Analysis of long-lived particle decays with the MATHUSLA detector

David Curtin
Maryland Center for Fundamental Physics, Department of Physics, University of Maryland, College Park, Maryland 20742, USA

Michael E. Peskin
SLAC, Stanford University, Menlo Park, California 94025, USA

(Received 23 May 2017; published 8 January 2018)

The MATHUSLA detector is a simple large-volume tracking detector to be located on the surface above one of the general-purpose experiments at the Large Hadron Collider. This detector was proposed in [J. P. Chou, D. Curtin, and H. J. Lubatti, Phys. Lett. B **767**, 29 (2017)] to detect exotic, neutral, long-lived particles that might be produced in high-energy proton-proton collisions. In this paper, we consider the use of the limited information that MATHULSA would provide on the decay products of the long-lived particle. For the case in which the long-lived particle is pair-produced in Higgs boson decays, we show that it is possible to measure the mass of this particle and determine the dominant decay mode with less than 100 observed events. We discuss the ability of MATHUSLA to distinguish the production mode of the long-lived particle and to determine its mass and spin in more general cases.

DOI: 10.1103/PhysRevD.97.015006

I. INTRODUCTION

Despite the successes of the standard model (SM) of particle physics, there are strong motivations to believe in new fundamental interactions that lie outside this model. The standard model does not contain a particle that could explain the dark matter of the Universe. Its theory of the Higgs boson and its symmetry-breaking potential is completely *ad hoc*. Many models have been proposed to generalize the standard model, but there is no compelling experimental evidence supporting any of these models. It is therefore important to propose additional windows through which to search for these new interactions.

A property shared by many models of supersymmetry [1–6], neutral naturalness [7–9], dark matter [10–15], baryogenesis [16–21], neutrinos [22–31], and hidden valleys [32–37] is that they contain long-lived particles (LLPs) with macroscopic decay lengths.

Searches for LLPs have typically involved the study of low-energy reactions, for example, using fixed target experiments with electron, proton, or neutrino beams. One strategy has been to position a detector behind a beam dump, where it can observe decays of neutral particles with weak interaction cross sections on matter. However, this approach to the search of LLPs is limited in mass scale. It is also limited because it requires the LLP to have large enough coupling to quarks and leptons.

The Large Hadron Collider (LHC) offers new mechanisms for the production of LLPs that are available only in high-energy collisions. These include production through W boson fusion, through the decay of heavy SM particles like the Higgs boson or Z, through the decay of new heavy parent particles such as squarks or gluinos, and through new, heavy scalar and vector bosons produced in the s-channel in quark-antiquark or gluon-gluon collisions. The most interesting and most highly motivated of these mechanisms is the exotic decay of the 125 GeV Higgs boson to a pair of LLPs [38–41]. However, though the LHC might have large production rates for LLPs, the ability of the LHC detectors to observe these particles is limited. As large as ATLAS and CMS are, the size of these detectors is a constraint. Furthermore, LLP events suffer from significant backgrounds, especially if the LLPs decay to hadrons [42].

The MATHUSLA detector was proposed in [43] to address this problem [44]. MATHUSLA is a large-volume detector on the surface above an LHC experiment. Essentially, it is an empty barn that provides a decay volume for LLPs, and, near its roof, is equipped with charged particle tracking to detect an LLP decay. It is shown in [43] that the limited instrumentation proposed allows one to reject cosmic-ray and other backgrounds with very high confidence. This dedicated detector would increase the sensitivity to LLPs over the capabilities of the current central detectors by several orders of magnitude. The comparison to ATLAS is shown in Fig. 1.

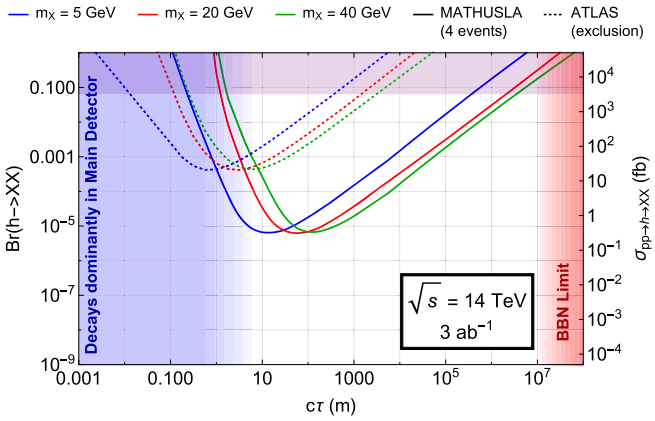

FIG. 1. Exclusion reach of MATHUSLA, corresponding to four expected decays in the detector (solid curves), compared to the best-case ATLAS projection (dotted curves), for pair production of LLPs in exotic Higgs decays $h \to XX$, from [43]. The three curves of each set correspond to three different values of the LLP mass. Sensitivity up to the big bang nucleosynthesis lifetime limit [45] is possible.

Because of its large size and because—as yet—there is no evidence for LLPs, the MATHUSLA detector must be built from relatively inexpensive components. The original concept for MATHUSLA in [43] imagined an empty building offering 20 m of decay space and, above this, ∼5 layers of resistive plate chambers (RPCs), along with some plastic scintillator for additional timing and veto information. This paper offered an explicit physics case for MATHUSLA, with estimates of its performance in the search for LLPs produced in exotic Higgs decays as a well-motivated benchmark model [46].

From this description, it is not obvious that MATHUSLA has any capability beyond the discovery of LLP events via the detection of decay vertices originating in its decay volume. However, we find that, by applying some simple arguments, it is possible to use the limited information provided by MATHUSLA to learn a surprising amount. In this paper, we analyze the performance of MATHUSLA for the most interesting and also most constrained situation—the decay of the Higgs boson to a pair of LLPs, such that the LLP has a dominant two-body decay mode. In Sec. II, we briefly review the design of MATHUSLA. In Sec. III, we show that, under the assumption of this production mode, it is possible to measure the mass of the LLP and to identify its most important decay modes, using only the information provided by MATHUSLA, with as few as 30–100 observed decays. In Sec. IV, we show how the production by Higgs decay may be distinguished from other hypotheses, and we discuss the generalization of this analysis to other LLP production modes.

II. DESIGN OF THE MATHUSLA DETECTOR

For concreteness, we define a simple design for the MATHUSLA detector that we use in our study. This closely follows the concept originally presented in [43], with one suggested modification for additional diagnostic capability.

The detector geometry relative to the LHC interaction point is shown in Fig. 2. MATHUSLA is an empty building of area 40 000 m^2 and height ∼25 m. The floor, ceiling, and walls contain a layer of scintillator to provide a veto for charged particles emerging from below. The veto is important in dealing with backgrounds, which, as is shown in [43], can be reduced to negligible levels. However, it plays no role in the analysis we present here.

At a height of 20 m, we place the first of five RPC tracking layers. These layers are spaced about 1 m apart, with the last layer just below the roof. The RPCs are arranged to record charged particle hits with a pixel size of 1 cm^2 and a time resolution of 1 ns. This allows the angle of charged tracks to the tracking planes to be determined with a precision of about 2 mrad.

On the right side of Fig. 2, we show schematically the pattern of charged tracks reconstructed for different final states of LLP decay. Muons show up as single tracks, taus as one or three collimated tracks, and jets as many tracks populating a relatively large solid angle.

Without additional material, electrons and muons may not be distinguishable, while photons are invisible. We therefore suggest a possible modification to the original design of [43] by inserting an uninstrumented steel sheet of several cm thickness between the first and second tracking layer. This provides 1–2 X_0 to convert photons and

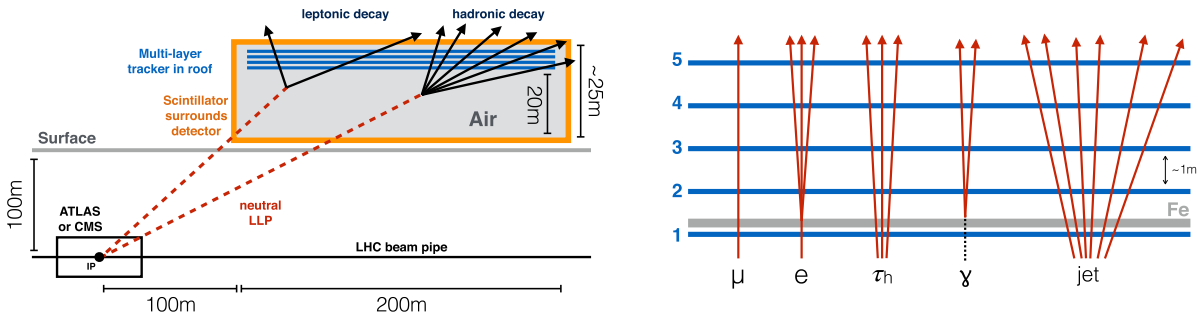

FIG. 2. We assume the MATHUSLA detector geometry shown on the left. On the right, we schematically show the patterns of charged tracks reconstructed for different final states of LLP decay.

electrons, producing visible electromagnetic showers. The thickness of the sheet for hadronic interactions is about 0.1 λ_i. This minimal detector does not allow measurement of the energy or momentum for any particles, but it does allow the various particle types to be distinguished qualitatively. The details of this possible modification, including the exact thickness, type, and location of material, as well as its viability in terms of cost and effect on tracking performance, are left for future investigation.

In our simulations, we assume the following minimum detection thresholds on particle three-momenta to ensure that the particles leave hits in all tracking layers: pions, 200 MeV; charged kaons, 600 MeV; muons, 200 MeV; electrons, 1 GeV; protons, 600 GeV; photons, 200 MeV. We ignore charged pion and kaon decays within the detector volume.

If the Higgs boson decays to a pair of LLPs with a branching ratio of 1%, the ATLAS or CMS detector will produce 1 500 000 pairs of LLPs with the 3 ab^{-1} of luminosity projected for the high-luminosity LHC. Assuming the best case of a ~100 m lifetime, we expect a sample of about 10 000 LLP decays within the detector acceptance. We see below that it is possible to draw interesting conclusions from samples with as few as 100 LPP events.

III. DIAGNOSING LLPS PRODUCED IN EXOTIC HIGGS DECAYS

In this section, we assume that LLPs are produced in pairs as the decay products of the Higgs boson. Our objective is to measure the LLP mass and determine the dominant decay modes, using only geometrical charged particle trajectories that can be measured by MATHUSLA. We postpone to the next section the problem of distinguishing this production mechanism from other possibilities.

To open up the largest possible number of decay final states, we study LLP masses above the $\bar{b}b$ threshold, $m_X \in (15, 55)$ GeV. The method is easily generalized to lower masses, though spatial track resolution may become more important for very light LLPs. We assume possible decays $X \to ee, \mu\mu, \tau\tau, \gamma\gamma$ or jj. We write $\tau_{h,\ell}$ to refer to an explicitly hadronic or leptonic τ.

In the simulations described below, we model Higgs production via gluon fusion at the HL-LHC, with subsequent decay to two LLPs XX and decay of one X in the MATHUSLA detector, by using the Hidden Abelian Higgs model [39] in MadGraph 5 [47], matched up to one extra jet, and showered in Pythia 8.162 [48,49]. In this model, the LLP X is modeled as a spin-0 particle, except in the case of "gauge-ordered" two-jet decay, described below, where it is spin 1. Only LLPs with an angle to the beam axis in the range [0.3, 0.8], corresponding approximately to the angular coverage of MATHUSLA, are analyzed [50].

A. Qualitative analysis

We have already illustrated in the previous section and in Fig. 2 that the various possibilities for the two-body decay of the X can be distinguished qualitatively from the pattern of tracks in the MATHUSLA detector. After requiring at least two detected charged tracks per displaced vertex, we can impose the following criteria to sort the events:

(i) two tracks, $\mu\mu$;
(ii) two tracks, which shower after the first layer, ee;
(iii) two showers but no hits in the first layer, $\gamma\gamma$;
(iv) between three and six tracks, at least partially hadronic $\tau_h\tau_{h,\ell}$;
(v) more than six tracks, jj.

Without the material layer, photons are undetectable and electrons look like muons, but all of our other conclusions are unaffected. Decays to $\tau^+\tau^-$ are recognized by the characteristic one-prong against three-prong topology that appears in 26% of $\tau^+\tau^-$ decays. The identification of subdominant decays modes and the measurement of their branching ratios depends strongly on what the dominant decay mode might be. In some cases, there is an obvious analysis using the criteria above. If the dominant decay mode is $b\bar{b}$, this generates backgrounds to other decay models that must be studied with care. The full analysis of that problem is beyond the scope of this paper.

The category of X decays to jj contains a number of more specific possibilities. Three simple benchmark scenarios are (1) decay to gluon jets, (2) gauge-ordered decay to $q\bar{q}$ jets with democratic flavor content, as would be generated by the decay of a dark photon [39], and (3) "Yukawa-ordered" decay to jets that are dominantly $b\bar{b}$, as would be generated by the decay of a dark singlet scalar that mixes with the SM Higgs boson. These possibilities cannot be distinguished on an event-by-event basis, but we show below that they can be distinguished in samples as small as 100 events.

B. Measurement of the LLP mass

Now we discuss the determination of the LLP mass. It is crucial that the decay vertex can be precisely located within the MATHUSLA decay volume. Since the LLP X originates from the nearby LHC collision region, the vector from the point of origin to the decay vertex is very well known. This allows the velocity β_X of the LLP to be found from the geometry of the decay.

Consider first a decay to two final-state charged particles, such as ee or $\mu\mu$. Let θ_1 and θ_2 be the angles of the two decay products with respect to the X direction, as shown in Fig. 3. The 4-vectors of the two products then have the form

$$p_i = E_i(1, \pm\beta_i \sin\theta_i, 0, \beta_i \cos\theta_i), \qquad i = 1, 2, \quad (1)$$

with θ_1 and θ_2 both being positive quantities and $E_1\beta_1 \sin\theta_1 = E_2\beta_2 \sin\theta_2$ by momentum balance. Since all components are known up to an overall prefactor, we can boost both p_i back along the direction of p_X until they are back to back, recovering the LLP rest frame. This yields

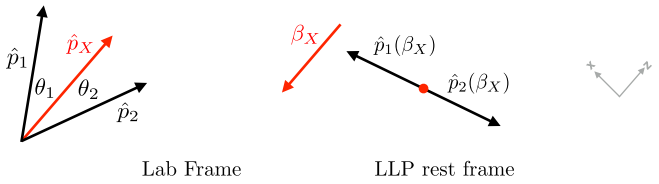

FIG. 3. Kinematics of the two-body decay of an LLP: The left-hand figure shows the angles θ_1, θ_2; the right-hand figure illustrates the boost back to the LLP rest frame in which the two products are back to back. Note that $\hat{p}_i$ or $\hat{p}_i(\beta_X)$ denote momentum vectors normalized to unit length *in each frame*. For convenience we work in the coordinate system where $\hat{p}_X$ is along the z-axis and the decay products are in the (x, z) plane (far right).

$$\beta_X = \frac{\beta_1\beta_2 \sin(\theta_1 + \theta_2)}{\beta_1 \sin\theta_1 + \beta_2 \sin\theta_2}. \tag{2}$$

Since the distance of the LLP decay to the LHC interaction point is much greater than the distance to the tracking planes, the precision of the measured angles θ_1, θ_2 is simply the precision of the measured angles between the tracks and the trackers, about 0.2% for $\theta_i \sim \mathcal{O}(1)$ and approximately independent of the uncertainty on the displaced vertex location. For the two-body decays we consider, the products are relativistic, with β_i close to 1. This makes the error induced by assuming that $\beta_i = 1$ negligible. In any case, the timing of the MATHULSA detector tracking elements allows each β_i to be measured to 5% or better.

For $h \to XX$, the transverse energy of the X with respect to the LHC beam direction should be roughly $m_h/2$, so the expected mean velocity of the produced X particles decreases as the mass increases. In Fig. 4, we show the expected distribution of $b = p_X/m_X$ from our simulation for three values of the X mass, illustrating this effect. From the figure, we estimate that a sample of 100 reconstructed events gives the X mass with a statistical error of about 1 GeV. The systematic error on this measurement, coming from the uncertainties in the measurement of lepton directions, is at the part-per-mil level.

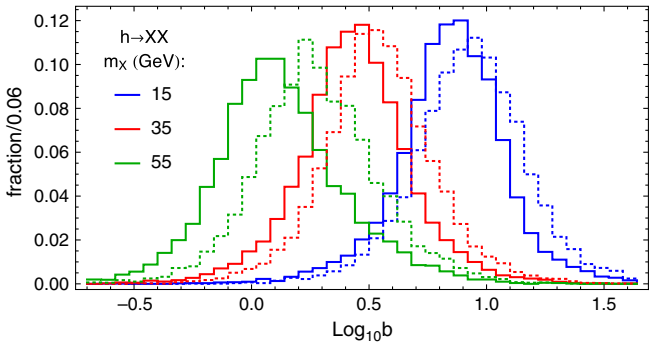

FIG. 4. Distribution of LLP boost $b = |\vec{p}_X|/m$ for different LLP masses. The solid histograms show the truth-level value of b, which is also close to the distribution of reconstructed boosts for X decay to two charged particles. The dotted histograms show the distribution of reconstructed boosts for hadronic LLP decays using the sphericity-based method of Sec. III D with only upwards going tracks.

The precise knowledge of the LLP speed in each event gives an error on the production time of a few ns, making it possible to identify the LHC bunch crossing in which the LLP was produced. This means that the event properties measured in the central detector can potentially be used to constrain hypotheses on the LLP production process.

This simple method of determining β_X fails for hadronic LLP decays, since jet axes cannot be reliably reconstructed from charged particle directions alone. Fortunately, we can achieve almost identical results using an only slightly more sophisticated method that we outline in Sec. III D.

We now discuss this mass measurement more quantitatively for the three cases of $X \to \mu\mu$, $X \to jj$, and $X \to \tau\tau$.

C. LLP decay to $\mu\mu$

For $\mu\mu$ decays, Eq. (2) gives the reconstructed LLP boost as long as both muons hit the roof of the detector. This is the case in about 95% (50%) of decays in MATHUSLA for $m_X = 15\,(55)$ GeV. This geometrical effect is the dominant factor in the efficiency ϵ for an event to be reconstructed by the MATHUSLA detector.

To determine the expected precision of a mass measurement with N reconstructed LLP decays, we conducted 1000 pseudoexperiments and made a maximum likelihood fit of the measured boost distribution to template functions obtained from the same boost distributions in the maximum-statistics limit. For a given pseudoexperiment, we define N_{obs} to be the number of decays in the MATHUSLA detector volume and $N_{\text{reconstructed}}$ to be the number of decays in which the tracks are oriented such that mass can be computed from the available information. The reconstruction efficiency $\epsilon = N_{\text{reconstructed}}/N_{\text{obs}}$ varies from 0.95 for $m_X = 15$ GeV to 0.55 for $m_X = 55$ GeV. The distributions of the reconstructed boosts are very close to the truth-level distributions shown in Fig. 4. We define the expected mass precision, $\langle \Delta m/m \rangle$, to be the average spread of best-fit mass values amongst the 1000 pseudoexperiments. In Fig. 5, we show the dependence of this quantity on the total number of LLPs decaying in the MATHUSLA detector volume. A 10% mass measurement requires only 20–30 observed decays.

For few reconstructed events, the precision of the mass measurement is better for heavier LLPs, while for many reconstructed events, it is better for lighter LLPs. This is not an artifact of the mass-dependent ϵ in Fig. 5, but is likely due to the fact that under the assumptions of our LLP production mode in Higgs decays, the LLP mass is bounded from above. Therefore, for very few observed events, measuring a handful of very low boosts has to be due to LLPs near threshold. As the number of reconstructed events becomes large, these parameter space "edge effects" become less important, and the slightly narrower boost distribution of light LLPs makes their mass measurement more precise.

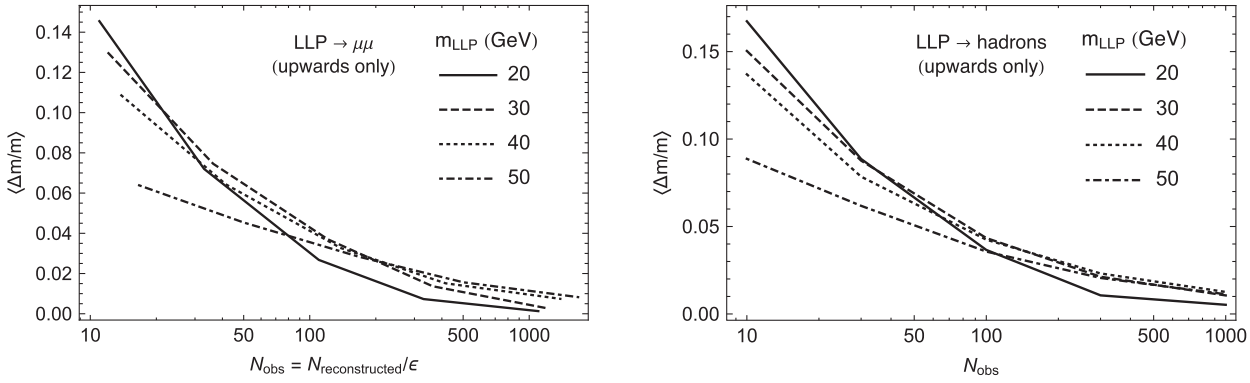

FIG. 5. The expected number of LLP decays in MATHUSLA required to measure the LLP mass to relative precision $\Delta m/m$ using only upwards-traveling tracks, for LLP decay to muons (left) and jets (right). $N_{\rm obs}$ refers to the number of decays in the detector volume regardless of whether the event is reconstructed. For muons, $N_{\rm reconstructed} = \epsilon N_{\rm obs}$ where $\epsilon \approx 0.95\,(0.55)$ for $m_X = 15\,(55)$ GeV. For hadrons, $\epsilon \approx 1$. The plots show the statistical uncertainty only.

D. LLP decay to jets

In our mass range of interest, LLP decays to jets produce events with 10–20 charged tracks in the detector. This is illustrated in Fig. 6 for $m_X = 15$ and 55 GeV and our three benchmark jet flavor compositions. This high multiplicity is a boon for several reasons. In terms of background rejection, a displaced vertex with full timing information and this many tracks is supremely difficult to fake by cosmic rays. In terms of signal analysis, the multiplicity distribution contains information about the jet flavor composition. Furthermore, the detection efficiency for LLPs decaying to jets (defined as the fraction of decays with at least six charged tracks hitting the roof) is very close to 100%.

On the other hand, the high multiplicity means it is not so simple to determine the jet directions that must be input into Eq. (2). One might, for example, try to extract two jet axes $\hat{p}_a, \hat{p}_b$ by maximizing the quantity

$$V_2 = \sum_i \max(\hat{p}_a \cdot \hat{p}_i, \hat{p}_b \cdot \hat{p}_i) \tag{3}$$

summed over charged track momentum unit vectors $\hat{p}_i$ in the event. The sum tends to be dominated by tracks carrying low fractions of the total jet momentum. In the absence of energy and momentum measurements, the jet axes cannot be reliably determined by this or similar methods.

Fortunately, we can exploit the high multiplicity for a different kind of approximate boost reconstruction. Naively, if the LLP decays to high multiplicity, the distribution of tracks in its rest frame should be spherically symmetric. Applying this assumption to individual events, we can estimate the LLP boost event by event by solving for β_X in the constraint

$$\hat{p}_X \cdot \sum_i \hat{p}_i(\beta_X) = 0. \tag{4}$$

The resulting sphericity-based boost distribution, using only upward-going tracks and assuming all final states to be ultrarelativistic in the lab frame, is shown as the dotted distributions in Fig. 4. In our simulation of LLP decays to jj, this method is surprisingly powerful. It gives a boost distribution very close to the original boost distribution from Monte Carlo truth. Even more importantly, its discriminating power for LLP mass is almost unaffected by the deviation between these distributions.

There are several reasons why the sphericity-based method might give an accurate result. If the parent of the LLP has spin 0 and CP violation can be ignored in its

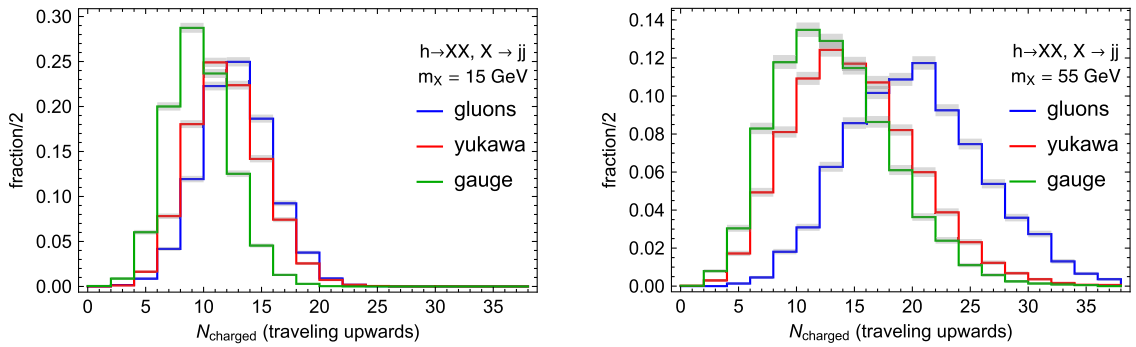

FIG. 6. Distribution of charged particle multiplicity (only counting upwards traveling tracks) for $m_X = 15, 55$ GeV and gluon jets and for flavor-democratic (gauge-ordered) and heavy-flavor dominated (Yukawa-ordered) quark jets.

decays, the distribution of tracks in the LLP rest frame will be front-to-back symmetric on average. The same result applies if the parent has spin but is produced with zero longitudinal polarization along the LLP direction.

For the decay to jj, though, a more important reason for the accuracy of the sphericity-based method is that the sum in (4) is dominated by hadrons that are soft in the rest frame of the LLP. The momentum distribution of these soft particles depends only on the color flow and is independent of any LLP polarization. Their high multiplicity ensures that the shift of any one sphericity-based boost measurement is much smaller than the width of the overall boost distribution, allowing the LLP mass to be accurately extracted. It should be noted that these same soft hadrons are also partially responsible for the noticeable positive bias of the sphericity-based boost distribution compared to the Monte Carlo truth. That effect deserves a dedicated discussion, which we present in the Appendix.

We can use the sphericity-based method to measure the mass of a LLP decaying as $X \to jj$ without having to determine the jet flavor content first. The event-by-event precision of the sphericity-based β_X measurement is sufficient to determine the LHC bunch crossing in which the LLP was created to within about two (six) bunch crossings for $m_X = 15\,(55)$ GeV.

We can estimate the required number of observed events N_{obs} for a given mass measurement precision in a fashion identical to that used for $X \to \mu\mu$ above. The only difference is the use of sphericity-based boost distributions as the templates. The efficiency for reconstructing these events is close to 1. The required number of observed events is shown in Fig. 5 (right). The result is very similar to that for $X \to \mu\mu$, with only about 20–30 events required for a statistical error on the mass measurement of 10%.

The mass measurement also has a systematic error from hadronization uncertainties in modeling the final state of the LLP decay. Varying Pythia tunes, we find this to be less than 1%, but a full analysis should also investigate the effect of using other generators such as Herwig [51,52] or Sherpa [53].

The different multiplicity of charged final states can be exploited to determine the flavor content of the $X \to jj$ final state. We make use here of the fact that a gluon jet has a higher multiplicity than a b quark jet, which in turn has higher multiplicity than a light quark jet. Although the differences in the multiplicity distributions are not large enough to identify the jet flavor on an event-by-event basis, this becomes an effective discriminator when applied to large enough samples. The effect is robust even taking in account the discrepancies in the predictions of different hadronization schemes [54]. A straightforward generalization of the mass measurement method to a two-dimensional likelihood fit in boost and multiplicity reveals that the different decay modes can be reliably distinguished with about 100 observed LLP decays.

The charged track distribution contains even more information, but it is likely more dependent on the hadronization model and assumed detector capabilities. For example, the minimum charged particle velocity in each event is significantly higher for gauge-ordered jets than for Yukawa-ordered or gluon jets for $m_X = 15$ GeV, while for $m_X = 55$ GeV the gluon jets have a higher fraction of slower particles than both types of quark jets. The angular correlations also contain information about the LLP spin. We have not exploited this property in our analysis, but with further study it would likely improve the diagnosis of hadronically decaying LLPs.

Finally, we point out that even though jet axes are not useful for measuring LLP boost, a rough determination of the LLP decay plane is possible by minimizing

$$\sum_i (a \cdot \hat{p}_i)^2 \tag{5}$$

for choice of plane normal vector $\vec{a}$. The resulting decay plane corresponds to the truth-level expectation up to a deviation angle $\Delta\theta \sim 0.2$–0.5. This is crude, but it could be useful for diagnosing significant invisible components in LLP decays.

E. LLP decay to $\tau\tau$

For $X \to \tau\tau$ events, each τ decay gives one track or three well-collimated tracks. For the events with two one-prong decays, we use the direction of the observed track as a proxy for the τ direction. For events with three-prong decays, maximization of the quantity V_2 in (3) provides good approximations to the two τ directions. Using the two τ vectors estimated in this way, we apply the method of Sec. III B. The fraction of events for which at least two charged particles hit the roof of the detector is about 90% (60%) for $m_X = 15\,(55)$ GeV. We note that the sphericity-based method described in Sec. III C gives slightly better results for the case of a spin-0 LLP; however, it is less robust with respect to the effects of possible LLP polarization. The event-by-event precision of the β_X measurement is sufficient to determine the LHC bunch crossing in which the LLP was created to within about two (four) bunch crossings for $m_X = 15\,(55)$ GeV. The required number of observed events for a given precision of mass measurement is very similar to the cases already presented in Fig. 5.

IV. DETERMINING THE LLP PRODUCTION MODE

The analysis of the previous section made explicit use of the assumption that the LLP is produced in pairs in Higgs decay. However, with enough events and a library of possible production mode hypotheses as templates, it may be possible that the LLP production mode, decay mode, and mass can all be independently determined in a global fit. In Fig. 7, we compare the boost distributions for a

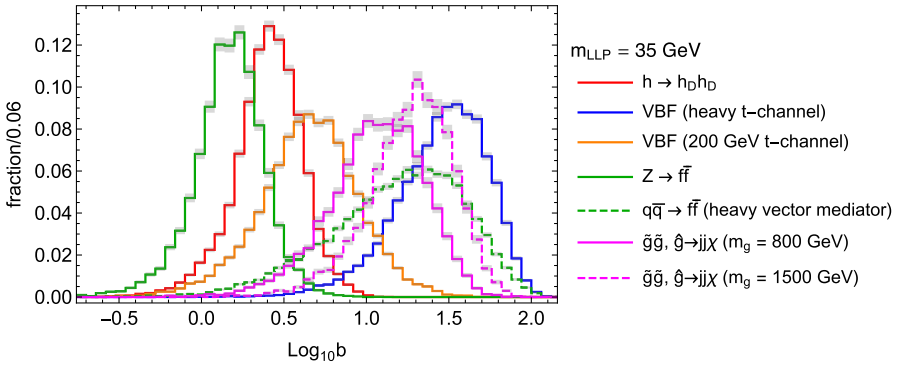

FIG. 7. Truth-level boost distribution of a 35 GeV LLP produced in a variety of different production modes, from left to right: $Z \to XX$, $h \to XX$, vector boson fusion through a 200 GeV mediator in the t channel, decay of gluinos with mass 800 and 1500 GeV, vector boson through a $WW \to XX$ contact interaction, and $q\bar{q} \to XX$ through a vector contact interaction.

35 GeV LLP produced through the following mechanisms at the 14 TeV LHC: $Z \to XX$, $h \to XX$, vector boson fusion through a 200 GeV mediator in the t channel, gluino decay, $q\bar{q} \to XX$ through a vector contact interaction, and vector boson fusion through a $WW \to XX$ contact interaction. For clarity of presentation, we have generated an equal number of events for each sample. These six cases give six shapes with different, distinguishable, features.

The event-by-event boost measurements also allow the LHC bunch crossing in which the LLP was produced to be narrowed down either uniquely or to one of a few choices. Given the low rate per bunch crossing of high momentum transfer events, it is likely that, if the production event is triggered on and recorded, it can be identified and studied. Even if only a small fraction of events were recorded during these bunch crossings, this still might put interesting limits on the energy spectrum of associated objects, distinguishing, for example, the hypotheses of Higgs or Z boson origin from hypotheses involving W fusion or contact interactions.

V. CONCLUSIONS

It is a real possibility that the Higgs boson decays to long-lived particles that couple very weakly to all other particles of the standard model, and that would be invisible to LHC detectors. In [43], a relatively simple large-volume detector was proposed to search for such particles. In this paper, we have explained that this simple detector nevertheless has the power to provide qualitative and even quantitative information about the nature of these long-lived particles. This could well be our first source of information on a new sector of particles that coexist with those of the standard model and open a new dimension into the fundamental interactions.

ACKNOWLEDGMENTS

We thank Henry Lubatti, John Paul Chou, and the other members of the MATHUSLA collaboration for helpful feedback during the completion of this paper. D. C. also thanks Raman Sundrum and Simon Knapen for useful conversations. We thank the referee for encouraging us to more closely examine the sphericity-based boost measurement, including systematic error, bias and spread. The work of D. C. is supported by National Science Foundation Grant No. NSF-PHY-1620074 and the Maryland Center for Fundamental Physics. The work of M. E. P. is supported by the U.S. Department of Energy under Contract No. DE–AC02–76SF00515.

APPENDIX: BIAS AND SPREAD OF THE SPHERICITY-BASED BOOST MEASUREMENT

The sphericity-based boost measurement discussed in Sec. III D and shown in Fig. 4 has three noticeable features.

(1) The width and shape of the sphericity-based distributions are about the same as the truth-level boost distributions.
(2) For lower masses, giving a high-velocity LLP, $\log_{10} b \gtrsim 0.5$, there is a positive bias of $\log_{10} b_{\text{measured}} - \log_{10} b_{\text{truth}} \sim 0.1$ which is approximately mass independent.
(3) For higher masses, giving a low-velocity LLP, the bias is again positive and significantly larger.

These points are important in preserving the sensitivity of the sphericity-based boost measurement to the LLP mass.

To investigate these effects, we found it useful to consider a toy model of LLP decay in which the charged final states are distributed isotropically in the LLP rest frame (without respecting momentum conservation, due to the undetected neutral hadrons). The charged particle multiplicity is sampled from a Poisson distribution centered on $N_{ch} = 10$. It is instructive to consider two possibilities for the charged final state momenta: either all lightlike, or with mass $m = m_\pi$ and energy distributed according to a thermal spectrum, $P(E) \propto \exp[-E/T]$ with $T = 140$ MeV as a crude model of soft pion emission [55].

We used this simple model to generate LLP decay events, boosted them to the lab frame by assuming a fixed

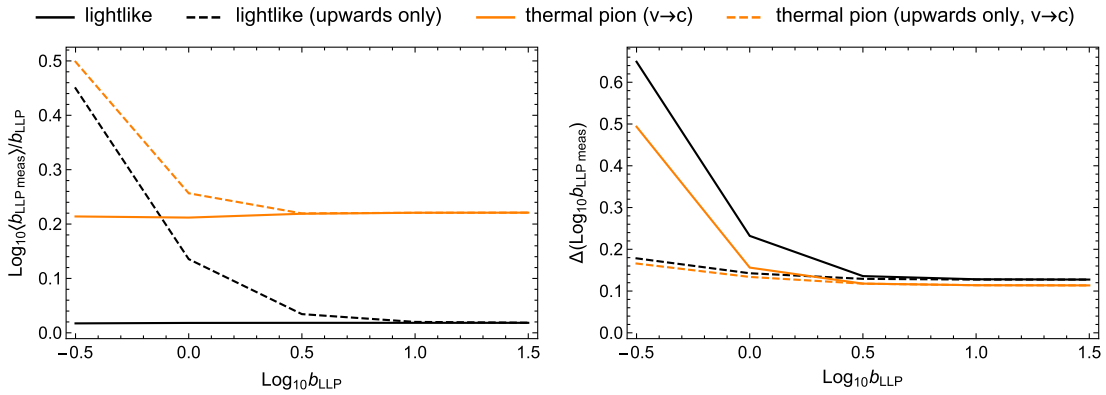

FIG. 8. Bias (left) and spread (right) of the sphericity-based boost measurement for average charged particle multiplicity $N_{\rm ch} = 10$.

LLP boost $b_{\rm LLP}$, and reconstructed the sphericity-based boosts. For simplicity, we neglected the horizontal offset of MATHUSLA from the LLP production point in this toy analysis. For each fixed central value $b_{\rm LLP}$, we defined the spread as the standard deviation of the resulting sphericity-based boost distribution $\Delta(\log_{10} b_{\rm LLP\,meas})$, and we defined the bias to be the deviation of the average boost, $\log_{10}\langle b_{\rm LLP\,meas}\rangle - \log_{10} b_{\rm LLP}$.

The results from the toy model are shown in Fig. 8. We compare four scenarios. The black lines show simulations in which all particles are generated with lightlike momenta. The orange lines show simulations in which particles are generated with a thermal energy distribution. The solid lines show analyses using all generated particles. The dashed lines show analyses in which only upward-going particles are included, as would be the case for the MATHUSLA detector.

Consider first the values of the bias shown in the left-hand plot of Fig. 8. The simulations with lightlike momenta that include all particles have essentially no bias. This makes sense, since the analysis boosts the lightlike particles correctly. The simulations with a thermal energy distribution show a significant upward bias of 0.2, independent of velocity. Lightlike momenta are stiffer under boosts, and the analysis method treats these massive momenta as lightlike, so a larger boost is needed to balance the longitudinal momenta. The value of the bias is of the same order but somewhat larger than that in Fig. 4, due to the fact that this simulation omits the leading particles in jets, which are very relativistic. For low LLP velocities, considering upward-going particles only removes a significant part of the track distribution. The removal of the downward tracks increases the upward bias to about 0.4 in the region $b_{\rm LLP} < 1$.

The spread, shown in the right-hand plot of Fig. 8, has a value of about 0.15, independent of velocity, for both types of simulation. This is consistent with Fig. 4. For analyses that consider all particles, the spread increases at low momenta: The true LLP velocity is close to 0, so the reconstructed LLP boost is dominated by random deviations of the charged track distribution from spherical symmetry in the LLP rest frame. For the analyses with upward-going tracks only, the reconstructed velocity is determined by the bias, mitigating this effect.

The MATHUSLA detector has the capability of measuring the velocities of charged particles. A velocity measurement of 5% is straightforward. This requires recording hits to about 1 nsec precision over the typical 10 m flight

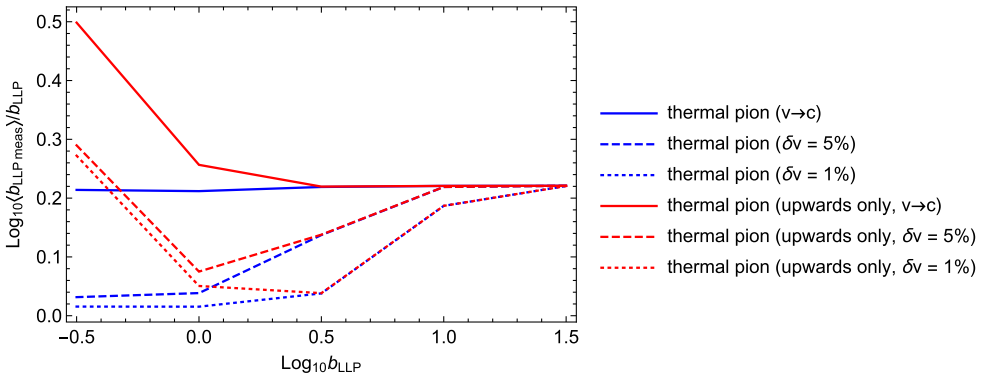

FIG. 9. Bias of the sphericity-based boost measurement for average charged particle multiplicity $N_{\rm ch} = 10$, showing the effect of speed measurement with precision δv in the thermal pion case.

path through the RPCs. A more aggressive design might allow a velocity measurement with 1% error. Thus, it is interesting to apply velocity information to the measured tracks to see if the bias can be reduced. The effect on the spread turns out to be quite small. The results on the bias are shown in Fig 9. The solid lines correspond to treating all tracks as lightlike as in the left-hand plot in Fig. 8. The dotted lines show the effect of using the velocity information for each track, assuming a velocity measurement with 5% or 1% error, spanning the range of capabilities estimated for MATHUSLA. Even with 5% errors, there is a significant effect for low LLP velocities. The bias returns at very low velocities when we restrict to upward-going tracks only. This effect of the track velocity measurement should be considered in more detailed design studies for MATHUSLA.

[1] A. Arvanitaki, N. Craig, S. Dimopoulos, and G. Villadoro, Minisplit, J. High Energy Phys. 02 (2013) 126.

[2] N. Arkani-Hamed, A. Gupta, D. E. Kaplan, N. Weiner, and T. Zorawski, Simply unnatural supersymmetry, arXiv: 1212.6971.

[3] G. F. Giudice and R. Rattazzi, Theories with gauge-mediated supersymmetry breaking, Phys. Rep. **322**, 419 (1999).

[4] R. Barbier *et al.*, R-parity violating supersymmetry, Phys. Rep. **420**, 1 (2005).

[5] C. Csaki, E. Kuflik, and T. Volansky, Dynamical R-Parity Violation, Phys. Rev. Lett. **112**, 131801 (2014).

[6] J. Fan, M. Reece, and J. T. Ruderman, Stealth supersymmetry, J. High Energy Phys. 11 (2011) 012.

[7] G. Burdman, Z. Chacko, H.-S. Goh, and R. Harnik, Folded supersymmetry and the LEP paradox, J. High Energy Phys. 02 (2007) 009.

[8] H. Cai, H.-C. Cheng, and J. Terning, A quirky little Higgs model, J. High Energy Phys. 05 (2009) 045.

[9] Z. Chacko, H.-S. Goh, and R. Harnik, The Twin Higgs: Natural Electroweak Breaking from Mirror Symmetry, Phys. Rev. Lett. **96**, 231802 (2006).

[10] M. Baumgart, C. Cheung, J. T. Ruderman, L.-T. Wang, and I. Yavin, Non-Abelian dark sectors and their collider signatures, J. High Energy Phys. 04 (2009) 014.

[11] D. E. Kaplan, M. A. Luty, and K. M. Zurek, Asymmetric dark matter, Phys. Rev. D **79**, 115016 (2009).

[12] Y. F. Chan, M. Low, D. E. Morrissey, and A. P. Spray, LHC signatures of a minimal supersymmetric hidden valley, J. High Energy Phys. 05 (2012) 155.

[13] K. R. Dienes and B. Thomas, Dynamical dark matter: I. Theoretical overview, Phys. Rev. D **85**, 083523 (2012).

[14] K. R. Dienes, S. Su, and B. Thomas, Distinguishing dynamical dark matter at the LHC, Phys. Rev. D **86**, 054008 (2012).

[15] I.-W. Kim and K. M. Zurek, Flavor and collider signatures of asymmetric dark matter, Phys. Rev. D **89**, 035008 (2014).

[16] A. Bouquet and P. Salati, R parity breaking and cosmological consequences, Nucl. Phys. **B284**, 557 (1987).

[17] B. A. Campbell, S. Davidson, J. R. Ellis, and K. A. Olive, Cosmological baryon asymmetry constraints on extensions of the standard model, Phys. Lett. B **256**, 484 (1991).

[18] Y. Cui and R. Sundrum, Baryogenesis for weakly interacting massive particles, Phys. Rev. D **87**, 116013 (2013).

[19] K. Barry, P. W. Graham, and S. Rajendran, Displaced vertices from R -parity violation and baryogenesis, Phys. Rev. D **89**, 054003 (2014).

[20] Y. Cui and B. Shuve, Probing baryogenesis with displaced vertices at the LHC, J. High Energy Phys. 02 (2015) 049.

[21] S. Ipek and J. March-Russell, Baryogenesis via particle-antiparticle oscillations, Phys. Rev. D **93**, 123528 2016.

[22] J. C. Helo, M. Hirsch, and S. Kovalenko, Heavy neutrino searches at the LHC with displaced vertices, Phys. Rev. D **89**, 073005 (2014).

[23] S. Antusch, E. Cazzato, and O. Fischer, Displaced vertex searches for sterile neutrinos at future lepton colliders, J. High Energy Phys. 12 (2016) 007.

[24] M. L. Graesser, Broadening the Higgs boson with right-handed neutrinos and a higher dimension operator at the electroweak scale, Phys. Rev. D **76**, 075006 (2007).

[25] M. L. Graesser, Experimental constraints on Higgs boson decays to TeV-scale right-handed neutrinos, arXiv:0705.2190.

[26] A. Maiezza, M. Nemevek, and F. Nesti, Lepton Number Violation in Higgs Decay at LHC, Phys. Rev. Lett. **115**, 081802 (2015).

[27] B. Batell, M. Pospelov, and B. Shuve, Shedding light on neutrino masses with dark forces, J. High Energy Phys. 08 (2016) 052.

[28] P. Bhupal Dev, R. N. Mohapatra, and Y. Zhang, Displaced photon signal from a possible light scalar in minimal left-right seesaw model, Phys. Rev. D **95**, 115001 (2017).

[29] P. S. B. Dev, R. N. Mohapatra, and Y. Zhang, Long-lived light scalars as Probe of low scale seesaw models, Nucl. Phys. **B923**, 179 (2017).

[30] M. Nemevek, F. Nesti, and J. C. Vasquez, Majorana Higgses at colliders, J. High Energy Phys. 04 (2017) 114.

[31] A. Caputo, P. Hernandez, J. Lopez-Pavon, and J. Salvado, The seesaw portal in testable models of neutrino masses, J. High Energy Phys. 06 (2017) 112.

[32] M. J. Strassler and K. M. Zurek, Echoes of a hidden valley at hadron colliders, Phys. Lett. B **651**, 374 (2007).

[33] M. J. Strassler and K. M. Zurek, Discovering the Higgs through highly displaced vertices, Phys. Lett. B **661**, 263 (2008).

[34] M. J. Strassler, Possible effects of a hidden valley on supersymmetric phenomenology, arXiv:0607160.

[35] T. Han, Z. Si, K. M. Zurek, and M. J. Strassler, Phenomenology of hidden valleys at hadron colliders, J. High Energy Phys. 07 (2008) 008.

[36] M. J. Strassler, Why unparticle models with mass gaps are examples of hidden valleys, arXiv:0801.0629.
[37] M. J. Strassler, On the phenomenology of hidden valleys with heavy flavor, arXiv:0806.2385.
[38] D. Curtin, R. Essig, S. Gori, P. Jaiswal, A. Katz *et al.*, Exotic decays of the 125 GeV Higgs boson, Phys. Rev. D **90**, 075004 (2014).
[39] D. Curtin, R. Essig, S. Gori, and J. Shelton, Illuminating dark photons with high-energy colliders, J. High Energy Phys. 02 (2015) 157.
[40] N. Craig, A. Katz, M. Strassler, and R. Sundrum, Naturalness in the dark at the LHC, J. High Energy Phys. 07 (2015) 105.
[41] D. Curtin and C. B. Verhaaren, Discovering uncolored naturalness in exotic Higgs decays, J. High Energy Phys. 12 (2015) 072.
[42] A. Coccaro, D. Curtin, H. J. Lubatti, H. Russell, and J. Shelton, Data-driven model-independent searches for long-lived particles at the LHC, Phys. Rev. D **94**, 113003 (2016).
[43] J. P. Chou, D. Curtin, and H. J. Lubatti, New detectors to explore the lifetime frontier, Phys. Lett. B **767**, 29 (2017).
[44] For an earlier proposal to detect long-lived particles copiously produced at hadron colliders, see K. Maki and S. Orito, Hadron colliders as the "neutralino factory": Search for a slow decay of the lightest neutralino at the CERN LHC, Phys. Rev. D **57**, 554 (1998).
[45] A. Fradette and M. Pospelov, BBN for the LHC: Constraints on lifetimes of the Higgs portal scalars, Phys. Rev. D **96**, 075033 (2017).
[46] A large-scale theory white paper on the MATHUSLA physics case, including the examination of many other signal models, is currently in preparation.
[47] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, J. High Energy Phys. 07 (2014) 079.
[48] T. Sjostrand, S. Mrenna, and P. Z. Skands, PYTHIA 6.4 physics and manual, J. High Energy Phys. 05 (2006) 026.
[49] T. Sjostrand, S. Mrenna, and P. Z. Skands, A brief introduction to PYTHIA 8.1, Comput. Phys. Commun. **178**, 852 (2008).
[50] We also consider the effect of the spatial tracking resolution and the momentum thresholds outlined in Sec. II, but unless otherwise stated, their the effect on our results is negligible.
[51] M. Bahr *et al.*, Herwig++ physics and manual, Eur. Phys. J. C **58**, 639 (2008).
[52] J. Bellm *et al.*, Herwig 7.0/Herwig++ 3.0 release note, Eur. Phys. J. C **76**, 196 (2016).
[53] T. Gleisberg, S. Hoeche, F. Krauss, M. Schonherr, S. Schumann, F. Siegert, and J. Winter, Event generation with SHERPA 1.1, J. High Energy Phys. 02 (2009) 007.
[54] P. Gras, S. Hoeche, D. Kar, A. Larkoski, L. Lnnblad, S. Pltzer, A. Sidmok, P. Skands, G. Soyez, and J. Thaler, Systematics of quark/gluon tagging, J. High Energy Phys. 07 (2017) 091.
[55] G. W. Van Apeldoorn *et al.*, Thermal emission of pions in antiproton proton annihilation at 12-GeV/c, Z. Phys. C **7**, 235 (1981).