

Long-lived particle decays at the proposed MATHUSLA experiment

David Curtin^{*} and Jaipratap Singh Grewal[†]

Department of Physics, University of Toronto, Toronto, Ontario M5S 1A7, Canada

 (Received 5 September 2023; accepted 16 January 2024; published 11 April 2024)

We carefully study the decay and reconstruction of long-lived particle (LLP) decays in the proposed MATHUSLA LLP detector for the HL-LHC. Our investigations are focused on three LLP benchmark models. MATHUSLA's primary physics target is represented by hadronically decaying LLPs with mass above ~ 10 GeV, produced in exotic Higgs decays. We also investigate GeV-scale scalar and right-handed neutrino LLPs, which are the target of many other proposed experiments. We first introduce a public MATHUSLA FastSim code to allow for efficient signal-only studies of LLP decays in MATHUSLA and general external LLP detectors. For each of our benchmark scenarios, we carefully simulate LLP production and decay, and make our simulation library publicly accessible for future investigations and comparisons with other experiments. We then systematically study the geometric acceptance of MATHUSLA for LLP decays in these scenarios, and present updated sensitivity projections that include these acceptances. Our results show that the idealized reach of MATHUSLA computed in earlier studies is mostly realized. We also investigate possible ways of increasing the signal acceptance using the inherent geometric flexibility of the FastSim, which will provide useful inputs for realistic experimental and engineering optimization of the detector in the future.

DOI: [10.1103/PhysRevD.109.075017](https://doi.org/10.1103/PhysRevD.109.075017)

I. INTRODUCTION

Many Beyond Standard Model (BSM) scenarios predict new physics at the Large Hadron Collider (LHC). This is most acute for solutions of the hierarchy problem and other extensions of the Higgs sector, but many other phenomena like dark matter freeze-out or the electroweak phase transition also motivate an experimental focus on the TeV scale. In light of null results to date, it is therefore crucial to extend our capabilities of detecting any signal of new physics that might be produced at the LHC and address the blind spots of the ATLAS, CMS, and LHCb detectors.

One of the most obvious such blind spots are some types of long-lived particles (LLPs) that decay a macroscopic distance away from their production point. LLPs are a very generic signature of hidden sectors and other new physics, which are explicitly predicted by many top-down and bottom-up BSM theories including solutions to the hierarchy problem like neutral naturalness, many dark matter and baryogenesis scenarios, as well as extensions of the neutrino sector, and general hidden valleys, see Ref. [1] for

a review. The lifetime τ of LLPs must, in general, be regarded as a free parameter up to at least the approximate bound set by big bang nucleosynthesis (BBN), $c\tau \sim 10^7$ m. The geometric nature of the displaced decay signal makes it a spectacular signature with potentially very low Standard Model (SM) background (depending on LLP mass and decay mode), but it is also generally missed by prompt searches. The dedicated LLP search program at the LHC main detectors has undergone rapid development in the last decade [2] to take advantage of this untapped discovery potential for new physics. Even so, intrinsic trigger and background limitations severely limit the ability of ATLAS, CMS, and LHCb to search for broad classes of neutral LLPs, especially for long lifetimes, where the tiny fraction of decays in the main detector makes high signal efficiency and near-complete background rejection a necessary condition for their discovery. (Missing energy searches can have some reach, but typical production rates are too small for this to be effective, see discussion in [1].) These limitations are most severe for LLPs that result in less than a few 100 GeV of dominantly hadronic SM final states from their production and decay, as well as GeV-scale LLPs with arbitrary decay modes [3].

This has motivated the proposal of the MATHUSLA (MASSive Timing Hodoscope for UltraStable neutral pArticles) experiment for the HL-LHC [4–6], as well as other transverse LLP detectors like CODEX-b [7,8] and ANUBIS [9], and forward LLP detectors like FASER/FPF [10–13] and FACET [14]. (Similar proposals have also

^{*}dcurtin@physics.utoronto.ca

[†]jp.grewal@mail.utoronto.ca

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

been made for future hadron colliders, see [4] and [15,16].) As we briefly review in Sec. II, MATHUSLA is envisioned as a ~ 100 m footprint building on CERN-owned land near CMS, with a mostly empty decay volume monitored by trackers for reconstruction of LLP decays into charged particles. Its shielded location on the surface eliminates QCD backgrounds and trigger bottlenecks that limit LLP reach for the LHC main detectors, while its large decay volume allows it to discover LLPs with lifetimes up to the BBN limit, extending sensitivity of the main detectors by three orders of magnitude in LLP lifetime and production rate [4].

The sensitivity of MATHUSLA to various BSM LLP scenarios has been extensively studied (see Refs. [1,17] and references therein). However, none of these studies took into account realistic geometric acceptances and experimental efficiencies for LLP decays in the MATHUSLA detector.¹ In general, under some reconstruction criterion, i.e., requiring some number of charged tracks of a certain quality to reconstruct a displaced vertex (DV), the predicted number of observed events at MATHUSLA (or any LLP detector) can be written as

$$N_{\text{obs},i} = \underbrace{(n\sigma\mathcal{L}) \otimes \xi_{\text{geo}}^{\text{LLP}} \otimes \bar{P}_{\text{decay}}}_{N_{\text{decay}}} \cdot \underbrace{\text{Br}(\text{vis}) \otimes \xi_{\text{geo},i}^{\text{decay}}}_{N_{\text{visible}}} \otimes \epsilon_i^{\text{recon}}. \quad (1)$$

Here, i labels the specific reconstruction criterion; $(n\sigma\mathcal{L})$ is the total number of LLPs produced at the HL-LHC (with n the average number of LLPs produced per event, production cross section σ , and luminosity \mathcal{L}); $\xi_{\text{geo}}^{\text{LLP}}$ is the fraction of those LLPs that fly through the MATHUSLA decay volume (typically $\sim 5\%$ but this depends on kinematics), \bar{P}_{decay} is the average chance for those LLPs to decay inside MATHUSLA, mostly determined by the LLP boost distribution and its lifetime, $\text{Br}(\text{vis})$ is the fraction of LLP decays that includes charged particles for vertex reconstruction in the final state, $\xi_{\text{geo},i}^{\text{decay}}$ is the fraction of those visible decays in the decay volume for which the SM final states satisfy the minimum geometrical requirements of the reconstruction criterion i (e.g., the charged final state trajectories intersect sufficiently many sensor planes), and finally, $\epsilon_i^{\text{recon}}$ is the experimental efficiency of actually reconstructing the displaced vertex for decays that satisfy these geometrical requirements. We use \otimes in the above product to emphasize that each term after the first one [except $\text{Br}(\text{vis})$, as indicated] should be interpreted as a conditional probability that depends on the preceding requirements, e.g., the average

¹A recent study did account for ceiling trackers in an approximate way [18].

chance of decaying in MATHUSLA for the subset of LLPs that actually flies through MATHUSLA.

Essentially all previous studies of MATHUSLA's sensitivity relied on predictions of N_{decay} or N_{visible} in Eq. (1). Furthermore, in cases where N_{visible} was used, decays to hadrons, including only neutral hadrons, may have been counted as visible. In this paper, we perform detailed simulation studies to determine $\text{Br}(\text{vis})$ and, most importantly, the geometric acceptance $\xi_{\text{geo},i}^{\text{decay}}$ for LLP decays in a variety of benchmark scenarios. This allows us to obtain much more realistic estimates of MATHUSLA's LLP reach based on $N_{\text{recon},i}$, the total number of potentially reconstructable LLP decays. This will be a very good approximation of the fully realistic reach, since the experimental efficiency $\epsilon_i^{\text{recon}}$ is expected to be very close to 1 owing to MATHUSLA's relatively clean and low-rate experimental environment, and the well-understood nature of particle tracking and vertex reconstruction.

To compute geometric acceptances for LLP decays we use a custom PYTHON-based detector simulation code called MATHUSLA FastSim [19], discussed in Sec. III. This geometry-only detector simulation is highly flexible, using input cards to specify the detailed LLP detector geometry, track and vertex reconstruction criteria, minimum energy thresholds and spatial resolution. We also make this code publicly available to facilitate future LLP sensitivity estimates for MATHUSLA (and possibly other LLP detectors).²

We then study three benchmark LLP models, with details of their simulation discussed in Sec. IV. MATHUSLA's primary physics target are hadronically decaying LLPs with $\lesssim \mathcal{O}(100 \text{ GeV})$ masses, with the most theoretically well-motivated version of this scenario being exotic Higgs decays [20–26] into a pair of LLPs XX with mass m_X , each of which has lifetime $c\tau_X$ and decays to a pair of SM jets. Since each such LLP decay produces $\mathcal{O}(10)$ charged particles for $m_X \gtrsim$ few GeV [27], we expect MATHUSLA to have very high geometric acceptance for reconstructing these decays. MATHUSLA's secondary physics target are GeV-scale LLPs [28]. These can also be searched for at a variety of fixed-target and forward experiments, but MATHUSLA's reach is competitive and complementary. Two benchmark models we focus on are a light singlet scalar LLP of mass m_S that has a tiny mixing angle $\sin\theta$ with the SM Higgs [29], referred to as “SM + S” and a right-handed Majorana neutrino (RHN) LLPs (see e.g., [30] for a review) of mass m_N , where we adopt the PBC benchmarks [28] that have dominant mixing $U_{e,\mu,\tau}$ with one active neutrino flavor, referred to as “RHN ($U_{e,\mu,\tau}$)”. Care must be taken in simulating production and decay of these light LLPs due to significant hadronic uncertainties. We make our complete library of LLP production and decay simulation events for these benchmark models

²Link (https://github.com/davidcurtin/MATHUSLA_FastSim) to MATHUSLA FastSim GitHub repository.

available in public repositories,³ to aid in future analyses and comparisons with other transverse LLP detector proposals.

Our results for the current baseline MATHUSLA geometry [17] are presented in Sec. V. We find that the geometric acceptance for hadronic LLP decays above a few GeV is very good, $\xi_{\text{geom}}^{\text{decay}} \sim 0.7\text{--}0.8$, with the slight reduction from unity almost entirely due to regions in the back of the detector with respect to the LHC interaction point (IP), where the decay products which are boosted along the LLP trajectory exit the detector through the rear wall without ever intersecting a tracking plane. For the decay of $\lesssim\text{GeV}$ LLPs without missing energy, as in the SM + S model, $\xi_{\text{geom}}^{\text{decay}} \sim 0.5$ across most of the relevant mass range once final state hadronization is taken into account. On the other hand, for RHN LLPs the final state multiplicity is significantly lower even at masses above a GeV, and the geometric acceptance is correspondingly lower, $\xi_{\text{geom}}^{\text{decay}} \sim 0.1\text{--}0.4$. This is also due to the fact that the active neutrino emitted in the LLP decay allows the charged decay products to more readily escape through the walls or floor of the detector. We illustrate the impact of these realistic geometric acceptances in updated reach plots for these benchmark models. The reach for LLPs from exotic Higgs decays and even the SM + S model closely resembles the previous idealized estimates. The reach for RHNs is slightly but not qualitatively reduced.

Finally, the flexibility of our detector FastSim allows us to easily investigate the impact of various MATHUSLA design aspects on the geometric acceptance for LLP decays and hence the overall BSM reach in Sec. VI. We find that the requirements of a local trigger in the MATHUSLA hardware are unlikely to impact sensitivity, and the gaps between detector modules in the ceiling also only have a minor effect. On the other hand, instrumenting the back wall or even all four walls with trackers would significantly enhance the sensitivity to low-multiplicity LLP decays. Such an upgrade compared to the baseline design proposal would carry some cost, which would have to be weighted against the expense of constructing the large 100 m baseline detector volume. We show that sensitivity scales better than linear with detector area (for fixed height), suggesting that one possible path towards optimizing the detector for a given cost may involve adjusting its size while enhancing tracker coverage. Our results provide a clear understanding of how the signal yield scales with different aspects of MATHUSLA's geometry and instrumentation. This will be one of many important inputs for future experimental and

engineering studies that include the full gamut of physics, technical and practical considerations to optimize the final design of the MATHUSLA detector.

II. THE MATHUSLA EXPERIMENT

The MATHUSLA detector proposal is described at length in Refs. [4–6] and most recently in [17], so we only briefly review the most salient details needed for our simplified analysis.

Figure 1 shows the overall geometry of the proposed detector relative to CMS. A large air-filled decay volume with a footprint of 100 m \times 100 m and a height of 25 m is instrumented with a ceiling tracker composed of six layers of plastic scintillator above the decay volume and two layers at a height of 20 m. The tracking layers are composed of scintillator bars with alternating orientations, supplying position information for charged particle hits with $\mathcal{O}(\text{cm})$, $\mathcal{O}(\text{m})$ resolution along the transverse, longitudinal bar direction. Together with $\mathcal{O}(\text{ns})$ timing resolution for each hit, this allows for highly robust four-dimensional tracking to construct displaced vertices in the decay volume and reject backgrounds like downwards-traveling cosmic rays.

Background rejection is aided by a double layer of detectors in the front wall and floor to help reject cosmic and LHC muons, which by rate are by far the dominant non-BSM processes in MATHUSLA. Other background processes include atmospheric neutrino scattering in the air and production of upwards-traveling K_L^0 in the detector floor from cosmic ray interactions, see e.g., [6] for preliminary discussions. While much rarer, these backgrounds could be very important since they are potentially more likely to produce displaced vertices in the detector volume. In general, these backgrounds will not matter for MATHUSLA's primary physics case of hadronically decaying LLPs with masses above the GeV scale, since no existing backgrounds can mimic the large multiplicity of tracks from the DV. On the other hand, for GeV-scale LLPs




FIG. 1. (a) MATHUSLA geometry relative to the CMS collision point, illustrating how LLPs can decay in the detector and be reconstructed as displaced vertices by the ceiling trackers. For clarity, the modular structure of the MATHUSLA detector is not shown. (b) Vertical structure of the 9 m \times 9 m footprint detector modules, which are arranged in a square grid with 1m gaps to make up the 100 m \times 100 m detector footprint.

³Links to event repositories: exotic higgs decay (https://github.com/davidrcurtin/MATHUSLA_LLFiles_HXX), SM + S (https://github.com/davidrcurtin/MATHUSLA_LLFiles_SMS), RHN (U_e) (https://github.com/davidrcurtin/MATHUSLA_LLFiles_RHN_Ue), RHN (U_μ) (https://github.com/davidrcurtin/MATHUSLA_LLFiles_RHN_Umu), RHN (U_τ) (https://github.com/davidrcurtin/MATHUSLA_LLFiles_RHN_Utau).

more care will have to be taken to understand these rare sources of SM DVs. Veto strategies are available, and preliminary estimates indicate that controlling these backgrounds is possible without significantly affecting MATHUSLA’s physics reach. Even so, the potential existence of these backgrounds makes it interesting to understand what the reach of 3+ pronged DV searches would be, given its inherent robustness against any known backgrounds.

III. MATHUSLA DETECTOR FASTSIM

We now introduce the MATHUSLA *FastSim* public LLP detector simulation code [19], which is available for download here.^{4,5} This relatively simple detector simulator is entirely focused on, effectively, obtaining $\xi_{\text{geo},i}^{\text{decay}}$ in Eq. (1). Detailed usage instructions and worked examples are given in the code repository; here we merely summarize the important functionality for our physics studies.

The MATHUSLA *FastSim* is a PYTHON module that accepts a `param_card` to specify the exact geometry of an external LLP detector (decay volume within which a displaced vertex is counted as signal, module dimensions and spacing, and location of horizontal and vertical detector planes) as well as (possibly multiple) reconstruction criteria for individual tracks and DVs. There are no material interactions, and no realistic track fitting or displaced vertex reconstruction is performed. Rather, a charged particle is said to be reconstructable as a track if it intersects some minimum number of detector planes.⁶ For example, our “default track” requires at least four hits in any of the sensor layers.⁷ The underlying assumption, just as explained for Eq. (1), is that the experimental efficiency for actually reconstructing a track will be very close to 1 if it satisfies these geometric requirements. Similarly, a reconstruction criterion for a displaced vertex merely requires that some minimum number of charged final states satisfy some track reconstruction criteria. Our default vertex requirement, called “DV2”, is that at least two observable charged particles from the LLP decay satisfy the default track criterion, but we also define an analogously defined “DV3” criterion to investigate the reach of a 3-pronged DV search. The vertex reconstruction criteria are summarized in Table I.

For our signal acceptance studies, we implement the current MATHUSLA benchmark geometry shown in Fig. 1 and present results for this geometry in Sec. V. However, it

TABLE I. Idealized geometrical reconstruction criteria for the decay products of an LLP decay inside the MATHUSLA decay volume. DV2 is the default, with DV3 included to investigate the robustness of LLP reconstruction with respect to an additional track requirement. An $e^\pm, \mu^\pm, \pi^\pm (K^\pm) \{p, \bar{p}\}$ is assumed to be observable if it has momentum above 200 MeV (400 MeV) {600 MeV}. The precise thresholds can be important for soft backgrounds, but have minimal impact on the acceptance for our LLP decays of interest.

Label	Idealized reconstruction criterion
DV2 (<i>default</i>)	At least 2 observable charged particles from the LLP decay each intersect at least 4 detector panes, across any of the modules.
DV3	At least 3 observable charged particles from the LLP decay each intersect at least 4 detector panes, across any of the modules.

is easy to specify different geometries to study their impact on BSM sensitivity, as we discuss in Sec. VI.

The MATHUSLA *FastSim* can be used to check whether a single charged particle (e.g., a cosmic ray muon) or a decaying LLP can be reconstructed as a track or a vertex respectively, according to the supplied geometrical criteria. In practice, the simulation includes “charged particle gun” and “LLP gun” functions. The charged particle gun is used to specify a starting point, momentum, and PID for some charged particle, and the *FastSim* returns which, if any, track reconstruction criteria the charged particle satisfies. The LLP gun takes as input an LLP starting point; an LLP decay position (therefore specifying the LLP momentum direction); the LLP boost $b = |\vec{p}|/m$ and a list of LLP decay final states in the LLP rest frame. The *FastSim* transforms the final states to the lab frame, places the corresponding outgoing particle trajectories at the chosen LLP decay point, and outputs which, if any, vertex reconstruction criteria this LLP decay satisfies. For convenience, the LLP decay gun can generate 2- and 3-body LLP decays on-the-fly for specified final state PIDs and masses using phase space only, but in general, the decay of LLPs is simulated previously by whatever method desired, saved to a file, and fed to the LLP decay gun as “ammo”.⁸ The momenta of produced LLPs are similarly simulated separately for a given signal model.

The way to actually compute the number of observed DVs satisfying a reconstruction criterion i is therefore as follows: (Rapidity and azimuthal angle are defined with respect to and around the beam axis, respectively.)

- (1) Start with a list of 4-vectors of LLPs produced at the HL-LHC. For generality, assume each LLP k carries

⁸Note that *FastSim* will not decay “detector-stable” states like π^\pm, K^\pm . Since their decay would typically yield another charged particle flying in close alignment with the original trajectory, this is unlikely to significantly affect our analysis.

⁴github.com/davidcurtin/MATHUSLA_FastSim.

⁵We would like to acknowledge the important contributions of Lillian Luo and Wentao Cui, who wrote an early version of the *FastSim* code for MATHUSLA geometry optimization studies [6].

⁶It is possible to specify that these detector planes must belong to a specific group, e.g., the ceiling stack rather than the floor.

⁷A “loose” or “super-tight” track might be defined to require fewer or more hits, but we use the default number of four minimum hits per track in the remainder of this analysis.

weight w_k , which is just the actual number of LLPs produced at the LHC this event represents.⁹

- (2) Define rapidity and azimuthal angle ranges $(\eta_{\min}, \eta_{\max}), (\phi_{\min}, \phi_{\max})$ that closely enclose the solid angle subtended by the entire MATHUSLA decay volume from the IP. For the benchmark geometry, suitable values are (0.66, 1.95) and $(-0.76, 0.76)$.
- (3) Discard LLPs outside of MATHUSLA's rapidity range. For each remaining LLP k , exploit the rotational symmetry of events around the beam axis and rotate the LLP 4-vector to a random angle in the range $(\phi_{\min}, \phi_{\max})$ while reducing its weight by a factor of $(\phi_{\max} - \phi_{\min})/2\pi$. This increases the numerical efficiency of our simulations by a factor of a few.
- (4) The FastSim will then determine the distances $L_{1,k}, L_{2,k}$ along the LLP's trajectory where the LLP enters and exits the MATHUSLA decay volume. The probability that the LLP decays in MATHUSLA is then

$$P_{\text{decay},k} = e^{-\frac{L_{1,k}}{b_k c\tau}} - e^{-\frac{L_{2,k}}{b_k c\tau}}, \quad (2)$$

where $b_k = |\vec{p}_k/m_k|$ and $c\tau$ is the LLP's decay length.¹⁰

- (5) The FastSim now chooses an actual decay position inside the decay volume, according to the exponential decay distribution between L_1 and L_2 normalized to unity. In the long lifetime limit, this approaches uniform decay probability along the trajectory. The LLP is then decayed according to the final-state ammo loaded into the decay gun, and the FastSim evaluates whether this particular LLP is reconstructed according to DV criterion i , giving $\zeta_k^{(i)} = 0$ or 1.
- (6) The total number of decays in the decay volume is then

$$N_{\text{decay}} = \sum_k w_k \frac{\phi_{\max} - \phi_{\min}}{2\pi} P_{\text{decay},k}, \quad (3)$$

while the number of reconstructable vertices is

$$N_{\text{recon},i} = \sum_k w_k \frac{\phi_{\max} - \phi_{\min}}{2\pi} P_{\text{decay},k} \zeta_k^{(i)}. \quad (4)$$

⁹If this list is extracted from unweighted LLP production events with total production cross section σ and N simulated events, then $w_k = (\sigma\mathcal{L})/N$. Note that multiple LLPs per event will each carry the weight of the full event.

¹⁰Some events will be inside of MATHUSLA's rapidity and angle range, but still not pass through MATHUSLA due to the irregular shape of MATHUSLA's projection into the η, ϕ plane from the IP. In that case, $P_{\text{decay},k} = 0$.

The sum \sum_k is only over events in the correct rapidity range.¹¹ For this LLP, the geometric acceptance is then $\xi_{\text{geo},i}^{\text{decay}} = N_{\text{recon},i}/N_{\text{decay}}$.

In addition to this basic functionality, the code has several other potentially useful features:

- (i) It is possible to specify an internal trigger criterion, in the format of requiring a track to be reconstructable using only hits within a 3×3 or larger subgroup of neighboring modules. This is not a fully realistic implementation of how the MATHUSLA data acquisition system may function in the future, but it can give some indication of whether a local-track-finding trigger requirement significantly impacts signal acceptance, as we investigate in Sec. VI;
- (ii) The LLP detector geometry can be quite general, including vertical detector panes, but it is limited by requiring each sensor pane to be lined up with the x , y , or z axis of the internal coordinate system, which does not have to line up with the beam line. (This greatly speeds up finding intersections with particle trajectories.) It is therefore possible to implement general "rectangular" LLP detector geometries, but not e.g., slanted sensor panes;
- (iii) We specify minimum momentum thresholds for different detector-stable charged particles for them to be detectable, see Table I. These thresholds have been implemented since in principle they can be important for certain low-energy backgrounds, but in practice they have no impact on our analysis;
- (iv) Some effects of finite spatial tracking resolution can be captured by specifying the dimensions and orientation of the scintillator bars making up each sensor layer, and merging charged particle "hits" if they occur in the same bar. For the $\sim\text{cm} \times \text{m}$ bar dimensions discussed in e.g., [17], this has no impact on any of our analyses due to the large size of MATHUSLA and the modest $\mathcal{O}(1-10)$ boosts of our LLP benchmark models, resulting in well-separated final state tracks.

More sophisticated studies that include backgrounds and material interactions will require the use of full GEANT simulations [31], which are much more resource intensive. The simplicity and flexibility of our geometry-only FastSim therefore makes it well-suited for scans over the parameter space of BSM LLP models and first optimization studies.

IV. LLP PRODUCTION AND DECAY SIMULATION

Studying the sensitivity of MATHUSLA using our FastSim requires separate simulation of LLP production and decays. For GeV-scale LLPs, hadronic uncertainties

¹¹These expressions are correct if the LLP decay files include invisible decays; otherwise, factors of $\text{Br}(\text{vis})$ must be included.

have to be carefully considered. We now discuss how these signal samples are generated for the three benchmark models under consideration; production of hadronically decaying LLPs in exotic Higgs decays, the SM + S model, and right-handed neutrinos. For what it's worth, we do believe that our simulations are the most accurate of their kind to date. We therefore make our entire LLP production and decay library available for public download, which may help with future sensitivity estimates for other transverse LLP detectors at the LHC.

A. Exotic Higgs decays

In this simplified model, a bosonic LLP X with mass m_X and decay length $c\tau_X$ is produced in exotic Higgs decays with some $\text{Br}(h \rightarrow XX)$. In addition to these three parameters, the decay mode of X must be specified. Different choices allow this model to stand-in for different more general BSM scenarios. Yukawa-ordered decays to SM fermions are one of the most motivated possibilities, since this would arise when identifying X with some hidden sector scalar that inherits its couplings from the Higgs via a small mixing angle. This can arise in the SM + S scenario below, but also in more complete BSM theories like the fraternal twin Higgs [21,22,32]. A dark photonlike [33,34] X with mass above $2m_\pi$ would decay dominantly to hadrons while still having a significant lepton fraction. However, a dark vector could also have more hadrophilic or leptophilic couplings (see e.g., [35]), resulting in light-flavor jets or leptons. An axionlike particle [36] X could have dominant decay modes to different SM gauge bosons, including gluons.

We will study two hadronic decay modes, since they are generally motivated and also represent LLP blind spots for the main detectors; $X \rightarrow \bar{b}b$ and $X \rightarrow gg$. Note that the former is the almost completely dominant decay mode for a scalar LLP with Yukawa-ordered SM fermion couplings for $m_X \sim 15\text{--}55$ GeV (see e.g., [37]).

We use MadGraph5 3.4.2 [38] + showering in Pythia 8 [39,40] to simulate SM Higgs production in gluon fusion and vector boson fusion. For gluon fusion, the effective ggh operator is added to the MadGraph model—jet matching between the hard matrix element and the shower is used to include events with up to one extra hard jet—and the events are slightly reweighted to 3 the NLO + NNLL Higgs p_T spectrum computed by HqT 2.0 [41,42]. For the total Higgs production cross section at the 14 TeV HL-LHC, the NNLO + NNLL values reported by the LHC Higgs Working group were used [43], 54.6 pb for gluon fusion and 4.27 pb for vector boson fusion. The Higgs boson 4-vectors were extracted from these produced events and decayed to two X LLPs for various m_X . The resulting list of X 4-vectors are then used to aim the FastSim LLP gun.

Decays of the X LLP to $\bar{b}b, gg$ are also performed in MadGraph5 and Pythia 8 using the HAHM model [34]. The detector-stable X decay products are extracted, boosted to

the X rest frame and saved to a file that supplies the decay ammo for the FastSim LLP gun.¹² The final state typically includes $\mathcal{O}(10)$ charged particles, with some dependence on m_X [27], and typically slightly more for the gg decay than for $\bar{b}b$, though that difference has very little impact as we discuss below.

Our library of LLP production and decay events for this model can be downloaded here.¹³

B. SM + scalar

The SM + S simplified model includes a SM singlet scalar S of mass m_S that has a small mixing $\sin\theta$ with the SM Higgs, setting both its production rate in exotic heavy meson decays (dominantly $B \rightarrow SK$) and its lifetime.¹⁴ We focus on the mass range $m_S \lesssim 5$ GeV. We closely follow [44] to compute predictions for production and decay of the S LLP, since it is to the best of our knowledge the most up-to-date treatment, taking into account final-state hadronization in S -decay and supplying the necessary information to compute exclusive decays $B \rightarrow X_k + S$ of B -mesons to hadrons X_k and the LLP. We now outline the details of our simulation strategy.

We use the public FONLL v1.3.2 code [45–47] to compute the fully differential $d\sigma(pp \rightarrow B + X)/dp_T^B d\eta^B$ cross section for B -meson production at the 14 TeV HL-LHC. This can be directly used to define a Monte Carlo generator of B -meson 3-momenta at the HL-LHC.¹⁵ These 3-momenta are assigned to be either $B^0, B^\pm, B_s^0, B_c^\pm$ with respective probabilities (0.448, 0.448, 0.103, 1.7×10^{-4}), which are the relative B -meson fractions produced by Pythia 8

¹²Note that B and D -mesons are not regarded as detector stable, and if they are produced in any LLP decay, their decay products are placed at the LLP primary decay vertex in MATHUSLA for all of our analyses. Given their tiny decay length compared to the size of MATHUSLA, its at-best $\sim\text{cm}$ DV spatial resolution, and the many meters of decay length for detector stable final states like π^\pm, K^\pm , this simplification will not affect any realistic analysis.

¹³github.com/davidrcurtin/MATHUSLA_LLPfiles_HXX.

¹⁴ S can also be produced in exotic Higgs decays, but that branching ratio depends on additional parameters in the model, namely the quartic $\lambda_S S^2 |H|^2$ coupling [28]. Since we separately study exotic Higgs decays, we do not include this production mode in our study of the SM + S model.

¹⁵This implicitly discards intraevent B -meson momentum correlations for $B\bar{B}$ production, which is irrelevant for our MATHUSLA signal estimate. The straightforward reason for this is that the long-lifetime limit and the geometrical (anti) correlation of B -meson momenta in the same HL-LHC event makes it exceedingly unlikely that two B -mesons from the same HL-LHC event result in a reconstructed LLP decay in MATHUSLA. The more complete explanation, which applies even if such a double decay were to take place, is that the MATHUSLA readout defines its own time structure that is independent of HL-LHC events, so detection and reconstruction of two LLPs in MATHUSLA is independent even if they are produced in the same HL-LHC event (neglecting the exceedingly rare case of spatial overlap of the decay vertices).

hadronization of $b\bar{b}$ events with the default tune, and are compatible with LHCb measurements [48]. Using the absolute exclusive decay widths $\Gamma(B \rightarrow X_k + S)$ computed in [44], these B -meson 4-vectors are decayed with full-phase space information to the daughter hadron X_k + the LLP, for each different m_S . This makes our treatment of LLP kinematics in this benchmark model the most accurate to date in any simulation study, as far as we are aware. The resulting LLP 4-vectors are then supplied to the FastSim LLP gun to specify LLP trajectories.

Special care must be taken to account for and minimize the impact of hadronization uncertainties in the decay of the S LLP. Our simulation strategy again uses the results of [44], adopting their lifetime calculation and exclusive decay rates of the S . Specifically:

- (i) For $m_S \leq 0.7$ GeV, the S is directly decayed to $ee, \mu\mu$, and $\pi\pi$ in accordance with the branching fractions in [44];
- (ii) For $m_S \geq 2$ GeV, we again use the HAHM MadGraph model and Pythia 8 to simulate $S \rightarrow gg, ss, \tau\tau, cc$, extract detector-stable SM final states, and weigh the different processes according to the branching fractions in [44];
- (iii) The intermediate mass region $0.7 \text{ GeV} < m_S < 2 \text{ GeV}$ is the most affected by hadronization uncertainties. We use MadGraph5 + Pythia 8 to simulate $S \rightarrow \mu\mu, ss, gg$, extract detector-stable final states, and weigh the different processes according to $\text{Br}(S \rightarrow \mu\mu)$, $\text{Br}(S \rightarrow KK)$ and $(1 - \text{Br}(S \rightarrow \mu\mu) - \text{Br}(S \rightarrow KK))$ in [44], respectively.

The resulting charged particle multiplicity distribution in S LLP decays is shown in Fig. 2 (top). For low mass, S always decays to two charged states. An invisible decay fraction turns on at the $2m_\pi$ threshold since π_0 cannot be detected by MATHUSLA. The fraction of multicharge production events increases slightly in the 0.7–2 GeV region where hadronization details are maximally uncertain, but this rate and therefore the resulting hadronization uncertainty of our analysis results is minor. [The exception is our exclusive prediction for the rates of multipronged vertices, which must be regarded as having $\mathcal{O}(1)$ uncertainty near masses of 1 GeV.] Production of 4+ charged states per S decay then increases sharply and monotonically (except for the 2τ threshold, which has lower-charged multiplicity than other hadronic decay modes that dominate below $2m_\tau$) for $m_S > 2$ GeV. As we will see, this robustly predicted production of multiple charged states from hadronization is very important since it significantly increases MATHUSLA’s geometric acceptance for S decays in the multi-GeV regime.

Our library of LLP production and decay events for this model can be downloaded here.¹⁶

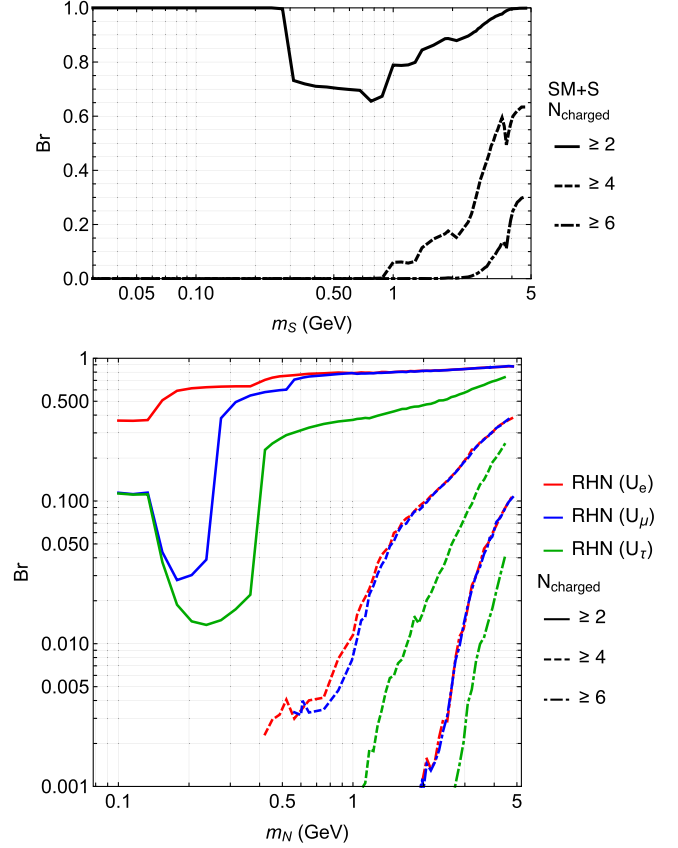


FIG. 2. Top: The fraction of LLP decays with $\geq 2, 4, 6$ (solid, dashed, dot-dashed) detector-stable charged particles in our simulations of the SM + S model. Bottom: The same for the three right-handed neutrino benchmark models with electron, muon and tau neutrino mixing dominance (red, blue, green). $\text{Br}(N_{\text{charged}} \geq 2)$ corresponds to $\text{Br}(\text{vis})$ in Eq. (1).

C. Right-handed neutrinos

The three simplified RHN models we consider are defined by a single Majorana RHN N with mass m_N and dominant mixing U_i^2 to one of the active neutrino flavor eigenstates $i = e, \mu, \tau$. We closely follow [49] to compute their production at the HL-LHC and decay in MATHUSLA. We focus on the mass range $m_N \lesssim 5$ GeV.

RHNs are produced in the decay of B and D mesons, as well as the decay of W, Z, τ . For production in meson decays, we use FONLL as described in the previous section to obtain 4-vectors of B and D mesons produced at the HL-LHC. These are decayed event-by-event into the various exclusive 2- and 3-body final states that include a RHN using full phase space information, in accordance with the exclusive decay rates computed by implementing the calculations of [49]. For production in W, Z, τ decays, we use the HeavyN MadGraph model [50–52] with showering in Pythia 8 to compute RHN production at leading order for the HL-LHC. The MadGraph cross section for RHNs is rescaled by a universal electroweak K -factor of 1.3, which is required to reproduce the experimentally measured W, Z

¹⁶github.com/davidrcurtin/MATHUSLA_LLFiles_SMS.

production cross section [53]. This supplies the LLP 4-vectors to aim the FastSim LLP gun.

The LLP decay length can be computed following [49] for each of the three RHN scenarios. To simulate the decay final states, we adopt the following strategy:

- (i) For low masses below the multihadron threshold $m_N < 0.42, 0.53, 0.42$ GeV for the electron-, muon-, and tau-mixing dominated scenarios, we compute the exclusive decay branching fractions to 2- and 3-body SM final states as outlined in [49] and decay the RHN 4-vectors using full phase space information;
- (ii) For higher masses above the multihadron threshold, we use MadGraph5 and the HeavyN model with showering and hadronization in Pythia8 to generate RHN decay events and extract the detector-stable SM final states. This is subject to the usual hadronic uncertainties of the predicted exclusive final state fractions near $m_N \sim$ GeV, but as we will see below we expect this to have very little effect on our analysis. For masses \gg GeV this recovers the partonic decay branching ratios computed in [49]. For each m_N in each RHN scenario, care is taken to only simulate those parton-level N -decays to SM quarks that are kinematically allowed when considering the minimum mass of the resulting final-state hadrons rather than bare quarks.

The resulting charged-particle multiplicity distribution in RHN LLP decays is shown in Fig. 2 (bottom). Several qualitative differences to the SM + S scenario are immediately apparent: the presence of fully or effectively invisible decay modes ($3\nu, \nu\pi^0\pi^0, \dots$) leads to a lower fraction of decays with 2+ charged final states at low masses, and even at a few GeV masses for the RHN (U_τ) scenario, owing to the large m_τ forbidding the charged-current decay at lower m_N . The lower level of hadronic activity in the decay compared to the SM + S scenario can be readily understood since the electroweak nature of the decay redirects a large fraction of the available mass energy to the production of a neutrino or lepton. Overall, this will lead to somewhat lower geometric acceptance for visible LLP decays in the baseline detector geometry.

Our library of LLP production and decay events for this model can be downloaded here¹⁷ for RHN (U_e), here¹⁸ for RHN (U_μ), and here¹⁹ for RHN (U_τ).

V. RESULTS

We now present the results of our MATHUSLA FastSim simulation studies, obtaining realistic geometric efficiencies for LLP decays as well as updated BSM sensitivity projections for the MATHUSLA baseline geometry of Fig. 1.

A. Geometric acceptance in long lifetime limit

The geometric acceptance for LLP decays is almost entirely independent of the LLP lifetime, except for very short lifetimes where the decays in MATHUSLA are dominated by the highly boosted subpopulation of produced LLPs. We therefore show the geometric acceptances for our three benchmark models in the long-lifetime limit in Fig. 3. This corresponds to lab-frame decay lengths of




FIG. 3. Geometric acceptances $c_{\text{geo},i}^{\text{decay}}$, see Eq. (1), for visible LLP decays in the MATHUSLA baseline detector geometry of Fig. 1 to satisfy the $i = \text{DV2}$ (solid) and DV3 (dashed) criteria in Table I for reconstruction of 2- and 3-pronged displaced vertices in the long-lifetime limit. Top: Exotic Higgs decays to LLPs that decay to either $b\bar{b}$ (red), gg (blue) or $\ell^+\ell^-$ (gray). Middle: LLP in the SM + S model (red). For comparison, we show the acceptance for a simplified S-LLP scenario produced only in B, D -decays and decaying only to e^+e^- (black). Bottom: LLPs in the three RHN scenarios with single-coupling dominance (red, blue, green). For comparison, we show acceptance for a simplified RHN-LLP produced only in B decays and decaying only to $e^+e^-\nu$ (black).

¹⁷github.com/davidrcurtin/MATHUSLA_LLFiles_RHN_Ue.

¹⁸github.com/davidrcurtin/MATHUSLA_LLFiles_RHN_Umu.

¹⁹github.com/davidrcurtin/MATHUSLA_LLFiles_RHN_Utau.

$bc\tau \gg 100$ m, where decay along any point of the LLP trajectory in MATHUSLA is equally likely for all LLPs.

1. Exotic Higgs decays

Figure 3 (top) shows the geometric acceptance for LLPs from exotic Higgs decays. There is very little difference between the two hadronic decay modes $\bar{b}b$ and gg . Furthermore, the large charged particle multiplicity means the acceptance for 2- and 3-pronged DVs is very similar. There is also very little dependence on m_{LLP} . Overall, geometric acceptances are very high, $\xi_{\text{geo}}^{\text{decay}} \sim 0.7\text{--}0.8$. As we discuss below, the reduction from unity is almost entirely due to LLP decays in the back of the detector, where all final states exit the rear wall without ever intersecting any tracking planes.

For comparison, we also show the geometrical acceptance for leptonic LLP decays, which is much lower, 0.45 for $m_{\text{LLP}} = 10$ GeV reducing to 0.2 for 55 GeV. This is again not surprising, given that one of the leptons has a high probability of escaping through the floor or walls undetected, especially as the mean LLP boost is reduced at higher masses.

2. SM + scalar

Figure 3 (middle) shows the geometric acceptance for scalar LLPs in the SM + S model. For comparison, we include $\xi_{\text{geo,DV2}}^{\text{decay}}$ for a simplified scalar LLP that is produced in the same way but only decays to two electrons. We see that the realistic S LLP closely tracks the simplified scenario, except for a slight enhancement as the mass rises above the dimuon and dikaon threshold (since the two relatively heavy final states are more collimated along the original LLP trajectory than final-state electrons) and a significant rise in both the 2- and 3-pronged DV acceptance as multihadron production becomes important for $m_S \gtrsim 2$ GeV. This enhancement at higher masses is important to keep the overall geometric acceptance across most of the mass range at a very reasonable $\xi_{\text{geo,DV2}}^{\text{decay}} \sim 0.5$.

3. Right-handed neutrinos

Figure 3 (bottom) shows the geometric acceptance for RHN LLPs in each of our three coupling-dominated scenarios. For comparison, we include $\xi_{\text{geo,DV2}}^{\text{decay}}$ for a simplified RHN-like LLP that is produced in B, D -decays and only decays to $e^+e^-\nu$.

The simplified comparison scenario immediately illustrates the somewhat lower geometric acceptance of RHNs compared to the other two benchmark models, with $\xi_{\text{geo,DV2}}^{\text{decay}} \sim 0.1\text{--}0.2$ since the invisible neutrino in the 3-body decay makes it more likely that one of the charged final states escapes through the walls or the floor undetected. Fortunately, the presence of a 2-body decay mode to lepton + charged hadron, as well as multihadron

production in the realistic model, enhances the geometric acceptance significantly, leading to $\xi_{\text{geo}}^{\text{decay}} \sim 0.1\text{--}0.4$ for the realistic RHN LLP.

B. Reach for exotic Higgs decays

Figure 4 shows the reach (four expected reconstructed decays) of the MATHUSLA baseline geometry for exotic Higgs decays into LLPs. The result is almost identical for LLPs that decay to $\bar{b}b$ and gg . Even a search that requires DVs with 3+ prongs essentially reproduces the idealized reach computed previously in [4]. Note we do not distinguish between N_{decay} and N_{visible} since $\text{Br}(\text{vis})$ is essentially 1.

The primary physics target of MATHUSLA is therefore extremely robust with respect to both signal reconstruction and potential backgrounds. MATHUSLA would extend the reach of the main detector searches and the reach of the proposed CODEX-b external LLP detector at LHCb by \sim three orders of magnitude in production rate and lifetime.

C. Reach for SM + scalar

Figure 5 shows the reach (4 expected reconstructed decays) of the MATHUSLA baseline geometry for scalar LLPs in the SM + S benchmark model. The idealized reach estimate, corresponding to requiring four visible decays in the detector volume (black dashed contour), agrees mostly with previous estimates [6,28] but displays some differences, notably an expanded reach at larger mixing




FIG. 4. Sensitivity of the MATHUSLA baseline geometry of Fig. 1 to hadronically decaying LLPs produced in exotic Higgs decays. The red curve shows the exclusion reach from a search for DVs with 3+ prongs, which is very close to the idealized estimate of four decays in the decay volume (black dashed). For comparison, we show the current $\text{Br}(h \rightarrow \text{invis})$ limit from ATLAS [54] (purple shading) and the HL-LHC projection [55] (purple line); current ATLAS constraints from searches for 1 or 2 DVs (blue shading) and 2 DVs (green shading) in the muon system [56,57]; projections for an ATLAS 1DV search in the muon system at the HL-LHC [3] (blue dashed), and the idealized sensitivity of CODEX-b for $m_{\text{LLP}} = 10$ GeV [8] (orange dashed).

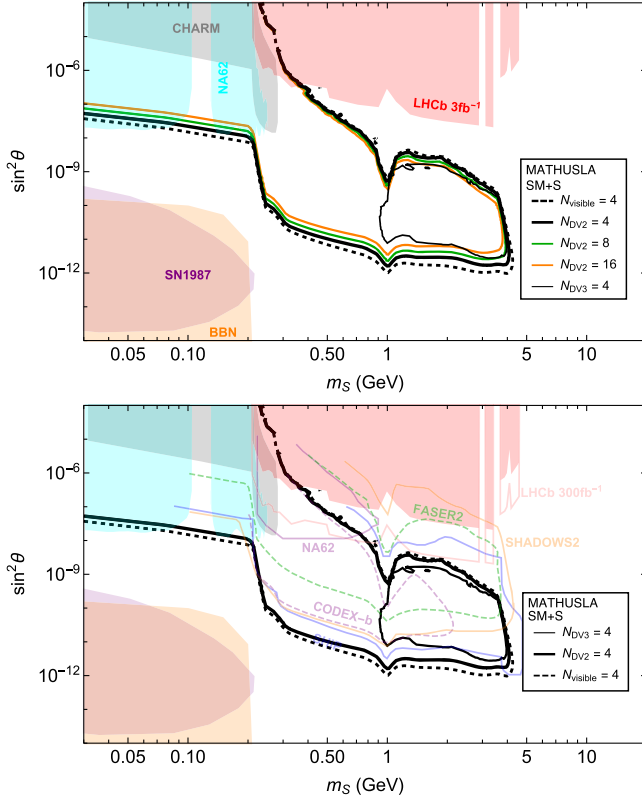


FIG. 5. Top: the sensitivity of the MATHUSLA baseline geometry in Fig. 1 to scalar LLPs in the SM + S model. Black dashed, thick black solid, thin black solid contours; four visible decays in decay volume, reconstructable as 2-pronged DVs, reconstructable as 3-pronged DVs. Hadronization uncertainties mostly affect the near-vertical left boundary of the thin $N_{\text{DV}3}$ contour at the $\mathcal{O}(0.1)$ GeV level. Colored contours show how DV2 signal scales in the parameter space. Shaded areas are bounds from CHARM [58,59], SN1987 [60], BBN [61], NA62 (K -decay) [62–64], and LHCb [65,66]. Bottom: MATHUSLA’s sensitivity compared to projections for NA62 (10^{18} Protons-on-Target beam dump mode) [28], FASER2 [13], LHCb [8], SHADOWS2 [67], CODEX-b [8], and SHiP [28].

angle. This arises due to the updated treatment of the S lifetime that more accurately takes into account final-state hadronic mass thresholds in the scalar decay. The reasonably high geometric acceptance of the MATHUSLA baseline geometry to scalar LLP decay final states means that the $N_{\text{DV}2} = 4$ contour of at least four reconstructable 2-pronged DVs very closely matches the idealized reach estimate. The reach for light scalars is therefore very robust.

In the top plot, we show contours for higher number of reconstructable 2-pronged DVs to illustrate how the number of events scales in this parameter space. We also show the region where four 3-pronged DVs, arising from multi-charged-particle production during hadronization, can be reconstructed. An $\mathcal{O}(1)$ hadronization uncertainty on the multi-hadron production fraction near a GeV would end up having only a very minor impact on these predictions: given the scaling of the signal rate and the small multihadron

production fraction, this would mostly affect the near-vertical left boundary of the thin $N_{\text{DV}3}$ contour at the $\mathcal{O}(0.1)$ GeV level. In the bottom plot, we compare MATHUSLA’s reach to the sensitivity projections for other experiments.^{20,21} MATHUSLA has world-leading reach for small mixing angles, and significant complementarity with detectors like FASER2 and SHADOWS2. Together, MATHUSLA and these proposals would cover the full range of mixing angles all the way up to present LHCb limits.

D. Reach for right-handed neutrinos

Figure 6 shows the reach (four expected reconstructed decays) of MATHUSLA’s baseline geometry for RHN LLPs in the RHN ($U_{e,\mu,\tau}$) benchmark models (solid black line), compared to the idealized reach estimate for four visible decays in the detector volume (black dashed line). As for SM + S, we also show how the number of events scales with parameter space, the reach of a 3-pronged DV search, and comparisons to other proposed experiments.²² RHN production from B , D -meson decays dominates throughout parameter space, except close to the upper edge of the $N_{\text{decay}} = 4$ contours, where W , Z -decay is important and ultimately extends the maximum mass reach. Applying a similar logic as for the SM + S scenario, an $\mathcal{O}(1)$ hadronization uncertainty on the multihadron decay fraction near a GeV would not affect the left boundary of the thin $N_{\text{DV}3}$ contour in the U_e, U_μ scenarios, since that is set by kinematic thresholds, but instead affect the lower reach in mixing angle of the $N_{\text{DV}3}$ contour by at most a similar $\mathcal{O}(1)$ factor, which does not significantly impact our conclusions.

Again, the idealized reach estimate agrees mostly with previous calculations [6,28], but the mass reach is higher while the reach in mixing angle is very slightly reduced, likely due to our more accurate treatment of B , D -meson and LLP decay kinematics, and our inclusion of the electroweak K -factor for RHN production from W , Z decay. Interestingly, MATHUSLA’s mass reach is higher in the RHN (U_τ) case than in the RHN ($U_{e,\mu}$) scenarios.

²⁰Note that amongst the different SM + S projections, there are some significant differences as to how the S production and lifetime is calculated and the exact number of events required for exclusion.

²¹In all our plots, dashed contours correspond to idealized sensitivity projections that are, to the best of our knowledge, based on the number of visible decays in the detector volume.

²²In the RHN (U_τ) scenario, the fraction of RHN decays yielding 4+ charged particles is so low that our simulations are statistically unable to obtain reliable estimates of $N_{\text{DV}3}$ for $m_N \lesssim 0.75$. To obtain a rough estimate, we use a power law to extrapolate $\xi_{\text{geom,DV}3}^{\text{decay}}(m_N) \cdot \text{Br}(\text{vis})$ from higher to lower masses and use it to estimate $N_{\text{DV}3}(m_N, |U_\tau|^2)$ by rescaling $N_{\text{decay}}(m_N, |U_\tau|^2)$. This is indicated by changing the $N_{\text{DV}3} = 4$ contour to a pink dashed line for this mass range.

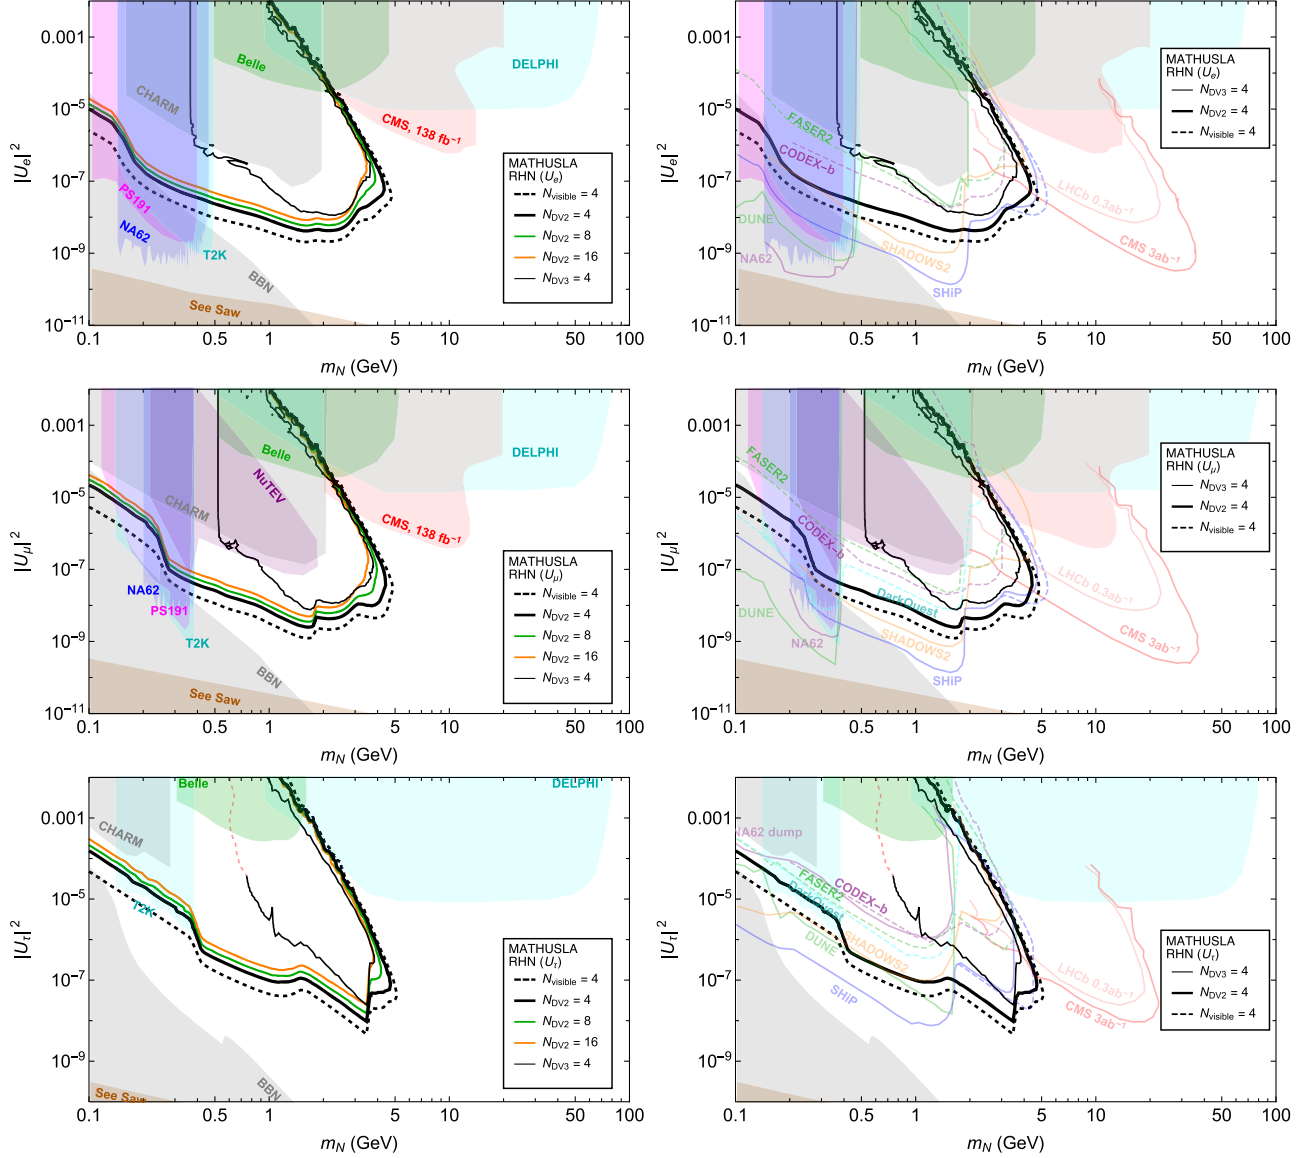


FIG. 6. Left: the sensitivity of the MATHUSLA baseline geometry in Fig. 1 to scalar LLPs in the RHN benchmark models with electron (top), muon (middle) and tau (bottom) coupling dominance. Black dashed, thick black solid, thin black solid contours; four visible decays in decay volume, reconstructable as 2-pronged DVs, reconstructable as 3-pronged DVs. [For RHN (U_τ), the $N_{DV3} = 4$ contour for $m_N < 0.75$ GeV is based on an extrapolation of $\xi_{\text{geom}, DV3}^{\text{decay}}(m_N) \cdot \text{Br}(\text{vis})$ from higher masses, indicated by the pink dashed contour, see text details.] Hadronization uncertainties mostly affect the lower boundary of the thin N_{DV3} contour near a GeV, at most at the $\mathcal{O}(1)$ level. Colored contours show how DV2 signal scales in the parameter space. Shaded areas are bounds from BBN [68], DELPHI [69], Charm [59], Belle [70], CMS [71], PS191 [72], NA62 (K -decay) [73,74], DUNE near detector, and T2K [75]. Right: MATHUSLA's sensitivity compared to projections for LHCb and CMS [76], FASER2 [13], SHADOWS2 [67], CODEX-b [8], SHiP [28], DarkQuest [77], NA62 (K -decay, 5×10^9 Protons-on-Target) [73,74,78], NA62 (10^{18} Protons-on-Target beam dump mode) [79], and the DUNE near detector [80].

This can be understood as arising from differences in the RHN lifetime near $m_N \sim 5$ GeV. In the U_τ case, the heavier mass of the tau in the final state reduces the RHN decay width, resulting in a $1.75\times$ longer decay length than in the U_e, U_μ cases. This in turn allows the mixing angle and hence production rate to be larger while keeping the decay length in the optimal range for decay in MATHUSLA.

The somewhat lower-geometric acceptance for RHN LLPs compared to the other benchmarks reduces the reach in $|U_{e,\mu}|^2$ compared to the idealized reach by a factor of $\sim 2-3$, which is consistent with the long-lifetime $\sim |U|^4$ scaling of the number of events near the lower sensitivity boundary. However, MATHUSLA's reach for all RHN scenarios is not qualitatively changed, and is amongst the most competitive of the various proposed experiments.

VI. IMPACT OF DIFFERENT MATHUSLA DESIGN ASPECTS

The results of the previous section indicate that the MATHUSLA baseline geometry of Fig. 1 has robust sensitivity to primary and secondary physics case LLP signals. However, the results also suggest that some parts of the detector design might be modified to significantly increase signal yield. Here, we investigate the impact of different detector design aspects on the geometric LLP acceptance and resulting BSM reach, in the hope that this can inform future realistic detector optimization by the experimental collaboration.²³

A. Local trigger

A local trigger is considered as part of the MATHUSLA design [17] to facilitate data acquisition and supply a trigger signal to CMS, which would allow for correlated analyses that can determine the LLP production, mass and other underlying parameters of the BSM scenario with as few as 10–100 detected LLPs [81]. While the details of this trigger are not yet determined, and investigating its realistic effect is beyond our scope, we can nevertheless ask whether a *local* version of the geometric track and vertex reconstruction criteria in Table I would significantly reduce the LLP acceptance.

As discussed in Sec. III, the MATHUSLA FastSim is able to implement a toy-trigger criterion, whereby a DV is only reconstructed if at least one of the tracks has some number of hits (in our case 4) within some local area of the detector, e.g., an arbitrary subset of 3×3 neighboring modules. We have repeated all the efficiency and reach studies of the previous section with this trigger criterion enabled, and find no significant difference in any of our results; the maximum reduction in $\xi_{\text{geo},i}^{\text{decay}}$ for some masses is of order of a few percent. We can therefore be reasonably confident that realistic local trigger strategies should not significantly impact the physics reach of MATHUSLA, and we can ignore trigger requirements in future phenomenological studies.

B. Spatial distribution of reconstructable decays

The maximum geometric acceptance $\xi_{\text{geo}}^{\text{decay}} \sim 0.7\text{--}0.8$ of the baseline detector geometry for hadronically decaying LLPs with $\mathcal{O}(10)$ charged final states is (very) naively puzzling; it is difficult to imagine how such a decay would not yield at least a few tracks. The culprit is obviously the back of the detector, where all daughter particles of LLPs exit through the rear wall without ever intersecting a tracking plane. This is a necessary consequence of the

²³We gratefully acknowledge the contributions of Lillian Luo and Wentao Cui to early versions of the investigations in this section.

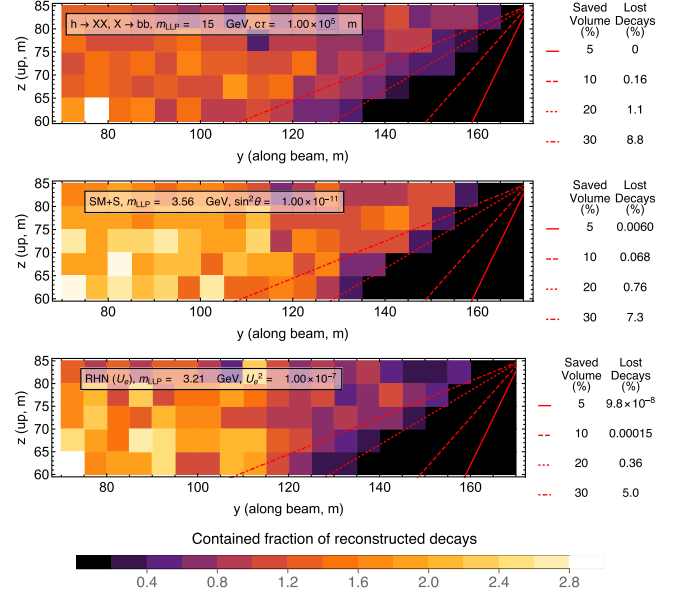


FIG. 7. Spatial distributions of reconstructed decay vertices (DV2 criterion) in the MATHUSLA baseline geometry of Fig. 1 for representative examples of the exotic Higgs decay, SM + S and RHN (U_e) benchmark models in the long lifetime limit, projected onto the vertical plane that intersects the LHC beamline. (The IP is situated at the origin of this coordinate system.) The red solid, dashed, dotted, dot-dashed lines represent optimized choices of volume cuts below lines starting on the floor and ending on the rear wall, which are constrained to reduce the detector volume by 5%, 10%, 20%, and 30% respectively while minimizing the fraction of “lost” reconstructed decays below the line.

fact that the LLPs we are most interested in are usually produced at the LHC with boost $b \gtrsim \mathcal{O}(\text{few} - 10)$.

We therefore expect this near-zero geometric acceptance for decays in the back of MATHUSLA to be a universal phenomenon for all benchmark models, and we can illustrate this with representative spatial distributions of LLP decays in Fig. 7 in the long lifetime limit.²⁴ The low acceptance is clearly visible by eye in the rear corner. We can investigate what fraction of the decay volume could be removed while minimizing the lost LLP signal. For each of the three scenarios in Fig. 7, the red lines indicate optimized cuts to reduce the decay volume by 5%, 10%, 20%, and 30%. The volume can be reduced by essentially the same 20% cut for all processes while only losing $\sim 1\%$ of the reconstructed signal, while a 30% volume reduction is possible with a $\sim 5\text{--}8\%$ signal loss. This indicates that about a quarter of the baseline MATHUSLA decay volume is not directly used for physics (though we emphasize that the tracking planes in the rear of the detector are very much used to reconstruct LLP decays in the detector’s interior).

²⁴For shorter decay lengths, vertices are concentrated near the front of the detector, and the impact of the low-acceptance part of the MATHUSLA decay volume is significantly reduced.

C. Tracker coverage

We now investigate how hypothetical extensions of MATHUSLA's tracker coverage could enhance LLP sensitivity. Specifically, we study three questions:

- (i) What is the impact of the gaps between detector modules? To this end, we can define an idealized detector that is identical to the baseline design of Fig. 1, except without any gaps in the eight ceiling tracker layers and compare its geometric acceptance and BSM reach to the baseline;
- (ii) How much can we gain by instrumenting the rear wall? Our understanding of the spatial distribution of reconstructable LLP decays in the previous subsection suggests that instrumenting the rear wall with a full tracker could restore $\zeta_{\text{geom}}^{\text{decay}} \rightarrow 1$ for hadronically decaying LLPs from exotic Higgs decays, and presumably significantly enhance the reach for the low-mass LLP benchmarks as well. To this end, we define a detector geometry that is identical to the baseline, except that the ten modules in the back row are equipped with a five-layer tracker, similar to the ceiling tracker stack, that covers their rear wall. Like the ceiling tracker, the rear tracker has 1m gaps between the modules;
- (iii) Taking this a step further, we can also ask how much one could gain by instrumenting all four walls with such a wall tracker.

These geometries are easily defined within the MATHUSLA FastSim, and we repeat all of our previous studies for these three comparison geometries.

The geometric LLP acceptances are compared in Fig. 8. Unsurprisingly, each of the alternate detector geometries with enhanced tracker coverage has higher geometric acceptance than the baseline. However, there are several important qualitative takeaways.

Eliminating the module gaps has only a very minor impact on all acceptances. The 1m gap between modules of the baseline design is therefore not a limiting factor on LLP sensitivity.

Both the back-tracker and full-wall tracker geometries significantly improve the reconstruction of all LLPs, but in different ways for the different benchmarks. For LLPs from exotic Higgs decays, the back tracker is sufficient to raise the geometric acceptance to close to unity. This is a modest but significant $\sim 30\%$ improvement in reach, which would allow MATHUSLA to realize the full idealized reach for LLPs from exotic Higgs decays in Fig. 4. For the RHN and SM + S models, the back tracker significantly improves $\zeta_{\text{geom}}^{\text{decay}}$, but a roughly equal factor of improvement is gained by the geometry with the full 4-wall tracker: it realizes acceptances close to 1 for the SM + S model, and doubles the acceptance for RHN LLPs compared to the baseline. This is reflected in the hypothetical reach of these three enhanced geometries to SM + S and RHN models shown

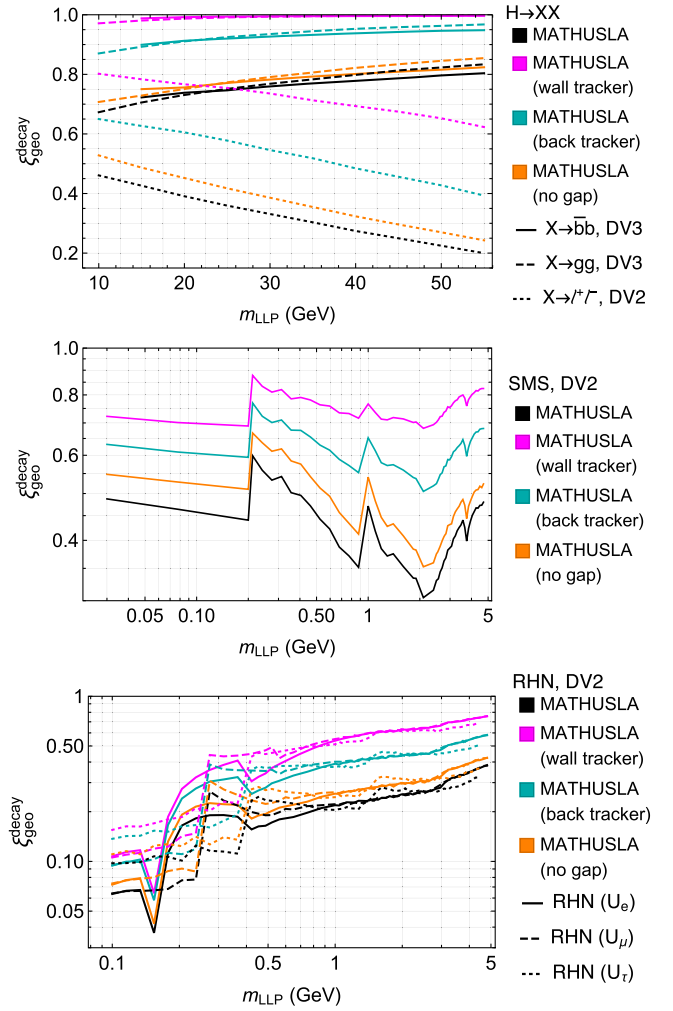


FIG. 8. Geometric acceptances for visible LLP decays, $\zeta_{\text{geo},i}^{\text{decay}}$, for LLP decays in the exotic Higgs decay, SM + S and RHN benchmark models, comparing the baseline MATHUSLA geometry of Fig. 1 (black) to the same geometry with a full-wall tracker (magenta), a back-wall tracker only (dark cyan) and no gaps between modules (orange).

in Fig. 9. In particular, the RHN sensitivities are moved much closer to the idealized $N_{\text{decay}} = 4$ limit.

Obviously, increasing tracker coverage would come with significant increases in cost and complexity, but it is important to understand that (1) the module gap is not a significant consideration, (2) there are highly significant sensitivity gains to be made by making judicious tracker additions, and (3) these gains arise partially by making decays in the back of the detector reconstructable. In a hypothetical detector optimization exercise, this would have to be weighed against other ways of affecting the detector physics reach and cost. One of the most obvious such factors is detector size.

D. Detector area

We first illustrate how the idealized reach of MATHUSLA scales with its area. We keep its footprint

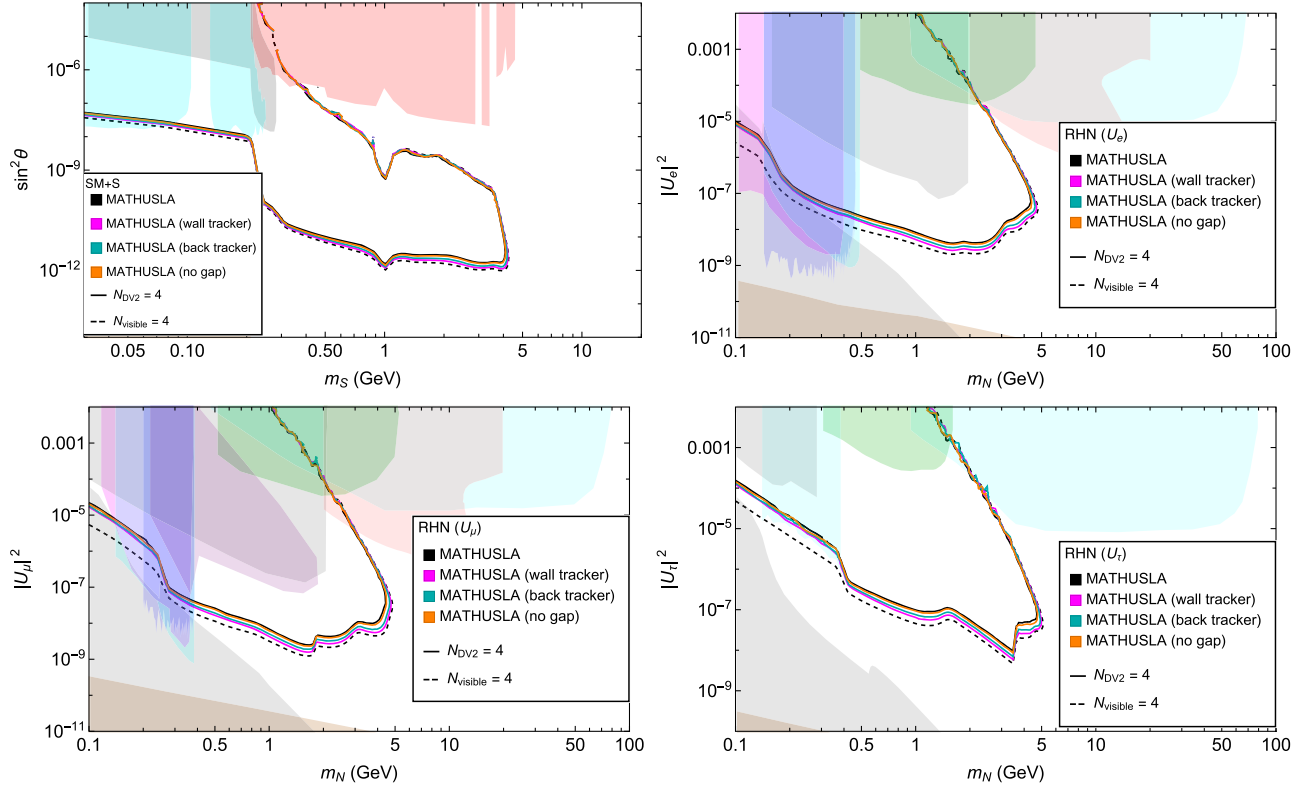


FIG. 9. Comparing the sensitivity of the MATHUSLA baseline geometry in Fig. 1 to the same geometry with a full-wall tracker (magenta), a back-wall tracker only (dark cyan) and no gaps between modules (orange) for the SM + S and RHN benchmark models. We do not show the corresponding comparison for hadronically decaying LLPs produced in exotic Higgs decays, since full and back wall tracker geometries essentially fully recover the ideal $N_{\text{decay}} = 4$ sensitivity in Fig. 4, and the geometry without module gaps performs identically to the baseline.

square, keep it centered on the beamline, keep the front wall at the same location on the surface, and imagine these smaller versions as still being made up of $9 \text{ m} \times 9 \text{ m}$ modules that are 1 m apart from each other. Again, these alternate geometries are easily simulated and studied in the MATHUSLA FastSim.

An obvious but important point is that MATHUSLA’s idealized sensitivity always scales better than linear with smaller detector area. This is illustrated in Fig. 10 for exotic Higgs decays, but similar scalings persist for the other models, since LLP decays are concentrated near the front of the detector, even in the long-lifetime limit [82]. For example, an $(80 \text{ m})^2$ detector has 0.64 times the area of the baseline but catches $>70\%$ of the LLP decays, with some slight variation depending on the model and LLP mass. This effect is even more pronounced for smaller areas, and makes for a strikingly modest reduction in idealized reach for the smaller MATHUSLA versions, as illustrated in Fig. 11.

We also explicitly computed the geometric acceptances of the $(60 \text{ m})^2$ and $(80 \text{ m})^2$ geometries with full-wall trackers, in analogy to the analysis of the previous subsection, and as expected we found that it very closely replicates the geometric acceptances of the full size

geometry with the full wall tracker shown in Fig. 8. This is useful information for future detector optimization, as the high signal acceptance of MATHUSLA with a wall tracker is roughly independent of detector size, allowing us

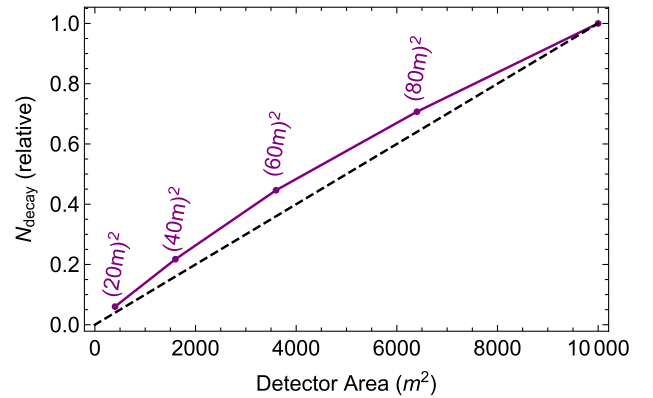


FIG. 10. Purple curve: number of LLP decays in the decay volume vs area of a MATHUSLA-like detector in the long lifetime limit, normalized to the $(100 \text{ m} \times 100 \text{ m})$ benchmark of Fig. 1. The dashed line indicates linear scaling. This is for a 15 GeV LLP produced in exotic Higgs decays, but similar scalings are observed for the other benchmark models across their parameter space.

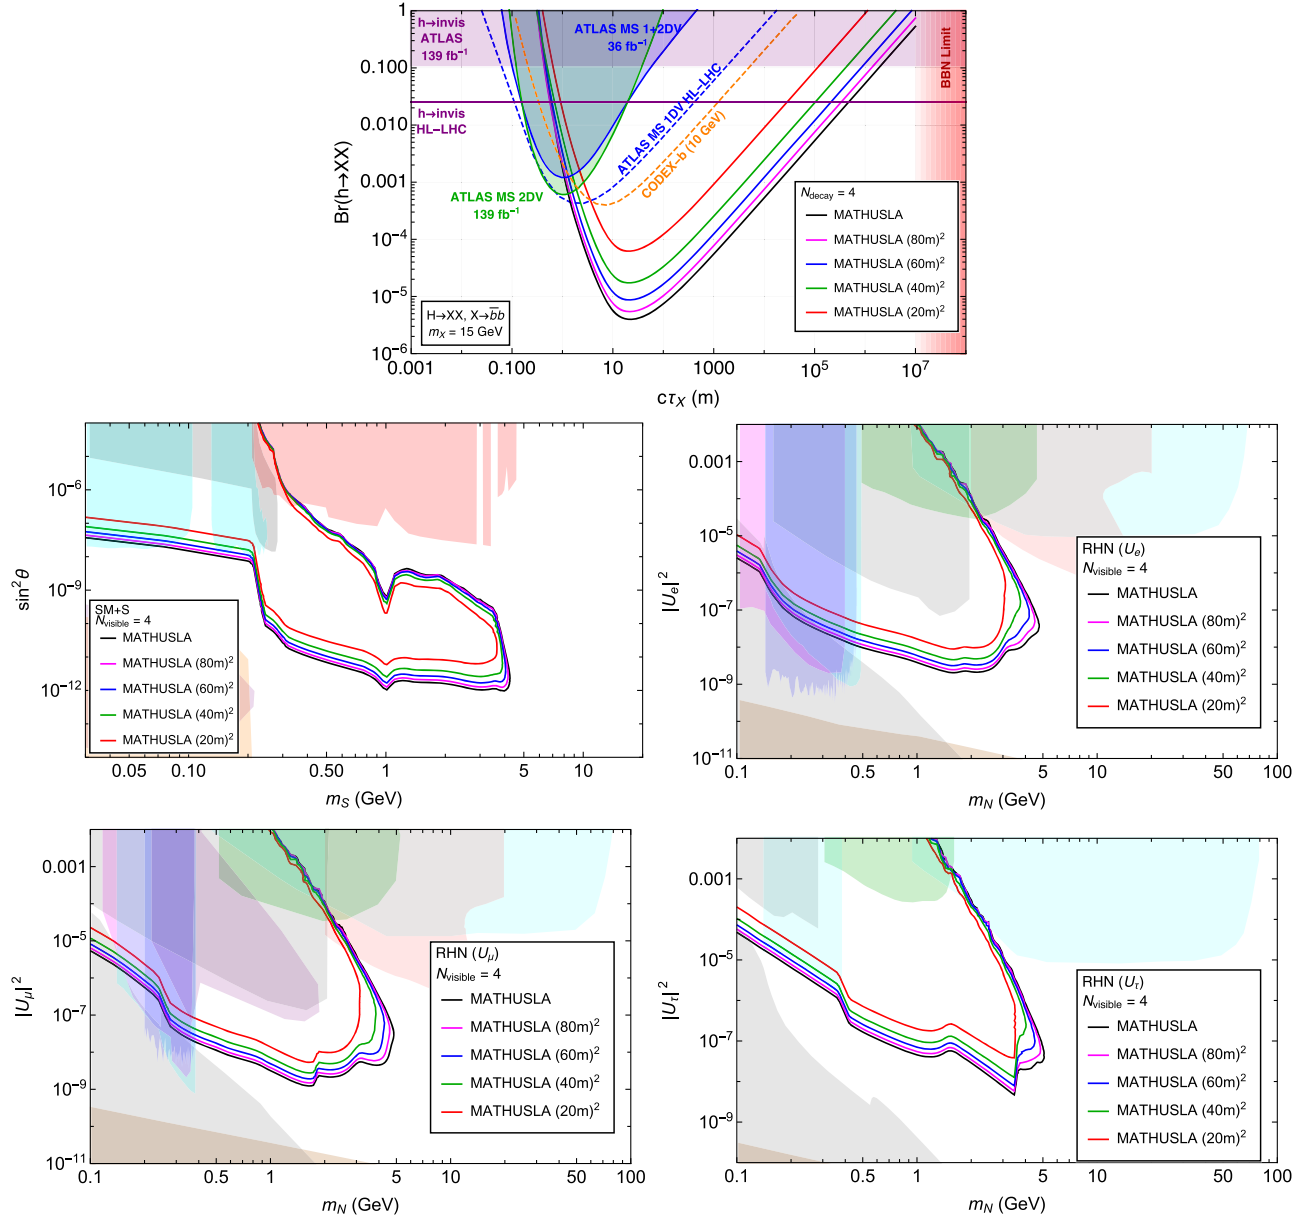


FIG. 11. Comparing the idealized reach of the MATHUSLA baseline geometry in Fig. 1 to similarly situated versions of MATHUSLA with smaller area.

to apply the rough scaling of Fig. 10 to obtain realistic sensitivity estimates for fully instrumented versions of MATHUSLA for different areas.

E. Discussion

Ultimately, the task of optimizing details of the detector design is a complicated exercise that has to take into account many practical considerations beyond those we can investigate here. Even so, we hope that the lessons learned from our purely geometrical detector simulations can inform these decisions.

Our results invite more detailed investigation of several possible tweaks to the MATHUSLA baseline geometry of

Fig. 1. Most importantly, one could maximize physics coverage by instrumenting the rear wall, or all four walls, with a multilayer tracker. It is conceivable that this could be done without increasing the overall cost of the experiment by accepting a smaller detector size. For example, an (80 m)² detector would catch roughly 70% of the LLP decays of the current baseline, but equipping all the walls with tracker would raise the geometric acceptance for hadronically decaying LLPs with $\mathcal{O}(10\text{--}100\text{ GeV})$ masses to order unity, essentially preserving the realistic reach of the baseline design shown in Fig. 4. The reach for SM + S models would be slightly enhanced, while the reconstructable signal yield for RHN models and other LLP decays

with invisible particles in the final state would be doubled. This makes it especially interesting to consider that the total amount of tracker material needed by a $(80\text{ m})^2$ detector with full tracker coverage for all four walls is the same as for the baseline geometry with the ceiling tracker only.

If such a path could be realized, it could make MATHUSLA a much more powerful general LLP discovery machine, while keeping the total amount of instrumentation equal to the baseline and potentially reducing civil engineering costs significantly. Optimizing the final design of the detector will rely on a myriad of experimental, technical, and practical engineering considerations that are far beyond our scope, but investigating this possibility should clearly be a high priority.

VII. CONCLUSIONS

In this paper, we studied several detailed aspects of LLP decays in the proposed MATHUSLA detector. We introduced the publicly available MATHUSLA FastSim for easy signal estimates at MATHUSLA and other external LLP detectors, taking realistic geometric acceptances for the LLP decay products into account, see Eq. (1). We studied three important benchmark models: hadronically decaying LLPs from exotic Higgs decays to represent MATHUSLA's primary physics target; and the SM + S and RHN simplified scenarios to represent the secondary physics target of GeV-scale LLPs that are also the target of many proposed fixed-target and intensity frontier searches. For all of these benchmark models, we carefully simulate their production at the HL-LHC and their subsequent decay, and our full simulation library is publicly available to facilitate future studies. We then systematically investigate the geometric acceptances of the MATHUSLA baseline design of Fig. 1 for these LLP scenarios. For GeV-scale LLPs, it is important to include multihadron production in their decay, since this can significantly increase acceptance.

We find highly robust acceptance for hadronically decaying LLPs with $\mathcal{O}(10)$ tracks, $\xi_{\text{geom}}^{\text{decay}} \sim 0.7\text{--}0.8$, good acceptance for decaying scalar LLPs $\xi_{\text{geom}}^{\text{decay}} \sim 0.3\text{--}0.5$ and somewhat reduced acceptance for RHN LLPs $\xi_{\text{geom}}^{\text{decay}} \sim 0.1\text{--}0.4$, see Fig. 3. We also find that hadronic uncertainties on LLP decays with GeV masses have minimal impact on our results, since the predicted branching fraction to 4+ charged states only becomes significant for masses above 2–3 GeV.

We present updated realistic sensitivity estimates for the MATHUSLA baseline design in Figs. 4–6. The idealized reach computed in earlier estimates is well-realized; MATHUSLA has world-leading reach for hadronically decaying LLPs from exotic Higgs decays by several orders of magnitude, and is highly competitive for GeV-scale scalar and RHN LLPs as well.

Finally, we work to understand which aspects of the MATHUSLA baseline design most influence the geometric acceptance for LLP decays. Trigger considerations and module gaps are found to not be limiting factors for LLP sensitivity. About a quarter of the baseline design's detector volume has near-zero acceptance due to the boost of LLPs decaying in the rear of the detector, and LLP decays with invisible particles in the final state generally have somewhat reduced acceptance due to the higher chance of charged particles escaping through the walls and floor. This invites consideration of instrumenting the walls with tracker, possibly accompanied by a modest decrease in detector size to respect cost and other constraints, in order to optimize the final design of MATHUSLA. This favorable flexibility arises due to the practically background-free nature of MATHUSLA's location on the surface, underlining the robustness of the basic detector concept and its generality as a world-leading future LLP discovery instrument.

ACKNOWLEDGMENTS

We would like to thank Brian Batell, John Paul Chou, Miriam Diamond, Abdulrahman Emad, Simon Knapen, Henry Lubatti, Runze Ren, Dean Robinson, Steven Robertson, Richard Ruiz, and Charles Young for helpful discussions during the completion of this project. We especially acknowledge the contributions of Lillian Luo and Wentao Cui, who wrote the first version of the MATHUSLA FastSim code used for earlier detector geometry comparison studies with D. C. The research of D. C. and J. S. G. was supported in part by Discovery Grants from the Natural Sciences and Engineering Research Council of Canada and the Canada Research Chair program. The research of D. C. was also supported by the Alfred P. Sloan Foundation, the Ontario Early Researcher Award, and the University of Toronto McLean Award. This research was enabled by computing resources and support provided by Compute Ontario (computeontario.ca, <http://computeontario.ca/>) and the Digital Research Alliance of Canada (alliance-can.ca, <http://alliance-can.ca/>).

- [1] D. Curtin *et al.*, *Rep. Prog. Phys.* **82**, 116201 (2019).
- [2] J. Alimena *et al.*, *J. Phys. G* **47**, 090501 (2020).
- [3] A. Coccaro, D. Curtin, H. J. Lubatti, H. Russell, and J. Shelton, *Phys. Rev. D* **94**, 113003 (2016).
- [4] J. P. Chou, D. Curtin, and H. J. Lubatti, *Phys. Lett. B* **767**, 29 (2017).
- [5] C. Alpigiani *et al.* (MATHUSLA Collaboration), arXiv:1811.00927.
- [6] C. Alpigiani *et al.* (MATHUSLA Collaboration), arXiv:2009.01693.
- [7] V. V. Gligorov, S. Knapen, M. Papucci, and D. J. Robinson, *Phys. Rev. D* **97**, 015023 (2018).
- [8] G. Aielli *et al.*, *Eur. Phys. J. C* **80**, 1177 (2020).
- [9] M. Bauer, O. Brandt, L. Lee, and C. Ohm, arXiv:1909.13022.
- [10] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, *Phys. Rev. D* **97**, 035001 (2018).
- [11] A. Ariga *et al.* (FASER Collaboration), *Phys. Rev. D* **99**, 095011 (2019).
- [12] L. A. Anchordoqui *et al.*, *Phys. Rep.* **968**, 1 (2022).
- [13] J. L. Feng *et al.*, *J. Phys. G* **50**, 030501 (2023).
- [14] S. Cerci *et al.*, *J. High Energy Phys.* **06** (2022) 110.
- [15] B. Bhattacharjee, S. Matsumoto, and R. Sengupta, *Phys. Rev. D* **106**, 095018 (2022).
- [16] B. Bhattacharjee, H. K. Dreiner, N. Ghosh, S. Matsumoto, R. Sengupta, and P. Solanki, arXiv:2306.11803.
- [17] C. Alpigiani *et al.* (MATHUSLA Collaboration), in Snowmass 2021 (2022), arXiv:2203.08126.
- [18] M. Ovchinnikov, J.-L. Tastet, O. Mikulenko, and K. Bondarenko, *Phys. Rev. D* **108**, 075028 (2023).
- [19] W. Cui, D. Curtin, L. Luo, and J. S. Grewal, MATHUSLA FastSim, <https://zenodo.org/records/8235028> (2023).
- [20] D. Curtin *et al.*, *Phys. Rev. D* **90**, 075004 (2014).
- [21] N. Craig, A. Katz, M. Strassler, and R. Sundrum, *J. High Energy Phys.* **07** (2015) 105.
- [22] D. Curtin and C. B. Verhaaren, *J. High Energy Phys.* **12** (2015) 072.
- [23] Z. Liu, L.-T. Wang, and H. Zhang, *Chin. Phys. C* **41**, 063102 (2017).
- [24] S. Alipour-Fard, N. Craig, M. Jiang, and S. Koren, *Chin. Phys. C* **43**, 053101 (2019).
- [25] J. Kozaczuk, M. J. Ramsey-Musolf, and J. Shelton, *Phys. Rev. D* **101**, 115035 (2020).
- [26] M. Cepeda, S. Gori, V. M. Outchoorn, and J. Shelton, *Annu. Rev. Nucl. Part. Sci.* **72**, 119 (2022).
- [27] D. Curtin and M. E. Peskin, *Phys. Rev. D* **97**, 015006 (2018).
- [28] J. Beacham *et al.*, *J. Phys. G* **47**, 010501 (2020).
- [29] D. O'Connell, M. J. Ramsey-Musolf, and M. B. Wise, *Phys. Rev. D* **75**, 037701 (2007).
- [30] M. Drewes, *Int. J. Mod. Phys. E* **22**, 1330019 (2013).
- [31] S. Agostinelli *et al.* (GEANT4 Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **506**, 250 (2003).
- [32] Z. Chacko, H.-S. Goh, and R. Harnik, *Phys. Rev. Lett.* **96**, 231802 (2006).
- [33] B. Holdom, *Phys. Lett.* **166B**, 196 (1986).
- [34] D. Curtin, R. Essig, S. Gori, and J. Shelton, *J. High Energy Phys.* **02** (2015) 157.
- [35] B. Batell, J. L. Feng, M. Fieg, A. Ismail, F. Kling, R. M. Abraham, and S. Trojanowski, *Phys. Rev. D* **105**, 075001 (2022).
- [36] J. Jaeckel and A. Ringwald, *Annu. Rev. Nucl. Part. Sci.* **60**, 405 (2010).
- [37] Y. Gershtein, S. Knapen, and D. Redigolo, *Phys. Lett. B* **823**, 136758 (2021).
- [38] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, *J. High Energy Phys.* **07** (2014) 079.
- [39] T. Sjostrand, S. Mrenna, and P. Z. Skands, *J. High Energy Phys.* **05** (2006) 026.
- [40] T. Sjostrand, S. Mrenna, and P. Z. Skands, *Comput. Phys. Commun.* **178**, 852 (2008).
- [41] G. Bozzi, S. Catani, D. de Florian, and M. Grazzini, *Nucl. Phys.* **B737**, 73 (2006).
- [42] D. de Florian, G. Ferrera, M. Grazzini, and D. Tommasini, *J. High Energy Phys.* **11** (2011) 064.
- [43] D. de Florian *et al.* (LHC Higgs Cross Section Working Group), 10.23731/CYRM-2017-002 (2016).
- [44] I. Boiarska, K. Bondarenko, A. Boyarsky, V. Gorkavenko, M. Ovchinnikov, and A. Sokolenko, *J. High Energy Phys.* **11** (2019) 162.
- [45] M. Cacciari, M. Greco, and P. Nason, *J. High Energy Phys.* **05** (1998) 007.
- [46] M. Cacciari, S. Frixione, and P. Nason, *J. High Energy Phys.* **03** (2001) 006.
- [47] M. Cacciari, S. Frixione, N. Houdeau, M. L. Mangano, P. Nason, and G. Ridolfi, *J. High Energy Phys.* **10** (2012) 137.
- [48] R. Aaij *et al.* (LHCb Collaboration), *Phys. Rev. Lett.* **124**, 122002 (2020).
- [49] K. Bondarenko, A. Boyarsky, D. Gorbunov, and O. Ruchayskiy, *J. High Energy Phys.* **11** (2018) 032.
- [50] A. Atre, T. Han, S. Pascoli, and B. Zhang, *J. High Energy Phys.* **05** (2009) 030.
- [51] D. Alva, T. Han, and R. Ruiz, *J. High Energy Phys.* **02** (2015) 072.
- [52] C. Degrande, O. Mattelaer, R. Ruiz, and J. Turner, *Phys. Rev. D* **94**, 053002 (2016).
- [53] G. Aad *et al.* (ATLAS Collaboration), *Phys. Lett. B* **759**, 601 (2016).
- [54] The ATLAS Collaboration, *Phys. Lett. B* **842**, 137963 (2023).
- [55] *CERN Yellow Reports: Monographs*, edited by A. Dainese, M. Mangano, A. B. Meyer, A. Nisati, G. Salam, and M. A. Vesterinen (CERN, Geneva, Switzerland, 2019), Vol. 7/2019.
- [56] M. Aaboud *et al.* (ATLAS Collaboration), *Phys. Rev. D* **99**, 052005 (2019).
- [57] G. Aad *et al.* (ATLAS Collaboration), *Phys. Rev. D* **106**, 032005 (2022).
- [58] M. W. Winkler, *Phys. Rev. D* **99**, 015018 (2019).
- [59] F. Bergsma *et al.* (CHARM Collaboration), *Phys. Lett.* **157B**, 458 (1985).
- [60] S. Balaji, P. S. B. Dev, J. Silk, and Y. Zhang, *J. Cosmol. Astropart. Phys.* **12** (2022) 024.
- [61] A. Fradette and M. Pospelov, *Phys. Rev. D* **96**, 075033 (2017).
- [62] E. Cortina Gil *et al.* (NA62 Collaboration), *J. High Energy Phys.* **03** (2021) 058.
- [63] E. Cortina Gil *et al.* (NA62 Collaboration), *J. High Energy Phys.* **06** (2021) 093.
- [64] E. Cortina Gil *et al.* (NA62 Collaboration), *J. High Energy Phys.* **02** (2021) 201.

- [65] R. Aaij *et al.* (LHCb Collaboration), *Phys. Rev. Lett.* **115**, 161802 (2015).
- [66] R. Aaij *et al.* (LHCb Collaboration), *Phys. Rev. D* **95**, 071101 (2017).
- [67] W. Baldini *et al.*, [arXiv:2110.08025](https://arxiv.org/abs/2110.08025).
- [68] A. Boyarsky, M. Ovchinnikov, O. Ruchayskiy, and V. Syvolap, *Phys. Rev. D* **104**, 023517 (2021).
- [69] P. Abreu *et al.* (DELPHI Collaboration), *Z. Phys. C* **74**, 57 (1997); **75**, 580(E) (1997).
- [70] D. Liventsev *et al.* (Belle Collaboration), *Phys. Rev. D* **87**, 071102 (2013); **95**, 099903(E) (2017).
- [71] C. Alpigiani *et al.* (MATHUSLA Collaboration), MATHUSLA conceptual design report, 2023.
- [72] G. Bernardi *et al.*, *Phys. Lett. B* **203**, 332 (1988).
- [73] E. Cortina Gil *et al.* (NA62 Collaboration), *Phys. Lett. B* **807**, 135599 (2020).
- [74] E. Cortina Gil *et al.* (NA62 Collaboration), *Phys. Lett. B* **816**, 136259 (2021).
- [75] K. Abe *et al.* (T2K Collaboration), *Phys. Rev. D* **100**, 052006 (2019).
- [76] M. Drewes and J. Hajer, *J. High Energy Phys.* 02 (2020) 070.
- [77] B. Batell, J. A. Evans, S. Gori, and M. Rai, *J. High Energy Phys.* 05 (2021) 049.
- [78] A. M. Abdullahi *et al.*, *J. Phys. G* **50**, 020501 (2023).
- [79] M. Drewes, J. Hajer, J. Klaric, and G. Lanfranchi, *J. High Energy Phys.* 07 (2018) 105.
- [80] M. Breitbach, L. Buonocore, C. Frugiuele, J. Kopp, and L. Mitnacht, *J. High Energy Phys.* 01 (2022) 048.
- [81] J. Barron and D. Curtin, *J. High Energy Phys.* 12 (2020) 061.
- [82] I. Alkhatib, [arXiv:1909.05896](https://arxiv.org/abs/1909.05896).