

Long-Lived Particles at the Energy Frontier: The MATHUSLA Physics Case

Editors:

*David Curtin¹, Marco Drewes², Matthew McCullough³, Patrick Meade⁴, Rabindra Mohapatra⁵,
 Jessie Shelton⁶, Brian Shuve^{7,8}.*

Contributors:

*Elena Accomando⁹, Cristiano Alpigiani¹⁰, Stefan Antusch¹¹, Juan Carlos Arteaga-Velázquez¹²,
 Brian Batell¹³, Martin Bauer¹⁴, Nikita Blinov⁸, Karen Salomé Caballero-Mora^{15,16}, Jae Hyeok
 Chang⁴, Eung Jin Chun¹⁷, Raymond T. Co¹⁸, Timothy Cohen¹⁹, Peter Cox²⁰, Nathaniel Craig²¹,
 Csaba Csáki²², Yanou Cui²³, Francesco D’Eramo²⁴, Luigi Delle Rose²⁵, P. S. Bhupal Dev²⁶, Keith
 R. Dienes^{27,5}, Jeff A. Dror^{28,29}, Rouven Essig⁴, Jared A. Evans^{30,6}, Jason L. Evans¹⁷, Arturo
 Fernández Tellez³¹, Oliver Fischer³², Thomas Flacke³³, Anthony Fradette³⁴, Claudia Frugiuele³⁵,
 Elina Fuchs³⁵, Tony Gherghetta³⁶, Gian F. Giudice³, Dmitry Gorbunov^{37,38}, Rick S. Gupta³⁹,
 Claudia Hagedorn⁴⁰, Lawrence J. Hall^{28,29}, Philip Harris⁴¹, Juan Carlos Helo^{42,43}, Martin
 Hirsch⁴⁴, Yonit Hochberg⁴⁵, Anson Hook⁵, Alejandro Ibarra^{46,17}, Seyda Ipek⁴⁷, Sunghoon Jung⁴⁸,
 Simon Knapen^{29,28}, Eric Kuflik⁴⁵, Zhen Liu⁴⁹, Salvator Lombardo²², Henry Lubatti¹⁰, David
 McKeen⁵⁰, Emiliano Molinaro⁵¹, Stefano Moretti^{9,52}, Natsumi Nagata⁵³, Matthias Neubert^{54,22},
 Jose Miguel No^{55,56}, Emmanuel Olaiya⁵², Gilad Perez³⁵, Michael E. Peskin⁸, David Pinner^{57,58},
 Maxim Pospelov^{59,34}, Matthew Reece⁵⁷, Dean J. Robinson³⁰, Mario Rodríguez Cahuantzi³¹,
 Rinaldo Santonico⁶⁰, Matthias Schlaffer³⁵, Claire H. Shepherd-Themistocleous⁵², Andrew Spray³³,
 Daniel Stolarski⁶¹, Martin A. Subieta Vasquez^{62,63}, Raman Sundrum⁵, Andrea Thamm³, Brooks
 Thomas⁶⁴, Yuhsin Tsai⁵, Brock Tweedie¹³, Stephen M. West⁶⁵, Charles Young⁸, Felix Yu⁵⁴, Bryan
 Zaldivar^{55,66}, Yongchao Zhang^{26,67}, Kathryn Zurek^{29,28,3}, José Zurita^{32,68}.*

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

² Centre for Cosmology, Particle Physics and Phenomenology, Université Catholique de Louvain, Louvain-la-Neuve, B-1348, Belgium

³ CERN, TH Department, CH-1211 Geneva, Switzerland

⁴ C.N. Yang Institute for Theoretical Physics, Stony Brook University, Stony Brook, NY 11794, USA

⁵ Maryland Center for Fundamental Physics, Department of Physics, University of Maryland, College Park, MD 20742-4111 USA

⁶ Department of Physics, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA

⁷ Harvey Mudd College, 301 Platt Blvd., Claremont, CA 91711, USA

⁸ SLAC, Stanford University, Menlo Park, California 94025 USA

- ⁹ School of Physics and Astronomy, University of Southampton, Highfield, Southampton SO17 1BJ, United Kingdom
- ¹⁰ Department of Physics, University of Washington, Seattle, WA 98195, USA
- ¹¹ Department of Physics, University of Basel, Klingelbergstrasse 82, CH-4056 Basel, Switzerland
- ¹² Instituto de Física y Matemáticas, Universidad Michoacana de San Nicolás de Hidalgo, Edificio C3, Ciudad Universitaria 58030, Morelia, Michoacán, México
- ¹³ Pittsburgh Particle Physics, Astrophysics, and Cosmology Center, Department of Physics and Astronomy, University of Pittsburgh, PA 15260, USA
- ¹⁴ Institut für Theoretische Physik, Universität Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany
- ¹⁵ Facultad de Ciencias en Física y Matemáticas, Universidad Autónoma de Chiapas (FCFM-UNACH), Ciudad Universitaria UNACH, Carretera Emiliano Zapata Km. 8 Rancho San Francisco, Ciudad Universitaria Terán, Tuxtla Gutiérrez, Chiapas C.P. 29050 Chiapas, Mexico
- ¹⁶ Mesoamerican Centre for Theoretical Physics, Universidad Autónoma de Chiapas (MCTP-UNACH), Ciudad Universitaria UNACH, Carretera Emiliano Zapata Km. 4, Real del Bosque (Terán), Tuxtla Gutiérrez, Chiapas C.P. 29050 Chiapas, Mexico
- ¹⁷ Korea Institute for Advanced Study, Seoul 02455, Korea
- ¹⁸ Leinweber Center for Theoretical Physics, University of Michigan, Ann Arbor, MI 48109, USA
- ¹⁹ Institute of Theoretical Science, University of Oregon, Eugene, OR 97403, USA
- ²⁰ Kavli IPMU (WPI), UTIAS, University of Tokyo, Kashiwa, Chiba 277-8583, Japan
- ²¹ Department of Physics, University of California, Santa Barbara, CA 93106, USA
- ²² Department of Physics, LEPP, Cornell University, Ithaca, NY 14853, USA
- ²³ Department of Physics and Astronomy, University of California, Riverside, California 92521, USA
- ²⁴ Dipartimento di Fisica ed Astronomia, Università di Padova, Via Marzolo 8, 35131 Padova, Italy
- ²⁵ INFN, Sezione di Firenze and Department of Physics and Astronomy, University of Florence, Via G. Sansone 1, 50019 Sesto Fiorentino, Italy
- ²⁶ Department of Physics and McDonnell Center for the Space Sciences, Washington University, St. Louis, MO 63130, USA
- ²⁷ Department of Physics, University of Arizona, Tucson, AZ 85721 USA
- ²⁸ Department of Physics, University of California, Berkeley, CA 94720, USA
- ²⁹ Theoretical Physics Group, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA
- ³⁰ Department of Physics, University of Cincinnati, Cincinnati, Ohio 45221, USA
- ³¹ Facultad de Ciencias Físico Matemáticas, Benemérita Universidad Autónoma de Puebla, Av. San Claudio y 18 Sur, Edif. EMA3-231, Ciudad Universitaria 72570, Puebla, México
- ³² Institute for Nuclear Physics (IKP), Karlsruhe Institute of Technology, Hermann-von-Helmholtz-Platz 1, D-76344 Eggenstein-Leopoldshafen, Germany
- ³³ Center for Theoretical Physics of the Universe, Institute for Basic Science (IBS), Daejeon, 34126, Korea
- ³⁴ Department of Physics and Astronomy, University of Victoria, Victoria, BC, V8P 5C2, Canada
- ³⁵ Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot 7610001, Israel
- ³⁶ School of Physics and Astronomy, University of Minnesota, Minneapolis, MN, 55455, USA
- ³⁷ Institute for Nuclear Research of the Russian Academy of Sciences, 60th October Anniversary prospect 7a, Moscow 117312, Russia
- ³⁸ Moscow Institute of Physics and Technology, Institutsky per. 9, Dolgoprudny 141700, Russia
- ³⁹ Institute of Particle Physics Phenomenology, Durham University, Durham, UK DH1 3LE
- ⁴⁰ CP³-Origins, University of Southern Denmark, Campusvej 55, DK-5230 Odense M, Denmark
- ⁴¹ MIT, 77 Massachusetts Ave, Cambridge, Ma 02139, USA
- ⁴² Departamento de Física, Facultad de Ciencias, Universidad de La Serena, Avenida Cisternas 1200, La

Serena, Chile

⁴³ Centro-Científico-Tecnológico de Valparaíso, Casilla 110-V, Valparaíso, Chile

⁴⁴ AHEP Group, Instituto de Física Corpuscular, C.S.I.C./Universitat de València, Edificio de Institutos de Paterna, Apartado 22085, E-46071 València, Spain

⁴⁵ Racah Institute of Physics, Hebrew University of Jerusalem, Jerusalem 91904, Israel

⁴⁶ Physik-Department T30d, Technische Universität München, James-Franck-Straße, 85748 Garching, Germany

⁴⁷ Department of Physics and Astronomy, University of California Irvine, 4129 Frederick Reines Hall, Irvine, CA 92617, U.S.A.

⁴⁸ Department of Physics and Astronomy, Center for Theoretical Physics, Seoul National University, Seoul 08826, South Korea

⁴⁹ Theoretical Physics Department, Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA

⁵⁰ TRIUMF, 4004 Wesbrook Mall, Vancouver, BC V6T 2A3, Canada

⁵¹ Department of Physics and Astronomy, University of Aarhus, Ny Munkegade 120, DK-8000 Aarhus C, Denmark

⁵² Particle Physics Department, STFC Rutherford Appleton Laboratory, Chilton, Didcot, Oxon OX11 0QX, UK

⁵³ Department of Physics, University of Tokyo, Tokyo 113-0033, Japan

⁵⁴ PRISMA Cluster of Excellence & Mainz Institute for Theoretical Physics, Johannes Gutenberg University, 55099 Mainz, Germany

⁵⁵ Departamento de Física Teórica and Instituto de Física Teórica, IFT-UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco, 28049, Madrid, Spain

⁵⁶ Department of Physics, King's College London, Strand, WC2R 2LS London, UK

⁵⁷ Department of Physics, Harvard University, Cambridge, MA 02138, USA

⁵⁸ Department of Physics, Brown University, Providence, RI 02912, USA

⁵⁹ Perimeter Institute for Theoretical Physics, 31 Caroline St. N., Waterloo, Ontario N2L 2Y5, Canada

⁶⁰ INFN Sezione di Roma Tor Vergata, Dipartimento di Fisica, Università di Roma Tor Vergata, Roma, Italy

⁶¹ Ottawa-Carleton Institute for Physics, Carleton University, 1125 Colonel By Drive, Ottawa, Ontario K1S 5B6, Canada

⁶² Instituto de Investigaciones Físicas (IIF) Laboratorio de Física Cósmica de “Chacaltaya” Universidad Mayor de San Andrés (UMSA), Campus universitario de Cota Cota, C/27 Av. Muñoz Reyes, La Paz, Bolivia

⁶³ Planetario “Max Schreier” Universidad Mayor de San Andrés (UMSA), C. Federico Zuazo N° 1976, La Paz, Bolivia

⁶⁴ Department of Physics, Lafayette College, Easton, PA 18042, USA

⁶⁵ Department of Physics, Royal Holloway University of London, Egham, Surrey, TW20 0EX, UK.

⁶⁶ Univ. Grenoble Alpes, CNRS, USMB, LAPTh, 74000, Annecy, France

⁶⁷ Service de Physique Théorique, Université Libre de Bruxelles, Boulevard du Triomphe, CP225, 1050 Brussels, Belgium

⁶⁸ Institute for Theoretical Particle Physics (TTP), Karlsruhe Institute of Technology, Egesserstraße 7, D-76128 Karlsruhe, Germany

Abstract

We examine the theoretical motivations for long-lived particle (LLP) signals at the LHC in a comprehensive survey of Standard Model (SM) extensions. LLPs are a common prediction of a wide range of theories that address unsolved fundamental mysteries such as naturalness, dark matter, baryogenesis and neutrino masses, and represent a natural and generic possibility for physics beyond the SM (BSM). In most cases the LLP lifetime can be treated as a free parameter from the μm scale up to the Big Bang Nucleosynthesis limit of $\sim 10^7\text{m}$. Neutral LLPs with lifetimes above $\sim 100\text{m}$ are particularly difficult to probe, as the sensitivity of the LHC main detectors is limited by challenging backgrounds, triggers, and small acceptances. MATHUSLA is a proposal for a minimally instrumented, large-volume surface detector near ATLAS or CMS. It would search for neutral LLPs produced in HL-LHC collisions by reconstructing displaced vertices (DVs) in a low-background environment, extending the sensitivity of the main detectors by orders of magnitude in the long-lifetime regime. In this white paper we study the LLP physics opportunities afforded by a MATHUSLA-like detector at the HL-LHC. We develop a model-independent approach to describe the sensitivity of MATHUSLA to BSM LLP signals, and compare it to DV and missing energy searches at ATLAS or CMS. We then explore the BSM motivations for LLPs in considerable detail, presenting a large number of new sensitivity studies. While our discussion is especially oriented towards the long-lifetime regime at MATHUSLA, this survey underlines the importance of a varied LLP search program at the LHC in general. By synthesizing these results into a general discussion of the top-down and bottom-up motivations for LLP searches, it is our aim to demonstrate the exceptional strength and breadth of the physics case for the construction of the MATHUSLA detector.

Contents

1	Introduction	7
2	The MATHUSLA Detector Proposal	11
2.1	Basic Principles and Simplified Detector Design	11
2.2	Discovering LLPs with MATHUSLA	13
2.3	Scalability and Modularity of a Realistic Detector Design	15
2.4	MATHUSLA Test Stand	16
2.5	Cosmic Ray Physics with MATHUSLA	17
3	Model-Independent Considerations	22
3.1	LLPs at MATHUSLA	22
3.2	Comparing LLP reach at MATHUSLA to the HL-LHC Main Detectors	28
3.3	Summary	38
4	Theory Motivation for LLPs: Naturalness	39
4.1	Supersymmetry	39
4.2	Neutral Naturalness	57
4.3	Composite Higgs	63
4.4	Relaxion Models	67
5	Theory Motivation for LLPs: Dark Matter	74
5.1	Coannihilation	75
5.2	Asymmetric Dark Matter	76
5.3	Freeze-In Scenarios	78
5.4	SIMPs, ELDERs and Co-Decay	92
5.5	Dynamical Dark Matter	96
6	Theory Motivation for LLPs: Baryogenesis	104
6.1	WIMP Baryogenesis	104
6.2	Baryogenesis via Exotic Baryon Oscillations	108
6.3	Leptogenesis	109
7	Theory Motivation for LLPs: Neutrinos	111
7.1	Minimal “Sterile” Right-Handed Neutrinos	111
7.2	The $B - L$ Gauge Portal	119
7.3	Left-Right Symmetric Model	123
7.4	Neutrinos from the Higgs Portal	128
7.5	Discrete Symmetries and FIMP Dark Matter	132
7.6	Enhanced Residual Symmetry (ERS) Scenarios and Freeze-Out Leptogenesis	133
8	Theory Motivation for LLPs: Bottom-Up Considerations	138
8.1	Hidden Valleys and High Multiplicity Scenarios	138
8.2	Exotic Higgs Decays	144
8.3	The BBN Bound	146
8.4	SM + S: Singlet Extensions	152
8.5	SM + V: Dark Photons	155
8.6	Axion-Like Particles	162
9	Executive Summary	169

Foreword

The MATHUSLA detector concept was proposed in [1] to detect neutral long-lived particles (LLPs) produced at the HL-LHC that could be missed by the main detectors. By providing a very low background environment without trigger limitations for the detection of ultra-long-lived particle decays, the MATHUSLA detector could greatly extend the LLP search capabilities of the LHC, at a relatively incremental cost. This idea led to the rapid formation of an experimental collaboration and the recent deployment of the small-scale MATHUSLA Test Stand detector at CERN to study backgrounds and help develop a full-scale detector proposal.

The present document is the result of an extensive study carried out by the high energy theory community to demonstrate the broad nature of the physics motivation for LLP searches at the LHC. We survey motivations for new physics at the LHC, including (for instance) many approaches to the hierarchy problem, dark matter, baryogenesis, and neutrino masses, and demonstrate how LLPs emerge as natural and generic predictions of these theories. While our treatment here focuses on the long-lifetime regime that especially motivates the construction of the MATHUSLA detector, our general discussion applies to LLP searches at the LHC more broadly. We additionally discuss how predictions from this wide range of theories map into the space of LLP signatures at both MATHUSLA and the LHC main detectors, and establish the regions in LLP parameter space where MATHUSLA offers unique sensitivity.

Many of the sections represent a collaboration of many authors and editors. Other sections contain primarily the work of a few authors, who are then indicated at the beginning of such sections.

This document focuses on the theoretical motivation for LLP signals at MATHUSLA without going into great detail regarding the detector design, since the theoretical arguments and signal estimates are relatively independent of the precise instrumentation. The MATHUSLA experimental collaboration aims to present a Letter of Intent in the second half of 2018, and develop detailed proposals for a full-scale MATHUSLA detector. It is the intention of the authors that the arguments presented in this document aid the preparation and theoretical justification of these proposals.

MATHUSLA also has significant capabilities to act as a cosmic ray telescope. Some of the potential physics studies in this field are briefly pointed out here. A companion document exploring this secondary physics case, which would represent a guaranteed return on the investment of building the detector, is currently in preparation.

1 Introduction

The quest for physics beyond the Standard Model (SM) encompasses many frontiers and employs a multitude of methods. While any genuine deviation from SM predictions is a sign of new physics, by far the most illuminating of these methods would be to directly produce new particle that can arise in theories Beyond the SM (BSM), and study their properties. The most basic properties of such new particles are their mass, charge, spin, and lifetime. While most direct searches for new physics often focus on short lifetimes, it is important to note that long-lived particles (LLPs), defined to have macroscopically detectable decay lengths, are ubiquitous in BSM physics. This can be trivially demonstrated by reference to known SM phenomena, which include a myriad of lifetimes ranging from $< 10^{-24}$ seconds for the top quark to at least $> 10^{41}$ seconds for the proton. Given this gargantuan range of lifetimes in the SM and potentially in BSM physics, it is important to understand the simple origin of LLPs.

The proper lifetime τ of any particle is given by the inverse of its decay width Γ , which can be calculated straightforwardly in Quantum Field Theory (QFT) as

$$d\Gamma \sim \frac{1}{M} |\mathcal{M}|^2 d\Pi. \quad (1)$$

M is the mass of the particle, \mathcal{M} is the matrix element governing the decay, and $d\Pi$ is the phase space for the decay. For a particle to be long-lived, the matrix element and/or the available phase space must be small. It is straightforward to characterize scenarios that lead to either case, and to find examples realizing these possibilities in the SM.

The matrix element for decay could be suppressed due to an approximate symmetry which would forbid the decay if it was precise, or simply a small effective coupling constant.¹ A small coupling in the matrix element can be further distinguished by whether it originates from a dimensionless coupling constant, or a dimensionful scale, larger than M , from a higher-dimension operator that mediates the decay. Phase space can also be suppressed due to the small breaking of an approximate symmetry that splits otherwise degenerate states, or it can arise due to accidental degeneracies in the spectrum.

Examples of these suppression mechanisms in the SM are plentiful. The proton is perhaps the most extreme example of an approximate symmetry giving rise to a very long lifetime, since proton decay is forbidden by baryon number, which is an accidental symmetry of the SM. The long lifetime of the μ arises from a small coupling corresponding to a large dimensionful scale, the Fermi constant G_F , arising due to the high mass of the W boson. The Higgs boson, while not macroscopically long-lived, has a lifetime significantly greater than the similarly massive top quarks or W/Z bosons. It propagates for about a proton diameter before decaying, due to the small dimensionless bottom Yukawa coupling $y_b \sim 0.02$ that dominates Higgs decay in the SM. The SM neutron is very long-lived, with a lifetime of about 15 minutes outside the atomic nucleus, in large part due to the phase space suppression of its weak decay to a proton and leptons. The b -quark's relatively long lifetime is due to a combination of effects: phase space suppression, approximate flavor symmetry and large dimensionful scale in the decay. It is not sufficiently long-lived to pass through the detector, but has a macroscopic decay length \sim mm at the LHC. The macroscopic decays of B -hadrons are utilized for particle ID, which is essential for studying properties of the Higgs and searches for BSM physics.

Given the abundance of LLPs in the SM, and that all classes of mechanisms for creating a long-lifetime occur within the SM, it is unsurprising that LLPs are commonplace in BSM theories as well. However, the discovery of LLPs presents experimental challenges stemming from both their production and detection. For instance, the nature of modern general-purpose particle physics detectors at the energy frontier is biased towards “prompt” particle production, assuming relatively short lifetimes $\lesssim \mu\text{m}$ and maintaining good geometrical acceptance for decays $\mathcal{O}(1 \text{ m})$ away from the IP. In part this is due to

¹While small effective couplings and approximate symmetries can coincide, this is not always the case and we distinguish them as logically distinct possibilities.

practical considerations, since requiring approximate 4π coverage makes detector sizes above $\mathcal{O}(10\text{ m})$ logistically and financially unfeasible.

If the LLP is charged or colored, it can be detected as it passes through the detector. However, the only way of directly detecting a *neutral LLP* is to observe its decay. An LLP with decay length $\lambda \gg 10\text{ m}$ only has small probability $\sim L/\lambda$ of decaying inside a detector of size L . *This small probability makes ultra-long-lived neutral particles inherently rare signals.* Background suppression and good trigger efficiency are therefore vital to their discovery.

Intensity frontier experiments are natural settings for neutral LLP searches, since their smaller scale means they can be customized to search for low-mass hidden sectors with very low backgrounds. Such experiments have much lower center-of-mass energies than the LHC, instead aiming for large rates of low-energy processes due to the intensity of their beams. If the LLP is light enough to be produced, one can try and exploit the inverse decay process of the LLP for its production. However, this precludes probing LLPs with masses above a few GeV, not to mention the weak or TeV scale. Furthermore, if the LLP is long-lived due to a higher-dimensional operator suppressed by a high scale, then processes at energies above that scale could lead to much larger LLP production rates than are possible at intensity frontier experiments. Finally, in many theories the LLP couples not just to its decay products but also other heavy SM or BSM particles. Only by accessing higher energies do these new production processes become available.

It is therefore clear that searching for ultra-long-lived neutral particles at energy frontier experiments like the LHC, HL-LHC or HE-LHC has many advantages. One can exploit the irreducible production mechanism which corresponds to the LLP's inverse decay for both dimensionless couplings and dimensionful scale suppressions. For dimensionful couplings it is trivial to understand why the energy frontier is more powerful than the intensity frontier, as the cross-section for production scales with energy. Typically the energy frontier is also more powerful for dimensionless couplings as well. For instance, although the well-motivated Higgs portal coupling to a hidden sector is fundamentally through a dimensionless coupling, at low-energy intensity frontier experiments this is effectively a dimensionful scale suppression set by the Higgs mass. Therefore any energy frontier experiment that can produce the Higgs will be a more powerful search tool for LLPs coupling to the Higgs. Even in the case of a dimensionless coupling and a low mass scale, for example hidden photon models, the energy frontier can in principle be just as powerful as intensity frontier experiments, since energy frontier experiments like the HL-LHC also utilize very high intensity beams. Energy frontier experiments are also well suited for more general models of LLPs, because the most efficient production mechanism for LLPs *is often completely unrelated* to the LLP decay mechanism. For instance, a large swath of BSM physics models described by supersymmetry [2] have LLPs. The R-parity symmetry structure inherently favors energy frontier production. Even if R-parity is slightly broken, allowing for macroscopic decays, the production of the heavy states is better suited for the energy frontier rather than using the inverse decay production mechanism typical of the intensity frontier.

The HL-LHC and HE-LHC would therefore be ideal tools for studying the *Lifetime Frontier*, supplying both the energy reach and the luminosity needed to study possibly rare LLP signals. However, as mentioned, the general purpose LHC detectors have intrinsic limitations that restrict their reach for very long-lived neutral particles. While LLP decays can be spectacular signals, the high-rate HL-LHC environment is unforgiving, and large QCD backgrounds as well as triggering limitations are major bottlenecks for many LLP searches. Furthermore, while missing energy (MET) searches have great utility in probing new physics giving rise to either more than several 100 GeV of MET or QCD-like production rates, sensitivity drops dramatically for rarer or even slightly softer signals. Finally, even if a very long-lived state were discovered via MET searches at ATLAS or CMS, a critical cosmological question remains: is the newly discovered state a dark matter candidate, or a meta-stable state?

To probe all accessible possibilities for physics beyond the SM, it would therefore be very useful to *combine* the reach of an energy frontier experiment with the shielding and lever arm of a long-lived

particle detector. This has been explored before in a number of cases for the LHC, but there is a tradeoff between the volume of the new detector and the distance from an LHC IP [3] that precludes the use of existing detectors and calls for a new design. The ideal LLP detector would be shielded from QCD backgrounds produced in LHC collisions, while being large and close enough to have sufficient acceptance for LLP decays.

Remarkably, a connection between collider physics and cosmology reveals a lifetime of interest that corresponds to a dedicated LLP detector of achievable size. If the LHC is capable of producing a certain LLP, the same LLP is likely to have been thermally produced in the early universe. As a result, its lifetime is bounded from above by the strict constraints on Big Bang Nucleosynthesis from primordial elemental abundances [4]. In most cases, this upper bound is about $\tau \lesssim 0.1\text{s}$. Fortunately, a detector at the surface above an LHC experiment, with a realistically sized, instrumentable volume, is precisely in the range to probe particles with such lifetimes if they are produced with plausible LHC cross-sections, such as in sizable exotic Higgs decays [1]. This connection led to the genesis of the proposed MATHUSLA experiment: a general-purpose LLP detector that exploits energy frontier particle production, coupled with shielding and size that is powerful enough to probe an enormous array of LLP scenarios.

The detector design for MATHUSLA is discussed in Section 2. In the simplest terms, MATHUSLA is a large tracker that can reconstruct displaced vertices on the surface near ATLAS or CMS. Its surface location provides shielding from QCD backgrounds at the interaction point. The remaining backgrounds can be rejected, allowing MATHUSLA to operate in the near-background-free regime without trigger limitations. The design is scalable, making it highly flexible from the budgetary standpoint as well as allowing for upgradability depending on the physics benchmarks of interest. The MATHUSLA program also provides a number of exciting possibilities and benefits beyond the HL-LHC LLP search program. As discussed in Section 2.3, the detector will also be able to act as a powerful and unique cosmic ray telescope, independent of the LHC and without interfering with the primary LLP search objective. Additionally, since MATHUSLA is proposed to be above an LHC IP, it will be useful for the entire lifetime of the HL-LHC program *and* a possible HE-LHC successor. MATHUSLA not only can extend the reach of an LHC experiment, it can also complement other discovery channels. If for instance a MET signature was discovered by the LHC, MATHUSLA could provide valuable additional information on the spectrum and properties of the hidden sector.

After the MATHUSLA detector was proposed in [1], its reach has been the subject of several studies [5–15], but this was usually done in the context of specific models. Recent complementary proposals for external LLP detectors at the LHC [16–19], as well as recent communal efforts to help guide the LLP search program at the LHC [20], underscore the need for a comprehensive, general examination of the physics motivation for LLP searches. The purpose of this document is therefore to explore, in detail, the physics case for neutral LLP searches in general and for construction of MATHUSLA in particular.

To this end, we develop in Section 3 a model-independent understanding of the MATHUSLA signal yield, the resulting mass and lifetime reach, and how to compare that reach to the ATLAS or CMS main detectors. Many of our methods can be applied, with minor modifications, to help understand other LLP detector proposals as well.

In Sections 4, 5, 6 and 7 we examine the top-down motivations for LLPs with long lifetimes in theories that address the fundamental mysteries of Naturalness, Dark Matter, Baryogenesis and Neutrino Masses respectively. Section 8 examines generic bottom-up scenarios, including hidden valleys [21, 22] and minimal extensions of the SM with additional scalars or vectors. Hidden valleys in particular deserve mention as one of the most generic bottom-up BSM possibilities that gives rise to LLPs. The possibility of a separate sector with its own particles and forces, only connected to the SM by a small portal coupling or a heavy mediator, is a straightforward consequence of the structure of gauge theories. Any massive states in the hidden sector that are not absolutely stable are natural LLP candidates. These and other

bottom-up possibilities are not only plausible on general grounds, they also arise as components of more complete theories, including those discussed in earlier sections.

Our investigation demonstrates the extremely broad motivation of LLPs and the foundational importance of searches for their signatures. LLPs not only arise ubiquitously in BSM theories; in many cases, they are intrinsic to the underlying theory mechanism as well. The lesson for the entire LHC search program is obvious: LLP searches need to be a priority, and they should be explored with the main detectors as well as dedicated experiments like MATHUSLA that take advantage of existing LHC facilities. For many broad classes of BSM scenarios with LLPs, MATHUSLA is the first or only discovery opportunity, being able to detect new physics with TeV scale masses and very long lifetimes with sensitivities that can exceed the cross-section reach of main detector searches by orders of magnitude. Clearly, the discovery potential of such a general-purpose LLP detector is enormous.

We prepare an Executive Summary of our findings in Section 9, which can be read as a stand-alone document and serves as a big-picture guide to the studies and important lessons of this white paper. MATHUSLA represents a unique opportunity for CERN. The collider to produce LLPs is already in place. A relatively incremental upgrade to maximize our chances of actually detecting these possible harbingers of new physics is not only feasible, but highly motivated from a vast and comprehensive range of bottom-up and top-down theoretical perspectives.

2 The MATHUSLA Detector Proposal²

Here we briefly summarize the MATHUSLA (MASSive Timing Hodoscope for Ultra-Stable neutral pArticles) LLP detector for the HL-LHC as first proposed in [1], and review why such a design would be able to search for LLPs with zero or very low backgrounds (Sections 2.1 and 2.2). While the original proposal (and this work) analyzes signal sensitivities in the context of the HL-LHC, it is important to note that MATHUSLA would perform its function as a dedicated LLP detector equally well for any future collider built in the same tunnel as the LHC, i.e. the HE-LHC, where higher LLP production rates lead to a correspondingly improved reach. We emphasize the modularity and scalability of the idea and discuss developments towards a more realistic detector design in Section 2.3. The recent deployment of a 5-meter scale MATHUSLA “test stand” detector at CERN is reviewed in Section 2.4.

It was also realized that MATHUSLA has impressive capabilities as a cosmic ray (CR) telescope. This secondary physics mission represents a guaranteed return on the investment of building the detector and will be the subject of its own dedicated studies [23]. For completeness, we qualitatively discuss in Section 2.5 why MATHUSLA can make unique CR measurements.

2.1 Basic Principles and Simplified Detector Design

The basic motivation for the MATHUSLA detector is the search for LLPs with lifetimes much greater than the size of the LHC main detectors, $c\tau \gg 100$ m. Any detector that can be reasonably constructed could only catch a small fraction of such LLPs decaying inside of its volume. Even with potentially large LLP production rates in LHC collisions, suppression of backgrounds is then crucial for discovery.

The primary signals of neutral LLPs in ATLAS or CMS are displaced vertices (DVs). A DV corresponds to two or more charged tracks that are reconstructed to originate from the same point in space (and in principle time as well, but this is highly detector dependent), a macroscopic distance displaced from the beam collision point where the LLP originated. Especially for LLP searches with high energy or leptonic final states, the spectacular geometrical nature of DVs generally leads to very low backgrounds. Any other class of LLP signature, including e.g. the anomalous energy deposits produced when a LLP decays within the calorimeters, suffers from backgrounds and triggering limitations that can be very significant. As we discuss further in Section 3, this greatly curtails the main detectors’ ability to discover LLPs with very long lifetimes.

To address this broad blind spot of the LHC, MATHUSLA is envisioned to be a (1) large, (2) relatively simple (3) surface detector that (4) can robustly reconstruct DVs with good timing resolution. This is to ensure that: (1) the detector has a similar geometric acceptance for LLP decays as the ATLAS or CMS main detectors, which makes it possible to detect LLPs with lifetimes near the generic BBN upper bound of $\sim 10^7$ m if there are no backgrounds; (2) it can be constructed in time for the HL-LHC upgrade with a realistic budget; (3) it is shielded from QCD backgrounds of the main collision by ~ 100 m of rock; (4) CR backgrounds to DV searches can be rejected with near-perfect efficiency.

A simplified detector design for MATHUSLA, showing its position on the surface near ATLAS or CMS, is shown in Fig. 1. (This is the geometry assumed for physics studies in subsequent Sections.) The main component of the detector is an approximately 5-m thick tracker array situated above an air-filled decay volume that is 20 m tall and $200 \text{ m} \times 200 \text{ m}$ in area. The tracker is envisioned to be composed of five planes to provide highly robust tracking with a timing resolution of ~ 1 ns. This allows the dominant background of downward going cosmic ray particles to be reliably separated from upward going LLP decay tracks. Each plane has a spatial resolution of ~ 1 cm in each transverse direction, providing the vertexing capability necessary to confirm the DV signal topology. The entire bottom and sides of the decay volume³ are covered with scintillator to veto incoming charged particles such as high-energy

²Martin Alfonso, Cristiano Alpighiani, Juan Carlos Arteaga-Velazquez, Mario Rodriguez Cahuantzi, David Curtin, Henry Lubatti, Caballero Mora, Karen Salome, Rinaldo Santonico, Arturo Fernandez Tellez, Subieta Vasquez, Charlie Young

³The diagram shows the top being covered in scintillator as well, but this might not be required depending on the triggering




Fig. 1: Simplified detector layout showing the position of the $200\text{ m} \times 200\text{ m} \times 20\text{ m}$ LLP decay volume used for physics studies. The tracking planes in the roof detect charged particles, allowing for the reconstruction of displaced vertices in the air-filled decay volume. The scintillator surrounding the volume provides vetoing capability against charged particles entering the detector.

muons coming from the primary pp interaction.

The sensor technology should be proven and cheap in order to achieve the requisite fiducial volume at a reasonable cost. Resistive Plate Chambers (RPC) is the current default detector technology, though other options are not excluded at this early stage of the design process. Its tracking performance has been proven in many earlier experiments. Indeed, the performance requirements for MATHUSLA are less stringent than what has already been achieved in large-scale deployments.

For example, ATLAS has achieved a timing resolution of 1 ns and a spatial resolution of 1 cm, while CMS has achieved a timing resolution of 1 ns [24] and a spatial resolution of 0.81 cm [25]. RPCs operating in streamer mode at the YangBaJing laboratory for cosmic ray studies have demonstrated the required rate capability [26]. Higher rates can be achieved by operating in avalanche mode. RPCs have also been deployed in detectors with similar geometry and areas greater than $\sim 7000\text{ m}^2$ [27, 28]. It is also worth noting that ARGO YBJ operated for 5 years almost unattended, testifying to the reliability of the technology. The construction procedure is straight-forward and has been industrialized, making its unit cost superior to the most obvious alternatives. There are no fundamental obstacles to achieve the production rate needed to match the HL-LHC time scale. Nevertheless, MATHUSLA will require a larger area of RPC than has been used in any single experiment before. Since the basic technology of RPC is well understood, the ongoing effort in exploring this detector option is focused on cost performance optimization.

This minimal detector design is sufficient for LLP discovery via geometrical DV reconstruction. As discussed in [9], it is even sufficient to measure the LLP boost event-by-event, which can reveal important information about the LLP mass and production mode, especially when correlated with main detector information. Distinguishing between leptonic and hadronic LLP decays is also straightforward (for LLP masses significantly above a GeV) simply by final state multiplicity.

It was also noted in [9] that MATHUSLA's capabilities could be extended by placing several cm of converter material like Iron between two of the RPC planes. This would allow photons to be detected by conversion and electrons to be distinguished from muons by the induced electromagnetic shower. As shown schematically in Fig. 2, this permits event-by-event particle identification of the LLP final states, supplying important information on the nature of a discovered LLP. Studies are underway on the amount and location of converter material to optimize particle identification performance while taking into account practical issues such as weight and cost. It should be noted that the outcome of this study

strategy.




Fig. 2: Possible particle ID in MATHUSLA with an extra layer of material between tracking layers [9].

has no impact on the design of tracker planes, and is unlikely to affect discovery prospects for LLPs decaying to charged particles.

MATHUSLA is a unique detector with unusual requirements, and its detailed design will require further study. However, its reliance on proven and cost-effective technology means there is no fundamental obstacle for its deployment in time for the HL-LHC upgrade.

2.2 Discovering LLPs with MATHUSLA

We now briefly summarize the arguments put forth in [1] that MATHUSLA could search for LLPs decaying into charged particles with little or no backgrounds. Fig. 3 (top) schematically shows the two main signals for MATHUSLA, LLPs decaying into at least two charged leptons (c), or into jets (d) that contain $\mathcal{O}(10)$ charged hadrons [9]. For leptonic decay, both charged particles hit the tracker in the ceiling in 50 – 90% of cases depending on the LLP boost, while for LLPs decaying to hadrons, almost all decays have 5 or more charged particles hitting the ceiling [9]. Since the tracker planes have \sim cm spatial and \sim ns timing resolution, the charged particle trajectories can be fitted to reconstruct a DV. Unlike traditional DV analyses in the main detectors, these DVs must satisfy the additional stringent requirement that all trajectories coincide *in time* at the DV. The scintillator is used as a veto to ensure that the charged particles *originated* at the DV. There should be no hits along the line between the vertex and the LHC IP, nor along the lines obtained by extrapolating the individual charged particle trajectories backwards. Taken together, these exhaustive geometric and timing requirements make it very difficult for backgrounds to fake the LLP signal.

The main backgrounds to LLP searches in MATHUSLA are represented in Fig. 3 (bottom). Cosmic rays (e) are by far the most dominant background and have a rate of $\sim \mathcal{O}(10\text{MHz})$ on the whole detector, resulting in $\sim 10^{15}$ charged particle trajectories over the whole HL-LHC run. They can be rejected because their downward direction of travel is extremely efficiently distinguished from the upwards-traveling LLP decay products. Furthermore, they rarely if ever reconstruct a displaced vertex inside the decay volume, not to mention one that satisfies the stringent signal requirements above. Finally, the large majority of charged particles from CRs arrive in a highly correlated manner, as part of an Extended Air Shower from a high-energy CR primary interacting in the atmosphere, see Section 2.5. As illustrated in Fig. 4, this leads to a large number of charged particles occupying MATHUSLA near-simultaneously and is easily rejected with very little “blind time” for the LLP search.

Muons from the IP traveling through MATHUSLA (f) or scattering inelastically in the decay volume (g) either do not satisfy the DV signal requirement or occur so rarely that they are easily vetoed by the scintillator. Neutrinos from atmospheric cosmic ray interactions and LHC collisions could be traveling upwards and scatter with a nucleus in the decay volume, producing $\lesssim 100$ genuine DVs. This seems like an irreducible background, but can be rejected because the large majority of final states that




Fig. 3: Schematic comparison of LLP Signal (top) and backgrounds (bottom) in MATHUSLA. Figure from [1]. The most stringent signal requirements are full 4-dimensional reconstruction of the DV from upwards traveling charged particle tracks measured with full spatial and timing information, as well as a veto on DVs involving charged tracks that originate outside the detector (provided by the scintillator). Note that the multiplicity of the LLP final states alone provides important information on the decay mode. The signal requirements are very difficult to fake by the dominant cosmic ray background (d), especially (but not exclusively) for hadronic decays. Muons (f) either do not satisfy the signal requirement or give rise to displaced vertices that are easily vetoed (g). Neutrinos from atmospheric cosmic rays (h) and the LHC (i) can be vetoed due to the presence of non-relativistic protons in the final state, as well as geometrical cuts on the final state cone.

can fake an LLP contain a non-relativistic proton that can be vetoed with timing cuts at minimal cost to signal efficiency. The fact that charged particles from cosmic ray neutrino scatterings form a narrow cone that does not point back to the IP allows for additional rejection, bringing the level of neutrino background events to $\lesssim \mathcal{O}(1)$.

These arguments and preliminary calculations suggest that near-perfect background rejection for LLP searches at MATHUSLA is highly plausible. Further work is needed to verify these estimates. In preparation for a Letter of Intent and the full detector proposal, the study of backgrounds and rejection strategies is currently being undertaken by the MATHUSLA experimental collaboration. If the zero-background regime can be reached, MATHUSLA can have up to $\sim 10^3 \times$ better sensitivity to LLP production cross-sections than ATLAS or CMS in the long-lifetime regime. This will be discussed in more detail in detail in Section 3.

Crucial to the background rejection strategies discussed above is the assumption that the LLP decays into at least two charged particles that can be well-separated by the MATHUSLA tracking system to form a DV. Other LLP signatures are possible. For example, if the LLP boost is higher than $\sim \mathcal{O}(1000)$, the charged particle tracks may not be well separated and the LLP would show up as a “one-pronged DV”, see Section 3.1.4. A search for this signal is still possible, but may suffer from higher backgrounds. If MATHUSLA can detect photons, LLP decays to two or one photon could be reconstructed, but the indirect nature of photon detection via conversion in material may lower the spatial and temporal resolution




Fig. 4: *Top:* GEANT-VMC simulation of a 30 GeV LLP (green dashed line) decaying hadronically inside of the simplified MATHUSLA detector layout from Fig. 1. Only charged hadrons (red lines) are shown. *Bottom:* CORSIKA + GEANT-VMC simulation of atmospheric muon bundle event from Air Shower due to a iron cosmic ray primary with energy about 3×10^{16} eV incident on the MATHUSLA decay volume (see sec. 2.5). White lines are atmospheric muons with energy threshold above 1 GeV. The total number of charged particles in the CR event is approximately a order of magnitude above the number shown in the image, illustrating the obvious differences between CR and LLP events in MATHUSLA.

of the DV reconstruction, with resulting higher levels of background. These alternative LLP signatures will be the subject of future study.

2.3 Scalability and Modularity of a Realistic Detector Design

The MATHUSLA detector idea is highly flexible. This allows for a large variety of possible implementations, depending on detector technology, available space, and budget.

The number of LLPs decaying in MATHUSLA depends only weakly on the precise geometry and location of the surface detector, as long as the decay volume is horizontally displaced from the IP by $\lesssim \mathcal{O}(100\text{m})$. As such, the sensitivity projections we present in this work will be valid (up to some $\mathcal{O}(1)$ factor which does not affect the conclusions) for any detector geometry which places $\sim 10^6\text{m}^3$ of fiducial volume near ATLAS and/or CMS. This leads us to three important conclusions:

- MATHUSLA lends itself to a *modular* implementation, for example by arranging many smaller detector modules at ATLAS, CMS, or both to reach the required decay volume. This greatly simplifies construction of the full detector and allows for iterative deployment. The layout of the full detector complex can then be adapted to the chosen experimental site.
- MATHUSLA is not an experiment with a fixed price tag, but rather a general detector concept which can be *rescaled* to whatever funding level is available. For example, one could imagine as a first stage of deployment a “mini-MATHUSLA” of 1/10 the full volume (assembled of one or several sub-modules) which would have $\sim 1/10$ the sensitivity of the full detector, at approximately 1/10 the cost, see Fig. 7. This would still improve LHC sensitivity to weak-scale, hadronically-decaying LLPs by orders of magnitude (even if the BBN lifetime limit cannot be reached with a smaller detector).
- A modular construction also makes it natural to equip certain modules with special capabilities




Fig. 5: (a): schematic view of the MATHUSLA test stand. (b): picture of the final assembled structure in his test area in the ATLAS SX1 building at CERN. The green dots identify the two scintillator layers used for triggering, while the red dots the three RPC layers used for tracking.

at a much lower cost than upgrading the whole detector. For example, some modules could be equipped with additional material between the tracking layers to add particle ID for a subset of observed LLP decays (or CRs, see Section 2.5). One could also equip one or more of the modules with much higher resolution trackers than the rest, to allow very low-mass LLPs to be searched for without background (see Section 3.1.4.)

While alternative technologies for MATHUSLA are not excluded, work is underway to finalize a realistic detector design using RPCs and possibly plastic scintillators. Taking advantage of the possible modularity, coverage of the full $200\text{ m} \times 200\text{ m}$ footprint would be achieved with a number of smaller identical modules, which will be entirely self-contained except for service connections. This facilitates construction, and the adoption of industrial practices for mass production is expected to reduce costs. The modularity allows easy adaption to a different-sized or different-shaped footprint so detector design and construction can proceed before a final decision on the experimental site. Furthermore, each module will be made weather-tight so there is no need for an experimental hall to house the MATHUSLA detector. Trigger information will be provided by the RPC tracking chambers, similar to what has been done in experiments such as ATLAS and CMS. It is anticipated that each module will contribute a local trigger signal to the overall event trigger. Timing stability over the large area of MATHUSLA requires care; however, it should be noted that tighter timing requirements and greater distances than MATHUSLA have been dealt with in accelerator facilities.

2.4 MATHUSLA Test Stand

A test stand has been assembled at CERN and it was installed in the surface area above the ATLAS detector in November 2017. Figure 5 shows the basic design of the test module and a picture of the final assembled structure in ATLAS SX1 building at CERN. The overall structure is $\sim 6.5\text{ m}$ tall, with an active area of $\sim 2.5 \times 2.5\text{ m}^2$.

Following the concept of the main detector, the test stand is made of three layers of RPCs between two layers of scintillator. Scintillator detectors are used to trigger upward and downward charged tracks. The top and bottom scintillator layers are comprised of 28 and 31 scintillators, respectively, recycled from the Tevatron D0 experiment. The RPCs are used for tracking and they were provided by Università

di Tor Vergata, Rome. They are the same type of chambers used in the Argo-YBJ experiment at the YangBaJing Laboratory in Tibet (4300 m a.s.l.).

Several efforts are underway to develop simulations of the backgrounds expected in MATHUSLA. For muons and neutrinos traveling upwards, the idea is to create a “MC particle gun” that shoots particles into MATHUSLA, while for cosmics the plan is to use the standard cosmic ray simulations (e.g. CORSIKA). Nevertheless, the simulations need to be validated and tuned using real data, and for this reason the test stand is crucial.

Since the main goal is to have a background-free MATHUSLA detector, the central purpose of the test stand is to measure the background from CRs and the LHC collisions in order to test the hypothesis that MATHUSLA could reject these most numerous of expected backgrounds. Nevertheless, the test stand should not be considered a prototype of the main detector: the layout could be further optimized, especially with custom-built rather than repurposed components, and detector technologies could be considered. Even so, it will provide very useful information for the design of the future MATHUSLA detector. All the tests that will be performed until the end of LHC Run 2 will be fundamental to understanding the cosmic ray rate in the test stand and to extrapolate the LHC-correlated background rate from the test stand to MATHUSLA. A precise measure of the charged particle flux in the test stand will provide the veto efficiency requirement for the main detector. The goal is to achieve a sufficient timing resolution to guarantee that no cosmic particles can fake a charged particle coming from LHC.

The on-going analysis of the data collected during 2017 and the beginning of 2018, along with all the experience gained from the construction, assembling and commissioning of the test stand, will be crucial for the preparation of the Letter of Intent that the MATHUSLA Collaboration plans to submit to the CERN Committee in late 2018.

2.5 Cosmic Ray Physics with MATHUSLA

The design of MATHUSLA is driven by the requirements of reconstructing upward-traveling displaced vertices and distinguishing them from downward-traveling cosmic rays. It therefore comes as no surprise that MATHUSLA has all the qualities needed to act as an excellent cosmic ray telescope. In fact, MATHUSLA’s particular combination of robust tracking and large area allow it to make many unique measurements that could address important and long-standing questions in astroparticle physics. The study of cosmic rays is therefore an important secondary physics goal of MATHUSLA. These measurements, which in no way interfere with the primary goal of LLP discovery, represent a “guaranteed physics return” on the investment of the detector, as well as an opportunity for CERN to establish a world-leading cosmic ray physics program.

The cosmic ray physics program at MATHUSLA warrants in-depth examination beyond the scope of this work. Some initial studies will be presented elsewhere [23]. Here we only briefly comment on the qualities that make MATHUSLA a uniquely interesting cosmic ray experiment, and outline some possible measurements that are of particular interest to the astroparticle physics community. To this end, we first review some basic facts about cosmic rays and how they are detected.

Cosmic rays, dominantly protons and heavier atomic nuclei, arrive at earth with an energy that spans some 12 orders of magnitude, from a few hundred MeV (10^8 eV) to 100 EeV (10^{20} eV) [29–31], see for example, Fig. 6. They are produced in violent astrophysical scenarios within our own galaxy (for energies $E \lesssim 10^{18}$ eV) and beyond (for $E > 10^{18}$ eV). At the highest energies, however, the origin of CRs is still mysterious, since propagation in galactic magnetic fields means CRs do not point back to far away sources [31, 32]. In general, details of CR acceleration mechanisms, composition, propagation through space, and features in their spectrum are not completely understood [31, 33–35]. The study of CRs offers a unique window on the most energetic natural phenomena of the cosmos [36–38], and their collisions with the atmosphere probe energies far in excess of the TeV scale [39–41].

Primary CRs with relatively low energy can be directly characterized by balloon- or space-born




Fig. 6: Global view of the all-particle cosmic ray energy spectrum (figure taken from [42]). The total spectrum decreases quickly according to a power-law formula $E^{-\gamma}$, where γ (the spectral index) varies from roughly 2.6 to 3.3 [31]. Between 10^{15} eV and 10^{19} eV, the spectrum exhibit three distinctive features created by a change in the value of γ : the knee, which is located at ~ 4 PeV [42, 43], the second knee, close to 100 PeV [44–46] and the ankle, around 4 EeV [47, 48]. There exists also a weaker structure called the low energy ankle [31] at ~ 10 PeV [44, 45]. MATHUSLA is expected to be sensitive to hadronic EAS with primary energies around the knee, i.e. in the interval $E \sim 10^{14} - 10^{17}$ eV according to its size 200×200 m² and the atmospheric depth at which it will be located (~ 1000 g/cm²).

particle physics experiments equipped with trackers and calorimeters, like AMS-02 [49, 50] and CREAM [51, 52] among others (see, for example, [53–55]). The small size of these detectors restricts this approach to energies below ~ 100 TeV – 1 PeV, both because the CR flux drops dramatically above this threshold and because the detector’s magnetic field and radiation depth limits the maximum energy which can be reconstructed.

To study higher energy CRs above $\sim 10^{14}$ eV, the atmosphere is used as a calorimeter [29, 31, 56, 57]. The detection technique consists in observing the Extensive Air Showers (EAS) of SM particles that CRs induce in collisions with the atmosphere. The EAS typically originates from 15 km to 35 km above sea level [57, 58] and spreads out as the particle front travels towards the ground. At lower altitudes, the EAS consists primarily of electrons and photons, with $\sim 10\%$ muons and a smaller fraction of hadrons (for vertical incidence) [59]. Depending almost linearly on the energy in logarithmic scale, the total number of particles in the EAS is in the range $\sim 10^4 - 10^{10}$, for primary energies $E = 10^{14}$ eV – 10^{20} eV, and it is mostly contained within a cone of $\mathcal{O}(0.1 - 1)$ km in diameter at ground level [32, 56]. Measurements of the EAS allow the primary CR’s direction, composition and energy to be reconstructed [59]. Earth-bound CR experiments employ two classes of techniques to probe the EAS: (1) particle detectors on the surface, like particle counters, trackers, and calorimeters to sample the air shower front and (2) various telescopes to observe electromagnetic emissions from the shower or from the interaction of the EAS with the atmosphere (Cherenkov radiation, radio, fluorescence light). Many experiments use a combination of both techniques like the Pierre Auger Observatory [60], TUNKA [46] and TALE [61]. EAS observatories monitor in general large areas to compensate the low CR flux at high

energies. For example, at $E \sim 10^{15}$ eV, CRs are received at a rate of ~ 1 particle/m² · year [29] and shower arrays with areas of order $\gtrsim \mathcal{O}(10^4 \text{ m}^2)$ are required, like the KASCADE air shower detector ($200 \times 200 \text{ m}^2$) [62]. On the other hand, at extremely high energies $E \sim 10^{20}$ eV, the CR flux is so low (~ 1 particle/km² · century [29]) that sufficient exposure requires installations with very large areas of the order of $\mathcal{O}(10^3 \text{ km}^2)$, like the 3000 km² Pierre Auger Observatory [60, 63].

EAS telescopes/antennas are able to observe the longitudinal development of the air shower, while surface detectors measure the lateral structure of the EAS at the atmospheric depth of the site [59]. In the latter, particle detectors are arranged in arrays and are spaced at regular intervals to optimize the measurements for the energy range of interest. Hence, in most cases, they sample only a small fraction of the shower at the observation level. For instance, in case of KASCADE, the main array of 252 e/γ detectors covered only 1.22% of the total surface, the muon array of 192 detectors, only 1.55%, while the muon tracking detector and the hadron calorimeter covered just 0.64% and 0.76%, respectively [62]. Just in a few cases, full coverage was achieved as in the case of the ARGO-YBJ detector, which consisted of a $74 \times 78 \text{ m}^2$ carpet of RPC's with an active area of almost 93% [64]. In general, most surface detectors are insensitive to the energy of a single charged particle (although a counter can be equipped with shielding or buried underground to implement a desired minimum energy threshold). Rather, the focus is on collective shower properties, like the spatial and temporal distribution of particles, as well as some basic particle ID to separate the e/γ , muon and/or hadron components in the EAS. This data is used to characterize the primary CR, assuming a certain hadron interaction model which governs the evolution of the EAS in the atmosphere [57, 65].

Hadron interaction models are a crucial part of air shower Monte Carlo simulations. They are tuned to available high energy physics data but rely on extrapolation in certain regions of phase space, in particular, the forward region at very high energies. This introduces unavoidable uncertainties in the determination of the properties of the primary cosmic rays [66–68]. Verifying and tuning these hadron interaction models is therefore of fundamental importance to CR physics [57]. Tests of hadron interaction models can be performed with the same data from EAS observatories. That requires, however, the simultaneous measurement of different observables of the air showers. KASCADE made important contributions to this topic [69–75], because of both the quality of its measurements and its different detection systems, like the full-coverage central tracker and calorimeter [62]. This central detector only had less than 1% the area of the full experiment but was crucial in allowing for more detailed analysis of the shower [69, 70, 72, 74, 75].

We can now understand why MATHUSLA would make important contributions to cosmic ray physics. MATHUSLA is the size of KASCADE and will be also located at a comparable atmospheric depth ($\sim 1000 \text{ g/cm}^2$). Therefore, it is expected to be sensitive to a similar CR energy range of $10^{14} - 10^{17}$ eV. After a few years of exposure ($\sim 3 \text{ yr}$), roughly 10^6 air showers with $E > 10^{15}$ eV would be recorded with their cores traveling through the MATHUSLA detector. Unique for a CR experiment of this size, MATHUSLA has full-coverage robust tracking with excellent position and timing resolution. The scintillator planes which enclose the LLP decay volume would supply additional information, especially for highly inclined showers. Even without track-by-track e/μ discrimination or calorimetry, this would allow for very detailed EAS measurements, including highly granular analysis of the shower's temporal and spatial structure, which has never been undertaken at this size scale at PeV energies. Furthermore, for roughly half of those CR events, the so-called “golden events”, part of the shower's high-energy muon component ($E_\mu^{th} \gtrsim 50 \text{ GeV} - 70 \text{ GeV}$ for vertical incidence) also passes through the underground ATLAS or CMS detector and could be simultaneously registered with special CR triggers during the LHC runtime. Therefore, whether working in standalone mode, or in tandem with the main underground detector, MATHUSLA constitutes a powerful cosmic ray experiment with unique capabilities that will open an era of precision EAS measurements in the PeV energy region.

The particular CR measurements which MATHUSLA can perform will be studied in more detail in a future document [23]. Some of the most compelling targets for CR measurements include:

- *Primary CR spectra and composition:* The spatial and temporal distribution of charged particles in the shower would probe the energy and composition of the primary CR, potentially addressing open issues like the exact position of the “light knee” (a change in spectral index in the spectrum of protons and helium, which ARGO-YBJ found at ~ 700 TeV [76] and KASCADE, around 3 PeV – 4 PeV [68, 77]) and the shape of the spectra of the heavy elemental components of primary CRs at PeV energies [68, 77].
- *Cosmic Ray Anisotropies:* MATHUSLA’s enhanced resolution, compared to previous CR experiments, could allow to improve the measurements on the dipole component of the anisotropy in the diffuse CR flux at PeV energies, which has been poorly investigated (see [31, 35] and references therein). Another important aspect is the search for local point sources in the northern celestial hemisphere, which would provide vital clues about the presence of nearby galactic accelerators of very high energy CRs [78–80].
- *Highly inclined showers:* The vertical scintillator planes enclosing MATHUSLA’s decay volume, together with its precise full-coverage tracking, allow for the study of highly inclined air showers at large zenith angles $\theta > 60^\circ$. These showers are interesting for a variety of reasons. If they originate from charged primary CRs, they are dominated by muons since their electromagnetic component is attenuated after traversing large distances in the atmosphere. Observations of such events could help to study the high-energy muon content of EAS and to test hadron interaction models by looking for anomalies in this sector at PeV energies, such as those which have been observed at higher energies by observatories like the KASCADE-Grande detector [81, 82], the Pierre Auger observatory [39, 83] or the Yakutsk experiment [84]. The former has observed, for example, that the actual attenuation length of shower muons at 10^{17} eV is bigger than the predictions of hadron interaction models [82], while the latter ones have measured an excess of muons in ultra-high-energy EAS in comparison with the models, a problem which is known as the *muon puzzle* [85]. Following [86], atmospheric and/or astrophysical neutrinos with energies above 10^{15} eV could also be detected in this way if they scatter deep in the atmosphere or interact with the rock of the nearby Jura mountains. Such neutrinos could induce very inclined young EAS, which could be distinguished from regular old showers produced by CRs due to MATHUSLA’s superior tracking resolution. Young EAS are characterized by a richer electromagnetic component, a broader time signal and a larger EAS front curvature. Thus measurements of the particle content and the spatial/temporal structure of inclined EAS in MATHUSLA could allow the search for neutrino signals. MATHUSLA might also offer a tool to look for upward-going EAS from Earth skimming ν_τ ’s [86, 87] and upward-going muons from ν_μ ’s interacting with the rock below the detector or inside the MATHUSLA’s instrumented volume [88, 89]. If MATHUSLA is able to detect also upward-going muons resulting from muon-flavored neutrinos interacting in rock, MATHUSLA could also provide complementary measurements for ν –oscillations (see, for example, [90–92]).
- *Study of EAS and tests of hadronic interaction models:* Detailed measurements of the spatial and temporal structure of EAS, as well as data on the charged particle attenuation length and the muon components of highly inclined showers, may provide a number of clues to understand several outstanding ambiguities in hadron interaction models [81–84, 86]. Additional constraints would be supplied by correlating MATHUSLA’s measurements with detection of the high-energy muon component in the underground detector.
- *High-multiplicity Muon Bundles:* Muons with an energy greater than ~ 50 GeV will penetrate down into the rock and reach the main detectors at CERN. ALEPH/DELPHI at LEP [93, 94] and ALICE at the LHC [95] have studied this high-energy muon component of EAS’s, observing events with more than 100 muons in the underground detector. During the LEP era, these high-multiplicity muon bundles could not originally be explained by hadron interaction models, and several BSM explanations were proposed [96, 97]. Today, the data by ALICE points towards iron-rich CR primary composition at energies above 10^{16} eV, but the data has very low statistics and

further measurements are needed to understand the origin of these muon bundles and their impact on primary CR studies.

Muon bundles could be detected by the LHC main detector and correlated with data from MATHUSLA. This would give a much more complete picture of these special CR events and could allow for their origin to be unambiguously determined. These events are also valuable for constraining hadron interaction models.

It is worth noting that these measurements could be significantly improved if e/μ discrimination capability were added to MATHUSLA. Apart from making measurements of CR primary composition and spectra less dependent on hadron interaction models, the separate muon data would allow for many additional detailed probes of hadron interaction models, with great benefit to all future CR measurements at other experiments, e.g. [61, 98–100]. This non-exhaustive lists of physics targets demonstrates that MATHUSLA could supply data which will be of unique value to the astroparticle and CR physics communities.

3 Model-Independent Considerations

In this section, we provide some model-independent information that allows us to understand MATHUSLA's sensitivity to LLPs produced at the HL-LHC main interaction point, what mass scales and lifetimes it can hence probe, and how its resulting capabilities for discovering new physics compare to those offered by both MET and LLP searches at ATLAS and CMS. This will provide important context for the signal estimates in the subsequent sections. We also comment on the importance of energy thresholds in the low-mass regime, which has implications for the final detector design.

3.1 LLPs at MATHUSLA

3.1.1 Signal Estimate

The probability for each LLP in a signal event sample to decay within MATHUSLA's assumed $(200m) \times (200m) \times (20m)$ decay volume of Fig. 1 is easily computed for a given proper lifetime $c\tau$. For an LLP that traverses the detector volume, this probability is given by

$$\begin{aligned} P_{\text{decay}}(bc\tau, L_1, L_2) &= e^{-\frac{L_1}{bc\tau}} - e^{-\frac{L_2}{bc\tau}} \\ &\approx \frac{L_2 - L_1}{bc\tau} \quad \text{for } (L_2 - L_1) \ll bc\tau, \end{aligned} \quad (2)$$

where L_1, L_2 are the distances from the IP where the LLP enters and exits the decay volume, and

$$b = \frac{|\vec{p}|}{m} \quad (3)$$

is the boost of the LLP.

The per-decay-detection-efficiency $\epsilon_{\text{LLP}}^{\text{MATH}}$ of the LLP within the detector volume will depend on the specific decay mode, as well the precise location of the decay within the detector. Since tracking has to be highly redundant to reject cosmic ray backgrounds, the dominant factor in determining signal efficiency for LLPs decaying into charged particles is simply geometric, i.e. whether the LLP decay products hit the tracker panes near the roof of MATHUSLA.

We concentrate on LLPs decaying into at least two charged particles, assuming at least two charged tracks have to be associated with a displaced vertex for signal reconstruction and background rejection. This was studied in [9], for LLPs with mass $\sim 10\text{-}50$ GeV having boosts of order $b \sim 1 - 10$, decaying to quarks, gluons or charged leptons. These hadronic decays produce $\sim 10\text{-}20$ charged hadrons, and the efficiency for more than 5 charged particles to hit the trackers is close to 100% inside the decay volume. For leptonic decays, the efficiency for both leptons to hit the tracker is better than 90% for lighter and more boosted LLPs, and about 50% for relatively heavy and slow-moving LLPs. These numbers help establish the expected range of the reconstruction efficiency $\epsilon_{\text{LLP}}^{\text{MATH}}$: approaching 1 for hadronic decays, and $\sim 0.5\text{-}1$ for leptonic 2-body decays. This efficiency may be somewhat reduced for LLP decays to soft ($|\vec{p}| < \text{GeV}$) or highly collimated ($\Delta\theta < 0.01$) final states (as we discuss in Section 3.1.4), or for decays that have a sizeable invisible component; however, this should not greatly affect the comparisons laid out below. In our signal estimates, we therefore either assume perfect efficiency, or quote all results normalized to an unknown $\epsilon_{\text{LLP}}^{\text{MATH}}$.

It is not yet established whether MATHUSLA will be able to detect photons, see Section 2.1. This would likely rely on a layer of material inserted between tracking panes to allow for conversion and subsequent detection of the electron-positron pair [9], see Fig. 2. While this layer would add cost and complication, it would also allow for some particle identification, which would greatly aid diagnosis of the LLP decay mode and is highly motivated from the cosmic ray physics point of view (in particular the electron vs muon discrimination, see Section 2.5). In the signal estimates for a few theories we therefore also examine LLPs decaying to one or two photons, to more closely examine the motivation for including this capability.

LLP production mode	effective parent mass scale m_{eff}
s-channel scalar $gg \rightarrow \Phi, \Phi \rightarrow XX$	$\sim 1.5m_\phi$
s-channel vector $\bar{q}q \rightarrow Z', Z' \rightarrow XX$	$\sim m_{Z'}$
s-channel vector $\bar{q}q \rightarrow Z', Z' \rightarrow XX$ with $m_{Z'} \gtrsim \text{TeV}$	$\sim 1.5 \text{ TeV}$ for $m_X \ll 700 \text{ GeV}$
$pp \rightarrow XXjj$ production via $W^\pm Y^\mp X$ coupling, $m_Y \lesssim \text{TeV}$	$\sim 2m_Y$
$pp \rightarrow XXjj$ production via $W^\pm W^\mp XX$ effective coupling	$\sim 2.5 \text{ TeV}$ for $m_X \ll 1.2 \text{ TeV}$
Heavy parent pair production $pp \rightarrow YY, Y \rightarrow X + \dots$	$\sim m_Y$
Exotic decays of B -mesons	$\sim 14 \text{ GeV}$

Table 1: The average boost \bar{b} of an LLP X which flies through MATHUSLA, produced at $\sqrt{s} = 14 \text{ TeV}$ in the above production processes, can be estimated using an effective parent mass scale m_{eff} (second column) using $\bar{b} = m_{\text{eff}}/2m_{\text{LLP}}$. The above table was empirically derived from simulation.

Analytical Approximation

It is very helpful to have an analytical approximation of the LLP signal yield at MATHUSLA. This is often sufficient for simple signal estimates, especially in the important limit of long lifetime, and gives a very general understanding of the cross-sections and mass scales MATHUSLA can probe. For a given LLP production process with cross-section $\sigma_{\text{sig}}^{\text{LHC}}$ in the 14 TeV pp collisions of the HL-LHC, the number of observed LLP decays over the HL-LHC run with $\mathcal{L} = 3000 \text{ fb}^{-1}$ can be estimated as

$$N_{\text{obs}}^{\text{MATHUSLA}} \approx (\sigma_{\text{sig}}^{\text{LHC}} \mathcal{L}) \epsilon_{\text{LLP}}^{\text{MATH}} n_{\text{LLP}} P_{\text{decay}}^{\text{MATH}}(c\tau) \quad (4)$$

where n_{LLP} is the number of LLPs produced per event, and $P_{\text{decay}}^{\text{MATH}}$ is the chance that an LLP decays in the MATHUSLA detector volume. It is given approximately by

$$P_{\text{decay}}^{\text{MATH}}(c\tau) \approx \epsilon_{\text{geometric}} P_{\text{decay}}(\bar{b}c\tau, L_1, L_2) \quad (5)$$

where $\epsilon_{\text{geometric}} \approx 0.05$ is the fraction of LLPs that fly through the MATHUSLA detector, and \bar{b} is the average boost of that fraction. The lengths (L_1, L_2) are taken to be (200m, 230m).

We have verified the above approximation for $c\tau \gg 200m$ by explicitly computing the signal acceptance for a variety of LLP production modes and masses in the range of several hundred to $\ll 1 \text{ GeV}$, simulated to lowest order in MadGraph 5 [101] and showered in Pythia 6 [102]. Eqn. 4 is very robust and agrees with the full simulation to within a factor of 2, usually underestimating the real signal yield. The average boost of LLP X can be estimated from an effective parent mass scale m_{eff} ,

$$\bar{b} = \frac{m_{\text{eff}}}{2m_{\text{LLP}}} \quad (6)$$

which depends on the production processes in a physically intuitive way, up to a numerical prefactor which can be determined from simulation.⁴ For the production processes we examined, m_{eff} is given in Table. 1.

For shorter decay lengths $c\tau \ll 200m$ the MATHUSLA signal is dominated by the tails of the LLP boost distribution. Due to the exponential dependence of Eqn. 2 on the LLP boost, the signal yield is not well captured by using an average boost. Therefore, Eqn. 4 will significantly under-estimate the signal in the short lifetime limit.

Note that in the above, we do not differentiate between two "simultaneous" displaced decays in MATHUSLA from the same LHC event, or two decays from different LHC events. For discovery, each DV is conspicuous enough at MATHUSLA that it can be treated as an independent signal.

⁴This is obviously reminiscent of the boost of a particle with mass m that is pair produced in the decay of a stationary parent particle with mass m_{parent} , $b = \frac{m_{\text{parent}}}{2m} \sqrt{1 - 4m^2/m_{\text{parent}}^2}$, when the parent mass is large.

3.1.2 Sensitivity Estimate

As argued in [1] and reviewed in Section 2, MATHUSLA can operate in the background-free regime for LLPs decaying to two or more well-separated charged particles. We therefore obtain projected exclusion limits on the LHC production cross-section of the LLPs by setting $N_{\text{obs}}^{\text{MATHUSLA}} = 4$ in Eqn. 4.

$$(\epsilon_{\text{LLP}}^{\text{MATH}} \cdot \sigma_{\text{sig}}^{\text{MATHUSLA limit}}) \approx \frac{4}{\mathcal{L} n_{\text{LLP}} P_{\text{decay}}^{\text{MATH}}(c\tau)}, \quad (7)$$

and analogously for estimates using full simulation of LLP acceptance. As we show below this gives sensitivity to cross-sections above a fb. For discovery, we assume $N_{\text{obs}}^{\text{MATHUSLA}} = 10$ is required.

Detailed study of backgrounds and signal reconstruction at MATHUSLA may slightly increase the required number of events for exclusion and discovery, but the above criterion is expected to be a good approximation for the majority of LLP decay modes and sufficient for studies to motivate the detector. Detailed studies may also reveal that the zero-background assumption does not hold for some final states, e.g. decays to photons (which may not be detectable, see Section 3.1.1), one-pronged decays of an LLP (e.g. to two collinear jets + an invisible particle, or photon + invisible particle), decays to electrons only (which may shower in the detector material, making exact DV reconstruction more challenging), or decays of LLPs resulting in very tightly collimated final states (which could mimic neutrinos scattering off air, see Section 3.1.4). In that case, sensitivity estimates for those final states may have to be adjusted accordingly.

3.1.3 Benchmark Signal Cross-Sections

Limits on LLP production cross-section σ as a function of lifetime $c\tau$ will have a minimum (best limit) at some $c\tau_{\text{best}}$. For $c\tau \gtrsim c\tau_{\text{best}}$ the limit depends linearly on $c\tau$ as long as the search requires only a single LLP decay (as is the case at MATHUSLA), while the short-distance behavior is slightly more complicated. In Eqn. (4) and in simulations, $P_{\text{decay}}^{\text{MATH}}(c\tau)$ is maximized for

$$[\bar{b}c\tau]_{\text{best}} \approx 200 \text{ m} \quad (8)$$

giving

$$P_{\text{decay}}^{\text{MATH}}([\bar{b}c\tau]_{\text{best}}) \approx 2 \times 10^{-3} \quad (9)$$

Assuming for simplicity that $\epsilon_{\text{LLP}}^{\text{MATH}} = 1$ (which is accurate enough for this estimate), this means that for *some range of lifetimes*, a model with LLPs will produce an observable MATHUSLA signal over the HL-LHC run if the LLP production cross-section is larger than

$$\sigma_{\text{sig}}^{\text{LHC}} \gtrsim \text{fb} \quad (10)$$

(This lower bound will be reduced if the number of produced LLPs per event is very large, as in some dark shower models.) In deriving this lower bound we required $N_{\text{obs}}^{\text{MATHUSLA}} = 4$, but at this level of precision, the distinction between exclusion and discovery is not important. If the LLP production cross-section is larger than $\sim \text{fb}$, then the maximum lifetime that can be probed is roughly

$$\bar{b}c\tau_{\text{max}} \sim (10^3 \text{ m}) \left(\frac{\sigma_{\text{sig}}^{\text{LHC}}}{\text{fb}} \right) \quad (11)$$

(assuming $\mathcal{O}(1)$ LLPs per production event) since the linear long-lifetime regime starts at a lifetime a factor of a few larger than $[bc\tau]_{\text{best}}$. This model-independent schematic sensitivity of MATHUSLA is shown in Fig. 7.

To emphasize the scalability of the MATHUSLA design (see Section 2.3) we also show the sensitivity of a detector with only 1/10 the volume of the $200\text{m} \times 200\text{m} \times 20\text{m}$ benchmark geometry, which




Fig. 7: Schematic order-of-magnitude sensitivity of MATHUSLA, assuming $\mathcal{O}(1)$ produced LLPs per production event at the HL-LHC. \bar{b} is the mean boost of the produced LLPs. The shape of the exclusion/discovery region at short lifetimes depends on the detailed boost distribution, but for long lifetimes $\bar{b}c\tau \gg 200m$ depends only on the mean boost and is very model-independent up to an $\mathcal{O}(1)$ factor. Note that LLPs near the BBN lifetime limit of $c\tau \sim 10^7m$ can be probed if they are produced with cross-sections in the pb range at the HL-LHC. To emphasize the scalability of the MATHUSLA design, we also show the reach achievable with a version of MATHUSLA with only 1/10 the detector volume of the $200m \times 200m \times 20m$ benchmark geometry.

is assumed throughout the rest of this paper. Since the LLP cross-section that can be discovered scales inversely with detector volume, all the expressions in this Sections and indeed the results of this entire whitepaper are easily rescaled for a smaller version of MATHUSLA. While such a mini-MATHUSLA may not probe BBN lifetimes, it would still extend the LLP sensitivity of the LHC main detectors by orders of magnitude.

This understanding of the model-independent LLP reach allows us to understand which BSM mass scales MATHUSLA can probe. In Fig. 8, we show benchmark LLP signal cross-sections at the 14 TeV HL-LHC, either as a function of parent particle mass m_{parent} or as a function of the LLP mass m_{LLP} , for the most important benchmark processes:

- Pair production of color octet fermions (gluinos $\tilde{G}\tilde{G}$), color fundamental scalars (stops $\tilde{t}\tilde{t}$) or fermions ($f_c f_c$) which can either be LLPs or decay to LLPs. See Sections 4.1.1, 4.1.3, 4.1.4, 4.3, 5.2, 5.4, 6.1, 8.1.
- Pair production of EW charged states like SUSY higgsinos ($\tilde{H}\tilde{H}$), and winos ($\tilde{W}_{\pm}\tilde{W}_{\pm}, \tilde{W}_{\pm}\tilde{W}_0$) which can either be LLPs or decay to LLPs. See Sections 4.1.1, 4.1.2, 4.1.5, 5.1, 5.3, 6.1, 7.5.
- S-channel resonance production of an Z' -type vector boson which can decay to LLPs. The plot assumes coupling $g = c_A g_W$, with $c_A = 0.1$ but the cross-section $\sigma \propto c_A^2$ can be easily rescaled for lower couplings. This also stands in for a kinetically mixed dark photon with $c_A \sim \epsilon$, and behaves similarly to the production cross-section for a charged W' . See Sections 5.4, 7.2, 7.2.2, 7.3.1, 7.3.2, 8.5, and also Sections 8.5, 8.6. for exotic decays of the SM Z boson.
- S-channel resonance production of a SM-like Higgs produced 1/10 the SM cross-section, which can decay to LLPs. This can be the case for e.g. singlet scalars which mix with the SM Higgs. Again, the cross-section can be rescaled to account for lower mixings. See Sections 6.1, 8.1, 8.2, as well as Section 8.2 and 4.2, 7.4, 8.4, 8.5 for exotic decays of the 125 GeV Higgs boson.




Fig. 8: Benchmark LLP production cross-sections at the 14 TeV HL-LHC, as a function of either parent particle or LLP mass.

The SUSY cross-sections are taken from [103], with the results available in tabulated form in [104]. The cross-sections for the Z' boson and coloured triplet fermion are calculated using MadGraph [101], with Z' coupling conventions as in [105]. The heavy Higgs cross-section is taken from [106], with the 14 TeV results rescaled by the luminosity ratio.

These processes are important for many theoretically motivated scenarios, as we discuss in more detail in the following sections. Some universally applicable statements can be made based on the mass for which the above cross-sections are \sim fb: for strong production processes, MATHUSLA can probe LLP or parent particle masses in the 1.3 - 2 TeV range. Depending on the coupling, vector mediators can be produced with masses of several TeV. Electroweak LLPs or parents, like higgsinos or winos, can be probed with masses up to \sim TeV.

Of course, exotic decays of the 125 GeV Higgs Boson are some of the most well-motivated and discoverable LLP production modes. At the HL-LHC, the Higgs production cross-section is about 50 pb, meaning branching ratios to LLPs of $\sim 10^{-5}$ can be probed for lifetimes near $c\tau_{\text{best}} \sim 200\text{m}$. Importantly, for branching ratios of $\sim 10\%$, which are not excluded by current measurements, LLP lifetimes near the BBN limit of $c\tau < 0.1$ seconds can be probed, see Section 8.2 and also 4.2, 7.4, 8.4, 8.5.

Finally, an important possible source of LLPs with masses in the GeV range or below are the decay of SM hadrons, especially B -mesons, which are produced at the LHC with ~ 0.6 mb cross-section [107]. The resulting $\sim 10^{15}$ produced B -mesons produced at the HL-LHC can give rise to displaced signals even for extremely tiny exotic branching fractions to LLPs, see Section 4.1.6, 6.2, 7.1, 8.4.

Note that for the case of LLP production in parent particle decays, these cross-sections have to be multiplied by the appropriate parent particle branching fraction. Similarly, in some cases production rates may be suppressed by small couplings or unknown mixing factors. We do, however, assume that the LLP decays to final states involving SM particles 100% of the time (with possibly different rates to different SM final states), since a displaced decay with a partial branching fraction into a hidden sector would require a seemingly unnatural coincidence of unrelated scales or couplings.

3.1.4 Impact of Detector Resolution and Thresholds

The possibility of background-free LLP detection in MATHUSLA relies on being able to assign at least two separate, upwards-going particle tracks to a DV in the detector’s decay volume. If the LLP daughter particles are too soft, the DV may not be reliably detected. If the LLP daughter particles are too collimated, it may only give rise to a ‘merged’ one-pronged DV. Detailed study of this signal is needed, but it would likely suffer significantly higher backgrounds than well-separated multi-pronged DVs. Physics reach is therefore maximized if well-motivated LLP scenarios can be detected as multi-pronged DVs. Here we discuss the impact of detector spatial resolution and energy thresholds to determine the regions of LLP parameter space that fall into the multi-pronged DV regime, with important implications for the final design of MATHUSLA.

Since LLPs produced in LHC collisions are generally very energetic compared to minimum ionization energies, energy thresholds of the MATHUSLA detectors are expected to play a less crucial role in determining sensitivity than spatial resolution. However, in cases where LLPs decay intrinsically to soft final states (e.g. dark shower models as in Section. 8.1), we assume the following minimum thresholds on charged particle three-momenta $|\vec{p}|$ for detection: pions, 200 MeV; charged Kaons, 600 MeV; muons, 200 MeV; electrons: 1 GeV; protons: 600 MeV; photons: 200 MeV.

To discuss the impact of spatial resolution, assume for simplicity that an LLP decays into two massless charged SM particles. (This discussion can be easily extended to higher-multiplicity final states or decays closer to kinematic threshold, but this does not qualitatively change the conclusions.) The characteristic opening angle of the decay products is then

$$\theta \sim \frac{1}{\bar{b}} \quad (12)$$

where \bar{b} is the average boost of the LLP. The spatial resolution of the tracker panes required to separate the decay products is therefore

$$\Delta x \sim (10m) \theta \sim \frac{10m}{\bar{b}}, \quad (13)$$

corresponding to a maximum LLP boost for detection of a multi-pronged DV,

$$b_{LLP}^{max} \sim 1000 \left(\frac{1cm}{\Delta x} \right) \quad (14)$$

If the LLP boost can be expressed in terms of an effective parent mass as in Eqn. (6), this can be translated into a *minimum LLP mass* for multi-pronged DV detection:

$$m_{LLP}^{min} \sim \frac{m_{parent}}{2b_{LLP}^{max}} \sim \left(\frac{m_{parent}}{2000} \right) \left(\frac{\Delta x}{1cm} \right) \quad (15)$$

For LLPs produced in the decay of a 0.1 - 1 TeV parent, this corresponds to a minimum LLP mass of $\sim 0.1 - 1$ GeV for cm spatial resolution. For light LLPs produced in B -hadron decays, the minimum mass is about 10 MeV, see Table. 1.⁵

The benchmark detector described in Section 2 assumes a $\sim cm$ spatial resolution. Clearly, one could lower the minimum discoverable LLP mass by improving this resolution, but this has to be balanced against cost. Options include having different resolutions in the horizontal x and y direction, or using finer segmentation in only a part of the MATHUSLA detector. However, even the baseline resolution allows MATHUSLA to discovery LLPs with very low masses below a GeV and perhaps even close to the MeV-scale.

⁵Note that if photon detection is possible in MATHUSLA, the minimum mass for multi-pronged DV detection of an LLP decaying to 2γ would be higher, since photons detection by conversion in material degrades angular resolution.

3.2 Comparing LLP reach at MATHUSLA to the HL-LHC Main Detectors

To understand the physics case for MATHUSLA, it is important to compare its capabilities for discovering very long-lived neutral BSM particles with those of the HL-LHC main detectors. ATLAS or CMS could discover such particles in two ways:

1. as missing energy in MET searches; or
2. through dedicated LLP searches.

We compare the projected reach of HL-LHC MET searches to the MATHUSLA reach in Section 3.2.1 for several important simplified models, and demonstrate that MATHUSLA can probe large regions of parameter space inaccessible to the main detectors. Further, even if a new particle is detected first as MET at the HL-LHC, MATHUSLA will still have an important role to play in characterizing its lifetime. This is obviously a question of great cosmological significance.

A quantitative comparison of MATHUSLA to direct HL-LHC LLP searches is much more challenging, as ultimate trigger capabilities and background rates for HL-LHC LLP searches are less well-established. The main detectors can search for neutral LLP decays as (i) displaced tracks or vertices in the tracker, (ii) isolated energy deposits in the calorimeters, and (iii) displaced vertices in the ATLAS Muon System. Since the detectors were not designed for LLP searches, reconstruction and triggering require dedicated algorithms and can be challenging. While LLP signals can be spectacular and are inherently low-background compared to prompt searches, the backgrounds that do exist are frequently non-collisional and hence difficult to characterize from first principles. These backgrounds will typically become increasingly important as the LHC luminosity increases, and must be taken into account in establishing the ultimate reach of LLP searches at the HL-LHC.

As we show in Section 3.2.2, MATHUSLA and the main detectors have very similar geometric acceptances for LLP decays in the long lifetime regime. The crucial advantage of MATHUSLA is thus not its enormous volume (which merely compensates for its distance from the IP), but that it operates almost entirely without backgrounds or triggering issues. Thus, for any LLP search which suffers backgrounds or is challenging to trigger on at the main detectors, MATHUSLA will beat the HL-LHC in cross-section sensitivity by up to several orders of magnitude.

There are very general and well-motivated classes of neutral LLP scenarios for which triggering and backgrounds are both major obstacles at the HL-LHC. For instance, many models where LLPs are produced in exotic Higgs decays [1] typically yield low-mass ($m \lesssim \mathcal{O}(100 \text{ GeV})$), hadronically-decaying LLPs without accompanying hard or leptonic prompt objects in the final state. The characteristic low-mass, hadronic final states in these models present challenges for both triggering and background rejection. Conversely, there are some LLP scenarios where the relative advantage enjoyed by MATHUSLA is much smaller: for example, an LLP with a TeV-scale mass decaying leptonically will likely be easy to trigger on and is unlikely to have much background. For scenarios in between, quantitative statements about the sensitivity gain offered by MATHUSLA are more difficult to make. They have to rely on results from the relatively small number of published LLP searches and studies, which can often be difficult to extrapolate to the different running conditions, detector capabilities, and search strategies available at the HL-LHC. Nevertheless, we can make very universal qualitative statements about how important MATHUSLA will be to cover the LLP parameter space, and parameterize our ignorance of LLP backgrounds at the main detectors in such a way that the results of future studies, or rough experimental intuition, can be utilized to understand the sensitivity gain from MATHUSLA in more detail. This is discussed in Section 3.2.3.

Our primary focus is neutral LLPs, but charged or colored LLPs can also be considered, since for masses above a few hundred GeV and sufficiently long lifetime, only a small fraction of such LLPs will be stopped in the rock before they reach the surface [108, 109]. Their decay can be reconstructed at MATHUSLA, possibly with different background considerations since the passage of such LLPs will register in the scintillator veto surrounding the decay volume. In general, the HL-LHC coverage for such

LLPs is quite good if they have long lifetime [110–112], since they are not invisible and leave signals in most detector subsystems. In that case, MATHUSLA will offer complementary information.

3.2.1 Comparing MATHUSLA reach to MET searches with the Main Detectors⁶

It is natural to ask whether missing energy searches could effectively probe LLPs with very long lifetimes. In this section, we present updated projections of the monojet + MET search reach at the HL-LHC for three canonical scenarios: exotic Higgs decays to invisible particles ($h \rightarrow \text{invis}$), DM simplified models, and supersymmetry. One can then compare the MET reach to the reach of MATHUSLA by assuming that the invisible neutral particle is instead unstable. We also compare the MATHUSLA reach to the reach of a simple $\text{MET}_{\text{PV}} + \text{LLP}$ search at the main detectors, which adds a DV requirement to the MET search and computes MET using only primary vertex information.

The $h \rightarrow \text{invis}$ projections are computed following current LHC practices for dark matter searches. We assume a generic LHC detector that stands in for either ATLAS or CMS, using a simple in-house detector simulation that models current running conditions and gives equivalent results to Delphes [113]. The MET trigger is assumed to be efficient above 200 GeV, and assumed to work in HL-LHC conditions. This is consistent with studies shown in [114], where a L1 Trigger with particle flow and PUPPI (for pile-up mitigation) is presented, and shown to give a consistent MET trigger across the full intensity range of the upgraded LHC.

The monojet search is very inclusive, requiring at least one jet and missing energy above the 200 GeV trigger threshold. Leptons are also vetoed with rapidity up to $|\eta| < 4$ assuming realistic inefficiencies that contribute a residual background, dominantly from $W \rightarrow \tau\nu$. The dominant Higgs signal comes from VBF, but gluon fusion (ggF) also contributes. VBF and ggF contributions are separately constrained using a two category fit of MET and m_{jj} . A series of five separate control regions consisting of a single muon/electron/photon and double muon/electron are used in a simultaneous fit *in situ* with the signal regions to constrain both the $Z \rightarrow \nu\nu$ and $W \rightarrow \ell\nu$ backgrounds. This method can be extended to constrain top background, but we do not make use of this method here: instead, we use standard MC predictions and apply the same extrapolation uncertainties as for the W background. To extrapolate from the control regions to the signal region we apply the extrapolation uncertainty scheme following the NNLO QCD+NLO EW predictions [115]. As a check, a more conservative uncertainty scheme is applied, which consists of the predicted uncertainties scaled up by an order of magnitude. This uncertainty scheme corresponds to the NLO QCD+NLO EW predictions where the full EW scale corrections are taken as uncertainty.

The resulting limit projection on $\text{Br}(h \rightarrow \text{invis})$ is shown in Fig. 9. With our assumptions about systematic uncertainty, the invisible branching ratio limit with 3000 fb^{-1} of luminosity is $\sim 1 - 2\%$, a significant improvement on the earlier projection of $\sim 7\%$ [116].

The corresponding reach at MATHUSLA on $\text{Br}(h \rightarrow XX)$ for LLPs X is readily computed by requiring 4 LLP decays within the detector volume, reproducing the analysis in [1]. MATHUSLA’s branching ratio reach as a function of LLP mass and lifetime is shown in Fig. 10 (left). MATHUSLA is orders of magnitude more sensitive than the MET search for a large range of lifetimes from meters to 100 km, probing branching ratios as small as $\sim 10^{-5}$.

A very powerful extension of the MET searches at the HL-LHC makes use of main detector upgrades that will allow some tracking information to be used at L1 [117]. This would allow tracks from the LLP decay to be removed from the MET at L1, either by explicit DV reconstruction if possible at L1, or simply because a L1 track trigger is likely to remove tracks not originating from the IP.⁷ The LLP

⁶Philip Harris

⁷Current upgrade trigger designs in CMS are considering a vertex constraint on the MET algorithm using Puppi, a PUPPI MET trigger. Through PUPPI this vertex constraint is capable of associating both neutrals and charged particles to the PV and neglecting any unassociated tracks or neutrals near these tracks. This would preserve the MET trigger even in the instance of displaced tracks, provided they are either not reconstructed or pointing sufficiently far away from the PV.




Fig. 9: Projected $\text{Br}(H \rightarrow \text{invis})$ monojet bounds as a function of HL-LHC luminosity. The blue curve is uses the systematic uncertainty extrapolated from the NNLO QCD+NLO EW predictions following [115]. The purple curve inflates that uncertainty by a factor of 10, while the orange curve assumes no systematic uncertainty.




Fig. 10: $\text{Br}(h \rightarrow XX)$ projected bounds at the HL-LHC as a function of LLP mass, from a background-free DV search at the MATHUSLA detector (left) or from the $\text{MET}_{\text{PV}} + \text{DV}$ search using the main detectors (right). In the shaded regions, the updated $\text{Br}(h \rightarrow \text{invis})$ bounds from Fig. 9 are stronger than the direct LLP bounds, but it is important to keep in mind that detection of an invisible Higgs decay would only add motivation to an LLP search at that signal rate.

would therefore not contribute to MET, allowing the MET trigger to be used for DV searches, where the LLP decay in the tracker is reconstructed in the off-line analysis. To distinguish this cleaned-up MET variable from the conventional MET discussed in this section, we refer to MET computed using only primary vertex information as MET_{PV} .

For the secondary (displaced) vertex identification, multiple schemes were considered consisting of progressively tighter secondary vertex identifications. The final scheme adopted follows the most recent displaced vertex search performed by ATLAS [118]. The final selection consists of a fit of the MET (without the LLP) and using the same control region constrained fit as used for the dark matter searches. The second vertex identification efficiency and background rates are taken from the efficiency maps and fake rate estimates shown in [118]. These results were cross checked on a related SUSY model and found to give consistent results, including for compressed spectra which somewhat mimic the kinematics of the exotic Higgs decay final state. The $\text{Br}(h \rightarrow XX)$ reach of this $\text{MET}_{\text{PV}} + \text{DV}$ search at the HL-LHC main detectors is shown in Fig. 10 (right). It is clearly highly complementary to

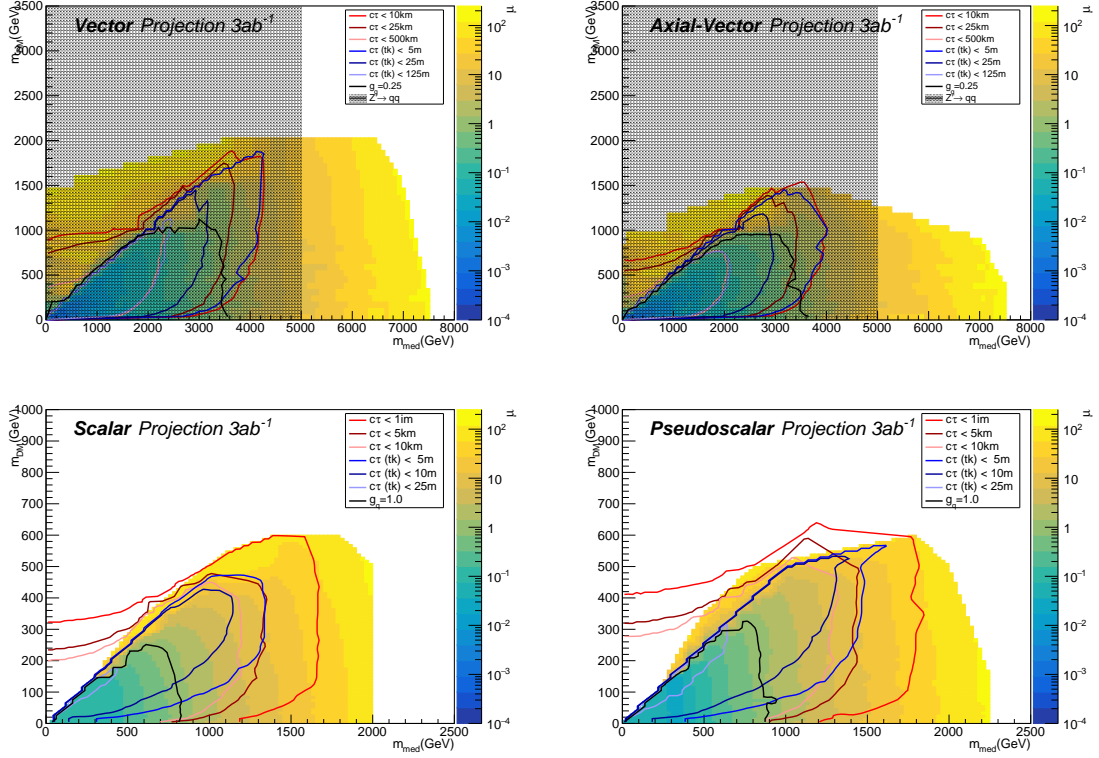


Fig. 11: Projected HL-LHC limits on simplified dark matter models from the MET-only monojet dark matter search (black), a $\text{MET}_{\text{PV}}+\text{DV}$ search assuming a range of lifetimes for the invisible particle (blue), and the MATHUSLA detector (red) for a different set of lifetimes. *Top:* The bounds are shown for a spin-1 mediator ($g_q = 0.25$ and $g_{DM} = 1.0$ with no assumed mediator to lepton couplings) for vector couplings (left) and for axial-vector couplings (right). The shaded black area corresponds to the expected bound from dijet searches projected out to 3000fb^{-1} (these searches lose sensitivity for $g_q \lesssim 0.1$), see text. *Bottom:* The bounds are shown for a spin-0 mediator ($g_q = 1.0$ and $g_{DM} = 1.0$ with SM Higgs-like Yukawa couplings rescaled by g_q to the SM fermions for scalar couplings (left) and for pseudoscalar couplings (right).

MATHUSLA, with great sensitivity for much shorter lifetimes. The trigger upgrades are crucial for this projected sensitivity, since it allows the MET trigger to be used even if the LLP decays in the detector.

Next we examine some canonical DM simplified models [119–124] where a fermionic dark matter candidate couples to a vector, axial vector, scalar or pseudoscalar mediator with coupling g_{DM} while the mediator couples to quarks with coupling g_q (spin 1) or Higgs-like Yukawa couplings (spin-0) scaled by a flavor-universal prefactor g_q . The analysis proceeds exactly like in the invisible Higgs decay case, as does the $\text{MET}_{\text{PV}}+\text{DV}$ search and the MATHUSLA DV search, assuming the invisible particle to be unstable instead of a DM candidate. Fig. 11 compares bounds in the mediator-DM mass plane, from the MET-only monojet search (black contours), the $\text{MET}_{\text{PV}}+\text{DV}$ search for a range of fixed lifetimes (blue) and from the MATHUSLA DV search for a different range of fixed lifetimes (red). In all cases, MATHUSLA significantly extends the mass reach for large ranges of lifetimes compared to the MET and $\text{MET}_{\text{PV}}+\text{DV}$ searches. While the $\text{MET}_{\text{PV}}+\text{DV}$ search is again complementary to MATHUSLA at shorter lifetimes, only MATHUSLA has sensitivity in the regime where the mediator is off-shell, since in that case the p_T spectrum of signal events is very difficult to distinguish from the $Z \rightarrow \bar{\nu}\nu$ background. The sensitivity gain is especially pronounced for spin-0 mediators, due to their lower mono-jet efficiency compared to vector mediators that are produced dominantly in valence quark collisions with more additional radiation.




Fig. 12: Expected minimum coupling g_q probed with $3000fb^{-1}$ of data for the optimal lifetime value (scanning all lifetimes) for spin-1 vector (top left) and axial vector (top right) mediators, and spin-0 scalar (bottom left) and pseudoscalar (bottom right) mediators, as a function of dark matter mass m_{DM} . In this plot $g_{DM} = 1$, and the mediator mass m_{med} is fixed to be exactly three times the mass of the dark matter. The black line corresponds to the minimum coupling probed in a MET only search, the blue line corresponds to the case where a secondary vertex is identified in the missing energy search, and the red line corresponds to the minimum coupling with the MATHUSLA detector. The dashed lines correspond to visible search in either a boosted jet (dashed-orange) or a di-jet resonance (dashed-green), for both searches the dark matter branching ratio is accounted for. Lastly the solid green line Ωh^2 corresponds to the minimum allowed coupling that will not overproduce dark matter (assuming a single dark matter candidate) for the shown spin-1 models where no lepton couplings are added, and the shown spin-0 models with a Yukawa coupling. While this coupling has no physical significance if the invisible particle is an LLP, it allows the shown coupling reaches to be compared to the sensitivity goal of the monojet searches.

For vector mediators, much of the mass range that is accessible by either MET, MET_{PV} + DV or MATHUSLA searches will also be covered by Z' dijet resonance searches. We obtain an HL-LHC sensitivity projection of dijet resonance search by rescaling the current limits [125–129] and show the resulting reach as the black shaded regions in Fig. 11. Note however that dijet searches have loose sensitivity for $g_q \lesssim 0.1$, while LLP searches would continue to see a signal (albeit with reduced mass range compared to the examples shown here), see Fig. 12. In that case, MATHUSLA may be the only way to see these models. Furthermore, a dijet resonance signal will not reveal that the produced resonance has a non-SM decay mode, let alone into unstable particles.

The sensitivity of dedicated LLP searches, either in the main detector or at MATHUSLA, is also illustrated by Fig. 12, where we show the minimum invisible particle coupling to the mediator (for the optimal range of lifetimes) that can be probed by the MET, MET_{PV} + DV and MATHUSLA searches. The spectacular nature of the LLP signal, with zero or greatly reduced background, means that the




Fig. 13: Projected HL-LHC limits on single squark pair production and decay into a neutralino, with bounds for projected limits to a total for 3000fb^{-1} of the monojet dark matter search (black), a search for a displaced vertex for fixed lifetimes (blue), and the MATHUSLA detector (red) for additional fixed lifetimes. The bounds are shown for squark pair production using a cross-section corresponding to a single light squark pair, typical bounds use four light squark flavors in the final state.

smallest couplings that can in principle be probed are one or two orders of magnitude smaller than the reach of the corresponding MET search.

Finally, we perform exactly the same analysis and comparison for a SUSY simplified model with a single light squark being pair-produced and decaying to two jets and two neutralinos. The mass reach of the MET-only monojet search, the $\text{MET}_{\text{PV}} + \text{DV}$ search and the MATHUSLA LLP search is shown in Fig. 13. This simple example of a SUSY decay chain is different from the scenarios discussed above, since the invisible particle is always produced in association with hard jets. As a result, the MET trigger is much more efficient than in the Higgs- and DM-related searches. For a heavy squark and a very light neutralino, there is little background and the reach of the MET search is signal-limited. Therefore, the $\text{MET}_{\text{PV}} + \text{DV}$ and MATHUSLA LLP searches do not significantly extend mass reach in the light neutralino case (though they would again be required to correctly diagnose the invisible particle as being an LLP). However, if the neutralino mass is even an $\mathcal{O}(1)$ fraction of the squark mass, the MET search becomes much less efficient and squark mass reach decreases drastically. On the other hand, the $\text{MET}_{\text{PV}} + \text{DV}$ and MATHUSLA LLP searches are unaffected, and greatly extend mass reach in regions of parameter space where the neutralino is of comparable mass to the squark.

While the SUSY scenario we studied was that of a single light squark species, broadly similar conclusions can be drawn for other searches where the LLP is produced in a decay chain, like gluinos, additional light squarks, EW SUSY partners, and non-SUSY scenarios with similar topologies. The MET search will be very efficient and signal limited if the invisible particle is very light, and at high parent particle masses. For moderately heavy LLPs compared to the parent mass, not to mention highly compressed regions, the MET search loses sensitivity. While strategies exist to probe these challenging spectra at the main detectors [130] it is likely for many scenarios that $\text{MET}_{\text{PV}} + \text{DV}$ and MATHUSLA LLP searches are the only discovery channels.

The benchmark scenarios studied here allow for some universal conclusions to be drawn. For invisible particles produced directly in Higgs decays or via mediators, MATHUSLA and other LLP searches significantly extend the mass range in both mediator/parent particle mass and invisible particle mass, since MET searches rely on additional radiation to trigger. For invisible particles produced in decay chains, LLP searches extend mass reach into (even very slightly) compressed regions. We also reiterate the fact that even if an LLP is first discovered in MET searches, detection of its decay is the only

way to ascertain whether the newly produced particle can have a cosmological role as a DM candidate, or whether it is part of a BSM spectrum of unstable states. It is clear that MATHUSLA offers great discovery potential that is inaccessible to HL-LHC MET searches.

3.2.2 LLP Signal Estimate for the Main Detectors

To understand the general HL-LHC sensitivity to LLP decays compared to MATHUSLA, it is helpful to begin with a few general comments about the size of a potentially observed LLP signal. The following discussion applies regardless of how a particular search is constructed (i.e., trigger requirements, prompt cuts, etc.) or how the LLP is reconstructed in detail (i.e., as a displaced vertex or displaced track, in the tracker or muon system, calorimeter deposition, etc.).

For a given LLP search, the number of observed signal events with a reconstructed LLP decay at the HL-LHC can be estimated as follows:

$$N_{\text{sig}}^{\text{LHC}} \approx (\sigma_{\text{sig}}^{\text{LHC}} \mathcal{L}) n_{\text{LLP}} \epsilon_{\text{LLP}}^{\text{LHC}} P_{\text{decay}}^{\text{LHC}}(\bar{b}c\tau) \epsilon_{\text{cuts}}^{\text{LHC}}, \quad (16)$$

where $\epsilon_{\text{LLP}}^{\text{LHC}}$ is the efficiency for reconstructing an LLP decay that occurred in the specified main detector subsystem (tracker, muon system, calorimeter, depending on the analysis), n_{LLP} is the number of LLPs produced in a single event, $P_{\text{decay}}^{\text{LHC}}(\bar{b}c\tau)$ is the chance that an LLP decays in that detector subsystem, and $\epsilon_{\text{cuts}}^{\text{LHC}}$ is the chance that *those events with reconstructed LLP decays* also pass the *trigger* and *off-line* analysis cuts which do *not* pertain to the displaced nature of the LLP decay. (Note the order in which these efficiencies are defined.)

Arriving at a realistic signal estimate through Eq. 16 for a specific signal requires a dedicated collider study. In the long-lifetime regime, however, several major simplifications occur, which helps make the comparison with MATHUSLA's capabilities transparent and robust.

- The chance $P_{\text{decay}}^{\text{LHC}}(\bar{b}c\tau)$ that an LLP with a lifetime $\bar{b}c\tau \gg 200m$ decays in a given main detector subsystem can be expressed relative to the corresponding MATHUSLA LLP acceptance in a nearly process-independent way. For the LLP production processes listed in Table 1, we observe:

$$\frac{P_{\text{decay}}^{\text{LHC}}}{P_{\text{decay}}^{\text{MATH}}} \approx \begin{cases} 2.2 & \text{ATLAS Muon System} \\ 0.8 & \text{ATLAS HCAL} \\ 1.0 & \text{ATLAS or CMS tracker (full volume)} \\ 0.25 & \text{ATLAS tracker (DV reconstruction volume)} \end{cases} \quad (17)$$

The DV reconstruction volume of the ATLAS tracker refers to the analysis in [131]. For this purpose, the ATLAS muon system barrel and endcap can be combined and their detection efficiencies $\epsilon_{\text{LLP}}^{\text{LHC}}$ averaged. The HCAL endcap has very low acceptance relative to the barrel and is neglected at this level of precision. The above relation holds for all examined LLP production modes at the $\sim 10\%$ level.

- The LHC LLP detection efficiency, relative to the corresponding efficiency in MATHUSLA, will be in the range

$$\frac{\epsilon_{\text{LLP}}^{\text{LHC}}}{\epsilon_{\text{LLP}}^{\text{MATH}}} \sim 0.1 - 1 \quad (18)$$

since $\epsilon_{\text{LLP}}^{\text{MATH}} \sim 1$ for hadronic decays given the redundant tracking required for cosmic ray background rejection, and $\epsilon_{\text{LLP}}^{\text{MATH}} \gtrsim 0.5$ for a LLP decaying to two widely separated leptons. On the other hand, characteristic reconstruction efficiencies for LLPs in the main detectors are $\epsilon_{\text{LLP}}^{\text{LHC}} \sim 0.3$ (ATLAS MS, [131]), 0.1 (DV in ATLAS tracker, [131, 132]) and 0.5 (CMS displaced jet in tracker, [133]).

- The final unknown factor, $\epsilon_{\text{cuts}}^{\text{LHC}}$, can be estimated using relatively simple simulations, provided the signal requirements are known. This requires an understanding of how various trigger and analysis thresholds would change at the HL-LHC compared to run 2. In many cases, especially if $\epsilon_{\text{cuts}}^{\text{LHC}}$ is dominated by a trigger requirement, it can also be analytically estimated from known kinematic distributions and branching ratios (which are important, for example, if the analysis relies on leptons which are only present in a fraction of signal events).

As we now discuss, this schematic understanding will be very helpful to understand MATHUSLA's advantage over the HL-LHC alone.

3.2.3 Comparing MATHUSLA reach to LLP searches with the Main Detectors

For purely geometric reasons, the HL-LHC main detectors will have superior sensitivity to MATHUSLA if the LLP has a relatively short lifetime $\bar{b}c\tau \ll 200\text{m}$. We are interested in understanding the relative sensitivity of the two experiments in the long-lifetime regime $\bar{b}c\tau \gtrsim 200\text{m}$. Below, we therefore discuss both hypothetical and performed searches that require a *single* observed LLP decay in the main detector, which will give the best limit in this regime.

If a given HL-LHC LLP search has $N_{\text{BG}}^{\text{LHC}}$ background events, let us parameterize this background simply by an effective background cross-section after analysis cuts:

$$N_{\text{BG}}^{\text{LHC}} \equiv \sigma_{\text{BG after cuts}} \mathcal{L}. \quad (19)$$

The exclusion limit of the LLP search at the HL-LHC can then be estimated by solving $N_{\text{sig}}^{\text{LHC}} / \sqrt{N_{\text{BG}}^{\text{LHC}}} = 2$, giving

$$\sigma_{\text{sig}}^{\text{LHC limit}} \approx \frac{2}{\epsilon_{\text{LLP}}^{\text{LHC}} P_{\text{decay}}^{\text{LHC}}(c\tau) \epsilon_{\text{cuts}}^{\text{LHC}} n_{\text{LLP}}} \sqrt{\frac{\sigma_{\text{BG after cuts}}}{\mathcal{L}}}. \quad (20)$$

In the absence of background, the HL-LHC sets an exclusion limit corresponding to $N_{\text{sig}} = 4$ in Eqn. (16):

$$\sigma_{\text{sig}}^{\text{LHC limit}} \approx \frac{4}{\mathcal{L} \epsilon_{\text{LLP}}^{\text{LHC}} P_{\text{decay}}^{\text{LHC}}(c\tau) \epsilon_{\text{cuts}}^{\text{LHC}} n_{\text{LLP}}} \quad (21)$$

We now define an important figure of merit, R_s , the *long-lifetime sensitivity gain of MATHUSLA*:

$$R_s \equiv \left. \frac{\sigma_{\text{sig}}^{\text{LHC limit}}}{\sigma_{\text{sig}}^{\text{MATHUSLA limit}}} \right|_{\bar{b}c\tau \gg 200\text{m}} \quad (22)$$

which is independent of lifetime in this regime. Making use of Eqns. (7), (20) and (21), we can express R_s in the following way:

$$R_s \approx \left(\frac{P_{\text{decay}}^{\text{MATH}}}{P_{\text{decay}}^{\text{LHC}}} \right) \left(\frac{\epsilon_{\text{LLP}}^{\text{MATH}}}{\epsilon_{\text{LLP}}^{\text{LHC}}} \right) \frac{1}{\epsilon_{\text{cuts}}^{\text{LHC}}} \times \text{Max} \left[1, \left(\frac{\sigma_{\text{BG after cuts}}}{10^{-3} \text{ fb}} \right)^{1/2} \right] \quad (23)$$

Eqn. (23) summarizes important information about the relative capabilities of MATHUSLA and the HL-LHC main detectors in the long lifetime regime:

- When the HL-LHC search is background-free, the arguments of Section 3.2.2 show that $R_s \gtrsim 1$, and possibly $R_s \gg 1$ if the trigger or reconstruction efficiency is low in the main detector, or if the analysis relies on a subdominant production or decay mode of the LLP (e.g. leptons). MATHUSLA will therefore never do much worse than the HL-LHC, and will do better in many cases.

- When there is *any* background above an ab level, $R_s > 1$ automatically. The relative sensitivity gain of MATHUSLA can then be very large, and can be estimated by inserting the effective background cross-section into Eqn. (23).

For many possible searches for a single LLP decay at the HL-LHC, the size of the backgrounds is still unknown, but once the corresponding experimental studies are completed, the relative sensitivity gain can be estimated using Eqn. (23). However we can already make some general statements based on information from experimental analyses at LHC Runs 1 and 2.

First, the LHC main detectors have excellent capabilities for LLPs that decay to well-separated pairs of leptons (e^+e^- , $\mu^+\mu^-$) in the tracker. Here lepton triggers can be used to record the event without any requirements on associated prompt objects [134]. This provides excellent and inclusive sensitivity to LLPs with masses as low as $\mathcal{O}(10 \text{ GeV})$. Backgrounds are negligible in the Run I searches even in the absence of additional cuts on the momenta of the displaced lepton vertex or on prompt objects [134], suggesting that the background cross-section for well-separated displaced lepton pairs will likely be very small even at the HL-LHC, in the ab-range or below.

If the LLP has a mass in the few GeV range or below, its decay to leptons gives rise to the displaced lepton jet (LJ) final state. To date, no search for a single displaced LJ has been performed, but ATLAS conducted a search [135] at 13 TeV for at least *two* displaced LJs decaying in the tracker, calorimeters, or lower regions of the muon system. The search is very inclusive, with no additional prompt signal requirements, and acceptance-times-efficiency (per event) is in the 0.1-0.3 range. However, even with the requirement of two reconstructed LLP decays instead of one, the background in the signal region corresponds to $\sigma_{\text{BG after cuts}} \sim 10\text{fb}$. This demonstrates that backgrounds for a single displaced LJ in the ultra-long lifetime limit would be much higher than for well-separated leptons, representing an opportunity for MATHUSLA to make very large sensitivity gains compared to the HL-LHC main detectors.

LLPs decaying hadronically in the tracker are more challenging for the main detectors, especially at low mass. This difficulty begins in the trigger. Triggering options for displaced decays are limited by the need to pass the Level 1 (L1) hardware triggers. While it may be possible to implement L1 triggers based on properties of the LLP decay itself, many analyses rely on L1 triggers optimized for prompt physics even if higher-level triggers are designed for LLP decays [133, 136]. Often in the long lifetime regime of interest, analyses using the MET and H_T triggers will be most important. This reliance on prompt triggers means that events must typically be relatively energetic to be recorded to tape, thus limiting LHC sensitivity to low-mass or low-energy ($\lesssim 100 \text{ GeV}$) final states. Trigger thresholds will generally rise at the HL-LHC to cope with the increased luminosity, exacerbating the issue. Even the dedicated displaced triggers at ATLAS targeting decays in the outer tracker and HCAL [136] are based on trigger objects with E_T thresholds of $\sim 50 \text{ GeV}$ or GeV that will also be expected to increase.

For this reason, the searches for single LLPs decaying hadronically in the tracker which have been carried out by ATLAS and CMS typically rely on high- p_T objects to clear Level 1 triggers [132, 133, 137]. These searches have typical signal reconstruction efficiencies in the 10–60% range, with acceptance and efficiencies increasing with increasing signal mass. In Run 1 these searches were effectively or nearly background free. The CMS displaced dijet search [133], which requires $H_T > 325 \text{ GeV}$ for triggering but does not require full displaced vertex reconstruction, sees $\sigma_{\text{BG after cuts}} \sim 60\text{ab}$; the ATLAS DV+MET search is less inclusive [132], imposing tighter vertex identification and a MET requirement of $\cancel{E}_T > 180 \text{ GeV}$, resulting in $\sigma_{\text{BG after cuts}} \sim 0.6\text{ab}$.

The dominant background for the ATLAS tracker searches for multi-track DVs occurs when a low-mass vertex is crossed by an unrelated high- p_T track, and will thus increase with luminosity [132, 137]. Indeed, the analogous ATLAS search at Run 2 maintained $\sigma_{\text{BG after cuts}} \sim 0.6\text{ab}$, at the cost of increasing the requirement on MET, $\cancel{E}_T > 250 \text{ GeV}$. This makes clear that backgrounds to LLP searches in the inner tracker will depend sensitively on the details of trigger and reconstruction, but will generically increase with luminosity. The already sizeable cuts imposed on event E_T scales in these analysis may also be

expected to increase at the HL-LHC, both to pass triggers and to control the increasing backgrounds. Thus we generically expect searches for hadronically decaying LLPs in the trackers with masses below a few 100 GeV to become increasingly challenging at the HL-LHC, while the prospects for high-mass LLPs will benefit from increasing luminosity.

On the other hand, LLPs decaying hadronically in the ATLAS Muon System can be triggered on directly using a L1 muon trigger seed and a dedicated higher-level LLP trigger [131, 136], with only weak dependence on the energy scale of the event for LLP masses $\gtrsim 5 - 10$ GeV. In the ultralong lifetime regime, a search for a single LLP decay in the Atlas Muon System is likely the best LHC search for low-mass, hadronically-decaying LLPs produced without additional prompt objects [138]. The background cross-section found by [138], derived using public data from [131], is of order $\sigma_{\text{BG after cuts}} \sim 100\text{fb}$. This is the leading way to probe hadronically decaying very long-lived particles produced in exotic decays of the 125 GeV Higgs.

Given this background cross-section and the reconstruction efficiencies $\epsilon_{\text{LLP}}^{\text{LHC}} \sim 0.5$, $\epsilon_{\text{LLP}}^{\text{MATH}} \sim 1$, along with $\epsilon_{\text{cuts}}^{\text{LHC}} \sim 1$, MATHUSLA’s sensitivity gain is substantial, corresponding to improving the cross-section reach by a factor of $R_s \sim 1000$. This agrees with the findings of the full study in [1]. Since the muon system is physically separated from the IP, and since the main background for displaced vertices in the muon system can be traced back to unusual high-energy QCD events, this analysis is somewhat unique in that its conclusions can probably be applied to the HL-LHC with a reasonable degree of confidence.

ATLAS and CMS also search for LLPs decaying in the calorimeters. Here a useful general lesson for very long-lived particles can be extracted by comparing the ATLAS search for two LLP decays in the HCAL [139] to the search for two LLP decays in the muon system (or one in the MS and one in the inner tracker) [138]. Both searches feature an inclusive search with no MET requirements that is sensitive to LLPs produced in exotic Higgs decays and in turn decaying hadronically. With 20.3 fb^{-1} of 8TeV data, the HCAL search has about 24 background events, while the MS search has about 2. This already indicates that for such LLPs, the ATLAS Muon System search described above provides the better inclusive sensitivity at the main detector. This is not surprising, since the ATLAS MS search has full displaced vertex reconstruction, while the information supplied by the HCAL is much less detailed. In general we expect that LHC LLP searches in the inner tracker or muon system will be the most powerful at the HL-LHC, as tracking capabilities will help control pileup events that spoil HCAL signal isolation and contribute non-collision backgrounds.

Depending on the final design chosen, it is possible that MATHUSLA may be able to detect photons from LLP decays by inducing conversion in material [9]. To date, CMS performed the only LHC search for a single LLP decaying to a single photon + MET [140] in the context of Gauge Mediation. Under the assumption that a neutralino NLSP is produced in the decay of heavy colored supersymmetric particles, and decays as $\tilde{\chi}_0^1 \rightarrow \gamma + \tilde{G}$, the search reconstructed a single LLP decay using timing measurements and required more than 60 GeV of MET, 2 jets with $p_T > 35$ GeV and the leading photon to have $p_T > 80$ GeV. Even with these additional kinematic cuts, the resulting background cross-section was non-negligible, $\sigma_{\text{BG after cuts}} \sim 0.5 \text{ fb}$. The importance of these non-LLP cuts to reduce backgrounds is illustrated by the ATLAS search [141], which looked for *two* LLPs decaying to $\gamma + \text{invisible}$, but imposed no additional cuts beyond $\text{MET} > 75$ GeV. That search had to contend with an effective $\sigma_{\text{BG after cuts}} \sim 20 \text{ fb}$. Given these examples, and the sensitivity of photon reconstruction to pile-up considerations, it is clear that searches for a single LLP decaying to a single photon will have orders of magnitude more background at the HL-LHC than searches for leptonic LLP decays.

Despite the difficulty of quantitatively extrapolating some of the above cases to the HL-LHC, these examples provide a useful point of reference for understanding MATHUSLA’s advantages compared to the main detectors. The greatest sensitivity gains, possibly by several orders of magnitude, apply for hadronically decaying LLPs with less than a few 100 GeV of visible energy (prompt or displaced), leptonically decaying LLPs with masses below 10 GeV, and (if detectable) LLPs decaying to photons.

3.3 Summary

The capabilities of MATHUSLA can be summed up in a few simple lessons:

1. If the LLP signal cross-section is greater than a fb, then MATHUSLA can see a signal for some range of lifetimes, see Section 3.1.3. In the long-lifetime limit, MATHUSLA's sensitivity to LLP production is readily estimated using Eqn. (4) and Table 1. The model-independent LLP cross-section sensitivity is shown in Fig. 7.
2. LLPs with average boosts lower than $\mathcal{O}(1000)$ can be reconstructed as DVs with two or more prongs, see Eqn. (14). This results in excellent rejection of cosmic rays and other backgrounds, justifying the zero-background assumption for LLP searches. Higher boosts (and hence lower LLPs masses) would require additional analysis and likely suffer higher backgrounds, unless the detector resolution is increased.
3. Geometrically, MATHUSLA has very similar acceptance for LLP decays with $c\tau \gg 200\text{m}$ as the HL-LHC main detectors, see Eqn. (17), though it may have significantly higher reconstruction efficiency. Therefore, if the corresponding LLP search at the main detectors has any appreciable background ($\sigma_{\text{BG after cuts}} > \text{ab}$) or low signal efficiency (small $\epsilon_{\text{cuts}}^{\text{LHC}}$, e.g. due to trigger issues, or requirements on the LLP production or decay mode), MATHUSLA will have better sensitivity.
4. The greatest sensitivity gains, possibly by several orders of magnitude, apply for dominantly hadronically decaying LLPs produced with less than a few 100 GeV of visible energy (prompt or displaced), leptonically decaying LLPs with masses below 10 GeV, and (if MATHUSLA can detect them) LLPs decaying to photons. See Section 3.2.3.
5. MATHUSLA is sensitive to mass regions that MET searches cannot reach, both in parent and LLP mass, and provides a means of diagnosing any MET signal that is found. See Section 3.2.1.

This will provide important intuition in assessing the physics case for the theoretically motivated LLP scenarios described in the following sections. Qualitatively, the above also applies to a mini-MATHUSLA with 1/10 the decay volume, see Fig. 7, though of course with the sensitivity gain reduced by one order of magnitude.

4 Theory Motivation for LLPs: Naturalness

With the discovery of the Higgs at the LHC, or something very similar to it, a resolution to the Hierarchy problem or Naturalness of the EW scale is an ever more pressing concern. From the modern Wilsonian understanding of renormalization, since the Higgs is a scalar field, the EW scale in the SM is in conflict with treating the SM as an EFT valid to energies greater than the EW scale. To avoid this inherent tension in any model where a scalar field is lighter than the cutoff of the theory requires some mechanism which makes the theory “natural”. This mechanism could be a symmetry; the Higgs might not be an elementary field; or there could be some sort of selection mechanism for the EW scale. There have been a proliferation of new ideas and more complicated models of naturalness in recent years given the dearth of BSM signals at the LHC typically associated with canonical models of naturalness. A common thread amongst all the classes of naturalness motivated models surveyed here is the presence of neutral long-lived states.

The nature of the Hierarchy Problem singles out the EW scale for new physics searches. In the theory frameworks of Supersymmetry (Section 4.1) and Neutral Naturalness (Section 4.2), this leads to predictions of LLP signatures that are intimately connected with the mechanism that stabilizes the weak scale. These LLPs would be produced with appreciable cross-sections at the LHC, and neutral long-lived states can naturally have lifetimes that are in the relevant range for MATHUSLA. In other models, like the Composite Higgs (Section 4.3) or Relaxions (Section 4.4), the existence of LLPs occurs as part of various possible complete models with lifetime as more of a free parameter. MATHUSLA then provides a new window beyond ATLAS and CMS into this theory space. In the rest of this section we discuss the motivation and predictions for LLPs in these theory frameworks.

4.1 Supersymmetry

Supersymmetry is the most well known and theoretically under control solution of the Hierarchy problem. It also has a wide array of naturally long-lived particles which are highly motivated. In particular, a mechanism which causes a long-lifetime, potentially measurable at MATHUSLA, exists solely within the MSSM alone without any additional model building or fine-tuning. This occurs when the gravitino is the LSP and as such the particles of the MSSM can decay to it. What is particularly compelling is that this mechanism only depends on the scale of SUSY breaking, \sqrt{F} . The gravitino mass in SUSY scales as

$$m_{3/2} \sim \frac{F}{M_{pl}}, \quad (24)$$

while the coupling of the NLSP to the gravitino scales as $1/F$. It turns out there is a range of F where $m_{3/2}$ is small enough and the coupling is sufficiently suppressed that the NLSP decays to the gravitino with a macroscopic lifetime. This region is commonly understood as Gauge Mediated SUSY breaking (GMSB), which is discussed in Section 4.1.2. More generally the concept can be extended to Sgoldstinos, the SUSY partner of the goldstino which is eaten by the gravitino in the super-Higgs mechanism (see Section 4.1.6). While the identity of the sGoldstino can be more general than the GMSB version of the MSSM, the ultimate cause of the long-lifetimes and scalings are the same.

Within the MSSM alone, it is also possible that there are long-lived particles coming from a hierarchy amongst SUSY particles. In particular if Gauginos are light and Sparticles are sufficiently heavy, then Gaugino decays are highly suppressed and have macroscopic lifetimes. This idea was originally put forth in the context of Split-Supersymmetry at the expense of tuning. However, in light of the Higgs discovery and other considerations the idea of Mini-Split supersymmetry has the same structure *and* Gaugino lifetimes in the MATHUSLA range (see Section 4.1.3).

Given that as of yet there are no excesses attributable to SUSY at the LHC, there has been increased interest in extensions of the MSSM. Nevertheless, these extensions have ubiquitous long-lifetime possibilities as well. The most commonly known possibility is R-parity violation (RPV), discussed in

Topologies		
LSP	Decay	Operator
\tilde{t}	$\bar{d} \bar{d}'$	λ'', η''
	$u \bar{\nu}$	η'
	$d \ell^+$	λ', η
\tilde{g}	$t \bar{d} \bar{d}' + c.c$	λ'', η''
	$t \bar{u} \bar{\nu} + c.c$	η'
	$t \bar{d} \ell^- + c.c$	λ', η
$\tilde{H}^0 / \tilde{H}^\mp$	$(t/b) \bar{d} \bar{d}' + c.c$	λ'', η''
	$(t/b) \bar{u} \bar{\nu} + c.c$	η'
	$(t/b) \bar{d} \ell^- + c.c$	λ', η

Table 2: Summary of various LSPs and their decay channels. The third column denotes the RPV operators from (25) or (26).

Section 4.1.1. To avoid the pitfalls of not-having a preserved R-parity, the dimensionless couplings of these operators are typically very small. Therefore they naturally give long-lifetimes when SUSY particles decay through an RPV operator.

Stealth SUSY (Section 4.1.4) is a more recent extension of the MSSM designed to avoid LHC bounds. It introduces another sector that the MSSM superpartners can decay to and which is approximately supersymmetric thus avoiding typical large MET signatures. However, the decay in this sector can also use the same mechanism as gravitino decays in GMSB and long lifetimes are a part of the experimentally preferred region.

Finally, there are natural extensions of the MSSM built to address the strong CP problem using an axion (Section 4.1.5). These models will naturally have SUSY partners of the axion, in particular the Axino. The Axino is a natural DM candidate to be the LSP. However, the decays of other SUSY particles are suppressed through the PQ breaking scale. Within the range of well-motivated PQ scales, the decay of SUSY particles to the Axino is also in the natural range of MATHUSLA.

4.1.1 RPV Supersymmetry⁸

Perhaps the most commonly studied BSM framework to stabilize the Higgs mass is the Minimal Supersymmetric Standard Model (MSSM). The MSSM predicts new scalar fields charged under SM gauge interactions, allowing one to write tree-level, renormalizable operators which violate baryon (B) and lepton (L) number. Since such flavor violation is highly constrained by low-energy experiments [142, 143], the most minimal solution to the B and L constraints is to forbid such operators by imposing a global Z_2 symmetry known as R-parity, under which all SM particles carry charge +1 and all supersymmetric partners carry charge -1. In this case, the lightest supersymmetric particle (LSP) is pair-produced, stable, and escapes the detector leading to MET signatures at the LHC.

However, given the null results for LHC SUSY searches thus far, perhaps SUSY manifests itself non-minimally, and one should consider R-parity violating (RPV) interactions. If R-parity is violated, the source must be small since these operators are highly constrained. If this is the case, the LSP may decay to SM particles via a hierarchically small RPV interaction, motivating one to search for LSPs with macroscopic lifetimes. For a review of RPV phenomenology, see [143–145].

The most general renormalizable Lagrangian allowing R-parity breaking is parametrized by the following superpotential written in terms of left-handed chiral superfields.

⁸Csaba Csaki, Eric Kuflik, Salvator Lombardo, Jared Evans, Brock Tweedie, Tim Cohen, Zhen Liu, Patrick Meade

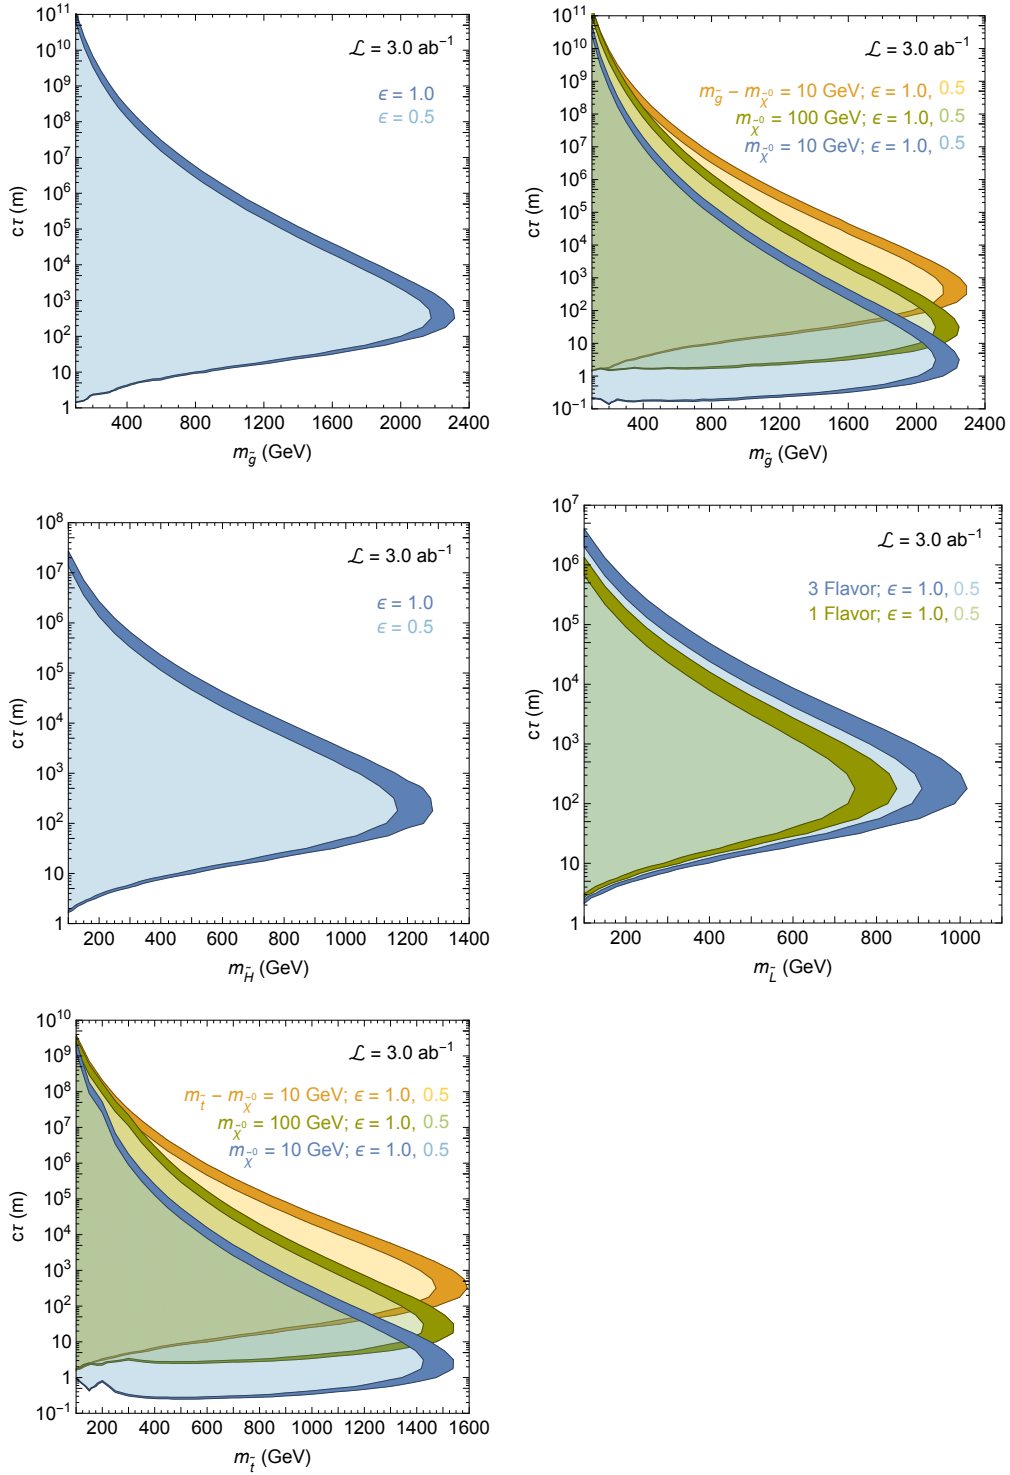


Fig. 14: Projected MATHUSLA reach (assuming the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1) for several LSP and production scenarios with 3 ab^{-1} . We assume that the decay of the LSP leads to visible particles in the MATHUSLA detector. In order: direct gluino pair-production, prompt gluino pair decay to long-lived binos (for discrete bino mass choices), direct Higgsino pair-production, direct sneutrino pair-production, and prompt stop pair decays to long-lived binos (for discrete bino mass choices). The efficiency is varied from 0.5-1 (darker bands) in order to take into account the difference in efficiencies for each particular LSP decay mode.

$$\mathcal{W}_{\text{RPV}} = \mu_i \ell_i h_u + \lambda_{ijk} \ell_i \ell_j \bar{e}_k + \lambda'_{ijk} \ell_i q_j \bar{d}_k + \lambda''_{ijk} \bar{u}_i \bar{d}_j \bar{d}_k \quad (25)$$

RPV interactions are usually parametrized by μ_i , $\lambda_{i,j,k}$, $\lambda'_{i,j,k}$, and $\lambda''_{i,j,k}$. The λ'' operator together with λ' lead to proton decay, while the μ term mixes neutrinos with the higgsinos and gauginos.

Depending on the UV completion, the usual RPV operators in Eq. 25 may not be the most important ones. For example R-parity may be a good symmetry in the UV of the visible sector but may be broken spontaneously by a SUSY-breaking sector and mediated to the visible sector dynamically [146, 147]. In this case, the dominant RPV operators could arise from the Kahler potential, allowing for the following non-holomorphic operators.

$$\mathcal{K}_{\text{nhRPV}} = \frac{1}{M} (\kappa'_i \ell_i^* h_D + \eta_{ijk} \bar{u}_i \bar{e}_j \bar{d}_k^* + \eta'_{ijk} q_i \bar{u}_j \ell_k^* + \eta''_{ijk} q_i q_j \bar{d}_k^*) \quad (26)$$

Here the RPV interactions are suppressed by the messenger scale M , explaining the small size of the interactions.

There are also cosmological considerations depending on the type of RPV operators that are active for the LLP decays. To avoid most searches at the LHC it is easier to have B violating operators which are potential hidden within QCD backgrounds. However, there is an interplay between introducing B violation and the baryon asymmetry that we observe in the universe. In particular, if the B violating RPV operator has a large coefficient, then its interactions could destroy any initial baryon asymmetry created through a standard baryogenesis mechanism. Whether the baryon asymmetry is washed out depends upon the reheat temperature which governs whether the interactions of the SM with SUSY particles are in equilibrium. Nevertheless, it is quite natural to expect long-lived LSPs when there is RPV from cosmological considerations [148].

The long-lived LSP has rich collider phenomenology since any of the superpartners could in principle be the LSP, and the largest RPV coupling determines the dominant final state of the LSP decay. Recently, the phenomenology and existing constraints on long-lived LSP scenarios was studied in detail [149, 150]. Motivated by naturalness, a summary of possible decay channels for stop, gluino, or higgsino LSP is provided in Table 2. Assuming the long-lived LSP decays to visible particles, we project the MATHUSLA reach for 3 ab^{-1} in Fig. 14 assuming a particular LSP and production channel (without assuming a particular LSP decay mode) for several scenarios including direct gluino pair-production, gluino pair decays to long-lived binos, direct Higgsino pair-production, direct sneutrino pair production, and pair-produced stops decaying to long-lived binos. The efficiency is varied from 0.5-1 in order to take into account the varying reconstruction efficiency for different LSP decay modes.

MATHUSLA can place strong constraints neutral LSPs which are competitive with MET searches performed by the LHC in the long lifetime limit. Even if a MET signal is observed at the LHC, the lifetime of the LSP can be probed by late decays in the MATHUSLA detector. Furthermore, there are regions of parameter space in which the MET in the event is suppressed, *e.g.* LSPs with mass below a few hundred GeV or scenarios with a compressed spectrum, where MATHUSLA can have more discovery potential than LHC MET searches. As argued in Section. 3, MATHUSLA would also have a significant advantage over LLP searches at the main detectors, especially for LLPs decaying hadronically with masses below a few 100 GeV. On the other hand, if the LSP is colored (*e.g.* a squark or gluino) and has a lifetime longer than $\Lambda_{\text{QCD}}^{-1}$, it hadronizes to form an R-hadron. LHC searches for heavy stable charged particles (*e.g.* [110, 151]) are sensitive to these scenarios. In this case, MATHUSLA can be complementary to these searches and can provide lifetime information of the R-hadron.

4.1.2 Gauge Mediation⁹

In gauge mediated supersymmetry breaking (GMSB), Standard Model superpartners decay with a potentially long lifetime to a much lighter gravitino \tilde{G} . A classic review of this scenario is [152]; a re-

⁹Matthew Reece

cent update on the experimental status is [153]. Here we will not assume gauge mediation in the strict sense [154], but any theory of low-scale SUSY breaking with a light gravitino, which could (for example) include Yukawa mediation effects as well [155, 156]. An important parameter dictating the long-lifetime phenomenology is the order parameter F_0 for supersymmetry breaking, which is related to the gravitino mass via

$$m_{3/2} = \frac{F_0}{\sqrt{3}M_{\text{Pl}}} \approx 1 \text{ keV} \left(\frac{\sqrt{F_0}}{2000 \text{ TeV}} \right)^2, \quad (27)$$

where $M_{\text{Pl}} \approx 2.4 \times 10^{18}$ GeV is the reduced Planck mass and $F_0 \equiv \sqrt{3}\langle e^{K/(2M_{\text{Pl}}^2)}W \rangle/M_{\text{Pl}}$ is determined by the full set of F -term VEVs of the theory. Standard Model superpartners acquire masses set by loop factors multiplying $F/M_{\text{mess}} = kF_0/M_{\text{mess}}$ where F is the dominant source of SUSY breaking mediated to the Standard Model, k is a model-dependent parameter satisfying $k \lesssim 1$, and $M_{\text{mess}} \ll M_{\text{Pl}}$ is the mass scale of “messengers” of SUSY breaking. The lightest MSSM superpartner, often called the LOSP (Lightest Ordinary SuperPartner), decays to Standard Model matter and the gravitino via higher-dimension operators that are suppressed by $1/F_0$. For example, for a neutralino LOSP, the lifetime is [157]

$$c\tau(\tilde{\chi}_1^0 \rightarrow \tilde{G} + \text{SM}) = \frac{16\pi F_0^2}{\xi m_\chi^5} \approx 100 \text{ m} \frac{1}{\xi} \left(\frac{\sqrt{F_0}}{10^7 \text{ GeV}} \right)^4 \left(\frac{250 \text{ GeV}}{m_\chi} \right)^5, \quad (28)$$

with ξ a constant depending on the neutralino mixing matrix and m_χ denoting the mass of the LOSP $\tilde{\chi}_1^0$. For example, for a pure bino, $\xi_{\text{bino}} = c_W^2 + s_W^2 \left(1 - \frac{m_Z^2}{m_\chi^2}\right)^4$; for a pure higgsino,

$$\xi_{\text{higgsino}} = \frac{1}{4} \left[(s_\beta + \text{sgn}(\mu)c_\beta)^2 \left(1 - \frac{m_Z^2}{m_\chi^2}\right)^4 + (s_\beta - \text{sgn}(\mu)c_\beta)^2 \left(1 - \frac{m_h^2}{m_\chi^2}\right)^4 \right]. \quad (29)$$

The large phase space suppression factors when $m_\chi \approx m_Z, m_h$ arise because the goldstino is derivatively coupled. The range of F_0 to consider ranges from low-scale SUSY breaking with $F_0 \sim 10^5$ GeV up to the intermediate scale $F_0 \sim 10^{10}$ GeV, where gravity-mediated effects become dominant. Across this range, the LOSP lifetime varies from prompt to macroscopic.

In Fig. 15, we show the number of events that MATHUSLA would detect for electroweak higgsino production at the LHC with a subsequent $\tilde{H}_1^0 \rightarrow (h, Z) + \tilde{G}$ decay in MATHUSLA. Kinematics of the higgsino events and leading order cross-sections were calculated using Pythia 8 [158]. Several points are noteworthy. First, the number of events observed at MATHUSLA could range up to several hundred, even for an electroweak production process with a relatively small cross-section. For $c\tau \sim 100$ m, MATHUSLA could exclude higgsino masses up to above 1 TeV. This may be compared to the LHC reach for higgsino production with a detector-stable higgsino LSP, which would not reach beyond masses of about 200 GeV even with optimistic assumptions about systematics [159, 160]. Furthermore, because the Z and h dominantly decay hadronically, HL-LHC searches for displaced higgsino decays using the main detectors will be less sensitive than MATHUSLA, especially for less than several hundred GeV of jet energy per decay (see Section 3.2.3 and also Section 4.1.5 for a discussion of this signal). It is well-established that measuring energies and angular distributions for macroscopically displaced decay events *inside* the LHC could provide a powerful handle on the underlying nature of GMSB physics, such as the mass and identity (e.g. wino or higgsino) of the LOSP [161–163]. A similar conclusion should be true of MATHUSLA as well: assuming a particular production mode, the mass and decay mode of the higgsino could be determined using the track geometry and multiplicity of the daughter products [9]. Complementary information from the HL-LHC main detectors could also help to unambiguously determine the properties of the higgsino. Production of strongly interacting particles (squarks or gluinos) that cascade down to LOSPs could lead to larger cross-sections and a greater discovery potential (cf. Section 3.1.3), in parameter space where new physics might first be spotted using the LHC main detectors (see the comparison of MATHUSLA to MET searches in Section 3.2.1).




Fig. 15: Number of $\tilde{H} \rightarrow \tilde{G} + (Z, h)$ events that MATHUSLA would detect from electroweak production of higgsinos at the LHC operating at $\sqrt{s} = 14$ TeV with an integrated luminosity of 3 ab^{-1} . Left: higgsino mass μ versus lifetime $c\tau$ in meters. Right: higgsino mass μ versus the SUSY breaking scale as parametrized by \sqrt{F} in GeV (label on left axis) or $m_{3/2}$ in keV (label on right axis). In a wide swath of parameter space with LOSP lifetimes ranging from 10 to 10^5 m, MATHUSLA could provide a discovery of new physics with electroweak cross-sections for which the LHC would fail to discover new physics. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

The right-hand panel of Fig. 15 shows the number of observed events at MATHUSLA plotted against $\sqrt{F_0}$ and $m_{3/2}$ on the left and right axes respectively. We see that MATHUSLA is sensitive to $\sqrt{F_0} \sim 10^6$ to 10^8 GeV and $m_{3/2} \sim 1$ keV to 1 MeV. This significantly extends the LHC's GMSB discovery potential toward larger SUSY-breaking scales.

4.1.3 Mini-Split Supersymmetry¹⁰

It is possible that nature could be fundamentally supersymmetric, but the weak scale could still be highly fine tuned.¹¹ This paradigm is known as Split SUSY [165–168]. Phenomenologically, this is motivated by the notion that SUSY would still explain dark matter via relic neutralinos while additionally accommodating gauge coupling unification. Furthermore, the R -symmetry implicit to the MSSM provides a rationale for superpartner fermion masses to be parametrically lighter than the scalar masses – this spectrum can be additionally motivated by models where the SUSY breaking is communicated to the scalar sector via gravity mediation, and the gauginos via anomaly mediation [169, 170] such that there is a loop factor difference between the scalar masses and gaugino mass [169–172] (this hierarchy can also occur in models of gauge mediation, *e.g.* [173]).

One of the interesting consequences of Split SUSY (as applied to the MSSM) is that the Higgs mass is predicted to lie within a finite range, with an upper bound given by ~ 140 GeV [174]. The discovery of the Higgs at 126 GeV, along with the lack of any BSM discoveries thus far, has reinvigorated interest in this paradigm, which has been re-coined Mini-split SUSY [175] due to the fact that the scalar superpartners cannot have their masses M_{SC} arbitrarily close to the Planck scale [175–179], see Fig. 16.

¹⁰Tim Cohen

¹¹Perhaps this tuning can be alleviated by additional physics, *e.g.* [164] in the case of a relaxion extension.




Fig. 16: The SM-like Higgs boson mass contours in the $\tan\beta$ versus M_{SC} plane. The parameter M_{SC} is the common mass scale of the scalar superpartners. This figure is taken from [177].

The only particles that are expected to be collider accessible for Mini-Split models are the gauginos.¹² As will be emphasized in the next paragraph, it is straightforward to find parameter space where the gluino lifetime is long, making it a natural target for MATHUSLA. For contrast, the electroweak gauginos all decay via the weak interactions, which do not lead to long enough lifetimes to be of interest without extremely small splittings. This case will be discussed briefly towards the end of this section.

The gluino decays via a higher dimension operator suppressed by inverse powers of the squark masses [181]:

$$\mathcal{O}_{\text{decay}}^{(6)} \sim \frac{g_s^2}{m_{\tilde{q}}^2} \bar{q} \tilde{g} \tilde{\chi} q \quad \mathcal{O}_{\text{decay}}^{(5)} \sim \frac{g_s^2}{16\pi^2 m_{\tilde{q}}} \tilde{g} \sigma_{\mu\nu} \tilde{\chi} G^{\mu\nu}, \quad (30)$$

where the superscript on the operator corresponds to the mass dimension. The dimension-6 operator comes from the tree-level exchange of an off-shell squark, while the dimension-5 operator comes from a one-loop diagram.

The gluino pair production cross-section is provided in Sec. 3. For reference, we provide an estimate of the lifetime using just the dimension-6 operator:

$$c\tau \sim 100 \text{ m} \left(\frac{m_{\tilde{q}}}{50 \text{ PeV}} \right)^4 \left(\frac{\text{TeV}}{m_{\tilde{g}}} \right) \quad (31)$$

Comparing this estimate to Fig. 16, we conclude that there is ample room for this model to be consistent and observable by MATHUSLA. There is clearly a competition between the two decay modes, as illustrated in Fig. 17. Furthermore, there are variations due to both the mass spectrum of the squarks (which effects which flavor dominates in the gluino decays) and the mass spectrum of the gauginos. For

¹²In principle the higgsinos could also be light enough to be collider accessible. For our purposes here, we will assume that the Higgsino mass scale is set by SUSY breaking in the context of the Giudice-Masiero mechanism [180]. This is in keeping with the desire to motivate Mini-split SUSY using minimality [177].




Fig. 17: Branching ratios for the gluino decay channels $\tilde{g} \rightarrow \tilde{\chi}g$ (dashed lines), $\tilde{g} \rightarrow \tilde{\chi}q\bar{q}$ (dotted) and $\tilde{g} \rightarrow \tilde{\chi}^{\pm}q\bar{q}'$ (dot-dashed), summed over all possible neutralino or chargino states, as a function of $m_{\tilde{g}}$ and with $m_{\tilde{g}} = 1$ TeV. This figure is taken from [181].

example, motivated by the renormalization group evolution of the squark masses, and taking the gaugino mass ordering to be $m_{\tilde{g}} > m_{\tilde{W}} > m_{\tilde{B}}$, the decay $\tilde{g} \rightarrow t\bar{t}t\bar{t}h\tilde{\chi}\tilde{\chi}$ could dominate [177].

Clearly the different channels involve search strategies for prompt searches at the LHC, which is not the focus of the study here. In order to estimate the reach for MATHUSLA, one must account for the interactions of the long-lived gluinos with matter, both to estimate the fraction of R -hadrons that are neutral, along with the energy one would expect them to have when they reach the MATHUSLA detector – for a discussion of these issues see Sec. 3. Then as long as the gluino decay products are energetic enough to light up the MATHUSLA scintillators, then there should be a detectable signal regardless of final state. The reach for a long-lived gluino is given in Fig. 14 and is in excess of 2 TeV, providing a complementary discovery channel to the HL-LHC main detectors. It is also worth noting that if we are in the exciting situation where a long-lived gluino has been discovered at ATLAS and CMS, MATHUSLA would also provide valuable information about the properties of R -hadron interactions with matter.

Finally, we note that there is a region of parameter space where the electroweak gauginos could be long-lived enough to be detectable at MATHUSLA [182, 183]. This can be additionally motivated by an interest in testing the bino-wino coannihilation region [182, 183]. The candidate long-lived state is a nearly pure neutral wino, which dominantly decays to the bino and either a photon or an off-shell Higgs, see Fig. 18 for a schematic of the relevant decay modes. For the version of Mini-split SUSY of interest here, where the Higgsino mass and scalar masses are all set by gravity mediated physics, a lifetime relevant for MATHUSLA can be achieved for $\mu \sim 100$ TeV [183]. Consistency with the measured Higgs boson mass then requires $\tan\beta \sim 1$. Exploring the detailed implications for the parameter space that could be probed at MATHUSLA, along with correlating these projections with the Higgs boson mass and relic density is an interesting subject that deserves further study. However, we expect MATHUSLA to significantly extend the sensitivity of the main detectors for these scenarios, for identical reasons as for the Higgsino LLPs discussed in Sections. 4.1.2 and 4.1.5.

4.1.4 Stealth Supersymmetry¹³

Stealth Supersymmetry is a scenario in which collider signals of supersymmetry involve very little missing momentum. It is thus less constrained than more typical models of supersymmetry, given stringent

¹³Matthew Reece and David Pinner




Fig. 18: An illustration of the possible decay patterns for long-lived electroweak gauginos in mini-Split SUSY scenarios. This figure is taken from [183]. The very long-lived particle visible at MATHUSLA would be \tilde{W}^0 for $\mu \gtrsim 100$ TeV.

LHC bounds on the rate of events with large missing momentum. Stealth SUSY phenomenology relies on the existence of particles that are nearly degenerate with their superpartners, naturally squeezing the phase space available for invisible particles [184]. The key decay step in any Stealth Supersymmetry model has the form

$$\tilde{X} \rightarrow X + \tilde{N}, \quad X \rightarrow \text{SM particles}. \quad (32)$$

Here \tilde{X} and \tilde{N} are R -parity odd, and $m_{\tilde{N}} \ll m_{\tilde{X}} \approx m_X$. The particle \tilde{N} is neutral and escapes the detector, carrying away missing momentum, but due to the kinematics of the decay it is necessarily soft (even in the presence of hard initial-state radiation). In order for the stealth mechanism to be operative, this decay must occur inside the detector; otherwise, \tilde{X} contributes a large missing momentum. In the simplest realization of Stealth SUSY, the neutral particle \tilde{N} is the gravitino \tilde{G} . However, this scenario is under some strain from searches for displaced decays at colliders: if $m_{\tilde{X}} \approx 100$ GeV, $\delta m = m_{\tilde{X}} - m_X \approx 10$ GeV, and $\sqrt{F_0} \approx 100$ TeV, the decay length of $\tilde{X} \rightarrow X\tilde{G}$ is about 8 cm [184]. Such macroscopic lifetimes will be in tension with a range of inclusive searches for long-lived particles that have already been carried out at ATLAS and CMS [131–134, 139]. Although these searches have not been specifically recast as constraints on Stealth SUSY, they are known to strongly constrain a wide variety of R -parity violating SUSY scenarios with broadly similar kinematics [149, 150, 185]. This disfavors the Stealth SUSY scenario with $\tilde{X} \rightarrow X\tilde{G}$ decay. Although we could consider smaller $\sqrt{F_0} \approx 10$ TeV, it is challenging to build a model of SUSY breaking that operates at such a low scale. For instance, two recent attempts to build models of very low-scale SUSY breaking achieve $m_{3/2} \sim 1$ eV, with $\sqrt{F_0} \gtrsim 65$ TeV [186, 187]. On the other hand, values of $\sqrt{F_0}$ below ~ 140 TeV are favored by cosmological constraints from CMB lensing and cosmic shear [188]. Thus the preferred range of $\sqrt{F_0}$ is around, but perhaps modestly smaller than, 100 TeV.

The theoretically simplest alternative decay is the case where \tilde{N} is an *axino* \tilde{a} [189]. This is also known as a Goldstone fermion (not to be confused with a goldstino): that is, \tilde{a} is the supersymmetric partner of a (pseudo-)Goldstone boson from a symmetry that is broken in an approximately supersymmetric

manner. In this case, \tilde{a} can enjoy derivative couplings to \tilde{X} , X , allowing for the desired decay $\tilde{X} \rightarrow X\tilde{a}$; for some range of symmetry-breaking scales, this decay is prompt enough on collider timescales to evade the ATLAS and CMS searches for displaced vertices. In that case, the bounds on natural realizations of Stealth SUSY can still be significantly weaker than bounds on standard SUSY scenarios [190].




Fig. 19: A full Stealth SUSY decay chain. It is crucial that the $\tilde{X} \rightarrow X\tilde{a}$ decay step happen promptly inside the LHC, to evade missing momentum and displaced vertex searches. Significantly later, the axino decay can produce a saxion which in turn can produce visible Standard Model particles inside MATHUSLA.

Interestingly, the axino scenario naturally predicts a *secondary* displaced decay, as depicted in Fig. 19, which may be visible to MATHUSLA. The axino is contained in a supermultiplet with two real scalar fields, s (the saxion) and a (the axion). It is possible to parametrize the interactions of this supermultiplet with a simple effective field theory [192, 193], from which one learns that s and \tilde{a} will generically obtain a mass on the order of the gravitino mass or larger, although a may remain lighter. The axino mass allows for a later decay

$$\tilde{a} \rightarrow a + \tilde{G} \text{ or } s + \tilde{G}. \quad (33)$$

In turn, a or s may decay to Standard Model particles.

Let us fill in some estimates of the relevant regimes of parameter space. The decay to the axino must be prompt. We can consider an effective theory in which the stealth field X couples to a chiral Goldstone superfield A with a shift symmetry,

$$K \supset \frac{1}{f}(A + A^\dagger)X^\dagger X. \quad (34)$$

In the approximation of negligible axino mass and a small stealth mass splitting $\delta m_X \equiv m_{\tilde{X}} - m_X$, this leads to a decay width (correcting a formula in [189] by a factor of 4):

$$\Gamma(\tilde{X} \rightarrow X\tilde{a}) \approx \frac{m_{\tilde{X}}\delta m_X^2}{\pi f^2} \approx \frac{1}{6 \mu\text{m}} \left(\frac{m_{\tilde{X}}}{100 \text{ GeV}} \right) \left(\frac{\delta m_X}{10 \text{ GeV}} \right)^2 \left(\frac{10^7 \text{ GeV}}{f} \right)^2. \quad (35)$$

Collider searches for displaced vertices will not be sensitive to this decay provided the decay length $bc\tau \lesssim 100 \mu\text{m}$. We require $\delta m_X \lesssim 10 \text{ GeV}$ for stealth phenomenology; furthermore, much larger splittings or much smaller values of f are potentially associated with sizable tadpole effects for the singlet X that could induce direct (non-stealthy) decays to axinos. Because of this target for δm_X , we require that the axino, saxion, and axion masses are all below 10 GeV.




Fig. 20: MATHUSLA discovery potential for Stealth SUSY. The signal arises in gluino pair production events (gluino mass is on the vertical axis), and the observed decay is a light saxion (mass on the horizontal axis) decaying to two gluons or photons. Blue curves are the reach (corresponding to 4 detected events) for various values of c_g and c_γ defined in equation (38). As the saxion mass approaches the axino mass (fixed here to 5 GeV), the axino begins to decay invisibly to an axion and a gravitino. Conversely, as the saxion mass falls below the two-pion threshold, the visible branching fraction of the saxion quickly becomes negligible, with the decay proceeding predominantly to a pair of axions. The red shaded regions show the current LHC exclusions for gluinos decaying to singlinos from the ATLAS RPV search [191]. The projected discovery reach at HL-LHC with 3 ab^{-1} is shown in green. The dashed green line and lighter red region are estimated discovery and exclusion curves, respectively, with the solid green line and darker red region corresponding to an efficiency reduced by 50%, shown here as a rough guide to possible systematic error in recasting. Note that even in the event of a positive signal at the HL-LHC, MATHUSLA would be required to discover the long-lived axino produced in the event. This estimate assumes the $200 \text{ m} \times 200 \text{ m} \times 20 \text{ m}$ benchmark geometry of Fig. 1.

The axino decay to gravitino has width

$$\Gamma(\tilde{a} \rightarrow (s, a) + \tilde{G}) = \frac{m_{\tilde{a}}^5}{32\pi F_0^2} \left(1 - \frac{m_{s,a}^2}{m_{\tilde{a}}^2}\right)^4 \approx \frac{1}{1 \text{ km}} \left(\frac{m_{\tilde{a}}}{4 \text{ GeV}}\right)^5 \left(\frac{100 \text{ TeV}}{\sqrt{F_0}}\right)^4, \quad (36)$$

where in the last step we have dropped the phase space factor and summed the two partial widths. We see that the requirement for the stealth decay to be kinematically accessible, $m_{\tilde{a}} < \delta m_X$, already pushes the lifetime of the axino decay to gravitino to order kilometers even when the SUSY breaking scale is low. This is in the right range for sensitivity at MATHUSLA. Much larger values of $\sqrt{F_0}$ will have decay lengths that are substantially longer, but the lowest possible values of $\sqrt{F_0}$ are favored for naturalness.

The final decay is that of the saxion or axion itself. In general, the axion may be much lighter than

the saxion. The saxion always has a decay width to axions arising from the kinetic term,

$$\Gamma(s \rightarrow aa) = \frac{m_s^3}{32\pi f_a^2} \approx \frac{1}{3 \text{ cm}} \left(\frac{m_s}{4 \text{ GeV}} \right)^3 \left(\frac{10^7 \text{ GeV}}{f_a} \right)^2, \quad (37)$$

so once the axino has decayed to a saxion and gravitino, the saxion will then decay relatively promptly. Of course, the $s \rightarrow aa$ decay will not be observable. But the saxion can additionally couple to Standard Model fields through couplings like

$$\frac{c_g \alpha_s}{\pi f} s G_{\mu\nu}^a G^{a\mu\nu} + \frac{c_\gamma \alpha}{\pi f} s F_{\mu\nu} F^{\mu\nu}. \quad (38)$$

Provided that the effective coupling scale $f/c_g, f/c_\gamma$ appearing in these couplings is not much larger than f_a , a fraction of the decays will be to gluons (producing hadrons in the MATHUSLA detector) or to photons (detectable through conversions).

The MATHUSLA reach for saxion decays in Stealth SUSY is depicted in Figure 20. These events originate with gluino pair production in the LHC, with the decay chain:

$$pp \rightarrow \tilde{g}\tilde{g}, \quad \tilde{g} \rightarrow g\tilde{X}, \quad \tilde{X} \rightarrow X\tilde{a}, \quad X \rightarrow gg, \quad \tilde{a} \rightarrow s\tilde{G}, \quad s \rightarrow gg. \quad (39)$$

In this figure we assume low-scale SUSY breaking with $\sqrt{F_0} = 50 \text{ TeV}$, along with $m_{\tilde{X}} = 100 \text{ GeV}$, $m_X = 90 \text{ GeV}$, an axino mass of 5 GeV and an approximately massless axion, corresponding to a rest-frame axino lifetime of approximately 20 m . Combined with an average boost factor of $\mathcal{O}(10)$ from the gluino production and decay, these choices approximately maximize the MATHUSLA sensitivity to Stealth SUSY. Nevertheless, these parameters fall within reasonable ranges for Stealth SUSY, given the requirements discussed above for prompt singlet decays in the collider, minimal MET, the naturalness pressure towards a low SUSY breaking scale, and cosmological constraints favoring a very light gravitino. On the other hand, because the axino lifetime is highly sensitive to $\sqrt{F_0}$ and $m_{\tilde{a}}$, much of the parameter space predicts a longer lifetime and a low acceptance for MATHUSLA. As an illustration, raising $\sqrt{F_0}$ to 100 TeV reduces the maximum MATHUSLA reach by approximately 250 GeV in the gluino mass, taking $c_g = c_\gamma = 1$.

Figure 20 also shows the current LHC exclusion and the HL-LHC discovery reach for this scenario. The gluino pair production events contain six energetic gluons (plus ISR or FSR) and very little missing energy. They thus have some similarity to R-parity violating decays in the MSSM, which can produce six quark jets from gluino decays. Thus we have relied on an ATLAS RPV search [191] to estimate the LHC's capabilities. Validating our recasting against the reported ATLAS RPV signal efficiency suggests that the solid red region, in which we have reduced our calculated Monte Carlo efficiency by a factor of 2, is approximately accurate. The lighter shaded red region is the unscaled result, which we expect overstates the current exclusion. To estimate the HL-LHC discovery reach, we have first computed the mass for which $\epsilon \times \sigma$ is $5/2$ the expected 95% CL_s exclusion at the current 13 TeV search, as a rough estimate for the gluino mass which could already have been discovered at 5σ confidence. We have then rescaled to the 14 TeV LHC with 3 ab^{-1} using the Collider Reach estimation procedure of Salam and Weiler [194], which has been argued to produce approximately accurate results based on simple considerations of parton luminosities.

From the figure, it is apparent that the HL-LHC has good discovery potential over a wide range of gluino masses. MATHUSLA would discover an additional displaced decay over a smaller region of parameter space. On the other hand, the LHC signal would be in a background-dominated region and could prove hard to interpret. For example, both R-parity violating SUSY and Stealth SUSY could produce the same several-jet final state. MATHUSLA could be crucial to avoid misdiagnosing the nature of the new physics, demonstrating that R-parity violation is not the origin of the physics and that a model containing a low mass, long-lived particle is necessary. Furthermore, even a handful of events at MATHUSLA would be sufficient to provide this information. The axino lifetime is long enough, and its

mass low enough, that MATHUSLA is much better suited than LLP searches at the HL-LHC to detecting it (see §3.2.3).

Notice that a decay chain along the lines we have discussed,

$$\tilde{\chi} \rightarrow \text{SM} + \tilde{a}, \quad \tilde{a} \rightarrow s + \tilde{G}, \quad s \rightarrow \text{SM}, \quad (40)$$

can happen in low-scale SUSY breaking scenarios completely independently of Stealth SUSY, and a MATHUSLA signal from a late-decaying saxion-like particle is even more generic. What Stealth SUSY adds is a strong motivation for having an axino \tilde{a} available for this decay chain to proceed. Precisely the challenge of obtaining a stealthy decay that is prompt on collider timescales leads to the added ingredient that gives us an auxiliary signal of a potentially very long lifetime.

4.1.5 Axinos¹⁴

Another concern for naturalness resides in the QCD sector, which is called the strong CP problem [195]. The gauge invariance does not forbid a CP-odd term:

$$\mathcal{L}_\theta = \theta \frac{g_s^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}_a^{\mu\nu} \quad (41)$$

where θ is a dimensionless parameter which is generically of order one. However, the actual value of θ is strongly constrained by measurements of nucleon electric dipole moments: $|\theta| < 10^{-10}$. The problem of such a (vanishingly) small θ is elegantly resolved by introducing an axion which is a Goldstone boson of a QCD-anomalous Peccei-Quinn (PQ) symmetry [196–198]. The PQ symmetry can be realized typically by introducing heavy quarks (KSVZ) [199, 200] or by extending the Higgs sector (DFSZ) [201, 202]. The PQ symmetry is supposed to be broken spontaneously by PQ-charged (dominantly SM-singlet) fields $\phi = v_\phi e^{ia_\phi/v_\phi}/\sqrt{2}$ carrying the PQ charge x_ϕ and thus the axion is a linear combination of the phase degrees of freedom a_ϕ : $a \equiv \sum x_\phi v_\phi a_\phi / v_{PQ}$ where the overall breaking scale is given by $v_{PQ} \equiv \sqrt{\sum_\phi x_\phi^2 v_\phi^2}$. Due to the QCD anomaly of the PQ symmetry, there arises an axion-gluon-gluon coupling

$$\mathcal{L}_{agg} = \frac{g_s^2}{32\pi^2} \frac{a}{f_a} G_{\mu\nu}^a \tilde{G}_a^{\mu\nu} \quad (42)$$

where $f_a \equiv v_{PQ}/N_{\text{DW}}$ and N_{DW} is the domain wall number counting the QCD anomaly. Note that the θ term can be absorbed into the dynamic degree a whose potential is generated after the QCD phase transition:

$$V_{\text{QCD}}[a] \approx m_\pi^2 f_\pi^2 \cos\left(\frac{a}{f_a}\right). \quad (43)$$

This settles the effective θ term to zero: $\theta_{eff} \equiv \langle a \rangle / f_a = 0$, and induces a non-vanishing axion mass $m_a \approx m_\pi f_\pi / f_a$.

The conventionally allowed window of the axion scale is $10^9 \lesssim f_a/\text{GeV} \lesssim 10^{12}$. The lower limit comes from star cooling processes [203, 204] and the upper limit from the axion cold dark matter contribution taking the initial mis-alignment angle θ_i of order one (see e.g. [205]). One may allow higher f_a if initial $\theta_i \ll 1$ is taken depending on cosmological scenarios.

The existence of such a high scale causes quadratic divergences to the Higgs boson mass and thus requires a huge fine-tuning to keep stable two scales, the electroweak scale and the axion scale (or a generic UV scale). Supersymmetry (SUSY) would be the best-known framework to avoid such a hierarchy problem. However, the electroweak symmetry breaking in SUSY suffers from a certain degree of fine-tuning to maintain a desirable potential minimization condition:

$$\frac{m_Z^2}{2} = \frac{m_{H_d}^2 - m_{H_u}^2 \tan^2 \beta}{\tan^2 \beta - 1} - \mu^2 \quad (44)$$

¹⁴Eung Jin Chun, Sunghoon Jung

where $m_{H_{u,d}}$ are the soft masses of the two Higgs doublets, $\tan\beta \equiv v_u/v_d$ is the ratio of their vacuum expectation values, and μ is the Higgs bilinear parameter in the superpotential. As LHC finds no hint of SUSY, it pushes up the soft mass scale above TeV range, the minimization condition (44) requires a fine cancellation among different terms. Barring too huge a cancellation, one may arrange $m_{H_{u,d}}$ and μ not too larger than m_Z . This has been advocated as “natural SUSY” [206–209] implying stops/sbottoms at sub-TeV and light higgsinos with $\mu \sim \mathcal{O}(100)$ GeV.

The origin of μ at the electroweak scale may be related to the PQ symmetry in the manner of DFSZ which introduces a non-renormalizable superpotential in the Higgs sector [210, 211]:

$$W = \lambda_\mu \frac{P^2}{M_P} H_u H_d \quad (45)$$

where P and thus $H_u H_d$ carries a non-trivial PQ charge and M_P is the reduced Planck mass. Upon the PQ-symmetry-breaking $v_{PQ} \sim \langle P \rangle$, a μ term is generated by $\mu = \lambda_\mu \langle P \rangle^2 / M_P$. Once PQ symmetry is broken, there appear the axion a , its scalar partner, the saxion s , and the fermion super-partner, the axino \tilde{a} . Forming an axion superfield $A = (s + ia, \tilde{a})$, one can schematically write down the effective μ -term superpotential;

$$W = \mu H_u H_d + c_H \frac{\mu}{v_{PQ}} A H_u H_d \quad (46)$$

where c_H is a parameter depending on the PQ symmetry breaking sector; we take $c_H = 2$ in this work. The axino mass is expected to be of order of the soft SUSY breaking scale, but it is in general model-dependent [212–214].

Although axino interactions are suppressed by the axion scale, cosmic axinos can be abundantly produced through thermal particle interactions (see Section 5.3 on Freeze-In Scenarios) like decays, inverse-decays and scattering of higgsinos and Higgs bosons into the axino:

$$\frac{\rho_{\tilde{a}}}{\rho_{DM}} \approx \xi_H \left(\frac{m_{\tilde{a}}}{\text{MeV}} \right) \left(\frac{\mu}{300 \text{ GeV}} \right)^r \left(\frac{10^{12} \text{ GeV}}{v_{PQ}} \right)^2 \quad (47)$$

ξ_H is a model-dependent parameter involving the soft SUSY breaking parameters. $r = 1$ if Higgsino decay to Higgs and axino is allowed only after the electroweak phase transition [215–217], and $r = 2$ if it is allowed before EWSB as well, see Section 5.3.3.1. Thus a stable axino has to be lighter than about MeV for $v_{PQ} \lesssim 10^{12}$ GeV, unless mechanisms to dilute its relic abundance are present. Such dilution could be provided, for example, by a decaying saxion condensate, see Section 5.3.3. We therefore consider axino masses above an MeV as well, though the precise axino mass will not greatly affect the LLP phenomenology we study.

Motivated by natural SUSY realizing the DFSZ axion, one could therefore have a Higgsino NLSP and an axino LSP [218]. Light higgsinos can be copiously produced at the LHC via Drell-Yan production (and even greater rates are possible if heavy colored particles decay to higgsinos, see Fig. 8). They decay to axinos plus the EW bosons h or Z through the coupling in Eq. (46):

$$\tilde{H}^0 \rightarrow \tilde{a}Z, \tilde{a}h \rightarrow \text{displaced dilepton/dijet} + \text{MET} \quad (48)$$

The typical decay rate of the Higgsino NLSP is estimated as

$$\Gamma_{\tilde{H}^0} \approx \frac{c_H^2}{32\pi} \frac{\mu^3}{v_{PQ}^2} \quad (49)$$

corresponding to

$$c\tau_{\tilde{H}^0} \approx \frac{620\text{m}}{(c_H/2)^2} \left(\frac{200\text{GeV}}{\mu} \right)^3 \left(\frac{v_{PQ}}{10^{12}\text{GeV}} \right)^2, \quad (50)$$




Fig. 21: Number of long-lived Higgsino decays $\tilde{H} \rightarrow \tilde{a}Z, \tilde{a}h$ observed in MATHUSLA for the DFSZ axino model, assuming only Drell-Yan like Higgsino production at the HL-LHC. A perfect LLP detection efficiency $\epsilon_{LLP} = 1$ is assumed. Dependence $\tan\beta$ and axino mass (except near kinematic threshold for the decay) is weak. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

where we assumed a massless axino (but the precise mass does not change this drastically, as long as the final-state EW boson is on-shell). Especially for PQ scales near the upper end of the motivated range, it is clear that higgsinos have the right lifetime to be detected by MATHUSLA.

We show the expected number of observed Higgsino decays in MATHUSLA at the HL-LHC in Fig. 21. MATHUSLA can probe a wide range of PQ scales and higgsino masses, $\mu \sim 100 - 1000$ GeV and $v_{PQ} \sim 10^{11} - 10^{14}$ GeV.

How does the MATHUSLA reach compare to the achievable reach of the HL-LHC main detectors in the long lifetime limit? For Higgsino masses below a few hundred GeV, or heavier higgsinos with axinos that are comparable in mass, the amount of visible hadronic energy per LLP decay is likely too low for background-free searches using the LLP decay along, see Section 3.2.2. Therefore, while MATHUSLA is sensitive to any visible final state of the long-lived higgsino decay, the most sensitive HL-LHC search would make use of the leptonic Z decay to suppress hadronic backgrounds and obtain a clean sample for offline LLP reconstruction. The signal for the main detector search is therefore suppressed, relative to MATHUSLA, by the branching ratio

$$\text{Br}_{\text{LHC}} = \text{Br}(\tilde{H}^0 \rightarrow \tilde{a}Z) \text{Br}(Z \rightarrow \ell^+\ell^-) \quad (51)$$

For high $\tan\beta = 50$, $\text{Br}_{\text{LHC}} \sim 0.05$ for either sign of μ , while for low $\tan\beta = 2$, $\text{Br}_{\text{LHC}} \sim 0.08$ ($\mu > 0$) or 0.02 ($\mu < 0$). As a result, the factor ϵ_{cut} in Sections 3.2.2 and 3.2.3 is $\text{Br}_{\text{LHC}} \times$ (various trigger and cut efficiencies) even without main detector backgrounds. For heavier higgsinos with more than a few hundred GeV of hadronic energy in their decay, displaced jet analyses at the main detectors may be able to achieve sensitivities closer to those of MATHUSLA.

MATHUSLA significantly extends our axino sensitivity to very high v_{PQ} scales. Given that the main detectors would have good sensitivity at shorter lifetimes, MATHUSLA would enable us to probe the entire v_{PQ} range motivated by axion DM (produced via mis-alignment) or axino DM (produced via freeze-in, see Sec. 5.3.3). Furthermore, as Higgsino NLSPs are produced in almost any cascade decays of heavier supersymmetric particles, the signal can be enhanced and probe heavier supersymmetric particles too.

4.1.6 Sgoldstinos ¹⁵

There is a lot of supersymmetric extensions of the Standard Model differing from each other in many respects, but exhibiting a common feature: in each model the supersymmetry must be spontaneously broken. The majority of models exploit for this purpose a chiral superfield,

$$\Phi = \varphi + \sqrt{2}\theta\psi + F_\varphi\theta\theta,$$

whose auxiliary component acquires non-zero vacuum expectation value (v.e.v.)

$$\langle F_\varphi \rangle \equiv F \neq 0,$$

that breaks the supersymmetry. In accordance with the Goldstone theorem, there is a massless particle in the spectrum, which is fermion in case of supersymmetry, *goldstino* ψ . Its superpartners, φ and φ^* form scalar and pseudoscalar *sgoldstinos*,

$$\frac{1}{\sqrt{2}}(\varphi + \varphi^*) \equiv S, \quad \frac{1}{i\sqrt{2}}(\varphi - \varphi^*) \equiv P,$$

respectively. As we discuss below, these sgoldstinos can be long-lived and are therefore a natural target for searches at MATHUSLA. They are an very challenging signal, but production in B -meson decays can give rise to LLP decays that MATHUSLA can detect, in a complementary manner to searches at SHiP [219].

4.1.6.1 Sgoldstino couplings and lifetime

The auxiliary field F_φ has a dimension of mass squared, and its v.e.v. is of the order of the squared energy scale of the supersymmetry breaking,

$$F \sim (E_{\text{SUSY-breaking-scale}})^2.$$

Goldstino couples to the non-conserved (super)current [220] as follows from the general Goldberger–Treiman formula,

$$\mathcal{L}_\psi \propto \frac{1}{F} J_{\text{SUSY}}^\mu \partial_\mu \psi. \quad (52)$$

When supersymmetry is promoted to the local symmetry, it becomes supergravity and goldstino gets eaten in the super-Higgs mechanism, giving mass to gravitino,

$$m_{\tilde{G}} = \sqrt{\frac{8\pi}{3}} \frac{F}{M_{\text{Planck}}},$$

and forming its longitudinal component.

Sgoldstinos remain massless at tree level as well, but gain masses due to higher order corrections. Their scale is very model-dependent and for phenomenological purposes sgoldstino masses m_S and m_P can be considered as free parameters. In particular, sgoldstinos are naturally expected to be light in models with no-scale supergravity [221] and models with gauge mediated supersymmetry breaking [152, 222]. Meanwhile, sgoldstino couplings to the Standard Model (SM) fields are fixed by supersymmetry and in most cases the coupling constants are proportional to the ratio of some supersymmetry breaking parameters (soft masses, trilinear couplings) and supersymmetry breaking parameter F [220]. The explicit expression may be obtained by either performing the supersymmetry transformation of goldstino interaction (52) or exploiting the spurion technique [223]. Then, to the leading order in $1/F$

¹⁵Dmitry Gorbunov

sgoldstino couplings to gravitino \tilde{G} , leptons f_L , up- and down-quarks f_U , f_D , photons $F_{\mu\nu}$ and gluons $G_{\mu\nu}^a$ read [224, 225]

$$\mathcal{L}_{eff} = -\frac{1}{2\sqrt{2}F} \left(m_S^2 S \bar{\tilde{G}} \tilde{G} + i m_P^2 P \bar{\tilde{G}} \gamma_5 \tilde{G} \right) - \frac{1}{2\sqrt{2}} \frac{M_{\gamma\gamma}}{F} S F^{\mu\nu} F_{\mu\nu} + \frac{1}{4\sqrt{2}} \frac{M_{\gamma\gamma}}{F} P \epsilon^{\mu\nu\rho\sigma} F_{\mu\nu} F_{\rho\sigma} \quad (53)$$

$$- \frac{1}{2\sqrt{2}} \frac{M_3}{F} S G^{\mu\nu a} G_{\mu\nu}^a + \frac{1}{4\sqrt{2}} \frac{M_3}{F} P \epsilon^{\mu\nu\rho\sigma} G_{\mu\nu}^a G_{\rho\sigma}^a \quad (54)$$

$$- \frac{\tilde{m}_{D_{ij}}^{LR 2}}{\sqrt{2}F} S \bar{f}_{D_i} f_{D_j} - i \frac{\tilde{m}_{D_{ij}}^{LR 2}}{\sqrt{2}F} P \bar{f}_{D_i} \gamma_5 f_{D_j} - \frac{\tilde{m}_{U_{ij}}^{LR 2}}{\sqrt{2}F} S \bar{f}_{U_i} f_{U_j} - i \frac{\tilde{m}_{U_{ij}}^{LR 2}}{\sqrt{2}F} P \bar{f}_{U_i} \gamma_5 f_{U_j} \quad (55)$$

$$- \frac{\tilde{m}_{L_{ij}}^{LR 2}}{\sqrt{2}F} S \bar{f}_{L_i} f_{L_j} - i \frac{\tilde{m}_{L_{ij}}^{LR 2}}{\sqrt{2}F} P \bar{f}_{L_i} \gamma_5 f_{L_j} . \quad (56)$$

Hereafter θ_W is the weak mixing angle, $M_{\gamma\gamma} = M_1 \sin^2 \theta_W + M_2 \cos^2 \theta_W$ and M_i , $i = 1, 2, 3$ are gaugino masses with index corresponding to SM gauge groups, $U(1)_Y$, $SU(2)_W$ and $SU(3)_c$, and $\tilde{m}_{D_{ij}}^{LR 2}$ and $\tilde{m}_{U_{ij}}^{LR 2}$ are left-right down- and up-squark mass terms. Above we omit two-sgoldstino coupling terms, see Refs. [224, 226–228] which are strongly suppressed by $1/F^2$.

These interactions determine sgoldstino production and decay rates. Their palletes depend on the patterns of MSSM soft terms. Without any specific hierarchy there, most naturally coupling to gluons dominate. Generically, sgoldstino couplings become weaker with increase of supersymmetry breaking scale, $\sim \sqrt{F}$, and ratio of superpartner mass and supersymmetry breaking scales, M_{soft}/\sqrt{F} , which must not exceed unity while the unitarity is conserved. The most attractive feature of sgoldstino phenomenology is that the measurement of sgoldstino couplings gives an opportunity to probe the scale of supersymmetry breaking in the whole theory. Sgoldstinos are R -even, contrary to R -odd gravitino, and so if SM superpartners are heavy (an assumption consistent with LHC results), their production is less suppressed as compared to (light) gravitinos, which couplings start at $\mathcal{O}(1/F^2)$ level only. Sgoldstino production at LHC with couplings (53)–(56) has been studied in Refs. [229–232].

For interesting sgoldstino masses above 1 GeV, sgoldstino decays into photons and into gluons are always open. An order-of-magnitude estimate of sgoldstino life-time,

$$\tau_{S(P)} \sim 4\pi \frac{F^2}{M_{soft}^2 m_{S(P)}^3} \sim 10^{-8} \times \left(\frac{\sqrt{F}/1000 \text{ TeV}}{M_{soft}/1 \text{ TeV}} \right)^2 \left(\frac{\sqrt{F}}{1000 \text{ TeV}} \right)^2 \left(\frac{10 \text{ GeV}}{m_{S(P)}} \right)^3 \text{ s}, \quad (57)$$

shows that in order to allow sgoldstino to cover the distance of a hundred meters and reach the MATH-USLA detector, supersymmetry breaking scale must be high and well-separated from the scale of MSSM soft terms. However, in this limit the sgoldstino coupling constants to SM fields become tiny. As we show below, its direct production rate is much below the critical 1 fb scale. However, sizable production rates are possible if its mass lies below the B -threshold.

4.1.6.2 Sgoldstino production mechanisms

Sgoldstinos can be produced in gluon fusion, the same way as the SM Higgs boson h . Hence, the sgoldstino production cross-section at LHC can be estimated as

$$\sigma_{gg \rightarrow S(P)}(m_{S(P)}) = \sigma_{gg \rightarrow h}(m_{S(P)}) \frac{\Gamma_{S(P) \rightarrow gg}(m_{S(P)})}{\Gamma_{h \rightarrow gg}(m_{S(P)})} = \sigma_{gg \rightarrow h}(m_{S(P)}) \left(\frac{3\pi}{\alpha_s(m_{S(P)})} \right)^2 \frac{\sqrt{2} M_3^2}{F^2 G_F} \quad (58)$$

where $\sigma_{gg \rightarrow h}(M)$ and $\Gamma_{h \rightarrow gg}(M)$ are production and decay width of the SM Higgs boson into gluons, if its mass would be M . Putting numbers into (57) one finds

$$\sigma_{gg \rightarrow S(P)}(m_{S(P)}) \simeq 0.1 \text{ fb} \times \left(\frac{\sigma_{gg \rightarrow h}(m_{S(P)})}{10 \text{ pb}} \right) \left(\frac{M_3}{1 \text{ TeV}} \right)^2 \left(\frac{100 \text{ TeV}}{\sqrt{F}} \right)^4. \quad (59)$$

It is clear from eqs. (57) and (59) that it is impossible to have both reasonably high production rate and long-lived sgoldstino. The product of eqs. (57) and (59) is

$$\sigma_{gg \rightarrow S(P)}(m_{S(P)}) \times \tau_{S(P)} \simeq 10^{-13} \text{ fb s} \times \left(\frac{10 \text{ GeV}}{m_{S(P)}} \right)^3 \left(\frac{\sigma_{gg \rightarrow h}(m_{S(P)})}{10 \text{ pb}} \right).$$

Sgoldstino can be also produced in decays of heavy SM particles emerged in proton-proton collisions at LHC: top-quark, Z -boson and Higgs boson are potential sources. The sgoldstino interaction with t -quark is governed by unknown flavor-violating structures of the squark squared mass matrix, see Ref. [229] for details, and top-quark can decay into sgoldstino and light quark with branching ratios as high as if $\sqrt{F} \lesssim 10 \text{ TeV}$, which implies short-lived sgoldstinos untestable at MATHUSLA, see eq. (57). Z -bosons are much more abundant at LHC than t -quarks, and can decay into sgoldstino and photon due to coupling [226]

$$\mathcal{L}_{Z\gamma} = -\frac{(M_2 - M_1) \cos \theta_W \sin \theta_W}{\sqrt{2} F} F_{\mu\nu} Z^{\mu\nu} S + \frac{(M_2 - M_1) \cos \theta_W \sin \theta_W}{2\sqrt{2} F} F_{\mu\nu} Z_{\lambda\rho} \epsilon^{\mu\nu\lambda\rho} P \quad (60)$$

with branching

$$\text{Br}(Z \rightarrow \gamma S(P)) \sim 10^{-7} \times \left(\frac{M_{soft}/1 \text{ TeV}}{\sqrt{F}/10 \text{ TeV}} \right)^2 \left(\frac{10 \text{ TeV}}{\sqrt{F}} \right)^2. \quad (61)$$

The SM Higgs boson can decay into a couple of sgoldstinos, if kinematically allowed. Sgoldstino interaction with the Higgs sector is considered in Refs. [232, 233], the relevant couplings are suppressed by $1/F^2$ (each sgoldstino leg brings factor $1/F$). Assuming all the MSSM massive parameters are of the same order M_{soft} we can estimate the branching as

$$\text{Br}(h \rightarrow SS) \sim 10^{-5} \times \left(\frac{M_{soft}/1 \text{ TeV}}{\sqrt{F}/10 \text{ TeV}} \right)^8. \quad (62)$$

In all the cases above with superpartner scale of order 1 TeV the sgoldstino is either too short-lived or too rare produced.

The best target for an ultra-long-lived particle search is the sgoldstino mass regime below B -threshold, $m_{S(P)} < 5 \text{ GeV}$, with sgoldstino production in beauty meson decays. In this case sgoldstino is much lighter, and so long-lived, and, for the production, the feeble sgoldstino couplings must compete not with strong but only with weak interactions. Thus, beauty mesons can decay into sgoldstino with branchings as large as 10^{-4} [234]. There may be contributions from quark flavor-conserving and flavor-violating sgoldstino couplings. In this case MATHUSLA will compete with SHiP experiment [219] operating on SPS 400 GeV proton beam in a beam-dump mode. The number of beauty-quarks at ATLAS/CMS is expected to exceed significantly that at the SHiP. Even though the SHiP geometry is optimized for the flux of outgoing particles, it has been shown that MATHUSLA has significantly better acceptance for LLP decays than SHiP, provided the lifetime is above $\sim 100 \text{ m}$ [18] (see also Section 8.4). The reach of SHiP has been investigated in detail in Ref. [234]. It was found, that depending on the MSSM soft term pattern and the scale of SM superpartners, SHiP can probe the scale of supersymmetry breaking as high as $\sqrt{F} \sim 10^2 - 10^4 \text{ TeV}$. Since the sgoldstino decay length is in the long-lifetime regime at the upper limit of this \sqrt{F} sensitivity range (and much shorter at the lower limit), we expect that MATHUSLA will be able to significantly extend the reach of SHiP and allow access to higher SUSY breaking scales.

4.2 Neutral Naturalness¹⁶

The discovery of a light, apparently elementary Higgs boson at the LHC has heightened the severity of the electroweak hierarchy problem, while increasingly severe bounds on new colored particles have begun to disfavor conventional solutions such as supersymmetry or compositeness. Models of neutral naturalness provide a compelling alternative, in which the lightest states protecting the weak scale are not colored (and, in some cases, entirely neutral under the Standard Model). Such protection of the weak scale is achieved primarily through discrete symmetries, rather than continuous symmetries. Realizations of neutral naturalness include the Twin Higgs [235], Orbifold Higgs [236], Quirky Little Higgs [237], and Folded Supersymmetry [238].

Addressing the hierarchy problem via discrete symmetries naturally leads to specific hidden valleys (Section 8.1) without Standard Model quantum numbers. Rather, the new states in these sectors primarily couple to the Standard Model through various portal-type interactions, including the Higgs portal $\lambda|H|^2\mathcal{O}$ and the neutrino portal $y(LH)\mathcal{O}$. If the discrete symmetry between the two sectors extends to hypercharge, the photon portal $\epsilon\tilde{F}_{\mu\nu}\hat{F}^{\mu\nu}$ may also mix the hidden and visible abelian gauge bosons. In contrast with generic hidden sectors, in models of neutral naturalness both the size of these portal interactions and the mass scale of hidden particles are typically dictated by naturalness considerations, providing a motivated range of rates and lifetimes that can be effectively probed by MATHUSLA.

4.2.1 LLP Signatures of the Hidden Sector

The discrete symmetries in successful models of neutral naturalness must be nearly exact in the top sector, given its relevance to the hierarchy problem, but may be approximate for states more remote from the Higgs sector [239]. In particular, since the QCD gauge coupling gives the dominant contribution to the renormalization of the top yukawa, the preservation of a near-exact discrete symmetry in the top sector requires the existence of one or more new QCD-like hidden gauge groups whose couplings are comparable to their Standard Model counterpart. Solution of the Hierarchy Problem via Neutral Naturalness therefore leads to the existence of specific confining Hidden Sectors. The confinement scales of these gauge groups are then typically of the same order as Λ_{QCD} , giving rise to hidden sector bound states whose masses range from $\mathcal{O}(1 - 100)$ GeV, allowing for their production at the LHC.

4.2.1.1 Production of Hidden Glueballs in Exotic Higgs Decays

The coupling between the Standard Model-like Higgs and the top partners required to address the hierarchy problem induces loop-level couplings to hidden gluons, much as the top quark generates the leading coupling of the Higgs to Standard Model gluons. The effective coupling between the SM-like Higgs h and the hidden gluon field strengths $\hat{G}_{\mu\nu}^a$ takes the general form

$$\mathcal{L} \supset \theta^2 \frac{\hat{\alpha}_3}{12\pi} \frac{h}{v} \hat{G}_{\mu\nu}^a \hat{G}_a^{\mu\nu} \quad (63)$$

where θ is a mixing angle that varies between different realizations of neutral naturalness. In Twin Higgs or Quirky Little Higgs models, $\theta^2 \simeq -v^2/f^2$, where f is a scale of spontaneous global symmetry breaking, while for Folded Supersymmetry $\theta^2 \simeq m_t^2/2\hat{m}_{\tilde{t}}^2$, where $\hat{m}_{\tilde{t}}$ is the mass of the QCD-neutral scalar top partners. Naturalness considerations bound the scales $f, \hat{m}_{\tilde{t}}$ to lie at or below the TeV scale, so that the relevant mixing angles are naturally $\mathcal{O}(0.1 - 1)$.

The coupling in Eq. (63) provides a predictive portal for the production of states in the hidden QCD sector, as well as an avenue for them to decay back to the Standard Model when kinematically allowed [239, 240]. In particular, for generic hidden sector masses and mixing angles, Eq. (63) predicts exotic Higgs decays into the hidden sector with branching ratios of order 0.01 – 1%, corresponding to rates on the order of 5-500 fb at the 14 TeV LHC.

¹⁶David Curtin, Nathaniel Craig, Yuhsin Tsai

Once produced, states in the hidden sector cascade down to the lightest accessible hidden sector state – typically a bound state of hidden QCD – which then can decay back to the Standard Model. The lifetime for these decays varies depending on the nature of the lightest bound states of hidden QCD. If the lightest hidden QCD states are mesons, or if even lighter hidden photons or neutrinos are part of the low-energy spectrum (as in the original Mirror Twin Higgs [235]), the hidden sector relics produced during the Big Bang are usually stable and abundant enough to be in conflict with ΔN_{eff} bounds from CMB measurements. We discuss this scenario in more detail below, but a straightforward way to avoid cosmological bounds is a hidden sector without such light states. This is always the case in Folded Supersymmetry, where the folded superpartners carry electroweak charge and have to be heavier than ~ 100 GeV to respect LEP limits. It may also be the case in Twin-Higgs-like models such as the Fraternal Twin Higgs [239]) where the discrete symmetry only applies to the third fermion generation, and there are no light mirror photons or neutrinos. In either scenario, the hidden QCD confines with zero light quark flavors, and the lightest accessible hidden sector states are hidden glueballs whose decay back to the Standard Model is governed by the dimension-5 operator in Eq. (63). Decays proceeding through this operator generically lead to lifetimes within the BBN limit.

In this case, the relevant process is the decay of a hidden glueball, typically the $J^{PC} = 0^{++} = G_0$ glueball, expected to lie at the bottom of the glueball spectrum with $m_{G_0} \approx 7\tilde{\Lambda}_{QCD}$, though decays of higher glueball excitations are also possible. G_0 decays to light Standard Model states via an off-shell Higgs: $0^{++} \rightarrow h^* \rightarrow YY$, where Y are SM fields. The amplitude for this process in terms of the glueball mass m_0 is [241]

$$\frac{\hat{\alpha}_3 \theta^2}{6\pi v^2} \langle YY | m_f \bar{f} f + m_Z^2 Z_\mu Z^\mu + 2m_W^2 W_\mu^+ W^{\mu-} | 0 \rangle \frac{1}{m_h^2 - m_0^2} \langle 0 | S | 0^{++} \rangle \quad (64)$$

which results in a width

$$\Gamma_{0^{++} \rightarrow YY} = \left(\frac{\hat{\alpha}_3 \theta^2}{6\pi(m_h^2 - m_0^2)} \frac{f_0}{v} \right)^2 \Gamma_{h \rightarrow YY}^{SM}(m_{0^+}^2) \quad (65)$$

where f_0 is the hidden 0^{++} decay constant. The corresponding mean decay length is

$$c\tau_0 \sim 20 \text{ m} \times \left(\frac{10 \text{ GeV}}{m_0} \right)^7 \left(\frac{f}{750 \text{ GeV}} \right)^4. \quad (66)$$

(A similar expression applies in FSUSY.) The decay width of these glueballs is a steep function of their mass and mixings. For typical values the mean decay length ranges from $10^{-6} - 10^7$ m, giving rise to the distinctive signal of exotic displaced decays where lifetime should be regarded as an almost free parameter within the BBN limit.

These displaced decays are sufficiently distinctive and occur with sufficient rate at the LHC that they may be distinguished from Standard Model backgrounds provided appropriate triggering strategies, which were explored in [240]. The projected reach is summarized in Fig. 22. The plot shows a simplified parameter space of lightest glueball mass m_{G_0} and top partner mass, either in Folded SUSY assuming no stop mixing (left vertical axis) or for a Fraternal Twin Higgs like model (right vertical axis). Glueball lifetime depends most dramatically on m_{G_0} , with proper decay lengths at the cm to sub-mm level for high glueball masses near $m_H/2$, and long lifetimes approaching the BBN limit for glueball masses below 10-20 GeV.

Uncertainties of hidden sector glueball hadronization make a precise prediction of glueball multiplicity in exotic Higgs decays difficult [242]. For the purpose of these estimates, we conservatively assume that only two glueballs are produced per Higgs decay. This is likely to be reasonably accurate at high glueball masses. At lower glueball masses, there many more glueballs, each of which has an accordingly lower boost than under our simple assumption. Since the long glueball lifetime in that regime is the




Fig. 22: Simplified Neutral Naturalness parameter space of lightest glueball mass m_{G_0} and top partner mass in Folded SUSY or the Fraternal Twin higgs. Shown is projected reach of HL-LHC LLP searches for glueballs produced in exotic Higgs decays [240] in the ATLAS Muon System (red) or in the tracker in association with VBF jets or leptons from Higgs production (blue, orange). The reach of MATHUSLA (assuming the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1) is shown in purple, and covers the regime of long-lived glueballs with masses $\lesssim 15$ GeV. Sensitivity in all searches is conservatively estimated by assuming two glueballs produced per Higgs decay, and dashed contours indicate uncertainties due to details of hidden sector hadronization. See text for additional details.

main bottleneck for searches with any detector, a more realistic treatment of hidden sector hadronization would only improve all reach projections. This makes our assumption suitable for a pessimistic estimate of the LHC's ability to probe Neutral Naturalness. Since the 0^{++} glueball is the state with the shortest lifetime, the fraction of glueballs that end up in the 0^{++} state is another important factor in estimating the LLP signal rate. However, the large ratio between the lightest glueball mass and the hidden QCD string tension suggests that 0^{++} states form a majority or at least a significant fraction of the produced states [243], based on modeling of hadronization processes as thermal emissions [244]. In Fig. 22 we therefore assume that the 0^{++} and other glueball fractions are given by spin-weighted Boltzmann factors for all kinematically available states. The dashed contours indicate the variation of reach estimates from varying that 0^{++} up or down by a factor of 2. Finally, vertical solid (dashed) lines show where the production rate of 0^{++} glueballs may be additionally enhanced or suppressed due to non-perturbative mixing effects [239].

This simplified model of Neutral Naturalness then produces the signal of LLP pair production in exotic Higgs decays, with subsequent LLP decay through the Higgs portal. (See also Sec. 8.2.) As explained in [240], there are three particularly promising search strategies using the HL-LHC main detec-

tors: (1) Search for two LLPs using the dedicated displaced decay trigger in the ATLAS Muon System, similar to [131] (red in Fig. 22). (2) Search for a single displaced vertex in the tracker, in association with VBF jets from Higgs production (blue). (3) Search for a single displaced vertex in the tracker, in association with a lepton from associated Higgs production (orange). In this case we optimistically assume a displaced vertex can be reconstructed less than a mm from the primary vertex to demonstrate how the short glueball lifetime regime may be probed. (Such a search would require significant further experimental study.) Each search requires either two observed LLP decays, or one LLP decay in association with a conspicuous prompt object like VBF jets or leptons. Search strategies along these lines are likely to have low or zero background. The sensitivity estimates for these three searches in Fig. 22 therefore show regions with at least 4 observed LLP decays passing the requirements of each analysis. Also shown for comparison is the $\text{Br}(h \rightarrow \text{invisible})$ bound achievable with a TLEP-like lepton collider [245] (green), which may provide clues for very production of very long-lived glueballs.¹⁷

It is clear that the HL-LHC has the capability to probe Neutral Naturalness at the TeV scale for hidden glueball masses above ~ 15 GeV using displaced vertex searches. Given the absence of conspicuous colored top partner signatures, this reach is impressive, but it misses the light glueball region of parameter space $m_{G_0} \lesssim 15$ GeV, corresponding to long glueball lifetimes $c\tau \gg 100\text{m}$. This is a very important region to probe, not only because the lifetime of the glueballs should be treated as an essentially free parameter in theories of Neutral Naturalness, but also because RG-arguments [240] can favor relatively low glueball masses for theories like the Fraternal Twin Higgs. This long-lifetime regime is an ideal target for the MATHUSLA detector, due to its low background and absence of trigger thresholds. We show the reach of MATHUSLA as the purple shaded region in Fig. 22. It provides the only direct probe of this important region of Neutral Naturalness parameter space.

4.2.1.2 Other LLP states and production modes

Additional processes may lead to the production of hidden sector bound states, though these channels are model-dependent. Here we briefly discuss LLP signatures due to hidden Bottomonia, Quirk production and production of heavy UV states.

Hidden Bottomonium production: In theories like the Fraternal Twin Higgs, where partner particles are entirely neutral under the Standard Model, hidden QCD bound states may also be produced in Higgs decays to bottom partners. Depending on $m_{\hat{b}}$ and $\hat{\Lambda}_{QCD}$, this may lead to the production of an excited quirky boundstate [247] of hidden $\hat{b}\hat{b}$, or the production of hidden bottomonium. Since this exotic Higgs decay proceeds through the mirror bottom Yukawa coupling and is otherwise only suppressed by the Higgs mixing factor v^2/f^2 , it can have $\sim 10\%$ branching ratio, leading to much larger hidden sector production rates than the hidden gluon coupling of Eqn. (63). If these mirror bottom states cascade down to lighter hidden glueballs, the LLP searches at MATHUSLA and the main detectors would cover even more of the parameter space shown in Fig. 22.

The lightest bottomonium states are the pseudoscalar $\hat{\eta}_b(0^{-+})$, the vector $\hat{Y}(1^{--})$, and the scalar $\hat{\chi}_{b0}(0^{++})$. These states have masses $\sim 2m_{\hat{b}} + \mathcal{O}(\hat{\Lambda}_{QCD})$. If the discrete symmetry is respected by the bottom Yukawa couplings, the Higgs coupling bound $f/v \gtrsim 3$ [248] implies they are heavier than ~ 35 GeV. While this makes it natural for glueballs to be lighter, it is also possible for threshold corrections to hidden QCD RG running near the discrete symmetry breaking scale to lift the glueball mass above $m_H/2$ [239]. In that case, hidden bottomonia that are produced in exotic Higgs decays can only decay back to the SM via the Higgs portal and would be detectable as LLPs. The lightest pseudoscalar state is extremely long-lived, but the scalar decays back to SM states through the Higgs

¹⁷In the short glueball lifetime regime (masses close to 60 GeV), prompt or displaced $h \rightarrow 4b$ searches at lepton colliders are likely to have better sensitivity [246].

portal, with a decay length that can be estimated as [249]

$$c\tau_{\hat{\chi}_{b0}} \simeq 8.3 \text{ cm} \left(\frac{m_b}{m_{\hat{b}}}\right)^5 \left(\frac{f}{1 \text{ TeV}}\right)^4 \left(\frac{5 \text{ GeV}}{\hat{\Lambda}_{QCD}}\right)^2 \left[\frac{9}{5} \left(\frac{\sqrt{s}}{3m_{\hat{b}}}\right)^2 - \frac{4}{5}\right]^{-1} \quad (m_{\hat{b}} \gg \hat{\Lambda}_{QCD}), \quad (67)$$

$$c\tau_{\hat{\chi}_{b0}} \simeq 3.8 \text{ cm} \left(\frac{m_b}{m_{\hat{b}}}\right)^2 \left(\frac{f}{1 \text{ TeV}}\right)^4 \left(\frac{5 \text{ GeV}}{\hat{\Lambda}_{QCD}}\right)^5 \left(\frac{\sqrt{s}}{3\hat{\Lambda}_{QCD}}\right)^{-2} \quad (m_{\hat{b}} \ll \hat{\Lambda}_{QCD}) \quad (68)$$

This estimate is valid in the range $2m_b < \sqrt{s} \ll m_h$, where \sqrt{s} is the mass of the hidden bottomonium (which may be somewhat greater than $m_{\chi_{b0}}$ if the state is excited).

It is also possible for the discrete symmetry to be broken in the bottom sector of the theory. In that case, hidden bottomonia could be much lighter than 35 GeV, with correspondingly longer lifetimes. In the original Fraternal Twin Higgs model, this would cause cosmological problems. An extremely large abundance of hidden sector states is produced during the Big Bang [250, 251], and if the scalar is much longer-lived than a meter, this abundance would no longer efficiently deplete itself via decays to the SM [249]. Late decays of the very long-lived pseudoscalar η would then disrupt BBN. However, this regime would still be permitted if there was a kinetic mixing between a massive hidden hypercharge gauge boson and SM hypercharge. This arises naturally, for example, if the UV completion of the theory contains particles charged under both hypercharges. The hidden vector meson \hat{Y} could then decay to SM particles, which sufficiently depletes the hidden hadron abundance at early times if it has a decay length below \sim meter [249]. Again, light bottomonia and cosmological constraints lead to some LLPs with relatively short lifetimes. Such lifetimes are well-suited to main detector LLP searches, and may also be good targets for LHCb [252]. MATHUSLA can then play a vital role searching for longer-lived hidden hadron states and diagnosing the connection between the hidden valley and naturalness.

Quirk Production: Models like Folded Supersymmetry and the Quirky Little Higgs feature top partners that are charged under the SM electroweak force. Drell-Yan production of top partners then provides an additional portal into the hidden sector. Once produced, these states remain connected by a hidden QCD flux tube and eventually undergo quirk-like annihilation decays [247] to hidden glueball states and possibly SM EW final states, depending on the details of the theory and spectrum. These additional processes can enhance the production of hidden sector states by an order of magnitude or more, as has been studied in [242]. This would greatly enhance the reach of LLP searches at MATHUSLA and the main detectors in the parameter space of Fig. 22.

Production of heavy UV states: The similarity between the SM and hidden gauge symmetries indicate the existence of a ‘‘unification’’ between the two sectors, and most of its UV-completion scenarios contain particles carrying both the SM and hidden charges and provide extra portals between the two sectors [235, 236, 249, 253–260].

For example, in many UV-completion of the twin Higgs models, there are exotic-fermions that carry either twin QCD + SM weak charge, or SM QCD + twin weak charge. These fermions are likely to carry masses close to the cutoff scale ($\gtrsim 5 \text{ TeV}$) near the scale of the UV-completion, and an observation of them will provide valuable information about the SM-twin unification [249]. We can produce these bi-charged fermions at the LHC either through strong or electroweak production. For the exotic-quark that carries SM color, once being produced inside detector, they promptly decay into a pair of SM tops and two twin Z bosons [261]. Each of the twin Z decays into twin quarks and form long-lived twin hadrons, and the whole event contains prompt high p_T leptons from the top decay plus displaced lepton or jet signals. The hard leptons provide simple triggering and background rejection, and the search at the LHC will be able to probe the exotic-quark mass up to $\approx 2.5 \text{ TeV}$ that is getting close to the UV scale of the model.

Besides the fermion that carries SM color, LHC can also produce the exotic-fermion that carry SM electroweak charge and twin QCD through Drell-Yan process [262]. The fermions can be produced

as a bound state bind by the mirror QCD force, which annihilates into SM fermions or gauge bosons for various resonance searches. Depending on details of the UV model and the SM charge of the produced exotic-fermion bound state, their decay can produce SM gauge bosons and/or hidden glueballs, giving different displaced signal topologies. If these bound states decay without leptons or hard SM jets in the final state, the main detector LLP searches will likely have at least some backgrounds or suffer from reduced trigger/cut efficiencies. In those scenarios, MATHUSLA will play an important role in discovering and diagnosing these LLP signals of the Neutral Naturalness UV completion.

4.2.2 LLP Signatures of Late-Time Reheating

In addition to the states directly connected to the stabilization of the weak scale, models of neutral naturalness predict a variety of additional light degrees of freedom whose couplings and masses are comparable to their Standard Model counterparts and may give rise to additional distinctive signals at MATHUSLA.

For example, models of neutral naturalness based on global symmetries (such as the Twin Higgs and its relatives) generically predict additional neutrino species in the hidden sector. The cosmology of these neutrinos gives rise to both observational constraints and potential discovery channels. In the case of the simplest realization of neutral naturalness, the mirror Twin Higgs [235], the energy density stored in twin neutrinos is in tension with CMB and BBN constraints on dark radiation [263–265]. One way to avoid these ΔN_{eff} constraints is to remove the light degrees of freedom in the hidden sector which are not instrumental for stabilizing the Higgs mass, leading to the Fraternal Higgs like models and their LLP signatures discussed in the above subsection. Another way of mitigating this tension is by diluting the energy density of the hidden sector at the time of BBN via late decays that preferentially reheat the visible sector [264, 265]. In particular, the late decay of right-handed neutrinos, motivated to explain the active neutrino masses and mixings (see Section 7), can reconcile cosmological constraints with the existence of light relics from the hidden sector which stabilizes the Higgs mass.

For the decay of right-handed neutrinos to sufficiently dilute the energy density stored in twin active neutrinos, the right-handed neutrinos must decay preferentially to the Standard Model. This is possible in a restricted region of parameter space where the right-handed neutrino masses m_N in the $\mathcal{O}(1 - 10 \text{ GeV})$ range (see [264] for details). However, the allowed parameter space of the theory opens up dramatically if the right-handed neutrinos in both sectors acquire part of their mass through Higgs portal-type couplings of the form

$$\mathcal{L} \supset \frac{1}{2\Lambda} (|H_A|^2 N_A^2 + |H_B|^2 N_B^2) \quad (69)$$

where H_A, H_B are respectively Higgs doublets in the Standard Model and mirror twin sectors, and N_A, N_B are corresponding right-handed neutrinos. The couplings in Eq. (69) ensure that the lightest right-handed neutrino mass eigenstates decay preferentially to the Standard Model, such that the energy density stored in the hidden sector is diluted consistent with observed limits.

The right-handed neutrino couplings required for viable cosmology in the mirror Twin Higgs also predict displaced decays with mean decay lengths relevant for MATHUSLA. In particular, the couplings in Eq. (69) give rise to rare displaced decays of the Higgs into right-handed neutrinos,

$$\Gamma(h \rightarrow NN) = \frac{m_h v^2}{16\pi \Lambda^2} \left(1 - \frac{4M_N^2}{m_h^2}\right)^{3/2} \quad (70)$$

In order for Eq. (69) to lead to an appropriate asymmetry in right-handed neutrino decays, the scale Λ must be of order $10 - 10^4 \text{ TeV}$, corresponding to an exotic Higgs branching ratio of $3 \times 10^{-5} - 3 \times 10^{-1}$ and therefore rates on order of femtobarn or larger at the 14 TeV LHC.




Fig. 23: Blue shaded regions: MATHUSLA reach for RH neutrino LLPs in the asymmetrically reheated MTH model with soft \mathbb{Z}_2 breaking in the neutrino sector as per Eqn. (69), setting $C = 1$ in Eqn. (70) and assuming the corresponding active neutrino is very light, with mass $m_\nu = 10^{-5}$ eV. Gray contours show the exotic Higgs decay branching ratio into RH neutrinos. In this scenario, cosmological constraints on the Mirror Twin Higgs model can be satisfied for $m_N \lesssim 30$ GeV [264]. These constraints are obtained by remapping the general $\text{Br}(h \rightarrow XX)$ MATHUSLA reach projections from Section 8.2. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

Once produced, decays of the right-handed neutrinos then proceed dominantly through the weak interactions back to light Standard Model fermions, with widths of order

$$\Gamma(N \rightarrow \text{SM}) \sim C \frac{G_F^2}{192\pi^3} \left(\frac{m_\nu}{M_N} \right) M_N^5 \quad (71)$$

where C is a $\mathcal{O}(1)$ number and m_ν is of order the masses of the appropriate active neutrino species. In the region of parameter space where Twin Higgs cosmology is consistent with CMB and BBN observables, the mean decay length of right-handed neutrinos ranges from $\sim 10^3$ m to the BBN limit at $\sim 10^7$ m. While this range of rates and lifetimes for exotic Higgs decays is challenging to probe at the LHC main detectors, it is ideally suited for MATHUSLA, as illustrated in Fig. 23. Given the long-lifetime regime of this Higgs-portal LLP search, MATHUSLA will have several orders of magnitude better reach than the HL-LHC main detectors, see Section 8.2.

In this section, we only discussed a few possible scenarios in which the states and couplings required by neutral naturalness give rise to displaced decays at the LHC. In all cases, production rates at the 14 TeV LHC are considerable, but constraining the parameter space is challenging using the LHC main detectors alone. In this respect, MATHUSLA would contribute significantly to comprehensive coverage of scenarios of neutral naturalness.

4.3 Composite Higgs¹⁸

An intriguing possibility is that the Higgs boson could be a composite state of some new, underlying strong dynamics that confines above the TeV scale. Assuming that a global symmetry of the strong dynamics is spontaneously broken at a scale f , the Higgs boson can then be identified as a Nambu-Goldstone boson, whose mass is protected by a shift symmetry. To actually generate a Higgs potential and mass, this symmetry must be explicitly broken. Motivated by partial compositeness, this is achieved via a linear mixing of elementary fields and composite operators with couplings that are related to the SM

¹⁸Peter Cox, Tony Gherghetta, Andrew Spray

gauge and Yukawa couplings (see Ref. [266] for a review). Currently, direct searches at the LHC for new resonances and deviations in Higgs couplings, which would provide evidence of the strong dynamics, place an approximate lower limit of $f \gtrsim \text{TeV}$, suggesting that composite Higgs models are becoming less natural. Furthermore, there are indirect limits from flavor and precision electroweak observables. While the precision electroweak constraints from the T parameter are avoided with a custodial symmetry, and those from the S parameter are ameliorated with sufficiently heavy vector resonances, the most stringent constraints actually arise from flavor observables, which give rise to an approximate lower bound on the scale of spontaneous symmetry breaking $f \gtrsim 10 \text{ TeV}$ [266]. It is therefore clear that composite Higgs models require additional model building in order to maintain naturalness and satisfy these constraints.

Instead, a more minimal approach is to simply assume that $f \gtrsim 10 \text{ TeV}$. Of course this simplicity comes at the price of a tuning in the Higgs potential, of order $v^2/f^2 \simeq 10^{-4}$. This meso-tuning is still a many orders of magnitude improvement compared to that encountered in the Standard Model with a Planck scale cutoff, and leads to an unnatural (or split) version of composite Higgs models, akin to models of Split Supersymmetry which preserve SUSY's various attractive features at the cost of meso-tuning (see Section 4.1.3). Interestingly, even though the resonances are now very heavy, these models can still give rise to distinctive experimental signals, such as LLPs. The crucial requirement involves improving gauge coupling unification due to a composite right-handed top quark [267]. The minimal coset preserving this one-loop result, together with a discrete symmetry needed for proton and dark matter stability, is $SU(7)/SU(6) \times U(1)$ with $f \lesssim 100 - 1000 \text{ TeV}$ [268]. This coset contains twelve Nambu-Goldstone bosons, forming a complex $\mathbf{5}$, comprising the usual Higgs doublet, H , a color triplet partner, T , and a complex singlet, S that can be a stable dark matter candidate. In addition, the composite right-handed top quark, needed for gauge coupling unification, is part of a complete $SU(6)$ multiplet containing extra exotic states, χ^c , that will be degenerate with the top quark. These states can be made sufficiently heavy by pairing them with top companions, χ , to form a Dirac mass of order f .

The particle spectrum of the unnatural (or split) composite Higgs model therefore consists of the pseudo Nambu-Goldstone bosons, H, T and S with masses $\ll f$, which are split from the resonances with masses $> f$, while the top companions have Dirac masses of order f . Thus for $f \gtrsim 10 \text{ TeV}$, the color triplet partner, T , of the Higgs doublet will generically be the lightest new colored state [268]. Its dominant decay mode is $T \rightarrow t^c b^c S S$ which arises from a dimension-six term, where $t^c(b^c)$ are the right-handed top (bottom) quarks and S is the singlet scalar. The decay length is given by

$$c\tau = 100 \text{ m} \left(\frac{1}{c_3^T} \right)^2 \left(\frac{8}{g_\rho} \right)^3 \left(\frac{3 \text{ TeV}}{m_T} \right)^5 \left(\frac{f}{200 \text{ TeV}} \right)^4 \frac{1}{J(m_t, m_S)}, \quad (72)$$

where c_3^T is an order one constant, g_ρ a strong-sector coupling, m_T (m_S) is the color triplet (singlet) scalar mass and $J(m_t, m_S)$ is a phase space factor (see Ref. [269] for details). Thus, since the scale $f \geq 10 \text{ TeV}$, the color triplet is long-lived and can decay via displaced vertices or outside the LHC main detectors with lifetimes in the most sensitive range for MATHUSLA.

The possible LLP signals at the LHC were analyzed in Ref. [269]. The color triplet scalar is pair-produced via QCD and then hadronizes to form an R-hadron. Roughly 50% of these are charged and can leave a track in the inner detector, and possibly the muon chambers. If the triplet is collider-stable (i.e. decays outside the main detectors), R-hadron searches at the LHC can be used to place limits on its mass. Limits from Run-I results [112] forbid a collider-stable color triplet with a mass below 845 GeV. At Run-II and beyond similar searches will be performed, and with 300 fb^{-1} of integrated luminosity triplet masses up to about 1.4 (1.5) TeV can be discovered (excluded) for lifetimes corresponding to $c\tau \gtrsim 10 \text{ m}$. These results are depicted in Figure 24. As shown, the efficiency of these searches decreases for lighter DM masses; displaced vertex searches have greater sensitivity when the triplet predominantly decays inside the LHC detectors. Larger values of f increase the triplet lifetime, and correspondingly enlarge the region where the R-hadron searches dominate. The mass reach remains unchanged at 1.5 TeV since it is set by the f -independent QCD production cross-section.




Fig. 24: Projections for the R-hadron (collider-stable) and displaced-vertex searches at the LHC with 300 fb^{-1} of integrated luminosity at $\sqrt{s} = 13 \text{ TeV}$, as functions of the scalar mass m_S and triplet mass m_T . The shaded regions can potentially be excluded at 95% CL and the dashed lines denote the 5σ discovery reach. The grey shaded region is excluded by current R-hadron searches at $\sqrt{s} = 8 \text{ TeV}$. This figure is taken from Ref. [269]

To be relevant for the MATHUSLA experiment the R-hadrons will need to pass through the substantial rock layer. An estimate of the survival probability suggests a large fraction will indeed make it to the surface. This is simply because the R-hadrons are heavy (much more than several hundred GeV), the energy loss per interaction is less than a GeV, and the interaction length is of order 0.5 m in rock. Furthermore, the charge of an R-hadron can change upon interactions with matter, such that the majority of R-hadrons formed by the color triplet are expected to be neutral when they reach the surface [109].

We use the procedure outlined in section 3, and in particular the estimate of Eq. (4), to project the expected number of R-hadron decays in the MATHUSLA detector volume. The color triplet T is pair-produced ($n_{\text{LLP}} = 2$) via QCD, with a production cross-section $\sigma_{\text{sig}}^{\text{LHC}} \sim 0.4 \text{ fb}$ at $m_T = 1.5 \text{ TeV}$ [103]. Production is dominantly near threshold, *i.e.* a boost factor $b \approx 1$. We show in Figure 25 the regions in the (m_T, m_S) plane where we expect $N_{\text{obs}}^{\text{MATHUSLA}}$ greater than 4 and 1, for two different values of f . The pattern of the exclusions is simple to understand. The number of R-hadrons produced is set by QCD, and so is a function of m_T (only). This determines the maximum mass that can be excluded to be $\sim 1.55 \text{ TeV}$. MATHUSLA has optimal sensitivity for decay lengths $c\tau \approx 200 \text{ m}$, which for smaller values of f requires an additional phase space suppression; the probed parameter space is then close to the kinematic boundary $m_T = 2m_S + m_t + m_b$. As f increases, less of a suppression is required and so the exclusions move to smaller values of the DM mass.

Comparing Figures 24 and 25, we see that MATHUSLA is unlikely to set stronger limits than the LHC in this minimal model. This is consistent with the arguments laid out in Section 3.2.3. The main detector searches for heavy hadronic particles that are either detector-stable or give rise to displaced vertices have very low background and suffer no particular trigger limitations. However, just as is the case for long-lived gluinos in split-SUSY (Section 4.1.3), MATHUSLA will provide an important complementary discovery channel for these long-lived hadronic particles. Any LLP search, when correlated




Fig. 25: Projected sensitivity of MATHUSLA to the minimal unnatural composite Higgs, for compositeness scale $f = 10$ and 30 TeV, and a total HL-LHC integrated luminosity $\mathcal{L} = 3000 \text{ fb}^{-1}$. In the shaded regions, we expect 4 signal events, while the dashed contours bound the region where at least one event is predicted. We also show the boundary where the four-body triplet decay is kinematically allowed. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

with the stable charged particle search that reveals mass and cross-section information, will reveal information on the lifetime of the LLP and hence the compositeness scale f . MATHUSLA will also provide information about the propagation of heavy hadronic states in ordinary matter.

4.3.1 Beyond the minimal model

The unnatural composite Higgs should be understood as a framework, rather than purely the specific minimal model discussed above. As such, while MATHUSLA has rather limited exclusion potential in the minimal scenario, it may be able to probe alternative implementations more effectively. The limitations found above derive from the LLP forming charged R-hadrons significantly often, about one-half of the time, which in turn is due to the LLP having the gauge quantum numbers of a down quark. Models with an electrically neutral LLP may be more promising, and we discuss several possibilities below.

As noted previously, all composite Higgs models consistent with gauge coupling unification will have top companion fermions χ, χ^c . In the above we took them to be heavy, $m_\chi \sim f \gg m_T$. However, their masses are determined by the Yukawa couplings, y_χ of the elementary states χ to the composite sector, $m_\chi \sim y_\chi f$. Taking all the $y_\chi \sim 1$ is the simplest possibility, but it is technically natural for one or more to be small. With a modest $y_\chi \sim 0.01 - 0.1$ and $f = 10 - 100$ TeV, top companions in the TeV range are possible. Their decays will proceed through the strong sector and be suppressed by f , so they will also be LLPs if $m_\chi < m_T$. They then determine the phenomenology in appropriate regions of parameter space.

Applying this reasoning to the minimal model, we note that the top companions comprise a **5** and

incomplete $\mathbf{10}$ of $SU(5)$. In particular, the former contains a lepton-like doublet \tilde{l} , which has electroweak production cross-sections and the charged component will decay quickly to the neutral one. SM gauge symmetries plus the requirement that the low-energy physics respects baryon and lepton number force the six-body decay $\tilde{l} \rightarrow S^\dagger S^\dagger S^\dagger qqq$ via an off-shell T , where q is the third generation quark doublet. In the strictly minimal model, a number of accidental symmetries suppress this decay further to the two-loop level, resulting in a lifetime too large for MATHUSLA to have any sensitivity. However, since these additional symmetries are not phenomenologically required, we can imagine breaking them so that the decay length is

$$c\tau_{\tilde{l}} \sim 100 \text{ m} \frac{1}{\epsilon^2} \left(\frac{8}{g_\rho}\right)^3 \left(\frac{800 \text{ GeV}}{m_{\tilde{l}}}\right)^{12} \left(\frac{m_T}{3 \text{ TeV}}\right)^4 \left(\frac{f}{60 \text{ TeV}}\right)^8, \quad (73)$$

where $\epsilon \ll 1$ parameterizes the size of the breaking of the accidental symmetries. We see that this is generically long-lived. Based on the production cross-sections for an electroweak doublet [270, 271] we expect MATHUSLA to have a potential sensitivity up to $m_{\tilde{l}} \lesssim 1.25 \text{ TeV}$. This is most relevant when the S is on the Higgs resonance, $2m_S \sim m_h$, since in that case DM searches have little sensitivity.

It is also possible to consider different global symmetries of the composite sector. This will lead to additional Nambu-Goldstone bosons that can potentially be long-lived for the same reasons that the triplet is in the minimal model. If any of these are color singlets, they will then be more susceptible to MATHUSLA searches. We outline two particularly promising possibilities. First, the symmetry breaking pattern $SU(7)/SU(6) \times U(1)$ is minimal in the number of Goldstones, but is non-minimal in the symmetries preserved by the strong sector. Gauge coupling unification requires only that we preserve an $SU(5)$ global symmetry; dark matter then requires an appropriate additional symmetry to distinguish the Goldstones. In particular, noting that $SU(7)/SU(6) \times U(1)$ breaking can be achieved by a spurion adjoint, we consider an alternative adjoint-induced breaking $SU(7)/SU(5) \times U(1) \times U(1)$. The Goldstones additionally include a second complex $\mathbf{5}$ of $SU(5)$, which is constrained to decay to DM plus visible sector particles. This will include an inert doublet that can potentially be a neutral LLP.¹⁹ Of particular note is that the SM couplings to the composite sector in this model are unchanged from the minimal scenario, while there are fewer top companions.

Second, recent developments [272–274] in the UV completions of composite Higgs models have identified a number of symmetry breaking patterns as particularly promising. The most relevant possibility for an $SU(5)$ GUT is $SU(N) \times SU(N)/SU(N)$, where symmetry breaking is induced by vector-like fermions in a complex representation of the confining gauge group. The smallest such group consistent with gauge coupling unification and where the Goldstones include a Higgs is $SU(6) \times SU(6)/SU(6)$; dark matter stability requires that we extend this to $U(6) \times U(6)/U(6)$. The pseudo Nambu-Goldstone bosons comprise a complex $\mathbf{5}$ of $SU(5)$ that contains the Higgs, as well as a real $\mathbf{24}$ and a real singlet. The details of the Nambu-Goldstone spectrum, top companions and elementary-composite couplings remain to be explored. However, we note that the $\mathbf{24}$ contains a neutral color octet charged under the DM symmetry. In some regions of parameter space, it can be an LLP, while its quantum numbers prefer the production of neutral R-hadrons, weakening LHC bounds. Thus the MATHUSLA experiment could provide valuable information on these types of models.

4.4 Relaxion Models²⁰

A novel idea to address the hierarchy problem introduces a new axion-like field, the relaxion, that couples to the Higgs field and dynamically relaxes to a field value that partially cancels the quadratic divergence in the Higgs mass squared [275]. The relaxion coupling to the Higgs induces a Higgs-relaxion mixing which can lead to long-lived decays of the relaxion. Furthermore, a two-field supersymmetric

¹⁹The third $SU(5)$ -consistent possibility for an adjoint spurion, $SU(7)/SU(5) \times SU(2) \times U(1)$, has the second $\mathbf{5}$ but no longer has the complex singlet. The inert doublet is then forced to be the dark matter.

²⁰Jason L. Evans, Tony Gherghetta, Natsumi Nagata

generalization of the relaxion mechanism [164, 276, 277], that addresses the little-hierarchy problem in supersymmetric theories, can also incorporate inflation where the second field is identified as the inflaton [278]. This has several possible signals at MATHUSLA, where the Higgs mixing with the inflation sector leads to a long-lived relaxion and the Higgsino-relaxino mixing leads to long-lived gauginos.

4.4.1 Higgs-Relaxion mixing

The relaxion can be produced through Higgs-relaxion mixing as discussed in [164, 275–278]. The relevant interaction, which is present in all of these models, is

$$\mathcal{L} \supset C \frac{h^2}{\Lambda} \Lambda_N^3 e^{i\frac{\phi}{f}} + h.c. = 2|C| \frac{h^2}{\Lambda} \Lambda_N^3 \cos\left(\frac{\phi}{f} + \delta\right), \quad (74)$$

where h is the real component of the SM Higgs field which obtains a VEV, ϕ is the relaxion, f is the decay constant associated with the breaking of a global U(1) symmetry, Λ is a UV scale typically associated with integrating out a heavy particle, Λ_N is the confinement scale of some strongly coupled gauge theory, and we have taken $C = |C|e^{i\delta}$ in the second equality. The Lagrangian term (74) leads to mixing with the Higgs boson under two different conditions. For single field relaxion models, we can take $\delta = 0$ and still have non-zero mixing because ϕ stops for $\sin\left(\frac{\phi}{f}\right) \sim 1$. For the two-field relaxion model, the relaxion mass continues to evolve after EWSB, due to the dynamical evolution of the second field. This causes $\sin\left(\frac{\phi}{f}\right)$, which determines the Higgs-relaxion mixing, to be quite small in our vacuum today. Therefore for some parameters of the two field relaxion model, the decays we discuss below may need $\delta \neq 0$, although for the two-field relaxion model in [277], $\delta \sim 0$ is needed in order to stop the relaxion and this generically leads to a long-lived relaxion. Constraints on this operator already exist and can be found in [279, 280], although Ref. [279] assumed a single field relaxion model.

Expanding the stopping potential (74) around the local minimum $\phi_0 = \langle \phi \rangle$ leads to the interactions

$$\mathcal{L} \supset 2\lambda' \sin\left(\frac{\phi_0}{f} + \delta\right) \frac{v^2}{f} h^2 \phi + \lambda' \cos\left(\frac{\phi_0}{f} + \delta\right) \frac{v^2}{f^2} h^2 \phi^2 + \dots \quad (75)$$

where we have assumed

$$|C| = \lambda' \frac{v^2 \Lambda}{\Lambda_N^3}, \quad (76)$$

with $\lambda' \lesssim 1$ in order that the relaxion VEV does not lead to a Higgs mass contribution larger than the weak scale. This gives a branching fraction

$$\text{BR}(h \rightarrow \phi\phi) = 4 \times 10^{-4} \left(\frac{\lambda'}{1}\right)^2 \left(\frac{c_\delta}{1}\right)^2 \left(\frac{10^4 \text{ GeV}}{f}\right)^4, \quad (77)$$

where $c_\delta = \cos\left(\frac{\phi_0}{f} + \delta\right)$ and we have assumed $m_\phi \ll m_h$. This branching fraction falls right in the middle of the range considered in [1].

Next we look at the decays of the relaxion. For the models we consider, the dominant decay modes of the relaxion arise via its mixing with the SM Higgs due to the term proportional to $\sin\left(\frac{\phi_0}{f} + \delta\right)$ in Eq.(74). Since this mixing is small compared to the Higgs mass, it has little effect on the relaxion and Higgs masses. In the small angle approximation we obtain²¹

$$\theta_{\phi h} \simeq 10^{-3} \left(\frac{\lambda'}{1}\right) \left(\frac{s_\delta}{10^{-2}}\right) \left(\frac{10^4 \text{ GeV}}{f}\right), \quad (78)$$

²¹In [277] for $m_{SUSY} \sim 10 \text{ TeV}$, this size of mixing is realized for $\delta \simeq 0$.

where $s_\delta = \sin\left(\frac{\phi_0}{f} + \delta\right)$ and we have assumed $m_\phi \ll m_h$. This scenario has several constraints coming from many different experiments as shown in Ref. [219], that constrains the scalar-Higgs mixing, $\theta_{\phi h} \gtrsim 10^{-3}$ for much of the parameter space. However, in the relaxion model, since the mixing angle $\theta_{\phi h}$ can be controlled by adjusting δ without affecting the production, much of the unconstrained parameter space from previous experiments can be realized. This means that relaxion models have mixing angles and production rates which can be seen at the MATHUSLA experiment.

Now we examine the decay length of the relaxion to verify that it does indeed decay inside the detector. If the mixing in Eq.(78) produces the dominant decay mode(s) for the relaxion, the relevant perturbative interactions are of the form

$$\mathcal{L} \supset \sin\theta_{\phi h} \frac{m_f}{v} \phi \bar{f} f, \quad (79)$$

where m_f is the mass of the fermion f , and v is the SM Higgs VEV. A perturbative calculation leads to the decay width

$$\Gamma_{\phi \rightarrow \bar{f} f} = \frac{1}{8\pi} \left(\frac{\sin\theta_{\phi h} m_f}{v} \right)^2 \left[1 - \frac{4m_f^2}{m_\phi^2} \right]^{3/2} m_\phi, \quad (80)$$

which will be the dominant decay mode for certain masses of the relaxion. However, for most masses the dominant decay mode will be to mesons, where non-perturbative effects dominate and the above perturbative approximation breaks down. The lifetime and branching fractions of a particle with a dominant decay mode coming from the interactions in Eq.(79) can be found in Ref. [18, 219]. Since this signal of the relaxion is identical to the LLP signal in the SM+S simplified model discussed in Section 8.4, the region which can be detected for this model at the MATHUSLA experiment can be read off of Fig. 55 by identifying $\theta \equiv \theta_{\phi h}$ and $S \equiv \phi$. It is clear that MATHUSLA can probe deep into the model's parameter space, using both exotic Higgs decays and Meson decays as a source of LLPs.

4.4.2 Mixing with the Inflation Sector

4.4.2.1 Higgs mixing

An interesting feature of the supersymmetric two-field relaxion model discussed in [278] is the identification of the inflaton with the second field that controls the amplitude of the stopping potential. The SM Higgs can then mix with fields in the inflation sector. In this model inflation occurs along a D -term flat direction which is lifted by radiative corrections, as is typical of D -term inflation. The inflation scale is pushed down to values less than a few GeV by taking the $U(1)$ gauge coupling associated with the D flat direction to be very small. One of the consequences of this very small gauge coupling is a very light waterfall field, that also has a very large VEV of order the GUT scale. This combination makes it hard to reheat since anything the waterfall field couples to strongly will be heavier than the waterfall field. This difficulty can be circumvented if the waterfall field is coupled to a supersymmetric F flat direction which involves the Higgs field [278]

$$W \supset \kappa_1 R \phi_+ \langle M_- \rangle + \kappa_2 R H_u H_d + m_R R \bar{R}. \quad (81)$$

In this expression ϕ_+ is the waterfall field, $\langle M_- \rangle$ is a residual VEV from some heavy field responsible for generating $\langle D \rangle$ during inflation, R is a singlet whose F -term contribution to the potential leads to interactions between the waterfall field and the SM Higgs, and \bar{R} has no other interactions and adjusts its VEV so that $F_R = 0$ is preserved²². This superpotential gives rise to the Higgs interaction

$$-\mathcal{L} \supset \kappa'_2 m_{\phi_+} \phi_+ h^2 + h.c. \quad (82)$$

²²This cancels large corrections to the Higgs B -term so the relaxion process still works.

where

$$\kappa_2 = \kappa_2' \frac{2m_{\phi_+}}{\sin 2\beta \kappa_1 \langle M_- \rangle}. \quad (83)$$

In the small angle limit this leads to a mixing between the SM Higgs and ϕ_+ given by

$$\theta_{\phi h} \simeq 10^{-1} \left(\frac{\kappa_2'}{1} \right) \left(\frac{m_{\phi_+}}{5 \text{ GeV}} \right), \quad (84)$$

and we have again assumed $m_{\phi_+} \ll m_h$. This situation is very analogous to the Higgs-relaxion mixing in Section 4.4.1. By varying κ_2 or κ_1 , we can reduce the mixing angle and extend the decay length. Furthermore, the ϕ_+ mass is independent of the mixing angle, so that our constraints are similar²³ to those found in Ref. [219]. Again, this relaxion implements the SM+S simplified model, and the region of parameter space that is discoverable at MATHUSLA is shown in Fig. 55 (Section 8.4).

Note also that while the decay width is similar to the Higgs-relaxion case, the production mode is very different. The waterfall field ϕ_+ is instead generated through B meson decays, and since the LHC will produce many B -mesons the production should be sufficient. More details can be found in Section 8.4 and Ref. [18].

4.4.2.2 Higgsino Mixing

Given the superpotential interactions in Eq.(81), we next describe the effect that these interactions have on the neutralino lifetimes. Since \tilde{R} will generally be much lighter than the MSSM higgsinos, the higgsinos can decay to \tilde{R} . However, the higgsinos of this model are much too heavy to be produced at the LHC. Because the interaction in Eq.(81) also generates mixing of \tilde{R} with the higgsinos,

$$\mathcal{L} \supset \kappa_2 \tilde{R} \left(v_u \tilde{H}_d + v_d \tilde{H}_u \right) + h.c., \quad (85)$$

interactions between \tilde{R} and the MSSM neutralinos will be generated. As we will see, for certain ranges of $\kappa_{1,2}$ these interactions can lead to events in the MATHUSLA detector.

There is one complication: since \tilde{R} is stable, decays of the neutralinos to \tilde{R} can overclose the universe. The dominant production mode of \tilde{R} is through Higgsino decays to \tilde{R} . Thus, the production of \tilde{R} can be drastically reduced by reheating below the Higgsino mass²⁴. For the model in [278], this means the reheat temperature can be as high as 10^6 GeV. If the reheat temperature is larger than the Bino and/or Wino mass, the mixing in Eq.(85) could still cause the Universe to overclose. However, if the decay width is sufficiently small that the decay of the Bino/Wino occurs after freeze out, then as long as the Bino or Wino thermal density is small enough, the decay to \tilde{R} will not overclose the Universe. The Bino or Wino will freeze out at a temperature of order $T_f \sim \frac{M_i}{25}$, where M_i for $i = 1, 2$ is the Bino or Wino mass respectively. Since the Wino or Bino will not decay until $\Gamma \sim H$, we can estimate the upper bound on the decay width of the Wino and Bino,

$$\Gamma(\chi \rightarrow XX) \lesssim H(T_f) = 7.5 \times 10^{-16} \text{ GeV} \left(\frac{g_\rho}{106.75} \right)^{1/2} \left(\frac{M_i}{10^3 \text{ GeV}} \right)^2, \quad (86)$$

$$= (0.26 \text{ m})^{-1} \left(\frac{g_\rho}{106.75} \right)^{1/2} \left(\frac{M_i}{10^3 \text{ GeV}} \right)^2, \quad (87)$$

²³There is another possible constraint on these models that we have not mentioned. If the coupling $\kappa_{1,2}$ causes the universe to reheat to too high of a temperature, the universe may become overclosed. This constraint can be avoided but is somewhat model-dependent.

²⁴The production of \tilde{R} can also be suppressed by taking κ_2 very small. For a reheat temperature above the weak scale, as was considered in [278], this method of suppression only works for some of the parameter space.

where g_ρ is the number of relativistic degrees of freedom and M_i is the LSP mass, and χ denotes either a Wino or Bino. Interestingly, this decay length falls right in the range where MATHUSLA is sensitive.

Therefore to obtain a viable reheating that does not over produce \tilde{R} , we need to reheat with a temperature less than the Higgsino mass while giving the LSP a decay length longer than about a meter. The reheat temperature which is generated through the interactions in Eq.(81) is

$$T_R = 153 \text{ GeV} \times \left(\frac{106.75}{g_\rho} \right)^{1/4} \left(\frac{\langle M_- \rangle}{10^{16} \text{ GeV}} \right) \left(\frac{10^3 \text{ GeV}}{m_{\phi_+}} \right)^{1/2} \left(\frac{\kappa_1}{10^{-13}} \right) \left(\frac{\kappa_2}{10^{-7}} \right), \quad (88)$$

where the normalizations of the couplings are typical values for an inflation scale of $H_I = 1 \text{ GeV}$.

Next we determine the decay length of the Wino and Bino. The easiest way to find the Bino/Wino interactions with \tilde{R} is to integrate out the higgsinos which gives

$$\mathcal{L} \supset \kappa_2 \frac{M_Z}{\mu} (s_\beta^2 - c_\beta^2) \left(-c_W \tilde{R} \lambda_3 + s_W \tilde{R} \lambda_0 \right) \frac{h}{\sqrt{2}} + h.c. \quad (89)$$

where h is the SM Higgs, λ_3 the neutral Wino and λ_0 the Bino. Using these interactions, the Wino decay width in the limit $M_2 \gg m_h, m_R$ is given by

$$\begin{aligned} \Gamma_{\lambda_3 \rightarrow h \tilde{R}} &= \frac{|\kappa_2|^2}{32\pi} \left(\frac{M_Z}{\mu} \right)^2 c_{2\beta}^2 c_W^2 M_2, \\ &\simeq (150 \text{ m})^{-1} \left(\frac{|\kappa_2|}{10^{-7}} \right)^2 \left(\frac{10 \text{ TeV}}{\mu} \right)^2 \left(\frac{c_{2\beta}}{0.5} \right)^2 \left(\frac{M_2}{1 \text{ TeV}} \right), \end{aligned} \quad (90)$$

where $c_{2\beta} = \cos 2\beta$. For a Bino LSP, the decay width is obtained by substituting $c_W \rightarrow s_W$ and $M_2 \rightarrow M_1$. As is clear from the above expressions, this is a prime candidate for detection at MATHUSLA.

Although the direct production rate of TeV neutralinos at the LHC is negligible, they can be produced through gluino decays. Since the decay rate of the gluino to \tilde{R} is quite small, all gluinos produced will eventually become the lightest neutralino. Because of this, the lightest neutralino production rate will be equal to the gluino production rate, which is large enough to produce MATHUSLA signals in a range of LLP lifetimes for gluino masses up to a few TeV, see Fig. 8. The reach projection in gluino mass will then be similar to the RPV search for gluinos decaying into neutralino LLP, shown in Fig. 14.

4.4.3 Decays to Relaxinos

Another possible signal of relaxion models which could be seen at MATHUSLA are neutralino decays to the relaxino. In supersymmetric relaxion models, the relaxion superfield couples to the field-strength superfield in a similar manner to the axion,

$$\mathcal{L} \supset \int d^2\theta \frac{C}{32\pi^2} \frac{S}{f} W^{a\alpha} W_\alpha^a + h.c. \quad (91)$$

where C is an order one coefficient²⁵. This superpotential term leads to the following interactions

$$\mathcal{L} \supset \frac{\sqrt{2} C g_a^2}{32\pi^2} \bar{\lambda}^a \Sigma^{\mu\nu} \gamma_5 \tilde{S} F_{\mu\nu}^a, \quad (92)$$

where λ^a is the gaugino Majorana fermion, \tilde{S} is the relaxino Majorana fermion, and $F_{\mu\nu}^a$ is the field-strength tensor with gauge group index a . This leads to the gaugino decay width

$$\Gamma_{\lambda^a \rightarrow \tilde{S} + \gamma/Z} = \frac{|C|^2}{128\pi} \left(\frac{g_a^2}{8\pi^2} \right)^2 \left(\frac{M_{\lambda^a}}{f} \right)^2 M_{\lambda^a}, \quad (93)$$

²⁵Decays from this operator were also discussed in [164] with S being the axion superfield.

$$= (171 \text{ m})^{-1} \left(\frac{|C|}{1} \right)^2 \left(\frac{M_{\lambda^a}}{1 \text{ TeV}} \right)^3 \left(\frac{3 \times 10^8 \text{ GeV}}{f} \right)^2 ,$$

where for the $\lambda^a \rightarrow \tilde{S}Z$ decay mode, we have absorbed some order one coefficients depending on the weak mixing angle into C . This decay length can be seen to be within the range of MATHUSLA, where the production mechanism occurs through gluinos, as discussed in Section 4.4.2.2.

4.4.4 Relaxion models with additional ultra long-lived BSM particles (ULLPs)²⁶

A heavy relaxion could yield interesting signatures in MATHUSLA if the relaxion can also decay into ultra long-lived BSM particles (ULLPs) X . The presence of such particles in relaxion models has been proposed in [281] in order to solve a cosmological problem. In some cases, reheating can de-stabilize the electroweak vacuum established by the relaxion mechanism. This occurs because the relaxion field starts to re-roll after re-heating, and adding a coupling of the relaxion to an ULLP, which in [281] is chosen to be a dark photon, can provide an additional friction term in the evolution equation of the relaxion which can stop the relaxion from rolling.

If such a ULLP is present and it is produced from relaxion decays, then its production cross-section is dictated by the relaxion production discussed in Ref. [280], while its lifetime follows from the ULLP couplings to Standard Model particles. Realizations of scalar or dark photon ULLPs are possible to construct. In the following we give estimates of the number of events which could be measured in MATHUSLA in a more general description, in terms of the branching ratios, life-times, and masses of the relaxion and the ULLP. This effectively realizes a well-motivated extension of the SM+S simplified model of Section 8.4, where S (the relaxion) decays into additional hidden-sector LLP states.

Production from B -meson decays

GeV-scale relaxions can be produced in $B \rightarrow K\phi$ decays. The $b\bar{b}$ production cross-section at the LHC is around $500 \mu\text{b}$ [282]. These B -mesons have a branching ratio into relaxions given by²⁷

$$Br(B \rightarrow K\phi) \approx 6.2 \sin^2 \theta , \quad (94)$$

where $\sin \theta$ is the Higgs-relaxion mixing angle. A current bound on $Br(B \rightarrow K\phi) \times Br(\phi \rightarrow XX)$ arises from the bound on $Br(B \rightarrow K\nu\bar{\nu}) \lesssim 1.4 \times 10^{-5}$ (see Eq.(14) of Ref. [283]). Thus, the largest potential LLP production cross-section from B-decays to a relaxion mixed with the Higgs is

$$\sigma_{pp \rightarrow XX} \lesssim 7 \text{ nb} , \quad (95)$$

corresponding to a relaxion-Higgs mixing angle $s_\theta \lesssim 10^{-3}$. This cross-section is well above the sensitivity expected for MATHUSLA, corresponding to $\sigma \sim 1 \text{ fb}$, demonstrating that MATHUSLA would be sensitive down to mixing angles $\sin \theta \sim 10^{-6}$.

Production from Higgs decays

The SM Higgs cross-section at 14 TeV is $\sigma_H \approx 50 \text{ pb}$. Furthermore, the anticipated limit on invisible Higgs decays from the main detectors at the end of HL-LHC running will be $\mathcal{O}(1 - 10\%)$. Thus the largest potential LLP production cross-section from Higgs decays to relaxions mixed with the Higgs which then subsequently decay to LLPs is

$$\sigma_{pp \rightarrow H \rightarrow XX} \lesssim 5 \text{ pb} . \quad (96)$$

The branching ratio into relaxions depends not only on the mixing angle, but also other scalar potential parameters, thus the cross-section cannot be simply related to the mixing angle as for B-meson decays. Nonetheless, the branching ratio can easily surpass $BR(H \rightarrow \phi\phi) \gtrsim 2 \times 10^{-5}$,²⁸ thus MATHUSLA has

²⁶Thomas Flacke, Claudia Frugiuele, Elina Fuchs, Rick S. Gupta, Gilad Perez and Matthias Schlaffer

²⁷See Sec. 8.4.

²⁸See Sec. 8.4.

sensitivity to new LLPs through relaxion-Higgs mixing.

5 Theory Motivation for LLPs: Dark Matter

The existence of dark matter (DM), comprising some 26% of the present-day energy budget of our universe [284], has been solidly established by several independent lines of gravitational evidence, and provides some of the sharpest evidence for new particle physics at potentially accessible mass scales. The particle nature of DM remains a mystery. Null results to date in indirect detection (ID), direct detection (DD), and missing energy searches at colliders have forced models of WIMP DM into severely constrained regions of parameter space, and have helped to stimulate a broader investigation into possible signals of particle dark matter. There are a wide variety of possible DM candidates, whose experimental signals are intimately connected to the mechanism responsible for populating DM in the early universe. These DM models often require new BSM states in addition to DM itself, or multi-component DM. In many cases, the mechanism that yields the correct relic density for DM then naturally and generically results in one or more of these BSM states having a proper decay length on collider scales. In other cases, long lifetimes are not a direct consequence of the mechanism that determines the DM relic abundance, but are a generic feature of models that implement it.

There are several reasons why a particle may have a collider-scale lifetime: a renormalizable coupling controlling its decay may simply be tiny; the phase space available for its decay may be small; and/or its available decays may be suppressed by high mass scales. As we demonstrate in this chapter, all of these mechanisms are naturally realized in well-motivated DM models. For instance, small phase space is a generic prediction of models where WIMPs coannihilate with an additional particle in the early universe. In this case the cosmically-mandated small mass splitting Δ between DM and its coannihilating partner can frequently result in long partner lifetimes, as discussed in Sec. 5.1. Decays suppressed by high mass scales naturally arise in theories of asymmetric DM, which, motivated by the apparent coincidence $\Omega_{DM} \approx 5\Omega_b$, relate SM baryon (or lepton) number to a conserved dark number D . Relating a baryon number asymmetry to a dark number asymmetry requires new interactions to transfer asymmetries between SM and dark sectors, which can be described through the introduction of transfer operators \mathcal{O}_{ADM} that carry both $B - L$ and dark number D . These operators are necessarily non-renormalizable, and can lead to displaced decays of (e.g.) SM superpartners, as demonstrated in Sec. 5.2.

Models like SIMPs and ELDERs, discussed in Sec. 5.4, require DM to have rapid number-changing self-interactions and to be (at least initially) in thermal contact with the SM. These models are perhaps most naturally realized when DM lives in a confining hidden sector. In this case many of the hidden sector hadrons, notably vector mesons, are natural and attractive targets for collider LLP searches; lifetimes are in this case rendered long thanks in large part to a small portal coupling between the SM and the hidden sector. Along with ADM scenarios, SIMPs and its relatives represent hidden valleys (Section 8.1) realizing both Dark Matter and collider signatures.

When DM lives in a sector that is not in thermal equilibrium with the SM in the early universe, a variety of novel possibilities open up for the thermal history of DM and thus for its signals today [285–287]. Such thermally decoupled hidden sectors require the leading coupling between the HS and the SM to be very small, and thus generically these hidden sectors will contain LLPs. However, the leading coupling between the SM and the HS is then generally too small to allow that same coupling to mediate production of HS particles at rates large enough to be observable at the LHC and MATHUSLA [288, 289]. Probing non-equilibrated dark states at both LHC and MATHUSLA thus requires the LLP to be produced in cascade decays of a BSM parent particle. This is the case in e.g. freeze-in scenarios (discussed in Sec. 5.3), as well as co-decaying models (Sec. 5.4 below). Conversely, the same observation allows us to immediately conclude that any LLPs observed at MATHUSLA whose production and decay are governed by the same couplings would have been in thermal equilibrium with the SM in the early universe. In that case, their lifetime is bounded from above so as not to disrupt Big Bang Nucleosynthesis, see Section 8.3. Hidden sector dark matter therefore provides strong motivations for LLPs: for instance, LLPs in the simple models SM+S, SM+V (see Secs. 8.4 and 8.5) are the leading collider signal of a class of secluded DM models [10], and a scenario where MATHUSLA can probe unique territory.

An important exception to the BBN constraint is provided by the Dynamical Dark Matter framework, see Section 5.5, which generalizes the notion of a single or a few hyperstable DM states to an ensemble of states with varying lifetimes. Some of these constitute the DM abundance today, while heavier states in the ensemble may be produced at the LHC and decay with lifetimes observable by MATHUSLA.

5.1 Coannihilation²⁹

Dark matter coannihilation [290] offers an attractive and useful twist to the standard dark matter relic abundance calculation. Typically, dark matter is assumed to be in thermal equilibrium with Standard Model particles in the early universe, and as the universe expands and cools, the comoving number density of the dark matter falls exponentially. As the temperature cools below about $T \sim m_{\text{DM}}/25$, however, the condition for chemical equilibrium fails, and the dark matter relic abundance is set. Moreover, in generic dark sectors, the additional particles in the model can significantly affect the ultimate dark matter abundance. Since the thermal freeze-out of dark matter occurs at finite temperature, these particles can have significant number densities as long as they are relatively close in mass to the dark matter, with lifetimes that are cosmologically negligible but very attractive targets for MATHUSLA. These particles, which are not the dark matter, will contribute new channels to the effective thermally averaged cross-section for dark matter annihilations to Standard Model particles. From Ref. [290], the effective annihilation cross-section is

$$\sigma_{\text{eff}} = \frac{g_{\text{DM}}^2}{g_{\text{eff}}^2} \left\{ \sigma_{\text{DM DM}} + 2\sigma_{\text{DM X}} \frac{g_{\text{X}}}{g_{\text{DM}}} (1 + \Delta)^{3/2} \exp(-x\Delta) + \sigma_{\text{X X}} \frac{g_{\text{X}}^2}{g_{\text{DM}}^2} (1 + \Delta)^3 \exp(-2x\Delta) \right\}, \quad (97)$$

where $\Delta = (m_{\text{X}} - m_{\text{DM}})/m_{\text{DM}}$, $x = m_{\text{DM}}/T$, g_a counts the number of degrees of freedom (spin, color, etc.) for $a = \text{DM}$ or X , the coannihilation partner, and $g_{\text{eff}} = \sum_{i=1}^N g_i (1 + \Delta_i)^{3/2} \exp(-x\Delta_i)$, $\Delta_i = \Delta$ for X and $\Delta = 0$ for DM , is the effective number of degrees of freedom in the dark sector. We assume that DM and coannihilation partners are in thermal and chemical equilibrium. In particular, we note that when the dark matter has a vanishingly small self-annihilation cross-section $\sigma_{\text{DM DM}}$, the dominant contributions to σ_{eff} are then from the coannihilation processes $\sigma_{\text{DM X}}$ and $\sigma_{\text{X X}}$.

The coannihilation process in the early universe can also play a dominant role in determining the phenomenology of dark matter production at colliders [291–293]. In particular, since the mass splitting between the dark matter and the coannihilation partner is typically small, the coannihilation partner can be long-lived on collider timescales. Even if the mass splitting is relatively large, the dominant Standard Model coannihilation products may be heavily kinematically suppressed, and this suppression can also lead to long-lived signatures. For MATHUSLA, we illustrate these ideas with a concrete simplified model.

5.1.1 Coannihilation through Higgs mediator

We study the following Lagrangian, in the broken phase of electroweak symmetry,

$$\mathcal{L} \supset \bar{\chi}(i\not{\partial} - m_{\chi})\chi + \bar{\psi}(i\not{\partial} - m_{\psi})\psi + (yh\bar{\chi}\psi + \text{h.c.}), \quad (98)$$

where $m_{\psi} > m_{\chi}$ and hence ψ is long-lived while χ is the dark matter. This structure can be realized in, for example, the minimal supersymmetric model when the lightest and second-to-lightest neutralinos are both dominantly bino-Higgsino admixtures³⁰, and the relevant mass range can be taken in $\mathcal{O}(100)$ GeV to several hundred GeV range (cf. Fig. 11 of Ref. [294]), corresponding to pair production cross-sections of

²⁹Felix Yu

³⁰In this case, the third lightest neutralino is also a bino-Higgsino admixture, and the lack of a direct Higgs-mediated $\bar{\chi}\chi$ annihilation channel can be motivated by taking a very small bino-Higgsino mixing angle for the lightest neutralino.

1-100 fb, see also Fig. 8. Separately, these states can be produced in cascade decays of heavier colored particles, as in supersymmetric quarks, extending the mass reach to roughly 1.2 TeV, but additional coannihilation channels should be included if the mass difference between the squarks and the dark matter is too compressed. In this simplified Higgs mediator scenario, as shown in Section 3.2.1, direct searches for the unstable coannihilation partner ψ at MATHUSLA can have much larger sensitivity than direct MET searches, especially if the squark-neutralino mass difference is even moderately compressed.

We have three free parameters: $\Delta = (m_\psi - m_\chi)/m_\chi$, m_χ , and y . The partial decay width of ψ is

$$\Gamma(\psi \rightarrow f \bar{f} \chi) = \frac{N_c y^2 y_f^2}{15\pi^3} \Delta^5 \frac{m_\chi^5}{m_H^4}, \quad (99)$$

where $N_c = 3$ for quark final states and 1 for leptons, and the SM final state masses are neglected. For simplicity, if we only consider one three-body decay mode, and if the fractional mass splitting $\Delta = 0.01$, the corresponding ψ decay length is macroscopic,

$$c\tau \sim 2.24 \times 10^1 \text{ m} \frac{1}{N_c y^2} \left(\frac{10^{-3}}{y_f}\right)^2 \left(\frac{10^{-2}}{\Delta}\right)^5 \left(\frac{100 \text{ GeV}}{m_\chi}\right)^5. \quad (100)$$

We remark that the decay length will be modified by $\mathcal{O}(1)$ factors depending on the interplay between the available phase space, given by $m_\chi \Delta$, and the QCD phase transition, but this estimate immediately points to the $\mathcal{O}(100)$ m lifetimes in the target zone for MATHUSLA. Moreover, in this benchmark model, the Higgs-mediated decays are dominantly hadronic and hence very difficult to detect from comparable LLP searches at the HL-LHC for masses below a few hundred GeV. MATHUSLA is therefore expected to increase sensitivity by orders of magnitude due to its background-free environment, as estimated in Section 3.2.3.

We also remark that when annihilating to heavy flavors, such as top quarks, the multiplicities of on-shell final state particles will counterbalance the large Yukawa enhancement and provide additional corrections to the above lifetime estimate. Nevertheless, this estimate demonstrates that macroscopic decays are a characteristic signature of the coannihilation partner with a small fractional mass splitting from the dark matter. If we generalize the Yukawa interaction to a singlet scalar mediator whose couplings to pairs of SM particles are free parameters, then the y_f Yukawa parameters above become model-dependent and new targets for MATHUSLA open up.

5.2 Asymmetric Dark Matter³¹

Asymmetric Dark Matter (ADM) is a class of hidden sector or hidden valley DM models where the DM density is set by its coupling to SM baryon or lepton number. Because $\rho_D M/\rho_b \sim 5$ observationally, the natural mass scale for DM in this model is $m_X \sim 5m_p$, though depending on the details of the model, other masses are possible. Higher dimension operators share a primordial asymmetry between the two sectors, and then decouple at low temperatures to separately freeze-in the asymmetry in the visible and dark (hidden valley) sector. See Ref. [295] for a review of these models. Asymmetric dark matter (ADM) arising from a hidden sector naturally gives rise to LLPs at the LHC [296], as we discuss here, following Ref. [297].

To transfer the asymmetry between sectors, we need a higher dimension operator connecting the SM operator \mathcal{O}_{B-L} which carries no Standard Model gauged quantum number, but carries $B - L$, to a hidden valley operator \mathcal{O}_X carrying DM number:

$$\mathcal{O}_{ADM} = \frac{\mathcal{O}_{B-L} \mathcal{O}_X}{M^{n+m-4}}, \quad (101)$$

³¹Kathryn Zurek




Fig. 26: *Left:* 3-body decay of a squark LOSP directly through the interaction of Eq. (102), and *Right:* 4-body decay of a neutralino LOSP through an off-shell squark for qld^c models. Here, the quark flavors q and q' are generically different. \tilde{x} denotes the scalar component of the ADM supermultiplet X . Decay of a slepton LOSP and a neutralino LOSP through an off-shell squark is also given by the same diagrams trading a squark and a lepton with a slepton and a quark, respectively. Figure from Ref. [297].

where \mathcal{O}_{B-L} and \mathcal{O}_X have dimension m , n , respectively. These interactions may be embedded in either a supersymmetric or non-supersymmetric theory. In the supersymmetric case, the simplest operators are

$$W_{\text{ADM}} = X\ell H, \quad \frac{Xu_i^c d_j^c d_k^c}{M_{ijk}}, \quad \frac{Xq_i \ell_j d_k^c}{M_{ijk}}, \quad \frac{X\ell_i \ell_j e_k^c}{M_{ijk}}. \quad (102)$$

These ADM interactions de-stabilize the lightest ordinary supersymmetric particle (LOSP) to decay into the X -sector plus additional SM particles, and can naturally lead to very long LOSP lifetimes. In non-supersymmetric implementations of ADM, the interactions of Eq. 101 can also naturally lead to LLPs, but the precise details of the LLP signatures will generally depend on the specific UV completion. For simplicity and generality, we will thus focus on the supersymmetric case, although qualitatively similar conclusions apply to non-supersymmetric models as well.

For example, as shown in Fig. 26, in the model with $\mathcal{O}_{B-L} = qld^c$, a squark LOSP decays to \tilde{X} , ℓ , q , while a neutralino LOSP decays via an off-shell squark to q , q' , ℓ , \tilde{X} , where the tilde denotes the scalar superpartner. As detailed in Ref. [297], the lifetime of the LOSP in the decay process depends on the scale M of the operator. This scale is in turn constrained by flavor physics.

The lifetime for the squark or slepton LOSP decay process, shown in the left side of Fig. 26, and ignoring masses of final state particles, is

$$c\tau \sim 100 \text{ m} \times \left(\frac{10^{-5} \text{ mm}}{F^{(3\text{-body})}} \right) \left(\frac{1 \text{ TeV}}{m_{\text{LOSP}}} \right)^3 \times \left(\frac{M_{ijk}}{3 \times 10^{11} \text{ GeV}} \right)^2. \quad (103)$$

The lifetime for the 4-body neutralino LOSP, as shown in the right side of Fig. 26 is

$$c\tau \sim 100 \text{ m} \times \left(\frac{100 \text{ mm}}{F^{(4\text{-body})}} \right) \times \left(\frac{M_{ijk}}{10^8 \text{ GeV}} \right)^2 \times \left(\frac{m_\phi}{1500 \text{ GeV}} \right)^4 \times \left(\frac{500 \text{ GeV}}{m_{\text{LOSP}}} \right)^7 \times (104) \\ \times x^5 [(10x^3 - 120x^2 + 120x) + 60(1-x)(2-x) \log(1-x)]^{-1},$$

where ϕ is the intermediate squark or slepton and $x = (m_{\text{LOSP}}/m_\phi)^2$. The 3-body and 4-body coefficients, $F^{(3\text{-body})}$ and $F^{(4\text{-body})}$, are calculated in Ref. [297]. Typical values are $(F^{(3\text{-body})})^{-1} \sim \text{few} \times 10^{-5} \text{ mm}$ and $(F^{(4\text{-body})})^{-1} \sim \text{few} \times 100 \text{ mm}$. Both processes give rise to macroscopic proper lifetimes depending on the supersymmetric particle masses and the scale of the ADM operator M_{ijk} .

Four-body neutralino decays in particular can easily give rise to proper decay lengths in the ~ 100 m range well-suited to MATHUSLA.

The LLPs in supersymmetric ADM theories arise naturally from the need to transfer the matter asymmetry using higher-dimension operators, combined with small phase space and the hierarchy of scales imposed by flavor physics. Consequently, the existence of LLPs is largely independent of the detailed superpartner spectrum, though the precise value of the lifetime can depend sensitively on the masses of the LOSP and the lightest relevant sfermion. Within the ADM scenario, the superpartner spectrum is mostly relevant for determining the LOSP lifetime and for setting the overall production of SUSY particles at the LHC. The collider signatures of ADM theories are broadly similar to those of RPV SUSY, for which MATHUSLA’s sensitivity can be read off from Fig. 14. Heavier sfermions translate into longer lifetimes for the LOSP, so for a neutralino ULLP either direct electroweak pair production or production from parent gluino cascades are especially well-motivated.

In the cases when the operator mediating LLP decay involves a light lepton ($W_{\text{ADM}} = X\ell H$, $\frac{Xq_i\ell_j d_k^c}{M_{ijk}}$, $\frac{X\ell_i\ell_j e_k^c}{M_{ijk}}$ in Eq. 102), the neutralino LOSP decay will be detectable at the HL-LHC with low background owing to the high- p_T lepton produced in the decay. In this case MATHUSLA will have sensitivity comparable to or up to ~ 10 times better than the HL-LHC detectors, as discussed in Sec. 3. Models that preferentially couple to taus through these leptonic operators will have higher backgrounds at the main detectors, and correspondingly greater relative advantage at MATHUSLA. For $W_{\text{ADM}} = \frac{Xu_i^c d_j^c d_k^c}{M_{ijk}}$, the LOSP has a purely hadronic decay and MATHUSLA can have much better sensitivity than the HL-LHC, by up to three orders of magnitude depending on the overall energy scale of the event.

5.3 Freeze-In Scenarios³²

Thermal freeze-out is one of the most popular mechanisms for dark matter (DM) production. DM particles have interactions with the thermal bath strong enough to be in thermal equilibrium at high temperatures. As the universe cools, and the radiation bath temperature drops below the DM mass, the expansion rate becomes larger than the annihilation rate and the DM particles go out of chemical equilibrium. This freeze-out happens at temperatures typically a factor of 20 below the DM mass, and therefore the final DM abundance is insensitive to the history of the universe before freeze-out.

Freeze-in is another motivated mechanism where DM production is dominated at IR temperatures of the order of the DM mass [298]. In such scenarios, DM particles are extremely weakly coupled to the thermal bath and never achieve thermal equilibrium. Bath particles scatter and/or decay to final states containing the DM particle, and these reactions proceed only in one direction: the DM abundance increases towards equilibrium, but never reaches it. Here, we discuss how DM is produced through freeze-in from the decay of parent particles in the thermal bath. If the interactions mediating this process are renormalizable, most of the DM is produced at temperatures of the order the bath particle mass. Interestingly, for a wide range of different cosmological evolutions, the decay length required for the observed dark matter abundance leads to displaced signals at colliders and in the MATHUSLA detector.

We review the set-up for freeze-in calculations in Sec. 5.3.1, and we evaluate the DM relic density for two different cosmological histories. First, we consider the standard cosmology where the universe snapshot at the time of Big Bang Nucleosynthesis (BBN) is extrapolated to high temperatures. The freeze-in calculation for this radiation-dominated (RD) universe is presented in Sec. 5.3.1.1. A motivated modification of the standard history, naturally arising in extensions of the standard model of particle physics, is an early matter-dominated (MD) epoch.³³ We compute DM freeze-in relic density for this case in Sec. 5.3.1.2. As a summary of this model-independent analysis, in Sec. 5.3.1.3 we present the

³²Raymond T. Co, Francesco D’Eramo, Lawrence J. Hall, Jose Miguel No, Stephen M. West, Bryan Zaldivar

³³Freeze-in taking place in a universe undergoing a *faster* than standard expansion rate also leads to displaced signatures, as studied in [299].

prediction for the decay length of the parent particle under the condition of reproducing the observed DM abundance [300]. Remarkably, our predictions are in the ballpark suitable for MATHUSLA.

In freeze-in scenarios, the DM parent particle could have a variety of couplings to the SM model, but here we focus on neutral parents, since neutral LLPs are of greatest interest to MATHUSLA. In Sec. 5.3.2, we consider a simplified model of fermionic dark matter, where freeze-in production can proceed via the Higgs portal. This minimal benchmark scenario demonstrates the range of neutral LLP signatures predicted by a standard radiation-dominated cosmology, and demonstrates that MATHUSLA can probe a wide portion of the freeze-in DM parameter space.

In Sec. 5.3.3 we discuss a more complete model which gives rise to freeze-in DM and the associated LLP signatures for a range of different RD and MD early universe cosmologies: the axino in supersymmetric DFSZ theories [5, 301]. In this scenario, the Higgsino is the parent LLP which can be observed at colliders and at MATHUSLA, realizing the signatures described in Sec. 5.3.2 as well as Sec. 4.1.5. Freeze-in on a RD background is discussed in Sec. 5.3.3.1, requiring very light axinos and hence a low inflationary reheating temperature to avoid the associated gravitino problem. This restriction can be avoided if there is a dilution mechanism to reduce the axino abundance to the observed value, realized via an early MD epoch. We show in Sec. 5.3.3.2 how supersymmetric DFSZ theories naturally incorporate a dilution mechanism through the saxion condensate. The diluted axino abundance is computed in Sec. 5.3.3.3, where we describe how this dilution effect leads to an axino abundance consistent with observations, and we generalize these conclusions to arbitrary dilution mechanisms in Sec. 5.3.3.4. We summarize the predicted LLP signals of the DFSZ freeze-in axino DM scenario in Sec. 5.3.3.5 and show that MATHUSLA puts almost the entire motivated parameter space within our reach.

5.3.1 Dark Matter Freeze-In

We consider DM freeze-in production through decays of a parent particle in thermal equilibrium with the plasma. Consistently with the notation in Ref. [300], we denote this process as follows:

$$B \rightarrow A_{\text{SM}} X, \quad (105)$$

where B is the decaying bath particle, X the DM and A_{SM} is one or more Standard Model particles. The DM abundance is initially negligible, and continuously increases as the bath particles decay. This production process is effective as long as the parent particle is relativistic. Once the temperature drops below m_B , the abundance of the decaying particle is exponentially suppressed and freeze-in is not effective anymore. Most X particles are produced at temperatures $T_{\text{FI}} \simeq m_B$.

The DM number density evolves according to the Boltzmann equation [298]

$$\frac{dn_X}{dt} + 3Hn_X = \Gamma_B n_B^{\text{eq}} \frac{K_1[m_B/T]}{K_2[m_B/T]}, \quad (106)$$

where Γ_B is the decay rate for the process in Eqn. 105 and $K_{1,2}[x]$ are the first and second modified Bessel functions of the 2nd kind. The bath particle equilibrium number density n_B^{eq} appearing on the right-hand side of the Boltzmann equation can be obtained using Maxwell-Boltzmann statistics

$$n_B^{\text{eq}} = \frac{g_B}{2\pi^2} m_B^2 T K_2[m_B/T]. \quad (107)$$

At high temperatures ($T \gg m_B$), we recover the T^3 dependence for a relativistic species, while at low temperatures ($T \ll m_B$) the number density has the Maxwell-Boltzmann exponential suppression.

The Boltzmann equation Eqn. 106 is general and its validity extends beyond the standard RD cosmology. The details of the cosmological history enters through the Hubble parameter H and the time vs temperature relation. In what follows, we present solutions to this Boltzmann equation for two different cosmological backgrounds.

5.3.1.1 Freeze-In for Standard Cosmology

In a standard RD cosmological background, the expansion of the universe is driven by its radiation content and the Hubble parameter reads

$$H = \frac{\sqrt{\rho}}{\sqrt{3}M_{\text{Pl}}} = \frac{\pi g_*^{1/2}}{3\sqrt{10}} \frac{T^2}{M_{\text{Pl}}}. \quad (108)$$

Here, g_* is the number of effective relativistic degrees of freedom, taken constant for this discussion. In such a background, the total entropy of the radiation bath is conserved, and therefore it is convenient to employ comoving variables. We define the comoving DM density $Y_X = n_X/s$, where s is the entropy density. Furthermore, we describe the evolution in terms of the inverse temperature $x = m_B/T$. The Boltzmann equation in terms of these dimensionless variables reads

$$\frac{dY_X}{dx} = \frac{\Gamma_B}{H x} Y_B^{\text{eq}}(x) \frac{K_1[x_B]}{K_2[x_B]}. \quad (109)$$

The equilibrium comoving number density for the bath particle is defined as $Y_B^{\text{eq}} = n_B^{\text{eq}}/s$. This differential equation can be integrated, with the initial condition that at very high temperatures the abundance of X is vanishing. The final comoving DM number density results in

$$Y_X^\infty = 4.4 \times 10^{-12} \left(\frac{g_B}{2}\right) \left(\frac{106.75}{g_*}\right)^{3/2} \left(\frac{300 \text{ GeV}}{m_B}\right) \left(\frac{\Gamma_B/m_B}{1.8 \times 10^{-25}}\right) \quad (110)$$

where g_B is the internal degrees of freedom of B .

This result has to be compared with the observed DM density, which can be expressed in terms of a comoving energy density

$$\frac{\rho_{DM}^{\text{obs}}}{s} = 0.44 \text{ eV}, \quad (111)$$

close to the temperature of matter radiation equality, $T_{eq} \simeq 1 \text{ eV}$. We can thus rewrite Eqn. 110 as follows

$$\frac{\rho_X}{\rho_{DM}^{\text{obs}}} = \frac{m_X Y_X^\infty}{\rho_{DM}^{\text{obs}}/s} = \left(\frac{g_B}{2}\right) \left(\frac{106.75}{g_*}\right)^{3/2} \left(\frac{m_X}{100 \text{ GeV}}\right) \left(\frac{300 \text{ GeV}}{m_B}\right) \left(\frac{\Gamma_B/m_B}{1.8 \times 10^{-25}}\right). \quad (112)$$

Upon requiring that this ratio is one, we obtain a prediction for the decay length of B

$$c\tau_B \sim 4 \times 10^6 \text{ m} \left(\frac{g_B}{2}\right) \left(\frac{106.75}{g_*}\right)^{3/2} \left(\frac{m_X}{100 \text{ GeV}}\right) \left(\frac{300 \text{ GeV}}{m_B}\right)^2. \quad (113)$$

The coupling λ , defined by

$$\Gamma_B = \frac{\lambda^2}{8\pi} m_B, \quad (114)$$

must be very small to avoid overclosure. For the benchmark points chosen in Eqn. 112, the observed DM density results for $\lambda \simeq 2 \times 10^{-12}$. Well-motivated examples of such feeble couplings include the gravitino, through interactions suppressed by the Planck scale, and the axino, through interactions suppressed by the Peccei-Quinn scale [5].

As discussed in Sec. 3.1.1, very long decay lengths near the BBN scale of $\sim 10^7 \text{ m}$ may be observable at MATHUSLA if the parent particle has a large enough cross-section near the pb range. However, the greatest chance for discovery exists for lifetimes below $\sim \text{km}$. In freeze-in scenarios, this can be realized in two ways: either either the DM candidate has a mass below the GeV scale (Sec. 5.3.2), or the relic density of the DM candidate is diluted by an earlier MD epoch (Sec. 5.3.3) as we describe below.

5.3.1.2 Freeze-In for an Early Matter-Dominated Epoch

The standard cosmology in Sec. 5.3.1.1 assumes that freeze-in occurs during the RD era and that total entropy is conserved after freeze-in, i.e. ρ_{DM}/s is constant between freeze-in and today. This conventional picture is drastically modified when there exists a late decaying matter field M . If M decays after dominating the energy content of the Universe, this late decay injects a large amount of entropy and dilutes the dark matter abundance. Examples of such matter include inflatons and moduli, and we discuss the saxion as a well-motivated candidate in Sec. 5.3.3.2.

We now elaborate on the details of the cosmological evolution. We begin with the case when M is not the inflaton so the matter energy density ρ_M is initially subdominant to radiation. Due to the scaling of ρ_M with the scale factor a^{-3} , ρ_M will eventually dominate over that of radiation, which scales as a^{-4} . This onset of the matter-dominated epoch occurs at temperature T_M . This MD era ends when M decays at the reheat temperature T_R , and is followed by a radiation-dominated era.

The MD epoch itself consists of two phases—adiabatic and non-adiabatic. When the Universe first enters the MD era, the decay of M is still inefficient, i.e. the Hubble rate is much larger than the decay rate of M ($H \gg \Gamma_M$). This means that the radiation energy density originated from the existing red-shifted radiation. The total entropy is conserved during this phase, which is thus adiabatic and called MD_A . On the other hand, the relativistic decay products of M will eventually outnumber the original radiation at the temperature we call $T_{NA} \sim (T_M T_R^4)^{1/5}$. Between T_{NA} and T_R , as radiation is constantly produced by M decay, the Universe is being reheated and therefore a large amount of entropy is actively injected. We call this a non-adiabatic phase, MD_{NA} .

In the case of M as the inflaton, inflation ends when M starts to oscillate around the true minimum and at this time M is the dominant contribution of the total energy density. Therefore, we enter a MD era immediately after inflation. The decay products of M also quickly becomes the dominant source of radiation so the Universe enters a MD_{NA} era without going through a MD_A era. The MD era ends when the inflaton completely decays away at T_R .

If dark matter is produced before and during the MD_A epoch, the abundance receives the full dilution factor $D \approx T_M/T_R \approx (T_{NA}/T_R)^5$. On the other hand, if the dark matter is dominantly produced during the MD_{NA} epoch at temperature T_X , the partial dilution factor is $D(T_X) = (T_X/T_R)^5$. In particular, for freeze-in $T_X = T_{FI}$. In addition to dilution, in calculating DM abundance, one needs to take into account the different Hubble rate during the MD epoch.

We now show the final results for freeze-in production in a matter-dominated background cosmology. The full derivations can be found in Ref. [300]. The dark matter particle X is constantly produced from the decays of B until B becomes non-relativistic with a Boltzmann-suppressed number density. As a result, the freeze-in process is IR-dominated at $T \sim m_B$. Dropping numerical factors of $\mathcal{O}(1)$, the final yield today reads

$$Y_i \sim \frac{\lambda^2 M_{Pl}}{m_B} \left(1, \quad 10^4 \frac{T_R^7}{m_B^7}, \quad \frac{T_R m_B^{1/2}}{T_M^{3/2}}, \quad \frac{T_R}{T_M} \right). \quad (115)$$

where $i = 1 - 4$ runs over the cases where freeze-in occurs in the (RD, MD_{NA} , MD_A , RD') eras, with RD the usual RD era at $T < T_R$ and RD' the early one at $T > T_M$. The first (RD) component reproduces the scaling behavior of Eqn. 110. The last three terms in the parentheses of Eqn. 115 are necessarily less than unity, and therefore the abundance is suppressed if freeze-in happens during MD_{NA} , MD_A , and RD' eras. This depletion results from a larger Hubble rate and/or large dilution due to the entropy production of M decays.

5.3.1.3 Displaced Signals at Colliders

As shown in Eqn. 115, freeze-in production is altered when a long matter-dominated era is present. This implies that the observed DM abundance also requires a different decay rate of B . In particular, once

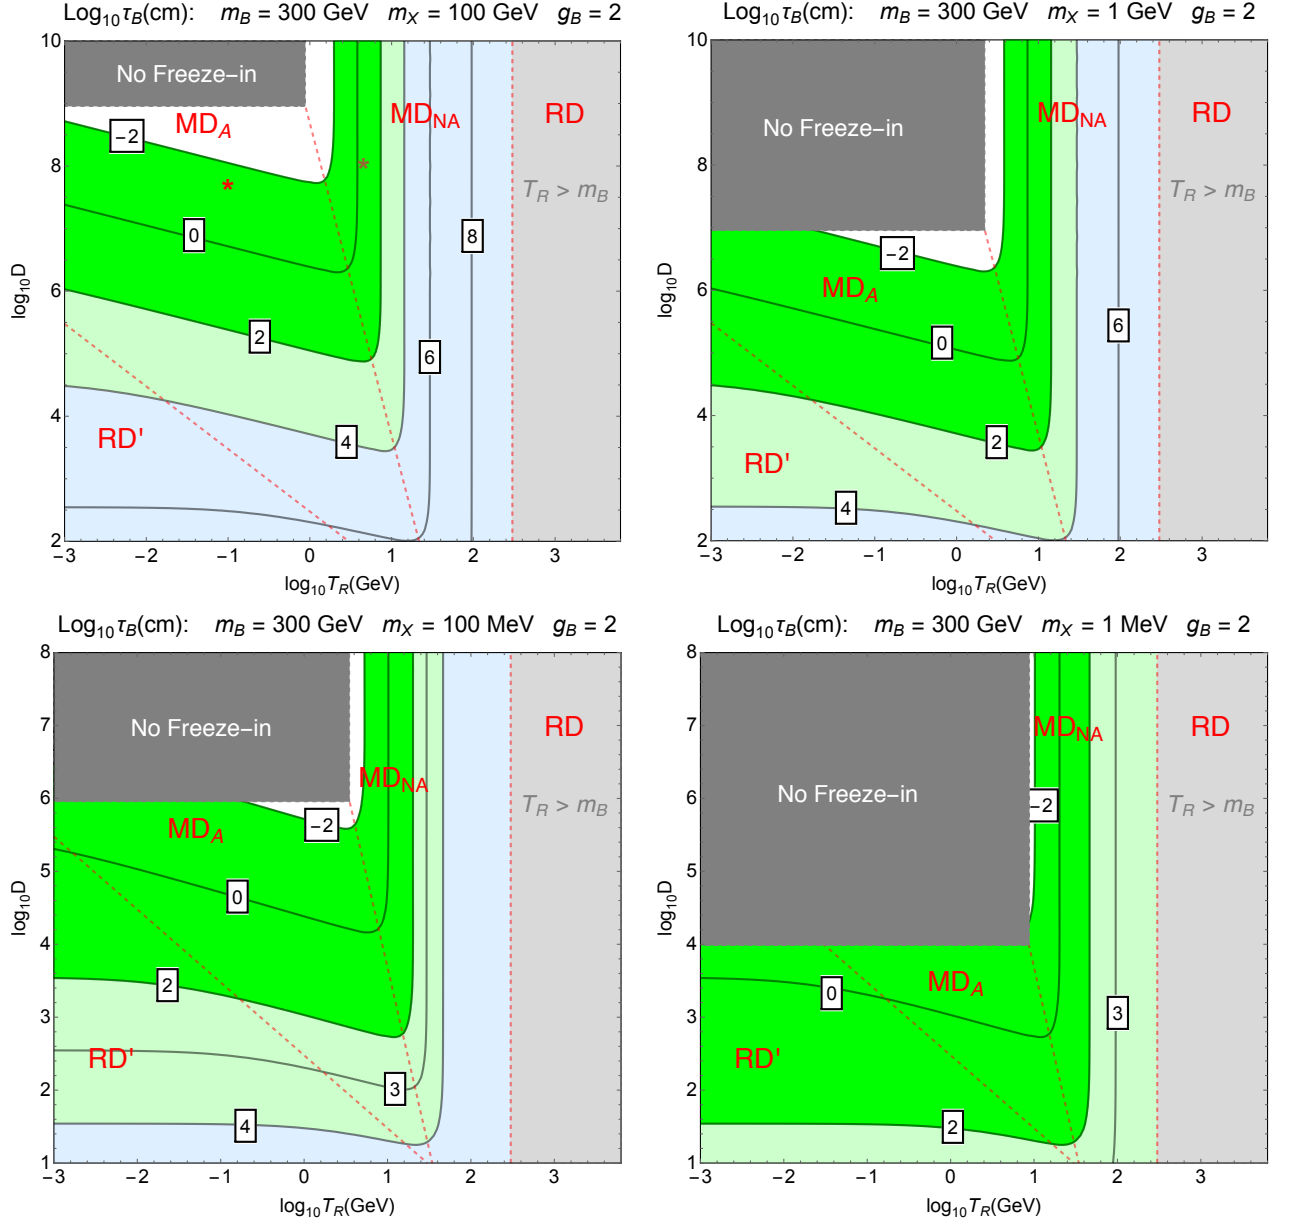


Fig. 27: Contours of the parent particle lifetime τ_B (in cm) that give the observed DM abundance via FI, assuming that some mechanism dumps entropy into the SM sector at temperature T_R to reduce the DM abundance by dilution factor D . We fix $m_B = 300$ GeV and in each panel we consider different values of m_X . The upper-left gray region does not give FI, as X thermalizes. The three dashed red lines separate four regions where FI occurs during (left to right) RD', MD_A, MD_{NA}, and RD eras. Displaced collider signals occur in almost the entire parameter space. Light green and blue shaded regions in particular are prime targets for MATHUSLA. (This figure is taken from [300], which studied the benchmarks marked by the red and brown stars.)

X is produced during the MD_{NA} , MD_A , or RD' era, the decay rate of B must be enhanced in order to compensate for the dilution effect, leading to a shorter decay length ideal for collider searches. Including all $\mathcal{O}(1)$ factors, the decay length of B predicted by our numerical results for production in various eras can be approximated by

$$c\tau_B \simeq 3 \times 10^6 \text{ m} \left(\frac{300 \text{ GeV}}{m_B} \right)^2 \left(\frac{m_X}{100 \text{ GeV}} \right) \quad (\text{RD}) \quad (116)$$

$$c\tau_B \simeq 10 \text{ m} \left(\frac{T_R}{10 \text{ GeV}} \right)^7 \left(\frac{300 \text{ GeV}}{m_B} \right)^9 \left(\frac{m_X}{100 \text{ GeV}} \right) \quad (\text{MD}_{NA}) \quad (117)$$

$$c\tau_B \simeq 1 \text{ m} \left(\frac{T_R}{\text{GeV}} \right) \left(\frac{10^5 \text{ GeV}}{T_M} \right)^{3/2} \left(\frac{300 \text{ GeV}}{m_B} \right)^{3/2} \left(\frac{m_X}{100 \text{ GeV}} \right) \quad (\text{MD}_A) \quad (118)$$

$$c\tau_B \simeq 1 \text{ m} \left(\frac{T_R}{10 \text{ MeV}} \right) \left(\frac{30 \text{ GeV}}{T_M} \right) \left(\frac{3 \text{ TeV}}{m_B} \right)^2 \left(\frac{m_X}{10 \text{ GeV}} \right). \quad (\text{RD}') \quad (119)$$

Figure 27 from Ref. [300] shows the numerical calculations of $c\tau_B$ in the (T_R, D) plane. For illustration, we fix the bath parent particle mass $m_B = 300 \text{ GeV}$ and vary the dark matter mass m_X from 1 MeV to 100 GeV in the four panels. Reheating temperatures above the BBN bound of $\sim \text{MeV}$ are considered. Remarkably, displaced collider signals can occur in almost the entire parameter space. We shade some regions dark green, light green and light blue to indicate proper decay lengths greater than 0.1 mm, 1 m and 100 m. The precise phenomenology depends of course on the LLP production mode, but broadly speaking green shaded regions might be probed at the LHC main detectors, while light shaded regions are prime targets for MATHUSLA.

In the gray region of Fig. 27, the dilution factor is so large that Γ_B required for ρ_{DM}^{obs}/s is already sufficient to thermalize X by scattering processes, and therefore freeze-in process does not occur. The red dashed lines separate the regions by the eras in which freeze-in occurs. In the light-gray RD region, freeze-in happens after the end of matter reheating and therefore the result reduces to the conventional case studied in Sec. 5.3.1.1. The prediction of $c\tau_B$ inside this RD region is the same as that on the left edge of the light-gray region. In the region labeled by MD_{NA} , freeze-in occurs during entropy production so the abundance receives only partial dilution and is insensitive to the total dilution D . As for MD_A , larger dilution D allows for a larger decay rate and thus smaller $c\tau_B$. Besides, larger T_R increases the Hubble rate at the freeze-in temperature, which also leads to smaller $c\tau_B$. Finally, when produced in RD' , DM abundance receives the full dilution factor with the usual Hubble scaling and is thus independent of T_R .

5.3.2 A Simplified Model: Freeze-In through the SM Higgs

We now consider a simplified model of fermionic DM which is populated by the freeze-in mechanism assuming a standard RD cosmology before BBN. This shows that even a minimal implementation of the freeze-in mechanism can give rise to neutral LLP signals at MATHUSLA across large regions of parameter space.

For the family of freeze-in models giving rise to processes like the one in Eq. 105, if A_{SM} is the SM Higgs, one of the simplest implementations is to add on top of the SM a Dirac fermion χ , singlet under the SM gauge group, and an $SU(2)_L$ Dirac doublet ψ :

$$\psi = \begin{pmatrix} \psi^+ \\ \psi^0 \end{pmatrix} \quad (120)$$

such that the Lagrangian reads:

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + i \bar{\chi} \gamma^\mu \partial_\mu \chi + i \bar{\psi} \gamma^\mu D_\mu \psi - m_s \bar{\chi} \chi - m_D \bar{\psi} \psi - y_\chi \bar{\psi} H \chi + h.c. \quad (121)$$

The system we consider could be regarded as a simplified model for a feebly interacting higgsino-bino or higgsino-singlino system (see also [302]) as well as for the higgsino-axino system considered in Sec. 5.3.3 for light axinos in an RD background, bearing in mind that here χ and ψ are Dirac fermions. The coupling y_χ is taken here to be very small (in the correct ballpark for the freeze-in regime, see below). Consequently, upon electroweak symmetry breaking, the neutral particles χ and ψ^0 acquire a tiny mixing, giving rise to mass eigenstates χ_1 (mostly singlet) and χ_2 (mostly doublet), with masses m_1 and m_2 , with $m_2 > m_1$ (assuming $m_D > m_s$ in Eq. 121). The mixing is simply given by

$$\sin\theta \simeq \frac{y_\chi v}{\sqrt{2}(m_2 - m_1)} \quad (122)$$

with $v = 246$ GeV the Higgs vev. The interactions of the DM candidate χ_1 are given by $h\chi_2\chi_1$ (from the Yukawa term in Eq. 121) and $Z\chi_2\chi_1$, $W^\pm\psi^\mp\chi_1$ (from the singlet-doublet mixing), of comparable strength. When kinematically possible, the decay widths for $\chi_2 \rightarrow h\chi_1$, $\chi_2 \rightarrow Z\chi_1$ and $\psi^\pm \rightarrow W^\pm\chi_1$ are given by

$$\Gamma(\chi_2 \rightarrow h\chi_1) = \frac{y_\chi^2}{32\pi m_2^3} [(m_2 + m_1)^2 - m_h^2] \lambda(m_2, m_1, m_h) \quad (123)$$

$$\Gamma(\chi_2 \rightarrow Z\chi_1) = \frac{y_\chi^2}{32\pi} \frac{[(m_2 - m_1)^2 - m_Z^2] [(m_2 + m_1)^2 + 2m_Z^2]}{m_2^3 (m_2 - m_1)^2} \lambda(m_2, m_1, m_Z) \quad (124)$$

$$\Gamma(\psi^\pm \rightarrow W^\pm\chi_1) = \frac{y_\chi^2}{16\pi} \frac{[(m_\psi - m_1)^2 - m_W^2] [(m_\psi + m_1)^2 + 2m_W^2]}{m_\psi^3 (m_\psi - m_1)^2} \lambda(m_\psi, m_1, m_W) \quad (125)$$

with

$$\lambda(x, y, z) = \sqrt{x^4 + y^4 + z^4 - 2x^2y^2 - 2x^2z^2 - 2y^2z^2} \quad (126)$$

As discussed in Sec. 5.3.1, the DM relic abundance is obtained via slow χ_1 production in the early Universe (during radiation domination, as discussed in Sec. 5.3.1.1) through the decays of the χ_2 and ψ^+ states, which are in equilibrium with the thermal bath, and the subsequent DM freeze-in when the abundance of χ_2 and ψ^+ becomes exponentially suppressed, around³⁴ $T \sim m_2/3$. Assuming that the reheating temperature $T_R \gg m_2$, the DM relic abundance can be estimated as:

$$\Omega_{\text{DM}} h^2 \simeq \frac{2m_1}{\rho_c/s_0} \frac{45 M_{\text{Pl}} \Gamma_{\text{FI}}}{4\pi^4 m_2^2} \int_0^\infty \frac{K_1(x)x^3}{[g_*(m_2/x)]^{3/2}} dx \simeq \frac{m_1}{\rho_c/s_0} \frac{135 M_{\text{Pl}} \Gamma_{\text{FI}}}{2\pi^3 m_2^2 [g_*(m_2/3)]^{3/2}} \quad (127)$$

with $\Gamma_{\text{FI}} = \Gamma(\chi_2 \rightarrow h\chi_1) + \Gamma(\chi_2 \rightarrow Z\chi_1) + \Gamma(\psi^\pm \rightarrow W^\pm\chi_1)$, $M_{\text{Pl}} = 10^{19}$ GeV the Planck mass, $\rho_c/s_0 = 3.6 \times 10^{-9}$ GeV the critical energy density over the entropy density today, and $g_*(T)$ the number of relativistic degrees of freedom in the early Universe at temperature T . We note an extra factor 2 in Eq. 127 due to the fact that both χ_2 , ψ^+ and their antiparticles $\bar{\chi}_2$, ψ^- are present in the early Universe plasma and contribute to the relic abundance through their decays. Demanding $\Omega_{\text{DM}} h^2 = 0.12$ fixes y_χ in terms of m_1 and m_2 .

5.3.2.1 Production of Dark Matter χ_1 at the LHC: MATHUSLA Sensitivity Estimate

The states χ_2 and ψ^\pm can be produced at the LHC via the Drell-Yan processes $pp \rightarrow \chi_2\chi_2$, $pp \rightarrow \chi_2\psi^\pm$, $pp \rightarrow \psi^+\psi^-$. The state ψ^\pm is short-lived and dominantly decays to $\chi_2\pi^\pm$, due to the electromagnetically induced radiative mass splitting $m_\psi - m_2 = 341$ MeV (see e.g. [303]). The short lifetime and very soft pion in the final state make direct detection of ψ^\pm very challenging at the LHC main detectors [159, 304, 305]. In contrast, the neutral state χ_2 is very long-lived: combining the freeze-in DM relic abundance

³⁴For the obtention of the DM relic abundance we use for simplicity $m_2 = m_\psi$, as their mass difference $m_\psi - m_2 = \mathcal{O}(100 \text{ MeV})$ does not play a role in the freeze-in mechanism.

condition, Eq. 127, with Eqs. 123-124 in the limit $m_2 \gg m_1, m_W, m_Z, m_h$ yields for the decay length $c\tau$ of χ_2

$$c\tau \simeq 3500 \text{ m} \left(\frac{m_1}{100 \text{ MeV}} \right) \left(\frac{500 \text{ GeV}}{m_2} \right)^2. \quad (128)$$

Despite the very large decay length of χ_2 , we show in the following that MATHUSLA can be sensitive to a wide range of freeze-in DM masses. We implement our model in FEYNRULES [306] and simulate the various Drell-Yan production processes for χ_2 and ψ^\pm (with subsequent decay into χ_2) at LHC 13 TeV in MADGRAPH_AMC@NLO [101], choosing to normalize the respective cross-sections to the corresponding NLO/NLL charged/neutral higgsino production cross-sections at 13 TeV LHC given by the CERN LHC SUSY XS Working Group [104]. The probability for an LLP χ_2 to decay inside MATHUSLA is computed directly via a convolution of the MATHUSLA detector geometry with the distribution of produced LLPs.




Fig. 28: Left: LLP MATHUSLA sensitivity in the $(m_2, c\tau)$ plane. Lines yielding the observed DM relic density are shown for $m_1 = 1 \text{ GeV}$ (black), 10 MeV (dark blue), 100 keV (light blue), with MICROMEGAS (solid) and through the analytic approximation Eq. 127 (dashed). Right: MATHUSLA sensitivity in the (m_2, m_1) plane. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

In Fig 28 (left) we show the MATHUSLA LLP “exclusion” (4 event reach) and “discovery” (10 event reach) sensitivity in the $(m_2, c\tau)$ plane for $\mathcal{L} = 3000 \text{ fb}^{-1}$ of integrated luminosity from HL-LHC, under the assumption of perfect MATHUSLA object reconstruction efficiency. We also show the curves in this plane yielding the observed DM relic abundance for $m_1 = 1 \text{ GeV}$ (black), $m_1 = 10 \text{ MeV}$ (dark blue) and $m_1 = 100 \text{ keV}$ (light blue), obtained both numerically using MICROMEGAS [307] (solid) and through the analytic approximation Eq. 127 (dashed), which are observed to agree to better than 20%. In Fig 28 (right) we directly show the MATHUSLA 4 and 10 event reach in the (m_2, m_1) plane, highlighting that MATHUSLA is sensitive to a large range of freeze-in DM masses for LLP masses below the TeV scale. As discussed in Section 4.1.5, which examines the same signature as it arises from higgsinos decaying to axinos, this MATHUSLA sensitivity is at least 1-2 orders of magnitude better in cross-section than the corresponding HL-LHC LLP search, especially for parent particle masses below a few hundred GeV.

5.3.3 A Complete Model: DFSZ axino

We now discuss a motivated freeze-in DM candidate that can arise as part of a more complete model with a variety of possible cosmological histories: the DFSZ axino. As we show below, this generates

displaced higgsino signals at colliders and at MATHUSLA.

DFSZ axion models are theories based on a Peccei-Quinn (PQ) solution of the strong CP problem, where the Higgs doublets carry a PQ charge. If the theory is made supersymmetric, the axion is promoted to a supermultiplet and appears with its partners: the saxion (CP-even scalar) and axino (Weyl fermion).

Stellar cooling bounds [203, 204] severely constrain the PQ scale to be high, $V_{\text{PQ}} \gtrsim 10^9 \text{ GeV}$. Thus we have a hierarchy between such a scale and the masses of the supersymmetric partners, which we consider around the scale $\tilde{m} \simeq \text{TeV}$. We take advantage of this hierarchy between scales to write down an effective field theory (EFT) where the PQ symmetry is non-linearly realized. The axion supermultiplet in the language of superfields reads

$$A = \frac{s + ia}{\sqrt{2}} + \sqrt{2}\theta\tilde{a} + \theta^2 F . \quad (129)$$

If we perform a PQ rotation with angle α , the axion superfield shifts

$$A \rightarrow A + i\alpha V_{\text{PQ}} . \quad (130)$$

Within this language, the QCD anomaly interaction is

$$\mathcal{L}_{AWW} = -\frac{g_3^2 N_{\text{DW}}}{32\pi^2 V_{\text{PQ}}} \int d^2\theta A W^\alpha W_\alpha + \text{h.c.} , \quad (131)$$

where N_{DW} is the color anomaly coefficient, also known as the domain wall number. (The suppression scale of this term is also often written as $f_a = V_{\text{PQ}}/N_{\text{DW}}$.)

We write down the most general set of interactions for the axion superfield, with all operators respecting the shift symmetry in Eqn. 130. Supersymmetric interactions include the Kähler potential and the superpotential

$$K = A^\dagger A + \frac{\kappa}{2V_{\text{PQ}}} A^\dagger A (A + A^\dagger) + \dots , \quad (132)$$

$$W = \mu H_u H_d + q_\mu \frac{\mu}{V_{\text{PQ}}} A H_u H_d + \dots . \quad (133)$$

The dimensionless quantity κ is a model-dependent coefficient, that depends on the charges and vacuum expectation values of the PQ breaking fields. For a single PQ-breaking field, we have $\kappa = 1$. The PQ charge of the μ term q_μ is also model-dependent, and in the minimal supersymmetric DFSZ theory it is $q_\mu = 2$. The renormalizable cubic coupling between the axion and the Higgs bosons is responsible for axino freeze-in production.

5.3.3.1 Axino Dark Matter in a RD Background

We review the calculation of the axino freeze-in for a standard cosmology. For the purpose of this illustration, we consider higgsinos mass eigenstates, but they need not be the lightest observable supersymmetric particles (LOSP). The final axino DM density can be derived from Eqn. 112, after we replace $X \rightarrow \tilde{a}$ and $B \rightarrow \tilde{h}$.

Accounting for both charged and neutral higgsinos leads to $g_{\tilde{h}} = 8$. The only missing information is the higgsino decay width. We can compute it from the cubic superpotential interaction in Eqn. 133, and we find

$$\Gamma_{\tilde{h}} = \frac{q_\mu^2 \mu^3}{32\pi V_{\text{PQ}}^2} . \quad (134)$$

corresponding to a higgsino lifetime of

$$c\tau_{\tilde{h}} = \Gamma_{\tilde{h}}^{-1} \simeq 180 \text{ m} \left(\frac{2}{q_\mu} \right)^2 \left(\frac{300 \text{ GeV}}{\mu} \right)^3 \left(\frac{V_{\text{PQ}}}{10^{12} \text{ GeV}} \right)^2 . \quad (135)$$

The resulting axino relic abundance in units of the observed density reads

$$\frac{\rho_{\tilde{a}}}{\rho_{DM}^{\text{obs}}} = 8 \times 10^4 \left(\frac{106.75}{g_*} \right)^{3/2} \left(\frac{q\mu}{2} \right)^2 \left(\frac{m_{\tilde{a}}}{100 \text{ GeV}} \right) \left(\frac{\mu}{300 \text{ GeV}} \right) \left(\frac{10^{12} \text{ GeV}}{V_{PQ}} \right)^2. \quad (136)$$

Without additional dilution mechanisms or a small initial misalignment angle, coherent oscillations of the axion field at a temperature near $T = 160 \text{ MeV}$ will overproduce axions unless $V_{PQ} \lesssim 10^{12} \text{ GeV}$ [205]. Imposing this constraint on the PQ scale has two consequences.

1. Eqn. (135) makes clear that the Higgsino lifetime lies in the range that is optimal for detection at MATHUSLA. The general collider phenomenology of Higgsino LLPs decaying into axions is studied in Sec. 4.1.5. As shown in Fig. 21, MATHUSLA can probe V_{PQ} up to $\sim 10^{13} \text{ GeV}$.
2. Reproducing the observed DM abundance in Eqn. (136) requires very low axino masses at or below the MeV scale. This essentially realizes the simplified scenario of Section 4.1.5.

Light axino freeze-in DM therefore represents an excellent target for LLP detection at MATHUSLA. There is, however, significant motivation to also consider heavier axinos with masses near the weak scale.

The axino mass is expected to be at least of order the gravitino mass [193]. This is due to non-renormalizable couplings between the PQ sector and the SUSY breaking sector that cannot be forbidden by symmetries (though certain extra-dimensional sequestering scenarios could change this argument). A light axino therefore implies a light gravitino, leading to the usual gravitino overclosure problem. One way to address this problem is a low inflationary reheating temperature $\lesssim 10^5 \text{ GeV}$, avoiding high-temperature overproduction of gravitinos. In that case, light axinos (or similar-mass gravitinos produced in axino decays) produced via freeze-in from higgsino decay would be a viable dark matter candidate.

While low reheating temperatures are a valid solution, they restrict the possible inflationary scenario and preclude any high-scale baryogenesis mechanism from accounting for the matter-antimatter asymmetry of the universe. Hence there is ample motivation for considering weak scale axinos to solve the gravitino overclosure problem. Eqn. (134) makes clear that this requires some dilution mechanism to reduce the axino relic density. As discussed in Section 5.3.1.2, this can be realized if axino freeze-in occurs during an early MD epoch. Remarkably, supersymmetric PQ theories incorporate such a dilution mechanism through the saxion condensate.

5.3.3.2 Saxion Cosmology

The CP-even scalar of the axion supermultiplet has no potential in the absence of SUSY breaking. In the vacuum today, PQ symmetry is broken, and the saxion takes a mass of the order of the superpartners scale. PQ may also have been broken during inflation and not restored afterwards, conveniently solving the domain wall problem for the minimal DFSZ model ($N_{DW} = 6$). If this is the case, the SUSY breaking vacuum energy during inflation would also provide a potential for the saxion. Unless theories with specific symmetries are considered [300], the minimum of the saxion potential today and during inflation are different. Inflation sets the initial condition for the evolution of the saxion field, displacing it from its current minimum by an amount $s_I \simeq V_{PQ}$ or $s_I \simeq M_*$, where the latter is the cutoff of the theory.

The subsequent evolution of the saxion condensate can be tracked by solving the equation of motion. Right after inflation ends, the saxion field does not evolve due to the Hubble friction. Once the universe slows down enough, the saxion condensate starts harmonic damped oscillations. This happens at a temperature T_{osc} found by solving the condition $3H(T_{\text{osc}}) \simeq m_s$, and it approximately reads

$$T_{\text{osc}} \simeq (m_s M_{\text{Pl}})^{1/2} \simeq 10^{10} \text{ GeV} \left(\frac{m_s}{1 \text{ TeV}} \right)^{1/2}. \quad (137)$$

The damped saxion oscillations at lower temperature red-shift with the expansion as non-relativistic matter. As this red-shifting is milder than the one for radiation, at some temperature T_M the saxion condensate energy dominates the universe:

$$T_M \simeq 10 \text{ TeV} \left(\frac{m_s}{1 \text{ TeV}} \right)^{1/2} \left(\frac{s_I}{10^{15} \text{ GeV}} \right)^2. \quad (138)$$

At temperatures below T_M the universe enters an early MD epoch where the saxion energy density controls the Hubble expansion. This epoch has to be terminated by the saxion condensate decay before the time of BBN. The possible saxion decay channels can be identified from the interactions in Eqs. 131-133. The QCD anomaly term mediates interactions to gluon final states,

$$\Gamma_{s \rightarrow gg} = N_{\text{DW}}^2 \frac{\alpha_3^2 m_s^3}{64\pi^3 V_{\text{PQ}}^2}, \quad (139)$$

while the cubic term in the superpotential induces decays to Higgs bosons and longitudinal electroweak gauge bosons,

$$\Gamma_{s \rightarrow hh, W_L W_L, Z_L Z_L} = \frac{q_\mu^2 \mu^4}{4\pi m_s V_{\text{PQ}}^2}. \quad (140)$$

Decay to EW bosons typically dominates over the loop-suppressed decay to two gluons.

In order for the condensate to dilute the axino abundance, saxion decay to axions and axinos must be small. These dangerous decay channels are generated by the model-dependent cubic self-interaction in the Kähler potential:

$$\Gamma_{s \rightarrow aa} = \frac{\kappa^2 m_s^3}{64\pi V_{\text{PQ}}^2}, \quad (141)$$

$$\Gamma_{s \rightarrow \tilde{a}\tilde{a}} = \frac{\kappa^2 m_a^2 m_s}{8\pi V_{\text{PQ}}^2}. \quad (142)$$

The second decay badly overproduces LSP dark matter, and we forbid it by assuming the saxion mass to be below twice the axino mass. The first decay produces axion dark radiation, which is severely constrained by BBN and CMB bounds. The corresponding largest allowed value of κ is [308]

$$\kappa < \kappa_{\text{max}} \approx 2.1 \left(\frac{q_\mu}{2} \right) \left(\frac{\mu^2}{m_s^2} \right) \left(\frac{\Delta N_{\text{eff}}}{0.45} \right)^{1/2}. \quad (143)$$

Given the current bound of $\Delta N_{\text{eff}} \lesssim 0.45$ [284], this constraint can be easily satisfied if $\kappa = 1$ for a single PQ-breaking field. $\kappa < \kappa_{\text{max}}$ then also guarantees that saxions decay dominantly to EW bosons.

Since the saxion decay width is dominated by Eqn. 140, we can use this expression to evaluate T_{Rs} , defined as the reheat temperature of the radiation bath after the saxion condensate has decayed:

$$T_{Rs} \simeq (\Gamma_s M_{\text{Pl}})^{1/2} \simeq 10 \text{ MeV} \left(\frac{\mu}{1 \text{ TeV}} \right)^{3/2} \left(\frac{\mu}{m_s} \right)^{1/2} \left(\frac{10^{15} \text{ GeV}}{V_{\text{PQ}}} \right). \quad (144)$$

For TeV scale SUSY parameters, the reheat temperature is in the MeV – GeV range. This range is interesting, since it is below the axino freeze-in production temperature while still being consistent with BBN bounds. Furthermore, for reheating temperatures below 160 MeV, any initial axion abundance produced from misalignment is diluted away, allowing V_{PQ} to be much larger than the usual limit of 10^{12} GeV [308].

As a result of the saxion early matter-dominated epoch, the axino freeze-in calculation presented in Sec. 5.3.3.1 has to be revisited. There are two reasons why the final result will be different: the different




Fig. 29: The axino yield from neutralino decays to SM + \tilde{a} (red) for $m_{\tilde{a}} \ll 1$ TeV, and neutralino inverse decays ($\tilde{\chi}^0 + \text{SM} \rightarrow \tilde{a}$, orange) for $m_{\tilde{a}} = 2$ and 20 TeV. In both panels, $2M_1 = M_2 = \mu = 1$ TeV, $m_s = 500$ GeV, $\tan\beta = 2$, $q_\mu = 2$, and $\mathcal{D} = 4$; while $s_I = V_{PQ}$ ($M_* \gtrsim 10^{17}$ GeV) for the left (right) panel. Figure taken from [5].

time vs. temperature dependence for an early MD era and the dilution due to the entropy dumped into the radiation bath from saxion decays. We present the results for the axino relic density in Sec. 5.3.3.3. Before we do that, we quantify the entropy due to saxion decays. As described in Sec. 5.3.1.2, the saxion MD epoch is made of two distinct phases. The temperature when we enter a non-adiabatic phase for the saxion condensate reads

$$T_{\text{NA}} \simeq (T_M T_{R_s}^4)^{1/5} \simeq 0.2 \text{ GeV} \left(\frac{\mu}{1 \text{ TeV}} \right)^{13/10} \left(\frac{\mu}{m_s} \right)^{3/10} \left(\frac{s_I}{V_{PQ}} \right)^{2/5} \left(\frac{10^{15} \text{ GeV}}{V_{PQ}} \right). \quad (145)$$

Axino produced through freeze-in at temperatures above T_{NA} get diluted by the full amount of entropy injected ($D \simeq T_M/T_{R_s}$). If the production happens at temperatures T_{FI} below T_{NA} , the dilution factor is only $D(T_{\text{FI}}) \simeq (T_{\text{FI}}/T_{R_s})^5$.

5.3.3.3 Freeze-In Axino Yield with Dilution

The final freeze-in axino abundance after saxion dilution can be obtained by using the general yield Y_i in Eqn. 115 with the saxion reheat temperature T_{R_s} in Eqn. 144 and T_{NA} in Eqn. 145. The coupling constant λ defined in Eqn. 114, in this case, is given by $q_\mu^2 \mu^2 / 32\pi V_{PQ}^2$ based on Eqn. 134.

The numerical result of the freeze-in axino yield is shown in Fig. 29. Though our focus in the context of LLP signals is freeze-in by decay, this figure was taken from [5] which also considered freeze-in by inverse decay ($\tilde{\chi}^0 + \text{SM} \rightarrow \tilde{a}$). The general features of this plot can be understood as follows. Freeze-in by inverse decay is less efficient than by neutralino decay, because in the former case the axino mass is higher, and the inverse decay process stops at $T_{\text{FI}} \sim m_{\tilde{a}}$, which is earlier than $T_{\text{FI}} \sim \mu$ for decay. In the left panel with $s_I = V_{PQ}$, the saxion condensate becomes insufficient to dominate the energy of the Universe when $V_{PQ} \lesssim 10^{13}$ GeV and therefore the result is identical to that of conventional RD cosmology. The effect of saxion dilution at higher V_{PQ} can be seen from the change of the slope at $V_{PQ} \simeq 10^{13}$ GeV. In the right panel, where $s_I \gtrsim 10^{17}$ GeV, freeze-in occurs during the MD_{NA} era so

the abundance becomes insensitive to s_I , as can be seen from Eqn. 115. The decay of a larger saxion condensate results in a much more severe dilution of the final yield than in the left panel.

Given that $m_{\tilde{a}} Y_{\tilde{a}} \simeq 0.44$ eV for the observed DM abundance, the freeze-in by decay production of axino DM heavier than \sim keV requires V_{PQ} in the range of $10^{10} - 10^{13}$ GeV. This corresponds to lifetimes in the ideal range for higgsino LLPs to be detected as displaced vertices at either the LHC main detectors or MATHUSLA.

It is worth briefly commenting on the cosmological interplay between gravitinos and axinos, especially since the gravitino problem was a major motivation to consider weak-scale axino masses. Since they are of similar mass, one can decay into the other via emission of an axion. Therefore, if the gravitino is lighter than the axino, the freeze-in mechanism effectively generates gravitino dark matter (with only $\mathcal{O}(1)$ modifications to the quantitative yields discussed here). Apart from allowing freeze-in weak-scale axino production, the saxion condensate has the additional feature that it allows for inflationary reheating temperatures in excess of 10^{12} GeV, at which even TeV-scale gravitinos would be overproduced. In that scenario, V_{PQ} would be fixed by the dilution of the thermal gravitino abundance required to reproduce the observed DM relic density, giving another motivation for higgsino LLPs. More details can be found in [5, 301].

5.3.3.4 General Dilution Mechanisms

Above, we used the specific dilution mechanism of the decaying saxion condensate to argue that freeze-in weak-scale axino dark matter, produced in Higgsino decays, motivates PQ-breaking scales in the range of $V_{PQ} \sim 10^{10} - 10^{13}$ GeV, resulting in observable displaced vertex signals. It is important to emphasize that similar V_{PQ} ranges, and hence higgsino LLP signals, are not limited to saxion dilution. In fact, any scalar condensate can provide the necessary dilution for axino dark matter from freeze-in production.

In the generic dilution case, once the dilution temperature T_R (corresponding to T_{R_s} in the saxion condensate case) and dilution factor D are computed for the scalar condensate, the PQ scale and hence the higgsino decay length are determined from the axino dark matter abundance. The resulting values of $f_a = V_{PQ}/N_{DW}$ are shown as solid contours in Fig. (30) for four different choices of higgsino and axino masses.

In any axino scenario, the possible contribution to the DM energy density from axions must also be considered. In the pre-inflationary scenario, PQ symmetry is broken above the energy scale of inflation and axions are produced via coherent oscillation of the axion field around $T = 160$ MeV. The axion relic density therefore depends on the initial misalignment angle θ_i , as well as f_a . The dash-dotted lines in Fig. (30) indicate where the axion relic density is equal to the observed DM abundance for a given θ_i . Note that they trace the f_a contours for $T_R > 160$ MeV, while for $T_R < 160$ MeV the axion contribution to the energy density is practically eliminated by the dilution. Similarly, the dashed line corresponds to the post-inflationary scenario, where PQ symmetry is broken below the energy scale of inflation. In most of the parameter space we show, the axion energy density can be negligible or subdominant to the axino freeze-in contribution. Observing the LLP decay into axinos can therefore provide direct information about the origin of dark matter.

The higgsino lifetime is indicated by the color shadings in Fig. (30), same as in Fig. (27): dark green, light green and light blue to indicate proper decay lengths greater than 0.1mm, 1m and 100m. The majority of the relevant parameter space corresponds to $V_{PQ} \lesssim 10^{13}$ GeV and lifetimes observable at MATHUSLA or the LHC main detectors.

5.3.3.5 Upshot: Long-Lived Higgsinos at MATHUSLA

Axinos can make up the observed DM abundance if they are produced via freeze-in from Higgsino decay. For a standard RD cosmology before BBN, the axino must have a mass around or below the MeV scale, generically calling for low inflationary reheating temperatures to avoid the gravitino problem. If




Fig. 30: Black solid contours: $f = V_{PQ}/N_{DW}$ required for freeze-in axinos from Higgsino decay to have the observed DM relic density for different values of μ and $m_{\tilde{a}}$, assuming that some mechanism dumps entropy into the SM sector at temperature T_R to reduce the axino abundance by dilution factor D . Axions could still be produced and constitute part of dark matter, but above and to the left of the dashed/dot-dashed lines, axions have less than the observed DM relic density for different effective misalignment angles, see text for details. If reheating takes place below $T = 160$ MeV, axions are diluted away. The yellow region is excluded by white dwarf constraints. The Higgsino lifetime in the light green and blue shaded regions is a prime targets for MATHUSLA, see Fig. 27. Figure taken from [300].

some dilution mechanism is present, whether the saxion or a more general scalar condensate, weak-scale axino masses and hence higher reheating temperatures are possible. As the arguments of the previous subsections show, all of these scenarios call for a PQ-breaking scale in the range $V_{PQ} \sim 10^9 - 10^{13}$ GeV. For Higgsino masses accessible at LHC energies, this leads to a wide range of possible lifetimes, from μm to 10^5 m.

The Higgsino decays to an invisible axino and a Higgs or EW gauge boson which in turn decays promptly to visible particles. At the lower end of the motivated lifetime range, the LHC main detectors will therefore be able to constrain this scenario with LLP searches. For larger lifetimes above ~ 10 m, MATHUSLA will be at least 1-2 orders of magnitude more sensitive than the main detectors due to small LLP branching ratios to leptons and non-negligible backgrounds for hadronically decaying LLPs with less than a few hundred GeV of visible energy, see Sec. 3.2.3.

Quantitative predictions for MATHUSLA’s reach depend on the LLP production mode. higgsinos can be produced at the LHC either through Drell-Yan processes, or via cascade decay of heavy colored particles like gluinos. By far the most pessimistic assumption is that only direct Drell-Yan production is present. In that case, we can refer to Sec. 4.1.5 which studied the general collider phenomenology of Higgsino LLPs decaying to axions. As shown in Fig. 21, MATHUSLA can probe V_{PQ} up to $\sim 10^{13}$ GeV even for TeV-scale higgsinos, putting almost the entire parameter space of freeze-in axino DM within our reach.

5.4 SIMPs, ELDERS and Co-Decay³⁵

In the WIMP paradigm, the dark matter relic abundance is entirely set by $2 \rightarrow 2$ annihilations of the dark matter, with no dependence on dark matter self-interactions or on its decay. Here we review three classes of dark matter candidates which differ from this vastly: Strongly Interacting Massive Particles (SIMPs), where the relic density is determined by the $3 \rightarrow 2$ annihilation rate within the dark sector; ELastically DEcoupling Relics (ELDERS), where the relic density is determined by the elastic scattering cross-section with the SM; and Co-Decaying dark matter, where the relic density is determined by strong interactions between the dark matter and other dark sector particles.

5.4.1 Strongly Interacting Massive Particles (SIMPs)

In the SIMP mechanism [309] (see also [310–312]), dark matter self-interactions play a crucial role in setting its relic density. Here the dark matter abundance is set by the freeze-out of number-changing self-annihilations, typically of $3 \rightarrow 2$ depletion processes where three dark matter particles collide and annihilate into two dark matter particles. In contrast to ‘the WIMP miracle’, which predicts weak-scale masses for weak couplings, such $3 \rightarrow 2$ processes point to strong scale masses for strong couplings, hence dubbed ‘the SIMP miracle’.

The self-annihilation process pumps heat into the system, and so the dark matter must be in thermal contact with the Standard Model bath in order to cool (or dump its entropy into other light degrees of freedom). Such interactions between the dark and visible sectors imply measurable signals that should be observable in a variety of upcoming experiments, including direct-detection, indirect-detection, and direct production at colliders. Moreover, the dark matter’s strong $3 \rightarrow 2$ self-interactions typically predict sizable contributions to $2 \rightarrow 2$ self-scattering processes which naturally address long-standing puzzles in structure formation. The SIMP setup robustly predicts light dark matter (typically in the MeV to GeV mass range), with strong interactions with itself and weak interactions with ordinary matter.

The SIMP mechanism can be realized in various different ways. In Ref. [310] it was found that SIMP dark matter emerges in generic classes of strongly coupled gauge theories. These are theories of dynamical symmetry breaking—resembling QCD—in which the pions play the role of dark matter, with the Wess-Zumino-Witten term generating the $3 \rightarrow 2$ interactions. $Sp(N)$, $SU(N)$ and $SO(N)$

³⁵Jeff Asaf Dror, Yonit Hochberg, Eric Kuflik

gauge theories are all viable, provided the number of confining quarks is large enough ($N_f \geq 2, 3, 3$ respectively). Dark matter then has the same type of mass, with the same type of interactions, in the same type of theory, as the strong force that binds nuclei together.

There are many potential ways to mediate the requisite thermalizing interactions between the dark matter and the Standard Model. An attractive possibility is the vector portal [313], where a kinetically mixed hidden photon (see Sec. 8.5) communicates between the two sectors. The kinetically mixed mediator can be embedded in the symmetry structure of the theory by identifying the appropriate $U(1)$ subgroup of the residual global symmetries to be gauged, leading to a calculable and predictive framework of SIMP dark sectors. Detailed constraints and prospects on the relevant parameter space can be found in Ref. [313].

From a phenomenological point of view, SIMP DM can therefore be seen as a strong theoretical motivation for the existence of a hidden valley that can be produced at the LHC through various portals. For example, a the kinetically mixed dark photon can be produced and decay to hidden quarks, which will shower and fragment into dark mesons. Depending on the masses and representations of the dark photon, the pions π and the ρ 's, some dark mesons will decay entirely visibly into leptons and/or hadrons, entirely invisibly, or via off-shell hidden photons to $\pi + \ell^+ \ell^-$ or $\pi + jj$. A host of signatures is thus expected, including a mix of missing energy and ρ -decays into narrowly collimated small invariant mass lepton pairs ('lepton jets') or jets. It is generic for some or all of these dark sector particles to be long-lived. The dark hadron decay width scales as

$$\Gamma \sim \frac{\alpha_D \alpha \epsilon^2 m_D^5}{18\pi m_V^4} \quad (146)$$

where α_D is the coupling of dark fermions to the dark photon, m_D is a dark hadron mass scale, ϵ and m_V are the kinetic mixing and mass of the dark photon. As discussed in Sec. 8.5, the dark photon mass and kinetic mixing could easily allow for copious production of dark photons at the LHC, while the dark sector parameters lead to some or all of the dark hadrons having suppressed decay widths to the SM, giving rise to LLP signatures at the LHC and MATHUSLA. (See also Sec. 8.1 for related signatures) Following the discussion in Sec. 3.2.3, MATHUSLA will have much better sensitivity than the main detectors if the dark hadrons decay hadronically or have GeV mass or below.

5.4.2 *ELastically DEcoupling Relics (ELDERs)*

Until recently, in all known examples of proposed dark matter frameworks, the relic abundance was determined by processes that change the dark matter number density. A novel alternative is the ELastically DEcoupling Relic proposal [314], in which the dark matter relic density is determined almost exclusively by the decoupling of the elastic scattering off Standard Model particles—a process that does not change dark matter number density. As was the case for WIMPs and SIMPs, couplings to the Standard Model are a necessary of the mechanism, and lead to observable predictions in a host of different experimental frontiers.

Much as in the SIMP scenario, $3 \rightarrow 2$ self-interactions of dark matter are required in order for the ELDER mechanism to be viable. As the temperature of the universe drops below the DM mass, $3 \rightarrow 2$ annihilation depletes DM number density, while elastic scattering to the visible sector dumps entropy into the SM bath, reducing energy density in the dark sector. In the SIMP scenario, the $3 \rightarrow 2$ annihilation freezes out while the dark sector is still in thermal equilibrium with the SM. The annihilation rate therefore determines the DM relic density. In the ELDER scenario, on the other hand, the elastic scattering between the two sectors freezes out before the $3 \rightarrow 2$ annihilation. The dark sector therefore enters a period of cannibalization after kinetic decoupling. Since it is in chemical equilibrium and its comoving entropy is conserved, the comoving dark matter density redshifts logarithmically with the expansion of the Universe. Therefore the DM density today is determined by the density at kinetic decoupling, which in turn depends on the strength of the elastic scattering interaction instead of the

$3 \rightarrow 2$ annihilation rate. This opens up different regions of parameter space for strongly coupled hidden sectors to produce viable dark matter candidates. Like SIMPs, ELDERs therefore represent another motivation for hidden valley type LLP signatures at both the LHC main detectors and MATHUSLA, as discussed above.

5.4.3 Co-decaying Dark Matter

Dark sector self-interactions, *i.e.*, strong interactions between the dark matter and other hidden sector particles, can also play a role in determining the dark matter relic abundance. Even though these dark sector interactions may not be directly observable in the lab, it can still be the case that the hidden sector interactions with the Standard Model play a critical role in determining the relic abundance, and are observable in current and future experiments. One example of dark matter freezeout where this is true is Co-decaying dark matter [315], which is generic in hidden sectors with approximate degeneracies between one or more LLPs with a stable particle. (See also Ref. [316, 317] for related ideas.) Crucial to the Co-decaying DM mechanism is that the dark sector includes an LLP. Either the dark sector LLP, or a visible sector LLP that is long-lived due to the same small portal, could then be produced at the LHC and observed in MATHUSLA.

Hidden sectors often have accidental symmetries as a residual of the symmetry used to keep the dark matter stable. This can lead to degenerate low-lying states with some of these particles remaining unstable. The decay of the unstable particles can efficiently deplete the dark sector, in analogy with co-annihilation (Section 5.1), leaving behind a dark matter candidate with the observed relic abundance.

As a concrete example, consider a system of two nearly-degenerate states A, B of mass m . A is the stable DM candidate, while B is meta-stable and decays to the SM with lifetime τ_B at temperature T_γ . The process $AA \leftrightarrow BB$ is active in the plasma with some cross-section σ that freezes out at temperature T_σ . Finally, there is a small interaction between the dark and visible sectors that keeps them in thermal equilibrium until it freezes out at temperature T_d . The Co-decaying DM mechanism corresponds to the regime $T_d > m > T_\gamma > T_\sigma$. In other words, as the universe cools, the hidden sector thermally decouples from the SM at T_d . A and B stay in equilibrium with each other as they cool and become non-relativistic. B starts decaying out-of-equilibrium into the SM at temperature T_γ , which depletes the dark sector energy density until A and B decouple at T_σ . The surviving relic density of A makes up DM today.

The decay of the LLPs play the role of the Boltzmann suppression in depleting the dark sector. This framework neatly evades the stringent bounds from direct-detection due to the small couplings to the Standard Model particles. Furthermore, the temperature where freeze-out occurs can be delayed by many orders of magnitude compared to WIMP dark matter, which leads to a relatively smaller relic density for a given cross-section. This results in a robust prediction of an enhanced indirect-detection signal that can be relevant even for dark matter masses above the TeV scale.

The parameter space of Co-decaying dark matter spans many orders of magnitude in both mass and lifetimes of the unstable particle(s), as illustrated in Fig. 31. At each point in the (m, τ_B) plane, σ is chosen to obtain the correct relic density. For simplicity, we work here in the limit where DM cannibalization effects are negligible.³⁶ In the light gray region, the decay of B does not occur out-of-equilibrium, *i.e.* T_γ is too high. The purple region is conservatively excluded by BBN constraints on LLP lifetime (the exact constraint depends on the B decay mode and may be more than an order of magnitude lower in τ_B). The light green shaded region is excluded because the σ required to avoid overclosure violates unitarity bounds.

Co-decaying DM is a very general framework for obtaining the observed DM relic density from dark sector dynamics. As such, a large variety of detailed phenomenologies are possible, depending on the model implementation. Interestingly, the co-decay mechanism can fit naturally in models which

³⁶See Ref. [315] for a detailed discussion on the differences when cannibalization is important.




Fig. 31: White region: the parameter space of co-decaying DM, assuming negligible cannibalization effects, that satisfies the out of thermal equilibrium (gray), unitarity (green) and Big Bang Nucleosynthesis (purple) constraints. Assuming LLP production in 1 TeV gluino decays, the most pessimistic MATHUSLA reach projection is shown as the dark gray shaded region. Details of the complete model implementation generically open up additional regions that MATHUSLA could probe, see text for details. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

address the hierarchy problem, such as supersymmetry or composite Higgs, since these frameworks can yield degenerate particles as the lowest lying particles in a new sector with weak couplings to the Standard Model. Furthermore, the same small portal between the visible and dark sectors that makes B long-lived can also result in a visible sector particle being long-lived (like the lightest parity-odd particle in SUSY) before it decays to the hidden sector. As a result, the co-decaying DM framework generically predicts LLPs production in the decay of heavier BSM particles. For states connected to the solution of the Hierarchy Problem, this leads to QCD-strength LLP production cross-sections at the LHC. In that case, the main detectors and MATHUSLA could probe large, complementary regions of Co-decaying DM parameter space.

As a relevant example, consider the case of LLPs produced from the production of 1 TeV gluinos. The most pessimistic reach estimate is obtained by assuming that the gluino decays directly to B and that most of the gluino mass is converted to LLP boost. In that case, the region of co-decaying DM parameter space where MATHUSLA sees at least 4 LLP decays is shown as the dark gray region in Fig. 31. If less than all the gluino rest energy is converted to LLP boost due to other energetic decay products, the lab-frame LLP lifetime decreases, shifting sensitivity to larger values of τ_B and probing more regions which are not already excluded by the out-of-equilibrium bound.

It could also be the case that the B LLP is not produced directly in gluino decays, but rather as part of a visible-sector decay chain that terminates with particle χ . If χ decays to the dark sector, the same small portal coupling which makes B long-lived can also make χ long-lived, though its decay width is expected to scale as $\Gamma_\chi/\Gamma_B \sim m_\chi/m > 1$ (or some higher power). Given that χ also has lower boost by a factor of $\sim m_\chi/m$, the lab-frame lifetime of χ is shorter than for B by at least a factor of $(b_\chi c\tau_\chi)/(b_B c\tau_B) \sim (m/m_\chi)^2$. MATHUSLA then has two chances of detecting an LLP: either B itself (produced in χ decays) or χ (produced in gluino decays). This allows MATHUSLA to probe both the dark-gray region in Fig. 31 and regions at larger τ_B where B escapes undetected but χ decays in MATHUSLA.

The comparison of MATHUSLA’s reach main detector LLP searches depends sensitively on model details, i.e. the exact production and decay modes. However, as argued in Section 3, the cross-section reach for LLP detection is likely to be better at MATHUSLA, possibly by orders of magnitude if the LLP is produced and decays without highly conspicuous jets, leptons or missing energy in the final state. Of course, the main detectors would have better sensitivity at lower lifetimes. MATHUSLA and the main detectors together would therefore probe deep into the center of the co-decaying dark matter parameter space.

5.5 Dynamical Dark Matter³⁷

In this section we study how MATHUSLA could discover DM within the *Dynamical Dark Matter* (DDM) framework [318–323], which intrinsically gives rise to a large number of dark sector states with varying lifetimes from collider-scales to cosmological hyperstability. This is a particularly dramatic example of LLP signatures giving direct insight into the nature, as well as the cosmological and astrophysical role, of the dark sector.

5.5.1 Introduction

In most models of DM, the dark sector is composed of one or several dark-matter particle(s) χ which carry the entire dark-matter cosmological abundance $\Omega_{\text{CDM}} \approx 0.26$ [324]. These particle(s) must be hyperstable, with lifetimes exceeding the age of the universe by many orders of magnitude: $\tau_\chi \gtrsim 10^{26}$ s. This stability is critical for traditional dark matter. Indeed, any particle which decays too rapidly into Standard-Model (SM) states is likely to upset BBN and light-element abundances, and also leave undesirable imprints in the CMB and diffuse X-ray/gamma-ray backgrounds. However, as a result of this stability, the resulting dark sector is then essentially “frozen” in time, with Ω_{CDM} remaining constant in our late-time matter-dominated universe. Moreover, as explained above, this stability also ensures that once such a dark-matter particle is produced in a collider, it escapes without any subsequent observable decay.

Dynamical Dark Matter [318–323] generalizes this assumption by positing that the dark sector consists not merely of one or more hyperstable DM particles, but *many* such particles which can have varying lifetimes. The number N of dark-matter states can be order 10, 1000, or even much larger in some scenarios. Thus, instead of having a single dark-matter particle χ , the dark sector contains an entire *ensemble* of dark-sector states χ_n ($n = 1, \dots, N$). In that case, no state individually needs to carry the full abundance Ω_{CDM} so long as the sum of their individual abundances Ω_n matches Ω_{CDM} . In particular, the individual dark components within the ensemble can carry a wide variety of abundances Ω_n , some relatively large but others relatively small. This is a critical observation, because *a given dark-matter component χ_n need not be stable if its abundance Ω_n at the time of its decay into SM states is sufficiently small*. Indeed, a sufficiently small abundance assures that all of the disruptive effects of the decay of χ_n into SM states will be minimal, and that all constraints from BBN, CMB, *etc.* will continue to be satisfied.

We are thus naturally led to an alternative concept [318]: balancing of SM-producing decay widths Γ_n against cosmological abundances Ω_n . Dark-matter states with larger abundances must have smaller decay widths and survive until (and potentially beyond) the present time, but states with smaller abundances can have larger decay widths and decay at earlier times. As long as decay widths are balanced against abundances in this way across our entire dark-sector ensemble, all phenomenological constraints can potentially be satisfied. Thus, dark-matter hyperstability is no longer required.

This is the basic principle of Dynamical Dark Matter: an alternative framework for dark-matter physics in which the notion of dark matter stability is replaced by a balancing of lifetimes against cosmological abundances across an ensemble of N individual dark-matter components χ_n with different

³⁷David Curtin, Keith R. Dienes, Brooks Thomas

masses m_n , lifetimes $\tau_n \equiv \Gamma_n^{-1}$, and abundances Ω_n . DDM is in some sense a natural generalization of the standard scenario of a single hyperstable DM species, which is recovered in the $N \rightarrow 1$ limit. In general, DDM can give rise to far richer cosmology and phenomenology. As its name implies, the dark sector becomes truly *dynamical*, with the different components of the DDM ensemble decaying before, during, and after the present epoch. Indeed, some portions of the DDM ensemble have already decayed prior to the present epoch, and are thus no longer part of the dark sector. However, other portions of the DDM ensemble have yet to decay. It is these ensemble constituents whose abundances Ω_n together comprise the specific dark-matter abundance $\Omega_{\text{CDM}} \approx 0.26$ observed today.

Since the original DDM proposal [318–320], there have been many explicit realizations of such DDM ensembles, i.e. different theoretical scenarios for BSM physics which give rise to a large collection of dark states in which the widths for decays into SM states are naturally inversely balanced against cosmological abundances. These include theories involving large extra spacetime dimensions [318–320], theories involving strongly-coupled hidden sectors [325, 326], theories involving large spontaneously-broken symmetry groups [327], and examples from string theory [325, 326, 328]. Indeed the dark states within these different realizations can accrue suitable cosmological abundances in a variety of ways, including not only through non-thermal generation mechanisms such as misalignment production [318–320] but also through thermal mechanisms such as freeze-out [329]. Mass-generating phase transitions in the early universe can also endow collections of such states with non-trivial cosmological abundances [330–332].

In these and other realistic DDM scenarios, the masses, lifetimes, and abundances of these individual particles are not arbitrary. Rather, these quantities are determined by the underlying physics model and take the form of *scaling relations* (either exact or approximate) which encode their dependence on each other and how they vary within the DDM ensemble. These scaling relations completely specify the properties of the ensemble constituents through a relatively small number of free parameters. Thus, even though the number of particles which contribute to the total dark-matter abundance is typically large, DDM scenarios can nevertheless be very predictive.

The most fundamental of these scaling relations governs the spectrum of masses for the DDM constituent particles χ_n . In general, we assume a constituent mass spectrum of the form

$$m_n = m_0 + (\Delta m) n^\delta \quad (147)$$

where $\{m_0, \Delta m, \delta\}$ are arbitrary parameters and where $\Delta m, \delta > 0$ (so that n labels the DDM constituents in order of increasing mass). Most concrete realizations of DDM ensembles have mass spectra which take this general form, either exactly or approximately. For example, if — as in Refs. [318, 319] — the ensemble constituents consist of the Kaluza-Klein (KK) excitations of a scalar field compactified on a circle of radius R (or a Z_2 orbifold thereof), we have either $\{m_0, \Delta m, \delta\} = \{m, 1/R, 1\}$ or $\{m_0, \Delta m, \delta\} = \{m, 1/(2mR^2), 2\}$, depending on whether $mR \ll 1$ or $mR \gg 1$, respectively, where m is the four-dimensional scalar mass. In general, for arbitrary mR , we find that the latter behavior holds for $n \ll mR$ and the former for $n \gg mR$. Likewise, if the ensemble constituents consist of the bound states of a strongly-coupled gauge theory, as in Refs. [325, 326], we have $\delta = 1/2$, where Δm and m_0 are related to the Regge slope and Regge intercept of the strongly-coupled theory, respectively. Thus $\delta = 1/2$, $\delta = 1$, and $\delta = 2$ may be considered as particularly compelling “benchmark” values.

Given a mass spectrum of this general form, we then typically take a scaling relation for the decay widths Γ_n of the form

$$\Gamma_n = \Gamma_0 \left(\frac{m_n}{m_0} \right)^y \quad (148)$$

where Γ_0 is the decay width of the lightest DDM state and y is an additional free parameter. Note that Γ_n is assumed to be the decay width of the n th ensemble constituent χ_n into SM states, and we disregard the possibility of intra-ensemble decays (or assume that the branching ratios for such decays are relatively small). The corresponding χ_n lifetimes are then given by $\tau_n \equiv \Gamma_n^{-1}$. In general, the scaling exponent

y can be arbitrary. For example, if we assume that the dominant decay mode of χ_n is to a final state consisting of SM particles whose masses are all significantly less than m_n , and if this decay occurs through a dimension- d contact operator of the form $\mathcal{O}_n \sim c_n \chi_n \mathcal{O}_{\text{SM}} / \Lambda^{d-4}$ where Λ is an appropriate mass scale and where \mathcal{O}_{SM} is an operator built from SM fields, we have

$$y = 2d - 7. \quad (149)$$

In general one finds $y > 0$ (such as $y = 5$ for hidden valley decays mediated by a dark photon, see Eqn. (233)) but this need not be a strict requirement. Indeed, since the fundamental couplings that underlie such decays can often themselves depend on n , the scaling exponents y could be large, depending on the scaling behavior of c_n and the dimensionality of \mathcal{O}_{SM} .

Another important quantity is the spectrum of cosmological relic abundances Ω_n associated with each DDM constituent. These are likewise assumed to satisfy an approximate scaling relation of the form

$$\Omega_n = \Omega_0 \left(\frac{m_n}{m_0} \right)^\gamma. \quad (150)$$

The precise value of the scaling exponent γ generally depends on the particular dark-matter production mechanism assumed. One typically finds that $\gamma < 0$ for misalignment production [318, 319], while γ can generally be of either sign for thermal freeze-out [329].

In a similar vein, for many investigations it is instructive to focus on the coupling coefficients $c_{m,n,\dots,p}$ of Lagrangian operators which involve multiple ensemble constituents $\{\chi_m, \chi_n, \dots, \chi_p\}$ together with some set of particles outside the ensemble. Such couplings can ultimately be relevant for dark-matter production, scattering, annihilation, and decay. In the analysis below, we are mostly interested in couplings that involve two dark-matter constituents χ_m and χ_n (or their antiparticles), and we further restrict our attention to the “diagonal” $m = n$ case. We then assume a scaling relation for such couplings of the form

$$c_{n,n} = c_0 \left(\frac{m_n}{m_0} \right)^\xi \quad (151)$$

where c_0 is an overall normalization and where ξ is a corresponding scaling exponent. For example, $\xi = 0$ corresponds to democratic decay into different final states that are much lighter than the parent particle, while $\xi = 1$ corresponds to a Yukawa-like coupling instead. Assuming a scaling relation of this form allows us to study a wide variety of underlying theoretical mechanisms that might generate such couplings. Once a particular scaling relation for the coupling is specified, the scaling behaviors of the corresponding production, scattering, or annihilation cross-sections are also determined. Since these cross-sections also depend on kinematic factors, their behavior across the ensemble can deviate significantly from the simple power-law couplings we have assumed for the underlying couplings. For example, the results in Ref. [329] can be interpreted as illustrating the large range of possible scaling behaviors that can be exhibited by an annihilation cross-section when the underlying couplings c_{nn} are held fixed ($\xi = 0$) across the entire DDM ensemble.

In general, the phenomenological viability of the DDM framework rests upon relations between these different scaling exponents. Two of the most important which underpin the entire DDM framework are the relations [318, 329]

$$\gamma y < 0 \quad (152)$$

and

$$-1 \lesssim \frac{1}{y} \left(\gamma + \frac{1}{\delta} \right) < 0. \quad (153)$$

The first of these relations simply ensures that states with larger abundances have smaller decay widths/longer lifetimes, as required within the DDM framework. The second relation ensures a suitable effective equation of state for the collective DDM ensemble, with an effective equation-of-state

parameter $w \approx 0$ that does not change appreciably over a significant portion of the recent cosmological past [329]. Moreover, this relation also ensures that the total energy density carried by the ensemble is finite in the $N \rightarrow \infty$ limit.

5.5.2 Detecting DDM at MATHUSLA

Scenarios within the DDM framework can give rise to distinctive signatures at colliders [333–335], at direct-detection experiments [336], and at indirect-detection experiments [337–340]. Such scenarios also give rise to enhanced complementarities [341, 342] between different types of experimental probes.

If a production mode for DDM states is available at the LHC, such as a heavy BSM state with SM charge that decays into DDM states, then the ensemble can give rise to a mixed variety of MET and DV signatures in the main detectors. Kinematic analysis of the visible final states could provide evidence for the multi-component nature of DM and the existence of DDM scaling relations. On the other hand, if the decay length of accessible states is much larger than the main detector size then MATHUSLA could be the best discovery opportunity for DDM. This is the scenario we examine here.

As evidenced by Eqn. (11), MATHUSLA is capable of detecting LLPs with lifetimes at or even exceeding the BBN limit depending on the LHC production cross-section. Even so, it is clear that no possible LHC production rate would allow MATHUSLA to detect the decay of the most stable DDM states which constitute the DM abundance today. Fortunately, the very nature of the DDM ensemble gives rise to dark states with a large variety of possible lifetimes, with decay lengths generally decreasing with increasing mass. Production of heavier and detectably meta-stable DDM states is therefore possible at energy frontier machines like the HL- or HE-LHC. This could give rise to MATHUSLA signatures alongside MET signatures from both the MATHUSLA-detectable states and longer-lived states in the ensemble. If the production process of DDM states does not give rise to hard SM final states in decays, then MET searches have to rely on ISR and MATHUSLA could be our only probe of these DDM scenarios, see Section 3.2.1. Careful analysis of the observed LLP decays within MATHUSLA could then reveal their varying masses and lifetimes [9] and provide evidence for the scaling relations of the DDM ensemble. On the other hand, if the DDM states are produced in the decays of heavy particles that also produce SM-charged final states, then correlating the MET and LLP signatures will clearly be an important tool to constrain the properties of the DDM ensemble. In either scenario, MATHUSLA would be invaluable to discover and diagnose DDM.

5.5.3 MATHUSLA reach in a benchmark DDM parameter space

In order to provide a quantitative assessment of the reach of MATHUSLA within the DDM parameter space, we conduct a toy study of the simplest scenario whereby DDM states χ_n are produced in the prompt two-body decay $\phi \rightarrow \chi_n \chi_n$ of a heavy state ϕ with mass m_ϕ and LHC production cross-section σ_ϕ . At our level of analysis, the spin of ϕ and χ_n is irrelevant, and we fix kinematics of ϕ by assuming the n^{th} DDM state has an average boost factor $b_n = \frac{|\vec{p}_n|}{m_n} = \frac{m_\phi}{2m_n} \sqrt{1 - 4m_n^2/m_\phi^2}$. The masses and decay widths of the χ_n are given by the scaling relations Eqns. (147) and (148), while their relative couplings to ϕ are determined by Eqn. (151). We therefore have a nine-dimensional parameter space $\{m_\phi, \sigma_\phi, m_0, \Delta m, \delta, \Gamma_0, y, c_0, \xi\}$. For each value of the chosen parameters, we can use the simple expressions in Section 3.1.1 to estimate the rate of observed decays within MATHUSLA for each state χ_n .

For concreteness, we take Γ_0 to be determined by the traditional dark-matter hyperstability bound, i.e., $\Gamma_0 = (10^9 t_{\text{now}})^{-1} = 10^{-26} \text{ s}^{-1}$ where $t_{\text{now}} = 10^{17} \text{ s}$ is the current age of the universe. (Larger values will simply linearly rescale the signal in the long-lifetime limit.) We also set $m_\phi = 2 \text{ TeV}$ as a concrete benchmark, to be discussed further below. If ϕ has couplings to SM or other non-DDM states, then the quantity c_0 determines the total branching fraction $\text{BR}_{\chi\chi} \equiv \sum_{n=0}^{\infty} \text{BR}(\phi \rightarrow \chi_n \chi_n)$ of ϕ into DDM states. Since MATHUSLA is ultimately sensitive not to σ_ϕ alone but to the product $\sigma_\phi \text{BR}_{\chi\chi}$,

we therefore now have a seven-dimensional parameter space $\{m_\phi, \sigma_\phi \text{BR}_{\chi\chi}, m_0, \Delta m, \delta, y, \xi\}$. We shall therefore quantitatively assess the reach of the MATHUSLA detector in terms of the minimum value of $\sigma_\phi \text{BR}_{\chi\chi}$ gives rise to four observed LLP decays within the MATHUSLA decay volume, given specific values of the remaining six parameters $\{m_\phi, m_0, \Delta m, \delta, y, \xi\}$. In the zero-background regime, this can be interpreted as an exclusion limit on $\sigma_\phi \text{BR}_{\chi\chi}$. In the event that LLP decays are observed, this would correspond roughly to the minimum cross-section required for DDM discovery.

For any value of m_ϕ , the decays of ϕ can potentially produce the ensemble constituents $\{\chi_0, \chi_1, \dots, \chi_{n_{\text{kin}}}\}$, where n_{kin} is the kinematic limit, defined as the maximum value of n for which $m_n \leq m_\phi/2$. We also define two further quantities n_{min} and n_{max} as those values of n which delimit the range of ensemble constituents χ_n whose subsequent decays into SM states are responsible for approximately 90% of the observed events within MATHUSLA. Thus $(n_{\text{min}}, n_{\text{max}})$ describes that subset of DDM ensemble states to which the MATHUSLA detector is most sensitive. Finally, we define n_{cs} as indicating the heaviest ensemble constituent χ_n which is cosmologically stable, with $\tau_n \equiv \Gamma_n^{-1} \geq t_{\text{now}}$. Thus only the ensemble constituents $\{\chi_0, \chi_1, \dots, \chi_{n_{\text{cs}}}\}$ will have survived to the present time and have the potential to contribute to the total present-day dark-matter abundance $\Omega_{\text{CDM}} \approx 0.26$. We have already noted that the MATHUSLA detector, while capable of probing large portions of the DDM ensemble, cannot actually probe those elements of the ensemble which constitute dark matter today, and $n_{\text{cs}} < n_{\text{min}}$. Therefore, the $n_{\text{cs}} < 1$ contour in our plots demarcates the area of parameter space in which DDM realizes a more traditional DM model with only a single hyperstable DM particle χ_0 . However, even in those scenarios the DDM ensemble could contain long-lived states that may have affected the early cosmological history of the universe.

Given these definitions, our results are as follows. In Fig. 32, we indicate the sensitivity of MATHUSLA by plotting contours of $\sigma_\phi^{\text{min}} \text{BR}_{\chi\chi}$, where σ_ϕ^{min} is the minimum production cross-section for the parent particle ϕ which will produce at least four signal events within MATHUSLA. In the left and right panels of Fig. 32, these contours (black curves) are plotted within the $(m_0/\Delta m, y)$ and $(m_0/\Delta m, \xi)$ planes, respectively. For each plot we have chosen the benchmark values $m_\phi = 2$ TeV, $m_0 = 100$ MeV, $\xi = 1$, and $\delta = 1.5$. Within each panel we also show contours of n_{min} (blue curves), $n_{\text{max}} - n_{\text{min}}$ (red curves), and n_{cs} (green curves). The orange shaded region in the left panel is the region in which at least one of the χ_n has a characteristic decay length $\beta\gamma c\tau_{\text{min}} < 1$ m. The results in Fig. 32 correspond to the case in which the χ_n are real scalars, but the results for spin-1/2 fermions are qualitatively similar.

In this connection, our benchmark value $m_\phi = 2$ TeV deserves further comment. This benchmark is motivated in part by a self-consistency requirement: in order for the ensemble to lead to a detectable signal at MATHUSLA during the HL-LHC run, the production cross-section σ_ϕ must exceed the sensitivity threshold σ_ϕ^{min} at any point within the DDM parameter space. For example, if ϕ is a real scalar that couples to quarks through a Yukawa-type interaction with a flavor-independent coupling constant g_q , the dominant production process for ϕ is resonant production of ϕ through quark fusion. In this case, we find that the product of the production cross-section and this branching fraction is $\sigma_\phi \times \text{BR}_{\chi\chi} \sim 100$ fb for the choice $m_\phi = 2$ TeV (with $g_q = 0.15$ and c_0 chosen such that the total branching fraction $\text{BR}_{\chi\chi}$ of ϕ to χ_n pairs is 0.5). As m_ϕ increases beyond this benchmark value, σ_ϕ rapidly decreases, rendering nearly all of the DDM parameter space beyond the reach of MATHUSLA during the upcoming LHC run. By contrast, while σ_ϕ can be significantly larger than 100 fb for m_ϕ below our 2 TeV benchmark, ATLAS and CMS searches for new physics in the monojet + MET [343, 344] and dijet [345, 346] channels impose stringent lower bounds on m_ϕ . Nevertheless, values of m_ϕ at or slightly below this benchmark are consistent with these constraints. Thus, we see that the choice $m_\phi = 2$ TeV corresponds to a MATHUSLA sensitivity in the range $\sigma_\phi^{\text{min}} \times \text{BR}_{\chi\chi} \sim 100$ fb and that values of m_ϕ near this benchmark are of particular phenomenological interest.

We see from the results shown in the left panel of Fig. 32 that there is indeed a substantial region of parameter space within which MATHUSLA is capable of detecting a DDM ensemble. As discussed in Section 3.2.3, the main detector reach for our simple scenario depends strongly on the decay mode

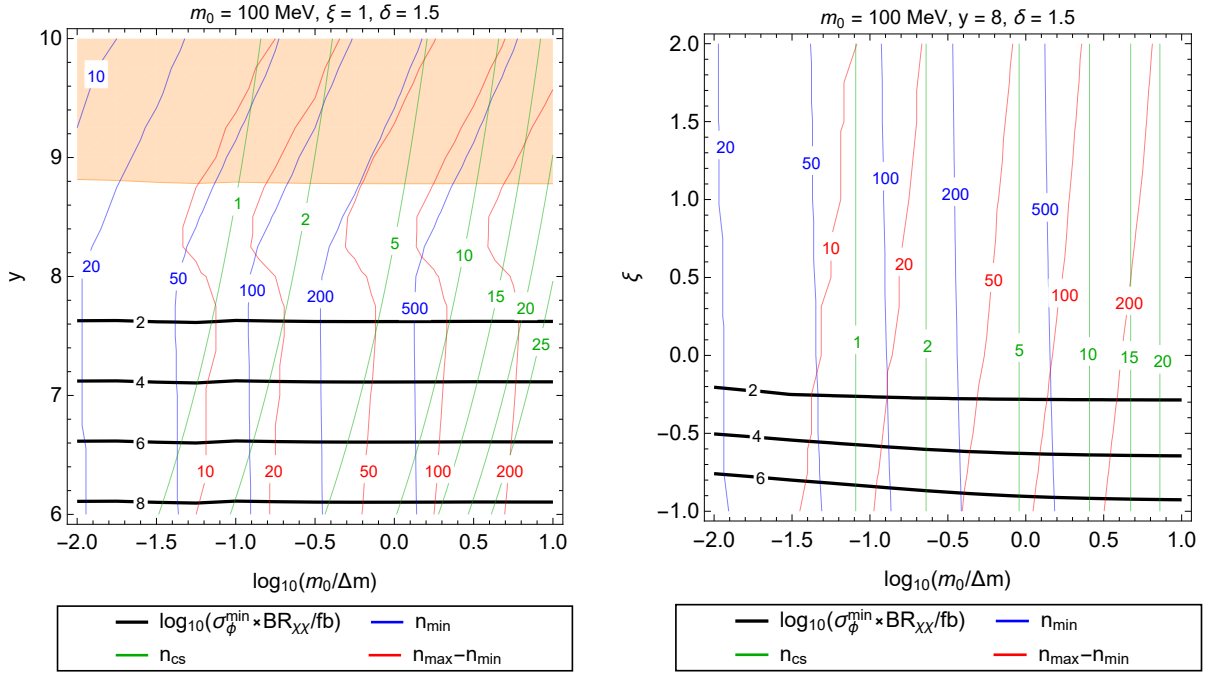


Fig. 32: The reach of MATHUSLA within the DDM parameter space for the benchmark values $m_0 = 100$ MeV and $\delta = 1.5$. Black curves indicate contours of $\sigma_\phi^{\min} \text{BR}_{\chi\chi}$, while blue, red, and green curves indicate contours of n_{\min} , $n_{\max} - n_{\min}$, and n_{cs} , respectively. Likewise, the orange shading indicates the region of DDM parameter space in which at least one of the ensemble constituents χ_n has a characteristic decay length $\beta\gamma c\tau_{\min} < 1$ m. As discussed in the text, the region with $m_0/\Delta m \gtrsim 0.1$, $7.5 \lesssim y \lesssim 8.8$, and $\xi \gtrsim -0.3$ is a particular “sweet spot” in which multiple light states within the DDM ensemble comprise the present-day dark matter while numerous heavier states within the same ensemble can lead to an observable signal at MATHUSLA. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

of the DDM states, which is not specified in our toy model. However, there are many general scenarios, like decay to hadrons or Yukawa- or gauge-ordered democratic decay to SM fermions, for which MATHUSLA is likely to exceed the main detector reach by orders of magnitude in cross-section.

The most obvious region for a MATHUSLA signal is $7.5 \lesssim y \lesssim 8.8$. For $y \gtrsim 8.8$, the characteristic decay lengths of the heaviest states in the tower fall below $\beta\gamma c\tau_n \lesssim \mathcal{O}(1)$ m. Since a significant number of particles with decay lengths in this regime would decay inside the main collider detector, ensembles with $y \gtrsim 8.8$ would either be detected at the high-luminosity LHC without the help of MATHUSLA or would already have been detected during the current LHC run. On the other hand, for $y \lesssim 7.5$, a parent-particle production cross-section $\sigma_\phi \text{BR}_{\chi\chi} \gtrsim 10^3$ fb is required in order for the expected number of signal events in the MATHUSLA detector to exceed the detection threshold. This is approaching the upper range of typical strong production rates for TeV-scale states, see Fig. 8. Furthermore, given the sensitivity of monojet searches to invisible Higgs decays, see Section 3.2.1, such large cross-sections are likely to be detectable (and possibly excluded by) current or future (HL-)LHC monojet searches.

The right panel of Fig. 32 indicates how the sensitivity of MATHUSLA depends on ξ , the scaling exponent for the couplings in Eqn. (151). For this plot where we have taken a fixed scaling exponent $y = 8$ for the decay widths of the χ_n . We see from this figure that there is generally a loss of sensitivity for MATHUSLA as ξ decreases. This behavior ultimately reflects the fact that for $\xi < 0$, the width of ϕ is dominated by decays to the lightest states in the DDM ensemble, which are also the states with the longest lifetimes.

We see from Fig. 32 that the reach of the MATHUSLA detector is not particularly sensitive to the




Fig. 33: Same as Fig. 32, except that we have now shifted m_0 from 100 MeV to 100 keV. This allows MATHUSLA to be sensitive to DDM ensembles with smaller values of y , leading to an even more compelling “sweet spot” with $m_0/\Delta m \gtrsim 0.01$, $4.3 \lesssim y \lesssim 5.0$, and $\xi \gtrsim -0.2$.

ratio $m_0/\Delta m$ — at least not within the region of parameter space shown. However, we see that this ratio nevertheless plays a crucial role in determining n_{cs} , the number of χ_n states which are cosmologically stable, with $\tau_n \gtrsim t_{\text{now}}$. Indeed, given the contours of n_{cs} shown in Fig. 32, we see that a significant number of ensemble constituents χ_n are cosmologically stable for $m_0/\Delta m \gtrsim 0.1$. By contrast, for $m_0/\Delta m \lesssim 0.1$, the only contribution from the ensemble to Ω_{CDM} is that associated with the single lightest particle species χ_0 . Thus, the region of parameter space in which $\xi \gtrsim -0.3$, $m_0/\Delta m \gtrsim 0.1$, and $7.5 \lesssim y \lesssim 8.8$ is of particular interest from a DDM perspective, with many individual dark-matter components χ_n potentially comprising the total present-day dark-matter abundance Ω_{CDM} .

Taken together, the results in Fig. 32 indicate that there exists a significant region of parameter space in which multiple light states in the DDM ensemble can contribute non-negligibly to the present-day dark-matter abundance — all while heavier states in the same ensemble can lead to an observable signal at MATHUSLA. This alone demonstrates that MATHUSLA can be highly relevant for collider-based probes of the DDM framework. However, these results also indicate that the value of the scaling exponent y in these regions is relatively large. Although there is no fundamental reason why such values are problematic, it would be interesting from a theoretical perspective to know whether the same successes can be achieved with smaller values of y .

Such regions of parameter space also exist, and correspond to very light hyperstable DDM states. In Fig. 33, we plot essentially the same information as we plotted in Fig. 32, the only change being that we have now taken $m_0 = 100$ keV rather than $m_0 = 100$ MeV. We see that this shift in m_0 has not changed the gross features of these plots relative to those in Fig. 32, but has shifted the regions in which MATHUSLA is most sensitive down to smaller values of y — precisely as desired. Indeed, we now see that MATHUSLA remains sensitive to the DDM ensemble even below $y \approx 5$ — a very natural value for y , given that this value corresponds to a dimension-six decay operator according to Eqn. (149), see also Eqn. (233). Once again, just as for greater m_0 , we find that taking $m_0/\Delta m \lesssim 0.01$ leads to situations in which only a single dark-matter component survives to the present day. However, for values

of $m_0/\Delta m \gtrsim 0.01$, we find that multiple components of the ensemble survive to the present day and can potentially contribute to Ω_{CDM} . Thus, for $m_0 \approx 100$ keV, MATHUSLA is sensitive to a theoretically particularly motivated region of DDM parameter space, $m_0/\Delta m \gtrsim 0.01$, $4.3 \lesssim y \lesssim 5.0$, and $\xi \gtrsim -0.2$. The low mass of the lightest DDM state(s) means that astrophysical and cosmological constraints may apply, but they are highly dependent on the mechanism which generates the DDM abundance. We leave analysis of these constraints for future work.

In summary, DDM is a framework for DM which generalizes the scenario of a single or a few hyperstable DM constituents. It arises naturally in a variety of top-down theoretical frameworks and gives rise to meta-stable states that are related to, and/or are a part of, the states contributing to the DM abundance today. Given this continuum of realizable lifetimes, MATHUSLA will be an important discovery and diagnosis tool for DDM, along with cosmological, astrophysical, and direct detection searches. If all accessible ensemble states have decay lengths exceeding the main detector size, MATHUSLA could easily be the first or only discovery opportunity for DDM.

6 Theory Motivation for LLPs: Baryogenesis

An overwhelming array of evidence indicates that the observable universe is expanding and originates from a very compressed dense state that was filled with a hot primordial plasma (“hot big bang”). If matter and antimatter had been present in exactly the same amount during that epoch, then they would have mutually annihilated, and no baryons would have been left to form galaxies, stars and eventually human beings. Hence, the presence of baryons in the present day universe clearly indicates a matter-antimatter asymmetry in the early universe, see e.g. [347]. The magnitude of this *baryon asymmetry of the universe* (BAU) can be determined from the baryon to photon ratio, or equivalently the ratio $Y_B = (n_B - n_{\bar{B}})/s$ of the total comoving baryon density $n_B - n_{\bar{B}}$, and the entropy density s . The value of Y_B can be consistently extracted from two very different measurements: the Cosmic Microwave Background anisotropy power spectrum [284], and the abundance of light elements in the intergalactic medium [348]. The current observed value for this ratio is $Y_B = (8.6 \pm 0.1) \times 10^{-11}$ [284]. In inflationary cosmology, this number cannot be set as an initial condition for the universe because the rapid expansion would have diluted any pre-inflationary asymmetry. Hence, the BAU must be generated dynamically by *baryogenesis* during or after inflation, but before the onset of Big Bang Nucleosynthesis when the Universe was at a temperature of $T \sim \text{MeV}$.

Baryogenesis requires fulfillment of the Sakharov conditions [349]: i) violation of baryon number B , ii) violation of charge conjugation C and charge+parity conjugation CP and iii) a deviation from thermal equilibrium. While all these ingredients are in principle provided in the SM, the numerical value of Y_B cannot be explained within the SM with a standard cosmological history because the amount of CP violation is too small [350] and the deviation from equilibrium is insufficient [351, 352]. The BAU is therefore a clear sign of BSM physics, and many possible extensions of the SM could generate the required baryon asymmetry.

Many BSM models of baryogenesis have been studied, including electroweak baryogenesis with modifications to the SM Higgs potential that induce a first-order phase transition [353], the Affleck-Dine mechanism [354], and Cogenesis mechanisms where the BAU is related to a dark matter asymmetry [355]. With the exception of a few cases (like electroweak baryogenesis, which a priori has to involve weak-scale degrees of freedom), baryogenesis models are generally very difficult to test, since the new physics can enter at many possible scales.

In this section, we concentrate on two baryogenesis scenarios that can give rise to LLP signals at the LHC and MATHUSLA: WIMP Baryogenesis, where the late-time decay of a weak-scale LLP is directly responsible for the production of the BAU, and Leptogenesis, which involves extensions of the SM neutrino sector (see Section 7) and can generate LLPs due to either the late-time decays required to generate lepton-number, or more generically due to the small couplings involved in the neutrino sector.

6.1 WIMP Baryogenesis³⁸

6.1.1 Introduction

In models of WIMP Baryogenesis [356] (WIMP BG), the WIMP miracle is leveraged to explain the observed BAU. Since the observed baryon and DM abundances in our universe are within an order of magnitude of each other, and since WIMPs can give rise to the observed DM abundance, the late-time decay of a WIMP-like parent particle into more baryons (or leptons) than anti-baryons (or anti-leptons) could easily give rise to the observed BAU. This scenario is particularly exciting, since the meta-stable WIMP-like progenitor particles can be produced at colliders and have to be long-lived in order to decay out-of-equilibrium in the early universe and satisfy the Sakharov conditions. Not only does this give rise to LLP signatures, we might even be able to observe a B or L -asymmetry in the decay of these LLPs, thereby allowing us to *directly* study the same primordial process of creation that gave rise to all the matter we are made of today.

³⁸Yanou Cui, Seyda Ipek

A particle decays out of equilibrium if its lifetime is longer than the Hubble time at a temperature comparable to its mass; $\tau_X > H^{-1}(T \sim M_X)$, where M_X is the mass of the particle X and $H(T)$ is the Hubble rate at temperature T . This gives a lower bound on the lifetime that depends on the particle mass. Since the baryon asymmetry needs to be produced before the BBN, there is also an upper bound on this lifetime: $\tau_X \lesssim 1$ s, that is $c\tau_X \lesssim 10^8$ m in terms of proper decay length. Hence baryogenesis requirements push us to an interesting region: a particle that could be produced at the LHC with mass $M_X \lesssim$ TeV could also have a lifetime long enough to travel $O(100\text{m})$ or longer.

The BBN limit on LLP lifetime is not the only cosmological reason why WIMP BG is a particularly attractive target for MATHUSLA. Relatively heavy LLPs at or above the TeV scale could decay out-of-equilibrium at lifetimes less than $\sim 10^{-11}$ seconds, before the time of the electroweak phase transition. These relatively short-lived LLPs could decay into either baryons or leptons, since a generated lepton asymmetry would be converted into the observed BAU via sphaleron processes that are still active in the plasma for $T \gtrsim 100$ GeV. On the other hand, progenitor LLPs with masses in the 100 GeV range have to have lifetimes longer than $\sim 10^{-10}$ seconds to decay out-of-equilibrium. It is also possible that heavier LLPs have a longer lifetime than the minimum required for WIMP BG. In either case, the LLP has to decay into baryons to generate a BAU, since the decay takes place after sphaleron processes have already switched off.

In other words, WIMP BG strongly motivates LLPs with decay lengths above \sim cm, and if they have such a long lifetime, they have to decay dominantly into hadrons. This is the perfect target for MATHUSLA, both because the lifetime above the \sim cm minimum is arbitrary and can easily exceed 100m, and because LLPs decaying hadronically are more difficult to search for in the main detectors, especially if their masses are around or below the 100 GeV scale. MATHUSLA would be orders-of-magnitude more sensitive to production rates of such LLPs than main detector searches, see Section 3.2.3, and could be our only option for discovering for many WIMP BG scenarios.

6.1.2 Model-independent Features and Phenomenology with MATHUSLA

Let us summarize at a more quantitative level the model-independent features behind the weak-scale baryogenesis scenario in which a weak scale particle X decays after its thermal freeze-out. The freeze-out temperature T_{fo} is given by

$$n_X^{\text{eq}} \langle \sigma_{\text{ann}} v \rangle = H(T_{\text{fo}}) \longrightarrow T_{\text{fo}} \sim M_X \ln^{-1} \left[10^7 \left(\frac{M_X}{100 \text{ GeV}} \right) \left(\frac{\sigma_{\text{ann}}}{\text{fb}} \right) \right]. \quad (154)$$

Here $H(T) = \sqrt{\frac{4\pi^3 g_*(T)}{45}} \frac{T^2}{M_{\text{pl}}}$ is the Hubble expansion rate with $M_{\text{pl}} \simeq 1.2 \times 10^{19}$ GeV is the Planck mass and $g_*(T)$ is the effective number of relativistic degrees of freedom at temperature T which we take to be ~ 100 for $T \sim 100$ GeV. For weak scale particles with annihilation cross-section $\sigma_{\text{ann}} \sim$ fb, the freeze-out temperature is $T_{\text{fo}} \sim \frac{M_X}{20} \ll M_X$.

Now let us assume X decays to a final state YZ , $X \rightarrow YZ$. For example, if X is a fermion, the final state YZ can either contain 3 fermions or a fermion and a scalar. We also assume X does not carry a lepton or baryon number while YZ has $+1$ baryon number³⁹. Furthermore, the decay needs to be CP violating. The out-of-equilibrium condition for the decays requires that

$$\tau_X > H^{-1}(T \sim M_X) \longrightarrow c\tau_X \gtrsim 1 \text{ cm} \left(\frac{100 \text{ GeV}}{M_X} \right)^2. \quad (155)$$

If the decay temperature is less than the freeze-out temperature, $T_{\text{fo}} > T_{\text{d}} > T_{\text{BBN}}$, and assuming that we can neglect washout processes, the baryon asymmetry is given by

$$\Delta_B = \epsilon_{\text{CP}} n_X(T_{\text{fo}}), \quad (156)$$

³⁹Since we are interested in particles with long lifetimes which would decay after the EW sphalerons shut off, we do not consider lepton number violating decays.




Fig. 34: MATHUSLA reach for the models described in Section 6.1.3 for long-lived particle masses $M_{\text{LLP}} = 0.1 - 1$ TeV and decay lengths $c\tau > m$ for 14 TeV HL-LHC, computed using the analytical approximations outlined in Section 3.1.1. (Regions between the solid lines give $N_{\text{obs}} > 4$ in MATHUSLA.) Note that these models give successful baryogenesis inside the full parameter space shown. The MRSSM reach projection assumes Bino LLP production in the decay of a 1.8 TeV gluino. The Higgs portal singlet projection is shown for $c_H v / \Lambda_H = 0.5$. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

where $\epsilon_{\text{CP}} < 1$ is a measure of CP violation in the decays and is model-dependent.

One important observation is that, in these models, the production mechanism for the new particles can be separate from the decay mechanism. Hence, even though their decays are suppressed, giving rise to long-lived particles, these particles could be copiously produced at the LHC. For example, if an approximately conserved Z_2 symmetry is responsible for the long lifetime, X particles can still be produced in pairs via Z_2 conserving interactions. A detailed outline for an LHC study of simplified models for such baryogenesis mechanisms, including various production and decay channels, can be found in [357], along with a recast of several dedicated displaced vertex searches by ATLAS and CMS.

6.1.3 Motivated Model Examples

Concrete, motivated models realizing the general idea outlined earlier have been proposed recently, with different mechanisms of generating the CP violation effect required by Sakharov conditions. Below we briefly describe three models, which correspond to the three benchmark cases illustrated in Fig.34. Despite different sources of CP violation and model setups, these models utilize B -violating, out-of-equilibrium decay of a weak scale particle in order to satisfy the Sakharov conditions for baryogenesis, and thus generically predict displaced vertex signals at the LHC.

6.1.3.1 WIMP baryogenesis-1: Higgs portal singlet

The WIMP baryogenesis mechanism proposed in [356] is in part motivated by providing a novel way of addressing the “coincidence” between the dark matter and baryon abundances today, while retaining the merit of WIMP miracle. The first example model proposed in [356] includes the following beyond-the-SM Lagrangian terms:

$$\begin{aligned} \Delta\mathcal{L} = & \lambda_{ij}\phi d_i d_j + \varepsilon_i \chi \bar{u}_i \phi + M_\chi^2 \chi^2 + y_i \psi \bar{u}_i \phi + M_\psi^2 \psi^2 \\ & + \alpha \chi^2 S + \beta |H|^2 S + M_S^2 S^2 + \text{h.c.} \end{aligned} \quad (157)$$

H is the SM Higgs, d, u are right-handed SM quarks, with family indices $j = 1, 2, 3$, ϕ is a di-quark scalar with same SM gauge charges as u . χ, ψ are SM singlet Majorana fermions, and S is a singlet

scalar. $\varepsilon_i \ll 1$ are formal small parameters leading to long-lived χ which triggers baryogenesis upon its late decay to udd . All new particles involved are assumed to be of weak scale. Flavor structure of the model is assumed to be third generation dominated so as to be consistent with various constraints. A concrete realization of such a flavor structure based on minimal flavor violation (MFV) is presented in [358]. The CP violation in this model is realized through the interference between the tree level decay of the χ and loop processes involving intermediate ψ . This model example [356] can be embedded in Natural SUSY framework with R-parity violation augmented with a singlet field as baryon parent, where ϕ is identified as the right-handed top squark and CP violation arises from the interference with loop processes involving an intermediate neutralino. In this class of models, the baryon parent χ can be pair produced through the Higgs portal (H - S mixing). The effective Lagrangian for LHC studies are as follows (this is just for illustration, the intermediate states may not be always heavier than χ and get integrated out):

$$\mathcal{L}_{\text{eff}} \supset \frac{c_H}{\Lambda_H} \chi^2 |H|^2 + \frac{c_q}{\Lambda_q} g_{ijk} \chi u_i d_j d_k. \quad (158)$$

In Fig.34 we show the MATHUSLA reach for a benchmark case for such a Higgs portal model where $m_\chi = 0.1 - 1$ TeV, $c_{Hv}/\Lambda_H = 0.5$ and all the other new particles decoupled.

6.1.3.2 WIMP baryogenesis-2: Split MSSM with RPV couplings

A later related work [359] demonstrates an embedding of the general WIMP baryogenesis idea in mini-split SUSY without an additional singlet, where bino plays the role of the baryon parent. The CP violation comes from interference with loop processes involving wino. Due to the very high mass of sfermions in such models, the rate of direct production of binos is negligible at the LHC. Nevertheless, (nearly) pure wino, which plays an essential role for baryogenesis is also expected to be long-lived with RPV decay. The relevant effective interaction is the following:

$$\mathcal{L}_{\text{eff}} \supset \frac{\sqrt{2}g_2 \lambda''_{ijk}}{3m_{\text{sq}}^2} T^a \tilde{W}^a \bar{u}_i \bar{d}_j \bar{d}_k + \text{h.c.}, \quad (159)$$

where heavy squarks with mass m_{sq} are integrated out. In Fig.34 we illustrate the MATHUSLA reach for a benchmark case for such a model where $m_{\tilde{W}} = 0.1 - 1$ TeV.

6.1.3.3 Baryogenesis via Pseudo-Dirac Bino Oscillations

In supersymmetric models with a global $U(1)_R$ symmetry (MRSSM), the gauginos are pseudo-Dirac fermions. Similar to neutral mesons, pseudo-Dirac gauginos go under particle–antiparticle oscillations. CP violation can be enhanced in this quantum-mechanical phenomena and could be $O(1)$. This has been studied in [360] for a pseudo-Dirac bino with RPV couplings. The spectrum of the model is such that the lightest neutralino is a pure bino and is the NLSP. (A keV gravitino is the LSP.) The bino mass is $O(100 \text{ GeV})$ while other neutralinos and gluinos are $O(\text{TeV})$. The sfermions and other (new) scalars in the model are heavier than a few TeV. The baryon-number violating, effective Lagrangian is given by

$$\mathcal{L}_{\text{eff}} \supset -\frac{\sqrt{2}g_Y \lambda''_{ijk}}{3m_{\text{sq}}^2} \tilde{B} \bar{u}_i \bar{d}_j \bar{d}_k - g''_{ijk} S \bar{u}_i \bar{d}_j \bar{d}_k + \text{h.c.}, \quad (160)$$

with heavy squarks with mass m_{sq} integrated out. λ'' is one of the usual RPV coefficients. S is SM a singlet fermion that is the Dirac partner of the bino and the coefficient g'' depends on the parameters of the model. (See [360] for details.) The proper decay length of the bino is

$$c\tau \approx 30 \text{ m} \left(\frac{100 \text{ GeV}}{M_{\tilde{B}}} \right)^5 \left(\frac{m_{\text{sq}}}{10 \text{ TeV}} \right)^4 \left(\frac{10^{-3}}{\lambda''} \right)^2 \quad (161)$$

Due to the $U(1)_R$ symmetry, supersymmetric particles would be pair produced at the LHC. Furthermore, some of the production channels are not available, *e.g.* s-channel Z exchange for neutralinos, which reduces the electroweak production cross-section for the wino. For a spectrum with $M_{\tilde{B}} < 1$ TeV, $M_g = 1.8$ TeV and squarks heavier than 5 TeV, we find that the main bino pair-production occurs via gluino decays to neutralinos and jets. (Note that gluinos are not long-lived.) For 14 TeV LHC, this cross-section is 4.9 fb. MATHUSLA reach for this model is shown in Fig.34.

6.2 Baryogenesis via Exotic Baryon Oscillations⁴⁰

In Refs. [361, 362] a new scenario for baryogenesis through particle-antiparticle oscillations of heavy flavor baryons was proposed. Parts of the mechanism can be related to WIMP BG models described in Section 6.1, but since baryon number is generated via exotic baryon oscillations late in the hadronization era, this mechanism occupies a very different low-mass region of parameter space and is compatible with low inflationary reheating temperatures. It is therefore compatible with a wider variety of inflationary scenarios, in particular avoiding some of the problems associated with high reheating temperatures in axion, gravitino and relaxion models.

The mechanism of baryogenesis via exotic baryon oscillations relies on the existence of GeV-scale Majorana fermions, $\chi_{1,2}$, coupled to SM heavy quarks. Stability of nuclear matter generically requires the χ_i to be long-lived, as discussed below. Focusing on a single Majorana fermion for simplicity, the relevant couplings are the four-fermion interactions with quarks,

$$\frac{g_{ijk}}{\Lambda} \chi u_i d_j d_k + \text{h.c.}, \quad (162)$$

where i, j, k label the up- or down-type quark flavors. This model is similar in field content and interactions to the baryogenesis scenario of Ref. [360] (as described above) but populates a lower-mass parameter space. The heavy flavor baryons must be produced at late times out of equilibrium, and this can be done in a number of ways, most simply through the late decay of a weak scale (or below) particle.

The coupling of χ to heavy flavor quarks in Eq. (162) leads to the dimension-9 baryon-number-violating operator $(u_i d_j d_k)^2$ that sources baryon-antibaryon oscillations which lead to the BAU. For these oscillations to be efficient, m_χ has to be comparable to that of a heavy flavor baryon. The decays of χ are mediated by the four-fermion interaction of Eq. (162) involving lighter quarks, with couplings suppressed relative to those of heavy flavors. For definiteness, focusing on decays to uss quarks, the χ decay length is

$$c\tau_\chi \gtrsim 100 \text{ m} \left(\frac{5 \text{ GeV}}{m_\chi} \right)^5 \left(\frac{\Lambda/\sqrt{g_{uss}}}{20 \text{ TeV}} \right)^4. \quad (163)$$

This lower bound on τ_χ , or equivalently on $\Lambda/\sqrt{g_{uss}}$, arises from avoiding dinucleon decay of ^{16}O at a rate above the limit observed at Super-Kamiokande. Note that this baryogenesis scenario relies on new particles that must be long-lived for reasons that are distinct from the requirement of a departure from thermal equilibrium in the early Universe. (If the exotic baryons χ are produced in out-of-equilibrium decay as in WIMP BG, then the WIMP-like parent would constitute an additional LLP signature of the model.)

At the LHC, these long-lived Majorana fermions are produced through the decay of heavy flavor baryons with branchings on the order of 10^{-3} . A detailed study of the phenomenology of these long-lived particles has not yet been performed. However, results shown in Sections 7.1 and 8.4 demonstrate that MATHUSLA has excellent sensitivity to LLPs produced in B decays, far exceeding the reach of the LHC main detectors and extending the reach of SHiP at long lifetimes. Beauty baryons Λ_b are produced at the LHC with orders of magnitude smaller cross-section than B -mesons [363], but compared to *e.g.* the SM+S model studied in Section 8.4, this is compensated by the much larger exotic branching ratio to

⁴⁰David McKeen

LLPs. Furthermore, the bound of Eqn. 163 places the LLP lifetime exactly in the MATHUSLA regime. We therefore expect that MATHUSLA would be sensitive to a large portion of the parameter space for this class of baryogenesis models.

6.3 Leptogenesis⁴¹

Leptogenesis is a particularly elegant mechanism that relates the BAU to the origin of neutrino masses [364]. The LLP phenomenology of neutrino extensions to the SM is discussed in detail in Section 7. Here, we briefly place leptogenesis scenarios in the context of these studied models.

If the SM is complemented by heavy right-handed Majorana neutrinos ν_R that give masses to the known light neutrinos via the seesaw mechanism [365–369], then the CP-violating decay of the very same particles ν_R can generate a matter-antimatter asymmetry amongst the leptons in the primordial plasma, which is then partly transferred into a $B \neq 0$ by weak sphalerons [370]. Meanwhile a plethora of leptogenesis scenarios has been proposed that adopt the idea that the C/CP violation and out-of-equilibrium Sakharov conditions are fulfilled by colour-neutral particles that may or may not be related to the origin of neutrino masses. A recent review can be found in refs. [371–376].

One way to classify leptogenesis scenarios is through the manner in which the nonequilibrium Sakharov condition is realised. In most scenarios this occurs due to the freeze-out and out-of-equilibrium decay of some heavy particle in the early universe at temperatures above the temperature $T_{\text{sp}} \sim 130$ GeV [377] when sphalerons freeze out ("freeze-out scenario").⁴² The standard thermal leptogenesis proposed in ref. [364] falls into this category. A review of the most studied scenarios of this kind can e.g. be found in ref. [378]. Another alternative is that the asymmetry is generated by feebly coupled particles that do not reach thermal equilibrium before the temperature drops below T_{sp} ("freeze-in scenario", see also Section 5.3). Leptogenesis from neutrino oscillations [379, 380] in the *Neutrino Minimal Standard Model* (νMSM) [380, 381] falls into this latter category. This possibility is particularly interesting in the context of MATHUSLA because the feebly coupled particles are generally long-lived. Both scenarios can be realised within the type-I seesaw model described by the Lagrangian Eq. (164) in Section 7.1 with experimentally accessible Majorana masses M_i , see e.g. [382] for a review. We shall use this well-known model as benchmark scenario in what follows.

In its minimal version, the type-I seesaw model only adds 2 right-handed neutrinos ν_R to the SM. In this case one can qualitatively distinguish between the cases with Majorana masses M_i below vs. above the electroweak scale.

For M_i above the electroweak scale, the "freeze-out scenario" is realised because the BAU is generated in the decay of the ν_R . This is impossible for experimentally accessible M_i [383] unless the M_i are quasi-degenerate ("resonant leptogenesis") [384]. Since the lifetime of ν_{Ri} scales as $\propto U^{-2} M_i^{-5}$ [385–387], searches for ν_{Ri} in this mass range usually focus on prompt decays [388–391]. Here $U^2 = \text{tr}\theta^\dagger\theta$ is the total heavy neutrino mixing. Numerous authors have proposed strategies to refine such searches, cf. e.g. [375, 387, 392, 393] for reviews, but it seems unlikely that MATHUSLA can access the viable leptogenesis parameter region of the minimal model in this mass range because the ν_R will either be too short lived or the number in which they are produced is too small. However, as we discuss below, details of the UV completion of such models may lead to LLP signatures.

For M_i below the electroweak scale, the seesaw relation Eq. (166) enforces comparably smaller Yukawa couplings, so that thermal ν_R production leading to equilibrium in the early universe is delayed and the freeze-in scenario can be realised. Two competing processes generate the asymmetry, the CP-violating oscillations of the ν_R [379, 380] and the decay of Higgs bosons into ν_R and SM leptons [394, 395]. The simplest model that realises this possibility is the νMSM , cf. [396] for a review, in which

⁴¹Marco Drewes

⁴²Contrast this to the WIMP Baryogenesis mechanism discussed in Section 6.1, where baryons can be created directly in out-of-equilibrium decay. This allows the mechanism to take place after the electroweak phase transition.

two ν_{R_i} generate the neutrino masses and the BAU while the third one is a viable Dark Matter candidate (cf. [397] for a recent review). The leptogenesis parameter region in this scenario [398–415] and its realisation within inverse and linear seesaw models [416, 417] has been studied by various authors, as well as the slightly more general case with 3 heavy Majorana neutrinos [403–405, 418–420]. The ν_{R_i} in this mass range tend to be long-lived and can be found in displaced vertex searches at ATLAS and CMS [421–424] or LHCb [425]. For lower masses, fixed target experiments like NA62 [426–428] can access the viable leptogenesis parameter range. In the future the proposed SHiP experiment [429, 430] or a similar detector at LBNE/DUNE [431] or T2K [432] can achieve a higher sensitivity for masses of a few GeV, while the future colliders FCC-ee [411, 415, 433–435], ILC [411, 434] and CEPC [411, 434] can probe heavier masses up to a few tens of GeV. Remarkably, MATHUSLA can access the viable leptogenesis parameter space, and has the potential to be the world’s most sensitive experiment in part of this mass range, see. Fig. 36 in Section 7.1.

Finally, we return to the freeze-out leptogenesis scenario. While the minimal version of these models does not predict LLPs, various UV completions which implement this scenario in the type-1 see-saw can give rise to LLPs due to the mechanism of discrete symmetry breaking at high scales.

In Sec. 7.6 we present a scenario that features resonant leptogenesis with ν_R of masses between 100 GeV to a few TeV [436], which also give rise to long-lived particles that could be detectable at MATHUSLA [437]. This scenario belongs to a class of models in which flavour and CP symmetries and their residual symmetries (G_ν and G_e in the neutrino and charged lepton sectors) explain the measured values of the lepton mixing angles, make predictions for leptonic CP violation in neutrino oscillations and neutrinoless double beta decay as well as connect low energy CP phases with those relevant for leptogenesis. The desired breaking scheme of the flavour and CP symmetries to G_ν and G_e can be realized in explicit models [438] in which flavour symmetry breaking fields obtain peculiarly aligned vacuum expectation values, achieved with an appropriately constructed potential. One can observe that the symmetry preserved by the latter vacuum can be larger than G_ν and G_e . In this case one encounters a point of *enhanced residual symmetry*. In Ref. [439], a model has been constructed in which the slight breaking of such larger symmetry via higher-dimensional operators can be connected to the smallness of the reactor mixing angle θ_{13} . In scenarios of resonant leptogenesis this type of breaking can not only be correlated with the possibility to maximize the CP asymmetries, but also with the longevity of heavy neutrinos. Additional effects, arising from the mixing of the different symmetry breaking sectors, disturb the mentioned symmetry breaking pattern and eventually lead to the breaking of G_ν and G_e . The size of the parameters, controlling the different symmetry breakings, is given in powers of the symmetry breaking parameter, whose size is expected to be a few percent. The effectiveness of symmetry breaking in the different sectors of the theory ultimately depends on the explicit model. However, the crucial insight is that details of the UV completion of freeze-out leptogenesis models may violate the intuition that the associated LLPs always have relatively short lifetime, leading to the possibility of probing these models at MATHUSLA.

7 Theory Motivation for LLPs: Neutrinos⁴³

Since the minimal standard model predicts that neutrino masses are zero, the observations of neutrino oscillations have provided the first definitive evidence for new particle physics beyond the standard model. However, neither the scale nor the detailed nature of the new states responsible for neutrino masses is known, and there are a variety of experiments under way to learn more about the physics of neutrino masses. Here we explore scenarios in which new states responsible for neutrino masses are long-lived.

A wide class of theories that explain the smallness of neutrino masses predict that neutrinos are their own anti-particles, *i.e.*, they are Majorana fermions. This implies that the interactions responsible for neutrino masses may break lepton number symmetry (or, more generally, $B - L$), a symmetry which is preserved in the SM. A widely discussed class of such models is based on the seesaw mechanism [365–369], in which a set of right-handed neutrinos with Majorana masses exist in addition to the SM leptons. In the simplest models, the only new particles are right-handed neutrinos and the tiny observed left-handed neutrino masses suggest that the right-handed neutrinos have very small couplings relative to SM fermions. If the right-handed neutrinos are within kinematic reach of current experiments, these small couplings can generically predict that the right-handed neutrinos have a long lifetime. Unfortunately, the same small couplings also tend to predict very tiny production rates, making the RHNs challenging to produce at colliders such as the LHC. We begin the study by considering the minimal scenario involving only the right-handed neutrinos in Section 7.1.

UV complete versions of the seesaw mechanism tend to have degrees of freedom beyond the new right-handed neutrino particles, including gauge bosons associated with a broken local $B - L$ symmetry or left-right symmetry. While each model can lead to differing signals at high-energy colliders depending on the details of the model, they naturally provide new production mechanisms for right-handed neutrino LLPs with masses in the GeV-TeV range. Enhanced right-handed neutrino production from $B - L$ gauge bosons and left-right symmetric models are discussed in Section 7.2.1 and Section 7.3.1 respectively. In these models, new particles associated with the seesaw mechanism can, in largely unexplored parameter ranges, lead to displaced vertex signatures at the LHC observable at the MATHUSLA detector. Observation of any of these signals will provide crucial insight into the origin of neutrino masses and potentially other physics beyond the standard model. Furthermore, the UV-completion of these scenarios can involve light scalar bosons which may also be produced via the gauge bosons and have their own displaced decays. Although such states are not directly connected with the neutrino mass generation mechanism, they may provide an avenue for discovery of the underlying framework. We study their phenomenology in Sections 7.2.2 and 7.3.2.

Following this we consider two alternative portals into the neutrino sector. In Section 7.4 the Higgs portal, motivated by unification, and in Section 7.5 the inert doublet portal, which connects a FIMP DM candidate (see Section 5.3) with discrete symmetries which generate the active neutrino masses. Finally in Section 7.6 we turn to global symmetry structures and consider points of enhanced residual symmetry which may lead to small couplings, in turn greatly suppressing decays and enhancing the lifetime of the RHNs.

7.1 Minimal “Sterile” Right-Handed Neutrinos⁴⁴

Type-I Seesaw mechanism and neutrino masses

The simplest implementation of the seesaw mechanism is the addition of n copies of right-handed neutrinos ν_R to the SM, which permits a mass term for all SM neutrinos. The renormalizable addition to the SM lagrangian reads

$$\mathcal{L}_N = i \bar{\nu}_{Ri} \not{\partial} \nu_{Ri} - f_{\alpha i} \bar{L}_\alpha \tilde{H} \nu_{Ri} - \frac{M_i}{2} \bar{\nu}_{Ri}^c \nu_{Ri} + \text{h.c.}, \quad (164)$$

⁴³Section Editors: Marco Drewes, Rabindra Mohapatra, Brian Shuve

⁴⁴S. Antusch, B. Batell, M. Drewes, O. Fischer, J. C. Helo, M. Hirsch, A. Ibarra, D. Gorbunov, J. M. No

with $\nu_{Ri}, i = 1, 2, \dots, n$ denoting n sterile neutrinos with Majorana masses M_i , H being the SM Higgs boson doublet⁴⁵ and $L = (\nu_L, e_L)^T$ is the SM lepton doublet; $f_{\alpha i}$ are the elements of a $3 \times n$ matrix of Yukawa coupling constants (the α indices take on the values of $e, \mu,$ and τ). When the Higgs field gains a non-zero vacuum expectation value, the Yukawa terms in Eq. 164 give rise to mass mixing between active and sterile neutrinos. In the field basis where the mass matrix is diagonal, the neutrino flavour states mix, which gives rise to neutrino oscillations among SM neutrinos. There is also mixing between active SM neutrinos and sterile Majorana neutrinos.

In the neutrino sector with mass term $(\overline{\nu_{L^c}}, \overline{\nu_R})\mathcal{M}^{(\nu)}(\nu_L^c, \nu_R)^T$, the $(3+n) \times (3+n)$ mass matrix $\mathcal{M}^{(\nu)}$ can be diagonalised by a unitary rotation

$$V^T \mathcal{M}_{(\nu)} V = \text{Diag}\{m_1, m_2, m_3, M_1 \dots, M_n\}. \quad (165)$$

The resulting spectrum contains the three very light active (SM) neutrinos ν_α ($\alpha = e, \mu, \tau$) and n heavy states, denoted by N_i with masses M_i . Since the mixing between ν_L and ν_R is tiny, it is convenient to perform the diagonalisation in two steps. We first block-diagonalise $\mathcal{M}_{(\nu)}$ into a light neutrino mass matrix M_ν and a heavy neutrino mass matrix M_N . It is convenient to introduce a so-called Dirac mass matrix

$$(M_D)_{\alpha i} \equiv \frac{v}{\sqrt{2}} f_{\alpha i}$$

where $v = 246$ GeV is the Higgs vacuum expectation value. Suppose, its elements are small compared to the sterile neutrino masses. Then after rotation the active neutrinos mix with one another and have a non-diagonal mass matrix

$$M_\nu = -M_D \frac{1}{M_N} M_D^T \quad (166)$$

while the masses in the sterile neutrino sector remain almost intact.

$$M_N \simeq \text{diag}(M_1, \dots, M_n). \quad (167)$$

The n heavy neutrino states therefore have mass

$$m_{N_i} = M_{N_i} \simeq M_i. \quad (168)$$

The rotation induces a mixing between active and sterile neutrinos parameterized by the mixing matrix θ ,

$$\theta = M_D M_N^{-1}, \quad (169)$$

so that the mass eigenstates of the active neutrinos are related to the SM and Majorana neutrino flavour states as follows,

$$\nu_i = (U_{PMNS}^\dagger)_{i\alpha} \nu_\alpha - (U_{PMNS}^\dagger \theta)_{ij} \nu_{Rj}^c. \quad (170)$$

with U_{PMNS} being the PMNS mixing matrix in the active neutrino sector. The formulae above show that the active-sterile mixing angles $\theta_{\alpha j}$ control the sterile neutrino contribution to the active neutrino masses M_ν . The ν_R are gauge singlets, but the heavy mass eigenstates

$$N \simeq \nu_R + \theta^T \nu_L^c \quad (171)$$

feel a θ -suppressed weak interaction due to their mixing with the SU(2)-doublet components ν_L . They also directly couple to the Higgs field via the Yukawa interaction. In summary, the sterile neutrino mass eigenstates N_i have the following couplings:

$$-\frac{g}{\sqrt{2}} \overline{N^c}_i \theta_{i\alpha}^\dagger \gamma^\mu e_{L\alpha} W_\mu^+ - \frac{g}{\sqrt{2}} \overline{e_{L\alpha}} \gamma^\mu \theta_{\alpha i} N_i^c W_\mu^-$$

⁴⁵The tilde denotes the usual contraction of SU(2) indices with the antisymmetric tensor.




Fig. 35: In the $n = 2$ model, the ratios U_{ai}^2/U_i^2 with $U_i^2 = \sum_a U_{ai}^2$ that are large enough to be tested in experiments can in good approximation be expressed in terms of the parameters in the light neutrino mixing matrix U_{PMNS} [406, 407]. With the exception of the Majorana phase, all parameters in U_{PMNS} are constrained by neutrino oscillation data. These constraints can be translated into probability contours for the heavy neutrino flavour mixing pattern [428]. The different shades indicate the 1σ (darkest), 2σ and 3σ (lightest) regions that can be obtained from the NuFIT 3.1 global fit to neutrino oscillation data [443, 444], assuming a flat prior for the unconstrained Majorana phase. The results depend only mildly on the choice of this prior. For $n = 3$ the U_{ai}^2/U_i^2 in general depend on more unconstrained parameters. However, neutrino oscillation data still allows us to constrain the heavy neutrino flavor mixing pattern. In particular, for a hierarchical spectrum of light neutrinos, values far outside the region displayed here can only be realized with considerable tuning. Figure taken from ref. [428].

$$\begin{aligned}
& -\frac{g}{2\cos\theta_W}\overline{N}_i^c\theta_{i\alpha}^\dagger\gamma^\mu\nu_{L\alpha}Z_\mu - \frac{g}{2\cos\theta_W}\overline{\nu}_{L\alpha}\gamma^\mu\theta_{\alpha i}N_i^cZ_\mu \\
& -\frac{g}{\sqrt{2}m_W}\theta_{\alpha i}h\nu_{L\alpha}N_i - \frac{g}{\sqrt{2}m_W}\theta_{i\alpha}^\dagger h\overline{N}_i\nu_{L\alpha}.
\end{aligned} \tag{172}$$

The last line is the Yukawa coupling to the physical Higgs field h in unitary gauge re-expressed using the definition of θ and the relation $m_W = \frac{1}{2}vg$. For $M_i \lesssim 5$ GeV, RH neutrino decay modes are complicated by hadronic effects, but decays with at least one charged lepton in the final state have $\mathcal{O}(1)$ branching fraction. For $M_i \gtrsim 10$ GeV, the branching ratios follow the perturbative prediction, so (partially) leptonic decays constitute a $\mathcal{O}(1)$ majority fraction.

There are generically $7n - 3$ new physical parameters in the seesaw model. These comprise the mixing angles and phase in the matrix U_{PMNS} , the n heavy neutrino masses, and additional mixing angles and phases amongst the sterile neutrinos. The connection between these parameters and observables has been studied by various authors [398, 406, 407, 428, 435, 440–442]. The present knowledge of light neutrino oscillation parameters also allows us to obtain probability distributions for the patterns of mixing with heavy RH neutrinos i.e., the relative size of the mixings $U_{ai}^2 = |\theta_{ai}|^2$, see Fig. 35. For the minimal model with $n = 2$, all parameters in the Lagrangian can be constrained from measurements of the U_{ai}^2 if in addition the Dirac phase δ in U_{PMNS} is measured in light neutrino oscillation experiments, making this a fully testable model of neutrino masses and (possibly, see Section 6.3) leptogenesis [406, 407].

Connection with cosmology

The heavy right-handed neutrinos ν_R in the type-I seesaw model can, in addition to generating the light neutrino masses, also help to address important outstanding questions in cosmology. In particular, they can explain the baryon asymmetry of the universe via *low-scale leptogenesis* (see Section 6.3). This can occur either during the RH neutrinos' freeze out and decay ("freeze-out scenario") [364] or during their production ("freeze-in scenario") [379, 380], cf. Section 6.3. For masses in the GeV range, where MATHUSLA is expected to have the highest sensitivity, the freeze-in scenario is at work. The minimal number n of ν_{Ri} for which this mechanism works is $n = 2$ [380], which is also the minimal number that is required to explain the observed neutrino oscillation data within the seesaw mechanism. This minimal scenario has been studied by a number of authors [380, 398–411, 416, 417]. For $n \geq 2$, the additional states can either have similar properties and participate in leptogenesis [403–405, 418–420], or some subset of RH neutrinos can have very different properties from those responsible for leptogenesis. An interesting possibility is that very "sterile" RH neutrinos with tiny mixing angles $U_{\alpha i}^2 < 10^{-8}$ and masses in the keV range are candidates for Dark Matter (DM) [445, 446], see [397] for a recent review. This possibility is realised, for example, in the *Neutrino minimal extension of the SM* (ν MSM) [380, 381], see [396] for a review. From the viewpoint of MATHUSLA, the ν MSM is equivalent to the minimal $n = 2$ scenario because the couplings of the DM candidate must be so feeble that it is not produced efficiently at colliders. Moreover, its contribution to the seesaw mechanism and leptogenesis is negligible. However, the GeV-scale neutrinos that can be probed at MATHUSLA do play an important role in generating the low-scale lepton asymmetry [398, 402] necessary for production of DM candidate [447]. Since this scenario has been studied in much detail, we use it as a benchmark in what follows.

One requirement for successful leptogenesis is that the oscillating sterile neutrinos must remain out of equilibrium down to temperatures near or below the electroweak phase transition; otherwise, the primordial lepton asymmetry is destroyed and no appreciable baryon number asymmetry is produced. This requirement places an upper limit on the sterile-active mixing $U^2 = \text{tr}(\theta^\dagger \theta)$, which for $n = 2$ can be roughly approximated as [399]

$$\left(\frac{m_N}{10 \text{ GeV}}\right) \left(\frac{U^2}{10^{-8}}\right) < 1, \quad (173)$$

where more recent studies suggest that an order of magnitude larger mixing angles are possible [406, 409]. For $n > 2$ this upper limit is believed to be weaker [420]. In particular, for $n = 3$, the leptogenesis region extends to relatively large mixing angles $U^2 \sim 10^{-5}$ [420, 448], all the way up to the DELPHI bounds in Fig. 36. A lower limit on U^2 comes from the requirement for the two sterile neutrino to give a contribution to active neutrino masses large enough to explain the light neutrino oscillation data in the active neutrino sector, namely the so-called atmospheric neutrino mass $m_{atm} \simeq 0.05 \text{ eV}$. This limit can be approximated as

$$\left(\frac{m_N}{10 \text{ GeV}}\right) \left(\frac{U^2}{0.5 \times 10^{-11}}\right) > 1. \quad (174)$$

Also this bound is significantly weakened for $n > 2$ [449, 450].

There is, however, a lower bound from cosmology that is independent of n , neutrino oscillation data and leptogenesis. The N_i are unstable particles, and their decay in the early universe can modify the formation of primordial elements or the Cosmic Microwave Background. The requirement to decay sufficiently long before primordial nucleosynthesis implies that they must be heavier than about 100 MeV [451] unless they are so feebly coupled that they are effectively stable in the early universe, in which case they cannot significantly contribute to the seesaw mechanism [452].

Experimental Projections and MATHUSLA Sensitivity Estimate

MATHUSLA can probe the sterile neutrinos N in a region of parameter space similar to that accessible with present and proposed fixed-target experiments (see e.g. [386, 453] for details) like the SHiP project

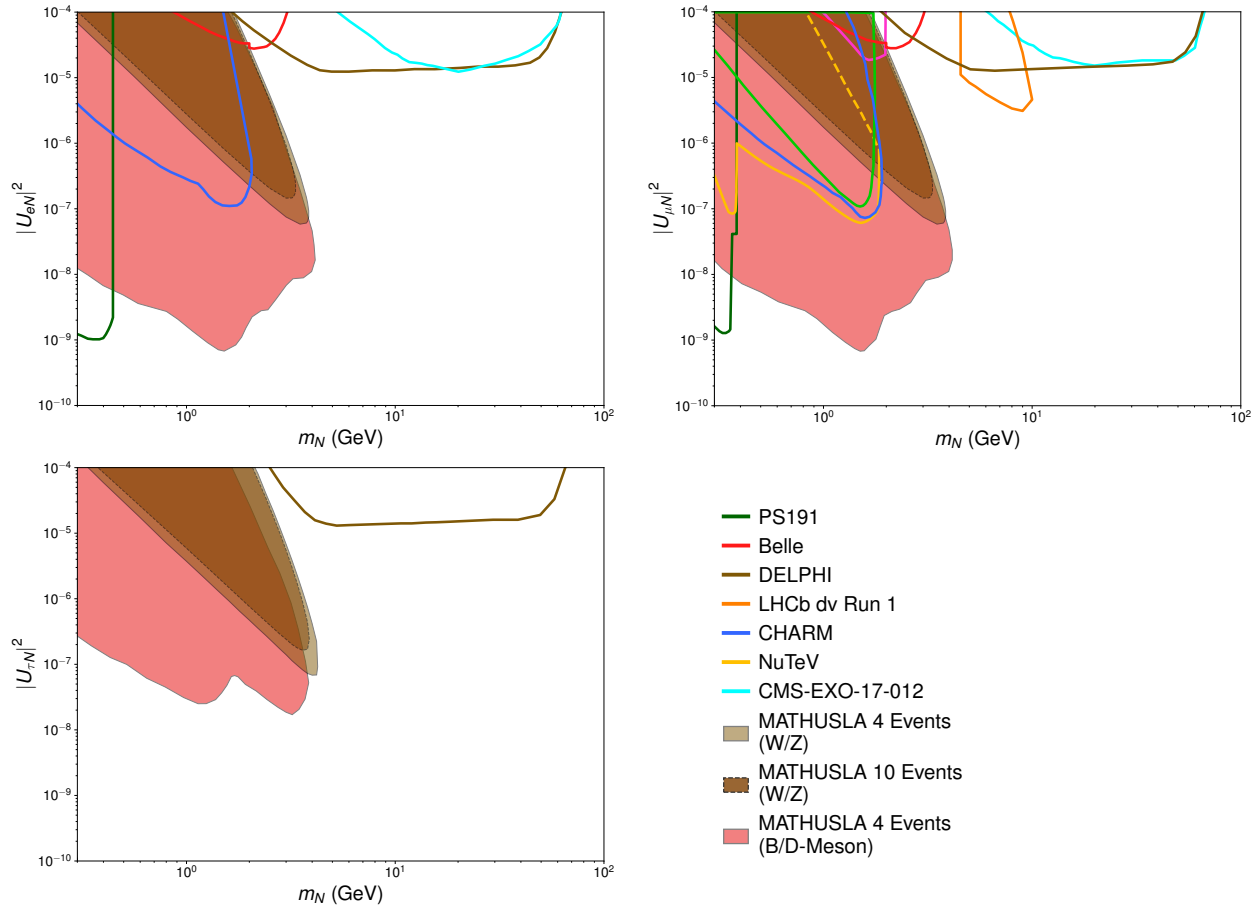


Fig. 36: Projected MATHUSLA sensitivity (assuming the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1) in the $(m_N, |U_{\alpha N}|^2)$ plane to sterile neutrinos, N , produced in W/Z decays (brown regions) and in B/D -meson decays (light red region) for $\alpha = e$ (top-left), $\alpha = \mu$ (top-right) and $\alpha = \tau$ (bottom-left). Also shown are the present exclusion limits (solid lines) at 90% C.L. from Belle (as given in the Erratum) [455] (red), PS191 [456] (using the reinterpretation from [441]) (dark-green), CHARM [457] (blue), NuTeV [458, 459] (yellow; the dashed line shows the decrease in sensitivity due to the short-lived nature of N , as taken from [455]), NA3 [460] (purple), BEBC [461] (light green), and at 95% C.L. from DELPHI [462] (brown), LHCb (using the reinterpretation from Ref. [425] of the displaced vertex search in Ref. [463]) (orange) and $\sqrt{s} = 13$ TeV CMS [391] (cyan).

at CERN [219, 453, 454] and similar facilities (e.g. based on DUNE or T2K). In such cases the sterile neutrinos are produced in leptonic meson decays, which are kinematically limited to be sensitive to masses $m_N \lesssim 5$ GeV. The mass region $m_N > 5$ GeV is inaccessible from meson decays (and in some experiments, even the region $m_N > 2$ GeV is inaccessible due to the low production rate of B -mesons).

At the LHC, the main production mechanism for the sterile neutrinos N are rare decays of heavy flavour mesons (B and D mesons for $M_N < 5$ GeV and $M_N < 2$ GeV, respectively) and decays of W, Z weak bosons into SM leptons and sterile neutrinos, $W^\pm \rightarrow \ell^\pm N$ and $Z \rightarrow \nu N$ (for $M_N < 80$ GeV and $M_N < 91$ GeV, respectively). In the following we estimate the sensitivity of MATHUSLA to sterile neutrinos in this mass range. (Note that the LHC can produce heavy neutrinos at higher masses via DY production, but in that case the heavy neutrino is typically not long-lived, with some exceptions, see Section 7.6.)

We adopt a phenomenologically driven approach to determining the MATHUSLA sensitivity, and we consider a simplified model of only one sterile neutrino N with mass m_N and mixing $U_{\alpha N}$ with the active flavours $\alpha = e, \mu, \tau$. The sterile neutrino decay rate Γ_N is given in Refs. [386, 464] (see also [465])

and scales parametrically like (neglecting phase space, color factors, and overall constants)

$$\Gamma_N \sim G_F^2 m_N^5 \sum_{\alpha} |U_{\alpha N}|^2. \quad (175)$$

For $m_N \ll m_W$ and small mixing $|U_{\alpha N}|^2$, the decay length is macroscopic. MATHUSLA is most sensitive to the parameters yielding a sterile neutrino decay length of ~ 200 m, which implies

$$L \simeq 2\gamma \left(\frac{3 \text{ GeV}}{m_N} \right)^5 \left(\frac{10^{-9}}{U^2} \right) 200 \text{ m},$$

where $\gamma = E_N/m_N$ is the sterile neutrino γ -factor and $U^2 = \sum_{\alpha} |U_{\alpha N}|^2$. See also Fig. 42 (a).

The strategy we follow in deriving the MATHUSLA sensitivity is similar to earlier proposals for sterile neutrino N searches via displaced vertices (see e.g. [421]). For the case of weak boson decays, we study the sensitivity to $\alpha = e, \mu, \tau$ via the processes $pp \rightarrow W^{\pm} \rightarrow \ell^{\pm} N$ and $pp \rightarrow Z \rightarrow \nu N$ (for $\alpha = \tau$ and $m_N < m_{\tau}$ the process $pp \rightarrow W^{\pm} \rightarrow \tau^{\pm} \nu$ with a subsequent tau decay into the sterile neutrino N also has to be taken into account) at LHC with $\sqrt{s} = 14$ TeV. We compute the cross-sections using Madgraph5_aMC@NLO [101] and Pythia 8.2 [158], and assume an integrated luminosity of $\mathcal{L} = 3000 \text{ fb}^{-1}$. The probability of detecting sterile neutrino decays at MATHUSLA is calculated using the corresponding simulated kinematic distributions and geometrical acceptance $\epsilon_{\text{geometric}}$ of MATHUSLA ($\epsilon_{W^+} \simeq 0.026$, $\epsilon_{W^-} \simeq 0.038$, $\epsilon_Z \simeq 0.029$ on average for $m_N \ll m_W$). The decay lengths are calculated from [421], and in determining the sensitivity to MATHUSLA including all the N decay modes that contain at least two charged particles.

We also estimate the sensitivity of MATHUSLA to sterile neutrinos produced in rare D and B meson decays, specifically focussing on the following four channels: $D \rightarrow K \ell N$, $D_s \rightarrow \ell N$, $B \rightarrow D \ell N$, $B \rightarrow \ell N$ (again, for $\alpha = \tau$ and $m_N < m_{\tau}$ the production of tau leptons with a subsequent tau decay into N also plays an important role). Charm and bottom production at $\sqrt{s} = 14$ TeV are simulated with Pythia 8.2 [158]. A dedicated simulation is then used to decay the mesons to sterile neutrinos and compute the probability for the sterile neutrinos to decay visibly within the MATHUSLA detector.

The MATHUSLA 4 event (“exclusion”) and 10 event (“discovery”) sterile neutrino sensitivities (under the assumption of zero background) in the $(m_N, |U_{\alpha N}|^2)$ plane are shown in Fig. 36 for $\alpha = e$ (top-left), $\alpha = \mu$ (top-right) and $\alpha = \tau$ (bottom-left). In each case, it is assumed that the N mixes only with a single flavour of SM neutrino and the other mixing angles are zero. The results are shown together with the present limits from Belle [455] (as given in its Erratum), DELPHI [462], CHARM [457], NuTeV [458, 459], PS191 [456] (using the reinterpretation from [441]), $\sqrt{s} = 13$ TeV CMS [391], LHCb (using the reinterpretation from Ref. [425] of the displaced vertex search Ref. [463] for masses $m_N > 4.5$ GeV), BEBC [461] and NA3 [460].⁴⁶ MATHUSLA is then projected to significantly surpass the present sensitivity to sterile neutrino masses in the few-GeV range, where the sterile neutrino is long-lived at sufficiently large mixing angles to produce an appreciable number of sterile neutrinos N at the LHC. For the minimal case $n = 2$, as in the ν MSSM, a sizable part of the parameter space with $m_N < 3$ GeV and an active-sterile mixing between Eq. (173) and Eq. (174), for which leptogenesis is possible, can be accessed with MATHUSLA.

For the case $\alpha = e$, we note that current constraints from the absence of neutrinoless double beta ($0\nu\beta\beta$) decay can place stringent limits in the $(m_N, |U_{eN}|^2)$, as shown e.g. in Fig. (1)-(3) from [421]. However, for $0\nu\beta\beta$ decay one needs to sum over all (virtual) mass eigenstates and in the presence of non-zero Majorana phases between different contributions this can lead to cancellations in the sum. This significantly weakens the $0\nu\beta\beta$ decay sensitivity compared to what can be achieved with MATHUSLA. This result holds, for instance, in the case of an approximate $B - L$ symmetry [467, 468], which both leads to larger mixing angles (increasing the MATHUSLA sensitivity) and suppressed contributions to $0\nu\beta\beta$ decay.

⁴⁶Note that the interpretation of several past experiments is a subject of controversy, cf. e.g. refs. [441, 466] for a discussion.

	z [m]	y [m]	x [m]
“standard”	[100,300]	[100, 120]	[-100, 100]
	z [m]	r [m]	ϕ [m]
“forward”	[20,40]	[5,30]	[0, 2π]

Table 3: Possible detector geometries for MATHUSLA at FCC-hh. The origin of the coordinate system is the IP, with $(z, y, x) = (0, 0, 0)$, with the z axis pointing along the direction of the beam, and y in the vertical and x in the horizontal direction. The “standard” geometry is the benchmark shown in Fig. 1 and assumed throughout this paper for HL-LHC. The “forward” detector variant is assumed to be symmetric in the angle ϕ (which rotates in the x - y plane) and with the fiducial detector volume starting outside of an inner circle with radius 5 m (to account for the beam pipe).

We also stress that both CMS and ATLAS detectors can perform searches for long-lived sterile neutrinos using displaced vertices. The main detectors will obviously have superior sensitivity for short RH neutrino lifetimes, but in the long-lifetime regime, the acceptance for LLP decays at MATHUSLA is about the same as for ATLAS or CMS. As discussed in Section 3.2.3, this means that MATHUSLA will have superior sensitivity if the main detector LLP search suffers from any bottlenecks due to triggering, cut efficiencies, requirements on the LLP decay, or backgrounds. For $m_N \lesssim 5$ GeV regime, many of these in some way restrict the sensitivity of main detector searches, since the low mass means that triggering and reconstruction of the DV will likely suffer from some inefficiencies and backgrounds, see Section 3.2.3.⁴⁷ The true performance of ATLAS/CMS searches will depend critically on details of the HL-LHC detector upgrades, but it is expected that MATHUSLA will have significantly better sensitivity to these light RHNs than the main detectors.

Regarding future pp colliders, the Future Circular Collider (FCC) in its hadron-hadron mode (FCC-hh) [470–472] (or, equivalently, the SppC [473]) would be excellent facilities to search for sterile neutrinos with very long lifetimes. We investigate the potential sensitivity of two variants of MATHUSLA at FCC-hh, namely the “standard” surface version used as a benchmark in this document, see Fig. 1, and an alternative “forward” version in the shape of a cylindrical ring aligned with the beamline, as defined in [1], with respective detector geometries given in Table 3. We consider sterile neutrino production through charged and neutral Drell-Yan processes. The cross-sections are evaluated for $\sqrt{s} = 100$ TeV at Leading Order with WHIZARD [474, 475] and Madgraph5_aMC@NLO [101] using the parton distribution function CTEQ6L (neglecting theoretical uncertainties and uncertainties on the input parameters), and we perform a similar analysis to the one carried above for the MATHUSLA sensitivity to sterile neutrinos from W , Z decays, considering in this case a total integrated luminosity of 20 ab^{-1} , as suggested in Ref. [476].

In Fig. 37 we show the sensitivity (4 events) in the $(m_N, |U_{\alpha N}|^2)$ plane for the two MATHUSLA variants at FCC-hh, i.e. the “standard” surface version (dotted) and the alternative “forward” version (dashed), for $\alpha = e$ (top-left), $\alpha = \mu$ (top-right) and $\alpha = \tau$ (bottom-left).⁴⁸ We also show the sensitivities for MATHUSLA at HL-LHC from Fig. 36, as well as the expected sensitivities from other proposed facilities. In the relatively near future, experiments like NA62 [427, 428] and SHiP [219, 429, 469] could explore new regions of parameter space. Note that the reach projections we show for SHiP must be regarded as preliminary. The solid blue line includes secondary B production in the fixed target. On the other hand, the contribution from B_c production is not yet understood and needs further study. The perturbative prediction for $\sigma(B_c)/\sigma(B)$ at SHiP is roughly two orders of magnitude below the measured value at the LHC [477, 478], but given the unknown non-perturbative effects, it is in principle possible

⁴⁷However, proposals exist for search strategies to minimize backgrounds for models with light sterile neutrinos [422].

⁴⁸For $\alpha = \tau$ we choose to only show the region $m_N > 2$ GeV. For lighter masses we expect a slight departure from the above sensitivities (in particular for $m_N < m_\tau$, due to the contribution from Drell-Yan tau lepton production).




Fig. 37: Projected sensitivity (4 events) in the $(m_N, |U_{\alpha N}|^2)$ plane to sterile neutrinos, N , produced in W/Z decays at FCC-hh for MATHUSLA “standard” benchmark surface version (dotted brown) and “forward” version (dashed brown), cf. Table 3 and Fig. 1, for $\alpha = e$ (top-left), $\alpha = \mu$ (top-right) and $\alpha = \tau$ (bottom-left, only shown for $m_N > 2$ GeV, see text for details). Also shown are the MATHUSLA sensitivities for HL-LHC from Fig. 36, as well as the projected future sensitivity of various facilities: SHiP [219, 429, 469], FCC-ee [433], CEPC [411, 434], ILC [411, 434] and NA62 [428] (see also [427]). The light blue shaded region indicates the uncertainty in the SHiP reach due to B_c production, with $\sigma(B_c)/\sigma(B)$ set to the LHC value at its outer boundary. *Note that for $\alpha = \tau$, some of these sensitivity projections assume τ -dominated mixing scenarios, rather than only τ -mixing. For a τ -dominated scenario, the MATHUSLA reach is very similar to the shown $\alpha = e, \mu$ sensitivity (see [428] for a demonstration of this effect).* The present limits on $(m_N, |U_{\alpha N}|^2)$ from Fig. 36 are shown as a light-grey region. The region excluded from primordial nucleosynthesis (BBN) is shown in medium-grey. The upper limit on $|U_{\alpha N}|^2$ from viable leptogenesis for the minimal case $n = 2$ (assuming normal neutrino mass hierarchy, see Ref. [406]) and the lower exclusion on $|U_{\alpha N}|^2$ from the active neutrino oscillation data for the minimal case $n = 2$ (for normal neutrino mass hierarchy, see Ref. [406]) are respectively shown as a black dotted line and a dark-grey region. For $n = 3$, the leptogenesis region extends up to the present DELPHI bounds [420, 448].

that this prediction is too small by up to two orders of magnitude. Therefore, the blue shading indicates the uncertainty in the SHiP reach due to B_c production, where $\sigma(B_c)/\sigma(B)$ is set to the measured LHC value as an absolute upper limit on the outer boundaries of the shaded region. A LBNE/DUNE-like facility [431] could have the best sensitivity for very small mixing angles at sub-GeV RH neutrino masses, but detailed estimates for DUNE’s updated detector design are not yet available. On time scales relevant for the FCC-hh, other future colliders like FCC-ee [433], CEPC [411, 434] and ILC [411, 434] would greatly extend sensitivity. The envelope of the excluded region from current experiments (from Fig. 36) is shown in grey, together with the constraint from the generation of light active neutrino masses via the see-saw mechanism (for normal neutrino mass hierarchy, see *e.g.*, Ref. [406]) and from the viability of leptogenesis (also for normal neutrino mass hierarchy) [406] for the minimal scenario $n = 2$. We also require the sterile neutrinos to decay before primordial nucleosynthesis ($\tau_N \leq 1$ s).

Clearly, a MATHUSLA-like detector at a future 100 TeV collider would probe previously unexplored regions of RHN parameter space. Furthermore, MATHUSLA at the HL-LHC and SHiP explore similar and complementary regions of parameter space.

7.2 The $B - L$ Gauge Portal

The active neutrino masses can also have weak-scale origins if the Majorana masses of the sterile neutrinos arise from a local $B - L$ symmetry that is broken at the weak scale, implementing the type-I seesaw at low energies [365–369]. We consider two scenarios: one with new vector bosons at masses well below the weak scale, with sensitivity to long-lived right-handed neutrino decays; and a UV-complete model of TeV-scale $B - L$ breaking with long-lived exotic Higgs scalar states.

7.2.1 Low-Mass Z' ⁴⁹

A simple and well-motivated extension of the SM is a model based on a local $U(1)_{B-L}$ symmetry. Neutrino masses naturally emerge in this model once the $U(1)_{B-L}$ symmetry is spontaneously broken, resulting in a type-I seesaw mechanism. In particular, unlike in the SM, three right-handed neutrinos are required in the $B - L$ model to cancel gauge anomalies. We assume that the $B - L$ gauge symmetry does not contribute to electric charge so that its coupling can be chosen arbitrarily small. As we demonstrate below, this gives rise to a particularly attractive discovery scenario for MATHUSLA. The RHN decays via the same small mixing angle as in the minimal model discussed in Section 7.1, but acquires an additional production mode through Drell-Yan like processes involving the on-shell Z' gauge boson of the $U(1)_{B-L}$ broken gauge symmetry. MATHUSLA and the main detectors will then cover different but equally motivated regions of the scenarios parameter space.

In this context, it is perhaps natural to expect the N mass to be correlated with the $B - L$ gauge boson mass since they are both governed by symmetry breaking in the $B - L$ sector. In this case, the $B - L$ gauge interaction opens up new production channels for the RHNs and potentially allows accelerator experiments to probe the seesaw mechanism in the laboratory. Here we explore the sensitivity of MATHUSLA to a particular phenomenologically viable benchmark scenario in this model. We focus on the following simplified approach following Ref. [479]. The effective interaction Lagrangian after electroweak and $B - L$ symmetry breaking is taken to be

$$\mathcal{L} = g' V_\mu \left(\sum_{\text{SM}} Q_{B-L} \bar{\psi} \gamma^\mu \psi + \bar{N} \gamma^\mu \mathcal{P}_L N \right) + U_{\mu N} \frac{g_W}{\sqrt{2}} (\bar{\mu}_L \gamma^\mu W_\mu^- \mathcal{P}_L N + \text{h.c.}) + \dots \quad (176)$$

We have included the sterile neutrino N , the $B - L$ gauge boson V_μ , along with the relevant SM fields. Under $B - L$, the SM lepton fields have charges -1 , SM quark fields have charges $+1/3$, and the N fields have charge $+1$.

⁴⁹Brian Batell




Fig. 38: Dark gray contour corresponds to 50 events in MATHUSLA from sterile neutrinos in the gauged $B - L$ model (gray) for the case $m_V/m_N = 3$, $g' = 10^{-3}$. Also shown are limits from DELPHI [462] (teal) and a projection from a proposed displaced vertex search at ATLAS and/or CMS during the high luminosity run (blue) [479]. The grey shaded region indicates the parameters favored in a minimal type-I see-saw to give light neutrino masses ranging from $\sqrt{\Delta m_{\text{sol}}^2}$ to the Planck upper limit [284]. Larger couplings are also allowed in $B - L$ models. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

For simplicity, we consider only one sterile neutrino which mixes exclusively with the muon flavoured SM neutrino ν_μ (similar results hold for mixing with ν_e). This Lagrangian has four unknown parameters: M_V , g' , M_N , and $U_{\mu N}$. Taking the simplest seesaw motivated parameter choices corresponding the scale $\sqrt{|\Delta(m_\nu^2)_{\text{atm}}|} \sim 0.05$ eV suggests a mixing angle

$$U_{\mu N}^2 \approx \frac{m_\nu}{m_N} \sim 5 \times 10^{-11} \times \left(\frac{1 \text{ GeV}}{m_N} \right). \quad (177)$$

Such a small mixing angle is difficult to directly probe if the only production channels for N occur through the weak interactions via mixing. However, in the $B - L$ gauge extension there are additional channels present. Here we consider $pp \rightarrow V \rightarrow NN$ at the LHC. The production rate can be dramatically enhanced over the usual weak-interaction production process. Once produced, the RHN will decay via mixing through the weak interactions. If the RHN is lighter the W boson, it will generically be long-lived. The mixing angle can, of course, also be larger than suggested by Eq. 177 due to cancellations in the active neutrino mass matrix, Eq. 166 (this is particularly true in $B - L$ models or other models with lepton flavour symmetries), in which case production modes from neutrino mixing can potentially be competitive.

In Fig. 38 we show the sensitivity of MATHUSLA in the $m_N - |U_{\mu N}|^2$ plane. For concreteness, we choose the parameters $m_V = 3m_N$ and $g' = 10^{-3}$, which is below the sensitivity of current direct dilepton searches for V [480]. Since the only role of g' is in the production rate of N , the sensitivity to $U_{\mu N}^2$ scales inversely to g'^2 . Thus, the 50 event contour in Fig. 38 is equivalent to a 5 event contour with $g' \sim 3 \times 10^{-4}$. This exceeds the sensitivity of optimistic projections for the high-luminosity LHC, which could discover Z' in Drell-Yan production for $g' \gtrsim 5 - 10 \times 10^{-4}$ [479, 480]. It also complements the parameter space for displaced vertex searches at the LHC [479].

We observe that MATHUSLA has the potential to probe the seesaw-motivated parameter space of theories with a local $B - L$ symmetry where both vector and RH neutrino masses are at or below the weak scale. We also remark that our findings apply quite generally to new light gauge bosons that couple to N , and in some of these cases the typical dilepton constraints for $B - L$ models are significantly relaxed,

giving MATHUSLA sensitivity to an otherwise uncovered parameter space. This parameter space could be otherwise challenging to probe at the LHC due to low reconstruction efficiencies for low-mass LLPs and possible sources of background in high-luminosity running (see Section 3.2.3).

7.2.2 TeV-Scale $B - L$ Symmetry Breaking and Long-Lived Scalars⁵⁰

In this section, we examine whether a Higgs boson that breaks the $B - L$ symmetry at the TeV scale can be accessible at MATHUSLA. We specifically consider the low-mass regime for this scalar field, which can arise if the $B - L$ symmetry is radiatively broken. Any evidence for such low mass scalars can play a crucial role in the elucidating the seesaw mechanism. The mass and couplings of this new Higgs field are, to a large extent, *a priori* unrestricted and we show that certain parameter ranges of this boson can be probed at the MATHUSLA detector.

Both the $U(1)$ and left-right symmetric completions of $B - L$ theories can have a light (\sim GeV-scale) neutral scalar field which will be long-lived. Depending on the details of the theories, the MATHUSLA detector may provide an appropriate venue for searching for these displaced vertices. In this section, we focus on simple $B - L$ models based on $SU(2)_L \times U(1)_{I_{3R}} \times U(1)_{B-L}$ local symmetry, and will discuss TeV-scale left-right models in Section 7.3.2 of which this is a subgroup. In this case the $B - L$ symmetry therefore clearly contributes to electric charge.

The SM fermions are charged under the gauge group $SU(2)_L \times U(1)_{I_{3R}} \times U(1)_{B-L}$ (with the gauge couplings g_L , g_R , and g_{BL} , respectively) as

$$Q = (u_L, d_L)^T : \left(\mathbf{2}, 0, \frac{1}{3} \right); \quad L = (\nu, e_L)^T : \left(\mathbf{2}, 0, -1 \right);$$

$$u_R : \left(\mathbf{1}, \frac{1}{2}, \frac{1}{3} \right); \quad d_R : \left(\mathbf{1}, -\frac{1}{2}, \frac{1}{3} \right); \quad e_R : \left(\mathbf{1}, -\frac{1}{2}, -1 \right).$$

Anomaly freedom requires that this model have three right-handed neutrinos (RHNs) N_i ($i=1,2,3$) with gauge quantum numbers $(\mathbf{1}, 1/2, -1)$. The minimal Higgs fields to break the symmetry to the SM gauge group are $H(\mathbf{2}, -1/2, 0)$ and $\Delta(\mathbf{1}, -1, 2)$ with the following Yukawa couplings:

$$\mathcal{L}_Y = h_u \bar{Q} H u_R + h_d \bar{Q} \tilde{H} d_R + h_e \bar{L} \tilde{H} e_R + h_\nu \bar{L} H N + f \bar{N}^c \Delta N + \text{H.c.} \quad (178)$$

Note that $\langle \Delta^0 \rangle = v_R$ breaks the gauge symmetry down to the SM gauge group, which in turn is further broken by $\langle H^0 \rangle = v_{\text{EW}}$ to $U(1)_{\text{em}}$. From the Yukawa interactions in Eq. 178 it is clear that after symmetry breaking this leads to the type-I seesaw formula for neutrino masses. In this model, $H_3 = \text{Re}(\Delta^0)$ is the light scalar candidate, though in principle it is allowed to have mass in a wide range starting from below GeV-scale to v_R [6, 7, 481, 482]. As argued in Ref. [6, 7], a small mass m_{H_3} can be stable under radiative corrections in the presence of direct couplings of H_3 to both the bosonic and fermionic particles.

We consider H_3 masses ≥ 100 MeV because lighter masses can be constrained by Big Bang Nucleosynthesis (BBN) (assuming $M_{H_3} \leq M_N$) and supernova energy considerations [7]. For M_{H_3} in the GeV range, the H_3 boson could be sufficiently long-lived and give rise to displaced vertex signatures accessible to the MATHUSLA detector.

The decay of the neutral scalar H_3 is dominantly governed by its mixing with the SM Higgs, parameterized by the angle $\sin \theta$. The tree level couplings of H_3 to the SM fermions are proportional to the SM Yukawa couplings, rescaled by the mixing angle $\sin \theta$, all of which are flavour conserving. If $m_{H_3} \lesssim$ GeV, it decays predominantly into the SM fermions at tree level, and into $\gamma\gamma$ and gg at one-loop level. The branching fractions do not depend on the mixing angles but only on H_3 mass, as all the couplings are universally proportional to the mixing angle. The H_3 lifetime can be long, when the couplings of the SM Higgs to sub-GeV particles are small.

⁵⁰P. S. Bhupal Dev, Rabindra Mohapatra, Yongchao Zhang

Flavour-changing couplings, such as $H_3\bar{s}b$, can arise at one-loop level through mixing with the SM Higgs, which leads to the flavour-changing rare decays of the K and B mesons mediated by the light scalar such as $B \rightarrow KH_3 \rightarrow K\mu^+\mu^-$ (see discussions of the flavour limits in the SM+S model in Section 8.4). However, the flavour limits in the $U(1)_{B-L}$ model are much weaker than those in the LR model, as in the latter case all these flavour-changing couplings occur at the tree level. The flavour limits in the $U(1)_{B-L}$ model could yet reach the level of 10^{-4} or below in future high-intensity experiments such as SHiP and DUNE; however, the detailed reach depends sensitively on the H_3 mass and the flavour decay modes. For details, see Ref. [7].

In the $U(1)_{B-L}$ model, the light scalar H_3 could be produced either from mixing with the SM Higgs (scalar portal) or through the gauge interaction with the heavy Z_R boson (gauge portal). The scalar portal production scenario is equivalent to the SM+S simplified model discussed in Section 8.4. As shown in Fig. 55, MATHUSLA can cover significant part of that scenario's parameter space that are inaccessible to the main detectors and highly complementary to the reach of proposed experiments like SHiP.

The light scalar H_3 could also be produced through the gauge portal, i.e. via interactions with the Z_R gauge boson. This provides another channel to probe the theory, where the light scalar H_3 could be produced in association with the heavy Z_R boson by analogy with the SM Higgs-strahlung process. The Z_R which further into the SM quarks and charged leptons (for simplicity we have neglected here the heavy and light neutrino decay modes), i.e.

$$pp \rightarrow Z_R^* \rightarrow H_3 Z_R, \quad Z_R \rightarrow q\bar{q}, \ell^+\ell^-. \quad (179)$$

H_3 could also be produced by the vector-boson fusion (VBF) of two heavy Z_R bosons, i.e. $pp \rightarrow Z_R^* Z_R^* jj \rightarrow H_3 jj$ (with $j = u, d, s, c$), which is subleading to the associated production mode with an on-shell $Z_R \rightarrow jj$; thus, we focus on the associated production mode.

With the heavy Z_R taking away most of the energy in the final state, the light scalar H_3 tends to be soft, with a transverse momentum $\lesssim 100$ GeV for most of the events. Therefore only a small portion could arrive at MATHUSLA, similar to the scalar portal. The high- p_T jets/leptons (typically \gtrsim TeV) allow events to pass the trigger in ATLAS and CMS. The combination of the high- p_T jets/leptons and the LLP which should also mean that backgrounds for this search will be low, but the low mass and high boost of the LLP might still mean that reconstruction is difficult (and there may still be some backgrounds; for now, we assume LHC has zero background, and for a fuller discussion, see Section 5.1). We take possible inefficiencies into account by showing LHC curves with representative LLP reconstruction efficiency of 1 and 0.1. In Fig. 39, we show the parameters giving a rate of at least four signal events after requiring that the light scalar decays in the tracker with approximate decay length of $1 \text{ cm} \lesssim \gamma c\tau_0 \lesssim 1 \text{ m}$, where the boost factor $\gamma = E_{H_3}/m_{H_3}$ has been taken into consideration. The ultimate LHC sensitivity likely lies somewhere between these curves.

A major limiting factor on the sensitivity to associated production of H_3 is due to dilepton limits on Z_R [483–486]. Considering only benchmark points that are not yet excluded by direct searches for Z_R , an optimistic benchmark scenario is $g_R/g_L = 0.835$, for which the Z_R mass limit is 3.64 TeV. When the g_R coupling becomes smaller, the gauge coupling $g_{BL} = g_Y g_R / \sqrt{g_R^2 - g_Y^2}$ becomes larger which would enhance the production cross-section of Z_R at the LHC and makes the dilepton mass limits on Z_R more stringent.⁵¹ In this optimistic benchmark scenario, the cross-section in the gauge portal is $\sigma(pp \rightarrow H_3 JJ) = 0.97 \text{ fb}$ after applying a k -factor of 1.2 (J runs over all the SM quark and charged leptons). This rate does not depend on the mixing angle $\sin\theta$ in the scalar sector. We apply simple cuts $p_T(J) > 25 \text{ GeV}$ and $\Delta\phi(JJ) > 0.4$ on all the quark and charged leptons in MadGraph5_aMC@NLO [101]. With the small production cross-section, suppressed by the large Z_R

⁵¹With the dilepton limits on Z_R becoming stronger at the LHC, it is very likely that the Z_R boson is so heavy that we could not have 4 events at MATHUSLA, even with the ultimate luminosity of 3000 fb^{-1} .




Fig. 39: LLP search sensitivities at LHC and MATHUSLA in the $U(1)_{B-L}$ model, with $\sqrt{s} = 14$ TeV and an integrated luminosity of 3000 fb^{-1} , in the gauge portals by coupling to the Z_R boson with a benchmark gauge coupling of $g_R = 0.835g_L$. (See also Section 8.4 and Fig. 55 for constraints from production through the scalar portal, as well as other limits on the mixing angle. For $m_{H_3} \lesssim 5$ GeV, those MATHUSLA sensitivity is better than the gauge portal sensitivity shown here.) For the LHC reaches we assume a signal efficiency factor of 1 or 0.1, with at least 4 signal events for both LHC and MATHUSLA. Due to the small mass, the lower efficiency is likely more realistic, but the assumption of no background at the main detector is likely justified due to the hard dilepton and dijet pair produced from decay of the on-shell Z_R . Note that MATHUSLA searches for the production of low-mass H_3 with $m_{H_3} \lesssim \text{GeV}$ from Z' decay may suffer from some backgrounds or lower reconstruction efficiency depending on the final detector design, see Section 3.1.4. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

mass, only a narrow region of $m_{H_3} - \sin \theta$ could be probed by the MATHUSLA detector, as shown in Fig. 39. With much more signal events expected at ATLAS/CMS, the displaced vertex searches is largely complementary to the ULLP searches at MATHUSLA. In other $U(1)$ models, *e.g.*, those motivated from Grand Unified Theories (GUTs) [487], the LHC constraints on Z' mass might be somewhat weaker, and the production cross-section $\sigma(pp \rightarrow Z'H_3)$ gets larger; in such a scenario, a broader region of the scalar mass m_{H_3} and mixing angle $\sin \theta$ could be probed at the HL-LHC. In addition, it might also be more promising to test the light scalar at the future 100 TeV collider and the dedicated forward detector, in searches of the ULLP events [472, 481].

In summary, long-lived scalars can arise as a result of $B-L$ symmetry breaking in UV completions of the minimal sterile RH neutrino scenario. For $m_{H_3} \lesssim 5$ GeV, the best small-mixing-angle sensitivity would likely come from MATHUSLA searches for LLPs produced in meson decays, see Fig. 55. For larger scalar masses, the gauge portal likely provides the best sensitivity, and while main detector searches have excellent sensitivity in this regime, MATHUSLA will likely be able to expand coverage to somewhat lower mixing angles.

7.3 Left-Right Symmetric Model

7.3.1 Long-Lived Right-Handed Neutrinos in the Left-Right Model⁵²

The standard left-right symmetric model [488–491] has the gauge group $SU(2)_L \times SU(2)_R \times U(1)_{B-L}$ and provides an ideal setting for a low scale seesaw model for neutrino masses. Gauge couplings are denoted respectively as g_L , g_R and g_{BL} . The gauge charges of quarks and leptons in the LR gauge group

⁵²P. S. Bhupal Dev, Juan Carlos Helo, Martin Hirsch, Rabindra Mohapatra, Yongchao Zhang

are, respectively,

$$\begin{aligned} Q_L &= (u_L, d_L)^T : (\mathbf{2}, \mathbf{1}, 1/3); & L &= (\nu, e_L)^T : (\mathbf{2}, \mathbf{1}, -1); \\ Q_R &= (u_R, d_R)^T : (\mathbf{1}, \mathbf{2}, 1/3); & R &= (N, e_R)^T : (\mathbf{1}, \mathbf{2}, -1). \end{aligned} \quad (180)$$

In this setup, RHNs N are automatically included in the theory. At a scale v_R , the LR symmetry is broken to the SM gauge group. The minimal Higgs sector includes the fields $\Phi(\mathbf{2}, \mathbf{2}, 0)$ and $\Delta_R(\mathbf{1}, \mathbf{3}, 2)$. The Yukawa couplings are given by

$$\mathcal{L}_Y = h_u \bar{Q}_L \phi Q_R + h_d \bar{Q}_L \tilde{\phi} Q_R + h_e \bar{L} \tilde{H} R + h_\nu \bar{L} \phi R + f \bar{R}^c \Delta R + \text{H.c.} \quad (181)$$

After symmetry breaking we get the seesaw formula for neutrino masses; however, we emphasize that the $B - L$ group in the left-right model is very different from that studied above.

The charged-current and neutral-current interactions relevant for our analysis are

$$\mathcal{L} = \frac{g_R}{\sqrt{2}} (\bar{d} \gamma^\mu P_R u + V_{lN}^R \cdot \bar{l} \gamma^\mu P_R N) W_{R\mu}^- + \text{H.c.}, \quad (182)$$

where V^R is the neutrino mixing matrix in the right-handed sector. The gauge $W_{L,R}$ -boson states can be rewritten in terms of the mass eigenstates as:

$$\begin{aligned} W_L^- &= \cos \zeta \cdot W_1^- - \sin \zeta \cdot W_2^-, \\ W_R^- &= \sin \zeta \cdot W_1^- + \cos \zeta \cdot W_2^-, \end{aligned} \quad (183)$$

whith the mixing angle given by

$$\begin{aligned} \tan 2\zeta &= \frac{2g_L g_R M_{W_L}^2 \cdot \sin 2\beta}{g_R^2 M_{W_L}^2 + g_L^2 (M_{W_R}^2 - M_{W_L}^2)} \\ &\approx 2 \frac{g_R}{g_L} \frac{M_{W_L}^2}{M_{W_R}^2} \sin 2\beta. \end{aligned} \quad (184)$$

Here $\tan \beta = \kappa'/\kappa$ is the ratio of the two vev's of the bidoublet Higgs Φ .

The RHNs are typically produced in the on-shell decays of W_R . The RHN subsequently decays back into an off-shell W_R^* , which decays exclusively into the light SM quarks with almost a BR of 100%, *i.e.*, $N \rightarrow \ell W_R^* \rightarrow \ell jj$ where j are the jets from the quarks u, d, s, c . In the mass range under consideration, the widths of the three heavy neutrinos of the LR models are approximately [421]:

$$\Gamma_N \approx \frac{3G_F^2}{32\pi^3} m_N^5 \left(\frac{M_{W_L} g_R}{M_{W_R} g_L} \right)^4 [1 + \sin^2 2\beta] \sum_l |V_{lN}^R|^2, \quad (185)$$

where we neglected the masses of all the final state particles. In our numerical study we will consider a simplified case with only one heavy neutrino in the relevant mass range. If the RHN mass is order GeV then its *proper* lifetime for a W_R mass of a few TeV would be at the 100 m level:

$$\tau_N^0 \simeq (290 \text{ m}) \left(\frac{m_N}{2 \text{ GeV}} \right)^{-5} \left(\frac{M_{W_R}}{3 \text{ TeV}} \right)^4 \left(\frac{g_R}{g_L} \right)^{-4}. \quad (186)$$

If the N mass is even lighter, then the RHN can be produced in meson decays such as $D_s \rightarrow \ell N$, with the subsequent decay $N \rightarrow \ell \pi$ [492]. Both the production of N from mesons as well as their decays into lighter states are mediated by the W_R gauge interaction. The masses M_{W_R} , m_N and the gauge coupling g_R can thus be probed at dedicated beam-dump experiments such as SHiP [219, 492, 493], as shown in Fig. 40, in addition to high-energy colliders.




Fig. 40: Light RHN sensitivity in the minimal LR model from the (U)LLP searches at the $\sqrt{s} = 14$ TeV LHC (red) and MATHUSLA (blue) with the RHN produced from (on-shell) W_R decay. We also show the MATHUSLA prospects (orange, assuming the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1) from the decays of B mesons via $B \rightarrow \ell N$, for three different values of $g_R/g_L = 0.6, 1$ and 1.5 . To take account possible signal inefficiencies in DV reconstruction to reject backgrounds, we show curves for an efficiency factor of $\epsilon_{DV} = 1$ (left) and 0.1 (right) for the LHC reaches. Below these curves we can have at least 4 signal events for both LHC and MATHUSLA. The purple lines indicate the DV prospect at SHiP for $g_R = g_L$ [219, 492, 493].

In the minimal LR group, where the $SU(2)_R$ gauge symmetry is broken by a RH triplet scalar Δ_R , we have $M_{Z_R} > M_{W_R}$. Thus the dominant production of RHNs at the LHC is through the s -channel (on-shell) W_R : $pp \rightarrow W_R^{(*)} \rightarrow \ell N$, followed by the three-body decay $N \rightarrow W_R^* \ell \rightarrow \ell jj$ [392, 494]. With $M_{W_R} \gtrsim 3 \text{ TeV} (g_R/g_L)^4$, as required to satisfy the direct LHC constraints [388, 495], as well as the low-energy flavour-changing neutral current (FCNC) constraints [496, 497], the production cross-section could reach few tens of fb, depending on the W_R mass as well as the gauge coupling g_R . The sensitivity contours for MATHUSLA are shown in Fig. 40, assuming at least 4 signal events, for three different values of the gauge coupling $g_R/g_L = 0.6, 1$ and 1.5 . For concreteness, we have assumed only the electron flavour $\ell = e$ without RH leptonic mixing. Though the effective cross-section is small, due to the small effective solid angle, MATHUSLA is sensitive to light RHN with mass as low as ~ 1 GeV. For the purpose of illustration, we also show the proper lifetime of RHN for $g_R = g_L$, estimated from Eq. (186); for the values of $g_R \neq g_L$ the lifetime should be rescaled via $(g_R/g_L)^{-4}$ accordingly.

If kinematically allowed, the light RHN could also be produced at the LHC from the decays of D and B mesons, as in the simplest seesaw mechanism in Section 7.1. In the minimal LR model, the RHN decay and production are mediated by a heavy W_R boson; by comparing the decay width in Eq. (175) and (185), the MATHUSLA prospects on the effective heavy-light neutrino mixing angles in Section 7.1 can be cast onto the W_R mass in the LR model, depending on the g_R coupling. To be specific, three benchmark values of $g_R/g_L = 0.6, 1$ and 1.5 are shown in Fig. 40, with the RHN from B meson decays (orange lines). It is apparent that the meson decay prospects are largely complementary to those from the (on-shell) W_R decay in Fig. 40.

For the sake of comparison we also estimate the sensitivity to RHN LLPs at the main ATLAS or CMS detectors. We require that the RHN decay inside the tracker with a decay length $1 \text{ cm} \lesssim \gamma c\tau_0 \lesssim$

1 m, as shown in Fig. 40. To account for possible inefficiencies in the reconstruction of the displaced vertex, possible benchmark values for the DV efficiency are set to $\epsilon_{\text{DV}} = 1$ and 0.1 in the left and right panels of Fig. 40. Note in Fig. 40 that, even when the heavy W_R is off-shell, *i.e.*, $M_{W_R} \gtrsim 6$ TeV, the light RHN could yet be produced abundantly. In fact, regardless of whether the W_R boson is on-shell or off-shell, the RHN tends to have a huge boost factor of $\gamma \simeq m_{W_R^{(*)}}/2m_N \sim 10^3$. The decay products of N are consequently highly collimated, and the reconstruction of the LLP events would be rather challenging. In the optimistic case, depending on g_R , the general-purpose detectors at LHC could probe the *proper* lifetime τ_N^0 from ~ 10 m to below 0.01 cm, and the RH sector can be probed up to $M_{W_R} \simeq 20$ TeV for a large $g_R/g_L = 1.5$, which is largely complementary to the ULLP searches at MATHUSLA. With a more pessimistic $\epsilon_{\text{DV}} = 0.1$, the probable regions shrink significantly. The reach of ATLAS and CMS may be even worse when realistic efficiencies of reconstructing boosted LLPs and their associated backgrounds are taken into account, see discussion in Section 3.2.3. However, even with these optimistic main detector projections, it is clear that MATHUSLA can provide the best sensitivity in the low-mass regime $m_N \lesssim 5$ GeV which are very difficult to probe at ATLAS or CMS. Even at higher masses, MATHUSLA would likely extend the coverage provided by the main detectors.

7.3.2 Long-Lived Scalars in the Left-Right Model⁵³

As in the $U(1)_{B-L}$ model, the symmetry-breaking sector of the left-right model can yield interesting dynamics that can be probed at the LHC and detectors such as MATHUSLA. Due to the expanded gauge structure $SU(2)_L \times SU(2)_R \times U(1)_{B-L}$, this model has a bidoublet Φ and a RH Higgs triplet Δ_R in the scalar sector that breaks the gauge symmetry and is responsible for implementing the seesaw mechanism:

$$\Phi = \begin{pmatrix} \phi_1^0 & \phi_2^+ \\ \phi_1^- & \phi_2^0 \end{pmatrix} : (\mathbf{1}, \mathbf{2}, \mathbf{2}, 0), \quad \Delta_R = \begin{pmatrix} \Delta_R^+/\sqrt{2} & \Delta_R^{++} \\ \Delta_R^0 & -\Delta_R^+/\sqrt{2} \end{pmatrix} : (\mathbf{1}, \mathbf{1}, \mathbf{3}, 2). \quad (187)$$

There are three physical neutral scalars in this model: the SM Higgs h , a new heavy Higgs field H_1^0 , and the remnant of the $SU(2)_R$ -breaking scalar H_3 (see Ref. [481] for nomenclature of these scalars):

$$h \simeq \text{Re } \phi_1^0, \quad H_1 \simeq \text{Re } \phi_2^0, \quad H_3 \cong \text{Re } \Delta_R^0, \quad (188)$$

in the limit of $\langle \phi_1^0 \rangle \gg \langle \phi_2^0 \rangle$. There is almost no absolute *lower* mass limit on H_3 (except those from the cosmological and astrophysical observations such as BBN and supernovae which requires that $m_{H_3} \gtrsim 100$ MeV), which renders it to be the only LLP candidate in the scalar sector of minimal LR model [6, 7, 481]. The role of H_3 is analogous to the light scalar in the $U(1)_{B-L}$ case but their properties are very different as we now show.

In analogy with the $B-L$ model, the smallness of m_{H_3} in the LR model is also stable against loop corrections due to heavy particles in the model; thus the H_3 field has a wide range of viable masses, and at low masses it could be sufficiently long-lived to give displaced signatures at MATHUSLA. The mixings of H_3 to the SM Higgs h and the heavy scalar H_1 are governed by two free parameters that represent the mixing between the $\text{Tr}(\Delta^\dagger \Delta)$ term with the $\text{Tr}(\Phi^\dagger \Phi)$ term in the scalar potential. However, as a result of the tree-level FCNC couplings of H_1 , it turns out that mixing of H_3 with h and H_1 are highly suppressed, $\lesssim 10^{-4}$ and $\lesssim 10^{-5}$, in the low mass range $M_{H_3} \leq 5 - 20$ GeV [6, 7].

The crucial difference from the $U(1)$ case is the presence of the right-handed W_R boson (along with additional charged Higgs bosons $H_1^\pm \simeq \phi_2^\pm$ and $H_2^{\pm\pm} \cong \Delta_R^{\pm\pm}$) at the TeV scale. Due to its suppressed coupling to quarks and leptons, the dominant decay mode of H_3 is to two photons from a W_R loop; this is analogous to the case of the W^\pm loop for the SM Higgs decay $h \rightarrow \gamma\gamma$ (with subleading contributions from the heavy charged scalars). When the H_3 is boosted and long-lived, this gives rise to two, collimated, displaced photons. As the mixing of H_3 to h and H_1 is tightly constrained by low-energy FCNC limits from B and K meson decays and oscillations, the long lifetime of H_3 in this case is guaranteed.

⁵³P.S. Bhupal Dev, Rabindra Mohapatra, Yongchao Zhang




Fig. 41: (U)LLP sensitivities at the $\sqrt{s} = 14$ TeV LHC (red) and MATHUSLA (blue) for the light scalar H_3 in the TeV LR model for three different values of $g_R/g_L = 0.6, 1$ and 1.5 . The grey contours show the decay lengths of H_3 in the laboratory frame with $g_R = g_L$; for $g_R \neq g_L$, the lifetime has to be rescaled by the factor of $(g_R/g_L)^{-2}$. To take into account the SM background, we assume an efficiency factor of $\epsilon_{\text{DV}} = 1$ (left) and 0.1 (right) for the LHC reaches. Below these curves we can have at least 4 signal events for both LHC and MATHUSLA. Note that MATHUSLA searches for the production of low-mass H_3 with $m_{H_3} \lesssim \text{GeV}$ from Z' decay may suffer from some backgrounds or lower reconstruction efficiency depending on the final detector design, see Section 3.1.4. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

In the LR model, the light scalar H_3 decays almost 100% into two photons via the W_R and charged scalar loops [6, 7]:

$$\Gamma(H_3 \rightarrow \gamma\gamma) = \frac{\alpha^2 m_{H_3}^3}{18\pi^3 v_R^2}. \quad (189)$$

with the factors in the parentheses from the loop functions for the vector bosons and scalars in the limit of $m_{H_3}/m_{W_R} \rightarrow 0$. In fact, the decay length of H_3 in the LR model is determined solely by the RH scale v_R , as well as the scalar mass m_{H_3} , when the scalar mixing angles are small. Therefore, in the presence of the extended gauge symmetry, the displaced (collimated) photon signal is rather unique. The effects at LHC and MATHUSLA might provide distinctive evidence of the parity-symmetric LR theories and neutrino masses beyond the SM via these clean displaced photon events. A cautionary note is that MATHUSLA may or may not be able to detect photons depending on the ultimate detector design (see Section 2.1 and Ref. [9]).

Turning to the production of H_3 in the minimal LR model, it can be produced from its coupling to the heavy RH gauge bosons W_R and Z_R , as well as through its coupling to the SM Higgs [6, 7, 481]. As the mixing of H_3 to the SM Higgs is severely constrained by the flavour data, we focus here only on the gauge portal production, which is dominated by the associated production of H_3 with a heavy W_R boson. The W_R subsequently decaying predominantly into the SM quarks ($J = u, d, s, c, b, t$):

$$pp \rightarrow W_R^* \rightarrow W_R H_3, \quad W_R \rightarrow JJ. \quad (190)$$

Here for simplicity we have assumed that the decay mode to on-shell heavy RHNs, *i.e.*, $W_R \rightarrow \ell N$, is kinematically forbidden. One should note that the $H_3 jj$ processes (with $j = u, d, s, c$) also receive (small) contributions from the heavy vector boson fusion (VBF) $pp \rightarrow W_R^* W_R^* jj \rightarrow H_3 jj$, which is however suppressed by the three-body phase space and the off-shell W_R propagator. At the LHC, limited by the total center-of-mass energy, the associated production with the Z_R boson is always highly suppressed, as it is heavier than the W_R boson in the minimal LR scenario. When $m_{H_3} \lesssim 10$ GeV, the

production rate is almost constant for a given v_R , and is sensitive only to the gauge coupling g_R (for the phenomenologically favored TeV range v_R , the production cross-section is at the fb level). For smaller $g_R < g_L$, the W_R boson is lighter and the production of H_3 can be significantly enhanced. As in the case of $U(1)$ model above, the associated jets from W_R decay tend to have a large p_T (typically $\gtrsim 1$ TeV), and can be easily used for triggering of the LLP events.

In Fig. 41, we give the range of H_3 masses and W_R masses that can be probed in the displaced diphoton channel at the LHC (red), as well as at MATHUSLA (blue) for three different values of $g_R/g_L = 0.6, 1$ and 1.5 , with $\sqrt{s} = 14$ TeV and an integrated luminosity of 3000 fb^{-1} , assuming at least 4 signal events. For the LHC case, we assume the scalar H_3 decays inside the tracker with a decay length of $1 \text{ cm} \lesssim \gamma c\tau_0 \lesssim 1 \text{ m}$. Note that with the TeV-scale W_R boson taking away most of the energy in the final state, reconstruction of the low-mass and highly boosted H_3 would be rather challenging. In addition, the displaced photons are more difficult than displaced charged particles, due to absence of tracking information. Therefore an efficiency factor of $\epsilon_{\text{DV}} = 0.1$ might be closer to be realistic in Fig. 41. The (U)LLP searches at LHC and MATHUSLA are largely complementary to each other, as in the case of $U(1)$ model, and could probe a W_R boson up to 6 TeV or so for $g_R = g_L$, which is complementary to the direct searches of W_R at the LHC in the same-sign dilepton plus jets events. As noted earlier, the virtually background-free environment at MATHUSLA might have superior sensitivity relative to the larger acceptance of the LHC for lighter H_3 masses.

7.4 Neutrinos from the Higgs Portal⁵⁴

In addition to supersymmetry, abelian extensions of the SM at the TeV scale represent an intriguing possibility among all the BSM scenarios. Indeed, the remarkable convergence of the gauge couplings, although only approximate, predicted by their renormalisation group evolution at around 10^{15} GeV, strongly hints in favour of GUTs. One of the main features of such theories is the appearance of an extra $U(1)'$ gauge group with the associated gauge boson Z' within reach of LHC energies; see, for instance, Ref. [498, 499] for a review.

The discovery of a new massive vector boson at the TeV scale would have further interesting implications. Indeed, the breaking of the extra abelian gauge symmetry would require the existence of an enlarged scalar sector with a heavy scalar field mixing with the SM Higgs doublet and giving mass to the Z' . Moreover, anomaly cancellation naturally predicts exotic fermionic states. These could be SM-singlet right-handed neutrinos, which give mass to the SM neutrinos through a seesaw mechanism. In the simplest realisation of a type-I seesaw scenario with one singlet fermion for each flavour generation (other seesaw realisations can be envisaged as well), heavy neutrinos typically have an extremely small coupling to the SM gauge bosons induced by their small mixing with the light, active neutrinos. Therefore, the decay width of a heavy-neutrino could be small and its lifetime particularly large over a substantial region of the available parameter space, making the $U(1)'$ extension one of the best BSM scenarios predicting LLPs. (See e.g. Section 7.2).

The heavy neutrino, N , has a decay length determined by its mass and by the neutrino mixing matrix, as outlined in Sec. 7.1. In the mass range $1 \text{ GeV} < m_N < 100 \text{ GeV}$, the N proper decay length spans from 10^9 m to few cm [479, 500, 501] for parameters motivated by the minimal type-I seesaw, reaching the BBN limit [4, 502] for m_N of few GeV; see Fig. 42 (a). In this mass range the heavy neutrinos decay through the processes $N \rightarrow l^\pm W^\mp$ and $N \rightarrow \nu_l Z^*$ with off-shell gauge bosons, leading to four available modes, $qql, l^+l^- \nu_l, qq\nu_l$ and $\nu_l \nu_l \nu_l$ with BRs approximately given by 50%, 24%, 18% and 8%, respectively. Long-lived heavy neutrinos are pair produced at the LHC via the s-channel exchange of the Z' , the 125 GeV Higgs h and its heavy partner H . The presence of these mechanisms is the main difference between $U(1)'$ extensions and the simple seesaw-extended SM. Indeed, in the minimal see-saw model, the production of N is suppressed by the square of the mixing between left-

⁵⁴Elena Accomando, Luigi Delle Rose, Stefano Moretti, Emmanuel Olaiya, Claire H. Shepherd-Themistocleous




Fig. 42: (a) Proper decay length of the heavy neutrino as a function of its mass m_N for three different values of the light neutrino mass. For illustrative purpose we have assumed a diagonal Dirac neutrino mass matrix. (b) $\sigma \times \text{BR}$ for the process $pp \rightarrow h \rightarrow \sum NN$ in the gluon channel at the 14 TeV LHC as a function of the heavy neutrino mass for three different values of the scalar mixing angle and $x = 4$ TeV. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

and right-handed neutrino components, whereas in $U(1)'$ models there is a large region of the parameter space where production can be much larger due to Z' and Higgs production of N . The decay modes of the heavy neutrinos remain the same in both scenarios.

The heavy scalar H is responsible for the dynamical generation of the N Majorana mass and, through the mixing α between H and the SM Higgs in the scalar sector, acts as a portal providing an exotic heavy neutrino coupling to the SM Higgs. As discussed in Sections 8.2, 8.5, and 8.4, such mixings are generic since it cannot be forbidden by any symmetry. The $B - L$ scenario discussed in Sec. 7.2 represents an explicit realisation of a heavy scalar portal motivated by abelian extensions of the SM. The mixing angle, α , scales all the interactions of the SM-like (heavy) Higgs to SM particles with $\cos \alpha$ ($\sin \alpha$). Interactions of scalars with other particles in the extended spectrum of the $U(1)'$ model, such as the Z' and RH neutrinos, are proportional to the complementary angle ($\cos \alpha$ for H , $\sin \alpha$ for h). In particular, this gives a $\cos^2 \alpha$ scaling of the cross-section of the standard h production mechanisms, and a $\sin^2 \alpha$ scaling of the partial width of the exotic h decay into heavy neutrinos.

Notice also that extensions of the SM scalar sector affect the running of the quartic scalar couplings and help in the stabilisation of the vacuum [503–509]. In particular, $\alpha \gtrsim 0.1$ may allow to achieve a stable and perturbative vacuum up to the GUT scale over a wide range of heavy Higgs masses while complying with LHC Higgs searches, see for instance [508, 509].

If $m_N < m_h/2$ and $\alpha \neq 0$, all three heavy neutrino production modes are kinematically accessible with the SM-like Higgs mediation being the dominant channel in a large region of the parameter space. The corresponding partial decay width is

$$\Gamma(h \rightarrow \sum NN) = \sum_i \frac{m_{N_i}^2}{x^2} \sin^2 \alpha \frac{m_h}{16\pi} \left(1 - \frac{4m_{N_i}^2}{m_h^2}\right)^{3/2}, \quad (191)$$

where i sums over the heavy neutrino families and x is the vacuum expectation value of the extra scalar. This expression is common to every extension of the SM in which the Majorana mass of the heavy neutrinos is generated by the vev x through spontaneous symmetry breaking of a SM-singlet scalar




Fig. 43: Sensitivity estimate for the MATHUSLA detector at the HL-LHC (assuming the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1) for the process $pp \rightarrow h \rightarrow \sum NN$. We have assumed a diagonal Dirac neutrino mass matrix and $m_{\nu_l} = 0.1$ eV to fix the RH neutrino lifetime. The cross-section required to see 4 events has been normalised to the detection efficiency $\epsilon_{\text{LLP}}^{\text{MATH}}$ for a displaced vertex. The 2% and 6.4% lines are representative of possible $\text{Br}(h \rightarrow \text{invis})$ exclusions achievable by the HL-LHC, see Sections 3.2.1 and 8.2.

mixed with the SM $SU(2)$ Higgs doublet. In a $U(1)'$ scenario in which the Z' mass $M_{Z'}$ arises by the same mechanism, the vev x is not a free parameter but is fixed by $x = M_{Z'}/(z_S g')$, where z_S and g' are, respectively, the $U(1)'$ charge of the scalar singlet and the $U(1)'$ gauge coupling, and, therefore, it is constrained by the Z' searches at the LHC in the di-lepton channel [483, 484] (see [500] for a reinterpretation of the limits in terms of a generic $U(1)'$ charge assignment).

Let us assume that the Z' mass arises from the same vev x . The $\text{BR}(h \rightarrow \sum NN)$ is constrained to be $\lesssim 1\%$, but it is compensated by a large Higgs production cross-section $\sigma_h = \cos^2 \alpha \sigma_{h_{\text{SM}}}$, Fig. 42(b), with $\sigma_{h_{\text{SM}}} = 54.67$ pb in the gluon-fusion channel at the 14 TeV LHC [510]. The shape of the N production cross-section in Fig. 42(b) is determined by the heavy neutrino mass. In particular its zeros are located at $m_N \simeq 0$ and $m_N = m_h/2$ where, respectively, the coupling of the Higgs to right-handed neutrinos is strongly suppressed or the process is kinematically closed. The dependence on α is mainly controlled by the scaling factors $\cos^2 \alpha$ and $\sin^2 \alpha$, affecting, respectively, the Higgs production cross-section and the partial decay width in Eq. (191). A mild residual dependence on α , which is only seen for large values of α , appears in the Higgs total decay width, $\Gamma_{\text{tot}} = \cos^2 \alpha \Gamma_{\text{tot}}^{\text{SM}} + \Gamma(h \rightarrow \sum NN)$, that normalises the $\text{BR}(h \rightarrow \sum NN)$. Moreover, if the vev x is not constrained by the Z' mass, the $\text{BR}(h \rightarrow \sum NN)$ is bounded from above only by the upper limit on the invisible Higgs decay [511, 512].

For $m_N > m_h/2$, the h channel is, instead, kinematically closed. For $M_{Z'} \gtrsim 3$ TeV, as required by the recent di-lepton searches at the LHC [483, 484], the dominant N production mode typically comes from decays of the heavy Higgs, H . However, if the heavy scalar sector is decoupled from the SM one, namely $\alpha \simeq 0$, both the $h \rightarrow NN$ decay and the heavy Higgs production from pp collisions are suppressed and the heavy neutrino pair production via the Z' remains the only available possibility over the entire range of the heavy neutrino masses. As an example, this provides, for $M_{Z'} = 4$ TeV and $m_N \ll M_{Z'}$, $\sigma_{Z'} \text{BR}(Z' \rightarrow \sum NN) \sim 0.6$ fb at the 14 TeV LHC. The production of N from lower-mass Z' with smaller gauge couplings g' was discussed in Sec. 7.2.1.

Here we consider long-lived heavy neutrino production from decay of the SM-like Higgs, and

we present in Fig. 43 an estimate of the sensitivity for the MATHUSLA detector at the HL-LHC. In particular, we show the $\text{BR}(h \rightarrow \sum NN)$ required to observe 4 signal events as a function of the heavy neutrino proper decay length $c\tau_0$. The BR has been normalised to the displaced vertex detection efficiency in MATHUSLA, $\epsilon_{\text{LLP}}^{\text{MATH}}$. The Standard Model Higgs production cross-section $\sigma_{h_{\text{SM}}}$ has been used. The $\cos^2 \alpha$ correction induced by the scalar mixing angle in the range $0 < \alpha \lesssim 0.3$ (larger values are disfavoured by the signal strength measurements of the Higgs) only reduces the Higgs cross-section by a factor less than 9% and does not considerably affect the result in Fig. 43. For the sake of simplicity we have also assumed a diagonal Dirac neutrino mass matrix and $m_{\nu_l} = 0.1$ eV. Notice that, for a given choice of the light neutrino mass m_{ν_l} , the mass of the heavy neutrino is completely determined by its decay length as shown in Fig. 42(a).

It is instructive to compare the MATHUSLA sensitivity to the Higgs portal N production with the capabilities of the ATLAS and CMS detectors at the HL-LHC. Since about 76% of the heavy neutrino decay channels contain a lepton in the final state and the heavy neutrinos are pair produced in the exotic Higgs decay, a search for two LLPs in the trackers of the main detectors may be characterised by low background contamination. Therefore, exploiting dilepton triggers may result in a better sensitivity than MATHUSLA in the short lifetime regime, $c\tau \lesssim 10$ m. On the other hand, the same search in the long lifetime regime would provide a very poor sensitivity. In that case, the most direct comparison to MATHUSLA would be a search for a single right-handed neutrino. In the most optimistic scenario one should focus on decays inside the ATLAS and CMS trackers and rely on the fully muonic decay mode, thus employing the dimuon trigger at Level 1. Assuming the presence of a muon and a DV would be sufficient to eliminate all backgrounds at the main detectors, we can follow the procedure outlined in Section 3.2.3 to compute the long-lifetime sensitivity gain R_s of MATHUSLA compared to the main detector. Taking DV reconstruction in the main detector tracker to have roughly a quarter the geometric acceptance and half the efficiency as MATHUSLA, along with the $\epsilon_{\text{cuts}}^{\text{LHC}} \sim 0.5$ requirement of a muon produced in RH neutrino decay, we arrive at $R_s \sim 15$. Therefore, even though this signature would be background free in the main detectors, the sensitivity to the cross-section (and hence long lifetime) achieved by MATHUSLA would be at least an order of magnitude better. Furthermore, it is worth noting that the hypothesis of negligible background for displaced vertices in the tracker may be too optimistic. The very low sterile neutrino masses in the long lifetime regime may make their reconstruction and the background rejection very challenging.

The study presented above relies on the existence of a portal mediated by a SM singlet scalar to heavy neutrinos for the SM Higgs. A natural realisation of this scenario is given by an extension of the SM with a spontaneously broken abelian symmetry where both states, the heavy neutrino and the extra scalar, are naturally required by the gauge symmetry and the generation of the Z' mass. Nevertheless, the $U(1)'$ symmetry is not mandatory and other scenarios with a global symmetry can still provide exotic Higgs decay to long-lived particles. A model-independent approach to such scenarios has been presented in [8] where the contribution of heavy new physics degrees of freedom (such as the heavy scalar and the Z' discussed above) has been parameterised in terms of a low-energy effective field theory (EFT) whose leading effects are encoded in dimension-5 operators. When the theory is augmented with SM singlet fermions with masses around or below the EW scale, new dimension-5 operators appear in the EFT [513–515] besides the well-known Weinberg operator [516]. One of them, $(\lambda_{ij}/\Lambda)\overline{\nu_{R_i}^c}\nu_{R_j}\Phi^\dagger\Phi$, contributes to the Majorana mass M_M of the right-handed neutrinos and provides additional couplings to the SM Higgs which are not necessarily proportional to M_M . In particular, for $m_{N_i} < m_h/2$, the Higgs can decay to heavy neutrinos via the coupling $v/(2\Lambda)h\overline{\nu_R^c}\lambda\nu_R$. After the identification $v/(2\Lambda)\lambda_{ij} \equiv M_{M_{ij}}/(2x)\sin\alpha$, one can recover the Higgs partial decay width in Eq. (191) and easily reach similar conclusions to those shown here.

7.5 Discrete Symmetries and FIMP Dark Matter⁵⁵

In the models discussed so far, we have focused on the type-I seesaw model for neutrino masses. An alternative explanation for the smallness of neutrino masses results from new particles charged under a discrete or continuous symmetry (global or local), such that the dimension-5 Weinberg operator responsible for SM neutrino masses does not arise at tree level but instead at the N -loop level. In this class of models, the active neutrino mass matrix can be expressed as:

$$\mathcal{M}_{ij}^\nu \sim \left(\frac{1}{16\pi^2} \right)^N \frac{\alpha_{ij}}{\Lambda} \langle \Phi^0 \rangle^2, \quad (192)$$

where Λ is the scale of the new physics and α_{ij} are effective couplings, accounting for the flavour structure of the neutrino mass matrix. Notably, the suppression by the loop factor of the radiatively generated neutrino masses lowers the mass scale of the new physics, Λ , thus opening the exciting possibility of producing in colliders the particles responsible for the neutrino mass generation. Furthermore, if the new symmetry is unbroken (or mildly broken) in the electroweak vacuum, the lightest particle of the sector responsible for neutrino masses is long-lived on cosmological time-scales and therefore constitutes a dark matter candidate. Collider experiments, in this case, may also produce dark matter particles, either directly or in cascade decays.

One of the simplest models of radiative neutrino mass generation incorporating a dark matter candidate is the so-called scotogenic model [517]. In this model, the particle content of the Standard Model (SM) is extended with one additional scalar doublet $H_2 \equiv (H^+, H_2^0)$ and at least two fermion singlets N_j ($j = 1, 2, \dots$). The model also postulates that the electroweak vacuum is invariant under a discrete Z_2 symmetry, under which all SM fields are even, whereas N_j and H_2 are odd. The Lagrangian of the model reads

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \mathcal{L}_{H_2} + \mathcal{L}_N + \mathcal{L}_{\text{int}}, \quad (193)$$

where \mathcal{L}_{SM} is the SM Lagrangian, \mathcal{L}_{H_2} and \mathcal{L}_N are, respectively, the terms in the Lagrangian involving only the fields H_2 and N_j ,

$$\mathcal{L}_{H_2} = (D_\mu H_2)^\dagger (D^\mu H_2) - \mu_2^2 (H_2^\dagger H_2) - \lambda_2 (H_2^\dagger H_2)^2, \quad (194)$$

$$\mathcal{L}_N = \frac{i}{2} \bar{N}_j \partial_\mu \gamma^\mu N_j - \frac{1}{2} M_j \bar{N}_j^c N_j + \text{h.c.}, \quad (195)$$

and \mathcal{L}_{int} contains the interaction terms of the Z_2 -odd fields with the Standard Model fields,

$$\begin{aligned} \mathcal{L}_{\text{int}} = & -\lambda_3 (H_1^\dagger H_1) (H_2^\dagger H_2) - \lambda_4 (H_1^\dagger H_2) (H_2^\dagger H_1) - \frac{\lambda_5}{2} \left[(H_1^\dagger H_2)^2 + \text{h.c.} \right] \\ & + \left[Y_{\alpha i}^\nu (\bar{\nu}_{\alpha L} H_2^0 - \bar{\ell}_{\alpha L} H^+) N_i + \text{h.c.} \right]. \end{aligned} \quad (196)$$

where H_1 is the SM Higgs doublet, ℓ are the charged leptons and ν are the active neutrinos. The parameters of the scalar potential are chosen such that $\langle H_1 \rangle = (0, v/\sqrt{2})$, with $v \simeq 246$ GeV, and $\langle H_2 \rangle = 0$, hence the minimum of the potential breaks the electroweak symmetry while preserving the Z_2 symmetry. A variant of the model has instead one Z_2 -odd fermion singlet, and at least two Z_2 -odd scalar doublets [518], which naturally lead to a mild hierarchy between the solar and atmospheric mass scales.

The conservation of the Z_2 symmetry ensures that the lightest Z_2 -odd particle is absolutely stable, which then constitutes a dark matter candidate if it is electrically neutral. The dark matter candidates of the model are the CP-even and CP-odd neutral scalars, H^0 and A^0 , and the lightest singlet fermion N_1 . Here we focus on the latter candidate, concretely in the region of the parameter space where it constitutes

⁵⁵Alejandro Ibarra

a Feebly Interacting Massive Particle (FIMP), see also Section 5.3. If this is the case, heavier Z_2 -odd particles can be very long-lived and decay inside MATHUSLA.

The signals of the scotogenic FIMP scenario crucially depend on the mass spectrum of the Z_2 -odd sector. Of particular interest for MATHUSLA is the scenario where $M_1 < M_2 < M_3 < m_H$, where m_H denotes the overall mass scale of the Z_2 -odd scalar sector. The Z_2 -odd scalars H^0 , A^0 and H^\pm , can be produced at the LHC via neutral and charged current Drell-Yan processes (as well as in gluon fusion with an off-shell Higgs in the s-channel [519]) and subsequently decay into singlet fermions, mostly N_2 and N_3 , due to the suppressed coupling of the FIMP to the Z_2 -odd scalars. As shown in Fig. 8, MATHUSLA could probe such EW production processes for LLPs for mass scales up to a TeV, depending on the lifetime. The singlet fermions N_2 and N_3 in turn decay producing visible particles in the final state with rates: [520]

$$\Gamma(N_j \rightarrow \ell_\alpha^- \ell_\beta^+ N_1) \simeq \frac{M_j^5}{6144 \pi^3 m_H^4} \left(|Y_{\beta 1}^\nu|^2 |Y_{\alpha j}^\nu|^2 + |Y_{\alpha 1}^\nu|^2 |Y_{\beta j}^\nu|^2 \right), \quad j = 2, 3, \quad (197)$$

$$\Gamma(N_3 \rightarrow \ell_\alpha^- \ell_\beta^+ N_2) \simeq \frac{M_3^5}{6144 \pi^3 m_H^4} \left(|Y_{\beta 2}^\nu|^2 |Y_{\alpha 3}^\nu|^2 + |Y_{\alpha 2}^\nu|^2 |Y_{\beta 3}^\nu|^2 \right). \quad (198)$$

Here, the masses of the final fermions have been assumed to be negligibly small compared to the mass of the decaying fermion.

For FIMP dark matter, the requirement of reproducing the observed dark matter abundance fixes the size of the Yukawa coupling as a function of the FIMP mass and the charged scalar mass, thus giving proper decay-lengths for N_2 and N_3 , given by [521]

$$c\tau(N_2) \approx 2 \times 10^{13} \text{ m} \left(\frac{M_1}{10 \text{ keV}} \right) \left(\frac{m_H}{500 \text{ GeV}} \right)^3 \left(\frac{100 \text{ GeV}}{M_2} \right)^5 \left(\frac{10^{-3}}{y_2} \right)^2, \quad (199)$$

$$c\tau(N_3) \approx 0.4 \text{ m} \left(\frac{100 \text{ GeV}}{M_3} \right) \left(\frac{m_H}{M_3} \right)^4 \left(\frac{10^{-3}}{y_2} \right)^2 \left(\frac{10^{-3}}{y_3} \right)^2, \quad (200)$$

where $y_k \equiv (\sum_\alpha |Y_{\alpha k}|^2)^{1/2}$. N_2 is then stable even at distance scales of the Solar System. However, N_3 can be stable within the ATLAS detector and decay some distance away, possibly inside MATHUSLA, producing two charged leptons (in general with different flavour) and missing energy, carried away by N_2 . The lifetime can easily be in the 100m or above range. MATHUSLA could then potentially supply the best sensitivity for production of such LLPs, in particular in the low-mass regime $m_{N_3} \lesssim 10 \text{ GeV}$ where searches for displaced lepton-jets at the main detectors suffer from some backgrounds, see discussion in Section 3.2.3. Furthermore, MATHUSLA may offer the possibility of discriminating the three body final state $\ell_\alpha^- \ell_\beta^+ + \cancel{E}_T$ from the two body final state $\ell_\alpha^- \ell_\beta^+$, from the angular distribution of the charged leptons inside the detector [522].

7.6 Enhanced Residual Symmetry (ERS) Scenarios and Freeze-Out Leptogenesis⁵⁶

The models considered in the previous sections treat neutrino masses and mixings as input parameters. In a complete theory, however, it is likely that symmetries play a crucial role in setting the observed pattern of masses, mixings and CP phases.

The scenario explored in this section belongs to a class of models which postulate a particular flavour symmetry breaking pattern to derive the low-energy parameters of the type I see-saw. It implements a type-I seesaw scenario with flavour symmetry G_f and a CP symmetry [437] that strongly constrain lepton mixing angles, and both low- and high-energy CP phases [523]. The three right-handed (RH) neutrinos N_i have (almost) degenerate masses. Their decays are responsible for the generation of the baryon asymmetry η_B of the Universe via resonant leptogenesis [384, 524].

⁵⁶P. S. Bhupal Dev, Claudia Hagedorn and Emiliano Molinaro

This not only explains the measured values of the lepton mixing angles, but also makes predictions for leptonic CP violation in neutrino oscillations, neutrinoless double beta decay, as well as connect low energy CP phases with those relevant for leptogenesis.

The dynamical mechanism of flavor symmetry breaking is not specified, but models which implement such mechanisms [438] often feature flavour symmetry breaking fields with aligned vacuum expectation values, resulting in a vacuum which respects an *enhanced residual symmetry* (ERS) compared to the naive expectation. This ERS is slightly broken by higher-dimensional operators, leading to observables which depend on various powers of a small breaking parameter.

Points of ERS are motivated for phenomenological reasons, offering an explanation for the smallness of the reactor mixing angle θ_{13} [439] and enhancing the yield of leptogenesis. It also leads to *some* of the RH neutrinos to have lifetimes much longer than the naive expectation of freeze-out leptogenesis.

In this scenario, it leads to a very long-lived RH neutrino N_3 that could be detected at MATH-USLA [1], while $N_{1,2}$ can be searched for via either prompt or displaced vertex signals at the LHC [421, 525].

Framework. A key feature of this scenario is the presence of a flavour G_f and a CP symmetry. Both symmetries are broken non-trivially to residual symmetries G_ν and G_e in the neutrino and charged lepton sector, respectively. We choose in the following $G_f = \Delta(3n^2)$ [526] or $G_f = \Delta(6n^2)$ [527] (n even, $3 \nmid n$, $4 \nmid n$). These groups are all subgroups of $SU(3)$ and allow the three generations of leptons to be unified in one representation $\mathbf{3}$ for $n \geq 2$, a hypothesis which is supported by the fact that two of the three lepton mixing angles are large. CP is an additional symmetry of this scenario as part of G_ν , since in this way low- and high-energy CP phases, of Dirac as well as Majorana type, can be predicted.

Left-handed (LH) lepton doublets l_α ($\alpha = e, \mu, \tau$) transform in an irreducible faithful representation $\mathbf{3}$, N_i are in an irreducible real representation $\mathbf{3}'$, whereas RH charged leptons α_R are assigned to $\mathbf{1}$ of G_f . The latter are distinguished by an auxiliary symmetry $Z_3^{(\text{aux})}$, under which l_α and N_i are invariant. The CP symmetry is given by the CP transformation $X(s)(\mathbf{r})$ in the representation \mathbf{r} and depends on the integer parameter s , $0 \leq s \leq n-1$, (see Case 1 in Ref. [528]). The forms of the neutrino Dirac mass matrix m_D and the charged lepton mass matrix m_l are determined by the residual symmetries $G_\nu = Z_2 \times CP$ and $G_e = Z_3^{(\text{D})}$ (the diagonal subgroup of Z_3 in G_f and $Z_3^{(\text{aux})}$), respectively. The Majorana mass matrix M_M of RH neutrinos is invariant under G_f and CP.

The matrix m_D can be written as [436]

$$m_D = v \Omega(s)(\mathbf{3}) R_{13}(\vartheta_L) \begin{pmatrix} y_1 & 0 & 0 \\ 0 & y_2 & 0 \\ 0 & 0 & y_3 \end{pmatrix} R_{13}(-\vartheta_R) \Omega(s)(\mathbf{3}')^\dagger, \quad (201)$$

where the unitary matrices $\Omega(s)(\mathbf{r})$ are determined by the CP transformation $X(s)(\mathbf{r})$ and $R_{13}(\vartheta)$ is a rotation in the (13)-plane through the angle ϑ . Note that here we use the convention that $v \approx 174$ GeV is the vacuum expectation value (VEV) of the Standard Model (SM) Higgs. There are five real parameters: y_i , ϑ_L and ϑ_R , in m_D . The charged lepton mass matrix m_l is diagonal with three undetermined entries corresponding to the charged lepton masses. The Majorana mass matrix M_M is of the form

$$M_M = M \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad (202)$$

and thus features three RH neutrinos with degenerate masses. We consider the two cases of strong normal (NO) and inverted ordering (IO): (a) strong NO arises for $y_1 = 0$ so that m_1 vanishes, $m_2 = y_2^2 v^2/M$ and $m_3 = y_3^2 |\cos 2\vartheta_R| v^2/M$, while (b) strong IO arises for $y_3 = 0$ so that $m_3 = 0$, $m_1 = y_1^2 |\cos 2\vartheta_R| v^2/M$ and $m_2 = y_2^2 v^2/M$. The two non-vanishing Yukawa couplings are fitted to the solar and the atmospheric mass squared differences Δm_{sol}^2 and Δm_{atm}^2 whose best fit values are taken from Ref. [443].

For particular values of ϑ_L and ϑ_R , the residual symmetry $G_\nu = Z_2 \times CP$ can be enhanced. If $\vartheta_L = 0, \pi$, the combination $m_D m_D^\dagger$ becomes invariant under a further Z_2 subgroup of G_f . Similarly, for the choices $\vartheta_R = 0, \pi/2, \pi, 3\pi/2$ the combination $m_D^\dagger m_D$ preserves a symmetry larger than G_ν . This symmetry is also larger than the one of $m_D m_D^\dagger$ for $\vartheta_L = 0, \pi$, since RH neutrinos transform as the real representation $\mathbf{3}'$ of G_f that is unfaithful for $n > 2$.

These points of ERS are of particular relevance for phenomenology, since ϑ_L deviating from $\vartheta_{L,0} = 0$ or π leads to a non-zero value of the reactor mixing angle θ_{13} . ϑ_R close to $\vartheta_{R,0} = 0, \pi/2, \pi$ or $3\pi/2$ makes it possible for the RH neutrino N_3 to be long-lived enough for being detected with the MATHUSLA detector (see Eq. 211 and Fig. 44), while simultaneously maximizing the CP asymmetries $\epsilon_{i\alpha}$ relevant for leptogenesis (see Eqs. 208 and 209). One can argue that the larger the ERS is, the smaller the deviation from points of ERS will be, i.e. ϑ_R is expected to deviate from $\vartheta_{R,0}$ by

$$\delta\vartheta_R = |\vartheta_R - \vartheta_{R,0}| \lesssim 0.01 \quad , \quad (203)$$

while ϑ_L can deviate from $\vartheta_{L,0}$ up to

$$\delta\vartheta_L = |\vartheta_L - \vartheta_{L,0}| \sim 0.2 \quad . \quad (204)$$

These splittings will thus determine the resulting phenomenology.

In one type of explicit model [438], the flavour and CP symmetry are spontaneously broken to the residual symmetries G_ν and G_e with the help of flavour symmetry breaking fields and a peculiar alignment of their VEVs, achieved with a potential with a particular form. Depending on the fields and the form of the potential, an ERS larger than G_ν and G_e can be preserved at leading order. Higher-dimensional operators then induce small deviations from these points of ERS, thus explaining the particular sizes of ϑ_L and ϑ_R . An example can be found in Ref. [439], where the correct size of ϑ_L and thus the reactor mixing angle θ_{13} is generated in this way.

Higher-dimensional operators connecting different sectors of the theory are responsible for the eventual breaking of the residual symmetries G_ν and G_e and thus affect the given form of m_D , m_l and M_M . In particular, they are also the source of corrections leading to a small splitting in the RH neutrino masses. This splitting is crucial for resonant leptogenesis. In the following, we focus on contributions to M_M that possess the residual symmetry G_e . These are proportional to κ , a positive power of the symmetry breaking parameter, measured in units of M . A small splitting of the RH neutrino masses therefore arises:

$$M_1 = M(1 + 2\kappa) \quad \text{and} \quad M_2 = M_3 = M(1 - \kappa) \quad . \quad (205)$$

Lepton Mixing and Low-Energy CP Phases. In the limit of residual symmetries G_ν and G_e , we obtain that the lepton mixing angles can be accommodated well for $\vartheta_L \approx 0.18$ (2.96) [528], i.e. $\sin^2 \theta_{13} \approx 0.0219$, $\sin^2 \theta_{12} \approx 0.341$ and $\sin^2 \theta_{23} \approx 0.605$ (0.395). Regarding the two physical CP phases in the cases of strong NO and IO, we find that the Dirac phase δ is trivial, whereas the Majorana phase α_2 depends on the chosen CP transformation $X(s)$

$$\sin \alpha_2 = (-1)^{k+r+s} \sin 6\phi_s \quad \text{and} \quad \cos \alpha_2 = (-1)^{k+r+s+1} \cos 6\phi_s \quad , \quad \text{with} \quad \phi_s = \frac{\pi s}{n} \quad , \quad (206)$$

where $k = 0$ ($k = 1$) for $\cos 2\vartheta_R > 0$ ($\cos 2\vartheta_R < 0$) and $r = 0$ ($r = 1$) for strong NO (IO). The value of the effective Majorana neutrino mass $m_{\beta\beta}$, accessible in neutrinoless double beta decay experiments, crucially depends on the choice of the CP symmetry and is in this scenario considerably restricted [436]. For $n = 26$, $\vartheta_L \approx 0.18$ and strong NO with $k = 1$, we get $0.0019 \text{ eV} \lesssim m_{\beta\beta} \lesssim 0.0040 \text{ eV}$, while for strong IO with $k = 0$, we find $0.016 \text{ eV} \lesssim m_{\beta\beta} \lesssim 0.048 \text{ eV}$, using the best fit values for Δm_{sol}^2 and Δm_{atm}^2 [443]. For strong IO, most of the admitted values of $m_{\beta\beta}$ can be tested with the proposed experiment LEGEND [529] and all of them can be explored with nEXO [530], whereas it is challenging to test the values of $m_{\beta\beta}$ predicted for strong NO with current and future experiments.

High-Energy CP Phases and the Leptogenesis Connection. Including the small mass splitting of the RH neutrinos, their out-of-equilibrium decays can generate the baryon asymmetry, η_B , via resonant leptogenesis [384, 524]. The CP asymmetries $\epsilon_{i\alpha}$ due to the decay of N_i and in the lepton flavour α read

$$\epsilon_{i\alpha} \approx \frac{1}{v^4} \sum_{j \neq i} \text{Im} (\hat{m}_{D,\alpha i}^* \hat{m}_{D,\alpha j}) \text{Re} \left(\left(\hat{m}_D^\dagger \hat{m}_D \right)_{ij} \right) \mathcal{F}_{ij}, \quad (207)$$

with \hat{m}_D being m_D in the RH neutrino mass basis and \mathcal{F}_{ij} related to the regulator that is proportional to the mass splitting of N_i [531].

We find the real part of $(\hat{m}_D^\dagger \hat{m}_D)_{ij}$ to be zero, if either $i = 3$ or $j = 3$. Hence, $\epsilon_{3\alpha} = 0$ for all α and $\epsilon_{i\alpha}$ only has one contribution for $i = 1, 2$. The imaginary part of $\hat{m}_{D,\alpha 1}^* \hat{m}_{D,\alpha 2}$ is proportional to $\sin 3\phi_s$ for even s and to $-\cos 3\phi_s$ for odd s , independent of the flavour α . If α is summed over, ϵ_1 and ϵ_2 both vanish. For strong NO and even s , the CP asymmetries $\epsilon_{1\alpha}$ read

$$\epsilon_{1\alpha} \approx \frac{y_2 y_3}{9} (-2y_2^2 + y_3^2 (1 - \cos 2\vartheta_R)) \sin 3\phi_s \sin \vartheta_R \sin \vartheta_{L,\alpha} \mathcal{F}_{12}, \quad (208)$$

and for strong IO, we find

$$\epsilon_{1\alpha} \approx \frac{y_1 y_2}{9} (-2y_2^2 + y_1^2 (1 + \cos 2\vartheta_R)) \sin 3\phi_s \cos \vartheta_R \cos \vartheta_{L,\alpha} \mathcal{F}_{12}, \quad (209)$$

with $\vartheta_{L,\alpha} = \vartheta_L + \rho_\alpha 4\pi/3$ and $\rho_e = 0, \rho_\mu = 1, \rho_\tau = -1$. For strong NO (IO) $\epsilon_{i\alpha}$ becomes very small, if $\vartheta_R \approx 0, \pi$ ($\vartheta_R \approx \pi/2, 3\pi/2$). In addition, \mathcal{F}_{ij} vanishes for $\cos 2\vartheta_R = 0$. The CP asymmetries $\epsilon_{2\alpha}$ are the negative of $\epsilon_{1\alpha}$ with \mathcal{F}_{12} being replaced by \mathcal{F}_{21} . We note that different values of s can lead to the same value of $\epsilon_{i\alpha}$. In particular, we find

$$\epsilon_{i\alpha}(s) = (-1)^s \epsilon_{i\alpha}(n-s) = \epsilon_{i\alpha}(n/2-s) = (-1)^{s+1} \epsilon_{i\alpha}(n/2+s) \text{ for } s \leq n/2. \quad (210)$$

Eqs. 206, 208 and 209 show the close correlation between CP violation at low and high energies.

Decay Lengths of RH Neutrinos. The decay widths $\Gamma_i \approx M_i (\hat{m}_D^\dagger \hat{m}_D)_{ii} / (8\pi v^2)$ of the RH neutrinos N_i are

$$\begin{aligned} \Gamma_1 &\approx \frac{M}{24\pi} (2y_1^2 \cos^2 \vartheta_R + y_2^2 + 2y_3^2 \sin^2 \vartheta_R), \quad \Gamma_2 \approx \frac{M}{24\pi} (y_1^2 \cos^2 \vartheta_R + 2y_2^2 + y_3^2 \sin^2 \vartheta_R), \\ \Gamma_3 &\approx \frac{M}{8\pi} (y_1^2 \sin^2 \vartheta_R + y_3^2 \cos^2 \vartheta_R). \end{aligned} \quad (211)$$

For M in the few hundred GeV range, we expect $y_i \sim 10^{-7}$ and thus mostly non-prompt decays at the LHC. If $\vartheta_R \approx \pi/2, 3\pi/2$ (for strong NO) or $\vartheta_R \approx 0, \pi$ (for strong IO), i.e. ϑ_R close to points of ERS, N_3 can have a very long lifetime, since Γ_3 tends to zero. Thus, N_3 could be searched for with the MATHUSLA detector, if it is produced at the LHC with sufficient cross-section. This is shown in Fig. 44 where we plot the decay length L of N_3 in the laboratory frame as a function of the deviation of ϑ_R from points of ERS for different values of the mass of N_3 . In doing so, we assume that N_i are produced in the decay of a parent particle with mass $m_{parent} = 4 \text{ TeV}$, corresponding to an average Lorentz boost factor of $\gamma = m_{parent}/(2M)$. For M in the few hundred GeV to TeV range and $10^{-4} \lesssim \delta\vartheta_R \lesssim 10^{-2}$, N_3 would decay, on average, within the MATHUSLA detector. If $10^{-3} \lesssim \delta\vartheta_R \lesssim 10^{-1}$, N_3 would decay on average within the LHC detectors, along with $N_{1,2}$ decays, where the latter giving rise to either prompt or displaced vertex signals at the LHC, depending on the choice of ϑ_R .

As these weak-scale right-handed neutrinos are not produced efficiently through the minimal interactions required for generating neutrino masses, to observe N_i decays at colliders, an efficient production mechanism is required. As in previous sections, there are a number of possibilities for additional UV states to be produced at colliders, which subsequently decay to final states including N_i 's. To mention just one example, one might consider theories related to the $B - L$ extensions discussed in the previous




Fig. 44: Decay length L of N_3 (in meters) in the laboratory frame as a function of $\delta\vartheta_R$ (defined in Eq. 203) for different values of M assuming production at the LHC in the decay of a 4 TeV parent particle. The shaded regions correspond to MATHUSLA or main detector average displacements (but not necessarily to the regime of best sensitivity).

sections, where the RH neutrinos could be produced in the decay of a TeV-scale Z' . This could produce these weak-scale LLPs in sufficient numbers for detection. In some models this may allow the region of parameter space relevant for leptogenesis to be probed.

Summary. We have presented a type-I seesaw scenario with a flavour and CP symmetry as well as three RH neutrinos with almost degenerate masses in the few hundred GeV to TeV range. This class of models can be connected to the observed baryon asymmetry through resonant leptogenesis. One of the RH neutrinos can be long-lived enough to be discovered at the MATHUSLA detector if the production cross-section is sufficient. The other two can be searched for at the LHC main detectors. This would allow both detectors working in conjunction to diagnose the mechanism of leptogenesis generating the baryon asymmetry of our universe.

8 Theory Motivation for LLPs: Bottom-Up Considerations

In this section we study a variety of well-defined, generic possibilities for physics beyond the Standard Model, and demonstrate that in many circumstances they naturally yield interesting LLP signals for MATHUSLA. The scenarios we study here represent plausible and consistent possibilities for physics beyond the SM, which can and should be studied independently of particular ‘top-down’ theoretical motivations for specific forms of new physics. We emphasize that this section is highly complementary to the previous sections. Two of the main topics covered here are hidden valleys (Section 8.1) and exotic Higgs decays (Section 8.2 and 8.3), and in both cases specific realizations of these general scenarios have appeared in multiple contexts in the previous sections. We also discuss MATHUSLA’s sensitivity to the minimal extensions of the SM with (1) a Higgs portal-mixed singlet scalar (‘SM+S’, Section 8.4); (2) a massive Abelian dark vector boson kinetically mixing with hypercharge (‘SM+V’, Section 8.5); or (3) an axion-like particle (ALP, Section 8.6). All of the above simple models are well-motivated from effective field theory considerations, and are, for example, realized in theories of Dark Matter (Section 5.4) or Neutral Naturalness (Section 4.2). The results of this section thus apply to a broad range of well-motivated theories of physics beyond the SM, and illuminate the essential features that make MATHUSLA a uniquely sensitive instrument for discovering SM-singlet LLPs in all of them.

8.1 Hidden Valleys and High Multiplicity Scenarios ⁵⁷

8.1.1 Motivation

In this section, we consider hidden valley (HV) models [21, 22] that give rise to high multiplicity final states. HVs are a class of models where there are relatively light states whose only coupling to Standard Model (SM) states is via a heavy mediator. This setup naturally allows for long lifetimes, while maintaining a sizable production cross-section at the LHC. The hidden valley framework is very general, and appears in many well-motivated scenarios. It also arises naturally from the structure of gauge theories, which makes disconnected sectors a straightforward possibility. For this reason, HV models are discussed throughout this report, including Section 4.2 on neutral naturalness and Section 5.4 on SIMPs/ELDERs.

In this section, we consider HVs with a confining gauge group so that showering and hadronization will lead to large particle multiplicity when the HV is accessed. Our only theory prior will be to assume a hadronization scale in the \sim GeV range, as well as a much heavier mediator which can be accessed at the LHC. We will show how the phenomenology depends on the strength of the hidden sector coupling by considering two limiting cases.

8.1.2 Hidden sector fragmentation and hadronization

The phenomenology of confining hidden valleys most notably depends on the strength of the hidden sector gauge coupling. In particular, if the ‘t Hooft coupling is relatively small, the fragmentation process is dominated by soft and collinear branchings, as these are typically enhanced by large logarithms. The result is a fairly collimated spray of particles with sizable hierarchies in their momentum distribution, where the average particle multiplicity scales as a powerlaw of the ratio of the hard scattering scale Q and the hadronizations scale Λ . Specifically, one can show [532]

$$\langle n(Q) \rangle \propto \left(\frac{Q}{\Lambda} \right)^{2\gamma_T(j)} \Big|_{j=1}. \quad (212)$$

where $\gamma_T(j) < 1/2$ the timelike (fragmentation) anomalous dimension. After hadronization, the dark hadrons decay to SM fields, but because the mediator is heavier, the decay length is often macroscopic

⁵⁷Simon Knapen, Dean Robinson, Daniel Stolarski

and can be quite long. Therefore a jet of dark hadrons becomes an ordinary jet at long distances. The phenomenology of these ‘emerging jets’ was studied [533] using benchmarks motivated by [534].

On the other hand, in the regime where ‘t Hooft coupling is large, the fragmentation process is very efficient regardless of the phase space configuration of the branchings. On average one therefore expects isotropic branchings with more or less equal energy sharing between the partons. Moreover, if the ‘t Hooft coupling remains large over a large energy window without triggering a mass gap, as can be the case in walking or quasi-conformal hidden sectors, a high multiplicity of hidden sector hadrons are produced, with momenta of the order of the hadronization scale. In this case the average multiplicity scales as

$$\langle n(Q) \rangle \propto \frac{Q}{\Lambda}. \quad (213)$$

Rather than a jet-like structure, the result of such efficient showering is an approximately spherically symmetric event, with an approximately democratic energy distribution [535–538]. This topology is referred to as SUEP (Soft Unclustered Energy Patterns). The energy distributions of the final states can be regarded as almost thermal, parametrized by an effective Hagedorn temperature [539], $T \sim \Lambda$, a feature which is also borne out by AdS/CFT calculations [540]. The case for which the hidden hadrons decay promptly poses significant trigger challenges at ATLAS and CMS, as was studied in [541].

The low-lying spectrum of hadrons can be very rich within an HV sector. Depending on the number of light and heavy HV quark flavors, as well as their HV charges, the spectrum may typically involve the analogues of glueballs, onium states, baryons, η' and lighter pions. In order to capture the leading features of a plausible hadronization model, we assume the low lying spectrum contains only a single flavor of a long-lived (pseudo)scalar ϕ , with $m_\phi \sim \Lambda$. MATHUSLA’s ability to reconstruct DVs for LLPs with masses below a GeV is likely to depend on the production mode and details of the detector design, see Section 3.1.4. Anticipating that the low lying HV hadron ϕ will decay either to two or four SM states, we therefore choose $\Lambda \simeq m_\phi = 1$ GeV for the hadronization scale and scalar hadron mass, in order to maximize the multiplicity of detectable final state particles. Assuming that detector efficiencies turn-on sharply at this detection threshold, this maximized multiplicity benchmark corresponds to the best-case scenario for detection.

The ϕ proper lifetime should fall between ~ 10 and $\sim 10^7$ meters in order to have the majority of ϕ ’s decay outside of ATLAS and CMS, and avoid potential BBN constraints, respectively. Since $\Lambda \sim m_\phi = 1$ GeV, possible decay products include SM leptons, pions or photons. In the rest of this work we take the lifetime to be a free parameter.

For the MATHUSLA detector, the phenomenology between both limiting cases differs in the following ways:

- For a similar confinement scale, the particle multiplicity for SUEP-like dynamics is higher than for jet-like dynamics. In both cases, the multiplicity is high enough such that two or more vertices can simultaneously occur in the detector, despite the relatively low geometric acceptance.
- For SUEP-like events, the multiplicity scales linearly with hard scattering scale, but the momentum distribution of the final state particles is determined only by the SUEP temperature. For jet-like events, increasing the hard scattering scale increases the boost of the particles, while the multiplicity of the events scales as a sublinear power law, as encoded in the anomalous dimension Eq. (212).

In what follows, we discuss both cases separately.

8.1.3 Emerging Jets

Inspired by [533, 534] we consider a QCD-like dark sector with $SU(N_d)$ confining gauge group and N_f flavours. It has a confinement scale near the QCD scale. The mediator to the dark sector is a heavy scalar




Fig. 45: Feynman diagram for one of the production processes for the emerging jet scenario at the LHC. The X scalar is pair-produced, and each X decays to an ordinary quark and a dark quark (Q , shown by the thick red line). This leads to events with two QCD jets and two emerging jets.

which is a bifundamental under both QCD and dark-QCD. The mediator, X_d , has a Yukawa coupling to quarks (q) and dark quarks (Q):

$$\kappa X_d \bar{Q} q + \text{h.c.} \quad (214)$$

Here we have suppressed flavour indices in both the SM and dark sectors, but we assume that the coupling of the X_d is dominantly to light flavours on the SM side. The production process for the mediator is shown in Fig. 45. The mediator is pair produced via its QCD interaction, and the production cross-section is that of a scalar top from supersymmetry [103] times N_d . In this work we take $N_d = 3$, and we take two benchmark masses for X_d , $M = 1000, 1500$ GeV.

Each mediator decays to one quark and one dark quark via the Yukawa coupling of Eq. (214). The quark and dark quark then shower and hadronize, so each event contains two ordinary jets from the quarks and two emerging jets from the dark quarks. As mentioned above, we take the simplifying assumption that the emerging jets contain only a single species of dark hadron whose mass is 1 GeV. The dark pions then decay to Standard Model quarks via a virtual heavy mediator. The fact that the mediator is much heavier than the dark hadrons gives the dark pions a naturally long lifetime. One can calculate the width of the dark hadron assuming it is a pion-like state using dark chiral perturbation theory⁵⁸ [534]:

$$\Gamma(\phi \rightarrow \bar{q}q) = \frac{\kappa^4 N_c f_\phi^2 m_q^2 m_\phi}{32\pi M_{X_d}^4} \quad (215)$$

where f_ϕ is the dark hadron decay constant and m_q is the mass of the Standard Model quark that participates in the decay. From this equation, we can get the lifetime

$$c\tau_0 = \frac{c\hbar}{\Gamma} \approx 300 \text{ m} \times \left(\frac{0.3}{\kappa}\right)^4 \times \left(\frac{1 \text{ GeV}}{f_\phi}\right)^2 \times \left(\frac{5 \text{ MeV}}{m_q}\right)^2 \times \left(\frac{1 \text{ GeV}}{m_\phi}\right) \times \left(\frac{M_{X_d}}{1 \text{ TeV}}\right)^4, \quad (216)$$

and we see that for the parameters chosen here, the dark pions will be quite long-lived for ATLAS and CMS, but in right in the ballpark for MATHUSLA. The lifetime is very sensitive to the Yukawa coupling κ , which is unknown and can vary widely, so we see that very short and very long lifetimes are possible.

We simulate emerging jets events at the 13 TeV LHC using the hidden valley implementation [542, 543] of Pythia 8 [544] with the implementation of gauge coupling running from [533]. The typical number of dark hadrons in an event depends strongly on the dark hadron mass, but only weakly on the mediator mass (see Eq. (212)). For the benchmark used here of a 1 GeV dark hadron and a 1 TeV mediator, the typical number of dark hadrons is ~ 100 , while for a 1.5 TeV mediator it is ~ 120 .

From the simulation we can estimate the fraction of emerging jets events which will have one or two dark pions decaying in MATHUSLA. This is shown in the left panel of Fig. 46, with the solid lines

⁵⁸This computation assumes that the SM final states can be treated as free quarks, so for dark hadrons with masses not too far above the QCD scale, there will be large corrections from strong QCD effects, but these will not change the qualitative picture.




Fig. 46: Sensitivity plots for the emerging jets scenario with a mediator mass of 1 TeV (red) and 1.5 TeV (blue). Left: Number of expected events in MATHUSLA with at least one (two) displaced vertices are shown in the solid dark (dashed light) lines for our two benchmark models. We also show the approximate number of expected decays in ATLAS or CMS tracker in the dashed lines. Right: MATHUSLA exclusion potential requiring four events with at least one displaced vertex (solid), compared to the projected jets plus missing energy limit in the long lifetime regime for the high luminosity LHC from [545] (dashed). This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

corresponding to one dark pion decay, and the dot-dashed line being two dark pion decays. We have used the squark production cross-section with $N_d = 3$ and $3,000 \text{ fb}^{-1}$ integrated luminosity. All the results are assuming perfect efficiency for detection of dark hadron decays MATHUSLA, $\epsilon_{DV} = 1$. This of course is an unrealistic assumption, but all results can simply be scaled by the actual value of ϵ_{DV} .

We see that the optimum lifetime for MATHUSLA is ~ 100 m, the distance it is from the interaction point, but that there are a few events in MATHUSLA for lifetimes as long as 10^6 m. We can turn this around and ask what the limit that MATHUSLA can place on this model assuming it does not see any events. This is shown on the right panel of Fig. 46, and computed by asking what cross-section would give 4 events in the detector. We see that MATHUSLA can be sensitive to cross-sections as small as 0.01 fb with the full run of the LHC. This is significantly below the benchmark minimum signal cross-section Eqn. 10 due to the very high dark hadron multiplicity. We also note that in this plane our two benchmarks are quite similar, meaning that their difference in number of observed events is due almost entirely to the different cross-section. The two benchmarks are similar because the number of hadrons in an emerging jet scales as a very small power of the mediator mass as seen in Eq. (212), so the difference of 50% in the mediator mass is barely visible on a log-log plot. This is to be contrasted with the SUEP scenario discussed in the next section where there is a much stronger dependence on the mediator mass.

This scenario of physics beyond the SM could also be discovered or bounded with more traditional search strategies at ATLAS and CMS. Because each event has two hard QCD jets, for sufficiently heavy mediators ($M_X \gtrsim 500 \text{ GeV}$), these events become trivial trigger on. There are then two obvious ways to look for the models. The first is to search for displaced decays of the dark hadrons in ATLAS or CMS. For the light hadron masses considered here, decays in the calorimeter or muon system become difficult to resolve (see Section 3.2.3 for more details). Therefore, we estimate the reach by requiring one event in tracker of ATLAS or CMS, and we very crudely model the tracker as a solid sphere of radius 1 meter. This estimate is shown as the dashed lines in the left panel of Fig. 46. We see that at short lifetimes, ATLAS and CMS see many more events, but for lifetimes longer than ~ 100 m, the MATHUSLA sensitivity becomes slightly better than that of current detectors. This ATLAS/CMS estimate assumes 100% signal efficiency, but it is not clear how good the efficiency would actually be in the high pile-up environment of the LHC. In the short lifetime regime, many events will have multiple decays within the tracker, which can make the event appear more spectacular, but which could also degrade the tracking and detection

efficiency. The estimate in Fig. 46 also assumes zero background, but this is also probably not realistic, as there will be secondary interactions of hadrons in material, as well as decays of ordinary hadrons. We leave all these questions to future study, but here we note that the estimate for the ATLAS/CMS reach for searching for decays is a *best case* estimate, and likely to be worse than what is given here.

For lifetimes $\gtrsim 10$ m, the majority of dark hadrons escape the detector and the emerging jets simply appear as missing energy in ATLAS and CMS. Therefore, the event topology for this model becomes jets and missing energy, an extremely well studied signature. The limits on this topology for the high luminosity LHC were studied in [545]⁵⁹, and we show them as the dashed lines in the right panel of Fig. 46. We see that for intermediate lifetimes MATHUSLA puts a significantly stronger limit than the jets plus missing energy search. Therefore, we see that MATHUSLA can extend the reach for this scenario significantly relative to jets plus missing energy, and is competitive or possibly superior to the search for displaced vertices within ATLAS and CMS depending on signal efficiencies.

8.1.4 Soft Unclustered Energy Patterns (SUEP)

Production of the SUEP event morphology can be encoded generally in an operator product of the form $M^{4-\Delta_{\text{vis}}-\Delta_{\text{HV}}}\mathcal{O}_{\text{vis}}\mathcal{O}_{\text{HV}}$, in which \mathcal{O}_{vis} is an SM neutral operator consisting of either SM degrees of freedom or heavy exotic mediators. While SUEP production need not be associated with the production of an intermediate resonant state between the visible and HV sectors, it is simpler and representative to focus on two cases, following the discussion in [541]:

- $\mathcal{O}_{\text{vis}} = S$, a heavy (pseudo)scalar field, that is singly produced by e.g. gluon fusion production, $SG_{\mu\nu}G^{\mu\nu}$ or $SG_{\mu\nu}\tilde{G}^{\mu\nu}$. For this scenario we consider two benchmarks, with scalar mass $M = 750$ GeV and $M = 400$ GeV.
- $\mathcal{O}_{\text{vis}} = h$, the SM-like Higgs scalar, produced by gluon fusion, vector boson fusion or associated production. The SUEP is then the result of an exotic Higgs decay into the strongly coupled hidden valley.

Other mediators are possible as well, and the phenomenology is largely independent of the choice of mediator. An important exception is the trigger efficiency by ATLAS and CMS, which greatly benefits from prompt, hard objects in the event, as discussed in the previous section. We assume a simplified fragmentation model, in which ϕ are produced spherically symmetrically, with a Maxwell-Boltzmann momentum distribution characterized by $T \sim \Lambda$. The multiplicity scales linearly with Λ , while for $\Lambda \sim T$ the typical boost remains roughly constant. This means that the sensitivity for different values of $\Lambda \sim T$ can be obtained by a simple rescaling.

There are many possibilities for the decay of the long-lived HV particles ϕ . Among many options for their decay, one might consider a kinetic mixing portal of the form $(\varepsilon/2)A'_{\mu\nu}B^{\mu\nu}$, with B the hypercharge field strength, and A' a hidden photon with mass $m_{A'} \ll m_Z$ that couples to ϕ via a chiral anomaly. Decays $\phi \rightarrow \ell\ell$ are not generated at tree level by this portal, but may occur at one-loop, assuming $m_{A'} > m_\phi$. The tree-level double-Dalitz process $\phi \rightarrow 2A'^* \rightarrow 4\ell$ may also proceed. However, the lifetime for these processes typically far exceeds 1 s for $\varepsilon \lesssim 10^{-4}$ and $m_{A'} \gtrsim 3$ GeV, which corresponds to the projected reach of other experiments sensitive to hidden photons, such as Belle II. A simple alternative is to consider the regime $2m_{\ell,\pi} < m_{A'} < m_\phi/2$, such that the $\phi \rightarrow 2A'$ proceeds promptly via the chiral anomaly, followed by $A' \rightarrow \ell\ell$ or $\pi\pi$ and so on, with rate $\Gamma \sim m_{A'}\varepsilon^2\alpha$. The hidden photon parametric region $\varepsilon < 10^{-8}$ and $1 \text{ MeV} < m_{A'} < 0.5 \text{ GeV}$ is unconstrained [547], and for example, for $\varepsilon \sim 10^{-8}$, this range of A' masses produces lifetimes $10^{-6} \text{ s} \lesssim \tau \lesssim 10^{-4} \text{ s}$ and lengths $1 \text{ km} \lesssim \lambda \lesssim 10^3 \text{ km}$. In this scenario, rather than ϕ , the long-lived A' effectively generates an effectively weak SM-HV portal, such that $\phi \rightarrow 2(A' \rightarrow \ell\ell)$ effectively has a very long lifetime. Based on this simple example, we therefore pick $\phi \rightarrow 4\ell$ as our benchmark ϕ decay mode. For $m_\phi = 1$ GeV,

⁵⁹We thank Mike Hance for providing us with more detailed results.




Fig. 47: Left: Number of expected events in MATHUSLA with at least one (two) vertices are shown in the dark (light) lines for our three benchmark models. For the $M = 400$ GeV and $M = 750$ GeV, we assumed a cross-section equal to that of a standard model Higgs boson with this mass [546]. The 125 GeV curve assumes production in exotic Higgs decays with branching fraction 10%. Also shown is the maximum number of events allowed by the expected mono-jet limits (dashed lines). Right: MATHUSLA exclusion potential (solid), compared to the projected mono-jet limits (dashed). This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

the main detector signal could therefore be that of a displaced lepton-jet. For such low LLP mass and assuming production in exotic Higgs decays (or another process without the guaranteed presence of conspicuous prompt objects to trigger on), this decay is actually the best-case scenario for a main detector LLP search. Even so, as discussed in Section 3.2.3, there are likely to be significant challenges triggering on and reconstructing these DVs without backgrounds. As such, MATHUSLA is likely to have significantly higher sensitivity than the main detectors in the long-lifetime limit. This advantage would be further compounded if the HV particles decayed dominantly into hadrons.

The resulting sensitivity for MATHUSLA as a function of the lifetime is shown in figure 47, as compared to the 3000 fb^{-1} projected jet+MET limits from ATLAS/CMS. For the $M = 750$ GeV and $M = 400$ GeV benchmarks, the jet+MET limits are adapted from [548]; for the Higgs portal benchmark we assume a maximum invisible branching ratio of 10% [549]. We find that MATHUSLA would significantly extend the reach of ATLAS and CMS for all benchmarks over almost all of the lifetime reach. At its peak sensitivity, MATHUSLA's detection efficiency for this signal is essentially order one, which means that the reach becomes luminosity limited. The reason is the very high multiplicity of spherically distributed, displaced decaying particles in each event, which effectively compensates for the loss in geometric acceptance as compared to a hermetic detector. Interestingly, there is also a sizable part of parameter space where one can expect events with more than one displaced vertex, which would be a smoking gun for a strongly coupled hidden valley. It is worth noting that a priori ATLAS and CMS themselves will have a significant geometric acceptance for displaced vertices from SUEP events. However with $\Lambda \sim T \sim 1$ GeV the average energy for the decay products, would be an extremely challenging, if not impossible search, and we do not attempt to estimate its sensitivity here. For $\Lambda \sim T \gtrsim 10$ GeV searches for displaced vertices at ATLAS and CMS could however become competitive with MATHUSLA if the final states are predominantly leptons. Hadronic final states pose a greater challenge for ATLAS and CMS in terms of triggering and background rejection, likely requiring a higher $\Lambda \sim T$. Finally, for $c\tau \lesssim 10$ m, it is likely to that ATLAS and CMS could constrain this scenario regardless the value of Λ , by searching for a specific pattern of hits in the inner detector [550].




Fig. 48: Solid lines show the MATHUSLA sensitivity to new particles X pair-produced in exotic Higgs decays, as a function of $c\tau_X$ and assuming the $200\text{ m} \times 200\text{ m} \times 20\text{ m}$ benchmark geometry of Fig. 1. The purple shading at the top of the plot shows projected exclusions from CMS $Br(h \rightarrow \text{invisible})$ searches [245] (although some projections are an $\mathcal{O}(1)$ factor better, see Section 3.2.1), which would also be sensitive to X s outside the blue shaded region. Dotted lines show projected ATLAS $Br(h \rightarrow XX)$ exclusions [138], which represent the best-case main detector reach projections for LLPs with very long lifetimes produced in exotic Higgs decays. Figure taken from Ref. [1].

8.2 Exotic Higgs Decays⁶⁰

One of the major discovery opportunities offered by the LHC is the search for new physics produced in exotic decays of the SM-like Higgs boson [551]. As for all newly discovered particles, a detailed characterization of the Higgs' decay modes is imperative. However, the SM Higgs is especially sensitive to the potential existence of new light degrees of freedom. The Higgs portal operator, $|H|^2$, is one of the two leading operators that can couple to new SM-singlet degrees of freedom, making the Higgs a natural window onto low-mass dark states. The very fact that $|H|^2$ is a low-dimensional operator and a singlet under all known symmetries of the SM, which lies at the root of the hierarchy problem, is what generically enables the Higgs to couple to all new physics to some degree. Discovery prospects are further enhanced thanks to the fortunate accident that all SM decay channels of the Higgs boson are suppressed, whether by phase space (WW^* , ZZ^*), loop factors (gg , $\gamma\gamma$, $Z\gamma$), or small Yukawa couplings ($b\bar{b}$, $c\bar{c}$, $\tau\bar{\tau}$, ...), resulting in an accidentally tiny SM Higgs width: $\Gamma(h \rightarrow \text{SM}) = 4.10\text{ MeV} \pm 0.73\%$ [510] for a $m_h = 125.09\text{ GeV}$ Higgs boson [552]. Thus even small couplings of the Higgs to new light degrees of freedom can easily yield substantial exotic branching fractions. The 3 ab^{-1} of data anticipated at HL-LHC will yield more than 10^8 Higgs bosons. This enormous data set could enable the discovery of exotic branching fractions as small as $\sim 10^{-6}$, provided that the signal can be both recorded and separated from background, which is frequently a stiff challenge at the LHC thanks to the low mass scales of Higgs events. MATHUSLA naturally provides a nearly background-free environment, enabling it to take full advantage of the Higgs sample produced at the HL-LHC. As shown in Fig. 49, MATHUSLA will be able to access Higgs branching fractions to LLPs down to the 10^{-5} level.

Exotic Higgs decays to LLPs appear throughout this document. In particular, they are one of the leading signals of many theories of neutral naturalness, extensively discussed in Sec. 4.2. Sec. 8.1

⁶⁰David Curtin, Jessie Shelton

discusses Higgs decays to (other) hidden valleys, while in Sec. 8.5 below we demonstrate that Higgs decays into dark photons offer a deep window into the parameter space of a kinetically mixed $U(1)$. The minimal Higgs portal model SM+S also has exotic Higgs decays as one of its leading signatures, and yields closely related signatures when the new scalar S is light enough to be produced in meson decays, as we discuss in Sec. 8.4 below. SM+S signatures arise naturally as signals of hidden sector dark matter or relaxation solutions of the hierarchy problem (Section 4.4). Finally, unification considerations motivate Higgs portal production of right-handed neutrino states (Section 7.4), giving access to much higher sterile neutrino masses than production in meson decays. In many of these examples, e.g. neutral naturalness, the produced LLPs decay back to the SM through Higgs portal couplings as well, predicting large LLP branching ratios to hadronic final states. These low-mass hadronic final states can be challenging at the LHC main detectors, but offer excellent prospects for MATHUSLA.

Higgs decays to ultra-LLPs X would dominantly appear as missing energy at the LHC when pair-produced in $h \rightarrow XX$. Given 3 ab^{-1} of data at 14 TeV, ATLAS and CMS have projected sensitivities to an invisible Higgs branching fraction of $\gtrsim \mathcal{O}(10^{-1})$ [245]; MATHUSLA would be able not only to establish the finite lifetime of the X particle but probe branching ratios four orders of magnitude smaller [1]. While proposed main detector searches for single LLPs that decay in the muon system can potentially approach the $Br(h \rightarrow XX) \gtrsim \mathcal{O}(10^{-3})$ level for X proper lifetimes in the range $1 \text{ m} \lesssim c\tau_X \lesssim 10 \text{ m}$ [138], Fig. 48 demonstrates that MATHUSLA will have orders of magnitude greater reach both in branching ratio and proper lifetime.

In Fig. 49 we show a general-purpose estimate of MATHUSLA’s sensitivity to a LLP X pair-produced in exotic Higgs decays. Contours show the branching fraction $Br(h \rightarrow XX)$ that can be tested at 95% CL, assuming a SM Higgs production cross-section. At fixed branching fraction and proper lifetime, lighter values for m_X result in a higher boost in the lab frame, and thus a suppressed probability of decaying within the detector volume. An exotic branching fraction of 10% can be tested for X lifetimes as long as $c\tau_X = 10^7 \text{ m}$, or $\tau_X = 0.03 \text{ s}$. As we discuss in the next section, MATHUSLA’s excellent reach allows it to approach the cosmic upper bound on possible X lifetimes. This also has important consequences for the interpretation of a $h \rightarrow$ invisible signal, if one is found at the HL-LHC: a (say) 10% invisible Higgs branching fraction could lead one to expect a signal at MATHUSLA if it is due to the production of LLPs with a lifetime below the BBN limit. Conversely, if no signal at MATHUSLA is found, this might add significant weight to the interpretation that the invisible Higgs decay did indeed produce cosmologically stable states, and hence candidates for DM.⁶¹

While for definiteness we show quantitative results for the simple decay $h \rightarrow XX$, when X is part of a larger hidden sector, it may dominantly appear together with additional particles, $h \rightarrow X + \dots$. Such multi-particle exotic Higgs decays can frequently be more challenging for the main detectors, as they distribute the relatively small Higgs energy among a large number of particles. For instance, ATLAS and CMS searches for $h \rightarrow$ invisible rely on large MET for triggering as well as background rejection, and additional soft visible particles appearing in the exotic decay can substantially suppress sensitivity. By contrast, MATHUSLA’s sensitivity does not strongly depend on the number or type of other particles produced together with X , and in fact tends to have *better* prospects for detecting such multi-species decays, since the ULLP X has a smaller boost and therefore a larger chance of decaying within the detector. Moreover, when large final state multiplicities arise through showering, this naturally results in a larger multiplicity of LLPs, and thus a larger probability of a decay within MATHUSLA.

⁶¹There are important caveats in this argument, most importantly the fact that for lighter LLPs, the resulting boost means that proper lifetimes right at the BBN ceiling are no longer probed. Nevertheless, the absence of a MATHUSLA signal would strengthen the case for the DM interpretation.




Fig. 49: MATHUSLA sensitivity to new particles X pair-produced in exotic Higgs decays. Contours indicate the value of $\log_{10} \text{Br}(h \rightarrow XX)$ that could be excluded at 95% CL, assuming SM Higgs production in gluon fusion with 3 ab^{-1} of data at $\sqrt{s} = 14 \text{ TeV}$. Note that MATHUSLA searches for the production of low-mass LLPs with $m_X \lesssim \text{GeV}$ from exotic Higgs decays may suffer from some backgrounds or lower reconstruction efficiency depending on the final detector design, see Section 3.1.4. This estimate assumes the $200 \text{ m} \times 200 \text{ m} \times 20 \text{ m}$ benchmark geometry of Fig. 1.

8.3 The BBN Bound⁶²

8.3.1 Introduction

The MATHUSLA proposal aims to search for long-lived exotic particles (LLPs) decaying away from the production point at the LHC [1, 9]. These LLPs are well-motivated theoretically (see *e.g.* [239, 422, 479, 553]) and both ATLAS and CMS provide robust lower bounds on their lifetimes from displaced vertices searches [131, 132, 134]. A natural place to look for an upper bound is cosmology, which together with the current collider constraints would define a clear band of interest in lifetimes for MATHUSLA. Ideally, for such an experimental search, the LLPs would have different coupling constants moderating their production and decays, with $\lambda_{\text{production}} \gg \lambda_{\text{decay}}$. Otherwise if one and the same λ were responsible for both the production and decay, large displacements would imply very inefficient production rates. Here we consider exotic decays of the Higgs boson to a pair of metastable states S .

Big bang nucleosynthesis (BBN), and its overall agreement with observations [554] (apart from the unclear status of ${}^7\text{Li}$) can provide a limit on the lifetimes of such particles with minimal assumptions. The thermal evolution to BBN temperatures involves self-depletion via $SS \rightarrow \text{SM}$ due to $\lambda_{\text{production}}$, in an expected WIMP-type annihilation process, and late-time decay of $S \rightarrow \text{SM}$ where depending on

⁶²Anthony Fradette, Maxim Pospelov

lifetimes and decay products the BBN outcome may get affected. These mechanisms are well-understood in the BBN literature (see *e.g.* [555, 556] for reviews) and a small and acceptable perturbation to the standard BBN (SBBN) outcome can be turned into a τ_S limit.

In this work, fully presented in Ref. [557], we analyze a fairly minimal model, where a new singlet scalar has predominantly a quadratic coupling to the Higgs boson that regulates both its production at colliders and its metastable cosmological abundance. We find that for most of the analyzed parameter space with $m_S < m_h/2$, the intermediate abundance of such particles is large enough to affect the neutron-proton freeze out ratios at relevant temperatures. This allows us to set fairly robust bounds on lifetimes of such particles, which come out to be remarkably strong, and shorter than 0.1 seconds, with mild dependence on the mass scale of S . In what follows we briefly review the model, its impact on the BBN; and present a summary of our results with a short discussion.

8.3.2 The minimal Higgs portal model

We consider the simplest extension of the SM by a singlet scalar field S . At the renormalizable level, the Lagrangian of the singlet sector (including the SM) generically takes the form

$$\mathcal{L}_{H/S} = \mu^2 H^\dagger H - \lambda_H (H^\dagger H)^2 - V(S) - ASH^\dagger H - \lambda_S S^2 H^\dagger H + \text{kin. terms.} \quad (217)$$

The self-interaction potential $V(S) = \lambda_4 S^4 + \lambda_3 S^3 + \frac{m_{S0}^2}{2} S^2$ can be redefined in such a way that the linear term is absent. It is important that the A , $\lambda_3 \rightarrow 0$ and $\langle S \rangle = 0$ limit would correspond to the case of stable S particles. To simplify the discussion without sacrificing much generality, we take $\lambda_{3,4} \rightarrow 0$ and assume $Av \ll m_{S0}^2$, $\lambda_S v^2$.

The physical mass of S receives a contribution from the electroweak symmetry breaking, $m_S = \sqrt{m_{S0}^2 + \lambda_S v^2}$. The two scalars develop a mixing angle and renders the S unstable via

$$\mathcal{L}_{\text{decay}} = S \times \theta \sum_{\text{SM}} O_h, \quad \theta = \frac{Av}{m_h^2 - m_S^2} \left(1 - \frac{\lambda_S v^2}{m_S^2} \right). \quad (218)$$

where O_h is the set of the standard Higgs interaction terms, with the Higgs field removed: *e.g.* $O_h = (m_f/v)\bar{f}f$ for an elementary SM fermion f .

Both θ and λ_S on their own are already subject to many observational constraints (see Refs. [219, 558, 559] for recent reviews). A generic feature is that λ_S is bounded by a maximal invisible Higgs branching ratio of 0.19 (at 2σ) [560]. With the well-predicted decay rate of the SM Higgs into SM particles of $\Gamma_{SM} = 4.08$ MeV [510] and

$$\Gamma_{h \rightarrow SS} = \frac{\lambda_S^2 v^2}{8\pi m_h} \sqrt{1 - \frac{4m_S^2}{m_h^2}}, \quad Br(h \rightarrow SS) = \frac{\Gamma_S}{\Gamma_S + \Gamma_{SM}} \simeq 10^{-2} \left(\frac{\lambda_S}{0.0015} \right)^2, \quad (219)$$

this translates into an upper bound on $\lambda_S \lesssim 0.007$ for $m_S \ll m_h$. If S is to be stable ($\theta \rightarrow 0$), such small λ_S would lead to an excessive abundance of S , which invalidates the Z_2 symmetric case, and forces us to include a non-zero decay term. From now on, we will consider $\theta \neq 0$, or in other words the case of unstable S particles. Since our analysis is motivated by the LHC physics, we will use $Br(h \rightarrow SS)$ as an input parameter, and substitute λ_S everywhere employing (219).

8.3.2.1 Decay products

Since S interacts with the SM in the same fashion as the Higgs with an additional θ mixing factor (218), its decay properties are similar to those of a light Higgs boson. For the derivations of the actual limits on the lifetime of S , we need to know its mesonic and nucleonic decay branching ratios, which are




Fig. 50: *Left:* Branching ratios of the scalar S in our baseline decay model. See text for details. *Right:* Scalar S lifetime of our baseline model and the spectator model for the mixing angle $\theta = 10^{-6}$.

still poorly understood and can vary by a few orders of magnitude near the di-kaon threshold [283]. In particular, the metastable mesons, such as π^\pm and K^\pm, \bar{K}^0, K^0 are “important” decay products, as they can participate in the charge-exchange reactions with nucleons and shift the $n - p$ balance, hence affecting the whole nucleosynthetic chain. We will show that two different decay models obtain similar constraints on the lifetime, thus minimizing the uncertainty in the decay rate.

The leptonic decay channels are straightforward, with the decay rate given by

$$\Gamma_{S \rightarrow l\bar{l}} = \frac{\theta^2 m_l^2}{8\pi v^2} m_S \left(1 - \frac{4m_l^2}{m_S^2}\right)^{3/2}. \quad (220)$$

If the decaying product is a pair of heavy quarks, there are $\mathcal{O}(1)$ corrections coming from the 1-loop QCD vertex correction [558]. For $m_S > 2.5$ GeV, we use the higher order perturbative results from the HDECAY code [561] which includes all necessary corrections.

In the mass range where the perturbative QCD calculations are no longer valid, we base our baseline calculations on Ref. [562]. The scalar-pion interaction can be extracted from the low-energy expansion of the trace of the QCD energy-momentum tensor (see for *e.g.* [563, 564]) yielding the effective decay rate [562]

$$\Gamma_{S \rightarrow \pi^+\pi^-} = 2\Gamma_{S \rightarrow \pi^0\pi^0} = \frac{\theta^2 m_S^3}{16\pi v^2} \left(\frac{2}{9} + \frac{11 m_\pi^2}{9 m_S^2}\right)^2 \sqrt{1 - \frac{4m_\pi^2}{m_S^2}}. \quad (221)$$

Final-state resonances however spoil this expression far from the threshold. To include kaons, we use an interpolation from Ref. [565], matching low-energy theorems to the dispersion results from the $\pi\pi$ phase-shift analysis above 600 MeV [566]. The photon decay channel is added with the prescription detailed in Ref. [567]. Finally, there is a gap for $1.4 \text{ GeV} < m_S < 2.5 \text{ GeV}$ where no analytical treatment is entirely trustworthy, we simply follow Ref. [562] and interpolate between the two regimes, under the assumption that there is no order of magnitude deviation in this mass range. The branching ratios and the lifetime for $\theta = 10^{-6}$ are displayed in Fig. 50.

As an alternative decay spectrum model, we also display the perturbative spectator approach [219, 568, 569], where the relative decay width above the kaon threshold are given by

$$\Gamma_{\mu^+\mu^-} : \Gamma_{KK} : \Gamma_{\eta\eta} = m_\mu^2 \beta_\mu^3 : 3 \frac{9}{13} m_s^2 \beta_K^3 : 3 \frac{4}{13} m_s^2 \beta_\eta^3, \quad (222)$$

with $\beta_i = \sqrt{1 - 4m_i^2/m_S^2} \Theta(m_S - 2m_i)$, Θ being the step-function, and we adopt the running of s quark mass following Ref. [567]. The pion contribution is kept as in equation (221) and then we use the HDECAY output at the c -quark threshold and above to match our baseline model.

For m_S of several GeV and heavier, decays with final state nucleon-antinucleon pairs are possible. Even though the branching to such states are generally lower than 10%, the effect on BBN can be quite significant, and therefore these are by far the most important channels for $\tau_S \gtrsim 1$ sec. On top of direct and for the most part subdominant contributions from $S \rightarrow \bar{n}n, \dots$, we need to take into account the (anti-)nucleon states that emerge from the hadronization of the quark decay products and heavy B -meson fragmentations.

8.3.2.2 Cosmological metastable abundance

Starting in thermal equilibrium with the SM, the S population eventually freezes out to a metastable abundance, via the s -channel annihilation $SS \rightarrow h^* \rightarrow XX$, where on the receiving end are the pairs of the SM states XX created by a Higgs-mediation process. The annihilation cross-section σv generically takes the form

$$\sigma v(s) = \frac{8\lambda_S^2 v^2}{(s - m_h^2)^2 + m_h^2 \Gamma_{\text{SM}+S}^2} \frac{\Gamma_{\text{SM}}^{m_h \rightarrow \sqrt{s}}}{\sqrt{s}}, \quad \langle \sigma v \rangle = \frac{\int_{4m_S^2}^{\infty} ds \sigma v(s) s \sqrt{s - 4m_S^2} K_1\left(\frac{\sqrt{s}}{T}\right)}{16T m_S^4 K_2^2\left(\frac{m_S}{T}\right)}. \quad (223)$$

This formula recasts the rate in terms of a SM Higgs width $\Gamma_{\text{SM}}^{m_h \rightarrow \sqrt{s}}$ evaluated at a fictitious mass of \sqrt{s} , thus encompassing both perturbative and non-perturbative channels in the h^* decay rate, which is the same as Γ_S with $\theta = 1$. Since the nonrelativistic annihilation cross-section in the minimal Higgs portal model ranges from 10^{-3} to 10^{-14} pb for $m_S \sim 1$ MeV – 60 GeV and $Br(h \rightarrow SS) \sim 0.1 - 0.001$, the standard nonrelativistic WIMP freeze-out approximation is not applicable and we numerically integrate the standard Boltzmann equation to determine the metastable S abundance. The results are shown in Fig. 51, normalized to the baryon number density for a more intuitive interpretation of its impact on BBN in the following section. Qualitatively, it is clear that the relative inefficiency of annihilation through the Higgs portal will leave behind a fairly significant population of S particles, which will eventually lead to strong constraints on τ_S .

For very light m_S , one can see that the freeze-out abundances are large, and the relative spread between different input values of $Br(h \rightarrow SS)$ gets smaller, as the annihilation cross-section becomes very small. Such small cross-sections mean that freeze-out happens in the semi-relativistic regime $x_{\text{f.o.}} \sim \mathcal{O}(1)$ and asymptote to the Y_{eq} relativistic plateau for small m_S . The only difference at the lightest masses is from $Y_{\text{eq}}^{\text{rel}} \propto 1/g_{*S}(T)$, where g_{*S} is the number of effective degrees of freedom appearing in the entropy density. Since g_{*S} is a monotonic function of temperature, weaker annihilation cross-sections freeze out earlier, at a higher temperature, thus yielding smaller abundances (as seen in the $m_S = 5$ MeV curves in Fig. 51). This is in contrast with the standard freeze out in the non-relativistic regime, with final abundances inversely proportional to the cross-section. We note in passing that the strong-interaction-related uncertainty “propagates” outside the $m_S \sim 2m_\pi - 2m_c$ window. For example, because of the relativistic freeze-out, for $m_S < 2m_\pi$ the hadronic channels may turn out to be important.

8.3.3 Big Bang Nucleosynthesis

The formation of light nuclei is one of the earliest probes of NP in cosmology and is well-understood within SM physics. Modulo the ${}^7\text{Li}$ discrepancy [554], the overall success of BBN in predicting the more populous element abundances can be used to constrain various types of NP [556].

The initial BBN stage is the neutron-proton ratio n/p freeze out. Maintained in equilibrium by electroweak interactions at high temperatures, the neutron abundance follows $n/p \sim e^{-Q/T}$, where $Q = m_n - m_p - m_e \simeq 1.293$ MeV, until the epoch when the weak processes decouple around temperatures of 0.7 MeV. The outcome, $n/p \simeq 1/6$, is quasi-stable, decreasing to $n/p \simeq 1/7$ at the end of the “deuterium bottleneck”. At $t_{\text{deut}} \sim 200$ seconds, ${}^4\text{He}$ formation is very efficient and most neutrons end up in the final ${}^4\text{He}$ abundance (expressed in mass fraction from the total baryon mass) $Y_p \simeq 2(n/p)/(1 + n/p) \simeq 0.25$.




Fig. 51: *Left:* Temperature evolution ($x = m/T$) of the Y_S intermediate abundance for $m_S = 5$ MeV and 500 MeV for the three benchmark higgs branching ratios. *Right:* Metastable abundance of S prior to its decay normalized over the baryon density. Values shown for $Br(h \rightarrow SS) = 10^{-1}, 10^{-2}$ and 10^{-3} . The dashed lines correspond to the perturbative spectator model.

For the problem at hand - the determination of the upper limit on the S lifetime - few of the finer BBN details matter. The ample decaying abundances of S particles ($n_S \sim 10^2 - 10^9 \times n_b$) will flood the neutron-proton bath with SM particles prior to the bottleneck, inducing charge-exchange reactions that will modify the n/p freeze out ratio. For each decay products X that can modify n/p , we solve the neutron-proton Boltzmann equation

$$\frac{dX_n}{dT} = \frac{\Gamma_{n\nu_e \rightarrow pe^-} + \Gamma_{ne^+ \rightarrow p\bar{\nu}_e}}{TH(T)} \left(X_n - (1 - X_n)e^{-Q/T} \right) + \frac{\Gamma_n X_n}{TH(T)} + \left. \frac{dX_n}{dT} \right|_X, \quad (224)$$

including the new charge-exchange term and require that Y_p does not deviate from SBBN by more than 4%,

$$\Delta Y_p \equiv |Y_p - Y_p^{\text{SBBN}}| < 0.01, \quad (225)$$

which is a rather generous allowance for the errors, considering the tight observational constraints on primordial helium abundance [554]. Consequently, it will result in conservative limits of τ_S . We show typical deviations from the SBBN case for pion, kaon, direct baryon⁶³ and neutrino (from muon decays) injections in figure 52, along with the maximal abundance (weighted by its S branching ratio ξ) that satisfy the $\Delta Y_p < 0.01$ requirement. We refer the reader to the complete paper [557] for additional detail, including direct SS annihilations to charged pions and N_{eff} deviations from energetic electron or muons.

8.3.4 Results and Discussion

Combining the constraints on each energy injection mode, the surviving parameter space of the minimal Higgs model is shown in Fig. 53 as a function of the scalar mass m_S and lifetime τ_s . The assumptions considered in each mass range, labelled from A to G are described in Ref. [557].

We find that throughout almost the whole mass range considered in this work, $2m_\mu < m_S < m_h/2$, the constraints on the lifetime of S particles are stronger than 0.1 seconds, with only a mild dependence on $Br(h \rightarrow SS)$ as the amount of decaying S is much larger than the baryon abundance in all cases. From the standpoint of LHC physics, the most notable portion of parameter space is at relatively large masses, where m_S is not far below $m_h/2$. In that case, the proper decay length has to be on the order of or smaller than 2×10^7 meters, and comparing with Fig. 49, this places the upper bound on X lifetimes in the same overall range that can be probed by MATHUSLA.

⁶³We tuned the injected baryons after hadronization to $N_n = \kappa N_p$ and $N_{\bar{n}} = \kappa N_{\bar{p}}$, where a phenomenological hadronization parameter κ is expected to be $\simeq 0.5$ from simple quark counting of the main weak decay chain [557].




Fig. 52: *Left:* X_n evolution for the SBBN and the injection of pions, kaons, baryons and muons (neutrinos) as described in the text for lifetimes of 0.05 seconds with the initial Y_S abundance tuned to yield $\Delta Y_p = 0.01$ (maximum allowed shift of Y_p). The baryonic injection is taken at $\kappa = 0.5$ (full line), the lines for $\kappa = 1$ (dashed) and $\kappa = 0.2$ (dotted) are also displayed. *Right:* Limit of injected pairs for each channel as a function of the S lifetime. The upper-right dotted line for $\kappa = 0.2$ is at $Y_p = 0.26$, the upper-left dotted island yields $Y_p = 0.24$.




Fig. 53: Lifetime constraint as a function of the S mass for three $h \rightarrow SS$ branching ratios. The lettered regions represent different assumptions or physics and are described in the text. The dotted lines correspond to the perturbative spectator model.

The above considerations can be generalized to other models of the Higgs-portal-coupled particles or even different types of interactions, via Z, Z' etc. For example, consider a fermion χ , coupled to the Higgs via $H^\dagger H(\bar{\chi}\chi)$ or $H^\dagger H(\bar{\chi}i\gamma_5\chi)$ dimension-five operators, and having a small decay term such as *e.g.* neutrino portal $LH\chi$. The main analysis of our work can be recast for that model, especially in the part that connects Higgs decays with a metastable abundance of χ . Evidently, for the same input values of $Br(h \rightarrow \chi\bar{\chi})$ and $Br(h \rightarrow SS)$, one will end up with $Y_\chi \sim Y_S$. The only change will be in the yields of mesons and baryons in the decays of χ compared to S . However, as the yield of pions and kaons in χ decays is already known to be substantial for $m_\chi > 250$ MeV [219], we expect that for the most part the constraints we have derived for τ_S will translate to similar limits on τ_χ . On the other hand, these constraints can be evaded if there are additional degrees of freedom to deplete the energy outside the BBN lifetime window, at the expense of additional complication of the model.

8.4 SM + S: Singlet Extensions⁶⁴

Hidden sectors populated with new particles that are only very weakly coupled to the standard model are well-motivated. A simple extension that includes new weakly-coupled scalars has been used to explain a wide variety of outstanding deficiencies in the standard model, such as dark matter [570–572], the $(g-2)_\mu$ anomaly [573, 574], inflation [562], naturalness [235, 238], neutrino masses [6], and the proton radius puzzle [575–577]. A new scalar, S , can be coupled to the standard model via the renormalizable Higgs portal interaction, $\epsilon |S|^2 H^\dagger H$. A minimal simplified model can be constructed with the scalar Lagrangian [18],

$$\mathcal{L}_{\text{scalar}} = \mathcal{L}_{\text{kin}} - \frac{1}{2}\epsilon S^2 H^\dagger H + \frac{1}{2}\mu_S S^2 - \frac{\lambda_s}{4!}S^4 + \mu_H^2 H^\dagger H - \lambda_H (H^\dagger H)^2 \quad (226)$$

for real scalar, S , with an imposed discrete symmetry $S \rightarrow -S$ to prevent all terms cubic and linear in S . Adding these terms complicates the sector, but does not qualitatively alter the physics from the story presented here. If both S and H have nonzero vacuum expectation values, $S = s + v_s$ and $H = (h + v_h)/\sqrt{2}$, the two scalar states will mix. For portal coupling $\epsilon \ll 1$, v_h and one of the mass eigenvalues can be identified as the usual Higgs vev, $v_h = 246$ GeV, and observed Higgs mass, $m_h = 125$ GeV, while the remaining three parameters in (226) can be identified with the mass of the scalar, m_s , a mixing angle between the two sectors, $\sin \theta = \frac{\epsilon v_h v_s}{m_h^2 - m_s^2} + \mathcal{O}(\epsilon^3)$, and the coupling of the Higgs with two hidden sector scalars,

$$\mathcal{L} \ni \frac{\kappa}{2} h s^2 = \frac{1}{2} \sqrt{\frac{\lambda_s}{3}} \sin \theta \left(\frac{m_h^2 + 2m_s^2}{m_s} \right) h s^2. \quad (227)$$

The hidden sector scalar couples to standard model fermions and vector bosons as a standard model Higgs, but with strength reduced by a factor of $\sin \theta$. The scalar width is thus reduced by $\sin^2 \theta$ from a standard model-like Higgs of the same mass, *i.e.*, $\Gamma_s = \sin^2 \theta \Gamma_{h,SM}(m_s)$, which, for sufficiently small mixings, results in a particle that is long-lived on collider timescales. The width of a hadronically decaying light scalar has a high degree of uncertainty for scalar masses between $2m_\pi$ and ~ 4 GeV (see refs. [283, 568] for more details). It is common in the literature to use a perturbative spectator model in this region, see *e.g.* [219, 557, 568, 569], but there are some reasons to suspect that this approach may be underestimating the scalar's partial width into hadrons. As the scalar mass is increased, it crosses through several dozen hadron mass thresholds that open up more and more accessible decay channels, and, as the Higgs-mixed scalar couples to mass and $\Lambda_{QCD} > m_s$, these channels may provide large corrections. Across this region there will be many scalar meson resonances for the state to mix with (including the observed $f_0(1370)$, $f_0(1500)$, $f_0(1710)$, and the $f_0(980)$, the latter of which is responsible for the extrema near 1 GeV in ref. [565], and shown in figure 54), and even rather broad states that have never

⁶⁴Jared Evans




Fig. 54: Branching ratios assumed in the SM+S model for the additional scalar in the light hadron region. For masses below ~ 1.4 GeV, the calculation of ref. [565] is used. We implement an extrapolation in the region from 1.4 GeV to $2m_D$ that yields a larger partial width into hadrons than predicted in the perturbative spectator model [568].

been observed could appreciably amplify the hadronic decay width [578] relative to the predictions of the perturbative spectator model. Motivated by these uncertainties, in this region we use the branching fractions in ref. [18], shown in figure 54.⁶⁵

At the LHC, the long-lived scalars can be produced in exotic Higgs decays. Allowing for m_s and $\sin \theta$ to assume any value still places a restriction on how large κ can be when one mandates perturbativity of λ_s ($\lambda_s < 16\pi^2$). The maximum allowed $h \rightarrow ss$ branching ratio is then,

$$\text{BR}(h \rightarrow ss)_{MAX} = \frac{\pi \sin^2 \theta m_h^3}{3m_s^2 \Gamma_{h,tot}} \left(1 + 2 \frac{m_s^2}{m_h^2}\right)^2 \sqrt{1 - 4 \frac{m_s^2}{m_h^2}} \quad (228)$$

where $\Gamma_{h,tot} \approx 4.15$ MeV. While it is possible that the addition of linear and cubic interactions could relax this constraint slightly, it cannot be modified parametrically without introducing additional particles, for instance, a two-site model with $h \rightarrow s_1 s_1$ followed by $s_1 \rightarrow s_2 s_2$.

Additionally, this scalar could be emitted in rare meson decays. For the purpose of MATHUSLA, B -mesons are the most relevant production mechanism, and the inclusive branching fraction can be expressed as [19, 580]

$$\text{BR}(B \rightarrow s X_s) = \frac{27\sqrt{2}G_F m_t^4}{64\pi^2 \Phi m_b^2} \left| \frac{V_{ts}^* V_{tb}}{V_{cs}} \right|^2 \left(1 - \frac{m_s^2}{m_b^2}\right)^2 \sin^2 \theta \text{BR}_{B \rightarrow X_c e \nu_e} \approx 6.2 \left(1 - \frac{m_s^2}{m_B^2}\right)^2 \sin^2 \theta. \quad (229)$$

where $\Phi \approx 0.5$ [581] is a phase space factor for the semi-leptonic decay.

Following [18], we can estimate the sensitivity of MATHUSLA to s particles produced in either meson decays or exotic Higgs decays. Events are generated for $h \rightarrow ss$ and $b\bar{b}$ production in Pythia 8 [158] at 14 TeV. From the kinematic distributions, we can compute the probability that s decays within the MATHUSLA detector volume. Joining these to either the 14 TeV Higgs production cross-section of 62.6 pb [510, 582] or an assumed $b\bar{b}$ production cross-section of 0.5 mb [107] and a projected luminosity of 3 ab^{-1} to determine the number of events that would decay in MATHUSLA. We project to 95% CL constraints assuming zero background and a 75% detection efficiency (4 scalar LLP decays observed). In the case of meson decays, most of the scalars have energies above 2 GeV [18], but for masses below \sim

⁶⁵The larger hadronic partial width and smaller muon branching ratio in this interpretation combine to enlarge the gap in coverage between LHCb (from above) and long lifetime experiments like MATHUSLA, SHiP [429], CODEX-b [19], or FASER [17] (from below), which makes this choice conservative with regards to ensuring complete coverage in the gap between the different experimental approaches.




Fig. 55: The projected sensitivity of MATHUSLA to scalar LLPs in the minimal SM+S extension after 3 ab^{-1} of 14 TeV LHC assuming 4 events, and assuming the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1. The red contour is the sensitivity to B -meson decays (Kaons would provide percent level corrections). The different blue-purple contours illustrate the minimum $\text{BR}(h \rightarrow ss)$ value to which MATHUSLA would be sensitive. The feature near 1 GeV is due to the peak in the partial width to hadronic states that appears in the Donoghue, Gasser, Leutwyler [565] modeling used in this region (between $1.4 - 4$ GeV, the interpolation from [18] is used). With decreasing (increasing) $\sin \theta$ or m_s the lifetime grows (shrinks). The overall shape of the sensitivity at higher masses is heavily sculpted by the maximum allowed value of $\text{BR}(h \rightarrow ss)$ consistent with perturbative couplings in the theory. The projected constraint contour for the SHiP experiment [219, 429, 579] is shown by the dashed orange contour. Current constraints are described in the text.

10 MeV, their detection efficiency may be significantly degraded and depends on details of the detector design, see Section 3.1.4.

The projected sensitivity for both meson decays and exotic Higgs decays are shown in Fig. 55 in the plane of m_s vs mixing angle $\sin \theta$. While the meson decay constraints are robustly determined by the position in this parameter space, the exotic Higgs decay constraints depends on an additional parameter, the $h \rightarrow ss$ branching ratio, which has no lower bound. For arbitrarily small λ_s (i.e., large v_s/m_s), all sensitivity to this channel can vanish. We show with the blue-purple contours in Fig. 55 constraints that will arise for different choices of the branching ratio. We additionally require that the maximum allowed branching ratio (228) is consistent with the resulting limit, which sculpts the shape of the contours at high mass. Also shown (computed with the same assumptions for scalar decay widths) are the projected constraint contour for the SHiP experiment [219, 429, 579] in dashed orange, and current constraints on

the parameter space from LEP Higgs searches (light red) [583, 584], $K^\pm \rightarrow \pi^\pm + \text{invisible}$ at E949 & E787 (light blue) [585], the CHARM beam dump (gold) [586], and rare B decays at LHCb (light green and brown) [587, 588]. Not shown are other proposals to find light, Higgs-mixed scalars at CODEX-b [19] and FASER [17].

In summary, MATHUSLA would allow us to peer deeply into the SM+S parameter space, both via exotic Higgs decays and via meson decays. In the former case, the sensitivity is orders of magnitude better than main detector LLP searches, as discussed in Section 8.2. The latter are even more challenging at the main detectors, and is the target of proposed experiments like SHiP. MATHUSLA would be able to extend the reach of these experiments to significantly smaller mixing angles.

8.5 SM + V: Dark Photons⁶⁶

Dark sectors can contain mediator particles that allow for interactions with SM particles through portals, see e.g. [547, 589–592] for recent reviews. If the dark sector contains a dark abelian gauge group $U(1)_D$, this may give rise to what is known as the “vector portal”, a renormalizable kinetic mixing between the *dark photon* and the SM hypercharge gauge boson: [593, 594],

$$\mathcal{L} \supset -\frac{\epsilon}{2 \cos \theta_W} F'_{\mu\nu} F_Y^{\mu\nu}, \quad (230)$$

Here ϵ is the kinetic-mixing parameter, θ_W is the Weinberg mixing angle, $F'_{\mu\nu} = \partial_\mu A'_\nu - \partial_\nu A'_\mu$ is the $U(1)_D$ field strength, and similarly $F_Y^{\mu\nu}$ denotes the SM hypercharge $U(1)_Y$ field strength. This mixing allows A' s to be produced in charged particle interactions. The value of ϵ is arbitrary, but a value of $\epsilon^2 \sim 10^{-8} - 10^{-4}$ is natural if generated by quantum effects of heavier particles charged under $U(1)_D$ and $U(1)_Y$. Since the operator is renormalizable, new physics effects at any scale can generate detectable kinetic mixings. If the SM forces unify in a Grand Unified Theory, then $\epsilon^2 \sim 10^{-12} - 10^{-6}$ is natural [595–597].

The dark photon can be massive if $U(1)_D$ is broken, the most obvious mechanism for this breaking being a dark Higgs mechanism at a scale close to the dark photon mass. This is called the Hidden Abelian Higgs Model, see e.g. [598] for a full description. The massive dark photon can then decay into SM particles through the small kinetic mixing, making it a possible LLP. The dark Higgs h_D is expected to have some degree of mixing with the SM Higgs via the Higgs portal

$$\mathcal{L} \supset \kappa |h_D|^2 |H|^2, \quad (231)$$

since such a term cannot be forbidden by symmetries. This mixing would provide another production mode for dark photons in exotic Higgs decays.

Fig. 56 shows constraints on an A' with mass between the MeV and the weak scale, assuming A' decays only to SM particles [574, 599–618]. The A' mass is arbitrary, but this range arises naturally in several models [595, 619–621]. An A' can also explain the discrepancy between the measured and calculated value of the muon anomalous magnetic moment [574, 622] (Fig. 56, green band), although non-SM decays are needed.

8.5.1 A' as a LLP

We first consider the case that the dark photon is the lightest (or only) dark-sector particle. In this case, once produced, it can decay only to three photons (for $m_{A'} < 2m_e$) or two charged SM particles (for $m_{A'} > 2m_e$). However, the decay length to three photons is much larger than 200 m even for $\epsilon \sim 1$. Since the reconstruction of such a light LLP that decays only to photons would anyway be very

⁶⁶Nikita Blinov, Jae Hyeok Chang, David Curtin, Rouven Essig, Brian Shuve




Fig. 56: Existing 95% confidence level limits (gray shaded regions) on dark photons (A') and proposed experimental searches. An A' inside the green band can explain the discrepancy between the measured and calculated value of the muon anomalous magnetic moment. For a figure showing all proposed searches and references, see e.g. [591].




Fig. 57: Dark photon lifetime for $\epsilon = 10^{-8}$ when A' can only decay to SM particles. $c\tau$ scales as ϵ^{-2} .

challenging at MATHUSLA, we only consider the regime $m_{A'} > 2m_e$. In this case, the decay width to electrons is given by

$$\Gamma(A' \rightarrow e^-e^+) = \frac{1}{3}\epsilon^2\alpha \left(1 + \frac{2m_e^2}{m_{A'}^2}\right) \sqrt{m_{A'}^2 - 4m_e^2} \Theta(m_{A'} - 2m_e). \quad (232)$$

with similar expressions for the other SM fermions at tree-level. Below the $\bar{b}b$ threshold, threshold effects and hadronic corrections cannot be neglected, but they can be accounted for using $e^+e^- \rightarrow$ hadrons experimental data, see e.g. [598] for details. For small ϵ , the dark photon will be long-lived, see Fig. 57.




Fig. 58: Examples of the Feynman diagrams for dark photon production at the LHC. Left: pion decay to the dark photon and a SM photon. Middle: bremsstrahlung of the dark photon during gluon-gluon collision. Right: exotic Higgs decay to two dark photons via mixing with the dark Higgs.

There are several production channels for dark photons at the LHC. Two important processes that assume only the presence of kinetic mixing are meson decays and bremsstrahlung processes, see Fig. 58. Assuming only kinetic mixing, the abundant meson production rates at the LHC in QCD jets mean their decay is the dominant production mechanism for dark photons, as long as the dark photon mass is below the meson mass. The most important meson decays are pion and eta decays, while other mesons contribute only a negligible amount to the total production. At higher masses, bremsstrahlung processes are the dominant source of directly produced dark photons. However, since both of these processes have dark photon production rates which scale as ϵ^2 , they lead to very small signal rates in the long-lifetime regime that is accessible by MATHUSLA. As a result, the only regions of $(m_{A'}, \epsilon)$ parameter space that lead to a MATHUSLA signal assuming dark photon production via kinetic mixing are already excluded from past beam-dump experiments [591].

Dark photons can also be produced in exotic Higgs decays, shown on the right in Fig. 58. This production rate depends on the mixing between the dark and the visible Higgs, leading to possible branching ratios as large as $\sim 10\%$. Since this production rate does not depend on ϵ , it allows for very small kinetic mixings to be probed if the search is sensitive to long lifetimes.

Similar to the LLP searches studied in [598], we can show the regions of $(m_{A'}, \epsilon)$ parameter space that can be probed assuming a certain exotic Higgs branching ratio $\text{Br}(h \rightarrow A'A')$ in Fig 59. The MATHUSLA sensitivity, corresponding to 4 decays in the detector, is shown as the blue contours. For comparison, we also show the sensitivity of an ATLAS search for a single LLP decay in the inner tracker (see definitions in Section 3.2.2), where the dark photon LLP is required to decay leptonically for triggering and background rejection purposes. This main detector search is also assumed to be background-free, but this is likely too optimistic, especially for dark photon masses below ~ 10 GeV, see discussion in Section 3.2.3. Even with these generous assumptions for the ATLAS search, MATHUSLA is able to probe about an order of magnitude smaller kinetic mixings down to $\epsilon \sim 10^{-12}$, representing the greatest sensitivity to small mixing possible at any experiment in that mass range.

8.5.2 A' decaying to LLP's

We next consider the highly generic possibility that there are additional dark-sector particles charged under $U(1)_D$ to which the dark photon can decay. For example, the $U(1)_D$ could be part of a confining hidden valley [21, 597] that gives rise to bound states of the hidden QCD-like interaction in the IR. If these hidden hadrons have mass scale m_D below the dark photon mass, the width for their decay to SM fermions via an off-shell A' is roughly

$$\Gamma \sim \frac{\alpha_D \alpha \epsilon^2}{18\pi} \frac{m_D^5}{m_{A'}^4}. \quad (233)$$




Fig. 59: Regions of dark photon parameter space that could be probed by an LLP search at MATHUSLA (blue) or the ATLAS inner tracker (orange) assuming dark photon production in exotic Higgs decays with the indicated $\text{Br}(h \rightarrow A'A')$. Note that MATHUSLA searches for the production of low-mass dark photons with $m_{A'} \lesssim \text{GeV}$ from exotic Higgs decays may suffer from some backgrounds or lower reconstruction efficiency depending on the final detector design, see Section 3.1.4. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

It is then possible for ϵ to be relatively large, leading to sizable dark photon production rates through the kinetic mixing operator, while the dark hadron decay length could easily be much larger than the size of the main detectors for modest hierarchies of $m_D/m_{A'}$. Just as was the case for dark photon LLP production via exotic Higgs decays, LLP lifetime is now largely decoupled from the LLP production rate (now via dark photon decays). This is a prime signal to search for at MATHUSLA, especially (but not only) in the regime where the hidden hadrons have mass below ~ 10 GeV, leading both to long lifetimes and making the background-free reconstruction of the associated displaced vertices at the main detectors more difficult.

To understand the number of LLPs that might be produced in dark photon decay, we first need the *total* dark photon production rate at the LHC as a function of ϵ and $m_{A'}$. (Here we assume only production processes that rely on kinetic mixing.) For $m_{A'} \lesssim m_\eta$, meson decay is the dominant dark photon production mode. At higher masses, bremsstrahlung production $pp \rightarrow A' + 2$ jets is the most important process (we checked that $pp \rightarrow A' + 1$ jet is smaller).

To estimate dark photon production in meson decays, we use the event generator SIBYLL2.3 [623, 624] to compute the total number of pions and etas produced at the HL-LHC:

$$N_\pi = 1.4 \times 10^{19} \quad \text{and} \quad N_\eta = 3.0 \times 10^{17}. \quad (234)$$




Fig. 60: Cross-sections for the bremsstrahlung of dark photon at LHC in terms of dark photon mass. Here, $\epsilon = 10^{-3}$ is chosen, and the cross-section is proportional to ϵ^2 .

The branching fraction to dark photons are

$$\text{Br}(\pi[\eta] \rightarrow A'\gamma) = 2\epsilon^2 \text{Br}(\pi[\eta] \rightarrow 2\gamma) \left(1 - \frac{m_{A'}^2}{m_{\pi[\eta]}^2}\right)^3 \Theta(m_{\pi[\eta]} - m_{A'}) \quad (235)$$

$$(236)$$

The dark photon production rate in bremsstrahlung processes is computed in MadGraph 5 [625] with a minimum jet p_T of 10 GeV.⁶⁷ It also scales with ϵ^2 and is shown as a function of dark photon mass in Fig. 60.

We now assume that the A' decays directly to two long-lived LLPs with mass m_{LLP} and lifetime $c\tau$. Given the possible high multiplicity of the produced dark mesons after hidden sector hadronization, this is a simplistic assumption that likely underestimate the LLP signal. Nevertheless, it is instructive to show what range of hidden sector lifetimes could be probed at MATHUSLA for different kinetic mixings.

We use simulations (as described above) to calculate the average boost of the LLP, and confirm that Eqn. (4) gives a very good analytical estimate of the number of observed decays in MATHUSLA in the long-lifetime regime if we use $\epsilon_{\text{geometric}} \sim 0.02$. Fig. 61 then shows the resulting number of LLPs that decay in the MATHUSLA detector for various $m_{A'}$. For $m_{A'} = 10$ MeV and $m_{A'} = 100$ MeV, production from meson decay dominates, so we use the results from SIBYLL2.3 with assumptions that the LLPs are co-linear to the decaying mesons and carry a quarter of the decaying meson's momentum. For the other values of $m_{A'}$ shown in the figure, bremsstrahlung is the only relevant production process, and we use MadGraph 5 to simulate $pp \rightarrow 2 \text{ jets} + (A' \rightarrow X \bar{X})$, where X is the LLP charged under $U(1)_D$, assumed to have mass $m_X \ll m_{A'}$. As expected, the number of events has a peak near $\bar{b}c\tau \sim 200$ m. In this model scenario, MATHUSLA can probe the parameter space between the intersections of the thick lines and the grid line, where the grid line indicates 4 signal events.

⁶⁷A more sophisticated matched calculation would give a somewhat higher dark photon yield, but our conservative estimate is sufficient to demonstrate the importance of dark photons as a potential LLP source.




Fig. 61: Projections at MATHUSLA for the case that the dark photon decays to two LLP's. The number of signal is presented in terms of $\bar{b}c\tau$ of the LLP for different dark photon masses. Here, $\epsilon = 10^{-3}$ is used, and N_{Signal} is proportional to ϵ^2 . For $m_{A'} = 10\text{MeV}$ and $m_{A'} = 100\text{MeV}$, the dominant dark photon production process is meson decay, and bremsstrahlung of dark photon for higher masses. The grid line is for $N_{\text{Signal}} = 4$. \bar{b} 's from the numerical simulations are $\bar{b} = 1.7m_\pi/m_{\text{LLP}}$ for $m_{A'} = 10\text{MeV}$ and $m_{A'} = 100\text{MeV}$, $\bar{b} = 2.0m_{A'}/m_{\text{LLP}}$ for $m_{A'} = 1\text{GeV}$, $\bar{b} = 0.80m_{A'}/m_{\text{LLP}}$ for $m_{A'} = 10\text{GeV}$, and $\bar{b} = 0.45m_{A'}/m_{\text{LLP}}$ for $m_{A'} = 100\text{GeV}$. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.




Fig. 62: Feynman diagram illustrating dark Higgs (h_{D}) and dark photon (A') production in rare Z boson decays. The dark photons in h_{D} decay can be on- or off-shell.

8.5.3 Long-lived dark Higgs production in exotic Z decays

An important test of the Hidden Abelian Higgs Model is the verification of the origin of symmetry breaking in the hidden sector. If a dark Higgs boson, h_{D} , is responsible for giving mass to a hidden photon, A' , then the dark Higgs can be produced in association with a dark photon: this is the dark Higgs-strahlung process [626, 627]. Depending on the magnitude of the hidden-sector gauge coupling, α_{D} , the dark Higgs-strahlung process can give the best sensitivity to hidden sector parameters. There exist searches for dark Higgs-strahlung at B -factories via the process $e^+e^- \rightarrow A'^* \rightarrow A'h_{\text{D}}$, which set the best limits on the hidden sector for large values of $\alpha_{\text{D}} \gtrsim 10^{-3} - 10^{-2}$ [628, 629]. However, there are currently no searches at the LHC that are sensitive to this process.




Fig. 63: Sensitivity of MATHUSLA to the dark Higgs-strahlung process is shown in bold solid lines for the case where $\alpha_D = 0.1$, $\varepsilon = 10^{-3}$, and $m_{h_D} = 15$ GeV with varying $m_{A'}$ mass. The colormap shows the approximate event yield and demonstrates that MATHUSLA has sensitivity to proper lifetimes in the $10 - 10^4$ m range. By comparison, the sensitivity of ATLAS and CMS is shown by the region within the gray dotted lines; their reach is optimal for lower lifetimes. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.




Fig. 64: Sensitivity of MATHUSLA to the dark photon parameter space for the particular case of $c\tau = 100$ m, $\alpha_D = 0.1$. MATHUSLA sensitivity is shown with solid lines while ATLAS/CMS sensitivity is shown with dashed lines. The h_D masses are given in GeV. For comparison, constraints from LHCb [630], electroweak precision observables (EWPO) [631], and at the far left, BABAR [614]. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

When kinematically accessible, the dominant Higgs-strahlung process at the LHC is in the rare Z boson decay, $Z \rightarrow A'h_D$, shown in Fig. 62. The dark Higgs subsequently decays into on- or off-shell A' , so that the final signature is $Z \rightarrow A'A'^{(*)}A'^{(*)}$, leading to final states with multiple resonances and high multiplicities of soft leptons. This scenario has been proposed and recently studied in detail in Ref. [627]. In particular, this rare Z decay mode is promising for moderately large $\alpha_D \gtrsim 0.05$, and small SM-dark Higgs mixing ($\kappa^2 \ll \alpha_D$). The branching fraction of $Z \rightarrow A'h_D$ is

$$\text{Br}(Z \rightarrow A'h_D) = \frac{2\alpha_D \varepsilon^2 \tan^2 \theta_W m_{A'}^2 m_Z}{3(m_Z^2 - m_{A'}^2)^2} \left[1 + \frac{(m_Z^2 + m_{A'}^2 - m_{h_D}^2)^2}{8m_Z^2 m_{A'}^2} \right] \beta \left(\frac{m_{A'}}{m_Z}, \frac{m_{h_D}}{m_Z} \right), \quad (237)$$

where $\beta(x, y) = [(1 - (x - y)^2)(1 - (x + y)^2)]^{1/2}$ and we assumed that $\varepsilon \ll 1$ and that the $h - h_D$

mixing angle is negligibly small.

Most relevant for MATHUSLA is the parameter regime where the dark Higgs, h_D , is long-lived. This occurs when $m_{h_D} < m_{A'}$, in which case the dark Higgs cannot decay to on-shell dark photons. In this case, the h_D typically decays radiatively via off-shell dark photons into two SM fermions $f\bar{f}$,

$$\begin{aligned}\Gamma(h_D \rightarrow f\bar{f}) &\sim \frac{\alpha^2 Q_f^4 \alpha_D \varepsilon^4}{32\pi^2} \left(\frac{m_f}{m_{A'}}\right)^2 m_{h_D} \\ &\sim \left(\frac{1}{100 \text{ m}}\right) \left(\frac{\alpha_D}{0.1}\right) \left(\frac{\varepsilon}{3 \times 10^{-3}}\right)^4 \left(\frac{15 \text{ GeV}}{m_{A'}}\right).\end{aligned}\quad (238)$$

Additionally, radiative corrections induce a mixing between the SM and dark Higgs even in the absence of a tree-level coupling. The result is UV-sensitive, and therefore the lifetime can take on a wide range of values from the mm-scale to $\gtrsim 100$ m depending on the couplings. Because of this dependence on otherwise unobservable UV model parameters, we adopt an approach where we take the dark Higgs lifetime, $c\tau_{h_D}$, to be a free parameter of the theory, while the production of the dark Higgs is governed by the kinetic mixing, ε , and the dark gauge coupling, α_D .

In Figs. 63 and 64, we project the sensitivity of the proposed MATHUSLA experiment to the dark Higgs-strahlung production of h_D , requiring four events in MATHUSLA with $\sqrt{s} = 14$ TeV and $\mathcal{L} = 3000 \text{ fb}^{-1}$ of integrated luminosity. In Fig. 63, we show the sensitivity to $m_{A'}$ and $c\tau$ for fixed m_{h_D} and dark photon couplings, while in Fig. 64, we show the sensitivity to $m_{A'}$, m_{h_D} and ε for fixed $c\tau$ and α_D . For comparison, we also show the sensitivity of the main ATLAS or CMS detectors to the long-lived h_D scenario, using the selection criteria from Ref. [627]. For the ATLAS/CMS sensitivity, we require that the A' produced directly in the $Z \rightarrow A'h_D$ process decays promptly to leptons, and that these leptons pass standard dilepton triggers of two opposite-sign, same-flavour (OSSF) muons with $p_T > 17, 8$ GeV, or two OSSF electrons with $p_T > 23, 12$ GeV, respectively [632]. We demand the impact parameters of tracks coming from the h_D decay to satisfy $1 \text{ mm} < |d_0| < 200 \text{ mm}$ and that the physical decay occur within 200 mm of the primary vertex. We additionally assign a 50% reconstruction penalty for the h_D displaced vertex; backgrounds are expected to be negligible due to the resonant reconstruction of the A' mass in the two prompt leptons. It is evident that, for long-lived h_D with $c\tau \sim 100$ m, MATHUSLA substantially outperforms ATLAS or CMS; the main LHC detectors have excellent, complementary coverage for lower lifetimes. Thus, MATHUSLA has unique sensitivity to parameters that are well-motivated by the Hidden Abelian Higgs Model and are otherwise unconstrained by existing experiments.

8.6 Axion-Like Particles⁶⁸

Axion like particles (ALPs) is a collective name for pseudo Nambu-Goldstone bosons with unspecified derivative couplings to Standard Model (SM) particles. The name is inspired by the axion which is the pseudo Nambu-Goldstone boson of the Peccei-Quinn symmetry [196–198, 633] introduced to solve the strong CP-problem, but an ALP appears in any theory with a spontaneously broken global symmetry. For some large breaking scale f , the ALP can be the harbinger of an ultraviolet sector of physics with masses $M_{UV} \propto f$ that is otherwise inaccessible by current and future collider experiments. Since ALP couplings instead scale as $1/f$, they can be long-lived if the New Physics is heavy, making them prime candidates for experiments probing displaced vertices. Measuring the ALP couplings to SM particles can therefore reveal non-trivial information about a whole New Physics sector. In addition, ALPs can be non-thermal candidates for Dark Matter [634]. In order for the decays of the ALP Dark Matter not to disturb cosmology, the ALP has to decay before Big Bang Nucleosynthesis [635] (see also Section 8.3). This means that the lifetime of the ALP must be $c\tau_a \lesssim 10^8$ m, providing additional motivation for

⁶⁸Martin Bauer, Matthias Neubert, Anson Hook, Andrea Thamm




Fig. 65: Production cross-section of ALPs in the decays of heavy SM particles.

displaced vertex searches.⁶⁹

Up to operators of dimension five, the couplings between the ALP and SM particles are given by the operators

$$\mathcal{L}_5 = c_G \frac{g_s^2}{16\pi^2} \frac{a}{f} G_{\mu\nu}^A \tilde{G}^{\mu\nu, A} + c_W \frac{g^2}{16\pi^2} \frac{a}{f} W_{\mu\nu}^A \tilde{W}^{\mu\nu, A} + c_B \frac{g'^2}{16\pi^2} \frac{a}{f} B_{\mu\nu} \tilde{B}^{\mu\nu} + \frac{\partial_\mu a}{f} \sum_i \frac{c_i}{2} \bar{\psi}_i \gamma_\mu \gamma_5 \psi_i, \quad (239)$$

where $c_{\gamma\gamma} = c_W + c_B$ and $c_{\gamma Z} = \cos^2 \theta_w c_W - \sin^2 \theta_w c_B$ and $c_{ZZ} = \cos^4 \theta_w c_W + \sin^4 \theta_w c_B$ are the relevant Wilson coefficients in the electroweak broken phase, and the couplings to fermions c_i are assumed to be flavour universal. Here, f sets the scale of the UV completion and is related to the ALP decay constant by $f = -2c_G f_a$. Operators that introduce couplings between the ALP and the Higgs boson H only arise at dimension six and higher,

$$\mathcal{L}_{>5} = \frac{c_{ah}}{f^2} (\partial_\mu a) (\partial^\mu a) H^\dagger H + \frac{c_{Zh}}{f^3} (\partial^\mu a) \left(H^\dagger i D_\mu H + \text{h.c.} \right) H^\dagger H + \dots, \quad (240)$$

where the Higgs portal allows for $h \rightarrow aa$ decays, whereas the coupling to the Higgs current introduces the decay $h \rightarrow Za$. A possible dimension five operator coupling the ALP to the Higgs current is redundant unless it is induced by integrating out new massive particles that obtain most of their mass from the electroweak scale [636–639]. An ALP mass can be generated through some external breaking of the corresponding symmetry, or can be dynamically introduced through its couplings to the QCD condensate. In the latter case, the ALP mass is directly related to the decay constant $m_a \propto f_\pi m_\pi / f_a$ with f_π and m_π the pion decay constant and the pion mass, respectively. In the more general case there is no such relation and m_a and f are independent parameters.

⁶⁹Due to their light masses, ALPs are generically displaced from their minimum during inflation. After reheating, their energy density behaves like dark energy until Hubble is of order their mass. Afterwards, they dilute away as normal matter. Because their energy density does not dilute away like matter until very late, they generically overclose the universe unless they decay or $m_a f_a^4 \lesssim (10^7 \text{ GeV})^5$ where we have made the optimistic assumption that the axions start oscillating as soon as they can. This assumption is not true for some ALPs, e.g. the QCD axion, where the mass term is not present at early times. See also discussions in Sections 4.1.5 and 5.3

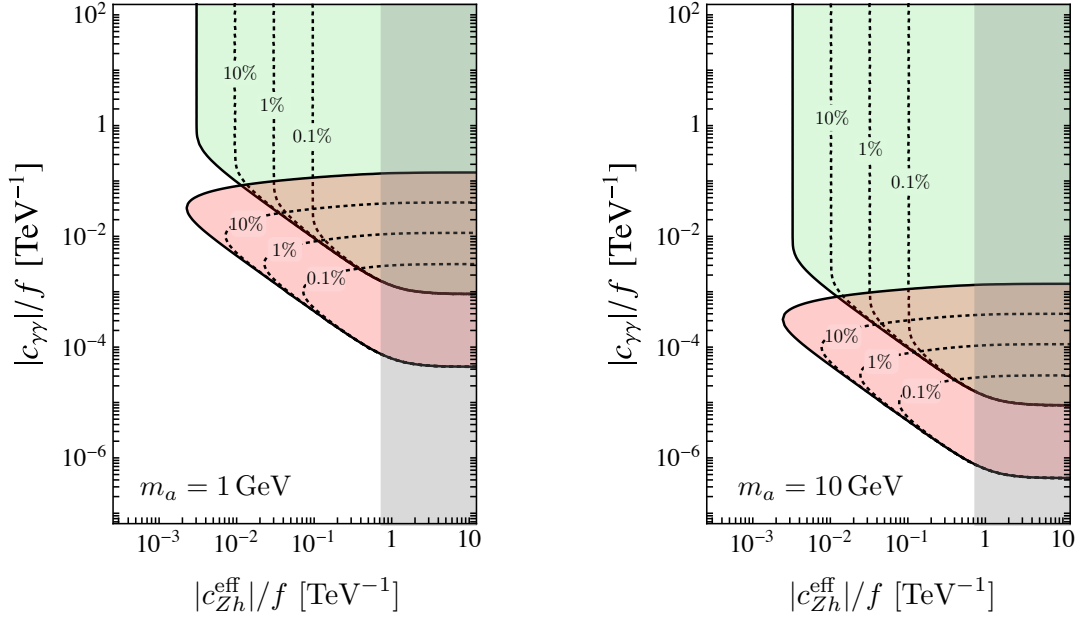


Fig. 66: Projected reach in searches for $h \rightarrow Za \rightarrow \ell^+ \ell^- + 2\gamma$ decays with ATLAS/CMS (green) and MATHUSLA (red) with $\sqrt{s} = 14$ TeV center-of-mass energy and 3000 fb^{-1} integrated luminosity. The parameter region with the solid contours correspond to a branching ratio of $\text{Br}(a \rightarrow \gamma\gamma) = 1$, and the contours showing the reach for smaller branching ratios are dashed. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

ALPs at the LHC suffer from a small production cross-section $\sigma(pp \rightarrow a)$ - if f is large - or decay promptly - if f is small. Beyond resonant production, ALPs can be produced in decays of heavy SM particles. In this case, larger scales f correspond to smaller branching ratios and delayed ALP decays. In the following, we will discuss the reach of the MATHUSLA detector for ALPs produced in the decays $Z \rightarrow a\gamma$, $h \rightarrow aZ$ and $h \rightarrow aa$. In Fig. 65, we show the corresponding production cross-sections in dependence of the breaking scale f for a mass $m_a = 0$, using the relevant branching ratios

$$\Gamma(h \rightarrow Za) = \frac{m_h^3}{16\pi f^2} |c_{Zh}^{\text{eff}}|^2 \lambda^{3/2} \left(\frac{m_Z^2}{m_h^2}, \frac{m_a^2}{m_h^2} \right), \quad (241)$$

$$\Gamma(h \rightarrow aa) = \frac{m_h^3 v^2}{32\pi f^4} |c_{ah}|^2 \left(1 - \frac{2m_a^2}{m_h^2} \right)^2 \sqrt{1 - \frac{4m_a^2}{m_h^2}}, \quad (242)$$

$$\Gamma(Z \rightarrow a\gamma) = \frac{\alpha \alpha(m_Z) m_Z^3}{96\pi^3 \sin^2 \theta_W \cos^2 \theta_W f^2} |c_{Z\gamma}|^2 \left(1 - \frac{m_a^2}{m_Z^2} \right)^3, \quad (243)$$

where we define $c_{Zh}^{\text{eff}} = c_{Zh}^5 + 2c_{Zh} v^2/f^2$ in order to take into account possible contributions from chiral new physics (that arise for example by integrating out the top quark). The cross-sections clearly show the different scaling for the dimension five, six, and seven operators. The shaded region is excluded by Higgs coupling measurements constraining general beyond the SM decays of the Higgs $\text{Br}(h \rightarrow \text{BSM}) < 0.34$ [640] and the error on the measurement of the total Z width, which corresponds to $\text{Br}(Z \rightarrow \text{BSM}) < 0.0018$ [641].

In the following discussion, in order to evaluate the reach of ATLAS, CMS and the MATHUSLA detector, we consider ALP decays into photons, leptons and gluons as exemplary final states, but other final states are equally interesting, if ALPs are heavy enough. Depending on their mass, ALPs from

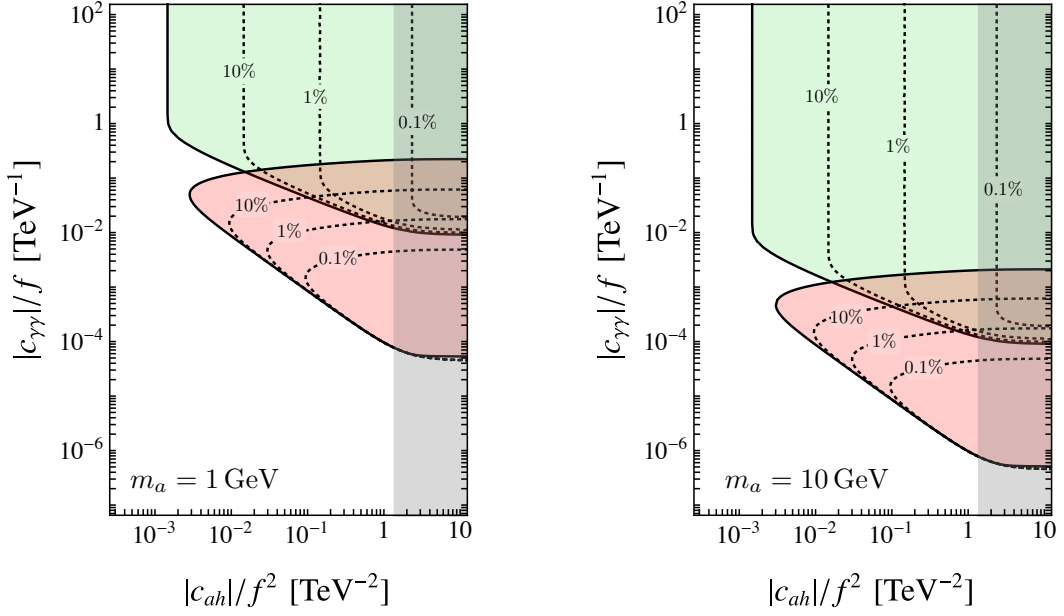


Fig. 67: Projected reach in searches for $h \rightarrow aa \rightarrow 4\gamma$ decays with ATLAS/CMS (green) and MATHUSLA (red) with $\sqrt{s} = 14$ TeV center-of-mass energy and 3000 fb^{-1} integrated luminosity. The parameter region with the solid contours correspond to a branching ratio of $\text{Br}(a \rightarrow \gamma\gamma) = 1$, and the contours showing the reach for smaller branching ratios are dashed. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

Higgs or Z decays can be highly boosted with the usual relativistic factor

$$\gamma_a = \begin{cases} \frac{m_h^2 - m_Z^2 + m_a^2}{2m_a m_h}, & \text{for } h \rightarrow Za, \\ \frac{m_h}{2m_a}, & \text{for } h \rightarrow aa, \\ \frac{m_a^2 + m_Z^2}{2m_Z m_a}, & \text{for } Z \rightarrow a\gamma. \end{cases} \quad (244)$$

For searches with ATLAS or CMS, we demand that all final state particles are detected in order to reconstruct the decaying SM particle and that the decays into photons occur before the electromagnetic calorimeter, $R = 1.5$ m, and the decays into leptons before the inner tracker $R = 2$ cm. For example, for $h \rightarrow Za \rightarrow \ell^+ \ell^- \gamma\gamma$ decays, we ask for the Z to be reconstructed in dileptons *and* the ALP to decay inside the detector. We therefore define the effective branching ratios

$$\text{Br}(h \rightarrow Za \rightarrow \ell^+ \ell^- + \gamma\gamma)|_{\text{eff}} = \text{Br}(h \rightarrow Za) \text{Br}(a \rightarrow \gamma\gamma) f_{\text{dec}}^a \text{Br}(Z \rightarrow \ell^+ \ell^-), \quad (245)$$

$$\text{Br}(h \rightarrow aa \rightarrow \gamma\gamma + \gamma\gamma)|_{\text{eff}} = \text{Br}(h \rightarrow aa) \text{Br}(a \rightarrow \gamma\gamma)^2 f_{\text{dec}}^{aa}, \quad (246)$$

$$\text{Br}(Z \rightarrow a\gamma \rightarrow \gamma\gamma\gamma)|_{\text{eff}} = \text{Br}(Z \rightarrow a\gamma) \text{Br}(a \rightarrow \gamma\gamma) f_{\text{dec}}^a, \quad (247)$$

for the different processes considered. $f_{\text{dec}}^a = f_{\text{dec}}^a(\gamma_a)$ is the fraction of axions decaying in the main detector, which is approximated as an infinitely long cylinder with the above mentioned inner and outer radii. Analogous expressions hold for the ALP decaying into leptons and gluons. We further do not distinguish displaced from prompt decays and derive the reach for a number of 100 signal events, which is typically needed to suppress backgrounds in searches for New Physics with prompt decays of h and Z bosons [640, 642, 643]. For MATHUSLA, it is impossible to detect both final state particles in $h \rightarrow Za$ and $Z \rightarrow a\gamma$ decays and highly unlikely to see both ALPs from $h \rightarrow aa$ decays in the decay volume.

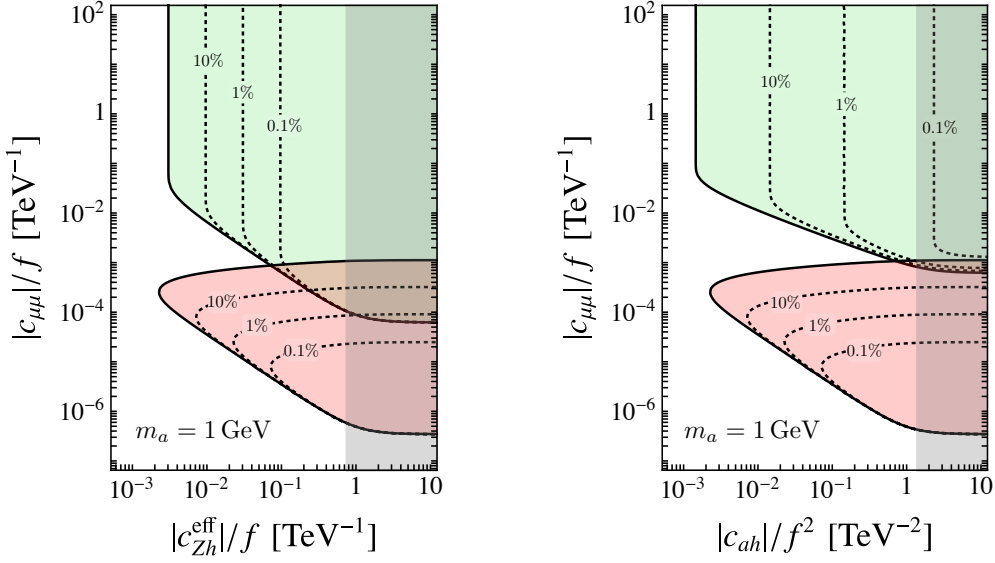


Fig. 68: Projected reach in searches for $h \rightarrow Za \rightarrow \ell^+\ell^- + \mu^+\mu^-$ (left) and $h \rightarrow aa \rightarrow \mu^+\mu^- + \mu^+\mu^-$ (right) decays with ATLAS/CMS (green) and MATHUSLA (red) with $\sqrt{s} = 14$ TeV center-of-mass energy and 3000 fb^{-1} integrated luminosity. The parameter region with the solid contours correspond to a branching ratio of $\text{Br}(a \rightarrow \mu^+\mu^-) = 1$, and the contours showing the reach for smaller branching ratios are dashed. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

However, because of the much lower background, single ALPs can be detected irrespective of their origin. We ask for at least four ALP decays within the MATHUSLA volume to derive the reach of the detector, so that the corresponding effective branching ratios for ALP decays in MATHUSLA read

$$\text{Br}(h \rightarrow Za \rightarrow Z + \gamma\gamma) \Big|_{\text{eff}}^{\text{M}} = \text{Br}(h \rightarrow Za) \text{Br}(a \rightarrow \gamma\gamma) f_{\text{M}}^a, \quad (248)$$

$$\text{Br}(h \rightarrow aa \rightarrow a + \gamma\gamma) \Big|_{\text{eff}}^{\text{M}} = 2\text{Br}(h \rightarrow aa) \text{Br}(a \rightarrow \gamma\gamma) f_{\text{M}}^a, \quad (249)$$

$$\text{Br}(Z \rightarrow a\gamma \rightarrow \gamma + \gamma\gamma) \Big|_{\text{eff}}^{\text{M}} = \text{Br}(Z \rightarrow a\gamma) \text{Br}(a \rightarrow \gamma\gamma) f_{\text{M}}^a. \quad (250)$$

Note that states to the left of the “+” on the LHS are visible not in MATHUSLA but in the main detector. $f_{\text{M}}^a = f_{\text{M}}^a(\gamma_a)$ is the fraction of LLPs which decay in the MATHUSLA decay volume. Again, the expressions for ALP decays into leptons and gluons are analogous to Eqns. (248)-(250). Multiplying the effective branching ratios Eqns. (245)-(247) and (248)-(250) with the production cross-sections at the LHC for a center of mass energy of 14 TeV, $\sigma_{14 \text{ TeV}}(pp \rightarrow h) = 54.61 \text{ pb}$ [582] or $\sigma_{14 \text{ TeV}}(pp \rightarrow Z) = 60.59 \text{ nb}$ (derived with MATRIX [644]) allows us to estimate the reach of the two search strategies.

We illustrate the reach for the ATLAS or CMS detector for discovering ALPs decaying into photons from $h \rightarrow aZ$ and $h \rightarrow aa$ decays in Fig 66 and Fig. 67, respectively. For the green parameter space with solid contours, ATLAS or CMS would see 100 events with a luminosity of $\mathcal{L} = 3000 \text{ fb}^{-1}$ and a branching ratio of $\text{Br}(a \rightarrow \gamma\gamma) = 1$. For smaller branching ratios, larger couplings $|c_{hZ}^{\text{eff}}|$ and $|c_{ah}|$ are required to obtain the same number of events. Dashed lines show the lower limit for $\text{Br}(a \rightarrow \gamma\gamma) = 0.1$, $\text{Br}(a \rightarrow \gamma\gamma) = 0.01$ and $\text{Br}(a \rightarrow \gamma\gamma) = 0.001$.⁷⁰ The red region with solid contours shows the parameter space for which four ALP decays are expected within the MATHUSLA detector for $\text{Br}(a \rightarrow \gamma\gamma) = 1$ and $\mathcal{L} = 3000 \text{ fb}^{-1}$. For $\text{Br}(a \rightarrow \gamma\gamma) = 0.1$, $\text{Br}(a \rightarrow \gamma\gamma) = 0.01$ and $\text{Br}(a \rightarrow \gamma\gamma) = 0.001$,

⁷⁰A smaller branching ratio for the given coefficients implies a larger total LLP width and hence shorter lifetime. In the long lifetime limit, the increased signal rate due to shorter lifetime offsets the lower branching ratio, making the lower boundaries of sensitivity independent of the branching ratio.




Fig. 69: Projected reach in searches for $h \rightarrow Za \rightarrow \ell^+\ell^- + jj$ (left) and $h \rightarrow aa \rightarrow 4j$ (right) decays with MATHUSLA (red) with $\sqrt{s} = 14$ TeV center-of-mass energy and 3000 fb^{-1} integrated luminosity. The parameter region with the solid contours correspond to a branching ratio of $\text{Br}(a \rightarrow \mu^+\mu^-) = 1$, and the contours showing the reach for smaller branching ratios are dashed. This estimate assumes the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1.

MATHUSLA therefore loses sensitivity for larger values of $|c_{\gamma\gamma}|/f$. In the case of $h \rightarrow aa$ decays, MATHUSLA will be able to probe smaller branching ratios than ATLAS or CMS. This underlines the complementarity between searches for prompt decays with ATLAS and CMS and searches for displaced ALP decays with MATHUSLA. In the event that MATHUSLA finds an LLP signal, event-by-event information on the LLP boost (and to some extent the final state, depending on the final detector capabilities) can be correlated with information prompt displaced object information from the main detector to elucidate the LLP production mode and eventually identify the LLP as an ALP.

In Fig. 68, we further show the reach for ALP decays into muons. Since at least approximate lepton flavour universality is expected in the couplings of the ALP, the muon decay mode is particularly well-motivated for $2m_\mu < m_a < 2m_\tau$.

Finally, we present the reach of MATHUSLA for ALP decays into jets in Fig. 69. In contrast to ALP decays into photons and leptons, we refrain from showing projections for LHC searches for $a \rightarrow jj$ decays, given the large backgrounds for the Higgs decays $h \rightarrow Za \rightarrow \ell^+\ell^-jj$ and $h \rightarrow aa \rightarrow 4j$. It has been shown in [1] that MATHUSLA has $1000\times$ better sensitivity to LLP production cross-sections than ATLAS or CMS for this channel. If decay to jets is the dominant axion decay mode, MATHUSLA provides by far the strongest sensitivity for the c_G coefficient. (Larger c_G values may be probed at the main detectors.)

For $Z \rightarrow a\gamma$ decays with subsequent ALP decays into photons, the relevant Wilson coefficients $c_{\gamma\gamma}$ and $c_{Z\gamma}$ are linear combinations of c_W and c_B . A scenario in which these coefficients are completely independent therefore appears to be fine-tuned.⁷¹ We therefore show the reach of ATLAS or CMS and MATHUSLA in the $m_a - |c_{\gamma\gamma}|/f$ plane under the assumption that $c_W = 0$ and $c_{Z\gamma} = -\sin^2\theta_w c_{\gamma\gamma}$, but impose a hard cut on $c_{Z\gamma}$ in the parameter space for which the constraint from the total Z width, $\text{Br}(Z \rightarrow \text{BSM}) < 0.0018$, is violated [646]. The corresponding exclusion region is shown in Fig. 70

⁷¹Integrating out a single electroweak multiplet for example always generate c_W and c_B with the same sign, resulting in $|c_{Z\gamma}| \lesssim \cos^2\theta_w |c_{\gamma\gamma}|$.

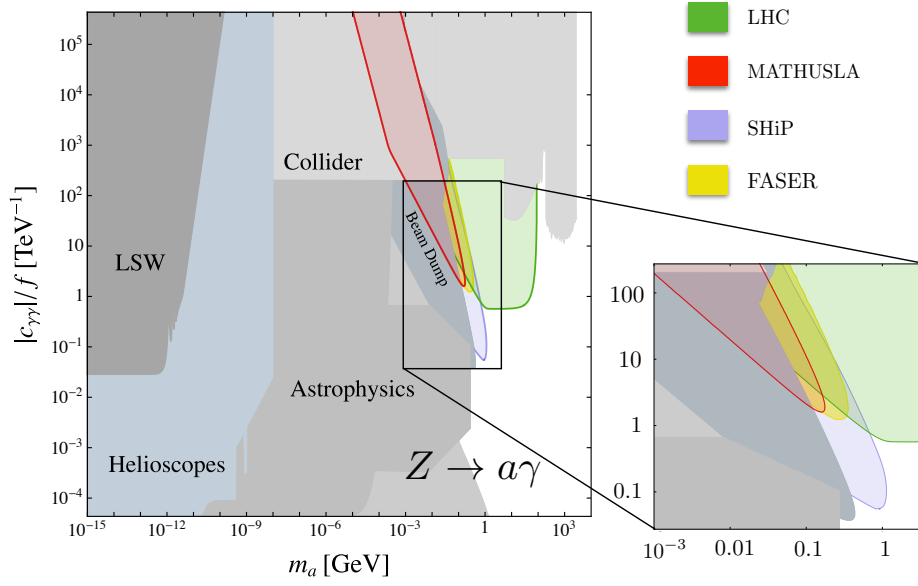


Fig. 70: Projected reach in searches for $Z \rightarrow a\gamma \rightarrow 3\gamma$ decays with ATLAS/CMS (green) and MATHUSLA (red, assuming the $200m \times 200m \times 20m$ benchmark geometry of Fig. 1) with $\sqrt{s} = 14$ TeV center-of-mass energy and 3000 fb^{-1} integrated luminosity. The area projected to be probed by the future experiments SHiP [219] and FASER [645] are shown in blue and yellow, respectively.

together with various other constraints which only depend on $c_{\gamma\gamma}$. The kink in the exclusion region at $|c_{\gamma\gamma}|/f \approx 1$ GeV and the resulting gap between the ATLAS/CMS and MATHUSLA reach occur because the values of $c_{Z\gamma}$ corresponding to smaller values of $c_{\gamma\gamma}$ would not yield enough events for small ALP masses. The reach of the MATHUSLA experiment overlaps with the existing limits from the E137 and E141 beam dump experiments (shaded light brown in Fig. 70) [599,600] and competes with the projected limits from the future SHiP experiment (shaded blue in Fig. 70) [219] and the future FASER experiment [645] (shaded yellow in Fig. 70), though these have higher sensitivity. It is also worth noting that the limits and sensitivity projections in Fig. 70 assume that all other coefficients in Eqn. (239) are zero. If any other coefficients are present, MATHUSLA could be sensitive to values of these coefficients (see Figs. 66 - 69) that are orders of magnitude smaller than the $c_{\gamma\gamma}/f$ range shown in Fig. 70.

In conclusion, Axion-like particles are a very general and well-motivated BSM scenario that can be probed by MATHUSLA at lifetimes much larger (and couplings much smaller) than possible with the LHC main detectors alone.

9 Executive Summary

This document has two main aims: to demonstrate that (1) neutral LLPs are broadly and fundamentally motivated in BSM theories, and (2) the construction of the MATHUSLA detector is necessary to fully leverage the LHC’s vast discovery potential for new physics. In this section we summarize how the results presented here fulfill these objectives.

We have discussed the most bottom-up motivations for LLPs in the Introduction. Many particles in the SM have much longer lifetime than a naïve expectation from dimensional analysis might suggest, in many cases out to macroscopic distances. A variety of mechanisms can suppress the decay width of an unstable particle: small couplings, heavy mediators, approximate symmetries, and/or phase space suppressions. These exact mechanisms can give rise to neutral LLPs in any BSM theory, including simple extensions of the SM such as hidden valleys (Section 8.1), which are a generic consequence of the structure of gauge theories, and minimal benchmark models like dark scalars (Section 8.4), dark photons (Section 8.5), and Axion-like particles (Section 8.6). These hypothesized new physics sectors are separated from the SM fields not (necessarily) by a mass hierarchy, but by an absence of large couplings between the hidden and visible sectors. The very nature of a (possibly confining) hidden sector, only connected to the SM via a tiny portal at low energies, makes neutral LLPs an obvious signal to search for. Regardless of the details of the new physics, exotic Higgs decays (Section 8.2) are one of the most motivated production modes for light new states including LLPs, due to the small SM Higgs width, its large LHC production rate, and the lack of symmetry protection for the $|H|^2$ operator allowing it to couple to any new physics. Furthermore, MATHUSLA would be sensitive to LLPs produced in exotic Higgs decays out to lifetimes near the upper limit from BBN (Section 8.3), provided they are not too boosted. Searches for these simplified IR scenarios are motivated in their own right to agnostically cover as much possible new physics theory- and parameter-space as we experimentally can, especially in light of recent LHC null results.

The bottom-up *plausibility* of LLP signatures is therefore well-established. However, one of the most important conclusions we can draw from the results presented in this document is that LLPs are not just plausible, but *strongly fundamentally motivated for a broad variety of top-down reasons*. They are ubiquitous in BSM scenarios that address longstanding mysteries like the naturalness of the weak scale (Section 4), Dark Matter (Section 5), Baryogenesis (Section 6) and Neutrino Masses (Section 7). Furthermore, they are often an intrinsic part of the theory mechanism which addresses the fundamental mystery in the first place.

One way to map out these top-down motivations is sketched in Fig. 71. This figure qualitatively illustrates which theories and frameworks discussed in the preceding sections give rise to which “IR LLP Scenarios”, broadly defined to include general classes of bottom-up theories like hidden valleys, simplified models like SM+S or SM+V, and specific LLP signatures like exotic Higgs decays. A very common example of a neutral LLP signature is simply the direct production at the LHC of a BSM state with sufficiently sizable couplings to the SM, which either is an LLP itself, or decays promptly to an LLP (“BSM= \rightarrow LLP”). The theories examined in this document are hardly exhaustive, but the ubiquity of LLPs in top-down motivated BSM theories is evident.

One might wonder why such a coarse-grained classification of signatures is even helpful. After all, the “BSM= \rightarrow LLP” scenario includes a wide variety of different LLP species with different production and decay modes.⁷² Essentially, this is because almost any theory with LLPs can give rise to very long lifetimes, either because the long-lifetime regime is specifically motivated, or because the lifetime is practically a free parameter.⁷³ As we review below, in this long-lifetime regime the discussion of LLP

⁷²This is to be contrasted with simplified models for LLP searches at the main detectors developed by the LHC-LLP Community working groups [20], which have to parameterize the large variety of displaced and associated prompt signals in considerable kinematic detail, in order to facilitate the development of concrete search and background rejection strategies.

⁷³There are a few exceptions which prove the rule, e.g. pure higgsinos with a tiny mass splitting from electroweak symmetry breaking [130, 159, 304, 305, 647], with lifetimes below a cm. We did not study these examples here, but they of course add

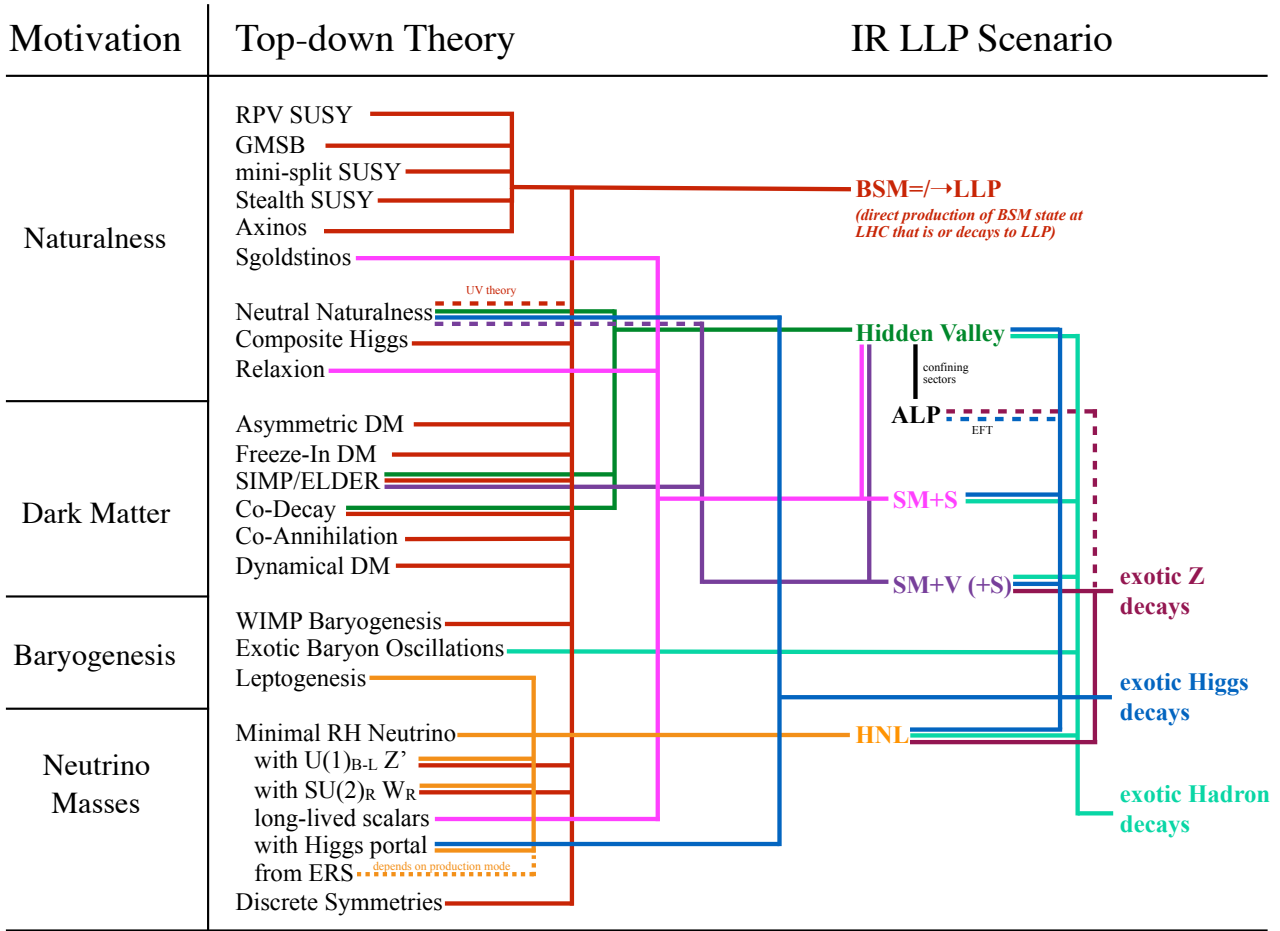


Fig. 71: Qualitative overview of the top-down theory motivations for neutral LLPs discussed in this document, with colored lines (from left to right) indicating which IR LLP scenario they motivate at the LHC. Some of the IR LLP scenarios or simplified models in turn motivate specific signatures like exotic Higgs decays. We stress that these top-down theories are not the only motivations for the IR scenarios or simplified models shown here: hidden valleys, exotic Higgs decays, etc., are also motivated in their own right on generic, bottom-up grounds.

signatures at MATHUSLA, and comparing sensitivity to the LHC main detectors, becomes quite simple, and leads to the conclusion that MATHUSLA has highly general and robust advantages when searching for LLPs.

The basic MATHUSLA detector concept is described in Section 2. The benchmark design is an empty box on the surface with trackers in the roof and active vetoes surrounding the $200m \times 200m \times 20m$ air-filled detector volume. Neutral LLP decays into two or more charged particles are reconstructed as displaced vertices with stringent geometric and timing requirements. MATHUSLA's position on the surface provides shielding from the deluge of SM particles produced at the collision point. The high-energy displaced signature of LLP decays is therefore even more distinctive in MATHUSLA than inside the LHC main detectors. The most important remaining backgrounds on the surface are cosmic rays, high-energy muons from the LHC, and neutrino scatterings. All of these can be rejected with extremely high fidelity, using simple requirements on the charged particle direction of travel as well as more elaborate geometrical and timing cuts. As a result, MATHUSLA can search for LLPs in effectively the background-free regime.

motivation for LLP searches at the LHC main detectors, and slight modifications of the model can yield longer lifetimes.




Fig. 72: (Identical to Fig. 7.) Schematic order-of-magnitude sensitivity of MATHUSLA, assuming $\mathcal{O}(1)$ produced LLPs per production event at the HL-LHC. \bar{b} is the mean boost of the produced LLPs. The shape of the exclusion/discovery region at short lifetimes depends on the detailed boost distribution, but for long lifetimes $\bar{b}c\tau \gg 200m$ depends only on the mean boost and is very model-independent up to an $\mathcal{O}(1)$ factor. Note that LLPs near the BBN lifetime limit of $c\tau \sim 10^7m$ can be probed if they are produced with cross-sections in the pb range at the HL-LHC. To emphasize the scalability of the MATHUSLA design, we also show the reach achievable with a version of MATHUSLA with only $1/10$ the detector volume of the $200m \times 200m \times 20m$ benchmark geometry.

In Section 3, we take a model-independent approach to assess the sensitivity of MATHUSLA to neutral LLP production rates, and compare its sensitivity to main detector LLP searches. In the long-lifetime regime $bc\tau \gg 100m$, MATHUSLA has comparable acceptance for LLP decays as ATLAS or CMS, with only very modest dependence on the production mode. However, unlike the underground detectors, which have to contend with a variety of backgrounds when searching for neutral LLPs, MATHUSLA can operate without backgrounds. This allows for the detection of neutral LLPs with lifetimes near the BBN limit of $c\tau \sim 10^7m$ (for order one boosts) if they are produced with \sim pb cross-section. Decay lengths of $\sim 100m$ can be detected for \sim fb cross-section at the LHC. This model-independent sensitivity is shown in Fig. 7, which we reproduce in this Section as Fig. 72 for convenience. For the purposes of neutral LLP discovery, MATHUSLA can therefore be thought of as a version of the main detectors that *sacrifices* sensitivity to shorter decay lengths in order to *gain* the ability to search for LLPs without backgrounds or trigger limitations.

With the motivation for neutral LLP searches established, we must therefore ask: (1) how important is MATHUSLA's advantage of zero background and no trigger issues compared to the main detectors, and (2) how motivated is the long-lifetime regime (so that there is any signal at MATHUSLA)?

The general discussion of the first issue is provided in Section 3.2 and can be summarized with a few simple qualitative conclusions. Missing energy triggers at ATLAS/CMS are generally inefficient for neutral LLP searches unless the production rates are sizable and sufficiently energetic prompt objects are also present. Since MATHUSLA has a similar acceptance for LLP decays as the main detectors, but operates essentially free from backgrounds or trigger limitations, it will have superior sensitivity for any neutral LLP signal where either backgrounds, cut efficiency (including requirements on LLP decay and production mode) or triggers impede the main detector search. This includes:

- LLPs that decay with less than a few hundred GeV of visible hadronic energy. For example,

MATHUSLA has 3 orders of magnitude better cross-section sensitivity to LLP that are produced in exotic Higgs decays and decay via the Higgs portal.

- LLPs that have subdominant leptonic branching fractions, for masses below a few hundred GeV.
- LLPs lighter than ~ 10 GeV that decay to lepton jets, where MATHUSLA may increase reach by 1-2 orders of magnitude in cross-section.

Another interesting scenario with potentially large sensitivity gains are LLPs that decay to photons, if MATHUSLA is configured for their detection and can search with low backgrounds. Conversely, if an LLP is always produced in association with a hard lepton or decays into jets with more than several 100s of GeV of energy, the event can pass L1 triggers and the main detector LLP search has very few backgrounds, resulting in likely similar sensitivity to MATHUSLA in the long-lifetime regime. These simple arguments illustrate why MATHUSLA has far superior sensitivity to the main detectors for large classes of important neutral LLP signals.

We now turn to the motivation for LLPs in the long-lifetime regime. Most of the theories discussed in this document feature LLP signals for which MATHUSLA could be our only discovery opportunity in large parts of parameter space. To more explicitly demonstrate the role of MATHUSLA in probing fundamentally motivated BSM theories, we provide Tables 4 and 5.

Table 4 summarizes those BSM scenarios where the discoverable LLP is a strongly motivated intrinsic part of the theory mechanism. We attempt to summarize the main role that LLPs play in each theory, the motivations for the long-lifetime regime if any, and the role MATHUSLA would play in their discovery. Examples include Neutral Naturalness (Section 4.2), where the very symmetry protection which stabilizes the weak scale gives rise to a hidden valley containing LLPs accessible via the Higgs portal; WIMP Baryogenesis (Section 6.1), where the LLP decay at long lifetimes is the very mechanism which generates the baryon asymmetry of the universe; and FIMP DM (Section 5.3), which can be produced both in the early universe and at the HL-LHC in the decay of a parent LLP with sizable SM couplings. In all these cases, not only would MATHUSLA have the greatest sensitivity to discover many classes of these BSM scenarios, there are often arguments why the long lifetime regime might even be theoretically preferred. This makes the search for very long-lived particles even more urgent.

For other theories, summarized in Tables 5, LLPs are no less motivated, but their existence and/or long lifetime is simply one part of a much larger space of possible signals, depending on the specific model details and parameters. This includes the ubiquitous hidden valley idea, the general new physics discovery channel of exotic Higgs decays, as well as broad classes of dark matter models.

MATHUSLA is also important for investigating other possible theories of new physics. Split versions of composite Higgs (Section 4.3) and supersymmetry (Section 4.1.3) generically give rise to long-lived colored particles which may be discovered at the main detectors, but MATHUSLA would provide an additional discovery channel as well as important information about the behavior of such R-hadrons in matter. Extended versions of these theories, which avoid some of the constraints suffered by the minimal models, can also give rise to neutral LLPs for which MATHUSLA is the prime discovery tool. In many cases, new physics might be discovered at the main detectors but be mis-diagnosed. This is generally true if a MET search discovers what is actually a very long-lived particle, but can also be true for resonance searches where a discovery of new physics obscures the existence of a hidden sector containing light LLPs that are important for diagnosing the complete theory, such as might be the case in versions of Stealth SUSY (Section 4.1.4). It is also possible for details of the UV theory to generate MATHUSLA signals in scenarios where we do not naively expect them from low-energy considerations, such as for neutrino models with Enhanced Residual Symmetry (ERS, Section 7.6). Finally, MATHUSLA has impressive capabilities as a cosmic ray telescope. This is briefly discussed in Section 2.5 and will be the subject of its own dedicated study. MATHUSLA's measurements could address many outstanding puzzles in cosmic-ray and astro-particle physics, and represent a guaranteed physics return on the investment of constructing the detector.

BSM Scenario	Role of LLPs	Typical $c\tau$	Role of MATHUSLA	Sec.	Fig.
Neutral Naturalness	Discrete symmetry stabilizing Higgs mass \rightarrow hidden valley with Higgs portal. Cosmology \rightarrow hidden valley particles are LLPs.	Any, but \mathbb{Z}_2 arguments favor lower $\hat{\Lambda}_{QCD}$ and hence long lifetimes.	Smoking gun signal are mirror glueball LLPs. For long lifetimes, they can only be discovered at MATHUSLA.	4.2	22, 23
WIMP Baryogenesis	Out-of-equilibrium decay of WIMP-like LLP produces baryon asymmetry.	\gtrsim cm for weak-scale LLP masses.	Decays to baryons \rightarrow MATHUSLA likely much greater sensitivity than main detectors. MCFODO	6.1	34
FIMP DM	Freeze-in via decay requires LLPs with SM couplings.	Fixed by masses & cosmology. Long lifetimes generic.	Model-dependent, but in long-lifetime regime MCFODO.	5.3	27, 28, 21,
Co-decaying DM	Out-of-equilibrium decay of hidden sector LLP determines DM abundance. Also, small portal \rightarrow visible sector LLPs.	For weak scale LLP masses, most of parameter space is long lifetimes.	Depending on model details (production & decay mode), MCFODO.	5.4.3	31
Co-annihilating DM	DM relic abundance relies on small mass splitting with another state \rightarrow other state is LLP.	Any, long lifetimes generic.	Depends on model details, but e.g. for Higgs Portal implementations, MCFODO.	5.1	
SUSY: Axinos	High PQ-breaking scale V_{PQ} suppresses axion/axino couplings, making LOSP an LLP	Any, long lifetimes generic.	For high V_{PQ} , MCFODO.	4.1.5	21
SUSY: GMSB	Low SUSY breaking scale F (motivated by flavor problem) leads to light gravitino and small couplings to LOSP, which can hence be LLP.	Any, long lifetimes generic.	MCFODO, depending on spectrum and lifetime.	4.1.2	15
SUSY: RPV	small RPV couplings (motivated by avoiding flavor violation, proton decay, baryon washout) \rightarrow LOSP can be LLP	Any, long lifetimes generic.	MCFODO, especially for EW-charged LLPs or squeezed spectra.	4.1.1	14
SUSY: Sgoldstinos	SUSY breaking scale F suppresses sgoldstino coupling to supercurrents \rightarrow can be LLP.	Any. Long lifetimes \rightarrow smallest production, hardest to probe.	Similar to SM+S. For masses $\lesssim 5$ GeV, MATHUSLA and/or SHiP may be only/first discovery opportunity.	4.1.6	
Exotic Baryon Oscillations	Exotic Baryon is LLP and induces oscillations that generate baryon number.	$\gtrsim 100$ m	Heavy baryon decays produce LLP. MATHUSLA and/or SHiP may be only/first discovery opportunity.	6.2	
minimal RH neutrino model	Type-1 see-saw \rightarrow tiny mixing between ν_L and $\nu_R \rightarrow \nu_R$ LLPs	Any, long lifetimes favor lower m_N	In long-lifetime/low-mass regime, MATHUSLA and/or SHiP may be only/first discovery opportunity.	7.1	36, 37
\leftrightarrow with $U(1)_{B-L} Z'$	Weakly gauged $B-L$ breaking generates M_N , additional ν_R production mode from Z' .	$m_N \sim 1-10$ GeV suggests long lifetime regime.	For sub-weak-scale m_N , MCFODO.	7.2.1	38
\leftrightarrow with $SU(2)_L W_R$	ν_R part of gauged $SU(2)_R$, breaking generates M_N . Additional ν_R production mode from W_R^\pm .	Any, long lifetimes favor lower m_N .	For $m_{W_R} \sim 10$ TeV: main detector probes weak-scale m_N . MATHUSLA/SHiP only discovery opportunity for $m_N \lesssim 5$ GeV.	7.3.1	40
\leftrightarrow with Higgs Portal	GUT motivates extra broken $U(1)$ gauge groups, extended scalar sectors mix with Higgs \rightarrow produce ν_R in Higgs and other scalar decays.	Any, long lifetimes favor lower m_N .	MCFODO, improves Br reach of main detectors by at least order of magnitude.	7.4	43
m_ν via discrete symmetries	Discrete sym. generates m_ν and stabilizes FIMP DM.	See FIMP DM.	LLPs with EW charge \rightarrow MCFODO, especially for $m \lesssim 10$ GeV	7.5	

Table 4: BSM scenarios discussed in this document where neutral LLP signals at MATHUSLA are a strongly motivated intrinsic part of the theory mechanism, *and* MATHUSLA Could be First or Only Discovery Opportunity (MCFODO). When discussing lifetimes, “any” means up to the BBN limit, “long” means the MATHUSLA regime. LOSP = lightest observable-sector supersymmetric particle.

BSM Scenario	Role of LLPs	Typical $c\tau$	Role of MATHUSLA (long $c\tau$)	Sec.	Fig.
Hidden Valleys (HV)	Small portal to visible sector and possibly hidden sector confinement \rightarrow meta-stable states.	Any.	MCFODO, especially if LLPs are significantly below the weak scale or decay hadronically.	8.1	46, 47
SM+S	Small mixing \rightarrow scalar LLP, produce in exotic Higgs decays for $m_S < m_H/2$. Large mixing \rightarrow S could decay to HV LLPs.	Any.	MCFODO. Complementarity with SHiP.	8.4	55
SM+V	Dark photon/dark Higgs LLP could be produced in exotic Higgs/ Z decays. Dark photon with non-tiny kinetic mixing could be copiously produced at LHC and decay to HV LLPs.	Any.	MCFODO. Significantly extends main detector long-lifetime reach for dark photons and dark Higgs produced in exotic H and Z decays. For LLPs produced in dark photon decays, see HV.	8.5	59, 61, 63, 64
Exotic Higgs decays	Higgs coupling to new states, like HV or other LLPs, is highly generic and leads to large production rates at LHC.	Any.	MCFODO for $\text{Br} \lesssim 0.1 - 0.01$. Higgs portal motivates hadronic LLP decays, for which MATHUSLA has 10^3 better Br reach than main detectors. MATHUSLA also has significantly better sensitivity for LLP masses $\lesssim 10$ GeV even if they decay leptonically, or for LLPs with subdominant leptonic decays.	8.2	48, 49
Asymmetric DM	Relating DM to baryon abundance requires operator connecting DM number and Baryon/Lepton number \rightarrow higher dimensional operator \rightarrow LLPs	Any, depending on kind and scale of physics generating the operator.	MCFODO (highly dependent on production and decay mode).	5.2	
Dynamical DM	Dark sector includes spectrum of states with varying life-time up to hyperstable DM states.	Any, DDM ensemble contains short to hyperstable $c\tau$.	MCFODO (highly dependent on production and decay mode).	5.5	32, 33
SIMP/ELDER DM	Strong dynamics of HV generate DM abundance. HV \rightarrow LLPs.	Any.	See HV.	5.4.1, 5.4.2	
Relaxion	Relaxion or other new scalars in theory generically mix with Higgs \rightarrow SM+S.	Any.	See SM+S.	4.4	
Axion-like particles	ALP couplings to h and Z are generic in EFT framework. $1/f$ suppression makes ALP an LLP.	Any.	MCFODO for low-scale f .	8.6	66, 67, 68, 69, 70
Leptogenesis	Motivates minimal RH neutrino model and other neutrino extensions, which generically feature LLPs.	Freeze-out LG favors weak-scale m_N but not so for other scenarios. Lower m_N favor long lifetimes.	Generally very difficult to probe, especially at high leptogenesis scale. In long-lifetime/low-mass regime, MATHUSLA and/or SHiP may be only/first discovery opportunity.	6.3	
Scalars in neutrino extensions	Gauge extensions in neutrino models give rise to new scalars that can mix with Higgs \rightarrow SM+S. Provides additional S production modes via heavy gauge boson decay.	Any.	See SM+S, with some additional production modes (new heavy gauge bosons).	7.2.2, 7.3.2	

Table 5: BSM scenarios discussed in this document where neutral LLP signals at MATHUSLA are a strongly motivated generic possibility (often as part of a broad parameter or theory space), *and* MATHUSLA Could be First or Only Discovery Opportunity (MCFODO). When discussing lifetimes, “any” means up to the BBN limit, “long” means the MATHUSLA regime. HV = hidden valley. Since lifetimes are mostly arbitrary here, we focus on the long lifetime regime when discussing the role of MATHUSLA.

We close by pointing out that MATHUSLA is not only a very strongly motivated and relatively affordable way of extending the capabilities of the LHC, the concept is also exceedingly flexible, general and scalable. Future proton colliders, like the 100 TeV FCC-hh [470–472] or SPPC [473] should include as part of their design an underground, shielded, dedicated displaced vertex detector to maximize their discovery potential for new physics. At the HL-LHC, MATHUSLA can be constructed incrementally in a modular fashion, and even a much smaller initial version than the $200\text{m} \times 200\text{m} \times 20\text{m}$ benchmark assumed in this document could quickly supply the world’s best sensitivity to many LLP physics scenarios, with the possibility of discovery within a few years. All of this makes MATHUSLA a uniquely exciting opportunity for the upcoming HL-LHC upgrade that would continue to yield physics dividends into the HE-LHC era and beyond.

Acknowledgements

The editors would like to thank David Morrissey for valuable feedback on a draft version of this manuscript. The work of D. Curtin was supported in part by the National Science Foundation under grant PHY-1620074, and by the Maryland Center for Fundamental Physics. The work of M. Drewes was supported by the Collaborative Research Center SFB1258 of the Deutsche Forschungsgemeinschaft and by the DFG cluster of excellence "Origin and Structure of the Universe" (universe-cluster.de). The work of P. Meade was supported in part by the National Science Foundation under grant NSF-PHY-1620628. R. N. Mohapatra, R. Sundrum and Y. Tsai acknowledge the support of NSF grant No. PHY1620074. The work of J. Shelton was supported in part by DOE grant DE-SC0017840. The work of B. Batell is supported in part by the US Department of Energy under grant DE-SC0015634. The work of N. Blinov is supported by the U.S. Department of Energy under Contract No. DE-AC02-76SF00515. The work of A. Fernández T  lez, K. S. Caballero-Mora, M. Rodr  guez Cahuantzi and J. C. Arteaga Velazquez is supported in part by CONACyT-M  xico, Grants CB 243290. The work of Jae Hyeok Chang is supported by DoE Grant DE-SC0017938. The work of T. Cohen was supported in part by the U.S. Department of Energy under grant number DE-SC0018191 and DE-SC0011640. The work of P. Cox is supported by the World Premier International Research Center Initiative (WPI), MEXT, Japan. The work of N. Craig is supported in part by the US Department of Energy under the grant DE-SC0014129. The research activities of K. R. Dienes are supported in part by the U.S. Department of Energy under Grant DE-FG02-13ER41976 (DE-SC0009913) and by the U.S. National Science Foundation through its employee IR/D program. The opinions and conclusions expressed herein are those of the authors, and do not represent any funding agencies. J. Dror is supported in part by the DOE under contract DE-AC02-05CH11231. R. Essig acknowledges support from DoE Grant DE-SC0017938. The work of J. A. Evans was supported in part by NSF CAREER grant PHY-1654502 and by DOE grant DE-SC0015655. The work of A. Fern  ndez T  lez, K. S. Caballero Mora, M. Rodr  guez Cahuantzi and J. C. Arteaga Velazquez is supported in part by CONACyT-M  xico, Grants 243290. The work of E. Fuchs was supported in part by the Minerva Foundation. The work of T. Gherghetta was supported in part by the U.S. Department of Energy under Grant DE-SC0011842. The CP3-Origins centre (C. Hagedorn) is partially funded by the Danish National Research Foundation, grant number DNRF90. J. C. Helo is supported by Chile grants Fondecyt No. 1161463, Conicyt PIA/ACT 1406 and Basal FB0821. M. Hirsch is supported by the Spanish MICINN grants FPA2017-85216-P, SEV-2014-0398 and PROMETEOII/2014/084 (Generalitat Valenciana). Y. Hochberg is supported by the Israel Science Foundation (grant No. 1112/17), by the Binational Science Foundation (grant No. 2016155), by the I-CORE Program of the Planning Budgeting Committee (grant No. 1937/12), by the German Israel Foundation (grant No. I-2487-303.7/2017), and by the Azrieli Foundation. A. Hook is supported by the NSF under grant NSF-PHY-1620074 and the Maryland Center for Fundamental Physics. A. Ibarra is partially supported by the DFG cluster of excellence EXC 153 "Origin and Structure of the Universe" and by the Collaborative Research Center SFB1258. S. Ipek is supported by the University of California President's Postdoctoral Fellowship Program and partially supported by NSF Grant No. PHY-1620638. S. Joon is supported by Korea NRF-2017R1D1A1B03030820 and NRF-2015R1A4A1042542. Simon Knapen is supported by the DoE under contract 5DE-AC02-05CH11231, and the National Science Foundation (NSF) under grants No. PHY-1316783 and No. PHY-1002399. The work of E. Kuflik is supported by the I-CORE Program of the Planning Budgeting Committee of the Israel Science Foundation (grant No. 1937/12), by the Binational Science Foundation (grant No. 2016153), and the Israel Science Foundation (grant No. 1111/17). This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. H. J. Lubatti and C. Alpigiani are supported in part by National Science Foundation grant PHY-1509257. The research of D. McKeen is supported in part by the National Research Council of Canada. The work of S. Moretti is supported in part through the NExT Institute, the STFC CG ST/L000296/1 and the H2020-MSCA-RISE-2014 grant n. 645722 (NonMinimalHiggs). The work of N. Nagata was supported in part by the Grant-in-Aid for Scientific Research No. 17K14270. The work of J. M. No was partially

supported by the European Research Council under the European Unions Horizon 2020 program (ERC Grant Agreement no.648680 DARKHORIZONS), and by the Programa Atraccion de Talento de la Comunidad de Madrid under grant n. 2017-T1/TIC-5202. The work of G. Perez is supported by grants from the BSF, ERC, ISF, Minerva, and the Weizmann-UK Making Connections Program. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Research and Innovation. The work of M. Reece was supported in part by the DOE Grant de-sc0013607. The work of D. J. Robinson was supported in part by NSF grant PHY-1720252. The work of A. Fernández Téllez, K. S. Caballero Mora, M. Rodríguez-Cahuantzi and J. C. Arteaga Velazquez is supported in part by CONACyT-México, Grants 243290. M. Rodríguez Cahuantzi also thankfully acknowledge the computer resources, technical expertise and support provided by the Laboratorio Nacional de Supercómputo del Sureste de México and VIEP-BUAP. Daniel Stolarski is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC). The work was supported by the Instituto de Investigaciones Físicas (IIF) and Universidad Mayor de San Andrés (UMSA) in collaboration with Benemérita Universidad Autónoma de Puebla (BUAP), Puebla, México. B. Tweedie was supported by DoE grant No. DE-FG02-95ER40896 and by PITT PACC. This research is supported by the Cluster of Excellence PrecisionPhysics, Fundamental Interactions and Structure of Matter (PRISMA-EXC1098). The work of B. Zaldivar was supported in part by the "Investissements d'avenir" program of the French ANR, Labex "ENIGMASS", and by the Programa Atraccion de Talento de la Comunidad de Madrid under grant n. 2017-T2/TIC-5455. The work of Yongchao Zhang was partially supported by the IISN and Belgian Science Policy (IAP VII/37).

References

- [1] J. P. Chou, D. Curtin, and H. J. Lubatti, *New Detectors to Explore the Lifetime Frontier*, *Phys. Lett.* **B767** (2017) 29–36, [arXiv:1606.06298 \[hep-ph\]](#).
- [2] S. P. Martin, *A Supersymmetry primer*, [arXiv:hep-ph/9709356 \[hep-ph\]](#). [Adv. Ser. Direct. High Energy Phys.18,1(1998)].
- [3] P. Meade, S. Nussinov, M. Papucci, and T. Volansky, *Searches for Long Lived Neutral Particles*, *JHEP* **06** (2010) 029, [arXiv:0910.4160 \[hep-ph\]](#).
- [4] K. Jedamzik, *Big bang nucleosynthesis constraints on hadronically and electromagnetically decaying relic neutral particles*, *Phys. Rev.* **D74** (2006) 103509, [arXiv:hep-ph/0604251 \[hep-ph\]](#).
- [5] R. T. Co, F. D’Eramo, and L. J. Hall, *Gravitino or Axino Dark Matter with Reheat Temperature as high as 10^{16} GeV*, *JHEP* **03** (2017) 005, [arXiv:1611.05028 \[hep-ph\]](#).
- [6] P. S. Bhupal Dev, R. N. Mohapatra, and Y. Zhang, *Displaced photon signal from a possible light scalar in minimal left-right seesaw model*, *Phys. Rev.* **D95** (2017) no. 11, 115001, [arXiv:1612.09587 \[hep-ph\]](#).
- [7] P. S. B. Dev, R. N. Mohapatra, and Y. Zhang, *Long Lived Light Scalars as Probe of Low Scale Seesaw Models*, *Nucl. Phys.* **B923** (2017) 179–221, [arXiv:1703.02471 \[hep-ph\]](#).
- [8] A. Caputo, P. Hernandez, J. Lopez-Pavon, and J. Salvado, *The seesaw portal in testable models of neutrino masses*, *JHEP* **06** (2017) 112, [arXiv:1704.08721 \[hep-ph\]](#).
- [9] D. Curtin and M. E. Peskin, *Analysis of Long Lived Particle Decays with the MATHUSLA Detector*, *Phys. Rev.* **D97** (2018) no. 1, 015006, [arXiv:1705.06327 \[hep-ph\]](#).
- [10] J. A. Evans, S. Gori, and J. Shelton, *Looking for the WIMP Next Door*, *JHEP* **02** (2018) 100, [arXiv:1712.03974 \[hep-ph\]](#).
- [11] J. C. Helo, M. Hirsch, and Z. S. Wang, *Heavy neutral fermions at the high-luminosity LHC*, [arXiv:1803.02212 \[hep-ph\]](#).
- [12] R. T. D’Agnolo, C. Mondino, J. T. Ruderman, and P.-J. Wang, *Exponentially Light Dark Matter from Coannihilation*, [arXiv:1803.02901 \[hep-ph\]](#).
- [13] A. Berlin, S. Gori, P. Schuster, and N. Toro, *Dark Sectors at the Fermilab SeaQuest Experiment*, [arXiv:1804.00661 \[hep-ph\]](#).
- [14] F. F. Deppisch, W. Liu, and M. Mitra, *Long-lived Heavy Neutrinos from Higgs Decays*, [arXiv:1804.04075 \[hep-ph\]](#).
- [15] S. Jana, N. Okada, and D. Raut, *Displaced vertex signature of type-I seesaw*, [arXiv:1804.06828 \[hep-ph\]](#).
- [16] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, *Dark Higgs bosons at the ForwArd Search ExpeRiment*, *Phys. Rev.* **D97** (2018) no. 5, 055034, [arXiv:1710.09387 \[hep-ph\]](#).
- [17] J. Feng, I. Galon, F. Kling, and S. Trojanowski, *ForwArd Search ExpeRiment at the LHC*, *Phys. Rev.* **D97** (2018) no. 3, 035001, [arXiv:1708.09389 \[hep-ph\]](#).
- [18] J. A. Evans, *Detecting Hidden Particles with MATHUSLA*, *Phys. Rev.* **D97** (2018) no. 5, 055046, [arXiv:1708.08503 \[hep-ph\]](#).
- [19] V. V. Gligorov, S. Knapen, M. Papucci, and D. J. Robinson, *Searching for Long-lived Particles: A Compact Detector for Exotics at LHCb*, *Phys. Rev.* **D97** (2018) no. 1, 015023, [arXiv:1708.09395 \[hep-ph\]](#).
- [20] The LHC Long-Lived Particle Community, *Searches for long-lived particles beyond the Standard Model at the Large Hadron Collider*, In preparation .
- [21] M. J. Strassler and K. M. Zurek, *Echoes of a hidden valley at hadron colliders*, *Phys. Lett.* **B651** (2007) 374–379, [arXiv:hep-ph/0604261 \[hep-ph\]](#).
- [22] M. J. Strassler and K. M. Zurek, *Discovering the Higgs through highly-displaced vertices*, *Phys.*

- Lett.* **B661** (2008) 263–267, [arXiv:hep-ph/0605193](#) [hep-ph].
- [23] *The MATHUSLA Cosmic Ray Physics Case*, Whitepaper, in preparation .
- [24] J. Ying, Y. L. Ye, Y. Ban, H. T. Liu, Z. M. Zhu, Z. Y. Zhu, T. Chen, J. G. Ma, and S. J. Qian, *Study of an avalanche-mode resistive plate chamber*, *J. Phys.* **G26** (2000) 1291–1298.
- [25] CMS Collaboration, F. Thyssen, *Performance of the resistive plate chambers in the CMS experiment*, *JINST* **7** (2012) C01104.
- [26] C. Bacci et al., *Performance of the RPCs for the ARGO detector operated at the YangBaJing laboratory (4300 m a.s.l.)*, *Nucl. Instrum. Meth.* **A508** (2003) 110–115.
- [27] Argo-YBJ Collaboration, G. Aielli et al., *Layout and performance of RPCs used in the Argo-YBJ experiment*, *Nucl. Instrum. Meth.* **A562** (2006) 92–96.
- [28] R. Iuppa, *Potential of RPCs in cosmic ray experiments for the next decade*, *JINST* **10** (2015) no. 04, C04044.
- [29] J. Bluemer, R. Engel, and J. R. Hoerandel, *Cosmic Rays from the Knee to the Highest Energies*, *Prog.Part.Nucl.Phys.* **63** (2009) 293–338, [arXiv:0904.0725](#) [astro-ph].
- [30] I. A. Grenier, J. H. Black, and A. W. Strong, *The Nine Lives of Cosmic Rays in Galaxies*, *Annu. Rev. Astron. Astrophys.* **53** (2015) 199–246.
- [31] S. Mollerach and E. Roulet, *Progress in high-energy cosmic ray physics*, *Prog. Part. Nucl. Phys.* **98** (2018) 85–118, [arXiv:1710.11155](#) [astro-ph].
- [32] Pierre Auger Collaboration, A. Aab et al., *Progress in high-energy cosmic ray physics*, *Science* **357** (2017) 1266–1270, [arXiv:1709.07321](#) [astro-ph].
- [33] G. Giacinti, M. Kachelriess, and D. V. Semikoz, *Escape model for Galactic cosmic rays and an early extragalactic transition*, *Phys. Rev.* **D91** (2015) 083009, [arXiv:1502.01608](#) [astro-ph].
- [34] T. K. Gaisser, T. Stanev, and S. Tilav, *Cosmic ray energy spectrum from measurements of air showers*, *Front. Phys.* **8(6)** (2013) 748–758, [arXiv:1303.3565](#) [astro-ph].
- [35] O. Deligny, *Cosmic-Ray Anisotropies: A Review*, XXV ECRS 2016 Proceedings , [arXiv:1612.08002](#) [astro-ph].
- [36] Fermi-LAT Collaboration, M. Ackermann et al., *Detection of the Characteristic Pion-Decay Signature in Supernova Remnants*, *Science* **339** (2013) 807–811, [arXiv:1302.3307](#) [astro-ph].
- [37] HESS Collaboration, A. Abramowski et al., *Acceleration of petaelectronvolt protons in the Galactic Centre*, *Nature* **531** (2016) 476–479, [arXiv:1603.07730](#) [astro-ph].
- [38] K. Fang and K. Murase, *Linking high-energy cosmic particles by black-hole jets embedded in large-scale structures*, *Nature Phys.* **14** (2018) 396–398, [arXiv:1704.00015](#) [astro-ph].
- [39] Pierre Auger Collaboration, A. Aab et al., *Muons in air showers at the Pierre Auger Observatory: mean number in highly inclined events*, *Phys. Rev. D* **91** (2015) 032003, [arXiv:1408.1421](#) [astro-ph].
- [40] Pierre Auger Collaboration, A. Aab et al., *Testing Hadronic Interactions at Ultrahigh Energies with Air Showers Measured by the Pierre Auger Observatory*, *Phys. Rev. Lett.* **117** (2016) 192001, [arXiv:1610.08509](#) [astro-ph].
- [41] Pierre Auger Collaboration, A. Aab et al., *Inferences on Mass Composition and Tests of Hadronic Interactions from 0.3 to 100 EeV using the water-Cherenkov Detectors of the Pierre Auger Observatory*, *Phys. Rev. D* **96** (2017) 122003, [arXiv:1710.07249](#) [astro-ph].
- [42] Particle Data Group Collaboration, C. Patrignani et al., *Review of Particle Physics*, *Chin. Phys.* **C40** (2016) no. 10, 100001.
- [43] G. Kulikov and G. B. Khristiansen, *On the size spectrum of extensive air showers*, *Sov. Phys. JETP* **35(8)** (1959) 441–444.

- [44] KASCADE-Grande Collaboration, W. D. Apel et al., *The spectrum of high-energy cosmic rays measured with KASCADE-Grande*, *Astropart. Phys.* **36** (2012) 183–194, [arXiv:1206.3834 \[astro-ph\]](#).
- [45] IceCube Collaboration, M. Aartsen et al., *Measurement of the cosmic ray energy spectrum with IceTop-73*, *Phys. Rev. D* **88** (2013) 042004, [arXiv:1307.3795 \[astro-ph.HE\]](#).
- [46] TUNKA Collaboration, V. V. Prosin et al., *Tunka-133: Results of 3 year operation*, *Nucl. Instrum. Meth. A* **756** (2014) 94–101.
- [47] Eur. Phys. J. Plus Collaboration, M. Settimo et al., *Measurement of the cosmic ray energy spectrum using hybrid events of the Pierre Auger Observatory*, *Phys. Rev. D* **127** (2012) 87, [arXiv:1208.6574 \[astro-ph.HE\]](#).
- [48] Telescope Array Collaboration, T. Abu-Zayyad et al., *The cosmic-ray energy spectrum observed with the surface detector of the Telescope Array Experiment*, *ApJL* **768** (2013) L1, [arXiv:1205.5067 \[astro-ph.HE\]](#).
- [49] AMS Collaboration, A. Kounine, *The Alpha Magnetic Spectrometer on the International Space Station*, *Int. J. Mod. Phys. E* **21** (2012) 1230005.
- [50] AMS Collaboration, M. Aguilar, *Precision Measurement of the Proton Flux in Primary Cosmic Rays from Rigidity 1 GV to 1.8 TV with the Alpha Magnetic Spectrometer on the International Space Station*, *Phys. Rev. Lett.* **114** (2015) 171103.
- [51] CREAM Collaboration, H. S. Ahn et al., *The Cosmic Ray Energetics And Mass (CREAM) instrument*, *Nucl. Instrum. Methods A* **579** (2007) 1034–1053.
- [52] CREAM Collaboration, H. S. Ahn et al., *Discrepant hardening observed in cosmic-ray elemental spectra*, *Astrophys. J.* **714** (2010) L89.
- [53] E. S. Seo, *Direct measurements of cosmic rays using balloon borne experiments*, *Astrop. Phys.* **39-40** (2012) 76–87.
- [54] E. Mocchiut, *Direct detection of cosmic rays: through a new era of precision measurements of particle fluxes*, *Nuc. Phys. B - Proc. Suppl.* **256-257** (2014) 161–172, [arXiv:1407.1143 \[astro-ph\]](#).
- [55] P. Maestro, *Cosmic rays: direct measurements*, Proc. of the 34th International Cosmic Ray Conference, PoS(ICRC2015)016, The Hague, The Netherlands, [arXiv:1510.07683 \[astro-ph\]](#).
- [56] K. H. Kampert and A. A. Watson, *Extensive Air Showers and Ultra High-Energy Cosmic Rays: A Historical Review*, *Eur. Phys. J. H* **37** (2012) 359–412, [arXiv:1207.4827 \[physics.hist-ph\]](#).
- [57] R. Engel et al., *Extensive Air Showers and Hadronic Interactions at High Energy*, *Ann. Rev. of Nucl. and Part. Sc.* **61** (2011) 467–489.
- [58] R. Engel, *Indirect detection of cosmic rays*. Springer-Verlag, Berlin, 2012.
- [59] P. K. F. Grieder, *Extensive air showers: High Energy Phenomena and Astrophysical Aspects (A tutorial, reference manual and data book) Vol. I and II*. Springer-Verlag, Heidelberg, 2010.
- [60] Pierre Auger Collaboration, J. Abraham et al., *Properties and performance of the prototype instrument for the Pierre Auger Observatory*, *Nucl. Instrum. Meth. A* **523** (2004) 50–95.
- [61] Telescope Array Collaboration, R. U. Abbasi et al., *The Cosmic-Ray Energy Spectrum between 2 PeV and 2 EeV Observed with the TALE detector in monocular mode*, [arXiv:1803.01288 \[astro-ph\]](#).
- [62] KASCADE Collaboration, T. Antoni et al., *The cosmic-ray experiment KASCADE*, *Nucl. Instrum. Meth. A* **513** (2003) 490–510.
- [63] Pierre Auger Collaboration, J. Abraham et al., *Measurement of the energy spectrum of cosmic rays above 10^{18} eV using the Pierre Auger Observatory*, *Phys.Lett.B* **685** (2010) 239–246,

- [arXiv:1002.1975](#) [[astro-ph](#)].
- [64] Argo-YBJ Collaboration, G. Aielli et al., *The cosmic-ray experiment KASCADE*, *Nucl. Instrum. Meth. A* **562** (2006) 92–96.
- [65] J. Knapp et al., *Extensive air shower simulations at the highest energies*, *Astrop. Phys.* **19** (2003) 77–99, [arXiv:0206414](#) [[astro-ph](#)].
- [66] S. Ostapchenko, *Models for cosmic ray interactions*, *Czech. J. Phys.* **56** (2006) A149–A159, [arXiv:0601230](#) [[hep-ph](#)].
- [67] T. Pierog et al., *Impact of uncertainties in hadron production on air-shower predictions*, *Czech. J. Phys.* **56** (2006) A161–A172, [arXiv:0602190](#) [[astro-ph](#)].
- [68] KASCADE Collaboration, T. Antoni et al., *KASCADE measurements of energy spectra for elemental groups of cosmic rays: Results and open problems*, *Astropart. Phys.* **24** (2005) 1–25, [arXiv:0505413](#) [[astro-ph](#)].
- [69] KASCADE Collaboration, T. Antoni et al., *Test of high-energy interaction models using the hadronic core of EAS*, *J. Phys. G: Nucl. Part. Phys.* **25** (1999) 2161–2175, [arXiv:9904287](#) [[astro-ph](#)].
- [70] KASCADE Collaboration, T. Antoni et al., *Test of Hadronic Interaction Models in the Forward Region with KASCADE Event Rates*, *J. Phys. G: Nucl. Part. Phys.* **27** (2001) 1785–1798, [arXiv:0106494](#) [[astro-ph](#)].
- [71] KASCADE Collaboration, T. Antoni et al., *The information from muon arrival time distributions of high-energy EAS as measured with the KASCADE detector*, *Astropart. Phys.* **18** (2003) 319–331, [arXiv:0204266](#) [[astro-ph](#)].
- [72] KASCADE Collaboration, T. Antoni et al., *Geometric structures in hadronic cores of extensive air showers observed by KASCADE*, *Phys. Rev. D* **71** (2005) 072002, [arXiv:0503218](#) [[hep-ph](#)].
- [73] KASCADE Collaboration, T. Antoni et al., *Comparison of measured and simulated lateral distributions for electrons and muons with KASCADE*, *Astropart. Phys.* **24** (2006) 467–483, [arXiv:0510810](#) [[hep-ph](#)].
- [74] KASCADE Collaboration, W. D. Apel et al., *Test of interaction models up to 40 PeV by studying hadronic cores of EAS*, *J. Phys. G: Nucl. Part. Phys.* **34** (2007) 2581–259.
- [75] KASCADE-Grande Collaboration, W. D. Apel et al., *Test of the hadronic interaction model EPOS with air shower data*, *J. Phys. G: Nucl. Part. Phys.* **36** (2009) 035201, [arXiv:0901.4650](#) [[hep-ph](#)].
- [76] ARGO-YBJ and LHAASO Collaborations, B. Bartoli et al., *Knee of the cosmic hydrogen and helium spectrum below 1 PeV measured by ARGO-YBJ and a Cherenkov telescope of LHAASO*, *Phys. Rev. D* **92** (2015) 092005, [arXiv:1502.03164](#) [[astro-ph](#)].
- [77] KASCADE-Grande Collaboration, W. D. Apel et al., *KASCADE-Grande measurements of energy spectra for elemental groups of cosmic rays*, *Astrop. Phys.* **47** (2013) 54–66, [arXiv:1308.2098](#) [[astro-ph](#)].
- [78] KASCADE Collaboration, T. Antoni et al., *Search for Cosmic-Ray Point Sources with KASCADE*, *Astrophys. J.* **608** (2004) 865–871, [arXiv:0402656](#) [[astro-ph](#)].
- [79] KASCADE-Grande Collaboration, S. Over et al., *Search for anisotropy and point sources of cosmic rays with the KASCADE-Grande experiment*, .
- [80] KASCADE-Grande Collaboration, D. Kang et al., *Search for gamma-ray point sources with KASCADE*, Proc. of the 34th ICRC **PoS(ICRC2015)812** .
- [81] KASCADE-Grande Collaboration, W. D. Apel et al., *Muon production height studies with the air shower experiment KASCADE-Grande*, *Astrop. Phys.* **34** (2011) 476–485.
- [82] KASCADE-Grande Collaboration, W. D. Apel et al., *Probing the evolution of the EAS muon*

- content in the atmosphere with KASCADE-Grande, *Astrop. Phys.* **95** (2017) 25–43, [arXiv:1801.05513 \[astro-ph\]](#).
- [83] Pierre Auger Collaboration, A. Aab et al., *Muons in air showers at the Pierre Auger Observatory: Measurement of atmospheric production depth*, *Phys. Rev. D* **90** (2014) 012012, [arXiv:1407.5919 \[astro-ph\]](#).
- [84] HiRes and MIA Collaborations, T. Abu-Zayyad et al., *A Multi-Component Measurement of the Cosmic Ray Composition Between 10^{17} eV and 10^{18} eV*, *Phys. Rev. Lett.* **84** (2000) 4276–4279, [arXiv:9911144 \[astro-ph\]](#).
- [85] A. A. Petrukhin, *Muon puzzle in cosmic ray experiments and its possible solution*, *Nucl. Instrum. Meth. Phys. Res. A* **742** (2014) 228–231.
- [86] Pierre Auger Collaboration, A. Aab et al., *An improved limit to the diffuse flux of ultra-high energy neutrinos from the Pierre Auger Observatory*, *Phys. Rev. D* **91** (2015) 092008, [arXiv:1504.05397 \[astro-ph\]](#).
- [87] J. L. Feng et al., *Observability of Earth-skimming Ultra-high Energy Neutrinos*, *Phys. Rev. Lett.* **88** (2002) 161102, [arXiv:0105067 \[hep-ph\]](#).
- [88] T. K. Gaisser et al., *Particle Physics with High Energy Neutrinos*, *Phys. Rept* **258** (1995) 173–236, [arXiv:9410384 \[hep-ph\]](#).
- [89] MACRO Collaboration, M. Ambrosio et al., *Atmospheric neutrino oscillations from upward through going muon multiple scattering in MACRO*, *Phys. Lett. B* **566** (2003) 35–44, [arXiv:0304037 \[hep-ex\]](#).
- [90] MACRO Collaboration, G. L. Fogli et al., *Upward-going muons and neutrino oscillations*, *Phys. Rev. D* **57** (1998) 5893–5902, [arXiv:9708213 \[hep-ph\]](#).
- [91] C. K. Jung et al., *Neutrino Oscillations in the Atmospheric Parameter Region: From the Early Experiments to the Present*, *Ann. Rev. of Nucl. and Part. Sci.* **51** (2001) 451–488.
- [92] G. Giacomelli et al., *Neutrino Oscillations in the Atmospheric Parameter Region: From the Early Experiments to the Present*, *Advances in High Energy Physics* **2013** (2013) 464926.
- [93] C. Grupen et al., *Measurements of the muon component of extensive air showers at 320 m.w.e. underground*, *Nucl. Instrum. Meth. A* **510** (2003) 190–193.
- [94] DELPHI Collaboration, J. Ridky et al., *Detection of muon bundles from cosmic ray showers by the DELPHI experiment*, *Nucl. Phys. Proc. Suppl. B* **138** (2005) 295–298.
- [95] ALICE Collaboration, J. Adam et al., *Study of cosmic ray events with high muon multiplicity using the ALICE detector at the CERN Large Hadron Collider*, *JCAP* **01** (2016) 032, [arXiv:1507.07577 \[astro-ph\]](#).
- [96] P. Kankiewicz et al., *Study of cosmic ray events with high muon multiplicity using the ALICE detector at the CERN Large Hadron Collider*, *ApJ* **839** (2017) 31, [arXiv:1612.04749 \[hep-ph\]](#).
- [97] J. Ridky, *Can We Observe the Quark Gluon Plasma in Cosmic Ray Showers?*, *Astropart. Phys.* **17** (2002) 355–365, [arXiv:0012068 \[hep-ph\]](#).
- [98] HAWC Collaboration, V. Joshi et al., *HAWC High Energy Upgrade with a Sparse Outrigger Array*, Proc. of the 35th ICRC **PoS(ICRC2017)806**, [arXiv:1708.04032 \[astro-ph.IM\]](#).
- [99] TAIGA Collaboration, N. Budnev et al., *The TAIGA experiment: From cosmic-ray to gamma-ray astronomy in the Tunka valley*, *Nucl. Instrum. Meth. A* **845** (2017) 330–333.
- [100] IceCube-Gen2 Collaboration, M. G. Aartsen et al., *The IceCube Neutrino Observatory - Contributions to ICRC 2017 Part VI: IceCube-Gen2, the Next Generation Neutrino Observatory*, Proc. of the 35th ICRC, [arXiv:1710.01207 \[astro-ph.IM\]](#).
- [101] 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*, *JHEP* **07** (2014) 079, [arXiv:1405.0301 \[hep-ph\]](#).
- [102] T. Sjostrand, S. Mrenna, and P. Z. Skands, *PYTHIA 6.4 Physics and Manual*, *JHEP* **05** (2006) 026, [arXiv:hep-ph/0603175 \[hep-ph\]](#).
- [103] C. Borschensky, M. Krämer, A. Kulesza, M. Mangano, S. Padhi, T. Plehn, and X. Portell, *Squark and gluino production cross sections in pp collisions at $\sqrt{s} = 13, 14, 33$ and 100 TeV*, *Eur. Phys. J. C* **74** (2014) no. 12, 3174, [arXiv:1407.5066 \[hep-ph\]](#).
- [104] M. Baak et al., *LHC SUSY Cross Section Working Group*, in *Proceedings, 2013 Community Summer Study on the Future Physics: Snowmass on the Mississippi 29-August 6, 2013*.
- [105] G. Busoni et al., *Recommendations on presenting LHC searches for missing transverse energy signals using simplified s-channel models of dark matter*, [arXiv:1603.04156 \[hep-ex\]](#).
- [106] M. Baak et al., *LHC Higgs Cross Section Working Group (2012-2013)*, in *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*.
- [107] LHCb Collaboration, R. Aaij et al., *Measurement of the b-quark production cross-section in 7 and 13 TeV pp collisions*, *Phys. Rev. Lett.* **118** (2017) no. 5, 052002, [arXiv:1612.05140 \[hep-ex\]](#).
- [108] R. Mackeprang and A. Rizzi, *Interactions of Coloured Heavy Stable Particles in Matter*, *Eur. Phys. J. C* **50** (2007) 353–362, [arXiv:hep-ph/0612161 \[hep-ph\]](#).
- [109] R. Mackeprang and D. Milstead, *An Updated Description of Heavy-Hadron Interactions in GEANT-4*, *Eur. Phys. J. C* **66** (2010) 493–501, [arXiv:0908.1868 \[hep-ph\]](#).
- [110] CMS Collaboration, V. Khachatryan et al., *Search for long-lived charged particles in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **94** (2016) no. 11, 112004, [arXiv:1609.08382 \[hep-ex\]](#).
- [111] ATLAS Collaboration, M. Aaboud et al., *Search for heavy long-lived charged R-hadrons with the ATLAS detector in 3.2 fb⁻¹ of proton–proton collision data at $\sqrt{s} = 13$ TeV*, *Phys. Lett. B* **760** (2016) 647–665, [arXiv:1606.05129 \[hep-ex\]](#).
- [112] ATLAS Collaboration, G. Aad et al., *Searches for heavy long-lived charged particles with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **01** (2015) 068, [arXiv:1411.6795 \[hep-ex\]](#).
- [113] DELPHES 3 Collaboration, J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi, *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057, [arXiv:1307.6346 \[hep-ex\]](#).
- [114] C. Collaboration, *The Phase-2 Upgrade of the CMS Tracker*, CERN-LHCC-2017-009. CMS-TDR-014, CERN, Geneva, Jun, 2017.
- [115] J. M. Lindert et al., *Precise predictions for V + jets dark matter backgrounds*, *Eur. Phys. J. C* **77** (2017) no. 12, 829, [arXiv:1705.04664 \[hep-ph\]](#).
- [116] CMS Collaboration, *Projected Performance of an Upgraded CMS Detector at the LHC and HL-LHC: Contribution to the Snowmass Process*, in *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*.
- [117] Y. Gershtein, *CMS Hardware Track Trigger: New Opportunities for Long-Lived Particle Searches at the HL-LHC*, *Phys. Rev. D* **96** (2017) no. 3, 035027, [arXiv:1705.04321 \[hep-ph\]](#).
- [118] ATLAS Collaboration, M. Aaboud et al., *Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, [arXiv:1710.04901 \[hep-ex\]](#).
- [119] Abercrombie, Daniel et al., *Dark Matter Benchmark Models for Early LHC Run-2 Searches: Report of the ATLAS/CMS Dark Matter Forum*, [arXiv:1507.00966 \[hep-ex\]](#).

- [120] Buchmueller, Oliver and Dolan, Matthew J. and Malik, Sarah A. and McCabe, Christopher, *Characterising dark matter searches at colliders and direct detection experiments: Vector mediators*, **JHEP** **01** (2015) 037, [arXiv:1407.8257 \[hep-ph\]](#).
- [121] Malik, Sarah A. *et al.*, *Interplay and Characterization of Dark Matter Searches at Colliders and in Direct Detection Experiments*, **Phys. Dark Univ.** **9-10** (2015) 51–58, [arXiv:1409.4075 \[hep-ex\]](#).
- [122] Harris, Philip and Khoze, Valentin V. and Spannowsky, Michael and Williams, Ciaran, *Constraining Dark Sectors at Colliders: Beyond the Effective Theory Approach*, **Phys. Rev.** **D91** (2015) 055009, [arXiv:1411.0535 \[hep-ph\]](#).
- [123] Haisch, Ulrich and Re, Emanuele, *Simplified dark matter top-quark interactions at the LHC*, **JHEP** **06** (2015) 078, [arXiv:1503.00691 \[hep-ph\]](#).
- [124] Abdallah, Jalal and others, *Simplified Models for Dark Matter Searches at the LHC*, **Phys. Dark Univ.** **9-10** (2015) 8–23, [arXiv:1506.03116 \[hep-ph\]](#).
- [125] CMS Collaboration, A. M. Sirunyan *et al.*, *Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV*, **JHEP** **01** (2018) 097, [arXiv:1710.00159 \[hep-ex\]](#).
- [126] CMS Collaboration, A. M. Sirunyan *et al.*, *Search for dijet resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter and other models*, **Phys. Lett.** **B769** (2017) 520–542, [arXiv:1611.03568 \[hep-ex\]](#). [Erratum: **Phys. Lett.** **B772**,882(2017)].
- [127] ATLAS Collaboration, M. Aaboud *et al.*, *Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV*, [arXiv:1804.03496 \[hep-ex\]](#).
- [128] ATLAS Collaboration, T. A. collaboration, *Search for new light resonances decaying to jet pairs and produced in association with a photon or a jet in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, .
- [129] CMS Collaboration, *Searches for dijet resonances in pp collisions at $\sqrt{s} = 13$ TeV using data collected in 2016.*, CMS-PAS-EXO-16-056, CERN, Geneva, 2017.
- [130] P. Schwaller and J. Zurita, *Compressed electroweakino spectra at the LHC*, **JHEP** **03** (2014) 060, [arXiv:1312.7350 \[hep-ph\]](#).
- [131] ATLAS Collaboration, G. Aad *et al.*, *Search for long-lived, weakly interacting particles that decay to displaced hadronic jets in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, **Phys. Rev.** **D92** (2015) no. 1, 012010, [arXiv:1504.03634 \[hep-ex\]](#).
- [132] ATLAS Collaboration, G. Aad *et al.*, *Search for massive, long-lived particles using multitrack displaced vertices or displaced lepton pairs in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, **Phys. Rev.** **D92** (2015) no. 7, 072004, [arXiv:1504.05162 \[hep-ex\]](#).
- [133] CMS Collaboration, V. Khachatryan *et al.*, *Search for Long-Lived Neutral Particles Decaying to Quark-Antiquark Pairs in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV*, **Phys. Rev.** **D91** (2015) no. 1, 012007, [arXiv:1411.6530 \[hep-ex\]](#).
- [134] CMS Collaboration, V. Khachatryan *et al.*, *Search for long-lived particles that decay into final states containing two electrons or two muons in proton-proton collisions at $\sqrt{s} = 8$ TeV*, **Phys. Rev.** **D91** (2015) no. 5, 052012, [arXiv:1411.6977 \[hep-ex\]](#).
- [135] ATLAS Collaboration, T. A. collaboration, *Search for long-lived neutral particles decaying into displaced lepton jets in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, .
- [136] ATLAS Collaboration, G. Aad *et al.*, *Triggers for displaced decays of long-lived neutral particles in the ATLAS detector*, **JINST** **8** (2013) P07015, [arXiv:1305.2284 \[hep-ex\]](#).
- [137] ATLAS Collaboration, *Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in 13 TeV pp collisions with the ATLAS detector*, ATLAS-CONF-2017-026, CERN, Geneva, Apr, 2017.

- [138] 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.* **D94** (2016) no. 11, 113003, [arXiv:1605.02742 \[hep-ph\]](#).
- [139] ATLAS Collaboration, G. Aad et al., *Search for pair-produced long-lived neutral particles decaying in the ATLAS hadronic calorimeter in pp collisions at $\sqrt{s} = 8$ TeV*, *Phys. Lett.* **B743** (2015) 15–34, [arXiv:1501.04020 \[hep-ex\]](#).
- [140] CMS Collaboration, C. Collaboration, *Search for long-lived neutral particles in the final state of delayed photons and missing energy in proton-proton collisions at $\sqrt{s} = 8$ TeV*, .
- [141] ATLAS Collaboration, G. Aad et al., *Search for nonpointing and delayed photons in the diphoton and missing transverse momentum final state in 8 TeV pp collisions at the LHC using the ATLAS detector*, *Phys. Rev.* **D90** (2014) no. 11, 112005, [arXiv:1409.5542 \[hep-ex\]](#).
- [142] M. Chemtob, *Phenomenological constraints on broken R parity symmetry in supersymmetry models*, *Prog. Part. Nucl. Phys.* **54** (2005) 71–191, [arXiv:hep-ph/0406029 \[hep-ph\]](#).
- [143] R. Barbier et al., *R -parity violating supersymmetry*, *Phys. Rept.* **420** (2005) 1–202, [arXiv:hep-ph/0406039 \[hep-ph\]](#).
- [144] V. Barger, G. F. Giudice, and T. Han, *Some new aspects of supersymmetry R -parity violating interactions*, *Phys. Rev. D* **40** (Nov, 1989) 2987–2996.
- [145] H. K. Dreiner, *An Introduction to explicit R -parity violation*, [arXiv:hep-ph/9707435 \[hep-ph\]](#). [Adv. Ser. Direct. High Energy Phys.21,565(2010)].
- [146] C. Csaki, E. Kuflik, and T. Volansky, *Dynamical R -Parity Violation*, *Phys. Rev. Lett.* **112** (2014) 131801, [arXiv:1309.5957 \[hep-ph\]](#).
- [147] C. Csaki, E. Kuflik, O. Slone, and T. Volansky, *Models of Dynamical R -Parity Violation*, *JHEP* **06** (2015) 045, [arXiv:1502.03096 \[hep-ph\]](#).
- [148] K. Barry, P. W. Graham, and S. Rajendran, *Displaced vertices from R -parity violation and baryogenesis*, *Phys. Rev.* **D89** (2014) no. 5, 054003, [arXiv:1310.3853 \[hep-ph\]](#).
- [149] C. Csaki, E. Kuflik, S. Lombardo, O. Slone, and T. Volansky, *Phenomenology of a Long-Lived LSP with R -Parity Violation*, *JHEP* **08** (2015) 016, [arXiv:1505.00784 \[hep-ph\]](#).
- [150] Z. Liu and B. Tweedie, *The Fate of Long-Lived Superparticles with Hadronic Decays after LHC Run 1*, *JHEP* **06** (2015) 042, [arXiv:1503.05923 \[hep-ph\]](#).
- [151] ATLAS Collaboration, M. Aaboud et al., *Search for metastable heavy charged particles with large ionization energy loss in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS experiment*, *Phys. Rev.* **D93** (2016) no. 11, 112015, [arXiv:1604.04520 \[hep-ex\]](#).
- [152] G. F. Giudice and R. Rattazzi, *Theories with gauge mediated supersymmetry breaking*, *Phys. Rept.* **322** (1999) 419–499, [arXiv:hep-ph/9801271 \[hep-ph\]](#).
- [153] S. Knapen and D. Redigolo, *Gauge Mediation at LHC: status and prospects*, *JHEP* **01** (2017) 135, [arXiv:1606.07501 \[hep-ph\]](#).
- [154] P. Meade, N. Seiberg, and D. Shih, *General Gauge Mediation*, *Prog. Theor. Phys. Suppl.* **177** (2009) 143–158, [arXiv:0801.3278 \[hep-ph\]](#).
- [155] Z. Chacko and E. Ponton, *Yukawa deflected gauge mediation*, *Phys. Rev.* **D66** (2002) 095004, [arXiv:hep-ph/0112190 \[hep-ph\]](#).
- [156] J. A. Evans and D. Shih, *Surveying Extended GMSB Models with $m_h = 125$ GeV*, *JHEP* **08** (2013) 093, [arXiv:1303.0228 \[hep-ph\]](#).
- [157] S. Ambrosanio, G. L. Kane, G. D. Kribs, S. P. Martin, and S. Mrenna, *Search for supersymmetry with a light gravitino at the Fermilab Tevatron and CERN LEP colliders*, *Phys. Rev.* **D54** (1996) 5395–5411, [arXiv:hep-ph/9605398 \[hep-ph\]](#).
- [158] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, *An Introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191**

- (2015) 159–177, [arXiv:1410.3012 \[hep-ph\]](#).
- [159] M. Low and L.-T. Wang, *Neutralino dark matter at 14 TeV and 100 TeV*, *JHEP* **08** (2014) 161, [arXiv:1404.0682 \[hep-ph\]](#).
- [160] A. Ismail, E. Izaguirre, and B. Shuve, *Illuminating New Electroweak States at Hadron Colliders*, *Phys. Rev.* **D94** (2016) no. 1, 015001, [arXiv:1605.00658 \[hep-ph\]](#).
- [161] K. Kawagoe, T. Kobayashi, M. M. Nojiri, and A. Ochi, *Study of the gauge mediation signal with nonpointing photons at the CERN LHC*, *Phys. Rev.* **D69** (2004) 035003, [arXiv:hep-ph/0309031 \[hep-ph\]](#).
- [162] P. Meade, M. Reece, and D. Shih, *Long-Lived Neutralino NLSPs*, *JHEP* **10** (2010) 067, [arXiv:1006.4575 \[hep-ph\]](#).
- [163] M. Park and Y. Zhao, *Recovering Particle Masses from Missing Energy Signatures with Displaced Tracks*, [arXiv:1110.1403 \[hep-ph\]](#).
- [164] B. Batell, G. F. Giudice, and M. McCullough, *Natural Heavy Supersymmetry*, *JHEP* **12** (2015) 162, [arXiv:1509.00834 \[hep-ph\]](#).
- [165] J. D. Wells, *PeV-scale supersymmetry*, *Phys. Rev.* **D71** (2005) 015013, [arXiv:hep-ph/0411041 \[hep-ph\]](#).
- [166] N. Arkani-Hamed and S. Dimopoulos, *Supersymmetric unification without low energy supersymmetry and signatures for fine-tuning at the LHC*, *JHEP* **06** (2005) 073, [arXiv:hep-th/0405159 \[hep-th\]](#).
- [167] G. F. Giudice and A. Romanino, *Split supersymmetry*, *Nucl. Phys.* **B699** (2004) 65–89, [arXiv:hep-ph/0406088 \[hep-ph\]](#). [Erratum: *Nucl. Phys.*B706,487(2005)].
- [168] N. Arkani-Hamed, S. Dimopoulos, G. F. Giudice, and A. Romanino, *Aspects of split supersymmetry*, *Nucl. Phys.* **B709** (2005) 3–46, [arXiv:hep-ph/0409232 \[hep-ph\]](#).
- [169] L. Randall and R. Sundrum, *Out of this world supersymmetry breaking*, *Nucl. Phys.* **B557** (1999) 79–118, [arXiv:hep-th/9810155 \[hep-th\]](#).
- [170] G. F. Giudice, M. A. Luty, H. Murayama, and R. Rattazzi, *Gaugino mass without singlets*, *JHEP* **12** (1998) 027, [arXiv:hep-ph/9810442 \[hep-ph\]](#).
- [171] M. Ibe, S. Matsumoto, and T. T. Yanagida, *Pure Gravity Mediation with $m_{3/2} = 10\text{--}100\text{TeV}$* , *Phys. Rev.* **D85** (2012) 095011, [arXiv:1202.2253 \[hep-ph\]](#).
- [172] B. S. Acharya, G. Kane, and P. Kumar, *Compactified String Theories – Generic Predictions for Particle Physics*, *Int. J. Mod. Phys.* **A27** (2012) 1230012, [arXiv:1204.2795 \[hep-ph\]](#).
- [173] T. Cohen, N. Craig, and S. Knapen, *Gauge Mediated Mini-Split*, *JHEP* **03** (2016) 103, [arXiv:1512.06128 \[hep-ph\]](#).
- [174] A. Arvanitaki, C. Davis, P. W. Graham, and J. G. Wacker, *One loop predictions of the finely tuned SSM*, *Phys. Rev.* **D70** (2004) 117703, [arXiv:hep-ph/0406034 \[hep-ph\]](#).
- [175] A. Arvanitaki, N. Craig, S. Dimopoulos, and G. Villadoro, *Mini-Split*, *JHEP* **02** (2013) 126, [arXiv:1210.0555 \[hep-ph\]](#).
- [176] G. F. Giudice and A. Strumia, *Probing High-Scale and Split Supersymmetry with Higgs Mass Measurements*, *Nucl. Phys.* **B858** (2012) 63–83, [arXiv:1108.6077 \[hep-ph\]](#).
- [177] N. Arkani-Hamed, A. Gupta, D. E. Kaplan, N. Weiner, and T. Zorawski, *Simply Unnatural Supersymmetry*, [arXiv:1212.6971 \[hep-ph\]](#).
- [178] E. Bagnaschi, G. F. Giudice, P. Slavich, and A. Strumia, *Higgs Mass and Unnatural Supersymmetry*, *JHEP* **09** (2014) 092, [arXiv:1407.4081 \[hep-ph\]](#).
- [179] J. Pardo Vega and G. Villadoro, *SusyHD: Higgs mass Determination in Supersymmetry*, *JHEP* **07** (2015) 159, [arXiv:1504.05200 \[hep-ph\]](#).
- [180] G. F. Giudice and A. Masiero, *A Natural Solution to the mu Problem in Supergravity Theories*, *Phys. Lett.* **B206** (1988) 480–484.

- [181] P. Gambino, G. F. Giudice, and P. Slavich, *Gluino decays in split supersymmetry*, *Nucl. Phys.* **B726** (2005) 35–52, [arXiv:hep-ph/0506214](#) [hep-ph].
- [182] C. Han, L. Wu, J. M. Yang, M. Zhang, and Y. Zhang, *New approach for detecting a compressed bino/wino at the LHC*, *Phys. Rev.* **D91** (2015) 055030, [arXiv:1409.4533](#) [hep-ph].
- [183] N. Nagata, H. Otono, and S. Shirai, *Probing Bino-Wino Coannihilation at the LHC*, *JHEP* **10** (2015) 086, [arXiv:1506.08206](#) [hep-ph].
- [184] J. Fan, M. Reece, and J. T. Ruderman, *Stealth Supersymmetry*, *JHEP* **11** (2011) 012, [arXiv:1105.5135](#) [hep-ph].
- [185] N. Zwane, *Long-Lived Particle Searches in R-Parity Violating MSSM*, *J. Phys.* **G44** (2017) no. 10, 105003, [arXiv:1505.03479](#) [hep-ph].
- [186] A. Hook and H. Murayama, *Low-energy Supersymmetry Breaking Without the Gravitino Problem*, *Phys. Rev.* **D92** (2015) no. 1, 015004, [arXiv:1503.04880](#) [hep-ph].
- [187] M. Ibe and T. T. Yanagida, *Lower limit on the gravitino mass in low-scale gauge mediation with $m_H \simeq 125$ GeV*, *Phys. Lett.* **B764** (2017) 260–264, [arXiv:1608.01610](#) [hep-ph].
- [188] K. Osato, T. Sekiguchi, M. Shirasaki, A. Kamada, and N. Yoshida, *Cosmological Constraint on the Light Gravitino Mass from CMB Lensing and Cosmic Shear*, *JCAP* **1606** (2016) no. 06, 004, [arXiv:1601.07386](#) [astro-ph.CO].
- [189] J. Fan, M. Reece, and J. T. Ruderman, *A Stealth Supersymmetry Sampler*, *JHEP* **07** (2012) 196, [arXiv:1201.4875](#) [hep-ph].
- [190] J. Fan, R. Krall, D. Pinner, M. Reece, and J. T. Ruderman, *Stealth Supersymmetry Simplified*, *JHEP* **07** (2016) 016, [arXiv:1512.05781](#) [hep-ph].
- [191] ATLAS Collaboration, T. A. collaboration, *Search for massive supersymmetric particles in multi-jet final states produced in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, .
- [192] T. Higaki and R. Kitano, *On Supersymmetric Effective Theories of Axion*, *Phys. Rev.* **D86** (2012) 075027, [arXiv:1104.0170](#) [hep-ph].
- [193] C. Cheung, G. Elor, and L. J. Hall, *The Cosmological Axino Problem*, *Phys. Rev.* **D85** (2012) 015008, [arXiv:1104.0692](#) [hep-ph].
- [194] G. Salam and A. Weiler, *Collider Reach ^{β} : Quick (and dirty) estimates for hadron machines. (Talk at “BSM physics opportunities at 100 TeV Workshop”, 10-11 Feb 2014)*, .
- [195] J. E. Kim and G. Carosi, *Axions and the Strong CP Problem*, *Rev. Mod. Phys.* **82** (2010) 557–602, [arXiv:0807.3125](#) [hep-ph].
- [196] R. D. Peccei and H. R. Quinn, *Constraints Imposed by CP Conservation in the Presence of Instantons*, *Phys. Rev.* **D16** (1977) 1791–1797.
- [197] F. Wilczek, *Problem of Strong p and t Invariance in the Presence of Instantons*, *Phys. Rev. Lett.* **40** (1978) 279–282.
- [198] S. Weinberg, *A New Light Boson?*, *Phys. Rev. Lett.* **40** (1978) 223–226.
- [199] J. E. Kim, *Weak Interaction Singlet and Strong CP Invariance*, *Phys. Rev. Lett.* **43** (1979) 103.
- [200] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, *Can Confinement Ensure Natural CP Invariance of Strong Interactions?*, *Nucl. Phys.* **B166** (1980) 493–506.
- [201] M. Dine, W. Fischler, and M. Srednicki, *A Simple Solution to the Strong CP Problem with a Harmless Axion*, *Phys. Lett.* **104B** (1981) 199–202.
- [202] A. R. Zhitnitsky, *On Possible Suppression of the Axion Hadron Interactions. (In Russian)*, *Sov. J. Nucl. Phys.* **31** (1980) 260. [*Yad. Fiz.*31,497(1980)].
- [203] G. G. Raffelt, *Astrophysical axion bounds*, *Lect. Notes Phys.* **741** (2008) 51–71, [arXiv:hep-ph/0611350](#) [hep-ph]. [,51(2006)].
- [204] J. H. Chang, R. Essig, and S. D. McDermott, *Supernova 1987A Constraints on Sub-GeV Dark*

- Sectors, Millicharged Particles, the QCD Axion, and an Axion-like Particle*, [arXiv:1803.00993 \[hep-ph\]](#).
- [205] P. Fox, A. Pierce, and S. D. Thomas, *Probing a QCD string axion with precision cosmological measurements*, [arXiv:hep-th/0409059 \[hep-th\]](#).
- [206] R. Kitano and Y. Nomura, *Supersymmetry, naturalness, and signatures at the LHC*, *Phys. Rev. D* **73** (2006) 095004, [arXiv:hep-ph/0602096 \[hep-ph\]](#).
- [207] C. Brust, A. Katz, S. Lawrence, and R. Sundrum, *SUSY, the Third Generation and the LHC*, *JHEP* **03** (2012) 103, [arXiv:1110.6670 \[hep-ph\]](#).
- [208] M. Papucci, J. T. Ruderman, and A. Weiler, *Natural SUSY Endures*, *JHEP* **09** (2012) 035, [arXiv:1110.6926 \[hep-ph\]](#).
- [209] H. Baer, V. Barger, P. Huang, A. Mustafayev, and X. Tata, *Radiative natural SUSY with a 125 GeV Higgs boson*, *Phys. Rev. Lett.* **109** (2012) 161802, [arXiv:1207.3343 \[hep-ph\]](#).
- [210] J. E. Kim and H. P. Nilles, *The mu Problem and the Strong CP Problem*, *Phys. Lett.* **138B** (1984) 150–154.
- [211] E. J. Chun, J. E. Kim, and H. P. Nilles, *A Natural solution of the mu problem with a composite axion in the hidden sector*, *Nucl. Phys.* **B370** (1992) 105–122.
- [212] T. Goto and M. Yamaguchi, *Is axino dark matter possible in supergravity?*, *Phys. Lett.* **B276** (1992) 103–107.
- [213] E. J. Chun, J. E. Kim, and H. P. Nilles, *Axino mass*, *Phys. Lett.* **B287** (1992) 123–127, [arXiv:hep-ph/9205229 \[hep-ph\]](#).
- [214] E. J. Chun and A. Lukas, *Axino mass in supergravity models*, *Phys. Lett.* **B357** (1995) 43–50, [arXiv:hep-ph/9503233 \[hep-ph\]](#).
- [215] E. J. Chun, *Dark matter in the Kim-Nilles mechanism*, *Phys. Rev.* **D84** (2011) 043509, [arXiv:1104.2219 \[hep-ph\]](#).
- [216] K. J. Bae, K. Choi, and S. H. Im, *Effective Interactions of Axion Supermultiplet and Thermal Production of Axino Dark Matter*, *JHEP* **08** (2011) 065, [arXiv:1106.2452 \[hep-ph\]](#).
- [217] K. J. Bae, E. J. Chun, and S. H. Im, *Cosmology of the DFSZ axino*, *JCAP* **1203** (2012) 013, [arXiv:1111.5962 \[hep-ph\]](#).
- [218] G. Barenboim, E. J. Chun, S. Jung, and W. I. Park, *Implications of an axino LSP for naturalness*, *Phys. Rev.* **D90** (2014) no. 3, 035020, [arXiv:1407.1218 \[hep-ph\]](#).
- [219] S. Alekhin et al., *A facility to Search for Hidden Particles at the CERN SPS: the SHiP physics case*, *Rept. Prog. Phys.* **79** (2016) no. 12, 124201, [arXiv:1504.04855 \[hep-ph\]](#).
- [220] E. Cremmer, B. Julia, J. Scherk, P. van Nieuwenhuizen, S. Ferrara, and L. Girardello, *Super-higgs effect in supergravity with general scalar interactions*, *Phys. Lett.* **79B** (1978) 231–234.
- [221] J. R. Ellis, A. B. Lahanas, D. V. Nanopoulos, and K. Tamvakis, *No-Scale Supersymmetric Standard Model*, *Phys. Lett.* **134B** (1984) 429.
- [222] S. L. Dubovsky, D. S. Gorbunov, and S. V. Troitsky, *Gauge mechanism of mediation of supersymmetry breaking*, *Phys. Usp.* **42** (1999) 623–651, [arXiv:hep-ph/9905466 \[hep-ph\]](#). [*Usp. Fiz. Nauk*169,705(1999)].
- [223] A. Brignole, F. Feruglio, and F. Zwirner, *Aspects of spontaneously broken $N=1$ global supersymmetry in the presence of gauge interactions*, *Nucl. Phys.* **B501** (1997) 332–374, [arXiv:hep-ph/9703286 \[hep-ph\]](#).
- [224] D. S. Gorbunov, *Light sgoldstino: Precision measurements versus collider searches*, *Nucl. Phys.* **B602** (2001) 213–237, [arXiv:hep-ph/0007325 \[hep-ph\]](#).
- [225] D. S. Gorbunov and A. V. Semenov, *CompHEP package with light gravitino and sgoldstinos*, [arXiv:hep-ph/0111291 \[hep-ph\]](#).
- [226] E. Perazzi, G. Ridolfi, and F. Zwirner, *Signatures of massive sgoldstinos at e^+e^- colliders*, *Nucl.*

- Phys. **B574** (2000) 3–22, [arXiv:hep-ph/0001025 \[hep-ph\]](#).
- [227] E. Perazzi, G. Ridolfi, and F. Zwirner, *Signatures of massive sgoldstinos at hadron colliders*, Nucl. Phys. **B590** (2000) 287–305, [arXiv:hep-ph/0005076 \[hep-ph\]](#).
- [228] S. V. Demidov and D. S. Gorbunov, *Flavor violating processes with sgoldstino pair production*, Phys. Rev. **D85** (2012) 077701, [arXiv:1112.5230 \[hep-ph\]](#).
- [229] D. Gorbunov, V. Ilyin, and B. Mele, *Sgoldstino events in top decays at LHC*, Phys. Lett. **B502** (2001) 181–188, [arXiv:hep-ph/0012150 \[hep-ph\]](#).
- [230] D. S. Gorbunov and N. V. Krasnikov, *Prospects for sgoldstino search at the LHC*, JHEP **07** (2002) 043, [arXiv:hep-ph/0203078 \[hep-ph\]](#).
- [231] S. V. Demidov and D. S. Gorbunov, *LHC prospects in searches for neutral scalars in $pp \rightarrow \gamma\gamma + \text{jet}$: SM Higgs boson, radion, sgoldstino*, Phys. Atom. Nucl. **69** (2006) 712–720, [arXiv:hep-ph/0405213 \[hep-ph\]](#).
- [232] M. Asano and R. Garani, *Sgoldstino search at the LHC*, [arXiv:1701.00829 \[hep-ph\]](#).
- [233] K. O. Astapov and S. V. Demidov, *Sgoldstino-Higgs mixing in models with low-scale supersymmetry breaking*, JHEP **01** (2015) 136, [arXiv:1411.6222 \[hep-ph\]](#).
- [234] K. O. Astapov and D. S. Gorbunov, *Decaying light particles in the SHiP experiment. III. Signal rate estimates for scalar and pseudoscalar sgoldstinos*, Phys. Rev. **D93** (2016) no. 3, 035008, [arXiv:1511.05403 \[hep-ph\]](#).
- [235] Z. Chacko, H.-S. Goh, and R. Harnik, *The Twin Higgs: Natural electroweak breaking from mirror symmetry*, Phys. Rev. Lett. **96** (2006) 231802, [arXiv:hep-ph/0506256 \[hep-ph\]](#).
- [236] N. Craig, S. Knapen, and P. Longhi, *Neutral Naturalness from Orbifold Higgs Models*, Phys. Rev. Lett. **114** (2015) no. 6, 061803, [arXiv:1410.6808 \[hep-ph\]](#).
- [237] H. Cai, H.-C. Cheng, and J. Terning, *A Quirky Little Higgs Model*, JHEP **05** (2009) 045, [arXiv:0812.0843 \[hep-ph\]](#).
- [238] G. Burdman, Z. Chacko, H.-S. Goh, and R. Harnik, *Folded supersymmetry and the LEP paradox*, JHEP **02** (2007) 009, [arXiv:hep-ph/0609152 \[hep-ph\]](#).
- [239] N. Craig, A. Katz, M. Strassler, and R. Sundrum, *Naturalness in the Dark at the LHC*, JHEP **07** (2015) 105, [arXiv:1501.05310 \[hep-ph\]](#).
- [240] D. Curtin and C. B. Verhaaren, *Discovering Uncolored Naturalness in Exotic Higgs Decays*, JHEP **12** (2015) 072, [arXiv:1506.06141 \[hep-ph\]](#).
- [241] J. E. Juknevich, *Pure-gluon hidden valleys through the Higgs portal*, JHEP **08** (2010) 121, [arXiv:0911.5616 \[hep-ph\]](#).
- [242] Z. Chacko, D. Curtin, and C. B. Verhaaren, *A Quirky Probe of Neutral Naturalness*, Phys. Rev. **D94** (2016) no. 1, 011504, [arXiv:1512.05782 \[hep-ph\]](#).
- [243] J. E. Juknevich, *Phenomenology of pure-gauge hidden valleys at Hadron colliders*. PhD thesis, Rutgers U., Piscataway, 2010.
- [244] G. W. Van Apeldoorn et al., *Thermal Emission of Pions in Anti-proton Proton Annihilation at 12-GeV/c*, Z. Phys. **C7** (1981) 235–239.
- [245] S. Dawson et al., *Working Group Report: Higgs Boson*, in *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*.
- [246] Z. Liu, L.-T. Wang, and H. Zhang, *Exotic decays of the 125 GeV Higgs boson at future e^+e^- lepton colliders*, Chin. Phys. **C41** (2017) no. 6, 063102, [arXiv:1612.09284 \[hep-ph\]](#).
- [247] J. Kang and M. A. Luty, *Macroscopic Strings and 'Quirks' at Colliders*, JHEP **11** (2009) 065, [arXiv:0805.4642 \[hep-ph\]](#).
- [248] G. Burdman, Z. Chacko, R. Harnik, L. de Lima, and C. B. Verhaaren, *Colorless Top Partners, a 125 GeV Higgs, and the Limits on Naturalness*, Phys. Rev. **D91** (2015) no. 5, 055007,

- [arXiv:1411.3310 \[hep-ph\]](#).
- [249] H.-C. Cheng, S. Jung, E. Salvioni, and Y. Tsai, *Exotic Quarks in Twin Higgs Models*, *JHEP* **03** (2016) 074, [arXiv:1512.02647 \[hep-ph\]](#).
- [250] I. Garcia Garcia, R. Lasenby, and J. March-Russell, *Twin Higgs Asymmetric Dark Matter*, *Phys. Rev. Lett.* **115** (2015) no. 12, 121801, [arXiv:1505.07410 \[hep-ph\]](#).
- [251] M. Farina, *Asymmetric Twin Dark Matter*, *JCAP* **1511** (2015) no. 11, 017, [arXiv:1506.03520 \[hep-ph\]](#).
- [252] A. Pierce, B. Shakya, Y. Tsai, and Y. Zhao, *Searching for Confining Hidden Valleys at the LHC(b)*, [arXiv:1708.05389 \[hep-ph\]](#).
- [253] S. Chang, L. J. Hall, and N. Weiner, *A Supersymmetric twin Higgs*, *Phys. Rev.* **D75** (2007) 035009, [arXiv:hep-ph/0604076 \[hep-ph\]](#).
- [254] N. Craig and H. K. Lou, *Scherk-Schwarz Supersymmetry Breaking in 4D*, *JHEP* **12** (2014) 184, [arXiv:1406.4880 \[hep-ph\]](#).
- [255] N. Craig and K. Howe, *Doubling down on naturalness with a supersymmetric twin Higgs*, *JHEP* **03** (2014) 140, [arXiv:1312.1341 \[hep-ph\]](#).
- [256] M. Geller and O. Telem, *Holographic Twin Higgs Model*, *Phys. Rev. Lett.* **114** (2015) 191801, [arXiv:1411.2974 \[hep-ph\]](#).
- [257] M. Low, A. Tesi, and L.-T. Wang, *Twin Higgs mechanism and a composite Higgs boson*, *Phys. Rev.* **D91** (2015) 095012, [arXiv:1501.07890 \[hep-ph\]](#).
- [258] R. Barbieri, D. Greco, R. Rattazzi, and A. Wulzer, *The Composite Twin Higgs scenario*, *JHEP* **08** (2015) 161, [arXiv:1501.07803 \[hep-ph\]](#).
- [259] P. Batra and Z. Chacko, *A Composite Twin Higgs Model*, *Phys. Rev.* **D79** (2009) 095012, [arXiv:0811.0394 \[hep-ph\]](#).
- [260] N. Craig, S. Knapen, and P. Longhi, *The Orbifold Higgs*, *JHEP* **03** (2015) 106, [arXiv:1411.7393 \[hep-ph\]](#).
- [261] H.-C. Cheng, E. Salvioni, and Y. Tsai, *Exotic electroweak signals in the twin Higgs model*, *Phys. Rev.* **D95** (2017) no. 11, 115035, [arXiv:1612.03176 \[hep-ph\]](#).
- [262] L. Li, E. Salvioni, Y. Tsai, and R. Zheng, *Electroweak-Charged Bound States as LHC Probes of Hidden Forces*, *Phys. Rev.* **D97** (2018) no. 1, 015010, [arXiv:1710.06437 \[hep-ph\]](#).
- [263] R. Barbieri, L. J. Hall, and K. Harigaya, *Minimal Mirror Twin Higgs*, *JHEP* **11** (2016) 172, [arXiv:1609.05589 \[hep-ph\]](#).
- [264] Z. Chacko, N. Craig, P. J. Fox, and R. Harnik, *Cosmology in Mirror Twin Higgs and Neutrino Masses*, *JHEP* **07** (2017) 023, [arXiv:1611.07975 \[hep-ph\]](#).
- [265] N. Craig, S. Koren, and T. Trott, *Cosmological Signals of a Mirror Twin Higgs*, [arXiv:1611.07977 \[hep-ph\]](#).
- [266] G. Panico and A. Wulzer, *The Composite Nambu-Goldstone Higgs*, *Lect. Notes Phys.* **913** (2016) pp.1–316, [arXiv:1506.01961 \[hep-ph\]](#).
- [267] K. Agashe, R. Contino, and R. Sundrum, *Top compositeness and precision unification*, *Phys. Rev. Lett.* **95** (2005) 171804, [arXiv:hep-ph/0502222 \[hep-ph\]](#).
- [268] J. Barnard, T. Gherghetta, T. S. Ray, and A. Spray, *The Unnatural Composite Higgs*, *JHEP* **01** (2015) 067, [arXiv:1409.7391 \[hep-ph\]](#).
- [269] J. Barnard, P. Cox, T. Gherghetta, and A. Spray, *Long-Lived, Colour-Triplet Scalars from Unnaturalness*, *JHEP* **03** (2016) 003, [arXiv:1510.06405 \[hep-ph\]](#).
- [270] B. Fuks, M. Klasen, D. R. Lamprea, and M. Rothering, *Gaugino production in proton-proton collisions at a center-of-mass energy of 8 TeV*, *JHEP* **10** (2012) 081, [arXiv:1207.2159 \[hep-ph\]](#).
- [271] B. Fuks, M. Klasen, D. R. Lamprea, and M. Rothering, *Precision predictions for electroweak*

- superpartner production at hadron colliders with RESUMMINO*, *Eur. Phys. J. C* **73** (2013) 2480, [arXiv:1304.0790 \[hep-ph\]](#).
- [272] J. Barnard, T. Gherghetta, and T. S. Ray, *UV descriptions of composite Higgs models without elementary scalars*, *JHEP* **02** (2014) 002, [arXiv:1311.6562 \[hep-ph\]](#).
- [273] G. Ferretti and D. Karateev, *Fermionic UV completions of Composite Higgs models*, *JHEP* **03** (2014) 077, [arXiv:1312.5330 \[hep-ph\]](#).
- [274] G. Cacciapaglia, H. Cai, A. Deandrea, T. Flacke, S. J. Lee, and A. Parolini, *Composite scalars at the LHC: the Higgs, the Sextet and the Octet*, *JHEP* **11** (2015) 201, [arXiv:1507.02283 \[hep-ph\]](#).
- [275] P. W. Graham, D. E. Kaplan, and S. Rajendran, *Cosmological Relaxation of the Electroweak Scale*, *Phys. Rev. Lett.* **115** (2015) no. 22, 221801, [arXiv:1504.07551 \[hep-ph\]](#).
- [276] J. R. Espinosa, C. Grojean, G. Panico, A. Pomarol, O. Pujolàs, and G. Servant, *Cosmological Higgs-Axion Interplay for a Naturally Small Electroweak Scale*, *Phys. Rev. Lett.* **115** (2015) no. 25, 251803, [arXiv:1506.09217 \[hep-ph\]](#).
- [277] J. L. Evans, T. Gherghetta, N. Nagata, and Z. Thomas, *Naturalizing Supersymmetry with a Two-Field Relaxion Mechanism*, *JHEP* **09** (2016) 150, [arXiv:1602.04812 \[hep-ph\]](#).
- [278] J. L. Evans, T. Gherghetta, N. Nagata, and M. Peloso, *Low-scale D-term inflation and the relaxion mechanism*, *Phys. Rev.* **D95** (2017) no. 11, 115027, [arXiv:1704.03695 \[hep-ph\]](#).
- [279] K. Choi and S. H. Im, *Constraints on Relaxion Windows*, *JHEP* **12** (2016) 093, [arXiv:1610.00680 \[hep-ph\]](#).
- [280] T. Flacke, C. Frugiuele, E. Fuchs, R. S. Gupta, and G. Perez, *Phenomenology of relaxion-Higgs mixing*, *JHEP* **06** (2017) 050, [arXiv:1610.02025 \[hep-ph\]](#).
- [281] K. Choi, H. Kim, and T. Sekiguchi, *Dynamics of cosmological relaxion after reheating*, [arXiv:1611.08569 \[hep-ph\]](#).
- [282] M. Pepe Altarelli and F. Teubert, *B Physics at LHCb*, *Int. J. Mod. Phys.* **A23** (2008) 5117–5136, [arXiv:0802.1901 \[hep-ph\]](#).
- [283] J. D. Clarke, R. Foot, and R. R. Volkas, *Phenomenology of a very light scalar ($100 \text{ MeV} < m_h < 10 \text{ GeV}$) mixing with the SM Higgs*, *JHEP* **02** (2014) 123, [arXiv:1310.8042 \[hep-ph\]](#).
- [284] Planck Collaboration, P. A. R. Ade et al., *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13, [arXiv:1502.01589 \[astro-ph.CO\]](#).
- [285] D. Pappadopulo, J. T. Ruderman, and G. Trevisan, *Dark matter freeze-out in a nonrelativistic sector*, *Phys. Rev.* **D94** (2016) no. 3, 035005, [arXiv:1602.04219 \[hep-ph\]](#).
- [286] A. Berlin, D. Hooper, and G. Krnjaic, *Thermal Dark Matter From A Highly Decoupled Sector*, *Phys. Rev.* **D94** (2016) no. 9, 095019, [arXiv:1609.02555 \[hep-ph\]](#).
- [287] N. Bernal, M. Heikinheimo, T. Tenkanen, K. Tuominen, and V. Vaskonen, *The Dawn of FIMP Dark Matter: A Review of Models and Constraints*, *Int. J. Mod. Phys.* **A32** (2017) no. 27, 1730023, [arXiv:1706.07442 \[hep-ph\]](#).
- [288] F. Kahlhoefer, *On the LHC sensitivity for non-thermalised hidden sectors*, *Phys. Lett.* **B779** (2018) 388–392, [arXiv:1801.07621 \[hep-ph\]](#).
- [289] M. Farina and D. Pappadopulo, *Private Communication*, .
- [290] K. Griest and D. Seckel, *Three exceptions in the calculation of relic abundances*, *Phys. Rev.* **D43** (1991) 3191–3203.
- [291] E. Izaguirre, G. Krnjaic, and B. Shuve, *Discovering Inelastic Thermal-Relic Dark Matter at Colliders*, *Phys. Rev.* **D93** (2016) no. 6, 063523, [arXiv:1508.03050 \[hep-ph\]](#).
- [292] M. J. Baker et al., *The Coannihilation Codex*, *JHEP* **12** (2015) 120, [arXiv:1510.03434 \[hep-ph\]](#).
- [293] M. Buschmann, S. El Hedri, A. Kaminska, J. Liu, M. de Vries, X.-P. Wang, F. Yu, and J. Zurita,

- Hunting for dark matter coannihilation by mixing dijet resonances and missing transverse energy*, *JHEP* **09** (2016) 033, [arXiv:1605.08056 \[hep-ph\]](#).
- [294] C. Cheung, L. J. Hall, D. Pinner, and J. T. Ruderman, *Prospects and Blind Spots for Neutralino Dark Matter*, *JHEP* **05** (2013) 100, [arXiv:1211.4873 \[hep-ph\]](#).
- [295] K. M. Zurek, *Asymmetric Dark Matter: Theories, Signatures, and Constraints*, *Phys. Rept.* **537** (2014) 91–121, [arXiv:1308.0338 \[hep-ph\]](#).
- [296] D. E. Kaplan, M. A. Luty, and K. M. Zurek, *Asymmetric Dark Matter*, *Phys. Rev.* **D79** (2009) 115016, [arXiv:0901.4117 \[hep-ph\]](#).
- [297] I.-W. Kim and K. M. Zurek, *Flavor and Collider Signatures of Asymmetric Dark Matter*, *Phys. Rev.* **D89** (2014) no. 3, 035008, [arXiv:1310.2617 \[hep-ph\]](#).
- [298] L. J. Hall, K. Jedamzik, J. March-Russell, and S. M. West, *Freeze-In Production of FIMP Dark Matter*, *JHEP* **03** (2010) 080, [arXiv:0911.1120 \[hep-ph\]](#).
- [299] F. D’Eramo, N. Fernandez, and S. Profumo, *Dark Matter Freeze-in Production in Fast-Expanding Universes*, *JCAP* **1802** (2018) no. 02, 046, [arXiv:1712.07453 \[hep-ph\]](#).
- [300] R. T. Co, F. D’Eramo, L. J. Hall, and D. Pappadopulo, *Freeze-In Dark Matter with Displaced Signatures at Colliders*, *JCAP* **1512** (2015) no. 12, 024, [arXiv:1506.07532 \[hep-ph\]](#).
- [301] R. T. Co, F. D’Eramo, L. J. Hall, and K. Harigaya, *Saxion Cosmology for Thermalized Gravitino Dark Matter*, *JHEP* **07** (2017) 125, [arXiv:1703.09796 \[hep-ph\]](#).
- [302] L. Calibbi, A. Mariotti, and P. Tziveloglou, *Singlet-Doublet Model: Dark matter searches and LHC constraints*, *JHEP* **10** (2015) 116, [arXiv:1505.03867 \[hep-ph\]](#).
- [303] M. Cirelli, N. Fornengo, and A. Strumia, *Minimal dark matter*, *Nucl. Phys.* **B753** (2006) 178–194, [arXiv:hep-ph/0512090 \[hep-ph\]](#).
- [304] M. Cirelli, F. Sala, and M. Taoso, *Wino-like Minimal Dark Matter and future colliders*, *JHEP* **10** (2014) 033, [arXiv:1407.7058 \[hep-ph\]](#). [Erratum: *JHEP*01,041(2015)].
- [305] R. Mahbubani, P. Schwaller, and J. Zurita, *Closing the window for compressed Dark Sectors with disappearing charged tracks*, *JHEP* **06** (2017) 119, [arXiv:1703.05327 \[hep-ph\]](#). [Erratum: *JHEP*10,061(2017)].
- [306] A. Alloul, N. D. Christensen, C. Degrande, C. Duhr, and B. Fuks, *FeynRules 2.0 - A complete toolbox for tree-level phenomenology*, *Comput. Phys. Commun.* **185** (2014) 2250–2300, [arXiv:1310.1921 \[hep-ph\]](#).
- [307] G. Bélanger, F. Boudjema, A. Goudelis, A. Pukhov, and B. Zaldivar, *micrOMEGAs5.0 : freeze-in*, [arXiv:1801.03509 \[hep-ph\]](#).
- [308] R. T. Co, F. D’Eramo, and L. J. Hall, *Supersymmetric axion grand unified theories and their predictions*, *Phys. Rev.* **D94** (2016) no. 7, 075001, [arXiv:1603.04439 \[hep-ph\]](#).
- [309] Y. Hochberg, E. Kuflik, T. Volansky, and J. G. Wacker, *Mechanism for Thermal Relic Dark Matter of Strongly Interacting Massive Particles*, *Phys. Rev. Lett.* **113** (2014) 171301, [arXiv:1402.5143 \[hep-ph\]](#).
- [310] Y. Hochberg, E. Kuflik, H. Murayama, T. Volansky, and J. G. Wacker, *Model for Thermal Relic Dark Matter of Strongly Interacting Massive Particles*, *Phys. Rev. Lett.* **115** (2015) no. 2, 021301, [arXiv:1411.3727 \[hep-ph\]](#).
- [311] S.-M. Choi, Y. Hochberg, E. Kuflik, H. M. Lee, Y. Mambrini, H. Murayama, and M. Pierre, *Vector SIMP dark matter*, *JHEP* **10** (2017) 162, [arXiv:1707.01434 \[hep-ph\]](#).
- [312] S.-M. Choi, H. M. Lee, P. Ko, and A. Natale, *Unitarizing SIMP scenario with dark vector resonances*, [arXiv:1801.07726 \[hep-ph\]](#).
- [313] Y. Hochberg, E. Kuflik, and H. Murayama, *SIMP Spectroscopy*, *JHEP* **05** (2016) 090, [arXiv:1512.07917 \[hep-ph\]](#).
- [314] E. Kuflik, M. Perelstein, N. R.-L. Lorier, and Y.-D. Tsai, *Elastically Decoupling Dark Matter*,

- [Phys. Rev. Lett. **116** \(2016\) no. 22, 221302, arXiv:1512.04545 \[hep-ph\]](#).
- [315] J. A. Dror, E. Kuflik, and W. H. Ng, *Codecaying Dark Matter*, [Phys. Rev. Lett. **117** \(2016\) no. 21, 211801, arXiv:1607.03110 \[hep-ph\]](#).
- [316] M. Farina, D. Pappadopulo, J. T. Ruderman, and G. Trevisan, *Phases of Cannibal Dark Matter*, [JHEP **12** \(2016\) 039, arXiv:1607.03108 \[hep-ph\]](#).
- [317] S. Okawa, M. Tanabashi, and M. Yamanaka, *Relic Abundance in a Secluded Dark Matter Scenario with a Massive Mediator*, [Phys. Rev. **D95** \(2017\) no. 2, 023006, arXiv:1607.08520 \[hep-ph\]](#).
- [318] K. R. Dienes and B. Thomas, *Dynamical Dark Matter: I. Theoretical Overview*, [Phys. Rev. **D85** \(2012\) 083523, arXiv:1106.4546 \[hep-ph\]](#).
- [319] K. R. Dienes and B. Thomas, *Dynamical Dark Matter: II. An Explicit Model*, [Phys. Rev. **D85** \(2012\) 083524, arXiv:1107.0721 \[hep-ph\]](#).
- [320] K. R. Dienes and B. Thomas, *Phenomenological Constraints on Axion Models of Dynamical Dark Matter*, [Phys. Rev. **D86** \(2012\) 055013, arXiv:1203.1923 \[hep-ph\]](#).
- [321] K. R. Dienes and B. Thomas, *Dynamical Dark Matter: A new framework for dark-matter physics*, [AIP Conf. Proc. **1534** \(2012\) 57–77](#).
- [322] K. R. Dienes and B. Thomas, *Dynamical Dark Matter: Introduction, Equation of State, and Cosmological Implications*, [PoS **ICHEP2012** \(2013\) 452](#).
- [323] K. R. Dienes and B. Thomas, *Phenomenology of Dynamical Dark Matter*, [PoS **ICHEP2012** \(2013\) 460](#).
- [324] WMAP Collaboration, G. Hinshaw et al., *Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results*, [Astrophys. J. Suppl. **208** \(2013\) 19, arXiv:1212.5226 \[astro-ph.CO\]](#).
- [325] K. R. Dienes, F. Huang, S. Su, and B. Thomas, *Dynamical Dark Matter from Strongly-Coupled Dark Sectors*, [Phys. Rev. **D95** \(2017\) no. 4, 043526, arXiv:1610.04112 \[hep-ph\]](#).
- [326] K. R. Dienes, F. Huang, S. Su, and B. Thomas, *Regge trajectories and Hagedorn behavior: Hadronic realizations of dynamical dark matter*, [AIP Conf. Proc. **1900** \(2017\) no. 1, 040003](#).
- [327] K. R. Dienes, J. Fennick, J. Kumar, and B. Thomas, *Randomness in the Dark Sector: Emergent Mass Spectra and Dynamical Dark Matter Ensembles*, [Phys. Rev. **D93** \(2016\) no. 8, 083506, arXiv:1601.05094 \[hep-ph\]](#).
- [328] D. Chialva, P. S. B. Dev, and A. Mazumdar, *Multiple dark matter scenarios from ubiquitous stringy throats*, [Phys. Rev. **D87** \(2013\) no. 6, 063522, arXiv:1211.0250 \[hep-ph\]](#).
- [329] K. R. Dienes, J. Fennick, J. Kumar, and B. Thomas, *Dynamical Dark Matter from Thermal Freeze-Out*, [Phys. Rev. **D97** \(2018\) no. 6, 063522, arXiv:1712.09919 \[hep-ph\]](#).
- [330] K. R. Dienes, J. Kost, and B. Thomas, *A Tale of Two Timescales: Mixing, Mass Generation, and Phase Transitions in the Early Universe*, [Phys. Rev. **D93** \(2016\) no. 4, 043540, arXiv:1509.00470 \[hep-ph\]](#).
- [331] K. R. Dienes, J. Kost, and B. Thomas, *Kaluza-Klein Towers in the Early Universe: Phase Transitions, Relic Abundances, and Applications to Axion Cosmology*, [Phys. Rev. **D95** \(2017\) no. 12, 123539, arXiv:1612.08950 \[hep-ph\]](#).
- [332] K. R. Dienes, J. Kost, and B. Thomas, *A new approach to the cosmological moduli problem*, [AIP Conf. Proc. **1743** \(2016\) 020003](#).
- [333] K. R. Dienes, S. Su, and B. Thomas, *Distinguishing Dynamical Dark Matter at the LHC*, [Phys. Rev. **D86** \(2012\) 054008, arXiv:1204.4183 \[hep-ph\]](#).
- [334] K. R. Dienes, S. Su, and B. Thomas, *Strategies for probing nonminimal dark sectors at colliders: The interplay between cuts and kinematic distributions*, [Phys. Rev. **D91** \(2015\) no. 5, 054002, arXiv:1407.2606 \[hep-ph\]](#).

- [335] K. R. Dienes, S. Su, and B. Thomas, *Beyond the bump-hunt: A game plan for discovering dynamical dark matter at the LHC*, *AIP Conf. Proc.* **1743** (2016) 020013.
- [336] K. R. Dienes, J. Kumar, and B. Thomas, *Direct Detection of Dynamical Dark Matter*, *Phys. Rev.* **D86** (2012) 055016, [arXiv:1208.0336 \[hep-ph\]](#).
- [337] K. R. Dienes, J. Kumar, and B. Thomas, *Dynamical Dark Matter and the positron excess in light of AMS results*, *Phys. Rev.* **D88** (2013) no. 10, 103509, [arXiv:1306.2959 \[hep-ph\]](#).
- [338] K. R. Dienes, J. Kumar, and B. Thomas, *The positron excess as a smoking gun for dynamical dark matter?*, *AIP Conf. Proc.* **1604** (2014) 22–33.
- [339] K. K. Boddy, K. R. Dienes, D. Kim, J. Kumar, J.-C. Park, and B. Thomas, *Lines and Boxes: Unmasking Dynamical Dark Matter through Correlations in the MeV Gamma-Ray Spectrum*, *Phys. Rev.* **D94** (2016) no. 9, 095027, [arXiv:1606.07440 \[hep-ph\]](#).
- [340] K. K. Boddy, K. R. Dienes, D. Kim, J. Kumar, J.-C. Park, and B. Thomas, *Boxes, Boosts, and Energy Duality: Understanding the Galactic-Center Gamma-Ray Excess through Dynamical Dark Matter*, *Phys. Rev.* **D95** (2017) no. 5, 055024, [arXiv:1609.09104 \[hep-ph\]](#).
- [341] K. R. Dienes, J. Kumar, B. Thomas, and D. Yaylali, *Dark-Matter Decay as a Complementary Probe of Multicomponent Dark Sectors*, *Phys. Rev. Lett.* **114** (2015) no. 5, 051301, [arXiv:1406.4868 \[hep-ph\]](#).
- [342] K. R. Dienes, J. Kumar, B. Thomas, and D. Yaylali, *Off-diagonal dark-matter phenomenology: Exploring enhanced complementarity relations in nonminimal dark sectors*, *Phys. Rev.* **D96** (2017) no. 11, 115009, [arXiv:1708.09698 \[hep-ph\]](#).
- [343] ATLAS Collaboration, M. Aaboud et al., *Search for dark matter and other new phenomena in events with an energetic jet and large missing transverse momentum using the ATLAS detector*, *JHEP* **01** (2018) 126, [arXiv:1711.03301 \[hep-ex\]](#).
- [344] CMS Collaboration, A. M. Sirunyan et al., *Search for dark matter produced with an energetic jet or a hadronically decaying W or Z boson at $\sqrt{s} = 13$ TeV*, [arXiv:1703.01651 \[hep-ex\]](#).
- [345] ATLAS Collaboration, M. Aaboud et al., *Search for new phenomena in dijet events using 37 fb⁻¹ of pp collision data collected at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev.* **D96** (2017) no. 5, 052004, [arXiv:1703.09127 \[hep-ex\]](#).
- [346] CMS Collaboration, V. Khachatryan et al., *Search for narrow resonances decaying to dijets in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. Lett.* **116** (2016) no. 7, 071801, [arXiv:1512.01224 \[hep-ex\]](#).
- [347] L. Canetti, M. Drewes, and M. Shaposhnikov, *Matter and Antimatter in the Universe*, *New J. Phys.* **14** (2012) 095012, [arXiv:1204.4186 \[hep-ph\]](#).
- [348] B. D. Fields, P. Molaro, and S. Sarkar, *Big-Bang Nucleosynthesis*, *Chin. Phys.* **C38** (2014) 339–344, [arXiv:1412.1408 \[astro-ph.CO\]](#).
- [349] A. D. Sakharov, *Violation of CP Invariance, c Asymmetry, and Baryon Asymmetry of the Universe*, *Pisma Zh. Eksp. Teor. Fiz.* **5** (1967) 32–35. [*Usp. Fiz. Nauk*161,61(1991)].
- [350] M. B. Gavela, P. Hernandez, J. Orloff, O. Pene, and C. Quimbay, *Standard model CP violation and baryon asymmetry. Part 2: Finite temperature*, *Nucl. Phys.* **B430** (1994) 382–426, [arXiv:hep-ph/9406289 \[hep-ph\]](#).
- [351] K. Kajantie, M. Laine, K. Rummukainen, and M. E. Shaposhnikov, *Is there a hot electroweak phase transition at $m(H)$ larger or equal to $m(W)$?*, *Phys. Rev. Lett.* **77** (1996) 2887–2890, [arXiv:hep-ph/9605288 \[hep-ph\]](#).
- [352] W. Bernreuther, *CP violation and baryogenesis*, *Lect. Notes Phys.* **591** (2002) 237–293, [arXiv:hep-ph/0205279 \[hep-ph\]](#). [,237(2002)].
- [353] D. E. Morrissey and M. J. Ramsey-Musolf, *Electroweak baryogenesis*, *New J. Phys.* **14** (2012) 125003, [arXiv:1206.2942 \[hep-ph\]](#).

- [354] I. Affleck and M. Dine, *A New Mechanism for Baryogenesis*, *Nucl. Phys.* **B249** (1985) 361–380.
- [355] J. March-Russell and M. McCullough, *Asymmetric Dark Matter via Spontaneous Co-Genesis*, *JCAP* **1203** (2012) 019, [arXiv:1106.4319 \[hep-ph\]](#).
- [356] Y. Cui and R. Sundrum, *Baryogenesis for weakly interacting massive particles*, *Phys. Rev.* **D87** (2013) no. 11, 116013, [arXiv:1212.2973 \[hep-ph\]](#).
- [357] Y. Cui and B. Shuve, *Probing Baryogenesis with Displaced Vertices at the LHC*, *JHEP* **02** (2015) 049, [arXiv:1409.6729 \[hep-ph\]](#).
- [358] Y. Cui, T. Okui, and A. Yunesi, *LHC Signatures of WIMP-triggered Baryogenesis*, *Phys. Rev.* **D94** (2016) no. 11, 115022, [arXiv:1605.08736 \[hep-ph\]](#).
- [359] Y. Cui, *Natural Baryogenesis from Unnatural Supersymmetry*, *JHEP* **12** (2013) 067, [arXiv:1309.2952 \[hep-ph\]](#).
- [360] S. Ipek and J. March-Russell, *Baryogenesis via Particle-Antiparticle Oscillations*, *Phys. Rev.* **D93** (2016) no. 12, 123528, [arXiv:1604.00009 \[hep-ph\]](#).
- [361] K. Aitken, D. McKeen, T. Neder, and A. E. Nelson, *Baryogenesis from Oscillations of Charmed or Beautiful Baryons*, *Phys. Rev.* **D96** (2017) no. 7, 075009, [arXiv:1708.01259 \[hep-ph\]](#).
- [362] D. McKeen and A. E. Nelson, *CP Violating Baryon Oscillations*, *Phys. Rev.* **D94** (2016) no. 7, 076002, [arXiv:1512.05359 \[hep-ph\]](#).
- [363] CMS Collaboration, S. Chatrchyan et al., *Measurement of the Λ_b cross section and the $\bar{\Lambda}_b$ to Λ_b ratio with $J/\Psi\Lambda$ decays in pp collisions at $\sqrt{s} = 7$ TeV*, *Phys. Lett.* **B714** (2012) 136–157, [arXiv:1205.0594 \[hep-ex\]](#).
- [364] M. Fukugita and T. Yanagida, *Baryogenesis Without Grand Unification*, *Phys. Lett.* **B174** (1986) 45–47.
- [365] P. Minkowski, *$\mu \rightarrow e\gamma$ at a Rate of One Out of 10^9 Muon Decays?*, *Phys. Lett.* **67B** (1977) 421–428.
- [366] R. N. Mohapatra and G. Senjanovic, *Neutrino Mass and Spontaneous Parity Violation*, *Phys. Rev. Lett.* **44** (1980) 912.
- [367] T. Yanagida, *HORIZONTAL SYMMETRY AND MASSES OF NEUTRINOS*, Conf. Proc. **C7902131** (1979) 95–99.
- [368] M. Gell-Mann, P. Ramond, and R. Slansky, *Complex Spinors and Unified Theories*, Conf. Proc. **C790927** (1979) 315–321, [arXiv:1306.4669 \[hep-th\]](#).
- [369] S. L. Glashow, *The Future of Elementary Particle Physics*, *NATO Sci. Ser. B* **61** (1980) 687.
- [370] V. A. Kuzmin, V. A. Rubakov, and M. E. Shaposhnikov, *On the Anomalous Electroweak Baryon Number Nonconservation in the Early Universe*, *Phys. Lett.* **155B** (1985) 36.
- [371] P. S. B. Dev, P. Di Bari, B. Garbrecht, S. Lavignac, P. Millington, and D. Teresi, *Flavor effects in leptogenesis*, *Int. J. Mod. Phys.* **A33** (2018) 1842001, [arXiv:1711.02861 \[hep-ph\]](#).
- [372] M. Drewes, B. Garbrecht, P. Hernandez, M. Kekic, J. Lopez-Pavon, J. Racker, N. Rius, J. Salvado, and D. Teresi, *ARS Leptogenesis*, [arXiv:1711.02862 \[hep-ph\]](#).
- [373] B. Dev, M. Garny, J. Klaric, P. Millington, and D. Teresi, *Resonant enhancement in leptogenesis*, [arXiv:1711.02863 \[hep-ph\]](#).
- [374] S. Biondini et al., *Status of rates and rate equations for thermal leptogenesis*, [arXiv:1711.02864 \[hep-ph\]](#).
- [375] E. J. Chun et al., *Probing Leptogenesis*, *Int. J. Mod. Phys.* **A33** (2018) no. 05n06, 1842005, [arXiv:1711.02865 \[hep-ph\]](#).
- [376] C. Hagedorn, R. N. Mohapatra, E. Molinaro, C. C. Nishi, and S. T. Petcov, *CP Violation in the Lepton Sector and Implications for Leptogenesis*, [arXiv:1711.02866 \[hep-ph\]](#).
- [377] M. D’Onofrio, K. Rummukainen, and A. Tranberg, *Sphaleron Rate in the Minimal Standard Model*, *Phys. Rev. Lett.* **113** (2014) no. 14, 141602, [arXiv:1404.3565 \[hep-ph\]](#).

- [378] S. Davidson, E. Nardi, and Y. Nir, *Leptogenesis*, *Phys. Rept.* **466** (2008) 105–177, [arXiv:0802.2962 \[hep-ph\]](#).
- [379] E. K. Akhmedov, V. A. Rubakov, and A. Yu. Smirnov, *Baryogenesis via neutrino oscillations*, *Phys. Rev. Lett.* **81** (1998) 1359–1362, [arXiv:hep-ph/9803255 \[hep-ph\]](#).
- [380] T. Asaka and M. Shaposhnikov, *The nuMSM, dark matter and baryon asymmetry of the universe*, *Phys. Lett.* **B620** (2005) 17–26, [arXiv:hep-ph/0505013 \[hep-ph\]](#).
- [381] T. Asaka, S. Blanchet, and M. Shaposhnikov, *The nuMSM, dark matter and neutrino masses*, *Phys. Lett.* **B631** (2005) 151–156, [arXiv:hep-ph/0503065 \[hep-ph\]](#).
- [382] M. Drewes, *The Phenomenology of Right Handed Neutrinos*, *Int. J. Mod. Phys.* **E22** (2013) 1330019, [arXiv:1303.6912 \[hep-ph\]](#).
- [383] S. Davidson and A. Ibarra, *A Lower bound on the right-handed neutrino mass from leptogenesis*, *Phys. Lett.* **B535** (2002) 25–32, [arXiv:hep-ph/0202239 \[hep-ph\]](#).
- [384] A. Pilaftsis and T. E. J. Underwood, *Resonant leptogenesis*, *Nucl. Phys.* **B692** (2004) 303–345, [arXiv:hep-ph/0309342 \[hep-ph\]](#).
- [385] R. Barbieri and A. Dolgov, *Bounds on Sterile-neutrinos from Nucleosynthesis*, *Phys. Lett.* **B237** (1990) 440–445.
- [386] D. Gorbunov and M. Shaposhnikov, *How to find neutral leptons of the ν MSM?*, *JHEP* **10** (2007) 015, [arXiv:0705.1729 \[hep-ph\]](#). [Erratum: *JHEP*11,101(2013)].
- [387] A. Atre, T. Han, S. Pascoli, and B. Zhang, *The Search for Heavy Majorana Neutrinos*, *JHEP* **05** (2009) 030, [arXiv:0901.3589 \[hep-ph\]](#).
- [388] ATLAS Collaboration, G. Aad et al., *Search for heavy Majorana neutrinos with the ATLAS detector in pp collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **07** (2015) 162, [arXiv:1506.06020 \[hep-ex\]](#).
- [389] CMS Collaboration, V. Khachatryan et al., *Search for heavy Majorana neutrinos in $\mu^\pm\mu^\pm$ + jets events in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *Phys. Lett.* **B748** (2015) 144–166, [arXiv:1501.05566 \[hep-ex\]](#).
- [390] CMS Collaboration, V. Khachatryan et al., *Search for heavy Majorana neutrinos in ee + jets and $e\mu$ + jets events in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **04** (2016) 169, [arXiv:1603.02248 \[hep-ex\]](#).
- [391] CMS Collaboration, A. M. Sirunyan et al., *Search for heavy neutral leptons in events with three charged leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. Lett.* **120** (2018) no. 22, 221801, [arXiv:1802.02965 \[hep-ex\]](#).
- [392] F. F. Deppisch, P. S. Bhupal Dev, and A. Pilaftsis, *Neutrinos and Collider Physics*, *New J. Phys.* **17** (2015) no. 7, 075019, [arXiv:1502.06541 \[hep-ph\]](#).
- [393] Y. Cai, T. Han, T. Li, and R. Ruiz, *Lepton Number Violation: Seesaw Models and Their Collider Tests*, *Front.in Phys.* **6** (2018) 40, [arXiv:1711.02180 \[hep-ph\]](#).
- [394] T. Hambye and D. Teresi, *Higgs doublet decay as the origin of the baryon asymmetry*, *Phys. Rev. Lett.* **117** (2016) no. 9, 091801, [arXiv:1606.00017 \[hep-ph\]](#).
- [395] T. Hambye and D. Teresi, *Baryogenesis from L-violating Higgs-doublet decay in the density-matrix formalism*, *Phys. Rev.* **D96** (2017) no. 1, 015031, [arXiv:1705.00016 \[hep-ph\]](#).
- [396] A. Boyarsky, O. Ruchayskiy, and M. Shaposhnikov, *The Role of sterile neutrinos in cosmology and astrophysics*, *Ann. Rev. Nucl. Part. Sci.* **59** (2009) 191–214, [arXiv:0901.0011 \[hep-ph\]](#).
- [397] M. Drewes et al., *A White Paper on keV Sterile Neutrino Dark Matter*, *JCAP* **1701** (2017) no. 01, 025, [arXiv:1602.04816 \[hep-ph\]](#).
- [398] M. Shaposhnikov, *The nuMSM, leptonic asymmetries, and properties of singlet fermions*, *JHEP* **08** (2008) 008, [arXiv:0804.4542 \[hep-ph\]](#).
- [399] L. Canetti and M. Shaposhnikov, *Baryon Asymmetry of the Universe in the NuMSM*, *JCAP* **1009**

- (2010) 001, [arXiv:1006.0133 \[hep-ph\]](#).
- [400] T. Asaka, S. Eijima, and H. Ishida, *Kinetic Equations for Baryogenesis via Sterile Neutrino Oscillation*, *JCAP* **1202** (2012) 021, [arXiv:1112.5565 \[hep-ph\]](#).
- [401] L. Canetti, M. Drewes, and M. Shaposhnikov, *Sterile Neutrinos as the Origin of Dark and Baryonic Matter*, *Phys. Rev. Lett.* **110** (2013) no. 6, 061801, [arXiv:1204.3902 \[hep-ph\]](#).
- [402] L. Canetti, M. Drewes, T. Frossard, and M. Shaposhnikov, *Dark Matter, Baryogenesis and Neutrino Oscillations from Right Handed Neutrinos*, *Phys. Rev.* **D87** (2013) 093006, [arXiv:1208.4607 \[hep-ph\]](#).
- [403] B. Shuve and I. Yavin, *Baryogenesis through Neutrino Oscillations: A Unified Perspective*, *Phys. Rev.* **D89** (2014) no. 7, 075014, [arXiv:1401.2459 \[hep-ph\]](#).
- [404] P. Hernández, M. Kekic, J. López-Pavón, J. Racker, and N. Rius, *Leptogenesis in GeV scale seesaw models*, *JHEP* **10** (2015) 067, [arXiv:1508.03676 \[hep-ph\]](#).
- [405] M. Drewes and S. Eijima, *Neutrinoless double β decay and low scale leptogenesis*, *Phys. Lett.* **B763** (2016) 72–79, [arXiv:1606.06221 \[hep-ph\]](#).
- [406] M. Drewes, B. Garbrecht, D. Gueter, and J. Klaric, *Testing the low scale seesaw and leptogenesis*, *JHEP* **08** (2017) 018, [arXiv:1609.09069 \[hep-ph\]](#).
- [407] P. Hernández, M. Kekic, J. López-Pavón, J. Racker, and J. Salvado, *Testable Baryogenesis in Seesaw Models*, *JHEP* **08** (2016) 157, [arXiv:1606.06719 \[hep-ph\]](#).
- [408] T. Asaka, S. Eijima, and H. Ishida, *On neutrinoless double beta decay in the ν MSM*, *Phys. Lett.* **B762** (2016) 371–375, [arXiv:1606.06686 \[hep-ph\]](#).
- [409] M. Drewes, B. Garbrecht, D. Gueter, and J. Klaric, *Leptogenesis from Oscillations of Heavy Neutrinos with Large Mixing Angles*, *JHEP* **12** (2016) 150, [arXiv:1606.06690 \[hep-ph\]](#).
- [410] T. Asaka, S. Eijima, H. Ishida, K. Minogawa, and T. Yoshii, *Initial condition for baryogenesis via neutrino oscillation*, *Phys. Rev.* **D96** (2017) no. 8, 083010, [arXiv:1704.02692 \[hep-ph\]](#).
- [411] S. Antusch, E. Cazzato, M. Drewes, O. Fischer, B. Garbrecht, D. Gueter, and J. Klaric, *Probing Leptogenesis at Future Colliders*, [arXiv:1710.03744 \[hep-ph\]](#).
- [412] J. Ghiglieri and M. Laine, *GeV-scale hot sterile neutrino oscillations: a numerical solution*, *JHEP* **02** (2018) 078, [arXiv:1711.08469 \[hep-ph\]](#).
- [413] S. Eijima, M. Shaposhnikov, and I. Timiryasov, *Freeze-out of baryon number in low-scale leptogenesis*, *JCAP* **1711** (2017) no. 11, 030, [arXiv:1709.07834 \[hep-ph\]](#).
- [414] J. Ghiglieri and M. Laine, *GeV-scale hot sterile neutrino oscillations: a derivation of evolution equations*, *JHEP* **05** (2017) 132, [arXiv:1703.06087 \[hep-ph\]](#).
- [415] T. Asaka and H. Ishida, *Lepton number violation by heavy Majorana neutrino in B decays*, *Phys. Lett.* **B763** (2016) 393–396, [arXiv:1609.06113 \[hep-ph\]](#).
- [416] A. Abada, G. Arcadi, V. Domcke, and M. Lucente, *Lepton number violation as a key to low-scale leptogenesis*, *JCAP* **1511** (2015) no. 11, 041, [arXiv:1507.06215 \[hep-ph\]](#).
- [417] A. Abada, G. Arcadi, V. Domcke, and M. Lucente, *Neutrino masses, leptogenesis and dark matter from small lepton number violation?*, [arXiv:1709.00415 \[hep-ph\]](#).
- [418] M. Drewes and B. Garbrecht, *Leptogenesis from a GeV Seesaw without Mass Degeneracy*, *JHEP* **03** (2013) 096, [arXiv:1206.5537 \[hep-ph\]](#).
- [419] V. V. Khoze and G. Ro, *Leptogenesis and Neutrino Oscillations in the Classically Conformal Standard Model with the Higgs Portal*, *JHEP* **10** (2013) 075, [arXiv:1307.3764 \[hep-ph\]](#).
- [420] L. Canetti, M. Drewes, and B. Garbrecht, *Probing leptogenesis with GeV-scale sterile neutrinos at LHCb and Belle II*, *Phys. Rev.* **D90** (2014) no. 12, 125005, [arXiv:1404.7114 \[hep-ph\]](#).
- [421] J. C. Helo, M. Hirsch, and S. Kovalenko, *Heavy neutrino searches at the LHC with displaced vertices*, *Phys. Rev.* **D89** (2014) 073005, [arXiv:1312.2900 \[hep-ph\]](#). [Erratum: *Phys. Rev.* **D93**, no. 9, 099902 (2016)].

- [422] E. Izaguirre and B. Shuve, *Multilepton and Lepton Jet Probes of Sub-Weak-Scale Right-Handed Neutrinos*, *Phys. Rev.* **D91** (2015) no. 9, 093010, [arXiv:1504.02470 \[hep-ph\]](#).
- [423] A. M. Gago, P. Hernández, J. Jones-Pérez, M. Losada, and A. Moreno Briceño, *Probing the Type I Seesaw Mechanism with Displaced Vertices at the LHC*, *Eur. Phys. J.* **C75** (2015) no. 10, 470, [arXiv:1505.05880 \[hep-ph\]](#).
- [424] G. Cottin, J. C. Helo, and M. Hirsch, *Searches for light sterile neutrinos with multitrack displaced vertices*, *Phys. Rev.* **D97** (2018) no. 5, 055025, [arXiv:1801.02734 \[hep-ph\]](#).
- [425] S. Antusch, E. Cazzato, and O. Fischer, *Sterile neutrino searches via displaced vertices at LHCb*, *Phys. Lett.* **B774** (2017) 114–118, [arXiv:1706.05990 \[hep-ph\]](#).
- [426] NA62 Collaboration, C. Lazzeroni et al., *Search for heavy neutrinos in $K^+ \rightarrow \mu^+ \nu_\mu$ decays*, *Phys. Lett.* **B772** (2017) 712–718, [arXiv:1705.07510 \[hep-ex\]](#).
- [427] G. Lanfranchi, *Search for Hidden Sector particles at NA62, to appear in proceedings for the EPS-HEP 2017 conference (2017)*, .
- [428] M. Drewes, J. Hajer, J. Klaric, and G. Lanfranchi, *NA62 sensitivity to heavy neutral leptons in the low scale seesaw model*, [arXiv:1801.04207 \[hep-ph\]](#).
- [429] SHiP Collaboration, M. Anelli et al., *A facility to Search for Hidden Particles (SHiP) at the CERN SPS*, [arXiv:1504.04956 \[physics.ins-det\]](#).
- [430] E. Graverini, *SHiP sensitivity to Heavy Neutral Leptons (Sep, 2016)*, .
- [431] LBNE Collaboration, C. Adams et al., *The Long-Baseline Neutrino Experiment: Exploring Fundamental Symmetries of the Universe*, [arXiv:1307.7335 \[hep-ex\]](#).
- [432] T. Asaka, S. Eijima, and A. Watanabe, *Heavy neutrino search in accelerator-based experiments*, *JHEP* **03** (2013) 125, [arXiv:1212.1062 \[hep-ph\]](#).
- [433] FCC-ee study Team Collaboration, A. Blondel, E. Graverini, N. Serra, and M. Shaposhnikov, *Search for Heavy Right Handed Neutrinos at the FCC-ee*, *Nucl. Part. Phys. Proc.* **273-275** (2016) 1883–1890, [arXiv:1411.5230 \[hep-ex\]](#).
- [434] S. Antusch, E. Cazzato, and O. Fischer, *Displaced vertex searches for sterile neutrinos at future lepton colliders*, *JHEP* **12** (2016) 007, [arXiv:1604.02420 \[hep-ph\]](#).
- [435] A. Caputo, P. Hernandez, M. Kekic, J. López-Pavón, and J. Salvado, *The seesaw path to leptonic CP violation*, *Eur. Phys. J.* **C77** (2017) no. 4, 258, [arXiv:1611.05000 \[hep-ph\]](#).
- [436] C. Hagedorn and E. Molinaro, *Flavor and CP symmetries for leptogenesis and $0\nu\beta\beta$ decay*, *Nucl. Phys.* **B919** (2017) 404–469, [arXiv:1602.04206 \[hep-ph\]](#).
- [437] P. S. B. Dev, C. Hagedorn, and E. Molinaro, *In preparation*, .
- [438] S. F. King and C. Luhn, *Neutrino Mass and Mixing with Discrete Symmetry*, *Rept. Prog. Phys.* **76** (2013) 056201, [arXiv:1301.1340 \[hep-ph\]](#).
- [439] F. Feruglio, C. Hagedorn, and R. Ziegler, *A realistic pattern of lepton mixing and masses from S_4 and CP*, *Eur. Phys. J.* **C74** (2014) 2753, [arXiv:1303.7178 \[hep-ph\]](#).
- [440] M. B. Gavela, T. Hambye, D. Hernandez, and P. Hernandez, *Minimal Flavour Seesaw Models*, *JHEP* **09** (2009) 038, [arXiv:0906.1461 \[hep-ph\]](#).
- [441] O. Ruchayskiy and A. Ivashko, *Experimental bounds on sterile neutrino mixing angles*, *JHEP* **06** (2012) 100, [arXiv:1112.3319 \[hep-ph\]](#).
- [442] T. Asaka, S. Eijima, and H. Ishida, *Mixing of Active and Sterile Neutrinos*, *JHEP* **04** (2011) 011, [arXiv:1101.1382 \[hep-ph\]](#).
- [443] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, and T. Schwetz, *Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity*, *JHEP* **01** (2017) 087, [arXiv:1611.01514 \[hep-ph\]](#).
- [444] *NuFIT 3.1*, 2017.
- [445] S. Dodelson and L. M. Widrow, *Sterile-neutrinos as dark matter*, *Phys. Rev. Lett.* **72** (1994)

- 17–20, [arXiv:hep-ph/9303287](#) [hep-ph].
- [446] X.-D. Shi and G. M. Fuller, *A New dark matter candidate: Nonthermal sterile neutrinos*, *Phys. Rev. Lett.* **82** (1999) 2832–2835, [arXiv:astro-ph/9810076](#) [astro-ph].
- [447] M. Laine and M. Shaposhnikov, *Sterile neutrino dark matter as a consequence of ν MSM-induced lepton asymmetry*, *JCAP* **0806** (2008) 031, [arXiv:0804.4543](#) [hep-ph].
- [448] M. Lucente, A. Abada, G. Arcadi, V. Domcke, M. Drewes, and J. Klaric, *Low-scale leptogenesis with 3 right-handed neutrinos*, poster online (Jun, 2018) .
- [449] D. Gorbunov and A. Panin, *On the minimal active-sterile neutrino mixing in seesaw type I mechanism with sterile neutrinos at GeV scale*, *Phys. Rev.* **D89** (2014) no. 1, 017302, [arXiv:1312.2887](#) [hep-ph].
- [450] M. Drewes and B. Garbrecht, *Combining experimental and cosmological constraints on heavy neutrinos*, *Nucl. Phys.* **B921** (2017) 250–315, [arXiv:1502.00477](#) [hep-ph].
- [451] O. Ruchayskiy and A. Ivashko, *Restrictions on the lifetime of sterile neutrinos from primordial nucleosynthesis*, *JCAP* **1210** (2012) 014, [arXiv:1202.2841](#) [hep-ph].
- [452] P. Hernandez, M. Kekic, and J. Lopez-Pavon, *N_{eff} in low-scale seesaw models versus the lightest neutrino mass*, *Phys. Rev.* **D90** (2014) no. 6, 065033, [arXiv:1406.2961](#) [hep-ph].
- [453] S. N. Gninenko, D. S. Gorbunov, and M. E. Shaposhnikov, *Search for GeV-scale sterile neutrinos responsible for active neutrino oscillations and baryon asymmetry of the Universe*, *Adv. High Energy Phys.* **2012** (2012) 718259, [arXiv:1301.5516](#) [hep-ph].
- [454] W. Bonivento et al., *Proposal to Search for Heavy Neutral Leptons at the SPS*, [arXiv:1310.1762](#) [hep-ex].
- [455] Belle Collaboration, D. Liventsev et al., *Search for heavy neutrinos at Belle*, *Phys. Rev.* **D87** (2013) no. 7, 071102, [arXiv:1301.1105](#) [hep-ex]. [Erratum: *Phys. Rev.* **D95**, no. 9, 099903 (2017)].
- [456] G. Bernardi et al., *FURTHER LIMITS ON HEAVY NEUTRINO COUPLINGS*, *Phys. Lett.* **B203** (1988) 332–334.
- [457] CHARM Collaboration, F. Bergsma et al., *A Search for Decays of Heavy Neutrinos in the Mass Range 0.5-GeV to 2.8-GeV*, *Phys. Lett.* **166B** (1986) 473–478.
- [458] NuTeV, E815 Collaboration, A. Vaitaitis et al., *Search for neutral heavy leptons in a high-energy neutrino beam*, *Phys. Rev. Lett.* **83** (1999) 4943–4946, [arXiv:hep-ex/9908011](#) [hep-ex].
- [459] A. G. Vaitaitis, *Search for neutral heavy leptons in a high-energy neutrino beam*. PhD thesis, Columbia U.
- [460] NA3 Collaboration, J. Badier et al., *Mass and Lifetime Limits on New Longlived Particles in 300-GeV/c π^- Interactions*, *Z. Phys.* **C31** (1986) 21.
- [461] WA66 Collaboration, A. M. Cooper-Sarkar et al., *Search for Heavy Neutrino Decays in the BEBC Beam Dump Experiment*, *Phys. Lett.* **160B** (1985) 207–211.
- [462] DELPHI Collaboration, P. Abreu et al., *Search for neutral heavy leptons produced in Z decays*, *Z. Phys.* **C74** (1997) 57–71. [Erratum: *Z. Phys.* **C75**, 580 (1997)].
- [463] LHCb Collaboration, R. Aaij et al., *Search for massive long-lived particles decaying semileptonically in the LHCb detector*, *Eur. Phys. J.* **C77** (2017) no. 4, 224, [arXiv:1612.00945](#) [hep-ex].
- [464] J. C. Helo, S. Kovalenko, and I. Schmidt, *Sterile neutrinos in lepton number and lepton flavor violating decays*, *Nucl. Phys.* **B853** (2011) 80–104, [arXiv:1005.1607](#) [hep-ph].
- [465] K. Bondarenko, A. Boyarsky, D. Gorbunov, and O. Ruchayskiy, *Phenomenology of GeV-scale Heavy Neutral Leptons*, [arXiv:1805.08567](#) [hep-ph].
- [466] B. Shuve and M. E. Peskin, *Revision of the LHCb Limit on Majorana Neutrinos*, *Phys. Rev.* **D94** (2016) no. 11, 113007, [arXiv:1607.04258](#) [hep-ph].

- [467] M. Shaposhnikov, *A Possible symmetry of the nuMSM*, *Nucl. Phys.* **B763** (2007) 49–59, [arXiv:hep-ph/0605047 \[hep-ph\]](#).
- [468] J. Kersten and A. Yu. Smirnov, *Right-Handed Neutrinos at CERN LHC and the Mechanism of Neutrino Mass Generation*, *Phys. Rev.* **D76** (2007) 073005, [arXiv:0705.3221 \[hep-ph\]](#).
- [469] SHiP Collaboration, *Sensitivity of the SHiP experiment towards heavy neutral leptons*. In preparation., .
- [470] T. Golling et al., *Physics at a 100 TeV pp collider: beyond the Standard Model phenomena*, *CERN Yellow Report* (2017) no. 3, 441–634, [arXiv:1606.00947 \[hep-ph\]](#).
- [471] M. L. Mangano et al., *Physics at a 100 TeV pp collider: Standard Model processes*, *CERN Yellow Report* (2017) no. 3, 1–254, [arXiv:1607.01831 \[hep-ph\]](#).
- [472] R. Contino et al., *Physics at a 100 TeV pp collider: Higgs and EW symmetry breaking studies*, [arXiv:1606.09408 \[hep-ph\]](#).
- [473] J. Tang et al., *Concept for a Future Super Proton-Proton Collider*, [arXiv:1507.03224 \[physics.acc-ph\]](#).
- [474] W. Kilian, T. Ohl, and J. Reuter, *WHIZARD: Simulating Multi-Particle Processes at LHC and ILC*, *Eur. Phys. J.* **C71** (2011) 1742, [arXiv:0708.4233 \[hep-ph\]](#).
- [475] M. Moretti, T. Ohl, and J. Reuter, *O’Mega: An Optimizing matrix element generator*, [arXiv:hep-ph/0102195 \[hep-ph\]](#).
- [476] I. Hinchliffe, A. Kotwal, M. L. Mangano, C. Quigg, and L.-T. Wang, *Luminosity goals for a 100-TeV pp collider*, *Int. J. Mod. Phys.* **A30** (2015) no. 23, 1544002, [arXiv:1504.06108 \[hep-ph\]](#).
- [477] S. P. Baranov, *Pair production of $B(c)^*$ mesons in $p p$ and $\gamma \gamma$ collisions*, *Phys. Rev.* **D55** (1997) 2756–2759.
- [478] K. Kolodziej and R. Ruckl, *On the energy dependence of hadronic $B(c)$ production*, *Nucl. Instrum. Meth.* **A408** (1998) 33–35, [arXiv:hep-ph/9803327 \[hep-ph\]](#).
- [479] B. Batell, M. Pospelov, and B. Shuve, *Shedding Light on Neutrino Masses with Dark Forces*, *JHEP* **08** (2016) 052, [arXiv:1604.06099 \[hep-ph\]](#).
- [480] I. Hoenig, G. Samach, and D. Tucker-Smith, *Searching for dilepton resonances below the Z mass at the LHC*, *Phys. Rev.* **D90** (2014) no. 7, 075016, [arXiv:1408.1075 \[hep-ph\]](#).
- [481] P. S. B. Dev, R. N. Mohapatra, and Y. Zhang, *Probing the Higgs Sector of the Minimal Left-Right Symmetric Model at Future Hadron Colliders*, *JHEP* **05** (2016) 174, [arXiv:1602.05947 \[hep-ph\]](#).
- [482] A. Maiezza, G. Senjanović, and J. C. Vasquez, *Higgs Sector of the Left-Right Symmetric Theory*, [arXiv:1612.09146 \[hep-ph\]](#).
- [483] ATLAS Collaboration, T. A. collaboration, *Search for new high-mass resonances in the dilepton final state using proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, .
- [484] CMS Collaboration, C. Collaboration, *Search for a high-mass resonance decaying into a dilepton final state in 13 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV*, .
- [485] S. Patra, F. S. Queiroz, and W. Rodejohann, *Stringent Dilepton Bounds on Left-Right Models using LHC data*, *Phys. Lett.* **B752** (2016) 186–190, [arXiv:1506.03456 \[hep-ph\]](#).
- [486] M. Lindner, F. S. Queiroz, and W. Rodejohann, *Dilepton bounds on left?right symmetry at the LHC run II and neutrinoless double beta decay*, *Phys. Lett.* **B762** (2016) 190–195, [arXiv:1604.07419 \[hep-ph\]](#).
- [487] P. Langacker, *The Physics of Heavy Z' Gauge Bosons*, *Rev. Mod. Phys.* **81** (2009) 1199–1228, [arXiv:0801.1345 \[hep-ph\]](#).
- [488] J. C. Pati and A. Salam, *Lepton Number as the Fourth Color*, *Phys. Rev.* **D10** (1974) 275–289. [Erratum: *Phys. Rev.*D11,703(1975)].

- [489] R. N. Mohapatra and J. C. Pati, *A Natural Left-Right Symmetry*, *Phys. Rev.* **D11** (1975) 2558.
- [490] G. Senjanovic and R. N. Mohapatra, *Exact Left-Right Symmetry and Spontaneous Violation of Parity*, *Phys. Rev.* **D12** (1975) 1502.
- [491] G. Senjanovic, *Spontaneous Breakdown of Parity in a Class of Gauge Theories*, *Nucl. Phys.* **B153** (1979) 334–364.
- [492] O. Castillo-Felisola, C. O. Dib, J. C. Helo, S. G. Kovalenko, and S. E. Ortiz, *Left-Right Symmetric Models at the High-Intensity Frontier*, *Phys. Rev.* **D92** (2015) no. 1, 013001, [arXiv:1504.02489 \[hep-ph\]](#).
- [493] J. C. Helo and M. Hirsch, *LHC dijet constraints on double beta decay*, *Phys. Rev.* **D92** (2015) no. 7, 073017, [arXiv:1509.00423 \[hep-ph\]](#).
- [494] W.-Y. Keung and G. Senjanovic, *Majorana Neutrinos and the Production of the Right-handed Charged Gauge Boson*, *Phys. Rev. Lett.* **50** (1983) 1427.
- [495] CMS Collaboration, V. Khachatryan et al., *Search for heavy neutrinos and W bosons with right-handed couplings in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *Eur. Phys. J.* **C74** (2014) no. 11, 3149, [arXiv:1407.3683 \[hep-ex\]](#).
- [496] Y. Zhang, H. An, X. Ji, and R. N. Mohapatra, *General CP Violation in Minimal Left-Right Symmetric Model and Constraints on the Right-Handed Scale*, *Nucl. Phys.* **B802** (2008) 247–279, [arXiv:0712.4218 \[hep-ph\]](#).
- [497] S. Bertolini, A. Maiezza, and F. Nesti, *Present and Future K and B Meson Mixing Constraints on TeV Scale Left-Right Symmetry*, *Phys. Rev.* **D89** (2014) no. 9, 095028, [arXiv:1403.7112 \[hep-ph\]](#).
- [498] P. Langacker, *Grand Unified Theories and Proton Decay*, *Phys. Rept.* **72** (1981) 185.
- [499] J. L. Hewett and T. G. Rizzo, *Low-Energy Phenomenology of Superstring Inspired E(6) Models*, *Phys. Rept.* **183** (1989) 193.
- [500] E. Accomando, L. Delle Rose, S. Moretti, E. Olaiya, and C. H. Shepherd-Themistocleous, *Novel SM-like Higgs decay into displaced heavy neutrino pairs in $U(1)'$ models*, *JHEP* **04** (2017) 081, [arXiv:1612.05977 \[hep-ph\]](#).
- [501] E. Accomando, L. Delle Rose, S. Moretti, E. Olaiya, and C. H. Shepherd-Themistocleous, *Extra Higgs Boson and Z' as Portals to Signatures of Heavy Neutrinos at the LHC*, [arXiv:1708.03650 \[hep-ph\]](#).
- [502] M. Kawasaki, K. Kohri, and T. Moroi, *Big-Bang nucleosynthesis and hadronic decay of long-lived massive particles*, *Phys. Rev.* **D71** (2005) 083502, [arXiv:astro-ph/0408426 \[astro-ph\]](#).
- [503] M. Gonderinger, Y. Li, H. Patel, and M. J. Ramsey-Musolf, *Vacuum Stability, Perturbativity, and Scalar Singlet Dark Matter*, *JHEP* **01** (2010) 053, [arXiv:0910.3167 \[hep-ph\]](#).
- [504] L. Basso, S. Moretti, and G. M. Pruna, *A Renormalisation Group Equation Study of the Scalar Sector of the Minimal B-L Extension of the Standard Model*, *Phys. Rev.* **D82** (2010) 055018, [arXiv:1004.3039 \[hep-ph\]](#).
- [505] J. Elias-Miro, J. R. Espinosa, G. F. Giudice, H. M. Lee, and A. Strumia, *Stabilization of the Electroweak Vacuum by a Scalar Threshold Effect*, *JHEP* **06** (2012) 031, [arXiv:1203.0237 \[hep-ph\]](#).
- [506] C. Coriano, L. Delle Rose, and C. Marzo, *Vacuum Stability in $U(1)$ -Prime Extensions of the Standard Model with TeV Scale Right Handed Neutrinos*, *Phys. Lett.* **B738** (2014) 13–19, [arXiv:1407.8539 \[hep-ph\]](#).
- [507] C. Coriano, L. Delle Rose, and C. Marzo, *Constraints on abelian extensions of the Standard Model from two-loop vacuum stability and $U(1)_{B-L}$* , *JHEP* **02** (2016) 135, [arXiv:1510.02379 \[hep-ph\]](#).

- [508] A. Falkowski, C. Gross, and O. Lebedev, *A second Higgs from the Higgs portal*, *JHEP* **05** (2015) 057, [arXiv:1502.01361 \[hep-ph\]](#).
- [509] E. Accomando, C. Coriano, L. Delle Rose, J. Fiaschi, C. Marzo, and S. Moretti, *Higgses and heavy neutrinos in $U(1)$ models: from the LHC to the GUT scale*, *JHEP* **07** (2016) 086, [arXiv:1605.02910 \[hep-ph\]](#).
- [510] LHC Higgs Cross Section Working Group Collaboration, D. de Florian et al., *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector*, [arXiv:1610.07922 \[hep-ph\]](#).
- [511] CMS Collaboration, V. Khachatryan et al., *Searches for invisible decays of the Higgs boson in pp collisions at $\sqrt{s} = 7, 8, \text{ and } 13 \text{ TeV}$* , *JHEP* **02** (2017) 135, [arXiv:1610.09218 \[hep-ex\]](#).
- [512] ATLAS Collaboration, T. A. collaboration, *Search for an invisibly decaying Higgs boson or dark matter candidates produced in association with a Z boson in pp collisions at $\sqrt{s} = 13 \text{ TeV}$ with the ATLAS detector*, .
- [513] M. L. Graesser, *Broadening the Higgs boson with right-handed neutrinos and a higher dimension operator at the electroweak scale*, *Phys. Rev.* **D76** (2007) 075006, [arXiv:0704.0438 \[hep-ph\]](#).
- [514] F. del Aguila, S. Bar-Shalom, A. Soni, and J. Wudka, *Heavy Majorana Neutrinos in the Effective Lagrangian Description: Application to Hadron Colliders*, *Phys. Lett.* **B670** (2009) 399–402, [arXiv:0806.0876 \[hep-ph\]](#).
- [515] A. Aparici, K. Kim, A. Santamaria, and J. Wudka, *Right-handed neutrino magnetic moments*, *Phys. Rev.* **D80** (2009) 013010, [arXiv:0904.3244 \[hep-ph\]](#).
- [516] S. Weinberg, *Baryon and Lepton Nonconserving Processes*, *Phys. Rev. Lett.* **43** (1979) 1566–1570.
- [517] E. Ma, *Verifiable radiative seesaw mechanism of neutrino mass and dark matter*, *Phys. Rev.* **D73** (2006) 077301, [arXiv:hep-ph/0601225 \[hep-ph\]](#).
- [518] D. Hehn and A. Ibarra, *A radiative model with a naturally mild neutrino mass hierarchy*, *Phys. Lett.* **B718** (2013) 988–991, [arXiv:1208.3162 \[hep-ph\]](#).
- [519] A. G. Hessler, A. Ibarra, E. Molinaro, and S. Vogl, *Impact of the Higgs boson on the production of exotic particles at the LHC*, *Phys. Rev.* **D91** (2015) no. 11, 115004, [arXiv:1408.0983 \[hep-ph\]](#).
- [520] E. Molinaro, C. E. Yaguna, and O. Zapata, *FIMP realization of the scotogenic model*, *JCAP* **1407** (2014) 015, [arXiv:1405.1259 \[hep-ph\]](#).
- [521] A. G. Hessler, A. Ibarra, E. Molinaro, and S. Vogl, *Probing the scotogenic FIMP at the LHC*, *JHEP* **01** (2017) 100, [arXiv:1611.09540 \[hep-ph\]](#).
- [522] A. Ibarra, E. Molinaro, and S. Vogl, *In preparation*, .
- [523] F. Feruglio, C. Hagedorn, and R. Ziegler, *Lepton Mixing Parameters from Discrete and CP Symmetries*, *JHEP* **07** (2013) 027, [arXiv:1211.5560 \[hep-ph\]](#).
- [524] A. Pilaftsis, *CP violation and baryogenesis due to heavy Majorana neutrinos*, *Phys. Rev.* **D56** (1997) 5431–5451, [arXiv:hep-ph/9707235 \[hep-ph\]](#).
- [525] F. F. Deppisch, N. Desai, and J. W. F. Valle, *Is charged lepton flavor violation a high energy phenomenon?*, *Phys. Rev.* **D89** (2014) no. 5, 051302, [arXiv:1308.6789 \[hep-ph\]](#).
- [526] C. Luhn, S. Nasri, and P. Ramond, *The Flavor group $\Delta(3n^2)$* , *J. Math. Phys.* **48** (2007) 073501, [arXiv:hep-th/0701188 \[hep-th\]](#).
- [527] J. A. Escobar and C. Luhn, *The Flavor Group $\Delta(6n^2)$* , *J. Math. Phys.* **50** (2009) 013524, [arXiv:0809.0639 \[hep-th\]](#).
- [528] C. Hagedorn, A. Meroni, and E. Molinaro, *Lepton mixing from $\Delta(3n^2)$ and $\Delta(6n^2)$ and CP*, *Nucl. Phys.* **B891** (2015) 499–557, [arXiv:1408.7118 \[hep-ph\]](#).

- [529] LEGEND Collaboration, N. Abgrall et al., *The Large Enriched Germanium Experiment for Neutrinoless Double Beta Decay (LEGEND)*, *AIP Conf. Proc.* **1894** (2017) no. 1, 020027, [arXiv:1709.01980 \[physics.ins-det\]](#).
- [530] nEXO Collaboration, J. B. Albert et al., *Sensitivity and Discovery Potential of nEXO to Neutrinoless Double Beta Decay*, [arXiv:1710.05075 \[nucl-ex\]](#).
- [531] P. S. Bhupal Dev, P. Millington, A. Pilaftsis, and D. Teresi, *Flavour Covariant Transport Equations: an Application to Resonant Leptogenesis*, *Nucl. Phys.* **B886** (2014) 569–664, [arXiv:1404.1003 \[hep-ph\]](#).
- [532] R. K. Ellis, W. J. Stirling, and B. R. Webber, *QCD and collider physics*, Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol. **8** (1996) 1–435.
- [533] P. Schwaller, D. Stolarski, and A. Weiler, *Emerging Jets*, *JHEP* **05** (2015) 059, [arXiv:1502.05409 \[hep-ph\]](#).
- [534] Y. Bai and P. Schwaller, *Scale of dark QCD*, *Phys. Rev.* **D89** (2014) no. 6, 063522, [arXiv:1306.4676 \[hep-ph\]](#).
- [535] M. J. Strassler, *Why Unparticle Models with Mass Gaps are Examples of Hidden Valleys*, [arXiv:0801.0629 \[hep-ph\]](#).
- [536] J. Polchinski and M. J. Strassler, *Deep inelastic scattering and gauge / string duality*, *JHEP* **05** (2003) 012, [arXiv:hep-th/0209211 \[hep-th\]](#).
- [537] D. M. Hofman and J. Maldacena, *Conformal collider physics: Energy and charge correlations*, *JHEP* **05** (2008) 012, [arXiv:0803.1467 \[hep-th\]](#).
- [538] S. Lin and E. Shuryak, *Toward the AdS/CFT gravity dual for high energy collisions. 2. The stress tensor on the boundary*, *Phys. Rev.* **D77** (2008) 085014, [arXiv:0711.0736 \[hep-th\]](#).
- [539] R. Hagedorn, *Statistical thermodynamics of strong interactions at high-energies*, *Nuovo Cim. Suppl.* **3** (1965) 147–186.
- [540] Y. Hatta and T. Matsuo, *Thermal hadron spectrum in $e+e-$ annihilation from gauge/string duality*, *Phys. Rev. Lett.* **102** (2009) 062001, [arXiv:0807.0098 \[hep-ph\]](#).
- [541] S. Knapen, S. Pagan Griso, M. Papucci, and D. J. Robinson, *Triggering Soft Bombs at the LHC*, [arXiv:1612.00850 \[hep-ph\]](#).
- [542] L. Carloni and T. Sjostrand, *Visible Effects of Invisible Hidden Valley Radiation*, *JHEP* **09** (2010) 105, [arXiv:1006.2911 \[hep-ph\]](#).
- [543] L. Carloni, J. Rathsman, and T. Sjostrand, *Discerning Secluded Sector gauge structures*, *JHEP* **04** (2011) 091, [arXiv:1102.3795 \[hep-ph\]](#).
- [544] T. Sjostrand, S. Mrenna, and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852–867, [arXiv:0710.3820 \[hep-ph\]](#).
- [545] T. Cohen, T. Golling, M. Hance, A. Henrichs, K. Howe, J. Loyal, S. Padhi, and J. G. Wacker, *SUSY Simplified Models at 14, 33, and 100 TeV Proton Colliders*, *JHEP* **04** (2014) 117, [arXiv:1311.6480 \[hep-ph\]](#).
- [546] LHC Higgs Cross Section Working Group Collaboration, J. R. Andersen et al., *Handbook of LHC Higgs Cross Sections: 3. Higgs Properties*, [arXiv:1307.1347 \[hep-ph\]](#).
- [547] J. L. Hewett et al., *Fundamental Physics at the Intensity Frontier*, [arXiv:1205.2671 \[hep-ex\]](#).
- [548] N. Zhou, D. Berge, L. Wang, D. Whiteson, and T. Tait, *Sensitivity of future collider facilities to WIMP pair production via effective operators and light mediators*, [arXiv:1307.5327 \[hep-ex\]](#).
- [549] *Projections for measurements of Higgs boson cross sections, branching ratios and coupling parameters with the ATLAS detector at a HL-LHC*, ATL-PHYS-PUB-2013-014, CERN, Geneva, Oct, 2013.
- [550] S. Knapen, S. Pagan Griso, M. Papucci, and D. J. Robinson, *In progress*, .

- [551] D. Curtin et al., *Exotic decays of the 125 GeV Higgs boson*, *Phys. Rev.* **D90** (2014) no. 7, 075004, [arXiv:1312.4992 \[hep-ph\]](#).
- [552] ATLAS, CMS Collaboration, G. Aad et al., *Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments*, *Phys. Rev. Lett.* **114** (2015) 191803, [arXiv:1503.07589 \[hep-ex\]](#).
- [553] P. W. Graham, D. E. Kaplan, S. Rajendran, and P. Saraswat, *Displaced Supersymmetry*, *JHEP* **07** (2012) 149, [arXiv:1204.6038 \[hep-ph\]](#).
- [554] R. H. Cyburt, B. D. Fields, K. A. Olive, and T.-H. Yeh, *Big Bang Nucleosynthesis: 2015*, *Rev. Mod. Phys.* **88** (2016) 015004, [arXiv:1505.01076 \[astro-ph.CO\]](#).
- [555] K. Jedamzik and M. Pospelov, *Big Bang Nucleosynthesis and Particle Dark Matter*, *New J. Phys.* **11** (2009) 105028, [arXiv:0906.2087 \[hep-ph\]](#).
- [556] M. Pospelov and J. Pradler, *Big Bang Nucleosynthesis as a Probe of New Physics*, *Ann. Rev. Nucl. Part. Sci.* **60** (2010) 539–568, [arXiv:1011.1054 \[hep-ph\]](#).
- [557] A. Fradette and M. Pospelov, *BBN for the LHC: constraints on lifetimes of the Higgs portal scalars*, *Phys. Rev.* **D96** (2017) no. 7, 075033, [arXiv:1706.01920 \[hep-ph\]](#).
- [558] J. M. Cline, K. Kainulainen, P. Scott, and C. Weniger, *Update on scalar singlet dark matter*, *Phys. Rev.* **D88** (2013) 055025, [arXiv:1306.4710 \[hep-ph\]](#). [Erratum: *Phys. Rev.* **D92**, no. 3, 039906 (2015)].
- [559] GAMBIT Collaboration, P. Athron et al., *Status of the scalar singlet dark matter model*, *Eur. Phys. J.* **C77** (2017) no. 8, 568, [arXiv:1705.07931 \[hep-ph\]](#).
- [560] G. Belanger, B. Dumont, U. Ellwanger, J. F. Gunion, and S. Kraml, *Global fit to Higgs signal strengths and couplings and implications for extended Higgs sectors*, *Phys. Rev.* **D88** (2013) 075008, [arXiv:1306.2941 \[hep-ph\]](#).
- [561] A. Djouadi, J. Kalinowski, and M. Spira, *HDECAY: A Program for Higgs boson decays in the standard model and its supersymmetric extension*, *Comput. Phys. Commun.* **108** (1998) 56–74, [arXiv:hep-ph/9704448 \[hep-ph\]](#).
- [562] F. Bezrukov and D. Gorbunov, *Light inflaton Hunter’s Guide*, *JHEP* **05** (2010) 010, [arXiv:0912.0390 \[hep-ph\]](#).
- [563] M. B. Voloshin, *Once Again About the Role of Gluonic Mechanism in Interaction of Light Higgs Boson with Hadrons*, *Sov. J. Nucl. Phys.* **44** (1986) 478. [*Yad. Fiz.* 44, 738 (1986)].
- [564] H. Leutwyler and M. A. Shifman, *Light Higgs Particle in Decays of K and η Mesons*, *Nucl. Phys.* **B343** (1990) 369–397.
- [565] J. F. Donoghue, J. Gasser, and H. Leutwyler, *The Decay of a Light Higgs Boson*, *Nucl. Phys.* **B343** (1990) 341–368.
- [566] B. Hyams et al., *$\pi\pi$ Phase Shift Analysis from 600-MeV to 1900-MeV*, *Nucl. Phys.* **B64** (1973) 134–162.
- [567] M. Spira, *QCD effects in Higgs physics*, *Fortsch. Phys.* **46** (1998) 203–284, [arXiv:hep-ph/9705337 \[hep-ph\]](#).
- [568] J. F. Gunion, H. E. Haber, G. L. Kane, and S. Dawson, *The Higgs Hunter’s Guide*, *Front. Phys.* **80** (2000) 1–404.
- [569] D. McKeen, *Constraining Light Bosons with Radiative $U(1)$ Decays*, *Phys. Rev.* **D79** (2009) 015007, [arXiv:0809.4787 \[hep-ph\]](#).
- [570] S. Tulin, H.-B. Yu, and K. M. Zurek, *Beyond Collisionless Dark Matter: Particle Physics Dynamics for Dark Matter Halo Structure*, *Phys. Rev.* **D87** (2013) no. 11, 115007, [arXiv:1302.3898 \[hep-ph\]](#).
- [571] A. Martin, J. Shelton, and J. Unwin, *Fitting the Galactic Center Gamma-Ray Excess with Cascade Annihilations*, *Phys. Rev.* **D90** (2014) no. 10, 103513, [arXiv:1405.0272 \[hep-ph\]](#).

- [572] G. Krnjaic, *Probing Light Thermal Dark-Matter With a Higgs Portal Mediator*, *Phys. Rev.* **D94** (2016) no. 7, 073009, [arXiv:1512.04119 \[hep-ph\]](#).
- [573] Y.-F. Zhou and Y.-L. Wu, *Lepton flavor changing scalar interactions and muon $g-2$* , *Eur. Phys. J.* **C27** (2003) 577–585, [arXiv:hep-ph/0110302 \[hep-ph\]](#).
- [574] M. Pospelov, *Secluded $U(1)$ below the weak scale*, *Phys. Rev.* **D80** (2009) 095002, [arXiv:0811.1030 \[hep-ph\]](#).
- [575] D. Tucker-Smith and I. Yavin, *Muonic hydrogen and MeV forces*, *Phys. Rev.* **D83** (2011) 101702, [arXiv:1011.4922 \[hep-ph\]](#).
- [576] V. Barger, C.-W. Chiang, W.-Y. Keung, and D. Marfatia, *Proton size anomaly*, *Phys. Rev. Lett.* **106** (2011) 153001, [arXiv:1011.3519 \[hep-ph\]](#).
- [577] B. Batell, D. McKeen, and M. Pospelov, *New Parity-Violating Muonic Forces and the Proton Charge Radius*, *Phys. Rev. Lett.* **107** (2011) 011803, [arXiv:1103.0721 \[hep-ph\]](#).
- [578] S. Raby and G. B. West, *The Branching Ratio for a Light Higgs to Decay Into $\mu^+\mu^-$ Pairs*, *Phys. Rev.* **D38** (1988) 3488.
- [579] SHiP Collaboration, G. Lanfranchi, *Sensitivity of the SHiP experiment to a light scalar particle mixing with the Higgs*, .
- [580] B. Grinstein, L. J. Hall, and L. Randall, *Do B meson decays exclude a light Higgs?*, *Phys. Lett.* **B211** (1988) 363–369.
- [581] A. Lenz, *Lifetimes and HQE*, [arXiv:1405.3601 \[hep-ph\]](#).
- [582] C. Anastasiou, C. Duhr, F. Dulat, E. Furlan, T. Gehrmann, F. Herzog, A. Lazopoulos, and B. Mistlberger, *High precision determination of the gluon fusion Higgs boson cross-section at the LHC*, *JHEP* **05** (2016) 058, [arXiv:1602.00695 \[hep-ph\]](#).
- [583] L3 Collaboration, M. Acciarri et al., *Search for neutral Higgs boson production through the process $e^+e^- \rightarrow Z^*H^0$* , *Phys. Lett.* **B385** (1996) 454–470.
- [584] ALEPH Collaboration, D. Buskulic et al., *Search for a nonminimal Higgs boson produced in the reaction $e^+e^- \rightarrow h Z^*$* , *Phys. Lett.* **B313** (1993) 312–325.
- [585] BNL-E949 Collaboration, A. V. Artamonov et al., *Study of the decay $K^+ \rightarrow \pi^+\nu\bar{\nu}$ in the momentum region $140 < P_\pi < 199$ MeV/c*, *Phys. Rev.* **D79** (2009) 092004, [arXiv:0903.0030 \[hep-ex\]](#).
- [586] CHARM Collaboration, F. Bergsma et al., *Search for Axion Like Particle Production in 400-GeV Proton - Copper Interactions*, *Phys. Lett.* **B157** (1985) 458–462.
- [587] LHCb Collaboration, R. Aaij et al., *Search for hidden-sector bosons in $B^0 \rightarrow K^{*0}\mu^+\mu^-$ decays*, *Phys. Rev. Lett.* **115** (2015) no. 16, 161802, [arXiv:1508.04094 \[hep-ex\]](#).
- [588] LHCb Collaboration, R. Aaij et al., *Search for long-lived scalar particles in $B^+ \rightarrow K^+\chi(\mu^+\mu^-)$ decays*, *Phys. Rev.* **D95** (2017) no. 7, 071101, [arXiv:1612.07818 \[hep-ex\]](#).
- [589] J. Jaeckel and A. Ringwald, *The Low-Energy Frontier of Particle Physics*, *Ann. Rev. Nucl. Part. Sci.* **60** (2010) 405–437, [arXiv:1002.0329 \[hep-ph\]](#).
- [590] R. Essig et al., *Working Group Report: New Light Weakly Coupled Particles*, in *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*.
- [591] J. Alexander et al., *Dark Sectors 2016 Workshop: Community Report*, .
- [592] M. Battaglieri et al., *US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report*, [arXiv:1707.04591 \[hep-ph\]](#).
- [593] B. Holdom, *Two $U(1)$'s and Epsilon Charge Shifts*, *Phys. Lett.* **B166** (1986) 196–198.
- [594] P. Galison and A. Manohar, *Two Z's Or Not Two Z's?*, *Phys. Lett.* **B136** (1984) 279.
- [595] N. Arkani-Hamed and N. Weiner, *LHC Signals for a SuperUnified Theory of Dark Matter*, *JHEP* **12** (2008) 104, [arXiv:0810.0714 \[hep-ph\]](#).

- [596] M. Baumgart, C. Cheung, J. T. Ruderman, L.-T. Wang, and I. Yavin, *Non-Abelian Dark Sectors and Their Collider Signatures*, *JHEP* **0904** (2009) 014, [arXiv:0901.0283 \[hep-ph\]](#).
- [597] R. Essig, P. Schuster, and N. Toro, *Probing Dark Forces and Light Hidden Sectors at Low-Energy $e+e-$ Colliders*, *Phys. Rev.* **D80** (2009) 015003, [arXiv:0903.3941 \[hep-ph\]](#).
- [598] D. Curtin, R. Essig, S. Gori, and J. Shelton, *Illuminating Dark Photons with High-Energy Colliders*, *JHEP* **02** (2015) 157, [arXiv:1412.0018 \[hep-ph\]](#).
- [599] E. M. Riordan et al., *A Search for Short Lived Axions in an Electron Beam Dump Experiment*, *Phys. Rev. Lett.* **59** (1987) 755.
- [600] J. D. Bjorken, S. Ecklund, W. R. Nelson, A. Abashian, C. Church, B. Lu, L. W. Mo, T. A. Nunamaker, and P. Rassmann, *Search for Neutral Metastable Penetrating Particles Produced in the SLAC Beam Dump*, *Phys. Rev.* **D38** (1988) 3375.
- [601] A. Bross, M. Crisler, S. H. Pordes, J. Volk, S. Errede, and J. Wrbanek, *A Search for Shortlived Particles Produced in an Electron Beam Dump*, *Phys. Rev. Lett.* **67** (1991) 2942–2945.
- [602] B. Batell, M. Pospelov, and A. Ritz, *Probing a Secluded $U(1)$ at B-factories*, *Phys. Rev.* **D79** (2009) 115008, [arXiv:0903.0363 \[hep-ph\]](#).
- [603] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, *New Fixed-Target Experiments to Search for Dark Gauge Forces*, *Phys. Rev.* **D80** (2009) 075018, [arXiv:0906.0580 \[hep-ph\]](#).
- [604] J. Blumlein and J. Brunner, *New Exclusion Limits for Dark Gauge Forces from Beam-Dump Data*, *Phys.Lett.* **B701** (2011) 155–159, [arXiv:1104.2747 \[hep-ex\]](#).
- [605] S. Andreas, C. Niebuhr, and A. Ringwald, *New Limits on Hidden Photons from Past Electron Beam Dumps*, *Phys.Rev.* **D86** (2012) 095019, [arXiv:1209.6083 \[hep-ph\]](#).
- [606] M. Reece and L.-T. Wang, *Searching for the Light Dark Gauge Boson in GeV-Scale Experiments*, *JHEP* **0907** (2009) 051, [arXiv:0904.1743 \[hep-ph\]](#).
- [607] BABAR Collaboration, B. Aubert et al., *Search for Dimuon Decays of a Light Scalar Boson in Radiative Transitions $\Upsilon \rightarrow \gamma A_0$* , *Phys. Rev. Lett.* **103** (2009) 081803, [arXiv:0905.4539 \[hep-ex\]](#).
- [608] KLOE-2 Collaboration, F. Archilli et al., *Search for a vector gauge boson in ϕ meson decays with the KLOE detector*, *Phys. Lett.* **B706** (2012) 251–255, [arXiv:1110.0411 \[hep-ex\]](#).
- [609] APEX Collaboration, S. Abrahamyan et al., *Search for a New Gauge Boson in Electron-Nucleus Fixed-Target Scattering by the APEX Experiment*, *Phys. Rev. Lett.* **107** (2011) 191804, [arXiv:1108.2750 \[hep-ex\]](#).
- [610] A1 Collaboration, H. Merkel et al., *Search for Light Gauge Bosons of the Dark Sector at the Mainz Microtron*, *Phys. Rev. Lett.* **106** (2011) 251802, [arXiv:1101.4091 \[nucl-ex\]](#).
- [611] H. Merkel et al., *Search at the Mainz Microtron for Light Massive Gauge Bosons Relevant for the Muon $g-2$ Anomaly*, *Phys. Rev. Lett.* **112** (2014) no. 22, 221802, [arXiv:1404.5502 \[hep-ex\]](#).
- [612] KLOE-2 Collaboration, D. Babusci et al., *Limit on the production of a light vector gauge boson in phi meson decays with the KLOE detector*, *Phys. Lett.* **B720** (2013) 111–115, [arXiv:1210.3927 \[hep-ex\]](#).
- [613] KLOE-2 Collaboration, D. Babusci et al., *Search for light vector boson production in $e^+e^- \rightarrow \mu^+\mu^-\gamma$ interactions with the KLOE experiment*, *Phys. Lett.* **B736** (2014) 459–464, [arXiv:1404.7772 \[hep-ex\]](#).
- [614] BaBar Collaboration, J. P. Lees et al., *Search for a Dark Photon in e^+e^- Collisions at BaBar*, *Phys. Rev. Lett.* **113** (2014) no. 20, 201801, [arXiv:1406.2980 \[hep-ex\]](#).
- [615] NA48/2 Collaboration, J. R. Batley et al., *Search for the dark photon in π^0 decays*, *Phys. Lett.* **B746** (2015) 178–185, [arXiv:1504.00607 \[hep-ex\]](#).
- [616] KLOE-2 Collaboration, A. Anastasi et al., *Limit on the production of a new vector boson in $e^+e^- \rightarrow U\gamma$, $U \rightarrow \pi^+\pi^-$ with the KLOE experiment*, *Phys. Lett.* **B757** (2016) 356–361,

- arXiv:1603.06086 [hep-ex].
- [617] KLOE-2 Collaboration, A. Anastasi et al., *Search for dark Higgsstrahlung in $e^+e^- \rightarrow \mu^+\mu^-$ and missing energy events with the KLOE experiment*, *Phys. Lett.* **B747** (2015) 365–372, arXiv:1501.06795 [hep-ex].
- [618] J. H. Chang, R. Essig, and S. D. McDermott, *Revisiting Supernova 1987A Constraints on Dark Photons*, *JHEP* **01** (2017) 107, arXiv:1611.03864 [hep-ph].
- [619] C. Cheung, J. T. Ruderman, L.-T. Wang, and I. Yavin, *Kinetic Mixing as the Origin of Light Dark Scales*, *Phys. Rev.* **D80** (2009) 035008, arXiv:0902.3246 [hep-ph].
- [620] P. Fayet, *U-boson production in e^+e^- annihilations, psi and Upsilon decays, and Light Dark Matter*, *Phys.Rev.* **D75** (2007) 115017, arXiv:hep-ph/0702176 [HEP-PH].
- [621] D. E. Morrissey, D. Poland, and K. M. Zurek, *Abelian Hidden Sectors at a GeV*, *JHEP* **07** (2009) 050, arXiv:0904.2567 [hep-ph].
- [622] S. N. Gninenko and N. V. Krasnikov, *The Muon anomalous magnetic moment and a new light gauge boson*, *Phys. Lett.* **B513** (2001) 119, arXiv:hep-ph/0102222 [hep-ph].
- [623] E.-J. Ahn, R. Engel, T. K. Gaisser, P. Lipari, and T. Stanev, *Cosmic ray interaction event generator SIBYLL 2.1*, *Phys. Rev.* **D80** (2009) 094003, arXiv:0906.4113 [hep-ph].
- [624] F. Riehn, R. Engel, A. Fedynitch, T. K. Gaisser, and T. Stanev, *A new version of the event generator Sibyll*, *PoS ICRC2015* (2016) 558, arXiv:1510.00568 [hep-ph].
- [625] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer, *MadGraph 5 : Going Beyond*, *JHEP* **06** (2011) 128, arXiv:1106.0522 [hep-ph].
- [626] B. Batell, M. Pospelov, and A. Ritz, *Multi-lepton Signatures of a Hidden Sector in Rare B Decays*, *Phys. Rev.* **D83** (2011) 054005, arXiv:0911.4938 [hep-ph].
- [627] N. Blinov, E. Izaguirre, and B. Shuve, *Rare Z Boson Decays to a Hidden Sector*, *Phys. Rev.* **D97** (2018) no. 1, 015009, arXiv:1710.07635 [hep-ph].
- [628] BaBar Collaboration, J. P. Lees et al., *Search for Low-Mass Dark-Sector Higgs Bosons*, *Phys. Rev. Lett.* **108** (2012) 211801, arXiv:1202.1313 [hep-ex].
- [629] Belle Collaboration, I. Jaegle, *Search for the dark photon and the dark Higgs boson at Belle*, *Phys. Rev. Lett.* **114** (2015) no. 21, 211801, arXiv:1502.00084 [hep-ex].
- [630] LHCb Collaboration, R. Aaij et al., *Search for dark photons produced in 13 TeV pp collisions*, arXiv:1710.02867 [hep-ex].
- [631] A. Hook, E. Izaguirre, and J. G. Wacker, *Model Independent Bounds on Kinetic Mixing*, *Adv. High Energy Phys.* **2011** (2011) 859762, arXiv:1006.0973 [hep-ph].
- [632] CMS Collaboration, A. M. Sirunyan et al., *Search for electroweak production of charginos and neutralinos in multilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV*, arXiv:1709.05406 [hep-ex].
- [633] R. D. Peccei and H. R. Quinn, *CP Conservation in the Presence of Instantons*, *Phys. Rev. Lett.* **38** (1977) 1440–1443.
- [634] J. Preskill, M. B. Wise, and F. Wilczek, *Cosmology of the Invisible Axion*, *Phys. Lett.* **120B** (1983) 127–132.
- [635] P. Arias, D. Cadamuro, M. Goodsell, J. Jaeckel, J. Redondo, and A. Ringwald, *WISPy Cold Dark Matter*, *JCAP* **1206** (2012) 013, arXiv:1201.5902 [hep-ph].
- [636] M. Bauer, M. Neubert, and A. Thamm, *The "forgotten" decay $S \rightarrow Z+h$ as a CP analyzer*, arXiv:1607.01016 [hep-ph].
- [637] M. Bauer, M. Neubert, and A. Thamm, *Analyzing the CP Nature of a New Scalar Particle via $S \rightarrow Zh$ Decay*, *Phys. Rev. Lett.* **117** (2016) 181801, arXiv:1610.00009 [hep-ph].
- [638] M. Bauer, M. Neubert, and A. Thamm, *LHC as an Axion Factory: Probing an Axion Explanation for $(g-2)_\mu$ with Exotic Higgs Decays*, *Phys. Rev. Lett.* **119** (2017) no. 3, 031802,

- [arXiv:1704.08207 \[hep-ph\]](#).
- [639] M. Bauer, M. Neubert, and A. Thamm, *Collider Probes of Axion-Like Particles*, **JHEP** **12** (2017) 044, [arXiv:1708.00443 \[hep-ph\]](#).
- [640] ATLAS, CMS Collaboration, G. Aad et al., *Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV*, **JHEP** **08** (2016) 045, [arXiv:1606.02266 \[hep-ex\]](#).
- [641] SLD Electroweak Group, DELPHI, ALEPH, SLD, SLD Heavy Flavour Group, OPAL, LEP Electroweak Working Group, L3 Collaboration, S. Schael et al., *Precision electroweak measurements on the Z resonance*, **Phys. Rept.** **427** (2006) 257–454, [arXiv:hep-ex/0509008 \[hep-ex\]](#).
- [642] CMS Collaboration, V. Khachatryan et al., *Search for neutral resonances decaying into a Z boson and a pair of b jets or τ leptons*, **Phys. Lett.** **B759** (2016) 369–394, [arXiv:1603.02991 \[hep-ex\]](#).
- [643] ATLAS Collaboration, G. Aad et al., *Search for new phenomena in events with at least three photons collected in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, **Eur. Phys. J.** **C76** (2016) no. 4, 210, [arXiv:1509.05051 \[hep-ex\]](#).
- [644] R. Hamberg, W. L. van Neerven, and T. Matsuura, *A complete calculation of the order $\alpha - s^2$ correction to the Drell-Yan K factor*, **Nucl. Phys.** **B359** (1991) 343–405. [Erratum: Nucl. Phys. B644,403(2002)].
- [645] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, *ALPs at FASER: The LHC as a Photon Beam Dump*, [arXiv:1806.02348 \[hep-ph\]](#).
- [646] J. Jaeckel and M. Spannowsky, *Probing MeV to 90 GeV axion-like particles with LEP and LHC*, **Phys. Lett.** **B753** (2016) 482–487, [arXiv:1509.00476 \[hep-ph\]](#).
- [647] H. Fukuda, N. Nagata, H. Otono, and S. Shirai, *Higgsino Dark Matter or Not: Role of Disappearing Track Searches at the LHC and Future Colliders*, **Phys. Lett.** **B781** (2018) 306–311, [arXiv:1703.09675 \[hep-ph\]](#).