

# A Guide to Hunting Long-Lived Particles at the LHC

Simon Knapen<sup>1,2</sup> and Steven Lowette<sup>3</sup>

<sup>1</sup>Lawrence Berkeley National Laboratory, Berkeley, California, USA; email: smknapen@lbl.gov

<sup>2</sup>Berkeley Center for Theoretical Physics, University of California, Berkeley, California, USA

<sup>3</sup>Inter-university Institute for High Energies (IIHE), Vrije Universiteit Brussel, Brussels, Belgium; email: steven.lowette@vub.be

ANNUAL  
REVIEWS **CONNECT**

[www.annualreviews.org](http://www.annualreviews.org)

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Annu. Rev. Nucl. Part. Sci. 2023. 73:421–49

First published as a Review in Advance on  
July 31, 2023

The *Annual Review of Nuclear and Particle Science*  
is online at [nucl.annualreviews.org](http://nucl.annualreviews.org)

<https://doi.org/10.1146/annurev-nucl-101920-013011>

Copyright © 2023 by the author(s). This work is licensed under a Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See credit lines of images or other third-party material in this article for license information.

## Keywords

long-lived particle, Large Hadron Collider

## Abstract

This article is a pedagogical review of searches for long-lived particles at the LHC. It is primarily aimed at experimentalists and theorists seeking to initiate and/or deepen their research in this field. We cover the general theoretical motivation and some example models, the main experimental techniques employed in searches for long-lived particles, and some of the important subtleties involved in estimating signal efficiencies and background rates.

## Contents

1. INTRODUCTION .....	422
2. THEORY PERSPECTIVES .....	423
2.1. General Theory Priors .....	423
2.2. Example Models .....	426
3. EXPERIMENTAL SIGNATURES .....	426
3.1. Displaced Tracks and Vertices .....	427
3.2. Calorimeter Signals .....	430
3.3. Other Experimental Handles .....	431
4. SIGNAL SELECTION .....	431
4.1. Signal Reweighting and Geometric Acceptance .....	431
4.2. Triggering on Long-Lived Particles .....	433
4.3. Offline Signal Reconstruction and Selection .....	435
4.4. New Developments .....	438
5. BACKGROUNDS .....	438
5.1. Standard Model Long-Lived Particles .....	438
5.2. Material Interactions .....	439
5.3. Fake Tracks and Vertices .....	440
5.4. External Particle Backgrounds .....	441
6. DEDICATED DETECTORS .....	442
7. EXAMPLE ANALYSES .....	442
7.1. Search for Long-Lived Particles in the CMS Endcap Muon Detectors .....	443
7.2. ATLAS Heavy Stable Charged Particle Search .....	444
8. CLOSING THOUGHTS .....	444

## 1. INTRODUCTION

The discoveries of the muon (1936) (1) and the kaon (1947) (2) marked the emergence of particle physics as a new field of physics: They were the first novel particles that did not fit into the contemporary framework of atomic and nuclear physics. Their discoveries famously hinged on the macroscopic lifetimes of both particles, which could be resolved beautifully with the cloud chamber technology of the early twentieth century. As the old cloud chambers gave way to modern-day silicon trackers and time projection chambers, the ability to accurately measure a decay length remained an important tool in particle spectroscopy, on par with accurate energy and momentum measurements. In the era of large particle accelerators, the focus shifted toward achieving higher center-of-mass energies and luminosities while maintaining good energy and momentum resolution. This approach led to the discovery of the top quark and the  $Z$ ,  $W$ , and Higgs bosons, all of which decay promptly. Measuring particle lifetimes nevertheless remained critical, in particular to identify decays of  $B$  mesons ( $b$ -tagging), and a number of important searches for exotic long-lived particles (LLPs) were carried out both at LEP (see, e.g., 3, 4) and at the Tevatron (see, e.g., 5, 6).

Given these historical precedents, exotic LLPs were always on the radar as a prime discovery mode for physics beyond the Standard Model (BSM) at the LHC. In the pre-LHC and early LHC eras, this primarily manifested itself in the form of signals of supersymmetry (SUSY), the dominant theoretical paradigm at the time. It is indeed remarkable how wide a range of long-lived signatures

even minimal SUSY can generate. One feature that SUSY signals have in common, however, is that they tend to be relatively hard, in the sense that they come with one or more high-energy photons, jets, or leptons. While there certainly existed models in the 2000s that generated much softer displaced signatures (7), this possibility was not yet as mainstream as it is today, both for practical reasons and because of the abovementioned theory priors before Run 1 of the LHC.<sup>1</sup> The second major evolution within the theory community has been an increased awareness of experimental subtleties, as recasting existing LLP searches or proposing new ones requires a much more detailed understanding of the capabilities and limitations of modern trigger and reconstruction algorithms as well as an intuition for the often very subtle backgrounds associated with the searches.

Since the start of the LHC, the ATLAS, CMS, and LHCb Collaborations have risen to the challenge and have produced an impressive set of new LLP results, often relying on highly innovative strategies. This has been possible because the lessons of the past have not been forgotten, even though the primary focus during the design phases of the ATLAS and CMS detectors was on the Higgs boson and SUSY discovery potential. Thanks to their hermeticity and excellent tracking capabilities, both ATLAS and CMS have proved to be powerful multipurpose detectors for long-lived signatures as well. They are further complemented by the superb tracking and vertexing capabilities of the LHCb detector in the forward regime. In those instances where the detector design appeared suboptimal for the signature of interest, the ingenuity of the analysis teams has enabled the collaborations to greatly surpass their design sensitivity. In some cases this was done by extending techniques originally developed as probes of the Standard Model (SM), such as *b*-tagging. Today, the LLP program is a major component of the LHC program as a whole, and experimental techniques that were once niche, such as  $dE/dx$ , displaced tracking, and time-delayed signals, are now powerful and common tools in the experimentalist's arsenal.

In this review, we aim to equip graduate students and postdocs entering the fascinating field of LLPs with the tools of the trade, on both the theoretical and experimental sides. By reading our review, we hope young theorists will learn the many new variables and experimental techniques that are needed in searches for LLPs and also develop a good intuition for what is and is not reasonable experimentally in terms of triggers, reconstruction, and backgrounds. Experimentalists getting started on LLPs will hopefully gain a better understanding of the modern theory priors, which types of models are of current interest, and which features they can and cannot predict. Finally, we hope that this review may also serve as a convenient bridge point for more senior colleagues on both sides of the LLP effort by summarizing the main physics points within a common and mutually intelligible vocabulary.

## 2. THEORY PERSPECTIVES

In this section, we briefly review some general theory priors and lessons from the SM, followed by some explicit examples. A more comprehensive overview of models featuring LLPs can be found in References 8 and 9.

### 2.1. General Theory Priors

As argued in Section 1, our (recent) historical focus on prompt signatures has an excellent theoretical and phenomenological justification: In all known examples, the natural width of a particle with mass  $m$  can be estimated by  $\Gamma \sim m/8\pi$  unless one or more of the following is true.

---

<sup>1</sup>As of today, the LHC's data-taking schedule is divided into runs with long shutdowns in between for maintenance and upgrades: Run 1 (2010–2012, 7 and 8 TeV), Run 2 (2015–2018, 13 TeV), and Run 3 (2022–2025, 13.6 TeV). From 2029 onward, the high-luminosity phase of the LHC project (i.e., the HL-LHC) will start.

1. The decay occurs through a heavy, off-shell particle, which implies that the width of the LLP is suppressed by a factor of  $(m/M)^n$ , where  $M$  is the mass of the heavy, off-shell state, and  $n$  is a positive, even integer that depends on the symmetries of the theory. In the SM, the role of the heavy scale  $M$  is usually played by the  $W$  mass ( $m_W$ ). The  $\pi^\pm$ , muons, and kaons are clear examples since, for instance,  $\Gamma_\mu \sim m_\mu^5/m_W^4$ .
2. The decay is subject to a severe phase-space suppression because the sum of the masses of the final states is very close to the mass of the LLP. The most spectacular example in the SM is that of the neutron.
3. The decay width is suppressed by a very small dimensionless coupling constant, which is associated with an approximate symmetry. While the smallness of  $V_{cb}$  in particular plays a role in the decay of  $B$  hadrons, there is no example in the SM for which a macroscopic lifetime can be understood solely in terms of a small coupling constant or mixing angle.

We can summarize these points in the following schematic formula:

$$\Gamma \sim \frac{\epsilon^2}{(8\pi)^{a-1}} \frac{m^n}{M^{n-1}}, \quad 1.$$

where  $n$  is always an odd, positive integer. The parameter  $a$  is also a positive integer and indicates the number of final-state particles, while  $\epsilon$  represents a potentially small, dimensionless parameter, such as a Yukawa coupling, CKM matrix element, or mixing angle. This toy formula clearly reflects the three suppression factors discussed above if we interpret  $m$  as the scale that characterizes the volume of the phase space. Usually, this means identifying  $m$  with the mass of the LLP, but for the neutron decay, for example, one would instead identify  $m \equiv m_n - m_p - m_e \ll m_W$ , which explains its extraordinarily long lifetime. Equation 1 therefore allows us to group the particles in the SM into equivalence classes, which are indexed by the integer parameter  $n$ . The higher the  $n$ , the longer the lifetime of the particle and the steeper the mass–lifetime dependence. This is illustrated in **Figure 1a** for a handful of example particles in the SM.

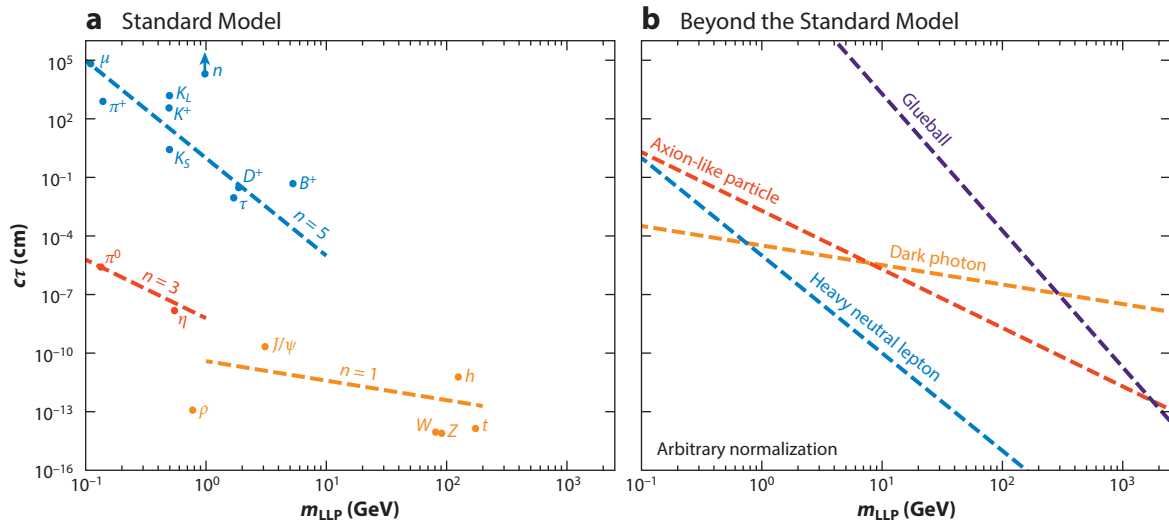
For  $\rho$  and  $J/\psi$  mesons, top quarks, and the Higgs,  $W$ , and  $Z$  bosons, none of the above suppression factors apply, and they naturally decay promptly. The  $\pi^0$  and  $\eta$  mesons decay to photons through the chiral anomaly, which induces the scaling  $\Gamma \sim m^3/f^2$ , where  $f$  is the scale of chiral symmetry breaking. Finally, the muons, tau leptons, and flavored mesons all decay through an off-shell  $W$ —hence the  $\Gamma \sim m^5/m_W^4$  trend. The neutron is the clear outlier given its enormous phase-space suppression.

Moving on to extensions of the SM, Equation 1 still applies. We can thus deploy it to critically examine a few pieces of conventional wisdom that are sometimes applied to LLP models. When inspecting **Figure 1a** without accounting for the color coding, one may be tempted to conclude that lifetime simply correlates with mass. This is the origin of a commonly held intuition:

**LoRE 1.** Heavy particles prefer to decay promptly, while light particles can have long lifetimes. Heavy LLPs therefore require “special” models and/or parameter choices. (**Misleading**)

While the proper lifetime ( $c\tau$ ) undoubtedly scales inversely with the particle mass, it is clear from Equation 1 that the above statement is too hasty: Since  $n = 5$  is typical in the SM and beyond, even a moderate hierarchy of scales in our BSM model suffices to suppress the decay width by orders of magnitude. The most famous example is presumably the long-lived gluino in a split-SUSY scenario (10), which arises from the  $m_{\tilde{g}} \ll m_{\tilde{q}}$  hierarchy.<sup>2</sup> Heavy neutral leptons

<sup>2</sup>As is the case in the SM, new LLPs may be charged or neutral, and even colored in the case of long-lived gluinos. Charged and/or colored LLPs are particularly prevalent in SUSY, while neutral LLPs tend to show up in models of dark matter, baryogenesis, and certain non-SUSY solutions to the hierarchy problem.



**Figure 1**

(a) Proper lifetime ( $c\tau$ ) versus mass for selected SM particles grouped according to the index  $n$  from Equation 1. The dashed trend lines are fits to the data, where we held  $n$  fixed but floated the intercept. The neutron is the clear outlier given its extreme phase-space suppression; to keep the axis range manageable, the neutron point was brought down and labeled with an arrow to indicate that its true  $c\tau$  is much larger than depicted here. It was also excluded from the fit that resulted in the blue curve. (b) Four examples of BSM models, where the dashed lines indicate the approximate scaling with their mass. (Small steps due to kinematic thresholds were neglected.) The overall normalization for each line was chosen to roughly correspond to the smallest  $c\tau$  allowed by current constraints, but it is otherwise arbitrary. The figure is therefore not meant to illustrate that a large  $c\tau$  at high  $m_{\text{LLP}}$  is disfavored.

(HNLs) (11) with mass  $m \ll m_W$  are another example in the same equivalence class, as shown in **Figure 1b**. In fact, just as for the SM, all BSM LLPs can be classified according to the index  $n$ . **Figure 1** shows the dark photon ( $n = 1$ ), the axion-like particle (ALP;  $n = 3$ ), HNLs ( $n = 5$ ), and dark sector  $0^{++}$  glueballs ( $n = 7$ ) (12).

This brings us to the second piece of conventional wisdom:

**Lore 2.** Theorists are not so good at predicting  $c\tau$ . (**Largely true**)

While there are important exceptions in very simple models such as HNLs or dark photons, the above statement is largely true, but not because of a lack of perseverance or cleverness: The heavy mass scale  $M$  can easily be outside the reach of the LHC and can be difficult to pin down theoretically. In some models it can be estimated using additional inputs, such as the dark matter relic density, but even in those cases, even modest uncertainties on  $M$  can be amplified significantly in Equation 1. This gets progressively worse for larger  $n$ ; the most extreme examples come from some composite particles such as dark sector glueballs. Fortunately, as discussed below in Section 4.1, the inability to predict  $c\tau$  more precisely is actually not so relevant experimentally.

All that said, there is a universal and relatively model-independent upper bound on  $c\tau$  from potentially spoiling big bang nucleosynthesis (BBN) (13). The argument goes as follows: Any particle with an observable direct production at the LHC should at some point be in thermal equilibrium with the cosmic plasma, as long as the Universe is hot enough to produce it. As the Universe cools, the LLP freezes out from the SM thermal bath and subsequently decays. If this decay takes place while BBN is occurring, the injection of additional particles and energy is likely to modify the primordial abundances, which is generally in conflict with observation. This gives us a range of upper bounds of  $\tau \lesssim 0.1\text{--}10^4$  s depending on the dominant decay modes and the abundance at the onset

of BBN. While interesting in its universality, even 0.1 s is an absurdly long time scale for any collider experiment, and this bound is therefore almost never experimentally relevant. Lower bounds on  $c\tau$  also exist, but they are always dependent on the model. This does not mean, however, that they are always trivial to evade, and they must be investigated on a case-by-case basis. Following Equation 1 and **Figure 1**, such lower bounds tend to exist for low  $m$  ( $m \lesssim 5$  GeV) and  $n \geq 5$ . Concretely, a smaller  $c\tau$  would require a larger  $(m/M)^n$ , while model-dependent lower bounds on the scale  $M$  exist from direct searches for other particles in the model in question. Important examples are HNLs and any light LLP that is a composite particle, such as in hidden valley models (7, 14).

## 2.2. Example Models

Rather than supplying a comprehensive overview of models, which can be found elsewhere (8, 9), here we touch upon three example models, one for each of the three general principles laid out in Section 2.1. We have deliberately picked examples that may be slightly less well known than the classic examples (e.g., split SUSY).

The most minimal way to obtain a macroscopic lifetime is to have the SM  $W$  boson serve as the heavy, off-shell particle. This is precisely what happens with HNLs, which mix with the SM neutrinos and undergo a three-body decay to an SM neutrino plus two other SM fermions. It is therefore no surprise that an HNL's lifetime obeys the  $c\tau \sim m^{-5}$  scaling law in **Figure 1**, similar to the SM flavorful mesons and heavy leptons. This also explains why with similar values for their mixing angles and masses, the lifetime of a low-mass HNL is much longer than that of a dark photon. For a recent and comprehensive overview of HNL phenomenology, we refer readers to Reference 11.

A strong phase-space suppression can be the result of an approximate symmetry (15), as is the case in the long-lived chargino models that generate the famous disappearing track signature. However, it can also be motivated by cosmological considerations: Suppose we have two dark sector particles  $\chi_1$  and  $\chi_2$  with  $m_{\chi_1} < m_{\chi_2}$ , which can annihilate to SM fermions ( $f$ ). Since dark matter ( $\chi_1$ ) freeze-out tends to occur at a temperature  $T \lesssim m_{\chi_1}/10$ , the  $\chi_1\chi_1 \rightarrow f\bar{f}$  freeze-out process in the early Universe receives an  $\exp(-2m_{\chi_1}/T)$  Boltzmann suppression. However, if a coannihilation process  $\chi_1 f \rightarrow \chi_2 \bar{f}$  exists, it is suppressed only by  $\exp(-m_{\chi_2}/T)$  and could easily dominate if  $m_{\chi_2} - m_{\chi_1} \ll m_{\chi_1}$  (16). The presence of a second particle  $\chi_2$  that is close in mass to  $\chi_1$  can therefore radically change the relic abundance of  $\chi_1$ . Moreover, if  $\chi_2$  can be produced at a particle collider, its decay ( $\chi_2 \rightarrow \chi_1 f\bar{f}$ ) can be very slow because of the small splitting mass between  $\chi_1$  and  $\chi_2$  (17, 18).

Though this possibility is not realized in the SM, there are many BSM models in which a particle acquires a macroscopic lifetime because of a very small coupling. Weakly interacting massive particle (WIMP) baryogenesis is such a model (19, 20). The idea is that a WIMP ( $\chi$ ) freezes out in the early Universe, but instead of being the dark matter, it is allowed to decay to SM baryons and the actual dark matter ( $X$ ). If baryon number,  $C$ , and  $CP$  are all violated in the decay, we expect  $\text{BR}(\chi \rightarrow BX) \neq \text{BR}(\chi \rightarrow \bar{B}\bar{X})$ , where  $B$  and  $\bar{B}$  represent SM final states carrying baryon number and antibaryon number, respectively. If the decay is fast compared with the expansion rate of the Universe, the inverse process where the  $B$  and  $X$  fuse back into  $\chi$  is also efficient, and we end up with no net asymmetry. We must therefore assume that  $\chi$  is an LLP in the context of the early Universe. For  $m_\chi$  in the range of hundreds of GeV, this implies that  $\chi$  is also long-lived at the LHC.

## 3. EXPERIMENTAL SIGNATURES

The experimental signatures we can use to distinguish LLPs from backgrounds strongly depend on the properties of the LLP: the  $c\tau$ , the decay modes, the LLP mass, charge, and so forth.

At the same time, the experimental capabilities of the detectors provide the boundary conditions under which these signatures can be used for signal selection. Luckily, the LHC detectors were instrumented for heavy flavor identification, have nanosecond-resolution timing capabilities, and possess other features useful for LLP detection; thus, they are well suited to search for LLPs.

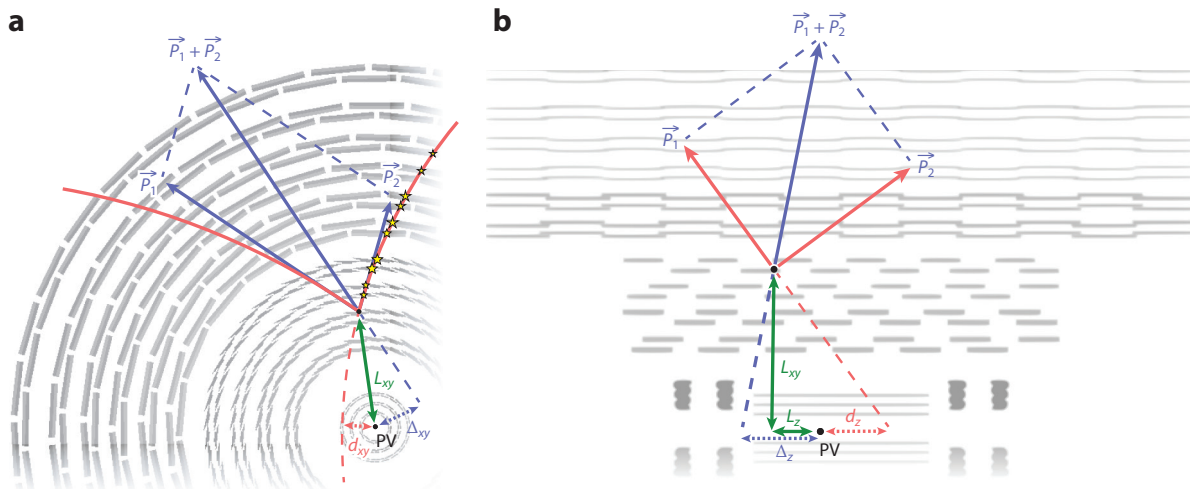
**Lore 3.** The LHC detectors are not equipped for LLP detection. (**False**)

This interplay of the characteristics of LLPs with the experimental context drives the selection choices made in each data analysis. In the subsections below, we review the various experimental tools at the disposal of the intrepid explorer hunting for new LLPs at the LHC.

### 3.1. Displaced Tracks and Vertices

Charged particles from the decay of an LLP are excellent experimental probes that can be used to discriminate a potential long-lived New Physics signal from the SM debris of the collisions, in which the charged particles mostly emerge promptly from the collision point. One of the experimental limitations restricting the identification of LLPs is the capability for the reconstruction of displaced charged particles. This reconstruction of displaced tracks is a trade-off between efficiency and purity, with computing capabilities restricting the ultimate performance. We discuss displaced track reconstruction in more detail in Section 4.3. Starting from the displaced tracks, there are several handles at our disposal for LLP identification.

**3.1.1. Track impact parameters.** The impact parameters of a track are defined by the vector from the proton collision point, which is called the primary vertex (PV), to the closest approach of the track to that point. Often this is split into its transverse component in the  $x$ - $y$  plane (where  $d_{xy}$  denotes its length) and its longitudinal component along the  $z$ -axis (where  $d_z$  denotes its length). **Figure 2** provides a graphical representation of  $d_{xy}$  and  $d_z$  for a displaced track, depicting the extrapolation toward the PV.



**Figure 2**

Illustration of the (a) transverse and (b) longitudinal cross sections of the current CMS inner tracker, with an example displaced vertex and the variables defined in the text. The energy loss ( $dE/dx$ ) in individual layers can fluctuate significantly, as indicated schematically by the varying sizes of the yellow stars in panel a. Abbreviation: PV, primary vertex.



The impact parameter is a simple yet powerful discriminator between prompt and displaced tracks. Before the advent of more powerful multivariate techniques,  $d_{xy}$  was a main ingredient in the identification of jets from heavy flavor, targeting the identification of displaced tracks coming from  $B$  and  $D$  meson decays. For this reason, the inner tracking detectors of modern particle physics detectors are designed for excellent  $b$ -jet identification, with a spatial segmentation in the inner silicon pixel detectors as small as  $50\ \mu\text{m}$ .

The resolution of the  $d_{xy}$  and  $d_z$  observables is driven by the granularity of the detector layers that the charged particle crosses first but also by the distance of the extrapolation to the beamline and by the number of hits on the track. Thus, the  $d_{xy}$  resolution of a track that emerges far from the collision vertex may be rather poor. This is particularly the case for tracks reconstructed in the muon spectrometer only, for which the pointing resolution is further deteriorated by multiple scattering in the material before the muon chambers. To account for this track-dependent varying resolution, some analyses employ the significance of the impact parameter,  $d_{xy}/\sigma_{xy}$  and  $d_z/\sigma_z$ , where  $\sigma_{xy}$  and  $\sigma_z$  are estimates of the impact parameter uncertainties as obtained from the track fit. These significance quantities are more effective in separating signal from background, at the expense of requiring more care when modeling or estimating their efficiencies and related systematic uncertainties.

For prompt background tracks, the impact parameter resolution  $\sigma_{xy}$  can reach below  $20\ \mu\text{m}$  for tracks with transverse momenta above  $10\ \text{GeV}$ , but it is closer to  $100\ \mu\text{m}$  for tracks with  $p_T \sim 1\ \text{GeV}$  or large  $\eta$  (21, 22). In the  $z$  direction, where the pixels are elongated, the resolution is somewhat worse, reaching at best  $\sigma_z \sim 30\ \mu\text{m}$  for high-momentum tracks. Another complication along the  $z$  direction is the spread of the collision vertices with a standard deviation of about  $11\ \text{cm}$ , while the transverse beam size is at the level of  $25\ \mu\text{m}$ . An accurate determination of  $d_z$  thus relies on an unambiguous identification of the PV in the presence of many simultaneous proton collisions, known as pileup. For this reason, analyses usually do not make use of  $d_z$  but instead use the already very powerful displaced track identification handles  $d_{xy}$  and  $d_{xy}/\sigma_{xy}$ .

**3.1.2. Track and vertex displacement.** If an LLP is light and has high momentum, the tracks may well be displaced but still have a small impact parameter and low significance. In such cases, the transverse ( $L_{xy}$ ) and longitudinal ( $L_z$ ) displacements of the decay vertex from the PV, as shown in **Figure 2**, may be better discriminators. For a displaced, isolated single track, the starting point cannot be unambiguously determined, but the missing inner hits still hold valuable information. In the case of two or more tracks emerging from the same LLP decay, a vertex can be fit from the tracks, yielding an often accurate three-dimensional estimate of the decay point, and thus of  $L_{xy}$  and  $L_z$ . Also, here the dimensionless significance is sometimes preferred. Particularly in the case of a decay with collimated displaced tracks, the position of the vertex in the direction of flight may be difficult to estimate, and the normalization of the displacement to its uncertainty can provide improved discrimination against poorly measured backgrounds.

**3.1.3. Track multiplicity and vertex mass.** Another discriminator that may be useful in selecting particular signals, such as displaced hadron jets, is the track multiplicity of a displaced vertex. The inefficiency of displaced tracking reconstruction, which increases with the displacement of the decay vertex, makes this variable sensitive to mismodeling in simulation. As a result, requiring a minimum or maximum number of tracks can induce a significant signal selection uncertainty. Nevertheless, a minimum number of tracks is often imposed as a preselection (having a vertex already implies at least two) as it strongly suppresses some of the backgrounds (see Section 5). Alternatively, the vertex mass (i.e., the invariant mass of the tracks that form a vertex) is a related observable that can be used for this purpose. Examples of the use of both variables can be found in References 23–26.



**3.1.4. Decay direction.** For New Physics scenarios where the LLP decays to charged as well as neutral particles (e.g., in the case of displaced tau leptons), the jet clustered from the charged particles may point off-axis with respect to the LLP flight direction. In such cases, both the impact parameters ( $\Delta_{xy}$  and  $\Delta_z$  in **Figure 2**) and the displacement of the jet can be used to separate the signal from SM and instrumental backgrounds, and in case of discovery, the direction information from the displaced tracks can be used to also characterize the neutral component. The utility of this observable is expected to decrease rapidly for larger displacements, as tracking inefficiencies wash out sensitivity from the decay direction. For signals without neutral particles in the final state,  $\Delta_{xy}$  and  $\Delta_z$  are expected to be consistent with zero, within the resolution. For those signals, they can therefore be useful variables to suppress backgrounds from fake vertices (see Section 5.3) since fake vertices tend to produce a flat distribution in  $\Delta_{xy}$  and  $\Delta_z$  (see, e.g., 27).

**3.1.5. Ionization loss.** Tracking detectors such as time projection chambers or silicon trackers can also measure ionization energy deposits per unit length along the track ( $dE/dx$ ).<sup>3</sup> For new charged particles with masses well beyond the charged pion, kaon, or proton mass, or with electric charge different from  $1e$ , the ionization deposits provide an additional handle to discriminate such signals from regular tracks from SM backgrounds.

The strongest discrimination may be achieved for particles with a boost ( $\beta\gamma$ ) that places them below the usual minimum ionizing plateau of the Bethe–Bloch curve. It is possible to discriminate different SM hadron or ion species in this manner, but only at very low momenta, of the order of 1 GeV (see, e.g., the supplementary material to Reference 29). New heavy stable charged particles (HSCPs), such as  $R$  hadrons, can have masses in the multi-TeV regime and therefore have a low enough boost to allow discrimination with  $dE/dx$  measurements (30, 31). For new particles with nonunit electric charge, on the other hand, the ionization mean free path dependence on the square of the charge makes the ionization loss a strong discriminator for a much broader mass range (32). Searches for ultrahighly ionizing particles such as magnetic monopoles are the most extreme examples in this category (33).

While ionization of material by particles is a well-understood phenomenon (28), the accurate simulation of the background particles that are produced in the collisions, as well as the description of the detector geometry, material, aging with radiation, electronics saturation, and so forth, renders a precise simulation of ionization loss in the detector very difficult for backgrounds as well as for signals. Analyses that use ionization loss as a discriminator between signal and background thus need to carefully calibrate energy loss—for instance, by using hadron tracks or muons from  $Z$  bosons, often combined with background predictions extracted directly using data.

Simulations of  $dE/dx$ , whenever possible, are performed with sophisticated software packages such as GEANT4 (34) and therefore may not be practical for theorists seeking to perform a reinterpretation. It is therefore tempting to rely on the Bethe–Bloch curve, which provides the mean energy loss of a particle through matter. This leads to an important pitfall, however, as the energy loss is a stochastic process with a highly skewed distribution, such that the mean energy loss is dominated by rare, high-energy collisions. The most probable energy loss is often a more useful estimator (28), especially for thin detectors such as the tracking layers.<sup>4</sup>

<sup>3</sup>Though its units are MeV/cm,  $dE/dx$  is usually reported in units of MeV cm<sup>2</sup>/g in the *Review of Particle Physics* (28). In these cases, one has divided out by the mass density of the target to allow for a more consistent definition for gas targets.

<sup>4</sup>Since energy loss in materials is a subtle matter, we advise theorists seeking to use analytic formulas to model  $dE/dx$  to verify all the limitations of the formulas provided in section 34.2 of Reference 28.

**3.1.6. Track timing.** Precise time measurements of the energy deposits in the detectors provide another source of information on charged particles. Inner tracking detectors close to the beamline do not provide a direct estimate of the peak time of the hits and are typically read out in a narrow time window around the expected arrival time for particles traveling at or close to the speed of light ( $\beta = v/c = 1$ ). For moderate  $\beta$ , this implies that an apparently smaller ionization is recorded, though estimating a time delay from these stochastic measurements is typically imprecise and is further confounded by the unknown LLP mass and charge. Track timing is more relevant for particles traversing the outer muon chambers, where the individual cluster timing measurements are controlled to the level of about 2 ns (35, 36). Furthermore, the track fit can be improved by making the hit position estimates from the gas ionization signals dependent on  $\beta$ . The accurate measurement of the particle's speed is leveraged by the long travel distance in the detector, which makes it possible to reach a resolution on  $\beta$  as small as 5% (31, 37).

## 3.2. Calorimeter Signals

The reconstruction of calorimetric signatures of LLPs is not hampered by low efficiency, as can be the case for the displaced tracking, but the lack of tracks means that identifying LLP-induced energy deposits is not trivial. There are, however, several powerful, advanced experimental handles, as described in the subsections below.

**3.2.1. Delayed calorimeter signals.** Calorimeters measure particle energies by observing the scintillation light that arises from electromagnetic or hadronic showers induced by the interactions of the incoming particle with an absorber. Depending on the calorimeter design, accurate measurements of the signal timing are available. In the ATLAS and CMS experiments, the electromagnetic calorimeters (ECALs) have an intrinsic time resolution as low as about 70 ps for energies larger than several tens of GeV (38, 39). In practice, the time resolution for measurements of energetic photons in the barrel is at best of the order of 200 ps (39, 40), dominated by a component that arises from the longitudinal spread of the LHC beams. For the hadron calorimeters (HCALs), the timing resolution is also rather accurate. The ATLAS Tile Calorimeter achieves a resolution as good as 0.4 ns for high-energy deposits (41), while the CMS Collaboration reports a time resolution of its HCAL of 1.2 ns for jet energies above 100 GeV (42).

These rather good time resolutions make it possible to also use the calorimeters to distinguish background low-mass ultrarelativistic particles from particles with low  $\beta$  or a delayed signal from an increased path length from a displaced decay. While a low  $\beta$  easily induces multnanosecond delays for new particles with large masses (e.g., in the TeV range), the path length can also bring about delays of several nanoseconds, well above the timing resolution (43).

**3.2.2. Displaced calorimeter signals.** The identification of signals that are unusually directed or displaced inside a calorimeter is another powerful experimental handle on LLPs decaying in the detector. Although calorimeters often sample the developing shower in several alternating layers of absorber and scintillator, a detailed segmentation of the shower is not always available for offline analysis. This is particularly the case for the CMS HCAL because the detector aggregates measurements along the shower depth. For most purposes, the shower energy estimate suffices, but for LLP identification, three-dimensional information on the shower profile is desired. It can then be used to search for decays that occur deep inside the calorimeter, using, for instance, the absence of an ECAL energy component in a hadronic shower or the detailed calorimeter cluster depth position (44, 45). Furthermore, fine-grained calorimeter shower information can help identify photons that impinge on the calorimeter at an angle due to a displaced decay, using the photon direction (46) or the (elliptical) shower shape (47).

### 3.3. Other Experimental Handles

Beyond the basic functionalities described above, detectors are being used well beyond their initial design. The recent CMS search for hadronic decays in the muon detectors is an excellent example (48): Thanks to the thick steel of the magnet's return yoke, the CMS muon system can be used as a calorimeter, clustering low-level muon chamber hits in the muon detectors. A related search also exists in the ATLAS experiment (49), which has a larger fiducial volume in its muon detector but has air as opposed to steel between its tracking detectors. As a result, the ATLAS approach relies on reconstructing displaced tracks and a displaced vertex, as opposed to searching for a calorimetric shower. Thus, the efficiency of the ATLAS search is more sensitive to the mass of the LLP, while the CMS search primarily depends on its energy. Combining information across subdetectors in the same experiment can also shed light on the individual measurements, as recently exemplified in Reference 30.

More detailed particle identification can furthermore be a powerful tool in specific exclusive searches, but it is harder to achieve without dedicated detectors. An example here is the use of the LHCb RICH detector to identify slow-moving HSCPs (50) and charged kaons coming from displaced low-mass scalar decays (51). With the advent of the HL-LHC, we will see an expansion of the experimenter's portfolio with additional subdetectors and detector capabilities. Most prominently, dedicated precise timing detectors with resolutions down to a few tens of picoseconds will be added in front of the calorimeters (52, 53). However, many of the other subdetectors are also getting upgrades that will provide opportunities for inventive new approaches to establish LLP particle signatures (see, e.g., 54, 55).

## 4. SIGNAL SELECTION

Quantifying the signal selection efficiency is more complicated for LLPs than for prompt particles because it can be difficult to simulate and parameterize. This challenge creates a number of important pitfalls for theorists trying to model an existing or proposed analysis. For experimentalists it is, moreover, important to simulate the desired signals as efficiently as possible because of the high computational complexity of the full detector simulation. In this section, we summarize a few tricks and point out possible subtleties.

### 4.1. Signal Reweighting and Geometric Acceptance

For prompt particles, the geometric acceptance is typically estimated by requiring that the final states all satisfy a set of relatively simple  $\eta$  and  $p_T$  cuts associated with the detector geometry. For unstable LLPs, this is more complicated because the trigger and reconstruction efficiencies depend strongly on the location of the decay vertex. Crucially, these efficiencies do not depend on the LLP's proper lifetime ( $c\tau$ ). This can be exploited to our advantage in a simple reweighting algorithm, as described below.

For simplicity, let us assume that we have a single LLP with momentum  $\vec{P}$  and denote the momenta of its decay products by the set of momenta  $\vec{p}_i$ . For a specific model, a sample of  $(\vec{P}, \vec{p}_i)$  can be generated with standard Monte Carlo simulation codes, such as MADGRAPH5\_AMC@NLO (56) or PYTHIA 8 (57). To calculate the efficiencies for low  $c\tau$  correctly, it is important to model the tail of the LLP  $p_T$  distribution carefully and with adequate statistics. Depending on the signal, this can mean including hard initial-state radiation (ISR) or simulating events that are weighted according to the LLP  $p_T$ . The latter is particularly important for light LLPs that are produced nonresonantly—for instance, in the Drell–Yan process or exotic  $B$  meson decays (see, e.g., 58). The reason is that the most boosted LLPs will dominate the efficiency for small  $c\tau$ .

When generating LLP events, it is most efficient to ignore the displacements of LLP decay vertices provided by the event generator and instead manually generate them along the halfline defined by  $\vec{P}$ . This translates to generating the distance the LLP traveled prior to its decay as a positive real number  $a$ , drawn from a probability distribution  $f(a)$ . The choice of  $f(a)$  is a priori arbitrary, but it should be such that the resulting vertex locations efficiently sample the detector volume of interest. For example, suppose that our sensitive detector element is a coaxial cylinder (e.g., a calorimeter) and that we care only about LLPs for which  $L_{xy}^- < L_{xy} < L_{xy}^+$ , where  $L_{xy}^-$  and  $L_{xy}^+$  denote the detector's inner and outer radii, respectively. It then makes sense to choose, for instance, the uniform probability distribution

$$f(a) = \frac{1}{L_{xy}^+ - L_{xy}^-} \quad \text{with} \quad L_{xy}^- < a < L_{xy}^+ \quad 2.$$

and define the vertex location as

$$\vec{x} = (a \cos \phi, a \sin \phi, a \sinh \eta), \quad 3.$$

where  $\phi$  and  $\eta$  are the azimuthal angle and pseudorapidity associated with  $\vec{P}$ . This ensures that as few events as possible are lost with decays outside the fiducial volume. These events can then be passed through the detector simulation to calculate the combined trigger and reconstruction efficiency  $\epsilon_r(\vec{x}, \vec{p}_i)$ , which depends only on the vertex location  $\vec{x}$  and final-state kinematics  $\vec{p}_i$  (for further discussion, see Sections 4.2 and 4.3). This calculation usually is computationally very expensive with a full-fledged detector simulation, but it needs only to be done once. With  $\epsilon_r(\vec{x}, \vec{p}_i)$  in hand, we can obtain the efficiency as a function of  $c\tau$  by assigning the following weight for each event:

$$w(\vec{P}, \vec{p}_i, \vec{x}, c\tau) \equiv \frac{e^{-|\vec{x}|/\beta\gamma c\tau} / \beta\gamma c\tau}{f(a)} \times \epsilon_r(\vec{x}, \vec{p}_i) \times w_{p_T}(\vec{P}), \quad 4.$$

where  $\beta\gamma$  is the boost factor of the LLP and  $w_{p_T}(\vec{P})$  is the  $p_T$ -dependent weight provided by the event generator. The total signal efficiency for a particular  $c\tau$  is then defined by simply averaging over the weights:

$$\epsilon(c\tau) = \frac{1}{N} \sum_{\text{events}} w(\vec{P}, \vec{p}_i, \vec{x}, c\tau), \quad 5.$$

where  $N$  denotes the number of events.<sup>5</sup> We are thus able to recycle the same set of events for any value of  $c\tau$ , greatly reducing the computational cost of the simulation.

To gain intuition for Equation 4, let us return to the simplified calorimeter example of the coaxial cylinder, such that  $\epsilon_r(\vec{x}, \vec{p}_i) = 0$  unless  $L_{xy}^- < L_{xy} < L_{xy}^+$ . For  $\beta\gamma c\tau \ll L_{xy}^-$ , we recover the expected exponential suppression of  $\epsilon(c\tau)$  regardless of the choice of  $f(a)$ . The events with the largest  $\beta\gamma$  carry an exponentially larger weight—hence the need to generate  $p_T$ -weighted events.<sup>6</sup> For  $\beta\gamma c\tau \gg L_{xy}^+ > |\vec{x}|/(1 + \sinh^2 \eta)$ , we can approximate  $e^{-|\vec{x}|/\beta\gamma c\tau} \approx 1$ , which implies that the dependence on  $c\tau$  factorizes from the efficiency. In other words,

$$\epsilon(c\tau) \approx \epsilon(c\tau') \times \frac{c\tau'}{c\tau} \quad \text{if} \quad \beta\gamma c\tau', \beta\gamma c\tau \gg L_{xy}^+. \quad 6.$$

<sup>5</sup>To generalize the algorithm to multiple displaced vertices per event, it suffices to define a weight for each displaced vertex as in Equation 4 and then combine these weights in Equation 5, including the appropriate combinatorics.

<sup>6</sup>This implies that for low enough  $c\tau$ , the efficiency will always be dominated by a handful of events, rendering the simulation unreliable. A plot of the cumulative distribution of the weights is a good way of verifying whether one is in this regime.

For this reason, the limits from all searches for a single displaced vertex scale as  $1/c\tau$  in the large lifetime limit regardless of the detector geometry and analysis details. Similarly, one can show that in the long lifetime regime, the sensitivity of searches for two independent displaced particles should always scale as  $1/(c\tau)^2$ . Moreover, it disproves the following common misconception:

**Lore 4.** If the detector is close to the interaction point (IP), it is exponentially difficult to detect LLPs with large  $c\tau$ . Therefore, placing the detector farther away leads to a higher signal efficiency. (**False**)

Indeed, Equation 6 shows that it is only linearly difficult to detect LLPs with long lifetimes. This is why some LLP searches can be sensitive to  $c\tau$  values as large as a kilometer (48). The scaling law in Equation 6 is, moreover, independent of the detector location, provided that the typical lab frame lifetime ( $\beta\gamma c\tau$ ) is much larger than the distance between the detector and the IP. On the other hand, for low  $c\tau$ , we have seen that the efficiency depends exponentially on the distance between the detector and the IP, such that a detector as close as possible to the IP always covers the largest range in  $c\tau$ . Of course, in practice some engineering and/or background considerations may require that the detectors be placed farther out, as discussed in Sections 5 and 6.

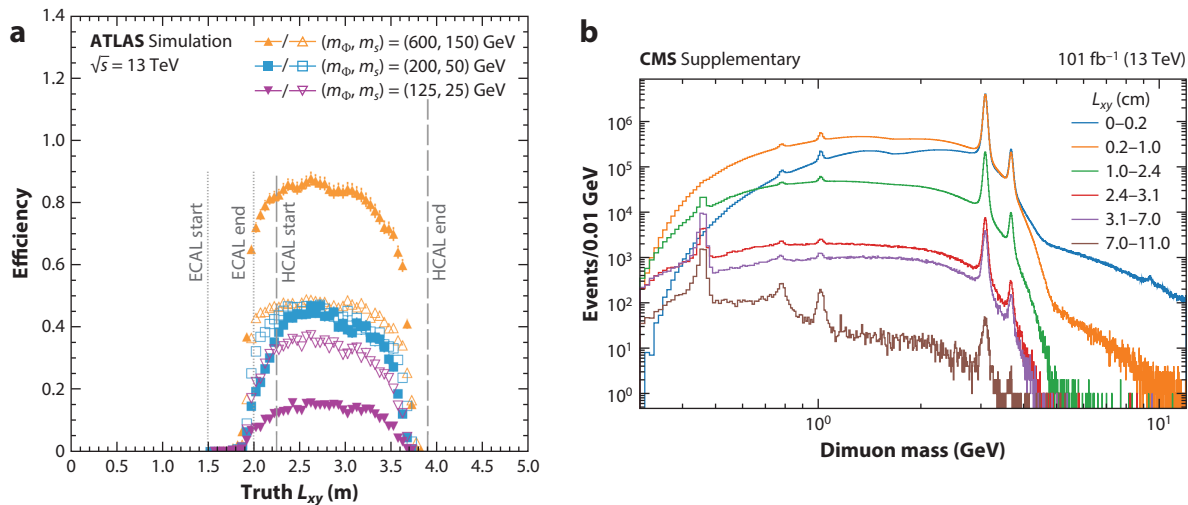
## 4.2. Triggering on Long-Lived Particles

The trigger infrastructure of any LHC experiment is both a critical and a complex component of the experiment. It is responsible for reducing the data intake from the 40-MHz beam-crossing frequency down to roughly 1 kHz, in very limited time with finite read-out and computing resources. As a result, it is in constant development through system upgrades, improved algorithms and calibrations, and shifting physics priorities. All of this particularly affects searches for LLPs because the event reconstruction is much more complex than for prompt particles. It is therefore important to study the trigger's capabilities and limitations before starting an analysis or theory study.

The ATLAS and CMS trigger systems consist of two stages: The Level 1 (L1 or LVL1) trigger (also called the hardware trigger) and the high-level trigger (HLT; also called the software trigger). The L1 trigger is responsible for reducing the rate from 40 MHz to a maximum of about 100 kHz, using coarse-grained, lower-resolution information from only a set of subdetectors. In particular, there is no time to read out and process the inner tracker data at this stage, and the calorimeters and muon systems are read out at reduced granularity. Thus, the L1 trigger makes a decision based on relatively simple observables, such as missing transverse energy (MET), the scalar sum of all reconstructed transverse energy ( $H_T$ ), a high-momentum lepton, or a jet. Events passing the L1 selection are handed over to the HLT, which reduces the rate further from about 100 kHz to about 1 kHz. The HLT performs a fairly faithful reconstruction of the full event, including tracking information. This allows for more sophisticated selection criteria, making possible many more trigger selection chains (also called paths). Nevertheless, the HLT is currently limited to an average latency of roughly 250 ms, which means that a very resource-intensive task such as reconstructing displaced tracks cannot be taken for granted for every input event. A good example is that of a purely hadronic displaced vertex in the tracker, for which displaced tracks currently cannot be efficiently reconstructed in the HLT. When a moderate  $H_T$  requirement is added, however, it is possible to trigger on a jet with an anomalously low number of prompt tracks (25). If such a prompt-veto strategy is not viable, then, depending on the signature sought, one must rely on the traditional triggers (e.g., MET,  $H_T$ , lepton).

Because of these complications, it could be tempting to conclude the following:

**Lore 5.** One cannot trigger on LLPs, so one always needs to use MET or another physics object such as a hard lepton, a jet, etc. (**False**)



**Figure 3**

(a) Trigger efficiency as a function of the truth-level  $L_{xy}$  for the calorimeter ratio trigger in the ATLAS experiment. Filled markers indicate high- $E_T$  calorimeter ratio trigger; open markers indicate low- $E_T$  calorimeter ratio trigger. Panel adapted from Reference 61 (CC BY 4.0). (b) Inclusive dimuon spectra for multiple  $L_{xy}$  bins and  $p_T^\mu > 3$  GeV, as measured in the CMS dimuon-scouting analysis. Displaced  $\omega/\rho$ ,  $\phi$ , and  $J/\psi$  mesons can be produced in the decays of boosted  $B$  hadrons—hence the persistence of those resonance peaks in the high- $L_{xy}$  bins. The  $K_S$  resonance peak is due to  $K_S \rightarrow \pi^+\pi^-$ , where both pions were misidentified as muons. Panel adapted from the supplementary material of Reference 27 (CC BY 4.0) (see [https://cms-results.web.cern.ch/cms-results/public-results/publications/EXO-20-014/index.html#Figure-aux\\_008](https://cms-results.web.cern.ch/cms-results/public-results/publications/EXO-20-014/index.html#Figure-aux_008)). Abbreviations: ECAL, electromagnetic calorimeter; HCAL, hadron calorimeter.

As we will see, however, this is not (anymore) the case since many exciting new developments and dedicated triggers have been or are being developed for LLPs. A nice example is that of displaced muons, for which the muon system can reconstruct muons that do not point to the beamline and/or do not have a matching track in the detector (59, 60). The absence of an inner detector track implies that the momentum resolution for displaced muons at the HLT is substantially worse than for prompt muons.

The ATLAS calorimeter ratio trigger is another successful example (61): It takes advantage of the fact that for an LLP decay in the hadronic calorimeter, the ratio of the energy deposited in the ECAL over that in the HCAL is much lower than that for a typical jet. This trigger has good efficiency for the range of  $L_{xy}$  that corresponds to the HCAL extent, as is nicely illustrated in **Figure 3a**.<sup>7</sup> Another interesting option is the possibility to trigger on LLPs decaying in the muon system, mainly because the muon systems have a large fiducial volume, are well shielded by the hadronic calorimeter, and are well suited for implementing new L1 trigger strategies. The ATLAS experiment already has in place such a trigger (49, 62), while CMS will have one in Run 3 (63). A more complete overview of existing and upcoming trigger strategies dedicated to LLPs can be found in the recent LLP working group report (63).

<sup>7</sup>**Figure 3** is an excellent example of a “high-value” plot for theorists seeking to understand or reinterpret an analysis: It presents an important but detector-specific quantity, the trigger efficiency, in terms of a model-independent, truth-level variable (truth-level  $L_{xy}$ ), which can easily be simulated by theorists using the recipe in Section 4.1. An analogous plot in terms of  $c\tau$  would be much less useful because one would need to reproduce the exact simulation settings of the collaboration to unfold the  $L_{xy}$  dependence of the efficiency, a process that is prone to potential pitfalls.

In some cases, it is possible to circumvent the standard bandwidth limitations by committing a reduced event format to storage at a high rate. The CMS and LHCb Collaborations recently employed this “data-scouting” technique (this is the CMS name; it is called Trigger-Level Analysis in the ATLAS experiment and Turbo Stream in the LHCb experiment) to search for displaced dimuon pairs with very low  $p_T$  thresholds (27, 64, 65), as shown in **Figure 3b**. For Run 3, the LHCb experiment entirely eliminates its hardware trigger, and thus the data-scouting concept can in principle be applied to the entire data set (66). This is expected to be a major boost to LHCb’s sensitivity for low-mass LLPs (67).

In addition to the trigger’s computing farms, the experiments also have extensive computing infrastructure, which promptly performs the more precise (and more demanding) offline reconstruction for all the events that pass the trigger (see Section 4.3). Some cases require a higher trigger rate than what this offline reconstruction can keep up with. To achieve such rates, a fraction of the data can be parked in raw format until more computing resources for their reconstruction are available (e.g., during LHC downtime). As an interesting example, in 2018 the CMS experiment registered a unique data set of about  $1.2 \times 10^{10}$  events containing a soft, displaced muon, the signature of a semileptonic  $B$  hadron decay (68). These data were recorded during the end of LHC fills when the instantaneous luminosity had dropped, resulting in lower collision rates, and when the overall demands on the trigger system were lower. Several LLP analyses that use this data set are ongoing, as it is indeed an interesting opportunity for low-mass LLP searches—that is, searching for New Physics in the approximately  $10^{10}$  inclusive  $B$  hadron decays in this sample.

Unlocking the full potential of the data-scouting and parking techniques remains an actively developing and exciting area of research during Run 3, as well as toward the HL-LHC phase. In general, for the HL-LHC, the detectors will be upgraded with capabilities that will also significantly boost the prospects to trigger on LLPs: track reconstruction in the Level 1 hardware trigger, ultraprecise timing with dedicated timing layers in front of the calorimeters, and impressively improved calorimetry in terms of, for example, depth segmentation (54, 55). As was the case in Runs 1 and 2, we fully expect that the ingenuity of the analysis teams will leverage these new hardware capabilities into often unexpected sensitivity gains for LLPs.

### 4.3. Offline Signal Reconstruction and Selection

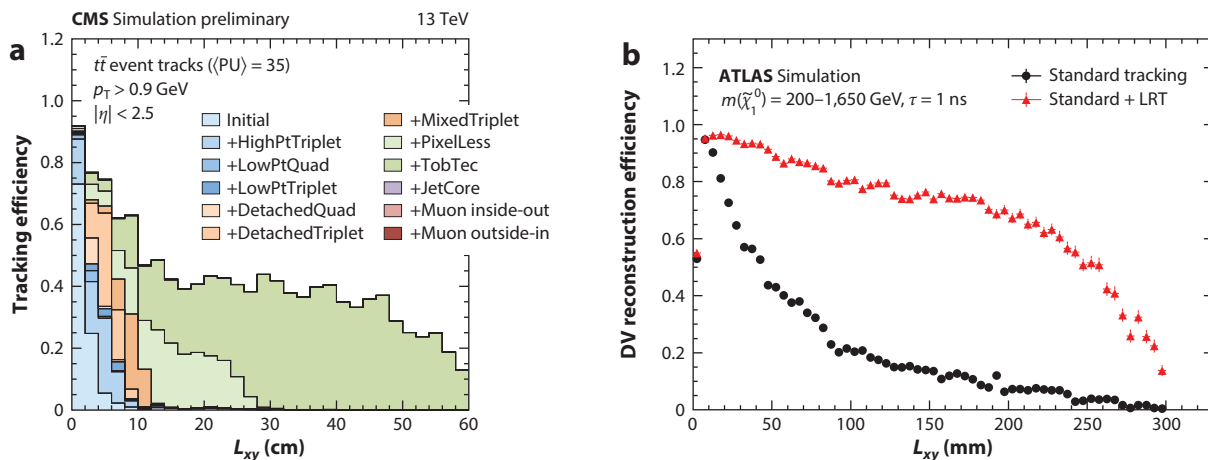
Once data events are selected by the trigger and the raw data are saved to storage, a more detailed offline event reconstruction takes place. This process uses more precise calibrations and algorithms of higher complexity, and thus requires more time, than what is possible online. The reconstruction starts from the electronic signals, local hit reconstruction, clustering in higher-level objects such as tracks and jets, and high-level object identification algorithms such as  $b$ -tagging. While standard reconstruction is described in detail in numerous references on the detectors and their physics objects, several of the reconstruction steps have interesting features that apply specifically to LLPs.

**4.3.1. Displaced tracks.** The arguably most important impact on LLP searches arises from reconstruction of the tracks of charged particles that are displaced from the PV of the event. As mentioned above, optimizing the reconstruction of displaced tracks requires trade-offs between efficiency and purity, and in practice it is limited by the available computing resources. The reason is the huge number of hits per event in the inner tracker, which makes the reconstruction of all charged particle tracks a major combinatorial challenge. In the standard track reconstruction, one requires each track to have hits on the innermost layers, which have the best spatial resolution, and to originate from near the beamline. These conditions provide an enormous reduction in the number of possibilities and therefore in the complexity of the computational task.



After the prompt tracks are reconstructed, further reconstruction of displaced tracks is attempted with the remaining, unused hits in the tracker. These unused hits are largely from particles produced in nuclear interactions of primary tracks in the detector material and from particles that were too soft to be considered in the standard reconstruction. For the current pileup conditions at the LHC, those unused hits are still in the range of several thousand per event. The further reconstruction of displaced tracks with looser constraints, in particular on  $d_{xy}$ , still involves a high level of combinatorial complexity, which in addition scales nonlinearly with the amount of pileup. As a consequence, the reconstruction of displaced tracks can be only partially efficient, to avoid picking up too many “fake” tracks from nuclear interactions or from combining unrelated hits into a displaced track. Such fake tracks can be a background in some LLP searches, as discussed in Section 5.3.

Concretely, the inner parts of the ATLAS and CMS detectors are equipped with several layers of high-precision silicon pixel detectors, starting from a transverse distance to the beam of about 3 cm and yielding typically at least three precise measurements for tracks with a transverse displacement of about 10 cm (69, 70). Current offline reconstruction algorithms operate at nearly full efficiency for tracks emerging before the first pixel layer—which is essential for excellent  $B$  hadron identification—and still show above 50% efficiency for tracks with  $p_T \gtrsim 1$  GeV produced just outside the outer radius of the pixel detector, for LHC Run 2 pileup conditions. For larger transverse displacements, an efficiency of about 40% is still achieved at a track production radius of 30 cm (21, 71), as shown in **Figure 4**. For the LHCb experiment, the inner VELO detector (72) is specifically constructed to be highly efficient for forward displaced charged particles arising in the VELO detector (73). During Run 3, LHCb will, moreover, have access to a new class of highly displaced tracks (T tracks) (74), which will enhance the experiment’s signal efficiency for LLPs with larger  $c\tau$  in particular.



**Figure 4**

(a) Track reconstruction efficiency in  $t\bar{t}$  events as a function of simulated track production vertex  $L_{xy}$  for the CMS tracker in its 2017 data-taking configuration. The various colors represent the different iterations in the reconstruction algorithm, where successively looser criteria are applied to increase the efficiency. Panel adapted from Reference 75 (CC BY 4.0). Green shading represents tracks reconstructed with algorithms specialized for displaced tracks. (b) Track reconstruction efficiency in the ATLAS experiment for displaced charged hadrons produced by the decay of long-lived gluino  $R$  hadrons, as a function of the displaced decay  $L_{xy}$ . The efficiencies of the standard tracking and of the standard plus dedicated displaced tracking (LRT) algorithms are shown. Panel adapted from Reference 24 (CC BY 4.0). Abbreviations: DV, displaced vertex; LRT, large radius tracking.

Measuring the efficiency of displaced tracking in data is a challenge. The SM  $K_S$  and  $\Lambda$  hadrons, commonly referred to as  $V^0$ s, are luckily modeled well in simulation (76–78) and can thus be used as a source of displaced tracks on which displaced tracking performance can be tested.

**4.3.2. Short tracks.** Short tracks are of great interest as a signature of displaced decays of new charged particles. Often, such particles decay into a nearly mass-degenerate neutral particle that escapes detection (see Section 2.2), along with a very soft track (79, 80), or they may continue into another charged track with a kinked signature (81).

In the former case, to reduce backgrounds, a track reconstruction is employed that—unlike standard tracking, which is more permissive—requires all consecutive hits on the track to be recorded; the absence of hits on several outer layers, in combination with a quiet calorimetric environment (82, 83), serves as confirmation that the track has stopped. The shortest tracks thus reconstructed are promptly produced and have four consecutive hits in the pixel detector. A strong confirmation of such a signal would come from matching a soft pion from the short track's end point. However, such a pion not only would appear displaced but also would circle on a helix in the tracker detector because of its low momentum. This reconstruction is in principle possible, but it would be a costly investment. At this point it is unnecessary since the background can be suppressed sufficiently in other ways.

Regarding kinked tracks, there is currently no dedicated analysis for this signature. Nonetheless, current disappearing track searches likely have excellent sensitivity already (84). This is because the outer part of the kinked track typically fails to be reconstructed, such that the signature is effectively identical to that of a disappearing track.

**4.3.3. Displaced muons.** For muons with a matched charged track in the inner tracker, the efficiency is set by the inner track reconstruction, as described above. Muons can also be reconstructed up to very large displacements beyond the inner tracker by using the muon system only (60, 85), albeit with a lower position and momentum resolution. This can be done efficiently up to track transverse displacements as large as 400 cm, imposing quality requirements that suppress backgrounds from hadrons punching through the calorimeters and other nonbeam backgrounds (see Section 5). For these muons, it is important to not use the collision vertex as a constraint, which would bias the reconstructed momentum to lower values, inducing inefficiencies at large impact parameter values.

**4.3.4. Displaced electrons, taus, and photons.** For displaced electrons originating in the first part of the tracker, the efficiency will mostly mirror the displaced track reconstruction performance. The reconstruction efficiency of such displaced electrons can be estimated in data by looking at photon conversions in the detector material (86). If the electron track is too displaced to be reconstructed but emerges before the EM calorimeter, the electron can still be reconstructed as a photon. While this signature comes at the cost of a higher background, it is particularly useful in the trigger (87, 88) prior to a more detailed offline selection.

The reconstruction of genuine photons from LLP decays benefits from a dedicated treatment, with adapted identification requirements on timing or direction (46, 47). Tau leptons emerging from displaced decays, either as displaced electrons or muons or as displaced jet-like signatures, are also an interesting target. Dedicated displaced hadronic tau identification is the most complex final state, which we expect to see develop strongly during LHC Run 3.

**4.3.5. Displaced vertices.** Once the displaced tracks have been reconstructed, the reconstruction of their corresponding displaced vertex is essentially fully efficient. Nevertheless, in cases where the tracking efficiency is somewhat inefficient, it may still be better to forgo vertex

reconstruction since the overall efficiency will scale as the tracking efficiency raised to the number of tracks required to reconstruct the vertex. In particular, if one expects multiple displaced decays in the event, it may be more beneficial to ask for a number of displaced tracks without requiring that they belong to a single vertex (89).

**4.3.6. Exotic objects.** In the above subsections, we have discussed aspects of offline reconstruction of the most common final-state objects. More specialized reconstruction furthermore targets specific experimental signatures, some of which, like showers in the muon system, are touched upon in Section 3.3. Also exciting is the potential for future development of advanced tracking for exotic tracks from monopoles, as they are bent along the magnetic field direction. Quirks oscillating in pairs when traversing the detector (90, 91) are another example.

## 4.4. New Developments

A recent avenue of significant progress is the development of more advanced techniques that aggregate various sources of experimental input into a multivariate discriminator. This can significantly boost the sensitivity of a search (92, 93). The potential downside of less detailed control of a selection requirement can be offset by using the neural network as a discriminator for selecting physics objects, coupled with more robust background predictions from the data. The simulation description of the multivariate discriminator, on the other hand, is difficult to assess because of the limited types of control samples with genuinely displaced tracks available. Techniques like domain adaptation (92) and adversarial networks (93) can alleviate this challenge.

As ever more advanced novel reconstruction algorithms are developed, lower-level detector features come into focus. This is at odds with the increasing computing complexity, which requires the use of high-level and thus compact objects. Therefore, advanced offline reconstruction benefits integration within the standard reconstruction chain (see, e.g., Reference 94 for the integration of ATLAS Large Radius Tracking into the main reconstruction from Run 3 onward). Alternatively, it must run on dedicated data streams in which low-level objects are available for such dedicated analyses. Either way, such efforts require sometimes heroic, long-term investments.

## 5. BACKGROUNDS

Searches for LLPs make use of nonstandard experimental signatures in both online and offline selection to drastically reduce prompt backgrounds at a modest cost to signal efficiency. This has led to the following general assumption in many LLP theory studies:

**Lore 6.** LLP searches are background free. (**False**)

Though in many analyses the backgrounds can indeed be reduced to negligible levels, this is possible only after extensive and subtle analysis efforts on the experimental side. Yet, even after extensive background reduction efforts, unusual irreducible backgrounds from instrumental, algorithmic, or other origins often remain. In the subsections below, we review various sources of backgrounds in searches for LLPs and some methods employed to eliminate them. Often, simulations are not reliable or sufficiently robust to estimate the remaining background, even if very small, and data-driven techniques are needed for their reliable quantification.

### 5.1. Standard Model Long-Lived Particles

Charm- and bottom-flavored hadrons are the most ubiquitous SM background for LLP searches as they can easily produce multitrack displaced vertices. A priori, these can be removed effectively by requiring a vertex mass of  $\gtrsim 5$  GeV (see, e.g., 95), though this severely limits the sensitivity

to low-mass LLPs. The proper decay lengths of charm- and bottom-flavored hadrons are also only  $\mathcal{O}(0.1)$  mm, such that cuts on  $d_{xy}$  and/or  $L_{xy}$  are also very effective (see, e.g., References 27 and 96 and **Figure 3b**). One must be aware, however, that there are still many events with highly boosted  $b$ - or  $c$ -jets, which could leak into the signal region in some cases. It is therefore essential to use  $p_T$ -weighted events when using simulation to estimate these backgrounds. Moreover, it is advisable, whenever possible, to also reweight the events in terms of the long-lived meson's decay vertex location, as described in Section 4.1.

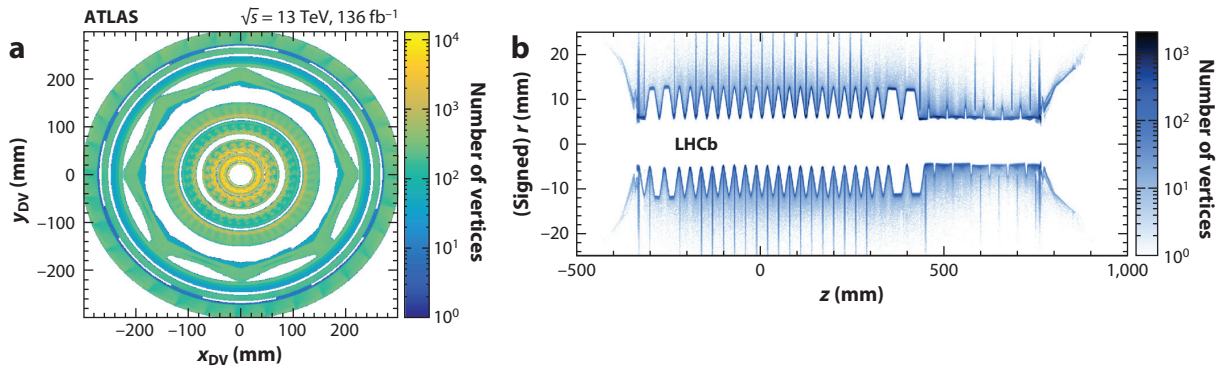
Charged pions have a proper decay length of 780 m, and the vast majority reach the calorimeter without decaying. Bearing in mind Section 4.1, we can roughly estimate the probability for a  $\pi^\pm$  to decay in the tracker: For a  $\pi^\pm$  with  $\beta\gamma \sim 10$ , the probability is  $1 \text{ m}/(10 \times 780 \text{ m}) \approx 10^{-4}$ . This seems very small, but one must keep in mind that every collision produces tens to hundreds of  $\pi^\pm$ , though most of them are soft. Multiplying this with the  $\mathcal{O}(100)$  for the number of pileup collisions per event, we see that an  $\mathcal{O}(1)$  fraction of all events will have a  $\pi^\pm$  that decays “early.” Similar considerations apply for  $K^\pm$ ,  $K_L$ , and especially  $K_S$ . Both pions and charged kaons are therefore a significant source of individually displaced muons. Fortunately, this background drops rapidly if isolation is imposed and if the  $p_T$  requirement on the muon is tightened. Moreover, kaons can produce a genuine displaced vertex. For example, **Figure 3b** shows a clear peak in the kaon mass range, which is due to  $K_S \rightarrow \pi^+\pi^-$  decays in which both pions are misidentified as muons. In general, such displaced vertices can easily be eliminated by requiring a vertex mass well above the kaon mass and/or by demanding a larger number of tracks.

Jets with an anomalously low number of tracks are an interesting handle for displaced decays near the back of the tracker or in the calorimeter (93) or for strongly interacting dark matter candidates (97). A priori, strong interaction processes may produce hadronic jets that consist primarily of  $K_L$  and neutrons, which do not leave traces in the tracking detectors but do deposit energy in the HCAL. While the corresponding probability per jet is very low, this background can nevertheless be important because of the huge cross section for QCD jet production.

For all the above backgrounds, the most important point is perhaps to be mindful of the limitations of all simulation codes since they are not designed to model very rare effects in specific corners of phase space. To some extent, this limitation can be offset with the weighting procedures described in Section 4.1. Nevertheless, more often than not, a data-driven validation is needed.

## 5.2. Material Interactions

The details of the detector material—that is, its density and geometry—are critical for LLP searches as this material can both reduce and generate backgrounds. Trackers are primarily always designed to have as little material as possible because particle scattering in the sensors or support structures adversely affects the momentum resolution of the tracks. This design driver is also helpful in LLP searches since inelastic collisions or photon conversions in the detector material can produce secondary vertices. And by no means is this a rare effect: Approximately 5% of all  $\pi^\pm$  with  $p_T \gtrsim 5$  GeV create a secondary vertex within the CMS inner tracker (98). Extrapolated to HL-LHC conditions, this implies secondary vertices at a rate of  $\sim 30$  MHz—that is, essentially in every event. Fortunately, this background is strongly dependent on the track  $p_T$ , the vertex mass, and the number of tracks per vertex (98). Moreover, the exquisite vertex resolutions of their trackers allow the ATLAS, CMS, and LHCb Collaborations to make beautiful maps such as those shown in **Figure 5**, pinpointing precisely where the vertices are produced. These can then be interpreted as radiographic images of the detector material and can be used as veto maps to suppress backgrounds efficiently in an LLP analysis. To avoid masking too much of the detector volume, which reduces signal efficiency, further track and vertex selections can be made to improve



**Figure 5**

Number of reconstructed secondary vertices in (a) the ATLAS inner tracker (23, 24) and (b) the LHCb Vertex Locator (VELO) (99). Panel a adapted from Reference 24 (CC BY 4.0). Panel b adapted from Reference 99 (CC BY 4.0).

the vertex resolution—for instance, avoiding vertices from very collimated tracks. Kinematic cuts may also prove useful against SM LLPs produced in material interactions, which may decay in the unmasked region.

Unlike the trackers, the calorimeters are designed to stop as many particles as possible. Especially the hadronic calorimeters are made out of heavy elements and are made as thick as can be accommodated by engineering and financial constraints. For example, the ATLAS HCAL in the barrel is made of  $\sim 2$  m of steel plus scintillator tiles. At  $\eta \approx 0$ , this corresponds to roughly 7.4 nuclear interaction lengths for a total of 9.7 nuclear interaction lengths when the liquid argon ECAL is included. We can therefore estimate the probability that a hard  $\pi^\pm$  will punch through the calorimeter into the muon system to be  $e^{-9.7} \approx 5 \times 10^{-5}$ . This is small, but not negligible, given the huge number of  $\pi^\pm$  that impinge on the calorimeter during a typical run. Thus, large numbers of hadrons still punch through the calorimeters into the muon chambers, where they may fake a muon or a displaced vertex. Fortunately, hard hadrons are typically found inside jets, and the calorimeter itself is therefore an effective veto for such punch-through events. This background cannot be simulated faithfully and must be estimated from data. A search along these lines was first performed by the ATLAS Collaboration, using a dedicated LLP trigger on activity in the muon system (49, 100). The CMS calorimeter is significantly thinner and thus suffers from more punch-through backgrounds. However, the level of punch-through that reaches the muon chambers is reduced by the steel in the return yoke of its magnet as an additional shield (48).

### 5.3. Fake Tracks and Vertices

While material interactions inject real displaced particles into the detector, tracks that do not arise from a real particle may also be reconstructed. Fake hits from detector noise, or more likely hits from the cloud of thousands that come from the pileup particles in the inner detectors, may align to form a track. For example, assuming an average of 50 pileup vertices, the four innermost tracking layers in the ATLAS barrel collect an average of  $\sim 2 \times 10^4$  hits per bunch crossing (101). Roughly 25% of these hits will be associated with reconstructed, prompt tracks, leaving several thousand unassociated hits in each event. While it is very unlikely that such unrelated hits will accidentally line up and be reconstructed as a promptly produced track with a good fit quality, this becomes more likely for displaced tracks for which only a handful of hits are required, as, for example, in searches for disappearing tracks (82, 83) or in cases of displaced muon reconstruction (60, 85).

Fake vertices are also common and can be composed out of a set of real but unrelated tracks. This frequently happens when a vertex is reconstructed from muons with very large displacements or when vertices with low displaced track multiplicity are considered in the inner detector. A particular source of fake vertices arises from overlapping tracks from different pileup vertices. While in this particular case such fakes can be identified from the tracks lining up in a plane with the beamline, more generally, fake vertices may be rejected by vetoing hits on the tracks upstream of the vertex.

Backgrounds from fake tracks and vertices are best measured using data, though there have been attempts to estimate them in phenomenological studies (102–104). In most cases, however, they must indeed be extracted from data, and caution is needed especially when a very high degree of background rejection is required.

## 5.4. External Particle Backgrounds

The above sources of background relate directly or indirectly to proton–proton collisions. There are, however, a few additional sources of background that may enter the signal phase space of LLP searches.

**5.4.1. Beam halo.** As protons travel through the LHC in dense bunches, some stray too far from their ideal trajectory, posing problems for the sensitive equipment of the LHC and the experiments. Although dedicated absorbers clean the proton beam of such beam halos, beam halo muons traveling along with the proton bunch parallel to the beam may be produced and may traverse the thick shielding in front of the detectors. The rate of such muons decreases strongly as a function of the radial distance to the beam (105).

These muons create some mostly harmless background hits in the tracking detectors, but in the calorimeters and muon spectrometers they can lead to unusual backgrounds. Since the muons do not originate from the proton collisions, their timing is asynchronous, though in a predictable manner. They may leave straight muon tracks in the forward muon tracker as well as significant energy deposits in the calorimeters. The parallel direction of the muon tracks, the potential match of these tracks with calorimeter deposits, the anomalous shapes of such deposits parallel to the beamline, the depth information, and the early or late timing of these signatures may be used to suppress this background. The azimuthal  $\phi$  distribution is another interesting handle: The beam halo rate spikes in the horizontal plane and is smallest near the bottom of the detector, as the floor of the LHC tunnel acts as a shield (105, 106).

**5.4.2. Cosmic muons.** Some of the relativistic muons that are created in cosmic ray showers reach underground and cross the detectors downward with a rate of roughly  $1 \text{ Hz m}^{-2}$  (105). Their energy spectrum falls exponentially and they exhibit a significant spread in direction as well as a geometric asymmetry, mainly due to the access shafts above the detectors. Only a small fraction is reconstructed as tracks or muons, since cosmic muons only rarely have a direction that is compatible with the constraints of the track reconstruction algorithms. In addition, the arrival time of the muons is random, and thus the timing of the energy deposits in the various subdetectors, including calorimeters, is not synchronized with the collision with which it may overlap. For downward-going tracks in the upper half of the detector, the travel direction is even opposite of what is expected for particles originating from collisions. Given these features of cosmic muons, they are only a potential background in LLP searches, where anomalous timing or track displacement is selected (see, e.g., 60, 107).

**5.4.3. Satellite collisions.** Though the LHC beam delivers its main bunches in 25-ns intervals, the beam inevitably also has “satellite bunches,” which follow and precede each main bunch with



a 2.5-ns time gap. These satellite bunches contain about  $10^{-5}$  times the number of protons of the main bunches, and their collisions can generate very rare, out-of-time backgrounds. Such backgrounds arrive with well-defined 5-ns delays compared with the primary collision and can be particularly important for analyses that focus on delayed calorimeter signals (e.g., 43).

## 6. DEDICATED DETECTORS

As discussed above, the ATLAS, CMS, and LHCb detectors are excellent tools in the hunt for LLPs. There are, however, some important cases where dedicated LLP detectors can perform qualitatively better. Concretely, there are three reasons to consider auxiliary detectors at relatively large distances from the IP: (a) to catch LLPs with very forward kinematics that are inaccessible to the main detectors, (b) to allow for a different detector technology, and (c) to allow room for additional shielding. All such detectors rely on the third of these to some extent.

The FASER experiment (108) makes use of a small service tunnel located 480 m forward from the ATLAS IP. Its extreme forward location allows it to search for LLPs produced through the huge pion flux that goes down the beam pipe with each collision. A larger version of FASER would need to be housed in a proposed, dedicated cavern: the Forward Physics Facility (109). The FACET experiment (110) would follow a similar philosophy and would be located roughly 100 m forward from CMS. It would cover a somewhat larger pseudorapidity range than FASER with beam backgrounds that may be more challenging. Finally, the SND@LHC detector (111) is meant to detect high-energy forward neutrinos from the ATLAS IP, but it may also have sensitivity to some low-mass LLPs (112). When modeling the acceptance of forward LLP detectors, it is important to remember that standard simulation tools, such as PYTHIA and MADGRAPH, are typically not appropriate in this kinematic regime. Instead, specialized tools are needed (see, e.g., 113).

The MoEDAL (114) and milliQan (115) experiments rely on alternative detector designs that are specialized for stable LLPs that leave anomalous tracks. MoEDAL has searched for magnetic monopoles trapped in sets of aluminum rods, which were scanned for monopoles with a SQUID sensor (116). The MoEDAL Collaboration also aims to install an extension to look for decaying LLPs at intermediate rapidities (117). Similar to FASER and SND@LHC, the milliQan experiment is housed in an existing service tunnel, but at moderate rapidity above CMS. It is shielded by 17 m of rock and designed to detect fractionally charged particles. It can do so by looking for coincident hits in four aligned bars of plastic scintillator. Currently, a first modular phase of the detector is under construction (118) with a sensitivity that will open up a large new phase space during LHC Run 3 and the HL-LHC phase.

Finally, some proposed experiments such as CODEX-b (119, 120) and MATHUSLA (121, 122) would look for displaced vertices in a similar manner as the ATLAS and CMS muon chambers but would have a much thicker shield between the detector and IP. This would result in much lower backgrounds, at a cost in geometric acceptance. MATHUSLA would be constructed on the surface above CMS, using the rock as shielding. CODEX-b would be installed directly in the LHCb cavern, and a suitable shield would need to be constructed. In either case, the shield would need to be equipped with an active muon veto to effectively reject muon-induced secondaries as potential backgrounds (119, 120, 123).

## 7. EXAMPLE ANALYSES

Having reviewed the basic techniques underlying LLP searches, we now present in more detail two example analyses where several of these aspects come together.

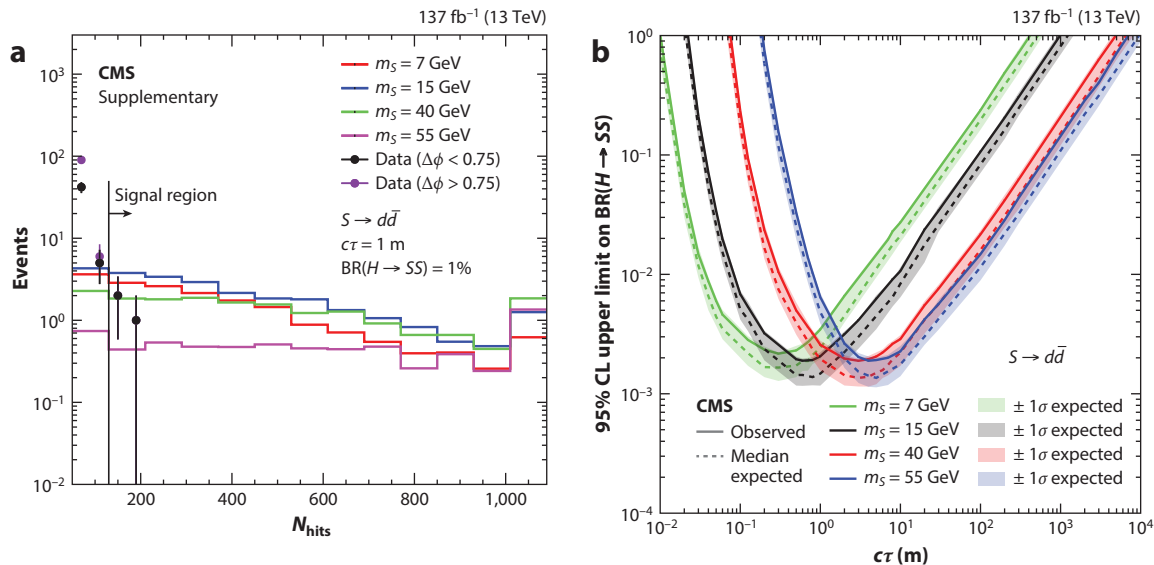


## 7.1. Search for Long-Lived Particles in the CMS Endcap Muon Detectors

In Reference 48, the CMS Collaboration used its muon detectors for the first time to identify an electromagnetic or hadronic shower of particles from an LLP decay in the muon detectors. This use of the muon system as a calorimeter is possible because of the up to 27 nuclear interaction lengths of detector material serving as a background absorber before the first layer of muon detection and also because the steel from the magnet return yoke interleaved with the several layers of muon detectors induces shower development. A single cluster of  $> 130$  hits is required, not aligned with the large missing momentum that is used to trigger the event, in addition to a jet that signals the presence of significant QCD ISR.

Various unusual backgrounds contribute to this signal selection: high-energy jets punching through the material before the muon detectors; muons that undergo bremsstrahlung, inducing an electromagnetic shower; and decays from SM LLPs like  $K_L^0$  particles, mostly from pileup collisions. These backgrounds are suppressed by vetoing events with jets or muons reconstructed in the direction of the hit cluster or with various types of activity in the first detection layers of the muon system. In addition, beam halo, calorimeter noise, and cosmic backgrounds require dedicated event vetoes to clean the selected event sample. Finally, the remaining background is dominated by SM LLPs from pileup collisions from the same and neighboring bunch crossings, and the accurate timing of the muon detectors is used to further reject events with out-of-time clusters of hits in the muon detector.

The distribution of the number of hits ( $N_{\text{hits}}$ ) of the selected cluster is shown in **Figure 6**, after the full event selection except for the  $N_{\text{hits}}$  requirement. One strength of the analysis becomes apparent: Requiring just a single cluster reduces backgrounds in the signal region to negligible levels, while the signal models show a high selection efficiency for clusters that generate showers



**Figure 6**

(a) Number of cluster hits ( $N_{\text{hits}}$ ) after full event selection, shown for background data and several signal models. Panel adapted from the supplementary material of Reference 48 (CC BY 4.0) (see [https://cms-results.web.cern.ch/cms-results/public-results/publications/EXO-20-015/index.html#Figure-aux\\_002-a](https://cms-results.web.cern.ch/cms-results/public-results/publications/EXO-20-015/index.html#Figure-aux_002-a)). (b) Upper limits on the branching fraction of the Higgs boson to a pair of scalars, each decaying further to quarks, as a function of the lifetime and for various mass assumptions of the scalar. Panel adapted from Reference 48 (CC BY 4.0).

in the considered fiducial detector volume. The remaining background cannot be accurately estimated from simulation, but it can be predicted from the data because the angle  $\Delta\phi$  between the cluster and the missing momentum is uncorrelated with  $N_{\text{hits}}$ . As an example of the sensitivity of the analysis, **Figure 6** shows the reach in branching fraction of an SM Higgs boson decaying to a pair of scalars that each decay further to a pair of quarks. There is no significant dependence on the mass of the scalar except for a shift in the lifetime probed as a function of scalar mass, which is induced by the boost of the scalars affecting their lifetimes in the lab frame. The behavior of the limits at high lifetime is also noteworthy, where the  $1/c\tau$  dependence in Equation 6 becomes very prominent, and lifetimes beyond 1 km are probed.

The selection of the signal in this analysis can be improved significantly by avoiding the limitation of requiring events to be triggered by a large momentum imbalance. In the LHC Run 3 data taking, this will be possible as new triggers are deployed (63) that directly target the identification and selection of clusters with many hits in the muon detectors already at the first-stage hardware trigger.

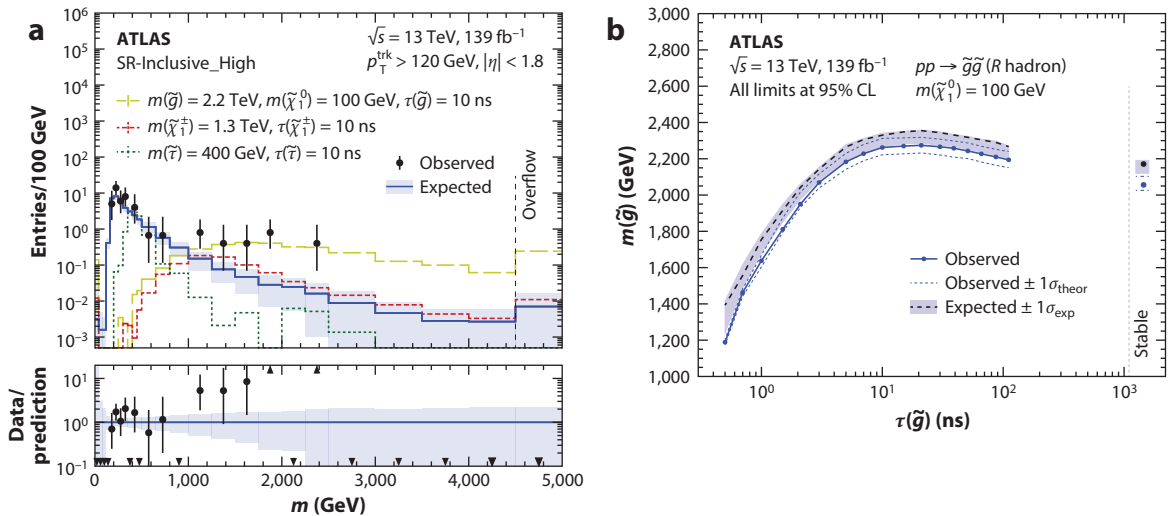
## 7.2. ATLAS Heavy Stable Charged Particle Search

The search in Reference 30 by the ATLAS Collaboration is strongly motivated by SUSY models, but it is broadly applicable to LLPs that are charged or colored. The gluino can be an example of such an HSCP as it can acquire a very long lifetime if the squarks are parametrically heavier (10). Because of the rapidly falling parton distribution functions, we expect heavy particles to be produced close to the kinematic threshold, especially when produced in pairs. Thus, HSCPs are expected to move substantially more slowly than the speed of light as they make their way through the detectors ( $\beta\gamma < 1$ ). This has two important consequences: First, they are expected to arrive later at various detector segments compared with particles that travel at the speed of light. Second, they are no longer minimally ionizing particles, and they will deposit an anomalous amount of energy as they pass through detector material.

The HSCP does not deposit much energy in the calorimeter, and it often arrives too late in the muon system to be reconstructed by the muon trigger. This means that the transverse momentum of the HSCP is often not reconstructed correctly with information from the calorimeters and muon systems at the first trigger level alone. The search in Reference 30 takes advantage of this feature by using the calorimetric MET trigger to record the events. Offline, the analysis hinges on the anomalous ionization yield of the track ( $dE/dx$ ) in the inner tracker, as discussed in Section 3.1. Combined with the usual  $p_T$  measurement of the track, this variable can be used to estimate the mass of the HSCP, as shown in **Figure 7a**. After backgrounds from multijets, leptonic  $W$ , and overlapping tracks have been mitigated as much as possible, the main residual background is thought to be from hard isolated tracks for which the ionization yield in several tracking layers has fluctuated upward following the Landau distribution associated with this process. It is not possible, however, to simulate this background from first principles, and it is therefore estimated from data. As one can see in **Figure 7**, the ATLAS experiment finds an excess of  $3.6\sigma$  local significance and  $3.3\sigma$  global significance (30). Time will tell whether this excess is just one of the many anomalies that are inherent to gargantuan data sets or whether it will be the first observable sign of BSM physics at the LHC.

## 8. CLOSING THOUGHTS

The ability to reconstruct and characterize displaced particle decays has historically delivered many discoveries and continues to be an essential tool for particle identification. It is a misconception that heavy, BSM particles are supposed to decay promptly: Not only do there exist many



**Figure 7**

(a) Invariant mass distribution for signal, expected background, and data in the high- $dE/dx$  bin category, labeled SR-Inclusive\_High. (b) Expected and observed lower limit on the gluino mass in a long-lived gluino model. Figure adapted from Reference 30 (CC BY 4.0).

counterexamples, but the general conditions for macroscopic lifetimes are very simple and likely generic, given our experience with the SM itself. Moreover, searches for displaced signatures benefit from a number of experimental handles that are unavailable for promptly decaying new particles and therefore can often maintain very low backgrounds, even with very high integrated luminosity. LLPs are thus one of the primary areas where we may achieve a major discovery at the HL-LHC.

Searches for LLPs are also fascinating because they are complicated, subtle, and often messy. The signal efficiency depends on the details of the detector design, the trigger capabilities and limitations, and the offline reconstruction methods available. The backgrounds are delicate and sometimes impossible to simulate; understanding them requires careful and clever estimation techniques. Moreover, the detector capabilities continue to evolve, especially with the upcoming HL-LHC upgrades, all while new trigger strategies are constantly added, unlocking qualitatively new searches. Thus, there is much room for new ideas that may affect the direction of the ongoing and future search programs.

The flip side is that the up-front investment can seem very high for those getting started in this field, especially as a good deal of the essential knowledge is often either unwritten or buried deep in thick technical design reports. In this review we have attempted to collect some of this knowledge in a manner that will be accessible to novices, along with some practical suggestions. We conclude with a few words of advice both for beginning theorists and for experimentalists. In particular for theorists, we suggest the following.

- It pays to learn how to read technical design reports and performance papers. While they may seem daunting at first, many are structured in a similar manner, and after a while you would be surprised how quickly you can mine them for the trigger and/or reconstruction efficiencies you need.
- Be realistic when it comes to your ability to accurately model complicated backgrounds. Rather than producing overly aggressive or overly conservative (projected) limits, consider plotting signal yields only, and let experimentalists follow up with a full analysis.

- Consult extensively with experimental colleagues, either by collaborating directly or by reaching out to the authors of the analyses you are studying. Figures may rely on important assumptions that are not obvious from the published material. The conveners of the analysis groups can put you in touch with the main analysis authors.

For experimentalists, we recommend keeping the following in mind.

- Connect with the experts in the individual subdetectors, who are often eager to understand together low-level issues in the detector data or simulation. Also share your challenges and progress regularly with your colleagues, as similar problems may find solutions in very different contexts.
- Plan ahead any special trigger or computing needs, including data storage and access. Also prepare to publicly release efficiencies, cut-flow tables, and so forth, to allow theorists to recast your analysis, which can greatly amplify its impact (see, e.g., 27).
- Keep abreast of the phenomenology literature and community. Theorists are generally happy to connect and follow up on their work. This may even lead to fruitful common research.

LLPs provide a compelling window on BSM physics and a rich research arena for both experimentalists and theorists. With this review, we hope to have conveyed this excitement, with links, tools, and advice that may help new physicists make the leap toward LLPs.

## DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

## ACKNOWLEDGMENTS

We would like to sincerely thank Dean Robinson, Simone Pagan Griso, Jessie Shelton, Matthew Citron, Daniele del Re, and Brian Shuve for their very valuable input regarding this review. We are particularly grateful to Jessie Shelton, Federico Leo Redi, and Alberto Escalante del Valle for providing comments on the manuscript. The work of S.K. was supported by the US Department of Energy, Office of Science (contract DE-AC02-05CH11231). Part of this work was performed at the Aspen Center for Physics, which is supported by the National Science Foundation (grant PHY-1607611). S.L. would like to acknowledge 2021 sabbatical leave support from Vrije Universiteit Brussel and the support of FWO Vlaanderen under the Excellence of Science (EOS) be.h project 3082081.

## LITERATURE CITED

1. Neddermeyer SH, Anderson CD. *Phys. Rev.* 51:884 (1937)
2. Rochester GD, Butler CC. *Nature* 160:855 (1947)
3. Abreu P, et al. (DELPHI Collab.) *Phys. Lett. B* 444:491 (1998)
4. Heister A, et al. (ALEPH Collab.) *Eur. Phys. J. C* 25:339 (2002)
5. Aaltonen T, et al. (CDF Collab.) *Phys. Rev. D* 78:032015 (2008)
6. Abazov VM, et al. (D0 Collab.) *Phys. Rev. Lett.* 97:161802 (2006)
7. Strassler MJ, Zurek KM. *Phys. Lett. B* 651:374 (2007)
8. Lee L, Ohm C, Soffer A, Yu T-T. *Prog. Part. Nucl. Phys.* 106:210 (2019)
9. Curtin D, et al. *Rep. Prog. Phys.* 82:116201 (2019)
10. Arkani-Hamed N, Dimopoulos S. *J. High Energy Phys.* 0506:073 (2005)
11. Abdullahi AM, et al. *J. Phys. G* 50:020501 (2023)

12. Craig N, Katz A, Strassler M, Sundrum R. *J. High Energy Phys.* 1507:105 (2015)
13. Fradette A, Pospelov M. *Phys. Rev. D* 96:075033 (2017)
14. Knapen S, Shelton J, Xu D. *Phys. Rev. D* 103:115013 (2021)
15. Sher M. *Phys. Rev. D* 52:3136 (1995)
16. Griest K, Seckel D. *Phys. Rev. D* 43:3191 (1991)
17. Bai Y, Tait TMP. *Phys. Lett. B* 710:335 (2012)
18. Izaguirre E, Krnjaic G, Shuve B. *Phys. Rev. D* 93:063523 (2016)
19. Cui Y, Sundrum R. *Phys. Rev. D* 87:116013 (2013)
20. Cui Y, Shuve B. *J. High Energy Phys.* 1502:49 (2015)
21. Chatrchyan S, et al. (CMS Collab.) *J. Instrum.* 9:P10009 (2014)
22. ATLAS Collab. ATLAS Note ATL-PHYS-PUB-2015-018, CERN, Geneva
23. Aaboud M, et al. (ATLAS Collab.) *Phys. Rev. D* 97:052012 (2018)
24. Aad G, et al. (ATLAS Collab.) *J. High Energy Phys.* 2306:200 (2023)
25. Sirunyan AM, et al. (CMS Collab.) *Phys. Rev. D* 104:012015 (2021)
26. Aaij R, et al. (LHCb Collab.) *Eur. Phys. J. C* 82:373 (2022)
27. Tumasyan A, et al. (CMS Collab.) *J. High Energy Phys.* 2204:62 (2022)
28. Workman RL, et al. (Part. Data Group) *PTEP* 2022:083C01 (2022)
29. Aaboud M, et al. (ATLAS Collab.) *Phys. Lett. B* 788:96 (2019)
30. Aad G, et al. (ATLAS Collab.) *J. High Energy Phys.* 2306:158 (2023)
31. Khachatryan V, et al. (CMS Collab.) *Phys. Rev. D* 94:112004 (2016)
32. Aaboud M, et al. (ATLAS Collab.) *Phys. Rev. D* 99:052003 (2019)
33. Aad G, et al. (ATLAS Collab.) *Phys. Rev. Lett.* 124:031802 (2020)
34. Allison J, et al. *Nucl. Instrum. Methods Phys. Res. A* 835:186 (2016)
35. Chiodini G, Spagnolo S. (ATLAS Collab.) *J. Instrum.* 8:T02004 (2013)
36. Chatrchyan S, et al. (CMS Collab.) *J. Instrum.* 8:P11002 (2013)
37. Aad G, et al. (ATLAS Collab.) *Phys. Lett. B* 703:428 (2011)
38. del Re D. *J. Phys. Conf. Ser.* 587:012003 (2015)
39. Mahon DJ. *J. Instrum.* 15:C06045 (2020)
40. Chatrchyan S, et al. (CMS Collab.) *J. Instrum.* 8:P09009 (2013)
41. Aaboud M, et al. (ATLAS Collab.) *Eur. Phys. J. C* 78:987 (2018)
42. Chatrchyan S, et al. (CMS Collab.) *J. Instrum.* 5:T03013 (2010)
43. Sirunyan AM, et al. (CMS Collab.) *Phys. Lett. B* 797:134876 (2019)
44. Aad G, et al. (ATLAS Collab.) *Eur. Phys. J. C* 77:490 (2017)
45. Isik C. *Nucl. Instrum. Methods Phys. Res. A* 1042:167389 (2022)
46. ATLAS Collab. arXiv:2209.01029 [hep-ex] (2022)
47. Sirunyan AM, et al. (CMS Collab.) *Phys. Rev. D* 100:112003 (2019)
48. Tumasyan A, et al. (CMS Collab.) *Phys. Rev. Lett.* 127:261804 (2021)
49. Aad G, et al. (ATLAS Collab.) *Phys. Rev. D* 106:032005 (2022)
50. Aaij R, et al. (LHCb Collab.) *Eur. Phys. J. C* 75:595 (2015)
51. Cid Vidal X, Tsai Y, Zurita J. *J. High Energy Phys.* 2001:115 (2020)
52. CMS Collab. Technical Design Report CERN-LHCC-2019-003/CMS-TDR-020, CERN, Geneva (2019)
53. ATLAS Collab. Technical Design Report CERN-LHCC-2020-007/ATLAS-TDR-031, CERN, Geneva (2020)
54. ATLAS Collab. Technical Design Report CERN-LHCC-2017-020/ATLAS-TDR-029, CERN, Geneva (2017)
55. CMS Collab. Technical Design Report CERN-LHCC-2021-007/CMS-TDR-022, CERN, Geneva (2021)
56. Alwall J, et al. *J. High Energy Phys.* 1407:79 (2014)
57. Sjöstrand T, et al. *Comput. Phys. Commun.* 191:159 (2015)
58. Evans JA, Gandrakota A, Knapen S, Routray H. *Phys. Rev. D* 103:015026 (2021)
59. Aad G, et al. (ATLAS Collab.) *Phys. Rev. D* 102:032006 (2020)

60. Tumasyan A, et al. (CMS Collab.) *J. High Energy Phys.* 2305:228 (2023)
61. Aaboud M, et al. (ATLAS Collab.) *Eur. Phys. J. C* 79:481 (2019)
62. Aad G, et al. (ATLAS Collab.) *J. Instrum.* 8:P07015 (2013)
63. Alimena J, et al. arXiv:2110.14675 [hep-ex] (2021)
64. Aaij R, et al. (LHCb Collab.) *Phys. Rev. Lett.* 124:041801 (2020)
65. Aaij R, et al. (LHCb Collab.) *J. High Energy Phys.* 2010:156 (2020)
66. LHCb Collab. Technical Design Report CERN-LHCC-2018-007/LHCb-TDR-017, CERN, Geneva (2018)
67. Borsato M, et al. *Rep. Prog. Phys.* 85:024201 (2022)
68. CMS Collab. CMS Detector Performance Summary CMS-DP-2019-043, CERN, Geneva (2019)
69. Aad G, et al. (ATLAS Collab.) *J. Instrum.* 3:S08003 (2008)
70. Chatrchyan S, et al. (CMS Collab.) *J. Instrum.* 3:S08004 (2008)
71. ATLAS Collab. ATLAS PUB Note ATL-PHYS-PUB-2017-014, CERN, Geneva (2017)
72. Alves AA Jr., et al. (LHCb Collab.) *J. Instrum.* 3:S08005 (2008)
73. Aaij R, et al. (LHCb Collab.) *J. Instrum.* 10:P02007 (2015)
74. Calefice L, et al. *Front. Big Data* 5:1008737 (2022)
75. CMS Collab. CMS Detector Performance Summary CMS-DP-2017-015, CERN, Geneva (2017)
76. Chatrchyan S, et al. (CMS Collab.) *Phys. Rev. D* 88:052001 (2013)
77. Aad G, et al. (ATLAS Collab.) *Phys. Rev. D* 85:012001 (2012)
78. Aaij R, et al. (LHCb Collab.) *J. High Energy Phys.* 1108:34 (2011)
79. Thomas SD, Wells JD. *Phys. Rev. Lett.* 81:34 (1998)
80. Dherghetta T, Giudice GF, Wells JD. *Nucl. Phys. B* 559:27 (1999)
81. Dimopoulos S, Dine M, Raby S, Thomas SD. *Phys. Rev. Lett.* 76:3494 (1996)
82. Aad G, et al. (ATLAS Collab.) *Eur. Phys. J. C* 82:606 (2022)
83. Sirunyan AM, et al. (CMS Collab.) *Phys. Lett. B* 806:135502 (2020)
84. Evans JA, Shelton J. *J. High Energy Phys.* 1604:56 (2016)
85. Aaboud M, et al. (ATLAS Collab.) *Phys. Rev. D* 99:012001 (2019)
86. Tumasyan A, et al. (CMS Collab.) *J. High Energy Phys.* 2207:81 (2022)
87. Tumasyan A, et al. (CMS Collab.) *Eur. Phys. J. C* 82:153 (2022)
88. Aad G, et al. (ATLAS Collab.) *Phys. Rev. Lett.* 127:051802 (2021)
89. Sirunyan AM, et al. (CMS Collab.) *J. High Energy Phys.* 1902:179 (2019)
90. Kang J, Luty MA. *J. High Energy Phys.* 0911:065 (2009)
91. Knapen S, Lou HK, Papucci M, Setford J. *Phys. Rev. D* 96:115015 (2017)
92. Sirunyan AM, et al. (CMS Collab.) *Mach. Learn. Sci. Technol.* 1:035012 (2020)
93. Aad G, et al. (ATLAS Collab.) *J. High Energy Phys.* 2206:5 (2022)
94. ATLAS Collab. Technical Report ATL-PHYS-PUB-2021-012, CERN, Geneva (2021)
95. Aaij R, et al. (LHCb Collab.) *Eur. Phys. J. C* 81:261 (2021)
96. Aaij R, et al. (LHCb Collab.) *Eur. Phys. J. C* 77:812 (2017)
97. Tumasyan A, et al. (CMS Collab.) *Eur. Phys. J. C* 82:213 (2022)
98. CMS Collab. CMS Physics Analysis Summary CMS-PAS-TRK-10-003, CERN, Geneva (2010)
99. Alexander M, et al. *J. Instrum.* 13:P06008 (2018)
100. Aad G, et al. (ATLAS Collab.) *Phys. Rev. D* 92:012010 (2015)
101. ATLAS Collab. Public pixel tracker plots for collision data. CERN. [https://twiki.cern.ch/twiki/bin/view/AtlasPublic/PixelPublicResults#PHYS\\_and\\_INDET\\_PUB\\_NOTES](https://twiki.cern.ch/twiki/bin/view/AtlasPublic/PixelPublicResults#PHYS_and_INDET_PUB_NOTES) (2023)
102. Gershtein Y. *Phys. Rev. D* 96:035027 (2017)
103. Gershtein Y, Knapen S. *Phys. Rev. D* 101:032003 (2020)
104. Hook A, Kumar S, Liu Z, Sundrum R. *Phys. Rev. Lett.* 124:221801 (2020)
105. Drollinger V. CMS Note CMS-NOTE-2005-012, CERN, Geneva (2005)
106. Sirunyan AM, et al. (CMS Collab.) *J. High Energy Phys.* 1710:73 (2017)
107. Aad G, et al. (ATLAS Collab.) *J. High Energy Phys.* 2306:153 (2023)
108. Feng JL, Galon I, Kling F, Trojanowski S. *Phys. Rev. D* 97:035001 (2018)
109. Feng JL, et al. *J. Phys. G* 50:030501 (2023)

110. Cerci S, et al. *J. High Energy Phys.* 2206:110 (2022)
111. Ahdida C, et al. (SHiP Collab.) arXiv:2002.08722 [physics.ins-det] (2020)
112. Boyarsky A, Mikulenko O, Ovchinnikov M, Shchutska L. *J. High Energy Phys.* 2203:6 (2022)
113. Pierog T, et al. *Phys. Rev. C* 92:034906 (2015)
114. Acharya B, et al. (MoEDAL Collab.) *Int. J. Mod. Phys. A* 29:1430050 (2014)
115. Haas A, Hill CS, Izaguirre E, Yavin I. *Phys. Lett. B* 746:117 (2015)
116. Acharya B, et al. (MoEDAL Collab.) *Phys. Lett. B* 782:510 (2018)
117. Acharya B, et al. (MoEDAL-MAPP Collab.) arXiv:2209.03988 [hep-ph] (2022)
118. Ball A, et al. (milliQan Collab.) *Phys. Rev. D* 104:032002 (2021)
119. Gligorov VV, Knapen S, Papucci M, Robinson DJ. *Phys. Rev. D* 97:015023 (2018)
120. Aielli G, et al. *Eur. Phys. J. C* 80:1177 (2020)
121. Chou JP, Curtin D, Lubatti HJ. *Phys. Lett. B* 767:29 (2017)
122. Alpigiani C, et al. (MATHUSLA Collab.) arXiv:1811.00927 [physics.ins-det] (2018)
123. Gligorov VV, et al. *Phys. Rev. D* 99:015023 (2019)



# Contents

Lepton Flavor Violation and Lepton Flavor Universality Violation in $b$ and $c$ Decays <i>Diego Guadagnoli and Patrick Koppenburg</i> .....	1
New Solutions to the Gauge Hierarchy Problem <i>Anson Hook</i> .....	23
COHERENT at the Spallation Neutron Source <i>P.S. Barbeau, Yu. Efremenko, and K. Scholberg</i> .....	41
Experimental Considerations in Long-Baseline Neutrino Oscillation Measurements <i>Francesca Di Lodovico, Ryan B. Patterson, Masato Shiozawa, and Elizabeth Worcester</i> .....	69
Detection and Calibration of Low-Energy Nuclear Recoils for Dark Matter and Neutrino Scattering Experiments <i>Jingke Xu, P.S. Barbeau, and Ziqing Hong</i> .....	95
Implications of Large- $N_c$ QCD for the $NN$ Interaction <i>Thomas R. Richardson, Matthias R. Schindler, and Roxanne P. Springer</i> .....	123
The Hubble Tension and Early Dark Energy <i>Marc Kamionkowski and Adam G. Riess</i> .....	153
High-Energy to Ultrahigh-Energy Neutrino Interactions <i>Mary Hall Reno</i> .....	181
Medium Response and Jet–Hadron Correlations in Relativistic Heavy-Ion Collisions <i>Shanshan Cao and Guang-You Qin</i> .....	205
Boson–Boson Interactions at the LHC <i>J. Manjarrés Ramos and Guillermo Gómez-Ceballos</i> .....	231
Physics of the Top Quark at the LHC: An Appraisal and Outlook of the Road Ahead <i>P. Ferreira da Silva</i> .....	255
Recent Progress in Leptonic and Semileptonic Decays of Charmed Hadrons <i>Bai-Cian Ke, Jonna Koponen, Hai-Bo Li, and Yangbeng Zheng</i> .....	285

The $s$ Process and Beyond <i>Maria Lugaro, Marco Pignatari, René Reifarth, and Michael Wiescher</i> .....	315
Ultra-High-Energy Gamma-Ray Astronomy <i>Zhen Cao, Songzhan Chen, Ruoyu Liu, and Ruizhi Yang</i> .....	341
Deep-Sea and Lunar Radioisotopes from Nearby Astrophysical Explosions <i>Brian D. Fields and Anton Wallner</i> .....	365
Physics Beyond the Standard Model Associated with the Top Quark <i>Roberto Franceschini</i> .....	397
A Guide to Hunting Long-Lived Particles at the LHC <i>Simon Knapen and Steven Lowette</i> .....	421

## Errata

An online log of corrections to *Annual Review of Nuclear and Particle Science* articles may be found at <http://www.annualreviews.org/errata/nucl>