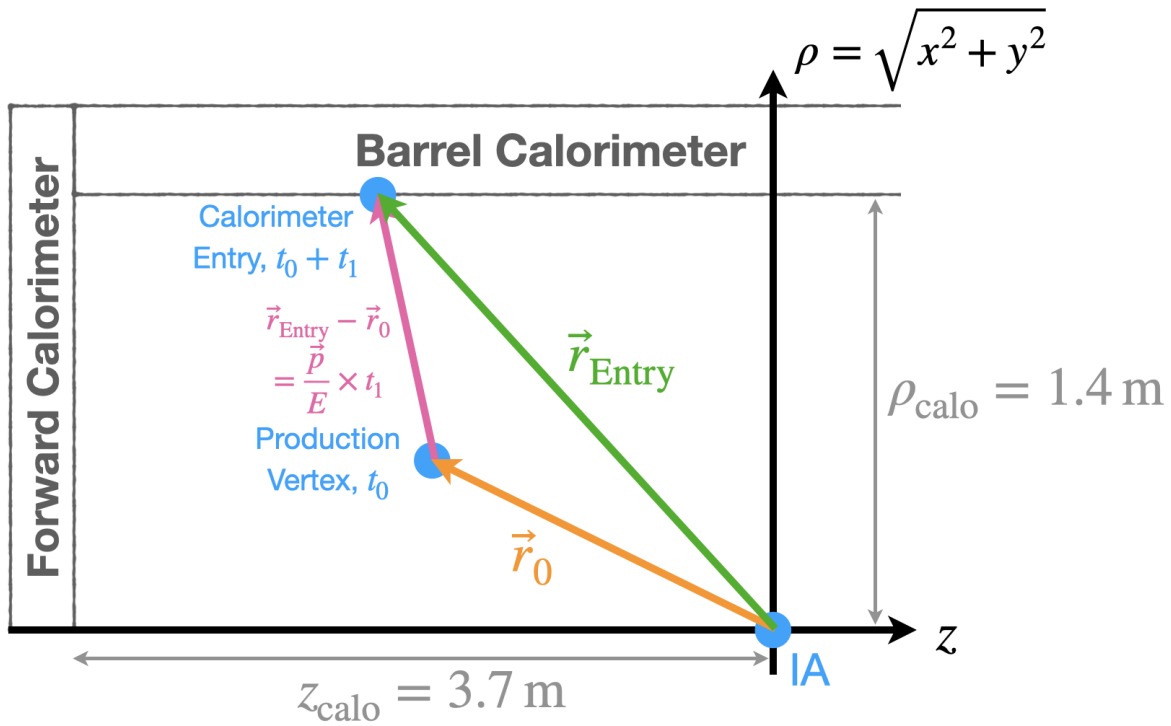

Figure A.2.: Propagation of truth particles. Visualisation of the extrapolation process in the $\rho \times z$ plane. In this example, a particle is produced displaced at $(\vec{r}_0, t_0)$ within the tracker, with a momentum $\vec{p}$. After a travel distance $|\vec{r}_{\text{Entry}} - \vec{r}_0|$ and a travel time t_1, it enters the calorimeter at $\vec{r}_{\text{Entry}}$ at the time $t_0 + t_1$. The readout time at the calorimeter entry is given by $t_0 + t_1 - \frac{|\vec{r}_{\text{Entry}}|}{c}$. The (η, φ) coordinates are defined by $\vec{r}_{\text{Entry}}$ rather than $\vec{p}$ which allows to assign the particle to the corresponding TT.

decision. This allows us to study correlations between two time windows, namely −3 to 15 ns (the sensitive time window of $N-1$), and 22 to 40 ns (the sensitive time window of N), where $t = 0$ is defined at the pp collision in $N-1$. Here, the quoted times refer to the *synchronised readout-time*, which is the actual time, t, corrected by the "time-of-flight", tof, of a prompt jet from the interaction point $\vec{0}$ to the calorimeter entry, or position in the calorimeter, respectively: $t_{\text{readout}} = t - \text{tof}$.

All detector-stable truth particles (which are visible in the calorimeters) are assigned to the corresponding time window **(i)** based on the extrapolated calorimeter entry $(\vec{r}_{\text{entry}}, t_{\text{entry}})$ if they are produced in the trackers, and **(ii)** based on their production vertex $(\vec{r}_0, t_0)$ if they are produced in the ECAL or HCAL. The readout time is then approximated by

$$(i) \quad t_{\text{readout}} = (t_0 + t_1) - \frac{|\vec{r}_{\text{entry}}|}{c}, \tag{A.3}$$

$$(ii) \quad t_{\text{readout}} = t_0 - \frac{|\vec{r}_0|}{c}. \tag{A.4}$$

The readout time of particles produced outside the calorimeters are set to infinity. In all cases, t_{readout} is used to assign the truth

particles to the corresponding time windows (either $N-1$, N, or none of both).

A.2. Trigger Strategy

As detailed in Chapter 5, the hardware trigger is searching for a small radius jet ($R = 0.4$) in BC N exceeding an E_{T} threshold of 40 GeV (the definition of a jet in the context of the hardware trigger is given below). If BC N contains such a jet (in fact, the six most energetic jets in $|\eta| < 3.2$ are considered), the hardware trigger investigates whether BC $N-1$ did contain missing transverse energy (MET definition is also provided below) correlated in φ (i.e. $\Delta\varphi$ between the jet and MET is smaller than 1.0).

- The MET object in $N-1$ must exceed the E_{T} threshold of 40 GeV, but be relatively soft (lower than ~ 100 GeV) — otherwise a MET trigger would fire in $N-1$.
- At least one jet is required within $|\eta| < 3.2$ in N, with a threshold of 40 GeV.
- The $\Delta\varphi$ separations between the MET object and any of the six most energetic jets are calculated. At least one jet with $\Delta\varphi(\mathrm{MET}, \mathrm{jet}) < 1.0$ (and E_{T}, η as defined above) is required. The most energetic jet with $\varphi < 1.0$ is denoted leading jet.

If these conditions are fulfilled, the BC N is further processed by the software-based high-level trigger as described in Section A.4, and potentially readout for off-line analysis.

A.3. Level1 Jets and MET

Small R Jets

The jet algorithm in the hardware-based calorimeter trigger contains the following steps:

1. Only consider TTs that exceed a so-called noise cut (basically an energy threshold to exclude noise). TTs below the noise-cuts are assigned an energy of zero. These noise cuts are of the order of few GeV, they are η dependent, and in general distinct for the ECAL and HCAL TTs.
2. Search for local energy maxima. For each TT in $\eta \times \varphi$, the energy of the corresponding ECAL and HCAL TTs are summed; subsequently, each TT is considered as central seed of a 5×5 search window (5×5 TTs correspond to roughly 0.5×0.5 in $\eta \times \varphi$). The energies of the $5 \times 5 = 25$ TTs within each search window are summed, and thereafter compared (in a

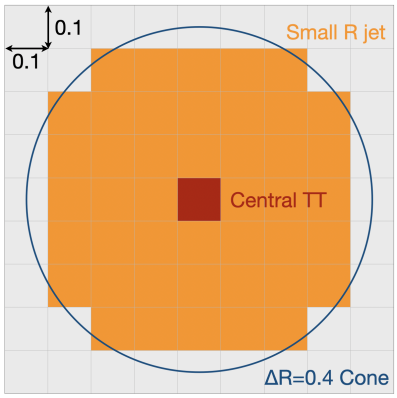

Figure A.3.: Visualisation of the shape of a small R jet in the $\eta \times \varphi$ plane.

so-called sliding window algorithm) to the 24 neighbouring search windows (those search windows with central seed in any of the 24 neighbouring TTs of the original search window's central TT). A local maximum is found when the energy in the original search window is larger than in the 24 neighbouring search windows.

3. Sum energy of neighbouring TTs in a 9×9 window[40] around the central seed to form a $R = 0.4$ cone. Again, the sums of the ECAL and HCAL layer are taken here. Eventually, the jet E_T is estimated based on the η of the central seed, and the summed energy: $E_T = E/\cosh \eta$.

40: In order to mimic an $R = 0.4$ circular shape, the edges of the 9×9 window are *not* considered, and also the corners of the remaining 8×8 window are not considered in the sum. The pattern is depicted in Figure A.3.

These steps have been implemented in a simplified version in the MC truth framework; Based on equations (A.1) to (A.4), all particles in the readout time-window N within the ECAL and HCAL TTs are considered in the jet algorithm:

1. Since we do not simulate pile-up etc. we do not expect noise, and can neglect the noisecuts.
2. After summing the ECAL and HCAL layers, we loop over all 66×64 TTs in $\eta \times \varphi$ and compare their energy, E_{central}, to the 80 neighbouring TTs in a 9×9 TT window. Note that the sliding window algorithm has been simplified, by just comparing the central TT instead of the complete search windows. If the energy in the central TT is greater than the energies in all neighbouring TTs, the central TT is stored as a local maximum.
3. Looping over all local maxima, the jet energies are computed by summing the corresponding TTs according to the circular shape in the 9×9 window. The jet η is taken from the central TT, and the transverse energy of the jet is calculated as $E_T = E/\cosh \eta$. The six most energetic jets in $|\eta| < 3.2$ are stored.

MET

The missing transverse energy at (hardware) trigger level is calculated by summing up the TT energies (weighted by $\sin\varphi$ and $\cos\varphi$) and calculating the energy balance (and computing the transverse part of it). In order to exclude noise, the individual TTs need to exceed a so-called noise-cut (basically an energy threshold), to be considered in the energy balance. TTs below the noise-cuts are assigned an energy of zero. These noise cuts are of the order of few GeV, they are η dependent, and in general distinct for the ECAL and HCAL TTs.

Similar to the jet algorithm in the truth framework, we can skip applying the noise cuts, and directly sum the energies of all particles for each TT, e.g. for TT J:

$$E_J^{\{x,y\}} = E_J/\cosh\eta_J \times \{\cos\varphi_J, \sin\varphi_J\}, \quad \text{with} \quad E_J = \sum_{i=0}^{N} E_i, \tag{A.5}$$

where i runs over the corresponding truth particles in TT J, and (η_J, φ_J) is the center of the TT J. The energy balance is computed by summing over all TTs:

$$E_{\{x,y\}} = \sum_{J=0}^{2\times66\times64} E_J^{\{x,y\}} \tag{A.6}$$

where J is running over the 2 (ECAL and HCAL) $\times$ 66 (in η) $\times$ 64 (in φ) TTs. Subsequently, MET and φ are calculated according to

$$\not{E}_{\mathrm{T}} = \sqrt{(-E_x)^2 + (-E_y)^2} \quad \text{and} \quad \varphi = \mathrm{atan2}(-E_y, -E_x). \tag{A.7}$$

A.4. Anti-k_t Jets

After the hardware based trigger has found a correlation between BCs $N-1$ and N, BC N is further processed by the software-based high-level trigger. The cutflow at this stage is pretty similar to the CalRatio analysis, i.e. an isolated jet with a low fraction of energy deposit in the ECAL is required. In particular, at least one jet with

- $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ and $|\eta| < 2.5$,
- EMF $= \frac{E_{\mathrm{EM}}}{E_{\mathrm{EM}}+E_{\mathrm{HAD}}} < 0.06$, and
- ΔR(jet,tracks) > 0.2 for (prompt) tracks with $p_{\mathrm{T}} > 2\,\mathrm{GeV}$

is required. At the software-based high-level trigger, jets are close to what we know from off-line jets, in our case $R = 0.4$ jets

clustered with the anti-k_t algorithm (without using particle flow algorithms).

All detector stable truth particles within $|\eta| < 2.5$ that are visible in the calorimeter (i.e. skipping muons, neutrinos, and the BSM particles), having a production vertex in the tracker, ECAL, or HCAL, and a calorimeter readout-time in the time window N are clustered with the fastJet framework using the anti-k_t algorithm with a radius of $R = 0.4$ using the Energy Recombination scheme. Jets with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ are stored as "off-line jets". Jets that pass additionally the following two criteria are denoted "HLT jets":

1. We approximate the EMF cut ($E_{EM} < 6.4\% E_{HAD}$), by checking if the jet is originating at most one interaction length, λ_{int}, away from the HCAL boundary. The jet origin is defined by looping over all jet constituents and sum their η and φ coordinates weighted by the constituent's energy. From Figure A.4, we know the material budget (in units of hadronic interaction length) at the entry (λ_{entry}) and at the exit (λ_{exit}) of the ECAL as a function of $|\eta|$. We can use this to estimate the "remaining" interaction length, $\lambda_{\mathrm{remaining}}$, which accounts for the material between the jet origin and the ECAL exit:

$$\frac{\lambda_{\mathrm{exit}} - \lambda_{\mathrm{entry}} - \lambda_{\mathrm{remaining}}}{\lambda_{\mathrm{exit}} - \lambda_{\mathrm{entry}}} = \frac{d - d_{\mathrm{remaining}}}{d}, \tag{A.8}$$

$$\Rightarrow \lambda_{\mathrm{remaining}} = (\lambda_{\mathrm{exit}} - \lambda_{\mathrm{entry}}) \times \frac{d_{\mathrm{remaining}}}{d}, \tag{A.9}$$

as visualised in Figure A.5a; here, d is the thickness of the ECAL, and $d_{\mathrm{remaining}}$ is the distance from the production vertex to the calorimeter exit:

$$d = \begin{cases} \dfrac{\rho_{\mathrm{exit}} - \rho_{\mathrm{entry}}}{\sin\theta} & |\eta| < 1.37 \\ \dfrac{z_{\mathrm{exit}} - z_{\mathrm{entry}}}{\cos\theta} & |\eta| > 1.70 \\ (1-\alpha)\left(\dfrac{\rho_{\mathrm{exit}} - \rho_{\mathrm{entry}}}{\sin\theta}\right) + \alpha\left(\dfrac{z_{\mathrm{exit}} - z_{\mathrm{entry}}}{\cos\theta}\right) & \text{else} \end{cases} \tag{A.10}$$

$$\text{with } \alpha = \frac{|\eta| - 1.37}{1.70 - 1.37}, \tag{A.11}$$

what is illustrated in Figure A.5b. Here, the first case corresponds to the barrel region, the second case to the end-cap region, and the third case to the transition region, where the jet travels through the barrel and the end-cap. A similar expression for $d_{\mathrm{remaining}}$ is obtained by substituting $(\rho_{\mathrm{entry}}, z_{\mathrm{entry}})$ by the production vertex coordinates (ρ_0, z_0). This allows to calculate the remaining interaction length (which is required

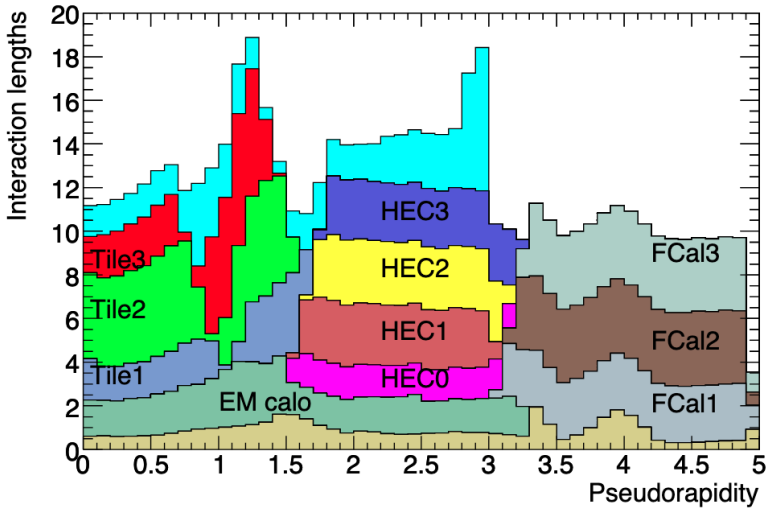

Figure A.4.: Material budget of the ATLAS detector in units of hadronic interaction length. Figure taken from [39].

to be less than 1):

$$\lambda_{\text{remaining}} = (\lambda_{\text{exit}} - \lambda_{\text{entry}}) \times \begin{cases} \dfrac{\rho_{\text{exit}} - \rho_0}{\rho_{\text{exit}} - \rho_{\text{entry}}} \\ \dfrac{z_{\text{exit}} - z_0}{z_{\text{exit}} - z_{\text{entry}}} \\ \dfrac{(1-\alpha)\left(\frac{\rho_{\text{exit}}-\rho_0}{\sin\theta}\right) + \alpha\left(\frac{z_{\text{exit}}-z_0}{\cos\theta}\right)}{(1-\alpha)\left(\frac{\rho_{\text{exit}}-\rho_{\text{entry}}}{\sin\theta}\right) + \alpha\left(\frac{z_{\text{exit}}-z_{\text{entry}}}{\cos\theta}\right)} \end{cases} \tag{A.12}$$

2. The jet is required to be isolated, i.e. there must not be tracks with $p_{\text{T}} > 2\,\text{GeV}$ in a $\Delta R = 0.2$ cone around the jet axis. Tracks are defined as charged detector-stable truth particles with a production vertex in the inner detector and a readout time within the time window N; here, the readout time and (η, φ) coordinates are calculated based on the production vertex. Eventually, the isolation cut is applied by requiring $\Delta R > 0.2$ for all tracks with $p_{\text{T}} > 2\,\text{GeV}$.

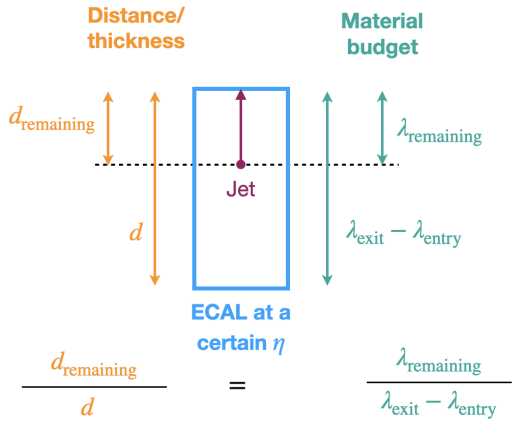

(a) Relation between the geometrical distance and the distance in hadronic interaction length.

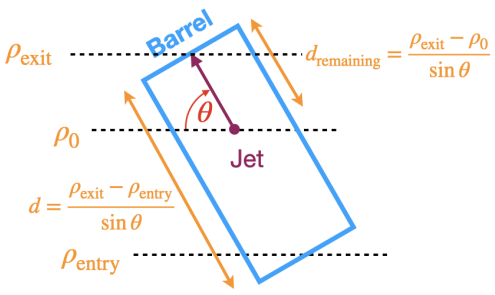

(b) Geometrical distance between jet origin and the calorimeter exit (exemplarily in the barrel).

Figure A.5.: Remaining interaction length in calorimeter. The remaining (hadronic) interaction length from the jet origin to the calorimeter exist is calculated based on the geometrical distance and the corresponding material budget of the ATLAS detector.

B. Data-Driven Background Projection

For illustration, we derive the background projection of the SM multi-jet background:

The CalRatio analysis observes

$$N^0_{\text{CR}} = 40\,743\,867 \tag{B.1}$$

SM multi-jet events after their trigger selection, and after requiring two clean jets. Here, a clean jet refers to a jet with $p_{\text{T}} > 20\,\text{GeV}$ and $|\eta| < 2.5$ that passes the dedicated CalRatio jet cleaning.

To get the according number of background events that do only contain one such clean jet, we calculate the ratio

$$\frac{\sigma(p_{\text{T}}^{j_1} > 40)}{\sigma(p_{\text{T}}^{j_1} > 40 \wedge p_{\text{T}}^{j_2} > 40)}, \tag{B.2}$$

using Pythia di-jet samples, which results in about 5.

That is, eq. B.1 needs to be increased by a factor of 5 to account for the fact that the B2F analysis only requires one jet above 40 GeV:

$$203\,719\,335. \tag{B.3}$$

However, as derived in Section 5.3, this number may be scaled by a factor of $0.029/\pi$ to account for the additional B2F background suppression through the required MET correlation:

$$N^0_{\text{B2F}} = 1\,880\,530 \tag{B.4}$$

Assuming that the ΔR_{min} cut is independent of the trigger differences between B2F and B2F and also independent of the jet multiplicity, we can assume that we will loose the same fraction of events by this cut as CalRatio does. They end up with

$$N^1_{\text{CR}} = 28\,248\,024. \tag{B.5}$$

We can therefore scale to the expected B2F yield at this stage:

$$N^1_{\text{B2F}} = N^0_{\text{B2F}} \times \frac{N^1_{\text{CR}}}{N^0_{\text{CR}}} = 1\,303\,785. \tag{B.6}$$

The low-E_{T} BDT > 0.05 cut is not applied, hence the background

numbers stay as they are:

$$N^2_{\text{B2F}} = N^1_{\text{B2F}} = 1\,303\,785, \tag{B.7}$$

while the CalRatio numbers are (according to Table 6.4) reduced to

$$N^2_{\text{CR}} = 1\,288\,596 \tag{B.8}$$

For the trigger matching, it is again assumed that the same fraction of events will be removed in the B2F analysis. That leaves us with

$$N^3_{\text{B2F}} = N^2_{\text{B2F}} \times \frac{N^3_{\text{CR}}}{N^2_{\text{CR}}} = 1\,152\,386, \tag{B.9}$$

where

$$N^3_{\text{CR}} = 1\,138\,961. \tag{B.10}$$

For the next three cuts (the jet timing, the jet EMF, and the exclusion of the transition region), we can not simply take the same fraction as above, since CalRatio applies these cuts on two signal-like (*i.e.* displaced) jets, whereas the B2F analysis only targets one of them. However, assuming that the two jets are independent from each other, we can approximate that the likelihood for two SM jets to pass the cuts is the square of the probability of one SM jet to pass the jets. Therefore, instead of scaling the B2F yield by the corresponding CalRatio fraction, we can "only" reduce by the square root of their fraction:

$$N^4_{\text{B2F}} = N^3_{\text{B2F}} \times \sqrt{\frac{N^4_{\text{CR}}}{N^3_{\text{CR}}}} = 1\,144\,405, \tag{B.11}$$

with

$$N^4_{\text{CR}} = 1\,123\,239. \tag{B.12}$$

Analogously, we obtain

$$N^5_{\text{B2F}} = 1\,100\,136 \qquad N^5_{\text{CR}} = 1\,038\,019, \tag{B.13}$$

$$N^6_{\text{B2F}} = 1\,067\,204 \qquad N^6_{\text{CR}} = 976\,805. \tag{B.14}$$

The next three cuts ($H^{\text{miss}}_{\text{T}}/H_{\text{T}} < 0.6$ and the $p_{\text{T}} > 80\,\text{GeV}$ cuts for both jets) are skipped. Therefore,

$$N^7_{\text{B2F}} = N^8_{\text{B2F}} = N^9_{\text{B2F}} = 1\,067\,204, \tag{B.15}$$

while the CalRatio rates are substantially reduced to

$$N_{\mathrm{CR}}^{7} = 965\,748, \tag{B.16}$$
$$N_{\mathrm{CR}}^{8} = 315\,530, \tag{B.17}$$
$$N_{\mathrm{CR}}^{9} = \ \ 73\,484. \tag{B.18}$$

Assuming that the two jets contribute evenly to the $\sum_j \log_{10}\left(\frac{E_{\mathrm{H}}}{E_{\mathrm{EM}}}\right)$ variable, the next cut can be disentangled:

$$\sum_j \log_{10}\left(\frac{E_{\mathrm{H}}}{E_{\mathrm{EM}}}\right) > 2 \Leftrightarrow \log_{10}\left(\frac{E_{\mathrm{H}}}{E_{\mathrm{EM}}}\right)_{j_1} > 1 \text{ and } \log_{10}\left(\frac{E_{\mathrm{H}}}{E_{\mathrm{EM}}}\right)_{j_2} > 1. \tag{B.19}$$

This can be used to estimate the probability of a SM jet passing the corresponding cuts:

$$\mathscr{P}_{\mathrm{CR}} = \mathscr{P}\left(\sum_j \log_{10}\left(\frac{E_{\mathrm{H}}}{E_{\mathrm{EM}}}\right) > 2\right) \tag{B.20}$$
$$\approx \mathscr{P}\left(\log_{10}\left(\frac{E_{\mathrm{H}}}{E_{\mathrm{EM}}}\right)_{j_1} > 1 \ \wedge \ \log_{10}\left(\frac{E_{\mathrm{H}}}{E_{\mathrm{EM}}}\right)_{j_2} > 1\right) \tag{B.21}$$
$$\approx \left[\mathscr{P}\left(\log_{10}\left(\frac{E_{\mathrm{H}}}{E_{\mathrm{EM}}}\right)_{j} > 1\right)\right]^2 = \mathscr{P}_{\mathrm{B2F}}^2. \tag{B.22}$$

Here, the square root of the RHS is the background rejection of the B2F cut, while the LHS is the background rejection of the CalRatio cut. That is, the background rates scale with the square root as above (for the jet timing, EMF, and the rejection of the transition region). This leaves us with

$$N_{\mathrm{B2F}}^{10} = N_{\mathrm{B2F}}^{9} \times \sqrt{\frac{N_{\mathrm{CR}}^{10}}{N_{\mathrm{CR}}^{9}}} = 228\,711, \tag{B.23}$$

using

$$N_{\mathrm{CR}}^{10} = 3\,375. \tag{B.24}$$

The cut on the NN signal score *product* is exchanged by a cut on the NN signal score of the leading jet. From Figure 6.17, we read off that cutting at a signal score of 0.95 allows to reduce the background by a factor of 138.89 while the signal is only reduced by a factor of 0.25. Therefore, we end up with

$$N_{\mathrm{B2F}}^{11} = N_{\mathrm{B2F}}^{10}/138.89 = 1\,647. \tag{B.25}$$

That is, the SM multi-jet background is expected to yield 1 647

events. Based on the CalRatio studies, we assume an uncertainty of 5% on this number.

The BIB background can be derived in analogy, ending up with 104 expected BIB events.

Bibliography

[1] Tobias Heintz. *Back-to-the-Future Triggers in L1Calo*. Jan. 2025. URL: https://indico.cern.ch/event/1472233/timetable/#4-back-to-the-future-triggers (cited on page v).

[2] Tobias Heintz. *Probing Delayed Jets with Dedicated Long-Lived-Particle Triggers at ATLAS*. Apr. 2025. URL: https://www.dpg-verhandlungen.de/year/2025/conference/goettingen/part/t/session/44/contribution/1 (cited on page v).

[3] Mark Thomson. *Modern Particle Physics*. Cambridge University Press, 2013 (cited on pages 3, 5).

[4] Peter Schmüser. *Feynman-Graphen und Eichtheorien für Experimentalphysiker*. Springer, 2011 (cited on page 3).

[5] Latham Boyle. *Mathematical formulation of the Standard Model*. https://en.wikipedia.org/wiki/Mathematical_formulation_of_the_Standard_Model. Accessed 15-02-2025. 2015 (cited on page 4).

[6] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to quantum field theory*. Reading, USA: Addison-Wesley, 1995 (cited on page 4).

[7] Steven Weinberg. 'A Model of Leptons'. In: *Phys. Rev. Lett.* 19 (21 Nov. 1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264 (cited on page 5).

[8] Abdus Salam. 'Weak and Electromagnetic Interactions'. In: *Conf. Proc. C* 680519 (1968), pp. 367–377. DOI: 10.1142/9789812795915_0034 (cited on page 5).

[9] Sheldon Glashow. 'Partial-symmetries of weak interactions'. In: *Nuclear Physics* 22.4 (1961), pp. 579–588. DOI: https://doi.org/10.1016/0029-5582(61)90469-2 (cited on page 5).

[10] Peter W. Higgs. 'Broken Symmetries and the Masses of Gauge Bosons'. In: *Phys. Rev. Lett.* 13 (16 Oct. 1964), pp. 508–509. DOI: 10.1103/PhysRevLett.13.508 (cited on page 6).

[11] F. Englert and R. Brout. 'Broken Symmetry and the Mass of Gauge Vector Mesons'. In: *Phys. Rev. Lett.* 13 (9 Aug. 1964), pp. 321–323. DOI: 10.1103/PhysRevLett.13.321 (cited on page 6).

[12] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. 'Global Conservation Laws and Massless Particles'. In: *Phys. Rev. Lett.* 13 (20 Nov. 1964), pp. 585–587. DOI: 10.1103/PhysRevLett.13.585 (cited on page 6).

[13] Janosh Riebesell. *Janosh Riebesell*. Apr. 2022. URL: https://tikz.net/higgs-potential/ (cited on page 7).

[14] S. Y. Lee. 'Higher-Order Corrections to Leptonic Processes and the Renormalization of Weinberg's Theory of Weak Interactions in the Unitary Gauge'. In: *Phys. Rev. D* 6 (6 Sept. 1972), pp. 1701–1717. DOI: 10.1103/PhysRevD.6.1701 (cited on page 7).

[15] J. E. Dodd and B. M. Gripaios. *The ideas of particle physics*. Cambridge University Press, 2020 (cited on page 8).

[16] ATLAS Collaboration. 'Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC'. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020 (cited on page 9).

[17] CMS Collaboration. 'Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC'. In: *Physics Letters B* 716.1 (2012), pp. 30–61. DOI: https://doi.org/10.1016/j.physletb.2012.08.021 (cited on page 9).

[18] J.A. Bagger et al. 'Precision electroweak measurements on the Z resonance'. In: *Physics Reports* 427.5 (2006), pp. 257–454. DOI: https://doi.org/10.1016/j.physrep.2005.12.006 (cited on page 9).

[19] Vincenzo Cirigliano and Michael J. Ramsey-Musolf. 'Low energy probes of physics beyond the standard model'. In: *Progress in Particle and Nuclear Physics* 71 (2013). Fundamental Symmetries in the Era of the LHC, pp. 2–20. DOI: https://doi.org/10.1016/j.ppnp.2013.03.002 (cited on page 9).

[20] Roberto D. Peccei. 'The Strong CP Problem and Axions'. In: *Axions*. Springer Berlin Heidelberg, 2008, pp. 3–17. DOI: 10.1007/978-3-540-73518-2_1 (cited on page 9).

[21] S. Navas et al. 'Review of particle physics'. In: *Phys. Rev. D* 110.3 (2024), p. 030001. DOI: 10.1103/PhysRevD.110.030001 (cited on page 9).

[22] Paul Adrien Maurice Dirac. 'The quantum theory of the emission and absorption of radiation'. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 114.767 (Mar. 1927), pp. 243–265. DOI: 10.1098/rspa.1927.0039 (cited on page 10).

[23] Simon Knapen and Steven Lowette. *A guide to hunting long-lived particles at the LHC*. 2022. URL: https://arxiv.org/abs/2212.03883 (cited on page 11).

[24] Marie-Helene Genest. 'Searching for long-lived BSM particles at the LHC'. In: *Scholarpedia* 17.12 (2022), p. 54697. DOI: 10.4249/scholarpedia.54697 (cited on page 11).

[25] Izaak Neutelings. *Izaak neutelings*. Oct. 2023. URL: https://tikz.net/bsm_longlived/ (cited on page 12).

[26] ATLAS Collaboration. 'Search for neutral long-lived particles in pp collisions at $\sqrt{s} = 13$ TeV that decay into displaced hadronic jets in the ATLAS calorimeter'. In: *Journal of High Energy Physics* 2022.6 (June 2022). DOI: 10.1007/jhep06(2022)005 (cited on pages 11, 33, 36, 38, 47, 50, 55, 60, 61).

[27] ATLAS Collaboration. 'Standalone vertex finding in the ATLAS muon spectrometer'. In: *JINST* 9 (2014), P02001. DOI: 10.1088/1748-0221/9/02/P02001 (cited on page 12).

[28] Henso Abreu et al. 'The FASER detector'. In: *JINST* 19.05 (2024), P05066. DOI: 10.1088/1748-0221/19/05/P05066 (cited on page 12).

[29] C. Ahdida et al. 'The SHiP experiment at the proposed CERN SPS Beam Dump Facility'. In: *Eur. Phys. J. C* 82.5 (2022), p. 486. DOI: 10.1140/epjc/s10052-022-10346-5 (cited on page 12).

[30] Martin Bauer et al. *ANUBIS: Proposal to search for long-lived neutral particles in CERN service shafts*. 2025. URL: https://arxiv.org/abs/1909.13022 (cited on page 12).

[31] CERN. *The Large Hadron Collider*. 2025. URL: https://home.cern/science/accelerators/large-hadron-collider (cited on page 13).

[32] ATLAS Experiment. *CERN accelerating science*. 2025. URL: https://atlas.cern/ (cited on page 13).

[33] CMS Experiment. *CERN accelerating science*. 2025. URL: https://cms.cern/ (cited on page 13).

[34] Tobias Heintz. 'Improvement of the template fit method for the fake photon background estimation at ATLAS'. Bachelorarbeit. Universität Heidelberg, Aug. 2023 (cited on pages 13, 21).

[35] Claudia Richoux. *Particle accelerators*. 2016. URL: https://www.terpconnect.umd.edu/~dsvolpe/TREND/2016/Multimedia/Richoux,%20Claudia/LHC%20&%20UMER.htm (cited on page 13).

[36] Daniel Boussard and Trevor Paul R Linnecar. *The LHC Superconducting RF System*. Tech. rep. Geneva: CERN, 1999 (cited on page 14).

[37] LHC Machine Outreach. *Collisions and Collision Rates*. 2025. URL: https://lhc-machine-outreach.web.cern.ch/collisions.htm (cited on page 14).

[38] ATLAS Collaboration. 'The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3'. In: *Journal of Instrumentation* 19.05 (May 2024), P05063. DOI: 10.1088/1748-0221/19/05/p05063 (cited on pages 14, 15, 17).

[39] ATLAS Collaboration. 'The ATLAS Experiment at the CERN Large Hadron Collider'. In: *Journal of Instrumentation* 3.08 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003 (cited on pages 14, 18, 92).

[40] Izaak Neutelings. *Izaak neutelings*. Feb. 2025. URL: https://tikz.net/axis3d_cms/ (cited on page 16).

[41] Hermann Kolanoski and Norbert Wermes. *Particle Detectors*. Oxford University Press, June 2020 (cited on pages 16, 17).

[42] Georg Viehhauser and T. Weidberg. *Detectors in particle physics: A modern introduction*. CRC PRESS, 2024 (cited on pages 17, 19).

[43] Wolfgang Lucha and Meinhard Regler. *Elementarteilchenphysik in Theorie und experiment*. Sappl, 1997 (cited on page 17).

[44] B. Aubert et al. 'Performance of a liquid argon electromagnetic calorimeter with an "accordion" geometry'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 309.3 (1991), pp. 438–449. DOI: https://doi.org/10.1016/0168-9002(91)90247-N (cited on page 18).

[45] Peter Loch. *Signal reconstruction in the ATLAS calorimeter: from particles to topological cell clusters*. 2024. URL: https://indico.cern.ch/event/1366444/ (cited on pages 18, 21).

[46] ATLAS Collaboration and Davidek, Tomas. *TileTimingCalibration*. 2021. URL: https://twiki.cern.ch/twiki/bin/viewauth/Atlas/TileTimingCalibration (cited on page 18).

[47] K. Anderson et al. 'Design of the front-end analog electronics for the ATLAS tile calorimeter'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 551.2 (2005), pp. 469–476. DOI: https://doi.org/10.1016/j.nima.2005.06.048 (cited on page 18).

[48] Danijela Bogavac. *Calibration of the ATLAS Tile Calorimeter*. Tech. rep. Geneva: CERN, 2024. DOI: 10.22323/1.476.0946 (cited on page 19).

[49] Imai Jen-La Plante and Maja Tylmad. 'Pulse shapes for signal reconstruction in the ATLAS Tile Calorimeter'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 617.1 (2010). 11th Pisa Meeting on Advanced Detectors, pp. 96–98. DOI: https://doi.org/10.1016/j.nima.2009.09.114 (cited on page 19).

[50] J. Stirling. *Hepforge*. URL: https://mstwpdf.hepforge.org/plots/plots.html (cited on page 19).

[51] ATLAS Collaboration. *Athena*. Version 22.0.1. Apr. 2019. DOI: 10.5281/zenodo.2641997 (cited on pages 20, 29).

[52] Sebastian Weber. 'Measurement of Vector Boson plus Jet Production Cross Sections and Dark Matter Interpretation with the ATLAS Detector'. PhD thesis. Universität Heidelberg, Jan. 2023 (cited on page 20).

[53] Sascha Mehlhase. *ATLAS detector slice (and particle visualisations)*. 2021. URL: https://cds.cern.ch/record/2770815 (cited on page 21).

[54] Matteo Cacciari, Gavin P Salam, and Gregory Soyez. 'The anti-ktjet clustering algorithm'. In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), pp. 063–063. DOI: 10.1088/1126-6708/2008/04/063 (cited on pages 20, 23).

[55] ATLAS Collaboration. 'Jet reconstruction and performance using particle flow with the ATLAS Detector'. In: *The European Physical Journal C* 77.7 (2017), p. 466. DOI: 10.1140/epjc/s10052-017-5031-2 (cited on page 21).

[56] ATLAS Collaboration. 'Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV'. In: *The European Physical Journal C* 3 (), p. 2304. DOI: 10.1140/epjc/s10052-013-2304-2 (cited on page 21).

[57] ATLAS Collaboration. 'Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1'. In: *The European Physical Journal C* 77.7 (July 2017). DOI: 10.1140/epjc/s10052-017-5004-5 (cited on page 21).

[58] ATLAS Collaboration. 'Improving topological cluster reconstruction using calorimeter cell timing in ATLAS'. In: *The European Physical Journal C* 84.5 (2024), p. 455 (cited on page 22).

[59] Matteo Cacciari. *Standard Model: Jet Physics*. 2018. URL: https://indico.mitp.uni-mainz.de/event/167/timetable/?view=standard (cited on page 23).

[60] Andy Buckley, Chris White, and Martin White. *Practical collider physics*. IOP Publishing, 2021 (cited on pages 23, 25, 27).

[61] William Balunas. *Making the Most of the Mess: Jets and MET at ATLAS*. July 2025. URL: https://indico.cern.ch/event/1552244/ (cited on page 23).

[62] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. 'FastJet User Manual'. In: *Eur. Phys. J. C* 72 (2012), p. 1896. DOI: 10.1140/epjc/s10052-012-1896-2 (cited on page 23).

[63] CMS Collaboration. *CERN accelerating science*. July 2025. URL: https://cms.cern/news/higgs-bosons-shadowy-cousin-cms-search-dark-higgs (cited on page 24).

[64] Alex Mitov. *QCD (in the LHC precision era)*. Nov. 2024. URL: https://www.hep.phy.cam.ac.uk/~amitov/QCD-2020-21-Lectures-1_5.pdf (cited on page 25).

[65] B. Webber. 'Parton shower Monte Carlo event generators'. In: *Scholarpedia* 6.12 (2011). revision #128236, p. 10662. DOI: 10.4249/scholarpedia.10662 (cited on pages 25, 26).

[66] G Peter Lepage. 'A new algorithm for adaptive multidimensional integration'. In: *Journal of Computational Physics* 27.2 (1978), pp. 192–203. DOI: https://doi.org/10.1016/0021-9991(78)90004-9 (cited on page 27).

[67] 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'. In: *Journal of High Energy Physics* 2014.7 (July 2014). DOI: 10.1007/jhep07(2014)079 (cited on page 27).

[68] Christian Bierlich et al. *A comprehensive guide to the physics and usage of PYTHIA 8.3*. 2022. URL: https://arxiv.org/abs/2203.11601 (cited on page 28).

[69] Nils Lavesson and Leif Lönnblad. 'Merging parton showers and matrix elements—back to basics'. In: *Journal of High Energy Physics* 2008.04 (2008), p. 085. DOI: 10.1088/1126-6708/2008/04/085 (cited on page 29).

[70] Leif Lönnblad. 'Correcting the Colour-Dipole Cascade Model with Fixed Order Matrix Elements'. In: *Journal of High Energy Physics* 2002.05 (May 2002), pp. 046–046. DOI: 10.1088/1126-6708/2002/05/046 (cited on page 29).

[71] ATLAS Collaboration. 'The ATLAS Simulation Infrastructure'. In: *The European Physical Journal C* 70.3 (Sept. 2010), pp. 823–874. DOI: 10.1140/epjc/s10052-010-1429-9 (cited on pages 29–31).

[72] J. Allison et al. 'Recent developments in Geant4'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 835 (2016), pp. 186–225. DOI: https://doi.org/10.1016/j.nima.2016.06.125 (cited on page 29).

[73] J. Allison et al. 'Geant4 developments and applications'. In: *IEEE Transactions on Nuclear Science* 53.1 (2006), pp. 270–278. DOI: 10.1109/TNS.2006.869826 (cited on page 29).

[74] S. Agostinelli et al. 'Geant4—a simulation toolkit'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. DOI: https://doi.org/10.1016/S0168-9002(03)01368-8 (cited on page 29).

[75] Zach Marshall and Andy Buckley. *Internal Communication*. Jan. 2025. URL: https://atlas-talk.web.cern.ch/t/issue-with-displaced-delayed-vertices-in-madgraph-long-lived-particle-simulation/45959 (cited on page 32).

[76] John Chapman. *Private Communication*. 2025 (cited on page 32).

[77] ATLAS Collaboration. 'Beam-induced backgrounds measured in the ATLAS detector during local gas injection into the LHC beam vacuum'. In: *JINST* 19 (2024). 61 pages in total, author list starting page 44, 13 figures, 9 tables, JINST 19 (2024) P06014. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PAPERS/DAPR-2021-02, p. 0614. DOI: 10.1088/1748-0221/19/06/P06014 (cited on page 39).

[78] ATLAS Collaboration. 'Performance of the missing transverse momentum triggers for the ATLAS detector during Run-2 data taking'. In: *Journal of High Energy Physics* 2020.8 (Aug. 2020). DOI: 10.1007/jhep08(2020)080 (cited on page 40).

[79] David Smith and Neal Weiner. 'Inelastic dark matter'. In: *Phys. Rev. D* 64 (4 July 2001), p. 043502. DOI: 10.1103/PhysRevD.64.043502 (cited on page 42).

[80] Krzysztof Jodłowski. 'Looking forward to inelastic DM with electromagnetic form factors at FASER and beam dump experiments'. In: *Physical Review D* 108.11 (Dec. 2023). DOI: 10.1103/physrevd.108.115025 (cited on page 42).

[81] Keith R. Dienes et al. 'Extending the discovery potential for inelastic-dipole dark matter with FASER'. In: *Physical Review D* 107.11 (June 2023). DOI: 10.1103/physrevd.107.115006 (cited on page 42).

[82] Andre Lessa and Jose Zurita. *Back to the Future Trigger*. Private Communication. 2024 (cited on pages 42–44, 59, 62).

[83] Andre Lessa. *back2futureTrigger*. 2025. URL: https://github.com/andlessa/back2futureTrigger (cited on pages 47, 49).

[84] Jackson Carl Burzynski. *Lifetime reweighting*. 2025. URL: https://gitlab.cern.ch/atlas-phys/exot/docs/lup/-/blob/master/docs/useful/lifetimerw.md?ref_type=heads (cited on page 48).

[85] ATLAS Collaboration. 'Luminosity determination in pp collisions at $\sqrt{s}$ = 13 TeV using the ATLAS detector at the LHC'. In: *The European Physical Journal C* 83.10 (Oct. 2023). DOI: 10.1140/epjc/s10052-023-11747-w (cited on page 57).

[86] Will Buttinger. *Background estimation with the ABCD method featuring the TROOFIT toolkit*. Oct. 2018. URL: https://indico.cern.ch/event/1122790/contributions/4713580/attachments/2381493/4270369/ABCDGuide_draft18Oct18.pdf (cited on page 58).

[87] A L Read. 'Modified frequentist analysis of search results (the CL_s method)'. In: (2000). DOI: 10.5170/CERN-2000-005.81 (cited on page 60).

[88] Lukas Heinrich, Matthew Feickert, and Giordon Stark. *pyhf: v0.7.6*. Version 0.7.6. https://github.com/scikit-hep/pyhf/releases/tag/v0.7.6. DOI: 10.5281/zenodo.1169739 (cited on page 61).

[89] Lukas Heinrich et al. 'pyhf: pure-Python implementation of HistFactory statistical models'. In: *Journal of Open Source Software* 6.58 (2021), p. 2823. DOI: 10.21105/joss.02823 (cited on page 61).

[90] Glen Cowan et al. 'Asymptotic formulae for likelihood-based tests of new physics'. In: *The European Physical Journal C* 71.2 (Feb. 2011). DOI: 10.1140/epjc/s10052-011-1554-0 (cited on page 61).

[91] ATLAS Collaboration. 'Search for events with a pair of displaced vertices from long-lived neutral particles decaying into hadronic jets in the ATLAS muon spectrometer in pp collisions at $\sqrt{s}$ = 13 TeV'. In: *Phys. Rev. D* 106 (3 Aug. 2022), p. 032005. DOI: 10.1103/PhysRevD.106.032005 (cited on page 63).

[92] Ralf Gugel. *Private Communication*. 2025 (cited on pages 64, 65).

[93] Benjamin Kerridge. *Private Communication*. 2025 (cited on page 65).

[94] Rachel Christine Rosten. *Exotics displaced hadronic jets informal analysis meeting*. 2016. URL: https://indico.cern.ch/event/507126/#2-calratio-trigger-and-pileup (cited on page 67).

[95] Tatiana Ovsiannikova and Emma Torro Pastor. *Private Communication*. 2025 (cited on page 68).

[96] Jeff Dandoy and Pierre Antoine Delsart. *Internal Communication*. Nov. 2025. URL: https://indico.cern.ch/event/1611130/#4-jet-energy-scale-for-delayed (cited on pages 69, 74, 75).

[97] Jackson Carl Burzynski. *Private Communication*. 2025 (cited on page 72).

[98] Falk Bartels. *Private Communication*. 2025 (cited on page 73).

[99] Pavol Strizenec et al. *Private Communication*. Nov. 2025 (cited on page 76).

[100] Jack Gargan. *L1-Tagged LLP Searches*. Dec. 2025. URL: https://indico.cern.ch/event/1599523/timetable/#28-l1-tagged-llp-searches (cited on page 78).

[101] Gessinger, Paul et al. 'The Acts project: track reconstruction software for HL-LHC and beyond'. In: *EPJ Web Conf.* 245 (2020), p. 10003. DOI: 10.1051/epjconf/202024510003 (cited on page 86).

List of Terms

B

B2F Back-to-the-Future. v, vi, xxii, 58–63, 65, 67, 68, 77–79, 95–97
B2FT Back-to-the-Future Trigger. xxi, 35, 38, 41, 46, 49, 58, 60, 65–68, 77, 79
BC bunch-crossing. i, vi, xx, 14, 18, 31–33, 35–39, 41, 49–58, 60, 62, 64–66, 68, 76–80
BDT boosted decision tree. 59, 78
BG-CR background-enriched control-region. 65
BIB beam-induced background. 31, 36, 39, 60, 61, 98
BSM Beyond the Standard Model. v, xxi, 1–3, 9–11, 14, 19, 24, 25, 32, 35, 42–45, 52, 60, 62, 63, 65, 68, 72, 75, 79, 80

C

CalRatio Calorimeter Energy Ratio. vi, xxi, xxii, 11, 33, 36, 38, 50, 55, 58–61, 63, 64, 66–68, 79, 95–98
CERN European Organization for Nuclear Research. 13
CTP central trigger processor. 20, 38

E

ECal electromagnetic calorimeter. 16, 38, 49, 52, 54, 59, 66, 67, 69–71, 74, 75, 77, 80
EMPFlow particle-flow jets clustered at the EM scale. 71–74, 76
EMTopo topological jets clustered at the EM scale. 71–76
EW electroweak. xx, 4–6, 8, 43, 44

H

HCal hadronic calorimeter. 19, 36–38, 49, 51, 52, 54–56, 66–68, 70, 71, 73–77, 80
HLT Higher-Level Trigger. vi, xxi, 20, 31, 33, 38, 40, 49, 51–55, 58, 60, 65–68, 77, 79

I

IDDM Inelastic Dipole Dark Matter. i, xx, xxi, 2, 41–46, 55, 59–65, 68, 80
ISR initial-state radiation. i, 28, 29, 36, 52, 54, 56, 57, 64, 70, 80

J

jFex jet feature extractor. 38

L

L1 Level-1. v, vi, xxi, 19, 20, 31, 33, 38–41, 49–53, 57, 58, 60, 65, 67, 77, 78, 80
L1A L1-Accept. 20
L1Topo Level-1 Topological Trigger. 19, 38, 50, 65, 66, 77, 79
LAr Liquid Argon. xx, 16–19
LHC Large Hadron Collider. i, xx, 1, 2, 8, 9, 12–14, 19, 25, 28, 31, 35, 39, 42, 62, 79
LLP Long-Lived Particle. i, v, xx, xxi, 2, 3, 9–12, 22, 25, 30, 32, 33, 35–38, 41, 46–48, 50–52, 55, 56, 63, 64, 66–70, 72–75, 78–80
LO leading order. 28, 29, 43

M

MC Monte Carlo. i, v, vi, xx, xxi, 23, 26, 27, 41, 48–50, 54, 58, 59, 64, 67, 68, 70, 72, 73, 75, 79, 80
MSVtx Muon Spectrometer Displaced Vertex. vi, 12, 62, 63, 78

N
NLO next-to-leading order. 28, 29
NN neural network. xxi, 59, 61, 78, 97

P
PMT photomultiplier. 18, 19
PS parton shower. 25, 26, 28, 29

Q
QCD quantum chromodynamics. 5, 7–9, 13, 14, 25, 28, 35, 44, 62, 74

S
SM Standard Model. xx, 1–4, 6, 9–11, 13, 14, 19, 32, 38–45, 54, 58, 60, 61, 63, 65, 69, 73–75, 79, 80, 95–97

T
TDAQ Trigger and Data Acquisition. 19, 20

List of Figures

List of Tables